EPISTEMIC METAPHYSICS

BY

SAMUEL JOHN ANDREWS

A thesis submitted to the University of Birmingham for the degree of DOCTOR OF PHILOSOPHY

European Research Institute
School of Philosophy, Theology and Religion
College of Arts and Law
University of Birmingham

November 2023

ABSTRACT

Metaphysical proposals are strongly influenced by epistemic considerations. Our capacity to know, explain or perceive influences theories about what is fundamental, exists or is true. These kinds of connections are pervasive as, for example, Fine (2012) and Schaffer (2009) base their metaphysics on explanation, Benacerraf (1973) and Peacocke (1999) base their metaphysics on our ability to know metaphysical truths, and Lewis (1986) and Jackson (1994) base their metaphysics on compatibility with contemporary physics. In each case, pertinent epistemic factors are used to motivate – provide reasons to accept or reject – particular metaphysical theories. This thesis examines significant epistemic factors, their connected metaphysical theories and the nature of the motivating connection exhibited between them to develop a novel approach to metaphysics: called 'Epistemic Metaphysics'.

To do this, in four papers, arguments are constructed in specific areas of metaphysics where epistemic considerations play central roles. Among others, topics covered include explanation, realism, normativity, laws, chance, dependence, empirical adequacy, and ontology. Then, by drawing inspiration from this wide array of philosophical arguments, a fine-grained account of the epistemic constraints on metaphysical inquiry can be composed; one that properly respects the intricacies of functional philosophical debate and provides general guidelines that are appropriate in any piecemeal application. Identifying the epistemic factors already employed in metaphysical arguments makes it possible to systematise the motivating connections into a set of general metaphysical maxims that constrain metaphysical theory choice: a framework for an explicitly epistemic approach to metaphysics.

ACKNOWLEDGEMENTS

I am deeply indebted to Alastair Wilson and Michael Hicks for being exceptional supervisors; their constant advice, guidance, and support has been thoroughly appreciated. I feel incredibly lucky to have had such enthusiastic mentors who made working through the stresses of a PhD manageable. Without doubt, both played a crucial part in the successful completion of this thesis for which I am sincerely thankful.

I am also appreciative to Nicholas Jones for overseeing the start of my PhD and for first directing me toward and repeatedly discussing the subject area that would later form the basis of my first publication. Gratitude is also due to past and present members of the Birmingham Philosophy faculty who played important roles in my professional development, including but not limited to Scott Sturgeon, Henry Taylor, Darragh Byrne, and Joaquim Giannotti. Thanks is also warranted to fellow graduate students Tom Baker, Will Sharp, Kathleen Murphy-Hollies, Harry Golborn, Tom Davies, Jessica Sutherland, John Murphy, and Nicholas Emmerson for their valuable discussion and moral support.

Additionally, this endeavour would have been much more challenging without the generous support of the Stapley Educational Trust, The Royal Institute of Philosophy, and the Jacobsen Foundation. The FraMEPhys research project also warrants a mention for hosting many insightful workshops in Birmingham over the years.

Lastly, it would be remiss of me not to mention my family for being a constant source of encouragement and support.

CONTENTS

LIST OF ILLUSTRATIONS

PAPER I:
EXPLANATORY METAPHYSICS

ABSTRACT

In our theories, certain elements are proposed to play certain roles. But simply being named as the element that plays a particular role does not make that element apt to play the role. As David Lewis argues, we ought to play fair when naming our elements as calling someone 'Armstrong' does not imbue them with the ability to curl heavy weights. This paper develops this constraint on theoretical elements into an explanatory principle that can serve as a criterion for metaphysical theory choice. The principle requires that there must be an explanation for how elements play their assigned roles. I show how the explanatory principle captures the force of influential arguments in the domains of laws, chances, mathematical truth, and grounding. In each case, certain views are rejected for positing elements that violate the explanatory principle. Finally, I address a potential conflict between other areas of Lewis's metaphysics and the explanatory principle before discussing the overarching meta-philosophy that results from accepting the explanatory principle: explanatory metaphysics.

