

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

Master of Arts – 120 Credit Thesis – 134.816 & 134.817



The Anatomy of Normative Reasons:

Tracking What Normative Reasons Are Via Appearances

Julian Shields

Student [REDACTED]
[REDACTED]
[REDACTED]

G. K. Harrison

Massey University Philosophy Department

Abstract

This paper sets out to provide an anatomy of normative reasons by showing that normative reasons are not facts, properties, or something else entirely, but rather that they are favouring relations, and that the accuracy conditions for our appearances of normative reasons include the existence of a sole non-natural source of this favouring. The method that this paper employs is to take our appearances of normative reasons, along with their content, as the starting point for analysis. From here, we find that all normative reasons have as their content the following relata: an action, belief, or attitude being favoured; an agent being favoured; the favouring itself; and the consideration in which the favouring is being made. I argue that a normative reason isn't any of these relata but is instead the relation itself. Since favouring requires a source, I move to find out what this source is. I show that the source is external to our minds, that there is a sole source, and that this source isn't a mind. The only meta-normative theory which is able to account for this is a monist account where the source of normative reasons is a non-natural entity which we can call the Good.

Acknowledgements

Firstly, I would like to thank my supervisor, Dr. Gerald. Harrison. Your faith in me has helped push me out of my comfort zone of purely literary review style work that I have done in the past to write something that seeks to provide original contributions to the literature. This thesis reflects something akin to what I want to produce in the future, and I could not have written this without your guidance, encouraging my best impulses and discouraging my worst. I hope the times I cite you throughout serves as the proper academic 'thank you'.

Secondly, I would like to thank my beautiful wife, Mrs Amanda Shields. Thank you for constantly being there to support me while I chase my academic dreams. I particularly want to thank you for the support you have given me over the last twelve months. This thesis was written while we were preparing for our wedding, and when we were enjoying our first year of marriage together. I know it wasn't easy having to sacrifice much of our time together for the sake of my university work.

Finally, I would like to acknowledge my father who unexpectedly passed away while finishing this thesis. Dad, I miss you, and hopefully I have done you proud.

Table of Contents

ABSTRACT	1
ACKNOWLEDGEMENTS	2
INTRODUCTION	5
CLARIFYING THE QUESTION AND METHOD	8
THE SUBJECT OF INQUIRY	8
<i>Normative Reasons vs Motivating & Explanatory Reasons</i>	<i>8</i>
<i>Basic vs Derivate Reasons</i>	<i>10</i>
<i>Normative Reason Domains</i>	<i>10</i>
<i>Normative Reasons Discourse</i>	<i>12</i>
THE QUESTION AT HAND	13
THE ROLE APPEARANCES PLAY IN ANSWERING THE QUESTION	16
WHOSE APPEARANCES ARE CONSIDERED	18
META-NORMATIVE THEORIES CONSIDERED	20
NORMATIVE REASONS ARE FAVOURING RELATIONS	25
THE NORMATIVITY OF NORMATIVE REASONS	25
<i>Understanding Normative Reasons Through Oughts</i>	<i>32</i>
THE STRUCTURE OF OUR APPEARANCES OF NORMATIVE REASONS	34
THE IDENTITY OF NORMATIVE REASONS	37
<i>Support For All Normative Reasons Being Favouring Relations</i>	<i>40</i>
THE FACTUALIST CHALLENGE	40
META-NORMATIVE IMPLICATIONS	43
NORMATIVE REASONS ARE EXTERNAL TO ME	46
A PERSONAL SOLE SOURCE OF NORMATIVE REASONS	46
NORMATIVE REASONS AND THEIR RELATIONSHIP WITH MOTIVATIONS AND DESIRES	47
REAL DISAGREEMENTS	50
WHAT THE POINTS RAISED HERE TELL US	52

NORMATIVE REASONS HAVE A SOLE SOURCE	54
THE OPPOSITION OF NORMATIVE REASONS	54
TYPES OF NORMATIVE OPPOSITIONS	55
THE POSSIBILITY OF NORMATIVE CONTRADICTIONS	56
THE POSSIBILITY OF NORMATIVE CONFLICTS	60
INTERDOMAIN AND INTRADOMAIN NORMATIVE CONFLICTS	62
THE PHENOMENOLOGY OF THE STRENGTH OF NORMATIVE REASONS.....	64
THE PARSIMONIOUS CASE FOR A SOLE SOURCE OF NORMATIVE REASONS	67
THE COHERENCE CASE FOR A SOLE SOURCE OF NORMATIVE REASONS	69
THE META-NORMATIVE IMPLICATIONS	71
THE SOURCE OF NORMATIVE REASONS IS NOT A MIND.....	72
NORMATIVE REASONS ARE NOT PERSONAL	73
NORMATIVE REASONS THAT INVOLVE OUR OWN WELL BEING DON'T SEEM TO BE ABOUT CARING	74
DIFFERENT SITUATIONS EXPLAIN DIFFERENT NORMATIVE REASONS FOR DIFFERENT AGENTS.....	75
NORMATIVE REASONS DO NOT VARY ACROSS TIME.....	77
TIME IS NOT NORMATIVELY RELEVANT	77
THE VARIABILITY OF NORMATIVE REASONS OVER TIME	80
NORMATIVE REASONS CANNOT VARY ACROSS TIME	82
THE SOURCE OF NORMATIVE REASONS IS NON-NATURAL.....	87
THE CASE FOR MONIST NON-NATURALISM.....	87
THE STRUCTURE OF NORMATIVITY	95
ACCOUNTING FOR BADNESS	97
THE GOOD AND MORAL PARTICULARISM	98
WHO SHOULD DEFEND MY VIEW?	100
CONCLUDING REMARKS.....	102
REFERENCES.....	103

Introduction

In both academic Philosophy and ordinary life, we talk about reasons. Sometimes we talk about the reason something occurred, and other times we talk about a motivation that we had as the reason we did something. This thesis is about *normative reasons*. A basic definition of normative reasons is that they are reasons for an agent to perform an action, believe something, or hold a particular attitude. When we tell a friend that they have a reason to meet a certain deadline with their university assignment, we are saying that they have a normative reason to meet this deadline. Likewise, when we console a friend and tell them that a B+ mark on an assignment is a reason to feel happy, we are again talking about normative reasons. To add to this basic definition, normative reasons are also reasons to *avoid* actions, beliefs, and attitudes. For instance, we refer to normative reasons when we tell someone that the evidence against a fringe scientific theory gives them a reason to avoid believing in that theory.

My aim in this thesis is to answer the following question: “what are normative reasons?”. When I ask this question, I am not seeking to provide a definition like the one just stated that tells us what role normative reasons are purported to play in the world, nor I am I seeking to say which things in the world, if any, are normative reasons. Instead, I am looking to find out what kind of thing a normative reason is – are they entities, properties, or something else entirely?

The method for answering this question is to appeal to my own appearances of normative reasons. When it seems to me that I or some other agent has a normative reason to

ϕ ¹, I look to the content of this appearance to find out what this can tell me about what normative reasons are.

In early sections of this work, I lay out briefly what appearances of normative reasons are like, and then apply analysis to these appearances. I find that my appearances of normative reasons all have the same four elements in their content. Firstly, normative reasons include probative force, which I refer to as a favouring element. Secondly, we have an agent or agents who this favouring is being directed towards. Thirdly, we have the action, belief, or attitude that the agent is being favoured for or against. Lastly, we have the consideration² in which the agent is being favoured for or against the action, belief, or attitude. From looking at these contents, I argue that normative reasons are not any of these four features in the contents of my appearances, but rather normative reasons are the relations that binds these contents together. Given that these relations have an element of favouring, I conclude that normative reasons are favouring relations.

Normative reasons being favouring relations means that they must be the favouring relations that have some favourer or favourers as their source. The middle section of this thesis looks to find out what the accuracy conditions are for our appearances of normative reasons – that is, what needs to exist as a source for at least some of our appearance of normative reasons to be accurate. In three separate sections, I make three conclusions that each build on the other. I first argue that by examining the nature two or more normative reasons opposing each other, we find that normative reasons are the favouring relations of some sole source. From there, I show that as normative reasons are often appear in disaccord with our desires or motivations, the source of these reasons are external to us. Finally, I find

¹ The phi symbol or ‘ ϕ ’ is a member of the Greek alphabet and is used in this thesis as a placeholder for any generic action, belief, or attitude.

² I will clarify this further in the thesis, but this ‘consideration’ is also referred to as a ‘fact’ and the ‘basis’ in which the agent is being favoured to ϕ .

that as normative reasons are impersonal and cannot vary across space and time, they are not the favouring relations of some mind.

Throughout the thesis, I evaluate how each conclusion impacts different meta-normative theories. I find that all subjectivist accounts – these being views which posit that the source of normative reasons is the subjective states of a mind or mind – are in disaccord with our appearances of normative reasons due to, of course, mind being ruled out as a source. Naturalism – the view that the source of normative reasons are some natural properties or facts – is shown to struggle to account for being the source of normative reasons as natural properties and facts do not favour things, and Naturalists would need to posit some Natural property or fact that can both favour things and be the sole source.

Given the elimination of other meta-normative positions, I spend the last section of the thesis making a case for the source of normative reasons being a sole non-natural entity which I term the Good.

Clarifying the Question and Method

In this section, I look to clarify the aim of this thesis by providing a basic question which I am setting out to answer and then providing the method that I implore to answer this question.

The Subject of Inquiry

My thesis is about *normative reasons*. While the term ‘normative reason’ is a technical term in academic Philosophy, it is one which is used to describe something that we appeal to in our day-to-day lives. When we tell a friend that they have a reason to meet a certain deadline with their university assignment, we are saying that they have a normative reason to meet this deadline. Likewise, when I see the number on the bathroom scales go up over the course of a few weeks, it seems to me that I have a normative reason to address my diet and exercise habits.

Normative Reasons vs Motivating & Explanatory Reasons

In everyday discourse, we refer to normative reasons simply as “reasons”. The prefix ‘normative’ is added by academic philosophers to help distinguish normative reasons from other uses of the word ‘reason’³. Other common uses of the word ‘reason’ occur when we talk about what philosophers deem explanatory reasons and motivating reasons. Explanatory reasons are what we invoke when we explain *why* something occurred or will occur. For example, we would be invoking an explanatory reason if we were explaining to someone that the reason a tree fell over in our backyard was because it was struck by lightning. Motivating reasons are held by some philosophers as being a type of explanatory reason that are to do with explaining why an agent took a particular action, adopted a particular belief, or chose to

³ To build on this etymology, Alvarez (2017) holds that the term ‘normative reason’ is derived from the idea that there are norms or codes that direct our actions and beliefs in particular ways.

hold a particular attitude⁴. For instance, when discussing the action someone took in donating a large sum of money to a cancer research charity, a motivational reason for doing so would be a strong desire to see a family member's illness to be treated. When looking at beliefs, someone who believes the proposition "Plato was Aristotle's teacher" on good evidence may have as their motivating reason the desire to hold true beliefs. Attitudes may have motivational reasons attached to them also. Take a child who finds that a group of prospective friends all share a love of football and judges that forming a positive attitude toward football would help them make friends. If the child does indeed form this attitude, we would refer to the motivation, or motivating reason, as something which included the desire to make friends.

While normative reasons are distinct from explanatory and motivating reasons, they do share a relationship with both. Motivating reasons may be invoked in certain scenarios where if one is motivated to ϕ then one has a normative reason to ϕ . Take someone with a strong desire to become a professional basketball player. This person's desire could give rise to a motivating reason for them to work towards becoming a professional basketball player, and this same desire could also give rise to a normative reason to follow this career path. We can cash this out as desires both motivating agents to outcomes that bring about their desire as well as desires serving as *prima facie* grounds for normative reasons. Furthermore, we could also talk about an explanatory reason as being the basis in which we have a normative reason to do or avoid something, such as if pushing a button causes a fire, then we have a normative reason to avoid pushing said button.

Yet, it is important to not conflate this relationship of normative reasons with both explanatory reasons and motivating reasons as one of identity. Focussing on motivational

⁴ There is debate in the literature on whether we should make a division between motivating reasons and explanatory reasons. For a summary of this debate, see Alvarez (2017).

reasons here to explain the point, while it does happen to be the case that many motivations we have are tied to normative reasons – as said that we could cash this out in terms of desires both motivating and providing *prima facie* grounds for normative reasons to work towards those desires – there are of course outlier scenarios, hence the term *prima facie*. You can take the stereotypical serial killer character from movies and television shows. This killer is often portrayed as wanting to kill purely due to the desire to inflict harm on others. In this scenario, we may say that their motivational reason for killing someone is a desire to inflict harm, but we may say that they had no normative reason to kill, and that they had a normative reason to act in disaccord with their desire.

Basic vs Derivate Reasons

Normative reasons can be both basic and derivative (Reisner, 2015, p. 13). We have a basic normative reason to ϕ if we have a normative reason to ϕ and we can act on this normative reason merely by ϕ -ing. To use an example, I may have a normative reason to believe that trees are made of wood and by believing that trees are made of wood, this normative reason is satisfied. Derivative reasons are when some action, belief, or attitude is a means to acting upon a normative reason. For instance, if I have a normative reason to go to work, then I have a derivative reason to get out of bed since getting out of bed is a means to getting to work.

Normative Reason Domains

We can divide normative reasons up in to at least three separate domains: actions, beliefs, and attitudes. While this isn't to deny that other domains exist, these three domains will be the focus of this thesis with all examples and thought experiments used throughout this thesis invoking at least one these domains.

Normative reasons for actions are invoked any time we talk about an agent having something that they are required to do or avoid in some way. We may have a friend who is thirsty, so we say to this friend that they need to go and get a glass of water. Perhaps this friend also needs a boost of energy as they have some tasks to do, so we may instead say that they have a reason to get a coffee instead of a glass of water. This same friend may be on a diet, meaning that, though they would really feel like one, they have a reason to avoid having a biscuit with their coffee.

Reasons for belief⁵ are invoked when we say that an agent either has a normative reason to believe something, or avoid believing something, or even dropping a particular belief that they hold. Most commonly, we seem to have normative reasons to believe something based on evidential grounds. For instance, there is ample scientific evidence that the earth revolves around the sun, so we all – or at least those of us who are aware of a reasonable amount of this evidence – have a normative reason to believe the heliocentric model of the solar system is correct. But despite evidence and beliefs seeming to be a natural pairing, not all normative reasons we have for belief seem to be given rise to by evidence. Take a situation in which a shy person would benefit from believing that they made a good impression during a social situation, as this would make them happy and increase their confidence. In lieu of any good evidence to the contrary, it seems that they have a normative reason to indeed believe that the impression they made was good. This reason would be given rise to in absence of evidence of the belief being true. Likewise, we often seem to have moral reasons to believe things, such as a mother believing that her son is not guilty of some crime after he tells her he is innocent. This is because, for many of us, it seems we have a moral

⁵ Often in the literature, normative reasons to believe something are referred to as ‘epistemic reasons’. I avoid this terminology as there is some confusion in the literature as to this term’s proper usage. Turri (2009) and Harrison (2018), for instance, refer to *any* normative reason for belief as an epistemic reason, while Bondy (2019) refers to epistemic reasons solely as evidence-based reasons for belief.

duty to believe the best of our loved ones even sometimes despite strong evidence to the contrary of what they are telling us.

We talk about normative reasons for attitudes when we talk about the ways in which someone should or shouldn't feel towards something or someone. A good example of this is when we tell a friend that they should feel positive about the grade that they got on a university assignment. What we are saying is that they should adopt a particular attitude toward this grade, one of positivity. Another example is in how we form attitudes towards people. Many of us have very negative attitudes towards Adolf Hitler which are generally based on his warmongering and genocidal actions. So, we have normative reasons for these attitudes. Reasons for attitudes also appear to be tied to aesthetics. For instance, the fact that a horizon is beautiful may be a reason to adopt an adoring attitude it, and the fact that a blanket feels soft may be a reason to take a favouring attitude of liking it over another.

Normative Reasons Discourse

While we often invoke normative reasons in our everyday lives, it is important to note that not all statements which invoke normative reasons make explicit reference to them. Sometimes normative reason discourse is hidden implicitly with other, often synonymous language. For instance, the phrases "it is raining, so take an umbrella" and "it is raining out there so you should take an umbrella" are elliptical to saying, "the fact it is raining provides a normative reason for you to take an umbrella". Similarly, when talking about beliefs, we may say "I cannot understand how you can believe in ghosts, there is no good evidence for their existence", which is again elliptical to saying, "there is no good evidence for the existence of ghosts, so you have no reason to believe in them".

Despite the lack of explicit mention in much of everyday speech, we often recognise language which implicitly involves reference to normative reasons by the involvement of

some markers. Firstly, normative reason discourse often has a directive nature to it. “You should ϕ because of x ” is just one example since it takes the form of directing or convincing an agent towards some action, belief, or attitude on some reason or reasons. Even a threatening remark such as “you should pass me the PlayStation controller because otherwise I will smack you” has an implicit appeal to normative reasons as the recipient is being advised that they have a reason to do as the speakers says.

But not all normative reason discourse takes place between interlocutors about each other’s normative reasons. Someone making the claim that “Politician P should appeal more to the working classes otherwise he isn’t going to get the votes needed” is, again, talking about normative reasons. Likewise, when someone says that “Footballer F should never have made that slide tackle as it was reckless and now his team are down a man”, a claim about an action an agent should have avoided but didn’t, they are talking about normative reasons.

Some normative reason discourse is hypothetical. We may say “if x then A should ϕ ”. Take the example of the weather where we may say that someone has a reason to take an umbrella *if* it is raining tomorrow. Or we may instead talk about normative reasons that were not the case but could have been in the past, such as “if A hadn’t lashed out violently then I would think more highly of them”.

What ties all normative reasons discourse together is that it is about rational data that agents can act upon, and I say ‘act’ to not only mean performing an action but also adopting a belief, or holding an attitude, as well as converse of these.

The Question at Hand

Now that I have outlined some details about normative reasons, it is time to talk about what sort of enquiry I will be making into them. If I had to phrase the question of this thesis in a simple sentence it would be “what is a normative reason?” This question, however, is

ambiguous and needs to be clarified. I follow my supervisor's warning on this that a question of this nature can be interpreted in both a first order and second order sense⁶.