OVERVIEW

In §1, examples of Lewis's fair play arguments are presented, and some general terminology is introduced. In §2, three approaches to the fair play argument are considered. The first two approaches are shown to face significant difficulties while the latter, the explanatory principle, is identified as the most fitting. In §3, the explanatory principle is shown to feature implicitly in a variety of metaphysical debates including laws, chances, mathematical truth, and grounding. Finally, in §4, a solution to the seeming conflict between other areas of Lewis's metaphysics and the explanatory principle is presented alongside a discussion of the underlying meta-philosophical proposal that results from accepting the explanatory principle.

1.0 NAMING WHATNOTS

1.1 FAIR PLAY

One recurring theme throughout David K. Lewis's metaphilosophy is the search for clarity on how certain parts of theories play their roles. This aim is best observed in what I will call the 'fair play' argument. The most renowned instance of this argument is directed against David Armstrong's (Armstrong, 1983) non-Humean proposal for laws to be understood as a second-order necessitation universal that links two first-order universals, expressed as the 'N' predicate N <F, G>. Lewis's objection is that just calling this predicate a necessitation predicate does not make the connection between F and G necessary. As he famously writes:

> 'But I say that N deserves the name of 'necessitation' only if, somehow, it really can enter into the requisite necessary connections. It can't enter into them just by bearing a name, any more than one can have mighty biceps just by being called 'Armstrong'' (Lewis, 1983a, p. 366)

Lewis employs a similar argument against conceptions of chance that include primitive chance-making relations. When we have evidence about their values, objective chances should, according to Lewis's Principal Principle (Lewis, 1994a, p. 476), be something that constrains rational credence. Again, Armstrong's (Armstrong, 1980) special chance-making relations do not meet this requirement and he fails to show how knowing these primitive relations constrains rational credence:

> 'Be my guest-posit all the primitive unHumean whatnots you like. But play fair in naming your whatnots. Don't call any alleged feature of reality "chance" unless

> you've already shown that you have something, knowledge of which could constrain rational credence.' (Lewis, 1994a, p. 484)

Lastly, when defending modal realism, Lewis objects to conceiving of possible worlds as abstract primitives because the relation between entities in the actual world and the abstract world would have to be fixed by 'magic' (Lewis, 1986, p. 182). Relying on a mysterious relation is deemed unacceptable and when considering whether 'entities' might be replaced with 'properties' to remedy this issue, Lewis once more complains that:

> 'If he tells us about some other entities he believes in that he deems suitable to play the role of properties, let him *really* tell us. It's unfair just to declare: the entities that play those roles are none other than the properties!' (Lewis, 1986, p. 189)

These arguments are versions of the same underlying methodological idea but directed toward different metaphysical proposals.

One way to capture the general sentiment is to consider theorists naming their whatnots as comparable to producers casting a play: actor names do not come into these deliberations. We could consider which thespian possesses the best combination of traits, skills, and likeness for the role and someone might interrupt our discussion to say: 'who shall be Polonius? - Let it be Polonius!' (Lewis, 1986, p. 184). The relevant matter for deciding whether someone is suited to play the role hinges not on whether the actor and character share a name but on whether the actor deserves to be called such a name in the theatrical setting; after all, if we cast Polonius as Polonius, and he plays the role with an uncanny aptness, he would not be denied an Academy Award simply because he was not christened 'Oscar'.

1.2 Elements, Roles, and Theories

Although each argument I will consider is directed against a different metaphysical entity, together they express a cohesive metaphilosophical request for specifying what roles the metaphysical entities play in their respective theories. To generalise these arguments, it will be helpful to clarify some terms. I will call the entities for which fair play is requested 'elements' and take this to be a neutral way of talking about various entities. For example, universals, properties, states of affairs, etc, are all elements. The request for fair play arises when an element is given a role to play in a theory. For example, when universals are proposed to play a role in a theory of laws, or properties proposed to play a role in a theory of chance. Importantly, Lewis is not saying that some unfairly named elements do not exist, only that they cannot adequately do the work they have been assigned.

With this brief clarification in place, we can now turn to the more substantive task of determining what it takes for an element to be named fairly; when exactly is an element suited to play a role in a theory? To answer this question, I will explore three approaches that distil the significance of the fair play argument into a single principle. The first two face significant challenges that render them unsuitable. As such, I will advocate for the last, the explanatory approach, as the preferable option. Moreover, although I will draw heavily on Lewis's work to inform a general principle, when understood as an explanatory constraint, I think it captures the force of many other arguments from the recent literature. Some examples explored here include (van Fraassen, 1989), (Benacerraf, 1973), and (Wilson, 2014). Each of these arguments makes the same kind of criticism – that elements cannot simply be given a role to play without any account of how they play that role.

2.0 A PRINCIPLE FOR FAIR PLAY

2.1 THE STRONG APPROACH

Perhaps the simplest reading of the fair play argument is as the rejection of primitive elements *tout court*. In each of Lewis's arguments doubt is raised with a primitive universal, a primitive chance-making relation, and a primitive modality relation. A common theme throughout is the presence of primitive elements and therefore, one apparent approach is to interpret the fair play argument as a wholesale dismissal of primitivism. Call this the strong fair play principle:

Strong Fair Play Principle: Primitive elements cannot feature in our theories.

Sticking with Armstrong's account of laws, on this approach Lewis's critique of Armstrong's necessitation universal is that because the necessitation universal is primitive it fails to adequately play the role of necessitating. The same applies in any other case: if an element is primitive, it cannot play its intended role.

Although the strong principle might appear intuitively plausible, it faces difficulties. Most prominently, the strong principle is explicitly not what Lewis has in mind. He thinks 'I accept it as primitive' (Lewis, 1983a, p. 352) is a perfectly legitimate way of making room for an element in a theory and utilises primitive elements in his systematic metaphysics (e.g., perfectly natural properties (Lewis, 1986, p. 63) and singletons (Lewis, 1991, p. vii)). As he writes, 'Not every account is an analysis! A system that takes certain Moorean facts as primitive, as unanalysed, cannot be accused of failing to make a place for them. It neither shirks the compulsory question nor answers it by denial. It does give an account' (Lewis, 1983a, p. 352). Either Lewis has calamitously undermined himself by both allowing for and utilising primitives immediately after excluding them in the fair play

argument (which appears to be the account endorsed by (Schaffer, 2016, p. 580)) or the strong principle is incorrect. To take the former stance, interpreting the fair play argument as a blanket ban on all primitive elements is uncharitable to Lewis.[1]

Moreover, every systematic theory ultimately requires some elements that are themselves not understood in terms of other elements. All theories need to start somewhere and posit basic elements that hold in virtue of nothing further. Setting aside esoteric cases of theories containing infinite elements or circles of dependence, where elements are forever understood in terms of other elements, almost every theory will fail to adhere to the standards outlined in the strong principle. The principle is too strong. If accepted, a serviceable distinction between permissible theories containing fairly named elements and illegitimate theories containing unfairly named elements is not forthcoming. Greater nuance is required to identify the target of the fair play argument.

2.2 The Weak Approach

With primitive elements permitted into theories, another reading of the fair play argument is as a request for clarification about what roles the elements play. For elements to be fairly named, the theory must also include axioms stipulating what roles they play. Call this the weak fair play principle:

> *Weak Fair Play Principle*: an axiom must specify what role the element plays in the theory.

This approach suggests that to play fair the role the element plays must be clearly articulated. In this case, a possible interpretation of Lewis's critique of Armstrong's

[1] Coates (Coates, forthcoming, p. 6) also suggests that Schaffer believes Lewis accepts something akin to the strong principle and similarly argues that it is unreasonably uncharitable.

necessitation universal is that he is suggesting Armstrong's theory of laws includes no axioms stating what role the necessitation universal plays. Equally so for other cases: if no axiom details what role the element plays, then it does not play that role.