First order questions about normative reasons are about what sorts of things that we have normative reason to do. The question then may be understood in a first order way as being elliptical to asking what I and other agents have reason to carry out or avoid. Harrison (2018) helpfully distinguishes first order questions as not being about what normative reasons are but *where* they are (p. 8). I find this phrasing quite helpful in that if I were to ask what a human being is, a first order response would be locate all the things in the universe that are human beings, so this would include yourself, myself, Joe Biden, and Hulk Hogan among many, many others. Here, what we are doing is pointing to *where* human beings are as opposed to finding out what all of these agents have in common that makes them human.

My thesis intends to answer a second order question. Second order questions are not about what instantiations there are of something, or where something is, but rather what instantiations of the thing in question are made of, or what it is about them that makes them that thing. If I were to ask again what a human being was but with clarifying that I wanted a second order response, I would expect to see something along the lines of a biological being that has *x* set of DNA or perhaps a being which fits in to a socially constructed category. This then isn't telling me where human beings are but rather what it is to be a human being, what human beings are *made of*. Likewise, when I ask what a normative reason is, I am not asking which things are normative reasons, or what we have normative reasons to do. Rather, I am asking what normative reasons are made of.

To find what a proper second order response to the question of what normative reasons are, I am looking for two components. The first component is the category that

⁶ See Harrison (2018, p. 8) for his outline of this.

normative reasons fit in to. Different things are made of different metaphysical material and if they exist they do so in different categories. Take the following list of things:

- *A human being.*
- *The Earth.*
- *Gravity.*
- *The colour red.*
- *The validity of an argument.*
- *The rightness of an action.*
- *The friendship between two people.*

All of these are ‘things’ in a generic sense – we identify each of them with a noun.

Yet all these things, if they do exist, do so in different categories of existence. Human beings and the earth are examples of entities, these being things which exist in their own right and do not appear to have their existence dictated by something else’s existence. Gravity is a force and doesn’t have a physical body to it. Redness, validity, and rightness could be considered to be properties in that entities bear properties and properties do not exist without being the property of an entity. Friendship may be considered an example of a relation in that to be friends with someone is to stand in a particular relation to that person.

The overall existence of each of these categories and how they cut off from one another is not of major importance here. What is important for my discussion is merely the category that normative reasons fit in to. So, the first part of providing a response to my thesis question is to answer which category normative reasons fit in to.

Providing just the ontological categorical designation is not enough. If a response to a second order question about human beings merely told me that they were entities, while I would feel somewhat fulfilled in that I could then eliminate numerous things which aren’t

entities, I would still be unsure as to what human beings are. Instead, I would expect to know what sort of entities they were. If someone were to hold that they were natural entities consisting of a biological body, I would want to know what these bodies are made of – I would want an anatomy of a human being. I am not only finding out what category normative reasons fit in to, but I will also be providing an anatomy of normative reasons.

The Role Appearances Play in Answering the Question

The method I have chosen for answering the question in this thesis is to appeal to the appearances that we all have of normative reasons.

The term ‘appearance’ is likely to conjure up the thought of visual perception, however I use the term a lot more broadly. For instance, we can touch something, and that thing can *appear* hard or soft to us, and we can hear something that *appears* high or low pitched. These along with other sensory data fall under the category of sense appearances. There are other types of appearances that do not appeal to our senses. While not an exhaustive list, these include mnemonic, introspective, and intellectual appearances (Huemer, 2005, p. 99). When we have a memory that something happened to us, or we seem to recall some piece of information, this is what is referred to as a mnemonic appearance. Introspective appearances occur when it seems that we feel a certain way. Some examples are when I find myself feeling happy, sad, irritated, or content with something. Intellectual appearances are to do with how we perceive more abstract matters. Take the equation ‘ $1 + 1 = 2$ ’ and the equation ‘ $2 + 2 = 5$ ’. The first of these equations seems to be right while the second seems to be wrong. But we do not see this rightness or wrongness as something out there in the world. Even if we read these two equations, this rightness or wrongness is not part of the visual appearance but is instead something that we see with our intellect or reason.

When trying to find out what type of appearance a normative reason is, I can look at when it is that I have appearances of normative reasons. Sometimes normative reasons appear to me to exist when I visually perceive something. It looks dark and cloudy outside and this seems to provide me with a normative reason to wear a waterproof coat on when I go outside. Other times, I may hear something, such as a fire alarm, and so it appears that I have a normative reason to exit the building I am in.

Yet, appearances of normative reasons are not identical with this sense appearances. My appearance of something that I see with my eyes is not the same as my appearance of a normative reason. To say that I witness something, and it appears to give me a normative reason to do something is not the same as saying I witness a normative reason where both are provided by our senses. Normative reasons, whatever they are, aren't the sorts of things that I see with my ordinary senses.

Further evidence for the appearances of normative reasons not being sense appearances is that my appearance of normative reasons are not confined to when I am having a sensory experience. Memories also prompt appearances of normative reasons, like when I remember back to a time when there was a choice to make and it appears to me, in hindsight, that I had a normative reason to make a decision I hadn't even considered. There are even instances of abstract thinking that are linked with normative reason appearances. For instance, when I try to think of a good logical syllogism to make a point in an essay and I notice a conclusion which doesn't follow from the premises. Here, it seems to me that I have a normative reason to believe that the syllogism is invalid.

Out of all the types of appearances that have been considered, the one that most captures what a normative reason is like is appearance of validity, as in the validity of an argument. This is a sort of intellectual appearance as propositions appear to us to be either

valid or non-valid. For instance, we can read an argument and it appears valid, and we can hear an argument that appears valid, or we can imagine an argument, and this also appears valid. In all these instances, it isn't that the validity appears on the page, through our hearing, or in our minds as the words or sounds in the argument, but rather validity is something else entirely, as it is something that appears to us to be the case via our intellect.

Appearances of normative reasons have similar phenomenological content to appearances of validity. Certain states of affairs, whether actual or hypothetical, true or false, give rise to the appearance of a normative reasons for either ourselves or others. Some appear to be fairly basic and require little other information. If I become aware that the spot I am presently standing in is going to have a large rock dropped on it in two minutes, it appears to me that I have a normative reason to move. The method in which I obtain this information doesn't impact the appearance, such as whether it is shouted at me, I read it, or I happen to remember a prophetic dream I had the night before, just as long as I seem to recognise the information as reliable or reliable enough to warrant caution. Sometimes states of affairs seem linked to normative reasons based on other bits of information. I may come across information that a television show has been renewed for an additional season and it may appear to me that my wife has a normative reason to feel excited as this information is also linked with other information that I hold, in this case that it is a show she highly enjoys.

Whose Appearances are Considered

At least some of us have different appearances when it comes to normative reasons. Many of us see the same situation and make different normative judgements, with thought experiments that commonly occur in undergraduate classes, such as the trolley problem, being an example of this. Some believe that we have reason to pull the lever to save five people with the sacrifice of one, while others think an intervention in fate is wrong and we have a normative reason to avoid pulling the lever. With that said, when we make normative

judgements, or when we engage in normative discourse with others, it does seem that we are all engaging with the same *kind* of appearance. Take the following two scenarios:

- *A and B are arguing as to whether or not the newest Star Wars film was cool. A says it was cool due to the lightsabre battles whereas B argues that it isn't cool as when he touches the disc the movie is on, it appears hot.*
- *A and B are arguing whether or not a shirt that B just purchased is blue or green. A says that it is blue as it appears to them as blue whereas B argues that it appears to be green.*

The first scenario involves a pseudo-disagreement. A is referring to the movie in question being cool, as in rad, awesome, etc, where B is referring to the temperature of the disc that the movie resides on. This isn't a real disagreement since they are involved in two different discursive areas and are merely talking past each other. The second scenario involves an actual disagreement. Both seem to be making a claim as to the objective status of an object's colour, though the disagreement is explained, in part at least, by A and B seeing the object as having a different colour property.

Our disagreements around normative reasons appear to be like the second case and not the first. When we navigate our day-to-day disagreements with others on what we have reason to do, there doesn't appear to be some confusion between talking about some abstract concept of a normative reason and, say, the colour purple that an interlocutor calls a reason. To evidence this further, we can look at the actual content of when normative reason disagreements are unable to be resolved versus when they are able to be resolved. Agents A and B may have a disagreement as to whether A should keep the promise they made to C about their keeping their knowledge of C's business affairs confidential, or whether A should break their promise to C and involve police due to illegal activity C is involved with. A and B

may disagree on many grounds, such as if the knowledge of illegal activity does provide a normative reason or whether there are multiple reasons and the ultimate reduction weighting of them. Sometimes disagreements can be resolved. For instance, if A is telling B about C's affairs and initially neglects to outline the harm caused by C's affairs, B may, with the addition of this new information, come to see the normativity of the situation identically. This is analogous then to the colour dispute where A and B may converge on a consensus with additional information, such as finding out that one of them is colourblind.

Despite disagreements around normative reasons, it appears to us that normative reasons are at least the same sort of thing. To find out what sorts of things normative reasons are, then it seems that an inquiry into our collective appearances of normative reasons will be fruitful. As the author, I only have clear evidence of my own appearances of normative reasons, but I don't consider myself atypical in anyway, and I expect that many of my appearances of normative reasons will concur with those of yours, the reader. For the sake of this thesis, I will be alternating between discussing my own appearances, as well as making cases where I think that you will see normative situations a certain way. This means deviating between statements such as "it appears to me that..." to statements such as "it appears to us that x in situation S". While not all intuitive appeals to the broader audience is likely to land, I expect that reader will be on board for most of the scenarios that I use.

What we can say then is that it is our appearances of normative reasons which are being considered here, and it is the accuracy of our normative reason appearances that we will at to find what the accuracy conditions of normative reasons happen to be.

Meta-Normative Theories Considered

This thesis will attempt to find out what normative reasons are by finding out what the accuracy conditions are for our appearances of normative reasons. To skip ahead a little, in

the next section, we will find that normative reasons, whatever they are, require a source for their existence. When providing an account of each normative reason, part of this will include providing a theory as to what this source is. We can refer to this as a meta-normative theory – i.e., a theory as to what the source of normativity is in a second order sense.

There are five meta-normative theories that I am looking at as part of this thesis. These theories are divided into two different categories. On the one hand, we have subjectivist theories. Subjectivism holds that there are no objective normative facts and that if normative facts exist these are constituted by subjective mental states. On the other hand, we have objectivist theories. Objectivists hold that if normative facts exist, these are not constituted by subjective mental states and that the normative facts exist irrespective of our attitudes towards them.

I will try to parse these explanations out further for those unfamiliar with the subjective and objective divide in meta-normativity. For subjectivists, if agent A has a normative reason to ϕ , this normative reason is provided by some mind or minds who favour or command A to ϕ . It could be the case that A has a normative reason to have a drink of water, so there must be at least some agent who either wants A to drink water or commands A to drink water. Under subjectivism, the normative reason is either identical with or is given rise to in some way by the mental states involved, these being the favouring or commands. For objectivists, if agent A has a normative reason to ϕ , this normative reason is not provided by some mind or minds and holds irrespective of an agent's attitudes or favouring towards it. Let us say for instance that A has a normative reason to keep a promise to B. It could be the case that A has a basic reason to keep their promise, if there is some objective feature about promise keeping – such as goodness, that makes it normative reason providing – or a derivative reason if the keeping of A's promise will lead to some outcome that has an objective feature that makes it reason providing – like if it will make B happy and happiness

itself was good. Regardless, it would be this objective feature of goodness that was reason providing, and this normative reason would still hold even if all agents took a negative stance towards promise keeping or happiness.

The division between the theories in the subjectivist category relates to whose mental states are the source of normative reasons and whom those normative reasons can be for, whereas the division between theories in the objectivist category relates to whether the objective feature or features of reality that are the source of normative reasons are part of the natural world or not.

The first meta-normative theory considered is Individualist Subjectivism. For individualist subjectivists, what an agent has normative reason to do is constituted by an agent's own subjective mental states. If A has a normative reason to ϕ , then it may be because A desires ϕ , is motivated to ϕ , holds a favourable attitude toward ϕ -ing, or ϕ -ing is a means to something they desire or hold a favourable attitude towards. Under this view, all agents are capable of being the source of normative reasons, but they are only the source of their own normative reasons. So, if A desires that B does ϕ , whether B has a normative reason to ϕ is not constituted by or brought about directly by A's desires. B may have a normative reason to ϕ , but this would be because B may desire themselves to ϕ or ϕ -ing is a derivative way to meet some desire, say if they want to appease A.

The next meta-normative theory considered is Normative Relativism. Under Normative Relativism, A has a normative reason to ϕ if A's culture or society favours people ϕ -ing or ϕ -ing is a means to something that A's culture or society favours. Relativists then hold that all agents' minds are capable of providing normative reasons, though only in conjunction with other minds in their same culture or society, and these normative reasons only apply to agents who are within this culture or society.

Divine Command Theory (or DCT for short) is the final subjectivist theory considered. Under DCT, an agent has a normative reason to ϕ *iff* a god or gods favour A ϕ -ing, or ϕ -ing is a means to bringing about something that a god or gods favour. According to Divine Command Theorists then, only the mental states of a god are possible of providing normative reasons, but they can provide normative reasons to all agents. Despite the name, the ‘favouring’ that is provides normative reasons is not necessarily literal commands. Agent A may have a normative reason to ϕ whether a god commands them to ϕ or a god would just like them to ϕ . To put this in other terms, take myself and my cat: I want my cat to not wake me up at 2am for food, and can either command him not to or just hold an attitude of not wanting him to. If I were a god (which I am not) and my cat were a rational agent (which he often seems not to be), my cat would have a normative reason to not wake me up at 2am whether I commanded him not to or just wanted him not to.

Naturalism is the first of two objectivist theories considered. To return to what was said earlier, objectivists hold that if agent A has a normative reason to ϕ , this normative reason holds irrespective of the attitudes and favouring of minds and instead is provided by some objective feature or features of reality. For naturalists, this objective feature or features are part of the natural world. What constitutes ‘natural’ here are those things which are discoverable by the empirical sciences as well as psychological states (Moore, 1903). The inclusion of psychological states may at first seem to confuse the theory with those under the category of subjectivism, however what naturalists have in mind is different to the favouring attitudes as being what gives rise to normative reasons. Instead, naturalists may hold that certain psychological states such as happiness are what provide normative reasons. For instance, naturalists may find that what we have normative reasons to do are those things which increase the net happiness in the universe. This would hold regardless of whether agents in the world took a dim view of increasing happiness or not.

The final meta-normative theory considered is Non-naturalism. Like naturalists, non-naturalists holds the view that what gives rise to normative reasons is some objective feature or features of reality, however this feature is not identical with or reducible to some feature or features of the natural world. So, whatever it is that gives rise to normative reasons, whether it is an entity, property, or fact, is not discoverable by the empirical sciences nor is it a psychological state. This may seem a little too abstract, but the view itself was first popularised by George Edward Moore in his 1903 *Principia Ethica* with the goal to explain what ‘good’ is from a second order perspective.

Normative Reasons are Favouring Relations

In this section, I look to our appearances of normative reasons to find which ontological category they fit in to. I do this by examining my own appearances and finding that any time that I have an appearance that there is a normative reason, the content of this appearance involves four elements: 1. An agent being favoured; 2. The action, belief, or attitude that the agent is being favoured towards; 3. The consideration in which they are favoured to act, believe, or hold an attitude; 4. The favouring itself. From here, I argue that normative reasons are not any of these elements but are instead the relation that exists between them. Finally, I show that the contemporary consensus position on what normative reasons is mistaken. This consensus view that normative reasons are facts is unable to account for the normativity of normative reasons and is thus false.

The Normativity of Normative Reasons

Normative reasons, whatever it is that they are, happen to be connected to actions, beliefs, and attitudes in a particularly normative way. The term normativity itself though is quite a slippery term, not only for laymen but also philosophers themselves. Stephen Finlay (2019), in a paper which seeks to find out whether philosophers use the term ‘normativity’ in a consensus way, found that there were seven different uses of the term ‘normative’ and ten different uses of the term ‘normativity’ throughout the literature he examined⁷. This is suggestive of many debates which use normative terms run the risk of the interlocutors talking past one another. To avoid walking into this murky world of unrefined terms, I will here attempt to provide an account of what the prefix ‘normative’ refers to in normative reasons by appealing to appearances of normative reasons and coaxing out what this term ‘normative’ is referring to.

⁷ Refer to the glossary section of the paper that holds all of the terms identified.

The appearance of a normative reason is not identical to an appearance of whether I did ϕ . I can look back on past events and it can appear to me that I had normative reasons to act, believe, or hold attitudes to certain things, but this appearance is not identical to whether I did any of these things. An example of such an appearance is when I look back on a particular essay I wrote and handed in. After receiving feedback, I may realise that there was a particular section of the essay that I had a normative reason to omit but that I ultimately included. So, it appears that I didn't ϕ and it appears that I had a normative reason to ϕ .

Likewise, when it appears to me that I have a normative reason to ϕ , this isn't an appearance that I *am* ϕ -ing. If I find that it appears I am ϕ -ing and that I have a normative reason to ϕ , these are two different appearances. I may find myself drinking a cup of tea and wonder whether I have any reason to be drinking this tea. It can then seem to me that I do in fact have a reason to drink this cup of tea, maybe on account of just the fact that I am thirsty, or possibly due to needing a small caffeine boost to assist in completing a particular task. Likewise, if it seems to me that I have a normative reason to drink a cup of tea, this obtains whether I am drinking a tea, drinking water, or sitting and doing nothing.

Furthermore, when it appears to me that I have a normative reason to ϕ , it doesn't appear to be identical to a prediction that I will ϕ . Much of the time, when it appears that I have a normative reason to ϕ , I tend to believe that I will ϕ , and often I find that I do ϕ . But there are instances where I do not believe that I will ϕ and the appearance of a normative reason still stands. Take for example when I accidentally promise two people that I will go to a different movie with each of them on the same night and at the same time. While it still seems that I have normative reasons to accompany both persons to their respective movies of choice, I don't believe that I will accompany both persons since it is physically impossible to be at two different movies at once, thus I know at least one of these normative reasons will go unacted on yet the appearance of it still remains.