The weak principle is based on the axiomatic proposals of Tooley (Tooley, 1988) and Schaffer (Schaffer, 2016). As Schaffer writes, 'one must outfit such a posit with some axioms that put it to work' (Schaffer, 2016, p. 581). Yet, where both Tooley and Schaffer propose that our theories must include role-specifying axioms, they disagree about whether this resolves the issue at hand. For the metaphysics of laws, van Fraassen outlines two closely related concerns (van Fraassen, 1989, p. 96) and Tooley aspires to resolve one while Schaffer aspires to resolve both. The first is the identification problem which can be generally understood as the task of identifying which element is the necessitation relation.[2] The second is the inference problem: the task of accounting for how the element plays the role a necessitation relation must play. Lewis's fair play argument is chiefly concerned with the latter.

Tooley maintains that the axiomatic proposal is only a solution to the identification problem. By equipping a theory with axioms that specify what work the element is doing it is clear to see which element does what work. For the inference problem, Tooley believes axioms alone are insufficient and embarks on a speculative discussion (Tooley, 1988, p. 123 ff) to resolve the inference problem (further exposition on Tooley's speculative discussion is presented on pp.13-14). Conversely, Schaffer proposes that his axiomatic solution by itself resolves both the identification and inference problem. He

[2] Specifically for van Fraassen, the identification problem is that even if granting that something alike a necessitation universal exists, there is no way of knowing in which cases it holds. The general version presented in the text is sufficient to mark the intended distinction between it and the inference problem.

claims to be unable to see what more there is to say beyond the axioms as 'one never needs to say anything further about how the posit does its stuff beyond saying that it is the business of the posit to do so' (Schaffer, 2016, p. 586). By including axioms that state what role the element plays, a full account is provided of how it plays its role. When considering Lewis's fair play argument, Schaffer believes that Lewis has simply misunderstood that one can and should simply state that the element is such that it plays the role it has been given.

Schaffer's outlook has one major problem that comes directly from Lewis and van Fraassen. They both consider responses like Schaffer's, one that aspires to resolve the problem by simply stipulating what role the element plays - Schaffer is not just giving the element a name but stipulating axioms that make the name appropriate - but deny that it offers a viable solution because this is ultimately no more helpful than just giving the element a name. As Lewis writes, 'what is not fair is just to declare: the entities that play the roles are none other than the entities that play the roles!' (Lewis, 1986, p. 186). van Fraassen goes slightly further and considers a similar line of thought to Schaffer, as we might think the minimum stipulation required pertains to what work the law does: if A necessitates B then any A is a B. But he rejects this suggestion by echoing the words of Lewis: 'its intuitive flavour may just result from the pleasing choice of words. To call the relation 'necessitation' no more guarantees the inference than being called 'Armstrong' guarantees having mighty biceps' (van Fraassen, 1989, p. 98).

Consider, for example, a regularity of F's and G's $(\forall x)\ (Fx \rightarrow Gx)$; for all cases, if Fa obtains then so too does Ga. The issue we need to resolve is that given this regularity, the law that expresses it needs to make it clear how there can never be an instance of Fa without Ga. For those invoking universals to account for the regularity, the axiomatic

solution is to stipulate an axiom that if a universal necessitating the connection between F and G obtains, then so too does that regularity:

$$N\langle F, G\rangle \rightarrow (\forall x)(Fx \rightarrow Gx)$$ (Schaffer, 2016, p. 578)

But the problem is not resolved. Aside from naming the N universal a necessitation universal, merely stipulating that a necessitation universal links together the first-order universals F and G does not make it clear how there can never be an instance of Fa without Ga. It is of no use to answer the question: why does the necessitation universal necessitate? With: because it is a necessitation universal! Likewise, for a view that posits fundamental and objective laws of the form $(\forall x)(Fx \rightarrow Gx)$ (Maudlin, 2007), the axiomatic solution is to posit that the inference holds because it is a law:

$$\text{Law}\ (\forall x)(Fx \rightarrow Gx) \rightarrow (\forall x)(Fx \rightarrow Gx)$$ (Schaffer, 2016, p. 578)

But again, the problem is not resolved. Stating that it is a law that if Fa obtains then so does Ga provides no account (even acknowledging the lawhood claim) of how being F makes something be G. It is of no use to answer the question: why does the law do the work of a law? With: because it is a law! This is just a pleasing choice of words. First outline why it is apt to be a law before calling it one.