Whether I did something, am doing something, or will do something isn't identical with a normative reason I had, have, or will have to do something, this much is clear, though it is of course connected to each of these actions in some way. Perhaps normative reasons are options that we have for actions, beliefs, and attitudes. We can define options as actions, beliefs, or attitudes which agents are both logically and physically capable of performing or adopting. The phrase "you have an option to draw a four-sided triangle" is something which violates both physical and logical laws, so while it has the same grammatical structure as, say, "you have an option to apply for that job", which does not, we can consider the former sentence as not outlining an actual option.

At first glance, normative reasons being options has some intuitive appeal. Looking back at my previous example of where I have mistakenly promised two people that I will go to the movies with each of them, both movie going experiences do seem like options that I can act upon as acting on either is both physically and logically possible. To further support this, there are claims about normative reasons that seem *prima facie* true, but that are ultimately contingent on physical and logical possibility. An extreme example is that it seems to me true that if there is a genocidal dictator that I have a normative reason to stop this person acting on their genocidal impulses. Yet even though Hitler was a genocidal dictator, I do not have a normative reason to stop Hitler since it is physically impossible for me to travel to the past to stop him. A less extreme example is that it seems true that if a task would net me millions of dollars, I would have a normative reason to pursue it, however if someone offered me millions of dollars to draw a three-sided hexagon, it doesn't seem that I have a normative reason to do so since this drawing this is not an option available to me.

Despite the initial intuitive appeal, normative reasons as being identical to options has a major problem: there is a plethora of physically and logically possible things that I can do at any one time that it does not seem that I have any normative reason to do. For instance,

imagine that I have saved up a modest amount of money over the course of a few years. It may seem to me that I have a normative reason to invest this money into buying a house, which is one option. It may also seem to me that I have a normative reason to invest these funds into the stock market, which is another option. But it does not seem to me that I have a normative reason to withdraw all my money and burn it all in my backyard, even though this would be both physically and logically possible to do so. The syntax of our language reflects this also. I may say to someone “I could ϕ ” and the response could be “yes, but you have no reason to ϕ ”. In this reply, it doesn’t seem that my interlocuter in any way is misunderstanding the concepts involved or making any error. This seems to be a perfectly coherent response, particularly if I insert in the place of ϕ some kind of ludicrous scenario, such as gluing a fake moustache to my face, writing a complaint letter to a food brand whose product I have never tried, or naming my first-born child after someone I dislike. Although all of these are options I could take, the response “you have no reason to ϕ ” after suggesting any of these seems entirely sensible.

Even when I look at options that do seem to be related to normative reasons, the normative reasons as options account does not explain the fact that it seems to me that some options appear to be more compelling than others. When I am at work and I am trying to finish an email and the fire alarm goes off, it seems that I now have two options which coincide with normative reasons: to finish the email and to evacuate the building. At that time, however, the option to evacuate the building seems much more compelling to me to act upon.

In this example provided, it is tempting to just cash this compulsion out as merely some connection to my own desires since, of course, it seems that my own safety is more important to me than the completion of a trivial work task. While it does often seem that I am compelled towards certain options that are attached to my desires, what I am getting at is a

certain type of compulsion where options are recognised as also being attached to normative reasons in some way. We can take an example where I have a strong desire to eat because I haven't eaten all day and I am placed in a room where there are two buttons, one which will give me food and another that will feed a hundred starving children for the rest of their lives but I will not be able to eat for another ten hours – I can only press one button. I recognise that I have both a normative reason to press the button that will feed me and one to feed the starving children. Despite being in heightened state of desire to satiate my own hunger, the normative reason to press the button to feed the starving children feels more compelling in a way disconnected with my desires. In this situation, I recognise that the normative reason to help the starving children has more *probative force* to it. In other terms, I feel *required* to push each button, but I am more strongly required to press the button to help the children than I am to press the button to help myself.

Finding out what this feeling of requirement or compulsion amounts to is still a difficult task. When I have the appearance that I have a normative reason to ϕ , this is not to say that it appears that I am subject to some sort of physical force which is pushing me to ϕ (Harrison, 2018, p 13). This is because normative reasons generally have the appearance of being things that I can physically avoid doing. For instance, it seems to me that I have a normative reason to go in to work on Monday, yet it seems entirely possible that I can just choose not to go in to work and choosing not to would not involve me battling against some physical force. Furthermore, most of my appearances of being pulled or pushed do not contain normative reasons in their content. For instance, it appears the force of gravity is pulling me to the ground, but it doesn't appear to me that I have a normative reason to stay on the ground, particularly if I am jumping in the air trying to catch a ball. Similarly, when it appears to me that I have a normative reason to ϕ , this normative reason can be to do something which is in complete opposition to some force. Imagine a dangerous tornado

ripping through town. It seems to me that I have a normative reason to fight against this force of nature and avoid being pulled away.

What this feeling of being compelled to ϕ or required to ϕ when it appears that I have a normative reason to ϕ is not explained by something physical pushing me in some direction, but rather this compulsion is something mental. This isn't to say, however, that when it seems that I have a normative reason to ϕ , that my mind is being controlled against my will. When it seems to me that I have a normative reason to ϕ , it doesn't seem like my free will is being subverted in any way. Instead, the appearance that I have a normative reason to ϕ is similar to another sort of appearance that I often seem to have, namely my own favouring of myself to ϕ . It seems to me that I can favour myself to ϕ , for instance ϕ -ing is something that I want to do or will bring about something that I want. I may favour myself eating a cheeseburger just because I really desire a cheeseburger. When this occurs, I feel compelled to eat a cheeseburger, with this compulsion being a direct result of my favouring to eat a cheeseburger. One may challenge this as perhaps it is my desire doing the work here, but I likewise seem compelled to ϕ when someone else favours me to ϕ , even if I don't myself desire ϕ -ing or if ϕ -ing brings about something I desire. This is similar to the phenomenology of a normative reason as when it seems to me that I have a normative reason to ϕ , the appearance involves feeling as if I am favouring to ϕ to someone or something.

I can say then that when I am favoured either by myself or others to ϕ , then I feel at least in some sense compelled to ϕ or required to ϕ . Sometimes I can recognise that I have no normative reason to ϕ , for instance, if the person wants me to do something that seems detrimental to myself or if it goes against a legitimate normative reason. An example would be if I am on a diet and though I feel compelled to eat a cheeseburger, I know that this compulsion needs to stay unacted on since it seems that I have a normative reason to avoid eating high-calorie food. Cashing this favouring or compelling aspect of normative reasons

out then, we may say that the ‘normativity’ around normative reasons is to do with a favouring element. It would then be the case if “I have a normative reason to ϕ ” is a true statement then “I am favoured to ϕ ” is a true statement also.

The normative element in normative reasons as being the result of favouring can be supported by several arguments. The first is that favouring tends to have degrees of strength, which is consistent with how normative reasons seem to be like. If we take favouring as a generic term for any degree of someone being directed to do something, this is certainly apparent. Someone can weakly want me to ϕ , like if my wife wants me to get her a glass of water from the kitchen. While she favours me to get her a glass of water, she certainly doesn’t favour it as strongly as she does for me to adhere to my wedding vows. Likewise, when it seems to me that I have a normative reason to pick up a small coin off the ground, this feels like a very weak favouring in that if the coin is in the middle of the road, the normative reason I have to avoid traffic seems a lot stronger, as if the source of this favouring values my own health more than it values me increasing my wealth by a modest at best sum.

Another reason to support normative reasons involving favouring is that normative reasons feel like things I am required to do. Requirements are things which seem intimately tied to favouring. For instance, my boss at work can require me to undertake a course – in other words, she favours me to do something. Take another example with the following basic requirement claim: “if you want to play football next season, you are required to be signed up by next Sunday”. This claim expresses that football registration closes on a particular date and that those that are required to register are those that want to play. If I favour myself playing football, then I am required to have my registration in, but if I don’t require myself to play football then I don’t. Here, this shows that the source of the favouring is myself as it seems that I am also the source of the requirement.

I have said that the appearance that I have a normative reason to ϕ involves me being compelled to ϕ . How I can explain this compulsion is the same way as I can explain my compulsion as a child to listen to my mother and now as an adult to listen to a safety warden during an emergency drill, in that whenever I am favoured by someone to do something, I feel at least somewhat compelled to do it, since I am reacting to a directive of some sort. In other words, me being favoured to ϕ is me being directed to ϕ , and compulsion to ϕ appears to be a by-product of me feeling as if I am being directed to ϕ . The level of favouring, from a light suggestion to a yelling directive seems to contain a mirrored compulsion where if I am lightly suggested to do something, I feel a bit compelled versus feeling very compelled if I am being yelled at to do that thing. This of course ties into the strength that normative reasons seems to have, with my compulsion to ϕ being tied to the level that I am favoured to ϕ .

Understanding Normative Reasons Through Oughts

The final support for normative reasons involving favouring is that normative reasons are invoked when we talk about ought or should claims. There are two senses of the word ‘ought’ and ‘should’ that we need to distinguish between. Firstly, there is the predictive sense. In this sense of the words, we say things such as, “your flu ought to subsist in a few days” and “the weather should pass by this afternoon”. Here, we are not making statements as to what your flu or the weather is directed or favoured to do, we are instead making a prediction as to what will occur⁸. The second sense of these words is normative. When we say things such as, “you ought to carry out your duty” or “you should believe the theory of evolution”, we are not making predictions, we are expressing that the subject of the sentence is directed or favoured to do or believe something, or to feel a certain way. Sometimes ought and should statements are ambiguous as to whether they are predictions or normative

⁸ Another lesser user predictive sense is when we talk about failed predictions, such as “I pushed the button so the kettle should have boiled, but it didn’t”.

statements. An example is if someone says that “the All Blacks ought to win by 20 points”. Normally within a conversation, we recognise the context that this is said in and the tone in which it is said and this tells us whether someone is making a predication or a normative statement, since they may be predicting the score of a game, believing that the All Blacks are 20 point favourites against their opposition, or saying that the All Blacks have a normative reason to win by 20 points, perhaps if they are requiring a dominant performance against weaker opposition to help bolster their standing. From here, I will be discussing ought and should statements in only their normative variety.

Ought and should statements are tied to normative reasons in an intimate way as an ought or should is made up of normative reasons. For instance, the sentence, “you ought to ϕ , but you don’t have any reason to ϕ ” seems confused. If someone were to say something like this to us, we would wonder whether they understood the concepts that they were invoking. We can conclude that a condition for an ought or should statement to accurately represent reality is the existence of a normative reason. Yet, it also seems apparent that not all normative reasons that I have for actions, beliefs, and attitudes can accurately be represented as things which I ought or should do. I can be in a scenario in which I want to finish an email to someone at work, with this email being of much importance, yet the fire alarm rings. It seems to me both that I have a normative reason to finish the email, but it seems to me that I have a stronger normative reason to follow the fire evacuation procedure. Looking at this scenario, it is appropriate to say that I ought to follow the fire evacuation procedure, but not that I ought to finish my email.

The difference between these two normative reasons and their ability to yield an ought or should seems to me to be decided fully on which normative reason is stronger. In the absence of the fire alarm or any other stronger normative reason, then I ought to finish that email. So, we may say that all normative reasons consist of what we *prima facie* ought to do,

believe, or feel. In other words, if we have a normative reason to ϕ then we *prima facie* should ϕ unless we find that we have another normative reason which conflicts and instead requires that we do something other than ϕ ⁹, then we ought to ϕ .

We can test this another way too by looking at a situation where an agent has a sole normative reason to do, believe, or feel something. It seems unnatural, confused, or just plain wrong for someone to appear to have a sole normative reason to ϕ but to claim it is wrong that to say that they ought to ϕ . If we therefore want to define an ought or should without using other synonymous normative terms such as ‘must’, we can call an ought or should what we have a decisive normative reason to do, believe, or feel.

How this supports normative reasons involving favouring is that ought and should statements represent favouring or commands¹⁰. When I say that you ought to ϕ , I am expressing that you are commanded to ϕ . A lot of time this is quite clear. If I am angry at someone, I may say that they should be quiet, which is a way of saying that I strongly favour them to be quiet. Other times, ought statements seem less like commands such as in the case of conditions. I may say to a child that if they want to be an actor when they grow up, they ought to take up acting classes. Here, there is no obvious favouring perhaps, but what I am expressing is that *if* they want to x then they will have a decisive normative reason to ϕ , or in other words more tied to the example, if they want to be an actor then they will be favoured to take up acting classes.

The Structure of Our Appearances of Normative Reasons

We have established that normative reasons involve favouring in some capacity, but the question now is whether normative reasons are identical with favouring, consist of some

⁹ where we are not able to do both ϕ and this other activity.

¹⁰ Put another way, Dancy (2006) states that an ought “is a directive; in a perfectly clear sense (at least for present purposes), it tells us what to do—what we should do.” (p. 40)

favouring, have appearances that cannot be untangled from the appearance of favouring, or have favouring as a condition for their existence. To help answer this question, I will search the content of my appearances of normative reasons and find what other elements exist within these appearances to see if I can draw a conclusion as to what they are.

Before getting into this, a word on the strategy here. Appearances of a thing do not always isolate the thing that they are an appearance of. For instance, I can look to my visual appearances, such as my appearance of a tree and find that the content of this appearance always involves more than just a tree. It involves the appearance of an external world – even when imagined, I am imagining the tree as a physical object in some external world – distance between myself and the tree, light, colours, etc. The point being that when trying to isolate my appearance of something, such as a tree, I cannot always fully isolate this thing in an appearance alone. It is a requirement for a visual appearance of a tree to occur that there must be light for this tree, so all appearances of a tree involve light, even though the light itself is not the tree. What this means is that though normative reason appearances always involve favouring, they may not be favouring and instead need favouring for their appearances in the same or a similar way as a tree does with light – further analysis of these appearances is required to draw a conclusion on this.

As stated, when it appears to me that I have a normative reason to ϕ , it appears to me that I am being favoured to ϕ . This favouring cannot be isolated on its own since this favouring clearly has a target, in this instance, it is me. When it appears to me that you have a normative reason to ϕ , it appears to me that *you* are favoured to ϕ . For an appearance of a normative reason to obtain, it seems that this favouring needs to be towards an agent. Trees, rocks, the weather, and time are all examples of things that are not agents that can be favoured. For instance, I may wish that time moved a bit more slowly at some points and at other times I wish it sped up, but it seems confused to me to say that time has a normative

reason to do any of these things, regardless of who is favouring it. Agents seem to be a requirement for normative reasons then on two grounds as shown by two pieces of evidence. The first is simply that when we search our appearances of normative reasons, all normative reasons relate to agents as being whom the action, belief, or attitude being favoured is for. The second is that it just seems conceptually confused to claim that non-minded things have normative reasons to do things since being able to respond to a normative reason, which only mind can do, looks to be a requirement for a normative reason to stick.

So, here we have two elements within the appearance of a normative reason, the favouring and the agent being favoured, neither of which can be separated from the content of my appearance of a normative reason, since normative reason appearances always have favouring as their content, and favouring, at least for normative reasons, is always agent directed. But this isn't the full story yet either. When an agent is favoured, they are always favoured to do something, believe something, hold an attitude towards something, or to do the converse and avoid an action, belief, or attitude.

We may then conclude that normative reasons always involve three elements in the content of their appearances, these being favouring, the agent being favoured, and the action (or belief or attitude) being favoured. Yet, this is still missing *something*. When it seems to me that I have a normative reason to ϕ , it seems to me that I am favoured to do so on some basis. For instance, it seems to be very much true that I have a reason to believe in my own existence – that I am favoured to believe that I exist. This normative reason seems to be based on something, in this case, the evidence of my own thoughts. Likewise, when it seems to me that I have a normative reason to avoid putting my hand on a hot stove top, it seems to me that I have this reason based on the fact that putting my hand on the stove top will cause a burn and that it will hurt me.

This final element is what we may refer to as the consideration in which an agent is favoured to ϕ . This consideration element may also be referred to as the ‘because’ element as when we talk about normative reasons in propositional form, we may say something akin to “you should ϕ because of C” where C stands for the consideration in which the agent is being favoured to carry out that action, belief, or attitude.

Given all of this, we can conclude that our appearances of normative reason always involve the following elements within their content:

1. *The favouring of A to ϕ .*
2. *The action, belief, or attitude being favoured in place of ϕ .*
3. *The agent who is being favoured to ϕ .*
4. *The consideration in which the action, belief, or attitude is being favoured – this is also sometimes called the ‘because’ element.*

The Identity of Normative Reasons

We have identified that the content of our appearances of normative reasons consist of a quadripartite structure in that all appearances of normative reasons contain four elements: favouring; an action, belief, or attitude being favoured; the agent being favoured; and the consideration in which the agent is being favoured towards this action belief or attitude. Given this, we can look to identify which if any of these share an identity with normative reasons.

There are two elements in our appearances that I feel we can dismiss easily. Normative reasons, whatever it is that they are, don’t seem to be either the action, belief, or attitude being favoured, nor do they seem to be the agent being favoured. Take first a normative reason to ϕ . It seems to me confused to say that ϕ is my normative reason to ϕ . We could substitute in any number of things in the place of ϕ and the sentence would read as

nonsense. “Playing video games is a normative reason to play video games” or “believing in noumena is a normative reason to believe in noumena” are just to examples to show this. Something similar seems to occur if we try and identify an agent as a normative reason. For instance, “I am my normative reason to eat healthy” seems odd and confused. Normative reasons pick out agents, but they aren’t the agents being picked out themselves. It seems more natural to say that my feeling unwell is a reason to eat well, than it is to say that I am the normative reason myself.

Perhaps then, normative reasons are the favouring element. This seems a good contender since favouring itself is a normative concept in that it directs action, beliefs, and attitudes of agents. Yet, this still doesn’t capture the meaning I have when I talk about normative reasons. When I say that agent A has a normative reason to ϕ , I am expressing more than just some favouring exists, rather I am expressing that A is favoured to ϕ based on some consideration. Narrowing normative reasons down to being purely favouring leaves out what I mean when I talk about normative reasons in that I mean that they are favoured to ϕ because of x.