The concern here is that appealing to role-specifying axioms is ultimately no more helpful than reiterating the element's name. Almost all theories that include some element and give it a name already implicitly affirm what role that element is supposed to play. Armstrong calls his N predicate a necessitation predicate; it is evident from being named as such that the element is supposed to play the role of necessitating. There has never been any ambiguity surrounding what role the N predicate is designed to play and restating that the N predicate necessitates because that is its role is of no help. We know

its role is to necessitate, the question is how it does that. It is not fair play to simply state that the elements that play the roles are just those the axioms stipulate as playing those roles.

The key problem with the weak principle stems from its triviality. The fair play argument is an objection to the idea that an element can play a certain role simply by being given a name. Likewise, it is also an objection to positing that an element can play a certain role by stipulating that it plays that role. Both approaches are unacceptable because they rely on asserting the connection between element and role to exist without providing a substantive (non-trivial) account of why this connection exists; they provide no further justification for why the element deserves its name. The weak approach does not require any such further justification and is, therefore, not a viable approach. The resulting principle is too weak. It allows for theories containing elements devoid of any non-trivial account of how they play their roles. A more robust fair play principle is needed that precludes this kind of triviality.

2.3 The Explanatory Approach

I will now propose an explanatory approach to the fair play argument that is of intermediate strength between the strong and weak approaches. In contrast to merely being a statement of what role the element plays, for an element to be fairly named, a further ingredient is required that *explains how* the element plays its role. Call this the explanatory principle:

> *Explanatory Principle (EP)*: how an element plays its role in a theory must be explained.

The chief development from the weak principle is that instead of merely requiring axioms that provide an account of *what* the element does, the explanatory principle requests an account of *how* the element does what it does.[3] To achieve this, an explanation of the means by which the element performs its role is required.

The explanatory principle is based chiefly on exploring the viable solutions proposed by those who endorse the fair play argument and the inference problem. Consider that Lewis proposes chance-making elements he believes warrant their names because they provide an account of 'how knowledge of frequencies and symmetries and best systems could constrain rational credence' (Lewis, 1994a, p. 484). Appealing to frequencies, symmetries, and best systems explains *how* the element plays its role and in doing so, deserves its name. The important focus here is that Lewis thinks his fair play demand is satisfied by the provision of an explanation of how the element plays its role and (following from the unacceptability of the weak approach) factors outside of what role the element has been assigned are required in this explanation.

Similarly, Van Fraassen explicitly defines his argument in terms of a request for the connection to be explained: 'it would be very disappointing if we find merely another postulate that asserts the connection to be there, and does not explain it.' (van Fraassen, 1989, p. 98). Sider also echoes this framing of the inference problem with the question being: 'How could any relation play the role that the relation of nomic necessitation must play[?]' (Sider, 1992, p. 262). Overall, interpreting the fair play requirement in terms of

[3] In a similar vein, Pagès (Pagès, 2002, p. 229) distinguishes between a validation requirement, that is satisfied by role specifying axioms, and an explanatory requirement, that requires the elements also have their 'nature explained.'

the explanatory principle remains charitable to these authors (and Lewis in particular) while capturing the key feature of their arguments: explaining how.

To be more precise, an explanation of how an element plays its assigned role can be understood as a request to describe some way the role is played. The means by which the element plays its role ought to be outlined; this typically requires elucidating something about the element that makes it apt to play that role. Such a request is satisfied by the provision of an explanation, φ, that provides an account of how the element, *a*, does the work, *b*; *a* plays the role of *b* by being such that φ. In other words, element *a* warrants being assigned the function *b* in the theory by having features φ. φ must then exhibit a connection between the element and its role such that it explains how it does its work. Importantly, the reasons present in φ cannot pertain solely to the role the element has been assigned as this does not explain how it does its work, nor a way for the element to play its role, nor an account of why the element is apt to play that role.