Finally, we have the consideration in which the agent is favoured to ϕ . This does seem to me to be a good contender since, as said earlier, it is natural for me to say that x is a normative reason to ϕ , such as when I say that the fact that I am thirsty is a normative reason to get myself a drink. Yet, my language on this point doesn’t seem to reflect what I mean when I talk about normative reasons since it seems just as natural for me to say that the fact that I am thirsty gives me a normative reason to get myself a drink. Language then seems to be imperfect for discussing normative reasons since what I mean when I discuss normative reasons is that someone is favoured to do something based on some consideration. To conclude that normative reasons share an identity with this consideration seems to be mistaking one element of a normative reason as the whole structure.

Now that we have found what normative reasons are not, it may seem that there is nothing left for them to be. After all, they are not the favouring, the agent being favoured, whatever standing in the place of the φ symbol, nor the consideration in which the favouring is based upon. However, left is something that hasn't yet been discussed, this being the relation that stands between all these elements.

A relation is an abstract entity that stands between different particulars, entities, or properties. For instance, I stand in a familial relation to my mother and my sister, and a marriage relation with my wife. When it comes to properties, it seems that the colour green is related to my eyes in a way in which my eyes share a property relation with the colour green. All the elements that we find in our appearances of normative reasons are related to each other in a favouring relation.

If all normative reasons are favouring relations, this does not entail that all favouring relations are normative reasons. Not all favouring relations involve the same relata as the content of our appearances of normative reasons. For instance, if I favour the weather being sunny tomorrow based on needing to put washing out, the object of my favouring is not an agent and thus this favouring relation does not constitute a normative reason. Even many of the favouring relations which have the same structure as the favouring relations which appear to be normative reasons are not normative reasons themselves. Take an example when someone favours me to sabotage my own life on the basis that they dislike me. This has the same structure that my appearances of normative reasons seem to contain – we have the favouring, the agent being favoured, an action that is favoured, and a consideration that the agent is being favoured to do the action – yet it doesn't seem to me that I have a normative reason sabotage my own life.

Support For All Normative Reasons Being Favouring Relations

The claim to entertain is that all normative reasons are favouring relations, but that not all favouring relations are normative reasons. There are two pieces of evidence that show us that normative reasons are in fact favouring relations.

Firstly, normative reasons being favouring relations is supported by something I have mentioned various times so far: meaning. When we make claims about someone having a normative reason, what we mean is that some agent stands in some favouring relation to carry out an action, belief, or hold an attitude based on some consideration. This position includes all elements of a normative situation.

Another piece of evidence has also been touched on, this being the phenomenology of normative reasons. When it seems to me that I have a normative reason to ϕ , it seems that I am in some favouring relation which involves all the discussed elements. Again, if normative reasons are favouring relations, this includes all of what I experience within my appearance.

Finally, this position is parsimonious in that it doesn't ascribe any more elements into a normative situation than are already captured in our general understanding of normative reasons. When it appears to us that there is a normative reason, all the elements identified are required for our appearance of a normative reason, including a relation. By holding that normative reasons are favouring relations, we are not inserting some further element into the mix, nor are we ascribing certain strange properties to any of the elements. We are not saying that the action, belief, or attitude, just for an example, does anything strange metaphysically.

The Factualist Challenge

The position that normative reasons are favouring relations is known as Relationism. While I have provided support for this view, there is another position that this discussion would feel incomplete without at least touching on due to its overwhelming popularity. The

view in mind is factualism. This view is or has been defended by a number of recent and contemporary philosophers including Joseph Raz (1975), Thomas Scanlon (1998), Stephen Finlay (2012), Susanne Mantel (2014), and Derek Parfit (2011). In fact, factualism is such a widely held position that Jonas Olson (2018) refers to this as the “common ground” in the current meta-normative discourse (p. 2). Factualists identify normative reasons as being the considerations in which we are favoured to ϕ , finding that all of these considerations which have the identity of being a normative reason are facts. An example of a factualist take is that the *fact* that I am running late *is* a normative reason to skip breakfast this morning, or the *fact* that it is snowing outside *is* a reason to take a jacket when I leave the house.

It could be the case that all normative reasons are indeed facts and that all considerations which sit in the structure of normative reasons are facts and they themselves are normative reasons. If this were the case, the challenge from the factualist would not be a challenge about what normativity consists of, but rather which element in a normative situation is a normative reason. The view would be entirely coherent when applied to the structure I have proposed where we could agree on what is involved in normativity, but disagree on what a normative reason is – I say it is the relation while they say it is a fact, where we both agree on the existence of each of these things, we just disagree on what a normative reason is.

Yet, the factualist position that I am responding to here is not just that normative reasons are facts, but that normativity itself involves no form of favouring. Thomas Scanlon (1998) perhaps most famously provides the quote that normative reasons are facts which count in favour of actions¹¹. The problem here is that reasons are then presented as counting in favour of some action, belief, or attitude. However, counting in favour of something does

¹¹ Scanlon (1998) only refers to actions, but the same could apply to beliefs and attitudes.

not capture normativity. Something can count in favour of in non-normative ways. For instance, the fact that it has been raining all evening counts in favour of it being wet outside in the morning. This kind of counting in favour expresses a kind of prediction, but it doesn't express normativity. Likewise, a red sheet's lack of being washed counts in favour of its colour running when being put in the washing machine. This, again, expresses a prediction. But giving other senses of the term 'counts in favour of' does not rule out this being a correct articulation if the sense in which they have in mind is appropriate. So, what do they have in mind then?

When saying that facts count in favour of us ϕ -ing, what seems to be in mind here is something along the lines of some rational constraints to reality. For instance, the fact that it is cold outside counts in favour of me wearing a jacket when leaving the house. On the face of it, such an articulation is not problematic. There are many things which seem to count in favour of us to ϕ or not ϕ . Other examples include the fact that there is sufficient evidence for the existence of dinosaurs counts in favour of us believing in the existence of dinosaurs, or the fact that something is offensive is a reason to hold a negative attitude towards it. But what explains this counting in favour of notion? We require some sort of 'push' or compulsion for this to capture what we mean when we invoke counting in favour of in a normative sense, at least when talking about our own actions, belief, or attitudes that tie into our desires. Say, for instance that the cold outdoors does seem to count in favour of me wearing a jacket when I go outside. Where does this counting in favour of come from? Well, it seems to me that I favour myself being warm, not catching a cold, not looking like a weirdo, etc. This counting in favour of then can be expressed as something that is connected to my what I favour in some way.

Not all facts which seem to count in favour of an agent to ϕ involve a connection between that fact and that agent's favouring. For instance, it seems entirely appropriate for

me to claim that the fact that something will cause pain counts in favour of some agent avoiding that action, regardless of whether the agent in question seems to favour avoiding pain or not. So, we may want to ask again what this counting in favouring is then and if it is connected to some sort of favouring? We can test whether this counting in favour of does have some favouring by proposing the following: imagine a scenario in which someone says agent A has a normative reason to ϕ , but no one favours them to ϕ and ϕ -ing is not connected to anything anyone favours in any way. We would likely want to know in what way this counts in favour of the agent ϕ -ing. The explanation seems to be something that they are required to do, but as we have seen earlier, requirements are the sort of things connected to favouring, so this scenario seems to us to be confused.

We can conclude that factualists that posit a position of normative reasons being merely ‘facts that count in favour’ as being mistaken. Factualism as purely an identitarian position that is compatible with the same structure may still be on the table, but I believe enough has been shown that when we speak about normative reasons, the relationism position – i.e., the position that normative reasons are relations – is the correct view.

Meta-normative Implications

Normative reasons being favouring relations has several meta-normative implications. For starters, all subjectivist views would be lent support, after all, subjectivism holds that normative reasons have their source in a mind or minds since minds are bearers of subjective mental states. A subjectivist may hold that normative reasons are either the favouring relations that exist between the favouring of some mind to itself, the collective favouring of a culture or society towards those in the culture or society, or some divine mind or minds towards agents in the world, with each of these favouring relations also encompassing something standing in the place of ϕ , and some consideration in which the source of favouring has made that favouring.

Naturalism, the view that the source of normative reasons is some natural or thing or things which isn't a mind, would be under threat. For naturalists, if normative reasons are favouring relations, some natural thing would need to be submitted that serves as the source of these normative reasons, this thing of course being not a mind and with the ability to favour. Normally speaking, natural non-minded things that we are aware of do not favour things. My chair does not favour me to sit on it, nor does the sea actually favour trying to pull me away when I swim.

The last position is Non-naturalism. It is unclear as to whether the position of normative reasons being favouring relations negatively impacts this view as a whole. Certainly, it does negatively impact non-naturalist accounts which take the view that the source of normative reasons are properties, ala some property of goodness. If happiness, for instance, is something we have a normative reason to promote because it is good, then this has some serious issues since properties, normally speaking, are not the sorts of things which can favour things. For instance, the property of red doesn't favour me seeing redness, I just happen to see it, nor does the property of justice favour any kind of action, instead this seems to be a consideration in which I am favoured to ϕ . Non-naturalist accounts that do not posit a property but instead some kind of non-minded non-natural entity may fair better. While normally we find that things which favour do have minds, a non-naturalist position could possibly provide some entity which can favour us to ϕ .

While meta-normative implications are had by normative reasons being favouring relations, it should be noted of course that this position itself does not completely vindicate or rule out any meta-normative position. If we want to find out what the accuracy conditions are for normative reasons, we need to find what this source of favouring is by further examining our appearances. The next few sections will be dedicated to looking further at our

appearances of normative reasons to find out what these appearances tell us about the source of normative reasons.

Normative Reasons are External to Me

In this section I argue that the favouring relations constitutive of normative reasons have a favourer or favourers which are external to us. This is done via the following arguments. The first argument is that normative reasons appear to be in disaccord with my desires and motivations. The second argument is that our disagreements about normative reasons appear to be meaningful disagreements and not just reporting our desires.

A Personal Sole Source of Normative Reasons

Perhaps normative reasons for myself have a sole source and normative reasons for you have a sole source, and both of these sources are completely different. This is the position that Personal Subjectivists take on the matter.

Personal Subjectivism is the meta-normative view that normative reasons or the source of normative reasons is constituted by our internal subjective states. Subjectivists of this sort generally fall into two camps: those that believe the source of normative reasons is our motivations and those which believe our source of normative reasons is our desires. Both conceptions of normative reasons are at least plausibly consistent with normative reasons being favouring relations. Desiring is of course a type of favouring. When we desire ϕ , it means that we have ϕ . Being motivated, on the other hand, while not as synonymous with favouring as desire, does account for the probative force that normative reasons have.

Under a subjectivist view then, the favouring relations which constitute normative reasons have their source in an agent's desires or motivations and these desires or motivations are always for that same agent to do, believe, or hold an attitude towards something. This means that a normative reason for A to ϕ includes A's own favouring of A to ϕ . This isn't to deny that agents can favour other agents to ϕ , but these favouring relations would not constitute normative reasons.

If personal subjectivists are correct, then there are as many sources of normative reasons as there are minds that desire or are motivated towards things, but each mind only has a sole source of its own normative reasons since this subjectivist view holds that to have a normative reason to ϕ one must be motivated to ϕ or desire ϕ -ing.

This view is not inconsistent with agents having normative reasons to ϕ where other agents favour them to ϕ . Take a scenario in which a worker is told by their boss that they are to finish a file by 4:30pm on Friday. This agent could be deemed to have a normative reason to finish the file if they either desire or are motivated to do so or doing so is a means to achieve another desire or motivation. In this case, they may desire or be motivated simply to obey their boss, or obeying their boss serves as a means to help increase their performance and thus future salary, so they desire or are motivated to listen to their boss to achieve some other objective. What isn't serving as the source of the normative reason is the favouring of the boss or anyone or anything other than the agent's subjective states.

Normative Reasons and Their Relationship with Motivations and Desires

The close relationship that normative reasons often share with our desires and motivations provides some credence to the position of personal subjectivism. Many times, when I am weighing up what I should do, I am weighing up different motivations and desires that I have to act upon. For instance, do I choose to go to the gym or rest and read a book, both being things which I desire to do, but with these desires being in conflict.

Problems arise for the personal subjectivist position as though normative reasons do often seem to be on track with our desires and motivations, there are instances when normative reasons are shown to be off-track with these subjective states, and we seem to be favoured to ϕ when we may want to not- ϕ . Before moving on to an example of this, it should

be stated that a case against normative reasons having a personal subjectivist ontology can be provided with the following argument:

1. *If Personal Subjectivism is true, all normative reasons for A will be on-track with A's desires (or motivations).*
2. *At least one normative reason for A is off-track with A's desires (or motivations).*
3. *Therefore, Personal Subjectivism is false.*

Given the above argument, only one normative reason for an agent needs to be found that we can show is off-track with an agent's subjective states. Instead of providing just one normative reason, a thought experiment can be given to show that a sizeable chunk of a domain of normative reasons can be off-track with these states given the right conditions.

What I am referring to is reasons for belief. Take the following scenario:

- *Roger has an unusual set of desires in that Roger wants to hold as many false beliefs as he can. Roger wants to believe that the Earth is flat, that the moon is a square, that the Nazis won World War 2, that a metre is 40cm, and so forth. While he doesn't believe these as facts now, he would like to by whatever means he can.*

This scenario, however unlikely, still seems to us to be entirely coherent and possible. Perhaps Roger is an eccentric or someone who enjoys holding dissenting beliefs to others as the mere act of argumentation makes him happy. Whatever the case, according to a personal subjectivist position, if Roger desires or he is motivated to adopt false beliefs, then he has a normative reason to do so. Yet, normative reasons for belief often seem completely unrelated to an agent's desires or motivations, at least not unless we consider fringe cases.

Ordinary cases for belief seem to be based primarily or solely on an evidential basis where A has a normative reason to believe ϕ if A has sufficient evidence supporting a belief in ϕ . This seems to include beliefs that we desire to be true, “painful” beliefs, and purely mundane beliefs. An example of a belief we may want to be true is something that vindicates us or our belief systems in some way. If someone were to base much of their political thought on the government being corrupt, then if there was evidence to support the corruption of the government, the normative reason this would provide to the agent would be in line with their desires. A painful belief may be a similar scenario where there is evidence to support the government not being corrupt or, perhaps more realistically, not corrupt in an area where an agent believes they are and has based much of their political thought around. Take 9/11 “truthers” for example, some of these truthers base much of their worldview on the United States government being corrupt in a particular kind of way – that the US government would commit a false flag attack on their own soil to persuade public interest. If there was evidence against this belief, it seems that the normative reason would stand, though it may go against the truthers’ desires. Moving on to mundane beliefs and away from government corruption examples, most people in the world do not care which year Socrates died. The death year of Socrates for many is just a mundane fact. Yet, provided with evidence for the fact that Socrates died in 399BC, it seems any individual has a normative reason to adopt this belief.

The cases above talk about normative reasons and their relationships with our desires. We could substitute in similar scenarios with motivation instead of desire and have the same sort of relationship. The point is that for these reasons for belief that have been considered, while some beliefs are in accordance with our desires or motivations, other oppose them, and some are neither here nor there, the favouring relation that constitutes these normative reasons do not seem to include any reference to our beliefs and desires in their content. It

seems we are favoured to believe things on the basis of evidence, irrespective of desires or motivations.

Real Disagreements

The fact that there is disagreement over normative reasons is a mere truism of day-to-day life. Our disagreements range from the most miniscule of tasks, such as who in our household has a reason to do a particular chore, to more wide-reaching reasons such as what we should do for the homeless in our countries.

If Personal Subjectivism is the correct meta-normative position, being consistent with our appearances of normative reasons, then disagreements would centre on what an agent desires or is motivated towards. This is because Personal Subjectivism holds that it is our own desires and motivations, directed towards ourselves, that normative reasons are made of. For a disagreement of this sort to be in accordance with our appearances, conversations would be like as follows:

- *Person A – “You ought to φ ”.*
- *Person B – “No, I would rather not φ ”.*
- *Person A – “But you really do want to φ as φ -ing will bring about x which you want”.*

Here, A seems to think that φ -ing will not satisfy any of their desires, yet B at least sees φ -ing as a means to one of A’s desires. Disagreements of this sort of course do seem to me to occur. Someone may propose that an agent has a normative reason to do some community work, which that agent denies wanting to do, but then it is pointed out that doing such community work will be a means of gaining an employment position that the agent is seeking, therefore linking up to the agent’s desires.

The problem for Personal Subjectivism and its accordance with our appearances is that while disagreements such as the above are common, they are by no means the only types of disagreements that we have about normative reasons. Disagreements instead usually take the following form:

- *Person A* – “*You ought to ϕ . You just have to*”.
- *Person B* – “*No, I don’t want to ϕ* ”.
- *Person A* – “*It doesn’t matter. You really ought to ϕ* ”.

This conversation highlights a basic disagreement about what A has a decisive normative reason to do. In this conversation, B dismisses what A desires and instead suggests that A has a normative reason to do what they may not actually desire to do. Disagreements of this sort are also common, and we can highlight how many discussions take this form from more contentious day-to-day disagreements to more extreme cases.

My first example comes from a regular household in which a husband, Ross, and his wife, Rachel, have a disagreement about what Ross has reason to do. Ross has promised Rachel that he will complete some chores around the house, but he has also made plans with his friend Joey to go to a coffee shop. While Ross may have some desire to keep his promise, he primarily wants to spend time with his friend Joey. Rachel believes that despite Ross’ desire to spend time with Joey, he has made a promise and thus has a stronger normative reason to complete the chores around the house. Ross disagrees and think that he should spend time with Joey.

Looking at Rachel’s point of view, it appears that her belief in Ross having a normative reason to do chores is given by the promise that he has made. For Rachel, it would seem that no reference needs to be made to Ross’ desires. He has made a promise to her to do something and the promise he has made seems to be the basis for a stronger normative reason

than any normative reason she may believe that Ross has to spend time with Joey. It seems in this situation then that the consideration in which Ross is favoured to do chores merely has as its content promise making.