My usage of the term 'how' in contrast with 'why' is intended to mark a distinction between two contrastive varieties of explanation. To plainly articulate the difference, if I asked the question: 'why did you wake up at seven this morning?' An explanation for *why* you woke up at seven this morning would pertain to the reason that you woke at that particular time, say because you had a lecture at nine. Distinguish the other kind of question: 'how did you wake up at seven this morning?' It would not answer this question to repeat the same explanation; having a lecture at nine does not explain *how* you woke up at seven. Instead, to properly answer you would need to provide an account of the mechanisms that allow you to wake at that particular time, say by setting your alarm for seven. The latter is explicitly an explanation about what enables you to do something; it describes a way you were able to play that role.

Applying this to metaphysics consider, for example, Tooley's speculative theory of the mereological features of universals (Tooley, 1988, p. 123 ff). In short, when proposing that the N predicate plays the role of necessitating in our theory of laws, the explanatory principle demands an explanation of how the N predicate necessitates. Where Lewis suggests no candidate could explain how the N predicate plays its role, Tooley proposes that because the N predicate possesses the unique mereological relations exhibited by universals it is apt to play the role of necessitating. Thus, element N warrants being assigned the function of necessitation in the theory by having distinctive mereological features; those mereological features explain how the necessitation universal necessitates. Explicitly, plugged into the structure outlined here:

a: N <F, G>

b: Making the connection between Fa and Ga necessary

EP: If *a* plays role *b*, there must be an explanation *φ* of how N <F, G> makes it impossible to have Fa without Ga

φ: The unique mereological relations exhibited by universals explain how N <F, G> makes it impossible to have Fa without Ga

It appears Tooley's theory explains how the N universal plays its role. But both Sider (Sider, 1992) and Hildebrand (Hildebrand, 2013) object to Tooley's speculative theory on similar grounds. Importantly, Sider argues that Tooley's theory fails to provide 'an illuminating solution to the inference problem' (Sider, 1992, p. 274). Having features φ, those mereological relations, does not adequately explain how the N universal necessitates. This highlights a further consideration at play when explaining. It is important to evaluate the features that are proposed to explain how the element plays its

role to ensure that they actually provide an illuminating account of how the elements play their roles.

Defining what counts as adequately illuminating is an expansive and difficult task but this poses no fatal issue for the explanatory approach so long as there is some reasonable consensus on what is and is not adequately explanatory. I will present this as an understanding-based condition: for the explanation to be adequate it must be capable of generating some degree of understanding about how the element plays its role. The intent here is to set a lower bound for adequate explanation as any explanation incapable of generating understanding illuminates nothing. So, in addition to being non-trivial, the explanation must be capable of producing some degree of comprehension about how the element plays its role. Disputes will likely remain concerning what level of understanding is adequate but irrespective, this need not undermine the framework outlined here.

Overall, the explanatory principle can be summarised as the demand to non-trivially and adequately explain how elements play their roles. To break this down, there are three connected ways an element can fail and be consequently disbarred from playing its intended role. First, there must be an explanation. Fail to present anything that could sit in place of φ, offer nothing which could be used to explain how the element plays its role, and the element has not been shown to play that role. Second, the explanation must be non-trivial. Offer only a reiteration of the element's name or what role it has been named to play, and the element has not been shown to play that role. Third, the explanation must adequately explain the connection between the element and its role. Present an explanation that is incapable of generating some understanding about how the element plays its role and the element has not been shown to play that role. Providing a bad, unilluminating, explanation is no better than providing no explanation at all.

For clarity, the requirements of the explanatory approach to the fair play principle can be framed as a hierarchy whereby each condition builds upon the previous:

(1) There must be an explanation of how the element plays the role.

(2) The explanation must be non-trivial.

(3) The explanation must be capable of generating some understanding of how the element plays the role.

These conditions are connected as (3) implies (1) and (2) – the must exist a non-trivial explanation for it to be capable of generating some understanding. The first two conditions are essentially the same as I see no discernible difference between a trivial explanation and no explanation. In fact, the term trivial explanation should be seen as oxymoronic; if an explanation is trivial then it explains nothing (I will use 'trivial explanation' merely as a placeholder for attempted explanations that fail for being trivial). The third condition is notably distinct as the impact of failure is lower due to the vagueness surrounding what counts as an explanation that generates understanding. In this sense, it is problematic if there are no explanations capable of generating understanding but exceptionally detrimental if there are only trivial explanations or no explanations at all. However, to adequately explain how the element plays its role, all three conditions must be satisfied.

Having outlined the explanatory principle and the conditions required to satisfy it, I next turn to the task of demonstrating its widespread application. Concerns similar to those discussed above are already present in many different metaphysical disputes; the general line of thought conforms not just with the consciences of a few metaphysicians but underlies a ubiquitous type of consideration underscoring numerous influential

arguments. As such, to further clarify and corroborate the explanatory principle, in the next section various debates are framed in terms of a failure to adequately explain how the elements can do their intended work.

3.0 EXPLANATORY PRINCIPLE IN PRACTICE

3.1 LAWS

To begin applying the explanatory principle to other metaphysical debates, it is first worth formalising the running example of laws of nature. To achieve this, I will explicitly spell out Armstrong's Universal-based account of laws: the element, its role, and what must be explained. First, Armstrong proposes that the laws are second-order universals of the sort: N <F, G>. Second, this element, N, is proposed to play the role of being a law. Here, an important clause is that laws necessitate; part of their role in the theory is to make it such that whenever Fa occurs so must Ga. This is where the explanatory principle applies as there ought to be a good, non-trivial, explanation of how N <F, G> makes it impossible to have F*a* without G*a*. For the element N <F, G> to warrant being assigned the function of necessitating the connection between Fa and Ga in this theory of laws, some features must explain how it plays this role:

a: N <F, G>

b: Making the connection between Fa and Ga necessary

EP: If *a* plays role *b*, there must be an explanation φ of how N <F, G> makes it impossible to have Fa without Ga

φ: ?

Lewis's critique of Armstrong's account is that φ is unfulfilled (Lewis, 1983a, p. 366), and so it fails on condition (1). Another route already considered is to stipulate trivial axioms that specify that the universal's role is to necessitate (Schaffer, 2016), but this does not resolve the problem as it fails condition (2) for being trivial.

Armstrong considers another route, as he writes that the 'inexplicability of necessitation just has to be accepted' (Armstrong, 1983, p. 92). One way of reading this is as a concession that there is no plausible account of how the element plays its role. Hence, it still fails condition (1). The ruling here is easy to make as Armstrong explicitly states that there is no explanation; φ is absent. Another interpretation is to take the fact that there is no account as the explanation for how the necessitation universal does its work; that it does so primitively. Amstrong might make the case that N's primitiveness is the explanation for how it plays its role.

Regardless of whether, strictly speaking, primitive explanations count as explanations, the same problems remain. Either primitive explanation is not an explanation; if φ is simply a fact pertaining to the inexplicable nature of how *a* plays role *b*, then there is still no account of how the element plays its role. So, remains a failure of condition (1). Or the explanation of how the primitive element does its work is that it does so primitively and therefore, explanation by primitiveness is trivial. So, there is a failure of condition (2). Or if for some reason we grant non-triviality to the primitive explanation, it remains incapable of generating any understanding about how N <F, G> makes it impossible to have Fa without Ga. It fails to adequately explain anything beyond the reason why there is no explanation. So, a failure of condition (3).

All in all, the main points here are first that Armstrong's proposal contains elements devoid of any explanation for how they play their roles, and primitivism about explanation is of no help. Second, Lewis's fair play argument can be properly understood as a charge of failure to explain how the laws necessitate.