Another example is of a serial killer who desires to kill as many people as possible. Not only does he desire killing people, but he is motivated to do so, and his motivation to kill cannot be changed. He both wants to kill and will continue being motivated to do so as long as he lives. He expresses his desire and motivation to kill to a police officer when he is caught. The police officer knowing of this killer's motivations does not think he has a normative reason to kill, but instead believes that he should refrain from murder, despite not being motivated to do so. We may then ask, is the police officer correct? It seems that the police officer has a point here in that certain acts like killing seem categorically wrong and thus we have *prima facie* reason to avoid any activity that involves killing¹². It seems then, there are certain normative reasons which fall outside of our motivational or desire sets.

What The Points Raised Here Tell Us

In this section, I have shown that we have appearances of normative reasons which are in disaccord with the desires and motivations of the agents who are the subject of these favouring relations. Desires and motivations do seem to be normatively relevant features of reality, but they only seem to be normatively relevant in some situations. An agent's desire to go skydiving seems to give that agent a normative reason to do so, but their desire to believe that the Earth is flat doesn't seem to provide them with a normative reason for this belief.

Given what has already been explored in the previous chapter, the fact that normative reasons appear to be in conflict with our desires and motivations means that this rules our

¹² I don't discount that for some there are justifiable killings, such as those in self defence or in certain war time scenarios. But I think most will find that just in general, if something involves killing then it probably should be avoided.

Personal Subjectivism as a correct Meta-normative position, based on the evidence of appearances. Motivations and desires conflict with many normative reasons that appear to us as many do not have these subjective features in the content of their considerations in which they are favoured.

This all tells us that my normative reason to ϕ is not sourced in myself and my own favouring to ϕ or to achieve some ends that ϕ -ing is a means to achieve.

Normative Reasons Have a Sole Source

In this section, I will provide arguments for normative monism. Normative monism is the view that there is one single source of normativity (Reisner, 2015, p. 13). The arguments go as follows: 1. Normative reasons can oppose, but these oppositions only appear to be conflicts and not contradictions; 2. Normative conflicts are to be expected in a case of one source and multiple sources, but contradictions would likely only occur with multiple sources.

The Opposition of Normative Reasons

It seems to me that at any given time, there are multiple normative reasons applying to me all at once. Sometimes these normative reasons can be carried out without impacting the other normative reasons. For example, it may seem to me that I have both a normative reason to believe that the Earth is round as well as a normative reason to make myself a cup of tea. Taking on this belief in the roundness of the Earth will not impact my ability to make myself a cup of tea and *vice versa*. Other times, normative reasons come into opposition.

I use the term ‘opposition’ to refer to when an agent has two or more normative reasons that cannot all be satisfied due to acting on one, meaning that they are unable to act on another. Think of a situation where I am both hungry and thirsty and am in a room where there are two buttons. The button on the right will grant me with a meal to quell my hunger whereas the button on the left will grant me with a drink to quench my thirst. Once one button is pressed, neither button will work again, so I can only get a meal or a drink. Given my state, it seems that I have both a normative reason to push the button on the right and a normative reason to push the button on the left, however I cannot push both buttons, so these normative reasons oppose each other.

We are aware of normative reason oppositions in our daily lives. Many of these oppositions seem like difficult to weigh up options. Does one choose love or a career in a situation where they may only be able to have one? Whichever option a person may choose though, we would likely deem them as being reasonable since it seems that they have a strong normative reason to pick either one (if the love interest and the career are both suitably 'good' for that person, that is).

Normative reason oppositions are not always difficult decisions though. The weightings in many of these oppositions are skewed in one way that acting on one normative reason would be seen as unreasonable. Take for instance my running to the elevator to make an important meeting. It seems I have a normative reason here to hurry to be on time, but if I spot a dropped coin on the ground, perhaps a measly 10c, it seems I also have a normative reason to pick that up. Here though, my normative reason to hurry to the elevator far outweighs my normative reason to pick up the stray coin.

Types of Normative Oppositions

There are two distinct ways in which normative reasons can oppose that I believe has a bearing on what the source of normative reasons is. These opposition types are defined as below:

- *Normative Contradictions* – When one fact provides two (or more) opposing reasons to the same agent (or agents) in the same situation. If fact *F* provides agent *A* normative reasons to do both *x* and *y* in *S* (where *x* and *y* are mutually exclusive actions, beliefs, or attitudes), then fact *F* is providing *A* with contradicting reasons.
- *Normative Conflicts* - When two (or more facts) provide two (or more) opposing normative reasons to the same agent (or agents) in the same

situation. If fact F1 provides agent A a normative reason to x in situation S and fact F2 provides agent A with a normative reason to y in situation S (where x and y are mutually exclusive actions, beliefs, or attitudes), these two normative reasons are in conflict.

Normative contradictions can appear in both a strong and a weak variety. We could call something a strong normative contradiction if the same fact provided reason to both do x and not do x in identical scenarios. Something is a weak normative contradiction when the same fact calls for an agent to do both x and do y where x and y are mutually exclusive behaviours.

Like normative contradictions, normative conflicts have both a strong and weak variety. Strong normative conflicts occur when there are two (or more) normative reasons for an agent where (at least) one is to do x and (at least) one is to refrain from doing x. Something is a weak normative conflict when one fact provides a normative reason to do x and another fact provides a normative reason to do y where x and y are mutually exclusive behaviours.

The Possibility of Normative Contradictions

Both opposition types are logically possible as we can imagine cases that are at least logically coherent which involve these oppositions. Looking at normative contradictions, though, while I can list hypothetical scenarios which have the structure of these oppositions, I cannot come to a hypothetical which I actually believe is the case.

To illustrate, let's take a look at a logically possible strong contradiction: I have a normative reason to get a drink and avoid getting a drink on the basis that I am thirsty. This proposition expresses that there are two normative reasons applying to me on the basis of the same consideration. Here, when I look at this proposition, it strikes me as false. This is based

on two pieces of evidence. Firstly, there is the phenomenological evidence in that this proposition provides me with an intuition that it is incorrect. My intuition seems to tell me that normativity does not function in this manner. Secondly, I can find no appearances which I believe to be accurate which show me the existence of a strong contradiction. With both pieces of evidence in place, I have to conclude that strong normative contradictions do not actually exist.

Before moving on from discussing strong normative contradictions however, there are a couple of aspects to them which I need to further address. The first is that I could have two simultaneous appearances of strong normative contradictions. While I cannot find an example of these from my own appearances, it is not inconceivable. Yet, as suggested above, the existence of these two simultaneous appearances would prompt the intuition that at least one of these appearances is false. The strongest belief I can form around strong normative contradictions is that I am *either* favoured to ϕ or not ϕ on the basis of C, which is itself a rejection of strong contradiction.

The second is more plausible and involves me non-consciously holding two beliefs that imply a strong normative contradiction. For this I will use an example. Take a case in which an agent holds two beliefs about normative reasons which imply a strong normative contradiction. Take an example of a person who was raised as a fundamentalist Christian but later in life became an atheist who was quite anti-Christian. We can call this person Chris. During Chris conversion time, while he rejected the theism of the Bible, there was a period in which he still believed that it was a good book and that you should obey its moral code. For Chris, if the Bible says you should ϕ , you have a normative reason to ϕ . Later though, his views about the Bible became less positive and Chris then adopted the belief that if the Bible says you should ϕ , you have a normative reason not to ϕ . Chris still holds both beliefs, but he isn't consciously aware of these contradicting beliefs. In some moods, he is more prone to act

on his more positive belief about the Bible, and in other moods he is prone to act on his negative belief. The question is now, what if Chris is made aware of his contradicting belief system? How does normativity seem to him? We would expect that Chris would consider one or both of these beliefs to be false, even if he wasn't sure which one was.

Moving on to weak normative contradictions, these at least seem to fare somewhat better at first glance. Take the following scenario:

- *Agent A is upstairs in her two-story house when she hears a smash of glass and sound of voices. It seems immediately clear to A that her house has been invaded by potentially violent criminals. A tries to get a grip of herself and think about what she should do next. She believes that she has a normative reason to call the police and another normative reason to escape out the window onto the large tree outside. Knowing that she cannot do both immediately, she has to choose which option to take.*

It seems to us that A has two normative reasons to escape her house based on the same consideration, that there are intruders. This seems to suggest that weak normative contradictions exist. Yet analysing these normative reasons in this way misses something out. When we discuss facts that provide reasons in everyday discourse, or we rationalise about them in our own minds, what we tend to do is omit certain details about these facts which may play into the normative relevancy of the situation. I will refer to this as the *omission of details tendency*.

The omission of details tendency can come about due to the fact that when we discuss normative reason in our daily lives, certain aspects of the facts that normative reasons are based upon are either implied non-verbally or already understood by the agent or speakers involved in a dialogue. For instance, when I say to my wife that we need to go grocery

immediately as we have no food, what I really mean is that we have a normative reason to go grocery shopping now on the basis that we have no food and we need to eat in the near future. To properly analyse this normative reason, I need to include the full fact which has parts made of smaller facts, but in my discourse, my wife would already know or understand the full fact as she already knows about the other part. The omission of details tendency creeps into how we think about normative reasons also. When we think about a scenario like A's house, we tend to shorten the considerations which the normative reasons seem to be based upon. The fact that there are intruders *and* that she has access to her phone provides a normative reason to call the police – this is one fact with multiple parts that forms the basis for the action being favoured – and the fact that there are intruders and there is a tree outside the window which she can climb onto provides a normative reason to escape out the window. These two outlines at least better capture how normative reasons seem to apply.

The above points and the highlighting of the omission of details tendency provides convincing evidence that weak normative contradictions do not exist. Whenever I search appearances of any apparent contradictions in normative reasons, I find that they all result in only apparent contradictions at first and that they are merely oppositions of a different variety. The lack of empirical evidence based on my appearances for weak normative contradictions means that I can rule out their existence as actually occurring.

Normative contradictions in both their strong and weak variety can only result from multiple sources of normative reasons. To test this statement, we only need to look at our own favouring relations that we have with different agents in our daily lives. Never do we find that some consideration gives us a reason to favour and disfavour an agent from the same action, belief, or attitude. The phenomenological evidence for weak normative contradictions however seems at least somewhat more plausible. Let's say that my house needs a paint job. In this scenario, it does seem plausible to me that based on the fact of my house needing a

paint job, I can both favour both my house being painted green, and my house being painted blue. Here then, we have the same consideration providing the consideration in which I favour something. Yet, something is still missing from this story, and this example is another case of the omission of details tendency that we have. For both the favouring relation in which I prefer green and the favouring relation in which I favour blue, the consideration that I favour my house being painted a particular colour is both that my house needs a paint job and that I have preference for a certain colour.

Given that it has been concluded that normative contradictions don't exist, can this rule out the existence of multiple sources of normative reasons? The answer to this question is "no". A pluralistic account of the source of normative reasons could hold that all sources of normative reasons happen to not contradict one another. There could be many sources of normative reasons that just always seem to cohere with what is favoured on the same basis. This would not lock such a position into stating that all sources of normative reasons favoured identically, as some may just not favour on some grounds while others do. Instead, the lack of normative contradictions provides us with evidence for a monist account of normative sources. We are intimately aware that many of the favouring relations with the same structure as normative reasons which are sourced from us can contradict one another, so we would expect at least some normative reasons to contradict, though the lack of contradictions would not be sufficient enough on its own to rule out multiple sources.

The Possibility of Normative Conflicts

Unlike normative contradictions, normative conflicts seem to me to exist. I can think of numerous hypothetical and actual normative conflicts that seem to apply to just myself during any given day. I wake up in the morning and it seems to me that I am tired and so I should stay in bed, but as I have work, it seems that I have a stronger normative reason to get up. When I have breakfast, I can often choose between something that I enjoy and something

that is healthy and sustaining. Again, it appears that I have a normative reason to eat something I enjoy and another to eat something healthy, but these normative reasons are in conflict. Before I even really start my day, there are various normative reasons that I need to weigh up as to which are stronger so to act upon them.

Both strong and weak normative conflicts seem to me to exist. We have already looked at some weak normative conflicts with the examples just given, so I will now turn to some strong conflicts. One actual real life example which I am dealing with currently is to do with the Covid-19 pandemic. My office at home is set up for me to be able to do my job, but it seems to me that due to better resources in my workplace and chance to connect with colleagues, I have a normative reason to go to the office. Yet, because of outbreaks of the virus, it also seems to me that I have a conflicting normative reason to avoid going to the office.

My appearances of normative reasons represent that strong and weak normative conflicts exist. In fact, they are so common, they seem to be just a trivial part of my intellectual life. Furthermore, unlike normative contradictions, my intuition about normative conflicts tells me that there is nothing troubling about these conflicts existing.

Normative conflicts in both their strong and weak variety are consistent with both a sole source of normative reasons as well as a multiplicity of sources. For the latter, we only need to look at the times in which our own favouring conflicts with that of another agent's. Many of us favour different foods, drinks, activities, which then create conflicts such as when we are planning out as something as mundane as a regular Saturday night. If A wants to eat Mexican and watch sports and B wants to eat Italian and watch a film, if these two are planning an evening together then their favouring relations are in conflict. Turning to ourselves, we see something of the same sort. I really favour a piece of cake since I enjoy

cake and am hungry for something sweet, yet I also favour avoidance of eating anything sweet as I am on a diet. I really want to play video games, yet I also know that I need to write a thesis with the time that I have free. In both of these scenarios, my favouring relations are in conflict. These conflicts can also occur when it comes to some source of favouring provide favouring of agents external to themselves. Parents often want their children to know everything there is about the world, but they also want them to be sheltered from certain painful truths, this is just one example.

The apparent existence of normative conflicts then does not tell us about whether normative reasons have a sole source or a multiplicity of sources since both are consistent with conflicts. To find out which is the case, we need to explore the nature of these conflicts further.

Interdomain and Intradomain Normative Conflicts

Throughout this thesis, I have referred to normative reasons as existing across various domains, these being actions, beliefs, and attitudes. So far with examples, I have only looked at how normative reasons for actions can conflict. Yet, normative reasons do seem to conflict across the belief and attitude domains also.

Conflicts for normative reasons to believe can exist when I have two pieces of strong evidence for contradicting conclusions. Perhaps I am curious about science, and I find two scientists who provide accurate seeming research to explain the same phenomena but with both coming to wildly different conclusions. When it comes to normative conflicts for feelings, I may receive two pieces of information about an individual which give me conflicting normative reasons as to how to feel about them. For instance, I may be told that this person gives a lot to charity and thus helps a lot of people, but also be told that they are

violent and vindictive so hurt those around them. Like normative conflicts for actions, these conflicts around beliefs and attitudes don't appear to be at all troubling.

The normative conflicts examined so far are what are referred to as interdomain conflicts. Interdomain conflicts occur when a normative reason for action conflicts with another normative reason for action, a normative reason for a belief conflicts with another normative reason for belief, and a normative reason for feeling conflicts with another normative reason for feeling. This is in contrast with intradomain normative conflicts where a normative reason from one domain comes into conflict with at least one normative reason from another domain.

This person may come across some empirical evidence that says couples who have the same background as him and his fiancé have an almost 100% chance of getting divorced. Though it would seem to him that he has a normative reason to believe that his marriage would fail, it seems that he has a normative reason make himself believe that his marriage will be successful. This is, he has a normative reason to act in such a way that conflicts with his beliefs. Another example can be given here also. An agent's relative has just passed away, and while they may seem to have a normative reason to feel sad, they may have a pragmatic reason to make themselves feel less upset about the situation, for instance, to help them manage the circumstances around the loss.

The apparent existence of intradomain normative conflicts tell us something important. It tells us that normative reasons from different domains are all made of the same material thus they are consistent with the same entity or property being the source of all normative reasons.

The Phenomenology of the Strength of Normative Reasons

It has already been established earlier in this thesis that normative reasons can have varying degrees of strength. What hasn't been explored to this point is how this strength is presented to us. What does it actually mean for normative reasons to be strong? The strength of normative reasons plays a close role with normative conflicts as we tend to evaluate which normative reason to act upon based on the strength of said normative reason. Here then is a good place to address the question of what this strength is.

The strength of normative reasons is clearly not something physical. We are not physically hit or pushed or pulled in certain ways by normative reasons. They are perhaps closer to a psychological strength in that they seem to compel us to act, believe, or feel certain things. Yet normative reasons themselves are not pushing us psychologically to behave in certain ways like an involuntary drive. Instead, this compulsion should be understood as like an inclination to choose to do something.

If this strength is not a physical or psychological force that pushes us, where does this compulsion come from? One possible avenue to explore is whether this strength comes from the actions, beliefs, or feelings that are within the content of normative reasons. It seems to me that the actions, beliefs, and feelings of which I am favoured can have varying degrees of gravitas in terms of how acting on these normative reasons can impact my life and the life of those around me. For instance, it seems to me that I am favoured to have breakfast but acting or not acting on this normative reason will not impact my life in any significant way. In contrast, a normative reason to believe a member of my family is a serial killer seems to impact me much more. Yet, I can think of situations where I have what seems to be stronger normative reasons for lesser impactful actions, beliefs, and feelings than I do for those that are impactful. To give an example, it seems perfectly plausible to me that I have a strong normative reason to save \$50 per week towards my retirement as opposed to a weaker

normative reason to give all my money to charity, despite the latter having more gravitas for me and those around me. Furthermore, it seems that I have more normative reason to accept Moore's response to radical scepticism than I do Descartes' argument to external world scepticism.

Another avenue is that the strength of normative reasons could come from the considerations on the basis that we are favoured to perform actions, hold beliefs, or feel different ways. Certain considerations seem to be weightier or have more gravitas just as these action, beliefs, or attitudes seem to. Normative reasons about my life seem on the face of it far more weighty than normative reasons about trivial day-to-day activities, like if I go to work. For instance, it seems to me that the consideration in which my life is in danger due to a fire is a big consideration as to whether I finish a piece of work. In the context in which there is a fire at my workplace, it then seems that the normative reason to escape this fire is stronger than the normative reason to finish my work. However, while this seems to be the case for myself, I value the order in that my life is more important to me than a piece of work I am completing, though both have value to me. It seems completely plausible to me that there is another agent who cares more about work than his own life. Though his outlook on these considerations differs for him, the normative reasons apply just as strongly in that if we are working in the same office, we have an equally strong normative reason to escape the building based on the fire threatening our lives. The consideration then, while important to me, are not what provides the strength of the normative reason.

The final avenue I want to explore is where the probative force of normative reasons comes from, this being the favouring. Normative reasons have probative force via a favouring element in their structure. It is through here that we find strength since the compulsion that I mentioned earlier to act on a normative reason is phenomenologically similar to the compulsion that we have to act on favouring from another agent. If someone favours me to

do something, there is a certain level of compulsion that I feel to follow through with that favouring, just like if the source or sources of normativity favour me towards some behaviour.

There are two explanations as to how the favouring element of normative reasons supplies the varying degrees of strength. The first explanation is that the source of an individual favouring relation that constitutes a normative reason has a certain level of authority based on the source. To liken this to a favouring relation which we are intimately aware of in our daily lives, many of us work in environments with hierarchical structures. If my colleague favours me to do some work task, while I may feel a certain level of compulsion, I will feel less compulsion than if someone more senior within this work structure favours me to do this same task. The second explanation is that the source of any given normative reason can vary by how much this source is favouring you. That is, the favouring itself can vary from something akin to “I would like you to ϕ ” to a “you *must* ϕ ”. This is similar to how we favour other agents in our daily lives. We may have a friend who we would prefer that they take their shoes off when they come into our house, but we strongly condemn them stealing from us.

Both explanations for the varying degrees of strength that normative reasons have are not necessarily independent. It could be the case that normative reasons gain their strength through both an authority and level of favouring from any given source. The phenomenological difference between the two explanations is whether strength is provided by a *who or what* favoured me versus *how much* am I favoured.

When I search my appearances of normative reasons, I do not find myself searching for an identity of the source of the favouring that I am subject to, nor when I evaluate normative reasons do I seem to be looking at who has favoured me to ϕ . Instead, it seems that

I am weighing up how much I am being favoured to ϕ . To return to the earlier example that I used of an office fire. It seems to me that I am favoured both to complete a work task as well as conserve my life. When I am evaluating these two normative reasons, it does not seem to me that I am evaluating who the author of the favouring is, but rather the amount that I am being favoured. In this example, it seems that my life is far more valuable than any work task and it is on this basis that I am being favoured more strongly to act in ways to save my life.

I liken the phenomenological aspect to a reverse situation in which I am favouring another agent. Let's say I am playing with my cat. I want my cat to have fun and exercise, but when my cat climbs up to a high, slippery surface, I now have a much stronger favouring that she act in a way to conserve her health and avoid hurting herself. This means that I have a conflicting favouring of her to play and favouring of her to be careful, but the latter favouring is much stronger than the former. The exact same structure seems to be the case for when it appears that I have a conflicting normative reason as it appears that, like my cat may experience that she has (and outright ignores), I stand in two favouring relations of which it seems to me that the one to act on is the one with the strongest favouring.

The Parsimonious Case for a Sole Source of Normative Reasons

Given the discussion so far, we can bring in various threads to make a case for normative reasons having a sole source of normative reasons. To begin with, we discussed that normative contradictions do not seem to me to exist. This provides some evidence for normative reasons having a sole source since we would at least expect some normative contradictions. In conjunction with this, normative conflicts do seem to exist, but these are consistent both with a monist and pluralist account of the normative source. Looking at the nature of normative conflicts, we see that these can be both intradomain and interdomain conflicts, suggesting that normative reasons for actions, beliefs, and attitudes are not made of

some different material or exist in some different logical space, but rather they are the same kind of thing and can then be from the same source.

Looking at the above, we have ingredients for a case of normativity that is consistent with both a monist and pluralist account. So, how do we decide which is the case, if we can? I turn to an argument for parsimony here. Arguments for parsimony hold that if you can explain some phenomena with a smaller number of entities than a larger one, you should favour the explanation which has the smallest number of things being posited. There are of course times when parsimonious accounts seem intuitively less plausible than those that posit a much higher number of entities. Let's take an example of external world scepticism in which we doubt the existence of everything outside of our minds. We can then further say that some external world sceptic lands on two possible scenarios: the first is that their mind is the only thing in existence, and the second is that external world scepticism is false. Even though the second account posits far more in the way of just minds, let alone trees, rocks, houses, the span of space, it seems to us that this less parsimonious account is the more plausible one. Without changing track and talking about external world scepticism more fully, what I will at least say is that an argument from parsimony does not seem like it is an intuitive decider in this case, at least not without further arguments.

An example of a case in which a parsimonious account seems intuitive and rational is with a house burglary. If someone comes home to find that their house is burgled, it seems more plausible for them to posit the existence of a sole burglar in first instance. If evidence is there to support a multiplicity of burglars, then a case could be made – i.e., if there was a small window of opportunity to break in versus things stolen, or if items stolen required more than one person to carry them. If we are to avoid scepticism over whether this burglary actually occurred, then it seems that we need to posit a number of entities which are the

source of this burglary from one to infinity. The more entities that we posit, however, the more it seems that we need a ‘why’ to explain each additional entity.

The case of the favouring relations which constitute normative reasons is more like the burglary example than it is of external world scepticism. We have some scenario which requires an explanation, and we need to posit whether this is explained by one thing or many. The difference, however, is that this is more akin to positing whether or not a string of burglaries is committed by the same or multiple persons. If we do have a situation in which multiple burglaries have occurred and the evidence is consistent with both sole and multiple perpetrators – i.e., they happened in different times, across a reasonable span of space, and had relatively consistent methods of entry – then it does seem sensible to suggest that these come from a sole source in absence of other evidence.

The Coherence Case for a Sole Source of Normative Reasons

I will admit a personal bias here: cases based on parsimony, while I believe valuable, are not definitive arguments in and of themselves. The parsimony argument should be viewed as just a support for my next argument, this being an argument from coherence. Furthermore, the parsimony argument that I have looked at merely lays the suggestion on rational grounds that normative reasons could have a sole source. A further argument can be presented here which builds on this and show why a sole source account is more coherent.

When looking at the parsimony argument, one piece of evidence has purposely been left out of the equation: the phenomenology of how we weigh up normative reasons. Without this piece of evidence, the argument from parsimony is the best argument available for normative reasons having a sole source, but inclusion of this provides us with another route.

To revisit, normative reasons appear to vary in different strengths. When we weigh up which normative reason to act on, we are not looking towards *who* is favouring us but rather

how strongly we are being favoured. Take any scenario where it seems to us that we have two different conflicting normative reasons, say we have a normative reason to go to bed early and a normative reason to stay up late and do some chores. What decides our choice often comes down to which of these is a stronger normative reason, but we are of course not looking towards who favoured us to do each of these things, but rather how strongly each favouring is being presented. It may seem to us that we have a stronger reason to go to sleep early, perhaps because this is connected to some other thing which this source favours, such as us being able to cognitively perform in some test the next day as just an example. The same kind of thing occurs when we think about normative reasons that do involve us being favoured by some agent. Take an example where someone's employer asks them to cover up a scandal. It seems to them that they may have a normative reason to undertake this action, but they also seem to have a moral obligation to make the matter public. Here, while they seem to have a normative reason to listen to their employer, the favouring of the employer is not the favouring that stands in the position of the favouring which is the source of the normative reason, but instead it stands in the position of the consideration in which the agent is favoured – in other words, we may say that the agent is favoured to ϕ because their employer favours them to do so.

In all instances then when it seems we have a normative reason to ϕ and a normative reason to not ϕ , we are not searching for who the source of the favouring is, but rather how strongly we are being favoured since, as we have seen, even normative reasons which involve us being favoured by some other agent seem to be situations in which their favouring is the consideration in which we are being favoured to ϕ . If normative reasons had multiple sources then this would be in contrast with our phenomenology of normative reasons. Having some sole source is coherent with how we weigh up and evaluate normative reasons since there doesn't appear to be an issue with who the source of normative reasons is; we do not need to

weight up if the source of normative reasons is of higher or lower authority, we merely need to just account for how strongly they favour us.

The Meta-normative Implications

A sole source of normative reasons rules out cultural subjectivism, since this position posits that normativity is sourced in a plurality of minds who favour certain actions, beliefs, and attitudes. It cannot account for a single mind.

Naturalism isn't ruled out but is met with a challenge that naturalists would need to answer to. This challenge is that they need to find some sole natural entity or property which is capable of favouring as is the only one of its kind. Does such a contender exist? It isn't clear at this stage.

Some forms of both Non-naturalism and Divine Command Theory are supported, these being versions which posit some sole source of normative reasons, such as the non-natural good, and those which posit a sole divine mind as the source of normative reasons.

The Source of Normative Reasons is Not a Mind

Minds favour things, so it seems to be a natural conclusion that if there is a sole source for all normative reasons, and that this source is external to us, then some mind would be this sole source of normative reasons.

In this section, I evaluate whether normative reasons are the favouring relations of some mind that is external to us. I do this by breaking the section into three parts. The initial two parts work together to build a cumulative case against the source of normative reasons being a mind by identifying two features of a mind that we find absent in either the phenomenology of normative reasons or the content of the favouring relations which constitute them. The first feature is what I call the personal nature of normative reasons or lack thereof. Minds often favour other agents to do, believe, or hold certain attitudes towards things out of a sense of caring about that agent, something I argue that is absent from the phenomenology of normative reasons. The second feature is the variation of what is favoured over time. Minds tend to change over time, so if the source of normativity were a mind, we would expect there to be at least one example of two situations which are identical in their non-normative features that are non-identical in their normative features.

In the final part of this section, I look to provide an argument that isn't just evidence-based against the source of normative reasons being a mind, but instead shows that the source of normative reasons can't be a mind. Here, I return to the discussion around the lack in variation of normative reasons over time showing not only that normative reasons *do not* vary over time, but that they *cannot* vary over time. This argument uses intuitions from each of the three domains focussed on in this thesis to show that there are at least some normative reasons we believe cannot vary over time, and that would be able to if the source of normativity were a mind.

Normative Reasons are Not Personal

Often when we favour another agent to do something, we do so due to some personal feelings that we may have towards that agent. For instance, when my wife goes on a road trip without me, I favour her to drive to the speed limit and the road conditions, not just because these are things all people should do, but because I care about her and her safety. This favouring then, when expressed to her, has the appearance of being an expression of care. When our anniversary arrives, my wife may favour me to get her an appropriate gift. This too feels personal in a way that isn't merely because I am her husband, but because she has certain expectations of me as a person. Not all favouring seems positive. If I do not like someone and tell them that I do not want them to come to my house, they would feel this as a favouring against them as a person.

What makes favouring seem personal then is that it feels emotive, either expressing implicitly or explicitly some form of caring or dislike, or some emotions of that kind. Returning to the example of my wife favouring me to get her a present, it doesn't just seem that this favouring is bought about by some cold and calculating mind that sees anniversaries as appropriate times to receive a gift, but instead it expresses that she expects me to do this due to who I am. Had she married someone else, her expectations may have been lower and her favouring expressed differently – in other words, her favouring is done in conjunction with her feelings towards me.

Not all agent-to-agent favouring is personal, however. It is true that a great deal of the favouring relations which we stand in feel completely impersonal. For instance, many people work for a boss and be hired to perform certain tasks. When their boss favours them to complete a particular task, this favouring may not express any personal expectation or care of them but be based entirely on the role that they are hired to carry out.

The phenomenology of the favouring relations that constitute normative reasons does not seem to me to express any certain care for me in any way.

Normative Reasons that Involve Our Own Well Being Don't Seem to Be About Caring

Though many normative reasons we have involve us being favoured to do, believe, or hold certain attitudes that seem counter to any kind of interest or ends of our own, there is still a large set of normative reasons which seem to us to be normative reasons to do things based on our own well being or interests. Take for example the appearance that I have of normative reasons to act in such a way that benefit my health or avoid damage to my health. If I find that my eating or drinking habits are becoming unhealthy, I find that I have normative reasons to change this behaviour and to act differently. In addition to this, as someone who has found a great interest and passion in Philosophy, it seems to me that I have a normative reason to follow this passion and to act in ways that would advance my knowledge in this subject.

These sorts of normative reason appearances are the best candidate for normative reasons which reflect some sort of personal favouring. After all, normative reasons of this sort do seem to be the sort of thing my loved ones favour of me. Yet, normative reasons of this sort do not seem to have the same phenomenological content that the favouring relations that I stand in when I am favoured to look after myself or follow my dreams by my loved ones. The source of normativity, what ever it may be, does not seem like it has an expectation of me or care about me in any way. Rather, even when normative reasons are focussed on my wellbeing or my interests, these normative reasons feel blunt and impersonal. If I had to liken this to anything, it would be like when I eat a food that provides me health benefits. It isn't that this food itself cares about me or expresses any good will towards me, but rather the food

benefits my health just because that is what it does in the situation when a human is eating it. This is the same with favouring relations which I appear to stand in with the source of normativity.

The difference between the personal and impersonal favouring relations is that when I am favoured personally to look after my health or to follow my interests, it is because my loved ones value me as an end in itself, however when I am impersonally favoured to look after my health or follow my interests, it is because these things are an end in themselves or means to an end.

Different Situations Explain Different Normative Reasons for Different Agents

It seems to be true that different agents have different normative reasons to act, believe, and hold certain attitudes. For instance, as I write this thesis, it seems to me that I have a normative reason to hand in a Philosophy thesis. My wife, on the other hand, does not seem to have a normative reason to do such a thing. Though, my wife, it seems, will have a normative reason to feel quite happy if I get a good mark on this thesis, whereas someone who doesn't know me wouldn't seem to have this normative reason. These two very basic examples show that normative reasons appear to vary across agents as it seems that all agents have different normative reasons applying to them.

What explains this variation across agents, however, isn't some personal favouring of one agent over another, but instead down to the different situations that agents find themselves in. For instance, it may seem to you that you have a normative reason to read this thesis, perhaps to gain interest in this topic, or perhaps you are giving it a grade. Regardless of what fact your normative reason seems to be based on, if I have a normative reason to read this thesis, it will be prompted by an entirely different consideration, say because I am

writing another thesis or dissertation and need to build upon this work or know where I can improve since the general content will already be in my head.

The variability of normative reasons between yourself and I here seem to be explained entirely by a situation that we are in. The situation you are in has different facts than the situation I am in as you may be a parent and I am not, so our situations are very different. Differences between our situations likely exist further, such as you may have children where I don't, and you may not have pets where I do, so certain reasons apply to you and not me, and *vice versa*. You may be favoured then to do certain things based on the fact that you are a parent, and I based on the fact that I am a pet owner.

In some cases, it seems that we have normative reasons to put ourselves in different situations. An example of this is when someone lives in a dangerous environment, and has certain normative reasons that apply while navigating their environment, but they also seem to have a normative reason to bring themselves into a situation where they are not in this environment. Other times, however, it doesn't seem that we have any normative reason to change our situations. For example, we could involve in someone's situation that they are a brother or sister. Due to being in a sibling situation, it seems that different normative reasons would apply to them that don't apply to only children. An only child doesn't seem to have any sort of normative reason to bring themselves into a situation where they are a sibling, and it seems pretty intuitively obvious that siblings have no normative reason to make themselves an only child! This similar situation applies when we consider different cultures that agents live in. Agents may have different normative reasons based on their culture, but no normative reasons to actually leave their culture to join another.

Normative Reasons Do Not Vary Across Time

If normative reasons were the product of some mind external to us, we would expect that we would find at least one case where a scenario at time T1 was identical in all its non-normative features to a scenario in time T2 but these scenarios were non-identical in their normative features. In other words, agents 100 years ago may have normative reasons to ϕ in situation S, but agents today have a normative reason to not- ϕ in situations non-normatively identical to S.

Why we would expect this is because minds tend to change what they favour over time. To use a very basic example, when I was an infant, I used to dislike bananas, now as an adult, I eat bananas most days of the week. Other examples include my changing taste in music, as well as the actions I favour myself and those around me doing.

In contrast to this, if the source of normativity were not a mind, we would expect it to play the same role as all other apparent ‘laws’ that exist in the world that are not sourced in a mind. The laws of physics, for instance, seem to remain static across space and time. If there were a planet identical to Earth in every single way, a twin-Earth so to speak, we would expect that our experience on that Earth in terms of how we navigate the physical terrain would be identical to our own – we would expect that the gravity being identical would mean that we wouldn’t float off into space or be pulled down to the ground. Likewise, abstract laws such as those of logic and mathematics seem to us to remain static and predictable also – we expect that “one plus one equals two” is a proposition that is true no matter if we are in the beginning middle or end times of the universe.

Time Is Not Normatively Relevant

Before discussing whether normative reasons vary over time in a way in which the source has changed what it favours, what I want to discuss is the normative relevancy of time

itself. This part of the thesis can be seen as trying to show whether the below principle is holds:

- *Two situations which are identical non-normatively will be identical normatively.*

A potential problem is that this principle threatens to be trivially true. There are no two identical situations that can exist since two elements of a situation are time and place, and no two events can occur at the same time and place together. Yet, our intuition about normative reasons tells us that not all facts in any given situation are normatively relevant. Take the example in which a police officer stops two different drivers on the motorway over the course of an hour, both of whom are driving 115km/h when the limit is 100km/h. Let us imagine that both persons are strangers (though it shouldn't matter) to the police officer, neither are driving fast due to any emergency, both drivers have the same level of understanding of the road rules, both have admitted the same amount of guilt, and the same level of traffic was on the road when both incidents occurred. On the face of it, these two scenarios are identical in all their normatively relevant facts, so the actions that the police officer should take are entirely identical. Differences in the colours of the cars, the ethnicity of the drivers, what the policeman ate for breakfast, and exact time on the clock do not factor into the consideration in which the officer is favoured to take the action that he should in each case, likely to issue a ticket. We can then amend the principle to read instead as below:

- *Two situations which are identical in all their normatively relevant facts will be identical normatively.*

A fact's normative relevancy is shown by whether it fits into the consideration that normative reasons are based upon. Normative reasons can be expressed as a favouring in the same way that any favouring can be expressed, i.e., "φ because of x", where x expresses the

basis in which someone is favoured. To use a very simple example, a normative reason may be expressed as “get yourself a drink of water because you are dehydrated”. In this example, anything outside of the agent’s dehydration is not mentioned and therefore is not normatively relevant, at least for this normative reason to stick to the agent. We can say then any agent that is dehydrated is in a normatively identical situation and then has a reason to get themselves a drink, though we may caveat this by saying that an addition of a normatively relevant fact such as the agent being a body builder who needs to dehydrate themselves for a competition would either cancel this normative reason out or provide a conflicting normative reason.

If variability of normative reasons is solely explained by the variation in normatively relevant facts, then normativity itself is static and predictable. It is static in that normativity functions like an ever present, non-changing law, and predictable that if we knew all the normatively relevant facts in the past and the content of which reasons they are related in, we would be able to predict many present and future normative reasons, sans any new normatively relevant situations.

Features of an agent can be relevant to a normative reason. To use another example of my wife and I, my skin is very fair while her is of a darker complexion. What this entails is that when we go to the beach, it seems that I have a normative reason to find shade while she does not. But this feature of myself does not seem to be favoured or disfavoured by the source of normativity. I am not favoured to get shade because some source dislikes my fair skin, but rather if I do not find shade, I will become sunburnt which will cause me pain. It is this pain that seems to be the consideration at the heart of the normative reason and not some feature about myself. It just happens to be the case that due to certain features about myself, I am caused pain in different scenarios to other persons, but in scenarios where they would experience pain and I wouldn’t, it seems that normativity favours them to act differently.

What this would mean is that there is nothing special about either of us in terms of how the source of normativity favours us. If I were in your situation, the normative reasons you have would apply to me, and likewise if you were in my situation mine would apply to you. This can be tested with a thought experiment. Let's say there is a mind swap machine that will swap my mind into your body and your mind into my body. We both agree to swap minds for a day, but on the basis that you live as me and I live as you, meaning that we need to live each other's lives as the other one would, so you inherit all my responsibilities and I inherit yours. For this day in which our minds are swapped, it seems to be the case that the normative reasons I now have would have applied to you if we hadn't agreed to mind swap.

The Variability of Normative Reasons Over Time

It is common within everyday discourse to talk about agents in the past engaging in what we would consider morally reprehensible or atrocious behaviour today and yet not consider their actions as immoral or irrational as we would for one of our contemporaries engaging in such actions. I am not immune to such thinking and an example of myself perhaps best fits how this sort of thinking functions. As an undergraduate student discovering Philosophy for the first time – I had no previous experience with any philosophical thought before taking up a Philosophy elective because it seemed interesting – I was enamoured with the works and thought of Plato. Yet, despite going through a period in which I idolised Plato, there were several positions of his that I was unable to support or condone in any way – in particular, his support of slavery. Even considering this, I could not let some of his views completely ruin the experience of his Philosophy. When I compared my feelings about Plato's positions to a hypothetical person today holding the same views, I found that while I couldn't endorse Plato, I didn't feel that I could condemn him to the same manner as a contemporary of mine.

My thoughts on agents in the past is not atypical either. Many of us have a tendency to accept certain historical figures as ‘good’ despite what we would now consider horrendous positions. One may argue, and I know that this at least seems true for myself, that our ability to ignore or not judge as harshly these horrendous actions and views tends to scale in a way that we are less forgiving or harsher in our assessment the closer a figure’s lifetime is to our own. For instance, while I can tend to overlook horrible views that Plato holds, I have a much harder time enjoying Martin Heidegger due to his Nazi affiliation.

This is a good starting point to really assess whether normative reasons do vary across time since these thoughts seem to imply that normativity itself changed from earlier times to the present. The phrase we often hear is “things were different back then”. This difference could be explained by a mind being the source of normative reasons, changing what it favours just as all minds seem to change what they favour. Changes in one’s favouring can consist of favouring something and then changing one’s mind, like my previously mentioned example of liking bananas now when I previously didn’t when I was younger, and even favouring something more or less intensely, like how we may enjoy a movie but we favour it less and less or more and more each time we watch it.

The question then that we want to answer is whether normative reasons do vary across time for non-normatively identical scenarios. There are two possible ways of coaxing out an answer which I am going to appeal to.

The first is that we can look to some thought experiments about the past that prompt whether the perceived variation in normative reasons is due to some difference in normativity or whether we need to account for some non-identical part of two scenarios which is making the normative difference. What we can do is the following time travel thought experiment:

Bill lives in the 2020s and has invented a time machine which allows him to travel back in time and interact with people in the past. Bill has decided that he will travel back to the year 400BC as he would like to meet Socrates, warn him of his impending demise, and generally get to know what life was like in ancient Athens. While in 400BC, Bill is confronted with many social situations which are contrary to his own. For instance, in Bill's time he understands that he should always avoid ϕ -ing in any situation, however in 400BC, ϕ -ing is deemed as acceptable. Given that not ϕ -ing will not offend anyone, bring him any danger, or cause any harm in any way, Bill believes that it still holds true for him to not- ϕ .

Do we agree with Bill's decision on this? It seems entirely plausible to me that Bill still has a normative reason to not ϕ .

The second way of answering the question is by looking at how we judge future normative reasons. It is not uncommon for us to talk about what reasons someone will have to ϕ in the future. For instance, there may be someone who is given medical advice that in six months they will need to start ϕ -ing. Likewise, if we were told that a particular head of state was going to command his subordinates to engage in genocide, it would seem to us that their subjects would have normative reasons to disregard this leader's orders. When we make statements about future normative reasons, we make these statements as if normativity itself is static.

Normative Reasons Cannot Vary Across Time

So far, the arguments against the favouring relations that are constitutive of normative reasons being sourced in a mind have been evidential and together form a cumulative case against the theory that normative reasons have a source which is a mind.

There are some responses that need to be considered. The first response is that the source of normativity need not be a presently favouring source such as some mind giving out

command. Instead, the source of normativity could have laid out the commands that give rise to normative reasons at some earlier point in time. Our intuitions of normative reasons may be like reading the pages of a book that expresses commands to us which are the source of the favouring relations constitutive of normative reasons. This view allows that the source of normativity is a mind, and this mind may be eternally existing yet changing what it favours over time, but that the book of commands remains with dried ink and no new editions, meaning that normativity appears to us to be static and predictable.

Another response open to Divine Command Theorists is that the source of normative reasons need not function as a mind like our own. Minds seem to change over time as evidenced by our own minds changing over time. If we found out that there were beings on other planets, we would expect that their minds would function like our own. For instance, if we found them to be hostile to us or cold to us, we would expect that there would be ways to change their disfavouring of us to a positive favouring. Yet, there is nothing at least logically impossible about a mind that is set in its ways in what it favours. After all, many of the things that we favour, we favour throughout our lives. Banging my knee on a desk elicits a pain response which I find unfavourable. I have never liked this pain response, and I would be shocked if I found myself ever favouring it. The source of normative reasons could be like this, but just have a broader set of fixed favouring relations than we ourselves do, including at least all favouring relations which are constitutive of normative reasons. I think many Divine Command Theorists would not have a problem accepting this since they can cash out static favouring relations as being a sign of a perfect or a divine mind, or something on that basis. At the very least, the fixed favouring of a mind doesn't seem entirely incompatible with their view.

Each of these responses I think are at least coherent and somewhat compelling. There may in fact be other non-anticipated responses that proponents of DCT can lodge which also

explain why a divine mind that issued commands would provide static commands across time. Yet, these responses all argue one thing: the source of normative reasons doesn't change its favouring over time. The arguments, however, do not entail that the source of normative reasons *cannot* change its mind over time. If at least one normative reason was shown by our appearances to be fixed across time, not merely by a mind's favouring, but due to some intrinsic nature of the considerations themselves, or some other explanation, this would show us that the source of normative reasons was not a mind. This case against normative reasons having a mind as a source could be presented as below:

1. *If the source of normative reasons is a mind, then all normative reasons could at least possibly vary over time.*
2. *There is at least one normative reason which cannot vary over time.*
3. *Therefore, the source of normative reasons is not a mind.*

The idea behind this is of course that it seems entirely possible for a mind to favour any assortment of options. Even options which a person would never want to favour seem *possible* to be favoured. For an example, I hate genocide, I can never imagine myself becoming a genocidal maniac, but I do deem it possible that I could have favoured genocide or that I still may in some extremely farfetched and unlikely scenarios. For one, I could have been born in a different environment to a family of genocidal maniacs who passed on their desires to me. For another, I could have an accident, lose all my personality, and likewise be adopted up by a new, genocidal group who, again, pass their desires down to me. So, however unlikely certain favouring is, it is possible that some mind may favour this given the right circumstances.

Here then, I want to focus on normative reasons which our appearances show to us as being unable to vary across time. While there are likely countless more examples, I have

narrowed down focus to three examples, one for each domain that have been the focus of this thesis: actions, beliefs, and attitudes.

When it comes to actions, it seems that if an action involves killing, this provides *prima facie* reason to avoid the action. Many of us do have intuitions that there may be some justified instances of killing, such as in self-defence, being in war, killing a mad tyrant, and so forth. However, there are some instances of killing that seem to always involve normative reasons to refrain from that action in all situations which involve them. Take infanticide. Regardless of the moral grey areas that some other types of killing bring us in, infanticide seems to us to always be wrong – that is, always involve a normative reason to avoid this. If someone, regardless of their authority, asks us to commit infanticide, it seems that we have a normative reason to disobey this person's request purely on the consideration that it is infanticide.

With attitudes, we find something similar. There are some scenarios where we seem to have reason to feel positive towards in all situations which it arises. Whenever we have met our goals, we should feel happy. We could have goals which may give us a reason to feel some sadness as well, such as after organising a funeral, but we cannot shake the idea that, unless we have some horrible outlier

Finally, we have beliefs. For many philosophers, there is a link between truth and reasons for belief. Some consider something's truth as sufficient for providing a normative reason to believe the content of that truth. There are some outliers, such as when some truth may make one feel upset or cause dire consequences, but *prima facie* if something is true then we ought to believe it.

There are certain things that one may argue that minds cannot do. For instance, when it comes to belief, a case could be made that minds cannot believe in their own inexistence.

For Descartes, evidence of one's own thinking gives indubitable evidence of one's own existence, just as an example. Favouring relations have some similarities with beliefs. In the last section, we discovered that a mind cannot have strong or weak normative contradictions in its set of favouring that it is consciously aware of. These points may be dubitable for any Divine Command Theorist, however. They may find that God as a divine mind is able to function in a way which is not possible for minds like our own. This argument, however, serves little strategic value for the proponent of DCT since it only allows that a divine mind could be able to do things that our minds cannot do, but it doesn't help a scenario where a divine mind is restrained more than our own, where we can disfavour truth or accomplishing our own goals.

All of this shows that, whatever the source of normative reasons happens to be, it is not some mind that is issuing favouring to us.

The Source of Normative Reasons is Non-natural

So far in this thesis we have seen that based on our appearances, if normative reasons exist, they are favouring relations, that the source of the favouring relations is external to us, that all normative reasons have a sole source, and that this source is not some mind. This has entailed that if normative reasons do exist then they do not have an ontology that is consistent with any of the meta-normative positions that have been considered. The only other positive meta-normative theory left to consider that may account for the source of normative reasons is Non-naturalism.

The Case for Monist Non-naturalism

As posited at the start of this thesis, Non-naturalism is the view that whatever it is that provides agents normative reasons to ϕ or not ϕ is not something that is natural – i.e., that can be known by the empirical sciences nor are psychological states. The 20th century philosopher George Edward Moore can be seen as either the founder or the rediscoverer of Non-naturalism for the modern world via his influential text *Principia Ethica* (1903). It has long been held in human societies that there exist things outside of the natural world.

My argument here is not that all forms of Non-naturalism accord with our intuitions, but at least one non-naturalist account can.

Non-naturalism by its very nature is objectivist. Non-natural facts or the things that they are facts about are not identical with any of our internal subjective mental states. Even if we found that happiness, a psychological state, was the sole bearer of non-natural ‘good’ properties, we would not be able to conclude that the source of normative reasons were internal, since your psychological states would be reason providing to me, thus they would be external to myself.

Non-naturalism is typically seen as a pluralist doctrine – i.e., it posits the existence of non-natural properties which things bear. This is the position held by George Edward Moore (1903).

To ask how many things provide normative reasons would be a similar question as to asking how many things are the colour red in the way that we would be counting the number of things in the world which hold the property of goodness or the property of redness. This isn't to say that Non-naturalism is *necessarily* pluralistic as a doctrine. The only essential features that something needs to have to fit into a non-naturalist ontology is that the entities or properties posited are not identical with or reducible to parts of the natural world. That said, I am not going to bat for Non-naturalism as a view, but rather another, less defended view than pluralistic Non-naturalism can accord with our appearances of normative reasons and still be described as 'non-natural' in sense understood in meta-normativity.

Where I am going with this is, of course, that our appearances of normative reasons suggest that there is a sole source for all normative reasons. A monist account of Non-naturalism then fits this description. Monist Non-naturalism is contrasted with pluralist accounts in that instead of there being a multitude of non-natural properties or entities which provide normative reasons, there is only one¹³. It is important to understand that a monist Non-naturalism is not consistent with the view that all things of one type, say all things that are x, have the property of goodness since this implies that there are multiple instantiations of good that only things that are x have. To put it in another way, it would not be correct to say that the colour orange had a monist ontology if we lived in a universe where the only things which are orange were carrots since, even though all things that are orange are the same kind of thing, there are as many instances of the colour orange as there are carrots. What we need

¹³ Hence the prefix 'monist'.

for an actual monist account is to find some sole thing in the world that either bears the property of being good, and no other thing has it, or we need to posit the existence of some entity that is itself good.

Earlier I discussed the problems facing Naturalism when it comes to accounting for our appearances evidencing that there is a sole source of normative reasons. Naturalists, based on the arguments provided, have the burden of positing that one thing in the world is natural and is the source of normative reasons. A similar burden applies to any Non-naturalist account that normative reasons have a sole non-natural source, and this is some property of a thing. Which thing, either natural or also non-natural bears this property? The answer is not intuitively obvious.

Furthermore, the idea that the source of normative reasons is a non-natural property is further exposed by the fact that, as we found earlier, normative reasons are favouring relations and thus require a favourer. We are not aware of any properties which have as a function any ability to issue commands or favour things in any way. While I do not think being a property precludes the ability to favour, without other intuitive appeals, this property-based monist account of Non-naturalism seems to fail.

Instead, we can look at a monist account where the source of normative reasons is an entity – something I defined earlier as a thing which exists in its own right and is the sort of thing which bears properties. To return to an earlier point, entities are things such as you, me, cats, tree, rocks, fish, mountains, and so forth. A non-natural entity is like any of these things but isn't part of the natural world which means it exists outside of the causal system that the natural world involves – it doesn't bump into things, nor is it pushed or pulled by anything else. This obviously brings about the question, what then is it?

To answer this question, we can look to our use of normative language for any clues. While earlier in this thesis I avoided defining the term ‘normative’, there are several bits of language, even sans a definition, that we recognise as being normative. According to Suikkanen (2017), this includes the following predicates: “ought”, “reason to”, and “good”, among others (p. 1). Within these normative predicates, philosophers tend to divide normative terms into two categories: deontic normative terms and evaluative normative terms¹⁴. The etymology of the word ‘deontic’ comes from the Greek meaning ‘to bind’ (Tappolet, 2014, p. 2). This category includes predicates which do suggest a binding of an agent’s behaviour such as ‘ought’, ‘should’, ‘permissible’, ‘obligated’, and so forth. Evaluative normative terms on the other hand include predicates to do with judging something’s value, including terms such as ‘good’, ‘bad’, ‘better’, and ‘best’.

The division between normative terms into the two categories may seem to suggest a different ontology for whatever it is that gives rise to these terms if they are indeed used to accurately reflect the world. For instance, it does seem that when we call an action ‘good’ or say it is ‘required’, we are saying two very distinct things. However, I want to show that deontic normative statements – e.g., “A is *required* to φ ” or “it is *obligatory* that A does φ ” – can ultimately be analysed as statements that include implicit use of evaluative terms.

Beginning this claim, we can return to what was already established earlier, that ‘ought’ and ‘should’ claims, which use deontic terminology, are to be understood as claims about what an agent has decisive normative reason to do. To put it basically, if you ought to φ , this means that you had a sole normative reason applying to you now to φ , or of the multiple normative reasons applying to you, then the one to φ is the strongest. Other deontic language can also be analysed in this way. For instance, the term ‘obligatory’ can be

¹⁴ Tappolet (2014) for example (p. 1-2).

understood as referring to decisive normative reasons, or an ‘ought’, applying all the time in a given situation. An example we may invoke here is that you are obligated you have to follow custom when

By itself, this does not show that deontic normative language is implicitly appealing to evaluative terms since there are perhaps countless normative reason claims that don’t make appeal to evaluative terms. Examples of this include the following:

- *“Our students have a normative reason to believe in Socrates is older than Plato because there is sufficient biographical evidence to support it”.*
- *“Jennifer has a normative reason to see a doctor since she is in pain”.*
- *“Luke has a normative reason to feel negative towards Anakin because of what Anakin did to Ben”.*
- *“Gaius Julius Caesar had a normative reason to march on Rome since he was likely to take control of the Republic if he did so”.*
- *“Charles Darwin had a normative reason to publish his findings since they were well researched”.*

While some of the sentence structures here may feel a little unnatural to a native English speaker, all sentences represent basic statements about normative reasons that seem intuitively true¹⁵ but do not, on the face of it, seem to require evaluative language. However, in each instance here, there is a very basic question that can be asked: why?

To see how to apply this ‘why’ question, let’s take the first example used. Students are learning about Ancient Greek Philosophy and when met with sufficient evidence for Socrates being older than Plato, it seems they have a normative reason to believe as the evidence suggests. But why is this? What makes it the case that they have a normative reason

¹⁵ Provided my oversimplified historical claims are accurate.

to believe said proposition? We may say that the evidence itself is normative reason providing, or that it is true that Socrates is older than Plato and that we have normative reasons to believe true things in light of evidence for those true beliefs. In one way this seems to answer the question since both evidence and truth do seem to be the sorts of things which provide us with normative reasons, however in another sense we are missing what it is about evidence as truth that do this.

The most obvious response to this is to hold that evidence – narrowing the field of discussion here and referring to just evidence and not truth for the sake of the example – is by its nature normative reason providing. We may hold that we need no other explanation to this than merely the nature of evidence as just being reason providing. But this runs into a few problems. Firstly, when looking at the structure of normative reasons, we note that evidence takes the place of the ‘because’ relata. Things sitting in this relata are the *why* we are favoured to φ – i.e., A is favoured to believe φ *because* of evidence E. Evidence itself does not do the favouring, it is the basis in which we are favoured. To further compound this, evidence does not seem to be the sole normative reason providing ‘thing’ as we have reasons to act, believe, and hold attitudes not related to evidence. For instance, the normative reason that Jennifer has to go to the doctor is due to the pain she is in with evidence not playing a role in this case. This then violates what we had found earlier: that normative reasons have a sole source. Finally, I want to look at the discrepancy with the source of normative reasons being favouring. Evidence, and other things we recognise as tied to normative reasons, are not the sorts of things which favour. You do not see a piece of evidence that commands you or wishes you to adopt a certain belief, nor does your pain want you to go to see a doctor. It is something other than this evidence or this pain that favours you to φ on those grounds.

It seems easy then to say that evidence is not normative reason providing itself, but this still doesn’t help the fact that evidence does seem to be linked to normative reasons in

some particular way, and what this link is we still need to find. Perhaps evidence itself is not normative reason providing, but instead it bears some *property* that is normative reason providing. When questioned on the why we should believe ϕ due to evidence, we do have other terms that we can use to explain this why that are intuitive. For instance, when we look at evidential based beliefs, it seems that evidence itself is not normative reason providing, but we have a derivative reason to believe things based on evidence due to the belief that we form being a justified true belief, and justified true beliefs are *good*.

Here then, we have a normative reason being explained ultimately in evaluative terms: “A should believe ϕ because of evidence E and basing beliefs on evidence is good”. This works for many of our other day-to-day normative judgements also, such as the judgement that we have a normative reason to keep our promises since promise-keeping itself is good, and that we have a normative reason to help those in need since it brings joy and joy itself is good. These evaluative terms don’t seem to be merely placeholder language, but rather they are doing the phenomenological leg work. To demonstrate that this is the case, we can take many of our normative judgements and place them into twin earth scenarios, such as a world where agent A has a belief set all based on evidence versus a world where all things are identical but A’s twin has the same number of beliefs in their belief set, but only half of those are based on evidence. It seems to us then that one world is *better* than the other, with the term ‘better’ being used here to describe a higher abundance of good in one world in comparison to the other.

If all normative reasons are ultimately provided by what is ‘good’, then we may ask what ‘good’ actually amounts to – what sort of thing is it? Given the number of things we have said are good in joy, promise-keeping, and evidence-based beliefs, the most obvious contender here is that ‘good’ is a property that each of these things, and many, many others hold. Yet, ‘good’ being a property is of course in disaccord with what was discussed earlier:

whatever the source of normative reasons is, it must be one thing, and it must be something which favours. Properties do not favour things, and if 'good' were a property, then there would be as many sources of normative reasons as there are things in the world which are good, meaning things which bear the property of 'good'.

Despite this, I think there is a way to account for the evaluative nature of normative judgements while holding all things are good but are not good in lieu of holding the property of goodness. What I am arguing here is for a meta-normative position that may on the face of it seem peculiar but can ultimately align with our normative intuitions.

Here, instead of accounting for individual instantiations of goodness or value providing normative reasons, we can instead hold that all normative reasons are provided by the Good. This capital 'G' Good is not some property but is instead an entity like human beings, the planet Mars, ducks, statues, and the north island of New Zealand. However, unlike each of these instances of entities, the Good is not part of the natural world, instead being non-natural – we do not bump into it, nor is it linked to the causal fabricate of reality. Referring to this entity purely as the Good may sound needlessly Platonic, but it doesn't matter what we call this entity, rather we note that the Good is value itself.

Positing the existence of some entity of value does give rise to a revisitation of what led us to this discussion point: certain actions, beliefs, and attitudes such as evidence-based beliefs, promise keeping, and others are linked to normative reasons in a way in which they appear to be good. If the Good is the sole evaluative entity in the existence, then it would seem to us that we have made a mistake. But I do not think this is the case entirely. Rather than accounting for individual instantiations of goodness, we can instead account for the multiplicity of good things by the Good itself standing in a valuing relation with each individual thing which is judged to be good.

We can look at this similarly to how we value things. For instance, I value tangible things such as my house and my wife, but I also value more abstract notions such as my marriage and the knowledge I have gained from the years I have spent studying Philosophy. When I say that I value these things, I am not making any kind of metaphysical claim about any of them individually, but rather stating that these things have value to *me*. The Good seems to have a similar role in that it values at least some things which we recognise as being good, such as health and happiness¹⁶. Unlike my individual valuing, the Good is not a mind so its valuing takes the form of being more impersonal and acting like a law, and would explain the intuition that things like happiness, health, promise keeping are linked with normative reasons in every case¹⁷.

The Structure of Normativity

Positing the existence of some entity that is itself value and provides value will only be intuitively appealing if we can have it accord with our intuitions of normative reasons as we have established throughout. So, I now turn to how the Good would take the role of the source of normative reasons, as per what we have sketched earlier. What we need to do is to fit the Good into the following relational structure:

1. *The favouring element.*
2. *The action, belief, or attitude favoured.*
3. *The agent being favoured.*
4. *The consideration that the agent is being favoured*

¹⁶ It at least seems to *me* at this point of the discussion.

¹⁷ While I will not make the case for it here, this would be a good starting point for arguing the necessary link between certain parts of the world and normative reasons, though such a strong case does not need to be made here. It is enough for the sake of this argument that normative reasons happen to always coincide with certain things in the world like a law and not that they necessarily do.

On relata 1, the favouring element is the Good. The Good not only values things but it is the source of favouring as a non-minded entity that favours us to do things. Relata 2 and 4 are linked. The Good favours us to promote what it values – i.e., if happiness is good, then we are favoured to act in ways which bring about happiness. Of course, then, relata 3 is any agent who is being favoured.

Given this account, we do not need to hold that the favouring of the Good is some kind of literal command, but rather an outgrowth of the Good favouring us to bring about evaluatively good things, or avoid lowering the amount of good things in the world. For instance, if something will hurt, the Good favours us not being in pain, thus the action being favoured, non-literally but still actually, is the avoidance of certain activities.

This then puts at a place of non-parsimony as we have to enter into our ontology for an accuracy condition for a normative reason, a valuing relation separate to the favouring relation. Despite this extravagance, this structure of normative reasons can be defended on the ground that it is able to account for our intuitions about certain good states of affairs. Take the hypothetical scenario where a scientist invents a pill that when taken, we will believe that everyone around us adores us and feel immense happiness. A side effect to this pill is that it will make our bodies secrete a chemical which makes all those around us will think about us negatively. Further to this scenario, we could weigh up that though people will feel mild annoyance due to our presence, the net happiness of taking this pill would far outweigh the strength of these negative feelings. Despite the increase in the good that taking this pill would net, it doesn't seem that we have a normative reason to take it. That is, there seems to be certain states of affairs which, even though acting in ways would increase the amount of value in the world, we do not seem to be favoured to do actions which increase this value.

One may be left with a question as to why it seems that most of the time we are favoured to ϕ if it increases value, but sometimes ϕ -ing will increase value and we are not favoured to ϕ . If the source of normativity is not a mind, as is posited here, then the answer really is that this just happens to be the case. There is no other explanation that can be given that will satisfy other than this is the way things are. Explanatory chains eventually do need to come to an end, so saying that certain things are favoured when they are valued and certain things just aren't is akin to saying that water just happens to be wet, or that water is made of hydrogen and oxygen atoms.

Accounting for Badness

One response to this conclusion is that normative reasons are often given not by something being good, but by something being bad. For instance, if something causes pain, it seems we have a normative reason to avoid it since pain seems evaluatively bad. If badness is the converse of goodness, then we may ask whether to account for normative reasons, we need to account the Bad as the opposite to the Good. What this would mean is that normative reasons would be in disaccord with our intuitions as there would be two sources of normative reasons: the Good and the Bad. A more parsimonious account however is to hold that the Good, which is the source of all value is also the source of all *disvalue* as well¹⁸. Taking disvalue as being the opposite to value, it seems both parsimonious and in accord with our intuitions that we hold the Good as having both the role of valuing and disvaluing things.

This is only one response available as we may instead want to hold that negative evaluative terms such as 'bad' do not pick out disvalue, but instead pick out a situation that

¹⁸ The term 'divalue' appears in at least two philosophical sources. The first is in Joseph Raz's 2009 paper *Reasons: Practical and Adaptive* (p. 7). The second is in Wlodek Rabinowicz's 2013 paper *Value: Fitting-Attitude Account Of* (p. 1). Both don't define disvalue, but they use it in a way to reflect an opposite of value, like what negative integers are like to positive integers, and not like the standard account of undervaluing that is heard in normal speak.

either lowers things which are evaluatively good, meaning valued by the Good, or that there are other options that could be chosen which would increase goodness. To look at this in an example, if I am in pain, we may look at my pain as being bad, which means I have a normative reason to see a doctor. This badness could be explained then as either lowering goodness in that it is decreasing my health, or that I have another option that I have a normative reason to undertake, namely seeing a doctor since it increases my health.

The Good and Moral Particularism

There is one further feature of the source of normative reasons being the Good that I want to cover, which is whether or not this meta-normative view can account for some of our intuitions that support normativity having a particularist ontology.

To understand what moral particularism is, we need to understand first what a moral principle is. According to Väyrynen (2018), moral principles are invoked in statements we make such as “you should always keep your promises”, “lying is wrong”, and “causing pain is bad” (p. 1). These statements here are used to reflect that all instances where these certain things are invoked, they have a certain evaluative stance, and this evaluative stance provides normative reasons always to do things. Particularism in part denies that any such principles exist (Dancy, 2009). Instead, each particular instance of, say, promise making may have a different evaluative status and thus provide different normative reasons.

There are several instances where I think that we recognise this to be the case. While it does seem that killing is something that we have *prima facie* reason to avoid, and keeping our promises is something we have *prima facie* reasons to do, it does seem that there are instances of both where the opposite reasons may hold. Take the following scenarios:

- *The only way to stop a nuclear bomb strike on a city is to kill the person holding the remote.*

- *Someone has made a promise to keep a secret, but by keeping the secret, they will be complicit in hurting others.*

In both of these scenarios, it may seem to us either that the normative reasons to avoid killing and to keep promises are outweighed by other stronger normative reasons, or that in these particular scenarios, killing and promise keeping no longer have their *prima facie* quality.

Some may deny this particularist account, and that is fine as this does not challenge the meta-normative claims made by myself here. They may take the view that, as said, these facts are normative reason providing as they normally are but are just offset by stronger normative reasons at different times, but a challenge could be lodged against the view that the Good is the sole source of normative reasons on particularist grounds. The objection I have in mind is that the Good is supposed to function like a law of nature. Things are supposed to be valued by what they are – the Good doesn't pick and choose. Particularist accounts could be argued to be more like the favouring of a mind than that of some non-natural entity as minds tend to favour things in some circumstances but not others.

An example that is open to someone making this claim is as follows: I favour the taste of curry leaves with some ingredients but not others. I may favour curry leaves with various Indian and Asian savoury dishes, but the taste of curry leaves in my coffee or with my lemon cheesecake is not appealing. Here then we have a structure similar to what is being posited from a moral perspective in that we have some things being favoured or valued in some circumstances but not others.

This possible objection I don't believe flies. In fact, I think particularism, if anything, supports the idea of the Good being this sole source. One argument for this support is that, as we saw earlier, normativity itself appears to be static. This is no different for particular

instances of favouring. We expect normativity remains static, and we still find it predictable, even if certain principles don't hold. What I mean is that an absolute principle seems to not be true such as "stealing is wrong", but it does strike us as true to say "stealing is wrong in all circumstances except S1, S2, S3,..." where S stands for the particular scenarios where stealing is right, such as in a situation where malevolent persons horde wealth to the detriment of others.

What I am getting at here is that if particularism were true, it would be more like a scientific equation than a mind favouring. There would just be something about x that in S1 it is normative reason providing to ϕ , but in S2 it is not, or even it is reason providing to not- ϕ . We can liken this to certain elements of the world that produce the colour red except when mixed with these other specific particles. Here, there would still be no room to move. If killing is wrong almost always, but it is right in circumstances where there is a nuclear bomb, then it is *always* right in that situation. This is not an intuition the other meta-normative theories can account for.

Who Should Defend My View?

Further work will likely need to be done to conclusively show that for normative reasons to exist, there must be a non-natural source that can be referred to as the Good. Strengthening will be needed to counter issues philosophers may take with points. I will be willing to build upon this work and respond in due course. However, I now want to make a case as to why I should not be the sole defender of this view from a strategic point. Of course, I believe all philosophers interested in normativity and normative reasons more narrowly have reason to defend my view on the basis that it is true, however I do want to appeal to certain strategic elements of thinking that philosophers may have. Here, I want to give some voice as to why this version of non-naturalism should be defend by both non-naturalists *and* error theorists.

To me, it seems that non-naturalists have shunned the metaphysics from their works. In two works which have had some influence on me, Russ Shafer-Landau (2003) and Michael Huemer (2005) made book length arguments for non-naturalism being the correct meta-normative position with neither tell us what kind of a ‘thing’ the source or sources of normative reasons are¹⁹. Shafer-Landau (2003), for instance, builds much of his case on why at least some of our normative beliefs are justified, referring to a realm of moral facts that is somehow ontologically separate from natural facts (p. 79). Leaving metaphysics behind stops Non-naturalism from being subject to debunking based on its metaphysical extravagance, but it also means that Non-naturalism stays as a view that is more about what the source of normative reasons is not as opposed to a position that has its proponents trying to accurately describe reality.

¹⁹ The books in questions are Russ Shafer Landau’s *Moral Realism: A Defence* (2003) and Michael Huemer’s *Ethical Intuitionism* (2005).

Concluding Remarks

In this thesis, I hope I have shown the reader two things: that based on our collective appearances, normative reasons are not entities, properties, or something else, but that they are favouring relations, and the source of this favouring is not some collection of things, not some subjective states, or not something in the natural world, but rather an entity which is non-natural and is the source of all value.

Much work needs to be done to further establish these claims, as intuitive as they may be to some, however I do think that given the resources this endeavour will be fruitful to those who are sympathetic to a meta-normative realist account where the source of normativity is objective. In this space, despite the recent breakthroughs in meta-ethics and meta-normativity more widely, non-naturalist account still hold weight, and I think this weight is held on something very intuitive about the position as whole. The separation of the normative from the natural world has a feeling of fittingness. It feels right to us.

For future research, I hope to build about the phenomenological lines drawn and consider further objections to my own view, for the purpose of this work is, I think, an area that is still wide open, where we are not debating what we should do, believe, or feel, but rather if we have any reasons to do, believe, or feel anything at all.

References

- Alvarez, Maria, "Reasons for Action: Justification, Motivation, Explanation", *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2017/entries/reasons-just-vs-expl/>.
- Bondy, Patrick (2019). The Epistemic Norm of Inference and Non-Epistemic Reasons for Belief. *Synthese* (2):1-21.
- Dancy, Jonathan (2006). What do reasons do? In T. Horgan & M. Timmons (eds.), *Metaethics after Moore*. pp. 39-60.
- Dancy, Jonathan (2009). Moral Particularism. In Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. Center for the Study of Language and Information, Stanford University.
- Finlay, Stephen (2012). Explaining Reasons. *Deutsches Jahrbuch Fuer Philosophie* 4:112-126.
- Finlay, Stephen (2019). Defining Normativity. In Kevin Toh, David Plunkett & Scott Shapiro (eds.), *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*. Oxford University Press. pp. 62-104.
- Harrison, Gerald K. (2018). *Normative Reasons and Theism*. Cham: Palgrave MacMillan.
- Huemer, Michael (2005). *Ethical Intuitionism*. Palgrave Macmillan.
- Moore, George Edward (1903). *Principia Ethica*. Dover Publications.
- Parfit, Derek (2011). *On What Matters: Two-Volume Set*. Oxford University Press.
- Rabinowicz, Wlodek (2013). Value: fitting-attitudes account of. In Hugh LaFollette (ed.), *International Encyclopedia of Ethics*. London, U.K.: Wiley-Blackwell.

- Raz, Joseph (1975). *Practical Reason and Norms*. Hutchinson.
- Raz, Joseph, *Reasons: Practical and Adaptive* (July 2007). Oxford Legal Studies Research Paper No. 12/2007, Available at SSRN: <https://ssrn.com/abstract=999870> or <http://dx.doi.org/10.2139/ssrn.999870>
- Reisner, Andrew (2015). Normative Conflicts and the Structure of Normativity. In Iwao Hirose & Andrew Reisner (eds.), *Weighing and Reasoning: Themes from the Work of John Broome*. Oxford University Press.
- Scanlon, Thomas (1998). *What We Owe to Each Other*. Belknap Press of Harvard University Press.
- Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford University Press.
- Suikkanen, Jussi (2017). Non-Naturalism and Reference. *Journal of Ethics and Social Philosophy* 11 (2):1-24.
- Susanne Mantel (2014) No reason for identity: on the relation between motivating and normative reasons, *Philosophical Explorations*, 17:1, 49-62, DOI: [10.1080/13869795.2013.815261](https://doi.org/10.1080/13869795.2013.815261)
- Tappolet, Christine (2014). The Normativity of Evaluative Concepts. In Anne Reboul (ed.), *Mind, Values, and Metaphysics. Philosophical Essays in Honor of Kevin Mulligan, Volume 2*. pp. 39-54.
- Turri, John (2009). The ontology of epistemic reasons. *Noûs* 43 (3):490-512.
- Väyrynen, Pekka (2018). Reasons and Moral Principles. In Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. pp. 839-61.