3.2 Chance

Another area of metaphysics where the explanatory principle applies is chance. To begin, just as part of a law's role is to necessitate, one of a chance-maker's roles is defined by the Principal Principle.[4] It states whatever the chance-makers are, they need to be such that if they are known they appropriately constrain our rational credence. If, for example, a rational agent knows a particular tritium atom has a 50% chance of decay in the next 12.26 years, then 'Whatever makes it true that the chance of decay is 50% must also, if known, make it rational to believe to degree 50% that decay will occur' (Lewis, 1994a, p. 476). Effectively, whatever makes chances be as they are should also (when known) make rational beliefs match those chances.

Consider another of Armstrong's proposals where, for example, the special chance-making universal N* expresses the chances that members of a class of Js are Ks: N*(J, K). Say there is a 50% chance that the next J will be a K, and due to the Principal Principle, if known, N*(J, K) must make it rational to set credence to the same degree. This is where the explanatory principle applies as there ought to be a good, non-trivial, reason for how knowing N*(J, K) makes it rational to believe there is a 50% chance the next J is a K. For the element N*(J, K) to warrant being assigned the function of making it rational to

[4] Schaffer (Schaffer, 2003) outlines two other features that are used to define chance: the Basic Chance Principle and the Humean Principle and considers several others in (Schaffer, 2007). The explanatory principle applies regardless of which is employed.

believe there is a 50% chance the next J is a K in this theory of chance, some features must explain how it plays this role:

a: N*(J, K)

b: Making the chance that the next J will be a K 50%

EP: If *a* plays role *b*, there must be an explanation *φ* of how knowing N*(J, K) makes it rational to believe there is a 50% chance the next J will be a K

φ: ?

Lewis's critique is that he cannot see how knowledge of this chance-making element has any bearing on our credence. What Lewis wants is to be convinced that the special chance-making relation between universals makes credence match objective chance without simply being told that it does. As he writes, 'You have to show me. Only if I already see—dimly will do!—how knowing the fact that N*(J, K) should make me believe to degree 50% that the next J will be a K, will I agree that the N* relation deserves the name of chancemaker that you have given it' (Lewis, 1994a, p. 485). Here Lewis argues that Armstrong's account only works if the chance-making element is trivially given the role of making chances; something precluded by condition (2) of the explanatory principle.

In contrast, Lewis suggests elements more suited to play the role of chance-makers in frequencies, symmetries, and best systems; each of which Lewis believes are capable of explaining how knowing them constrains rational credence to the appropriate degree (Lewis, 1994a, p. 484). For example, consider a (simple) frequentist account that equates chance to the frequency of outcomes, *f**, in cases of the same type: *f**(J, K):

a: *f**(J, K)

b: Making the chance that the next J will be a K 50%

EP: If *a* plays role *b*, there must be an explanation *φ* of how knowing *f**(J, K) makes it rational to believe there is a 50% chance the next J will be a K

φ: ?

Lewis claims he can see, 'dimly but well enough' (Lewis, 1994a, p. 484) how knowledge of frequencies constrains rational credence before this element is called a chance-maker. As he writes, 'we can well understand how frequencies, if known, could constrain rational credence' (Lewis, 1994a, p. 476). Yet, precisely how frequencies satisfy the fair play demand is not spelled out in any detail and, even when Lewis advances his preferred account of chance (Lewis, 1994a, pp. 485-489) he remains quiet about exactly how it satisfies the fair play demand. Hall picks up on Lewis's imprecision and argues that an element can only claim the advantage of being named fairly if one completes the substantive task of precisely explaining how it does its work: 'accomplishing this task requires real work; no dialectical advantage worth mentioning accrues to the bare claim that one can "dimly see" how it can be done' (Hall, 2004, p. 108). Moreover, Hall argues that when one does take a closer look at Lewis's approach, it fails to be adequately illuminating; only through an indistinct gaze is one able to see, but when examined in detail, 'they reveal that what one "dimly saw" was only a mirage' (Hall, 2004, p. 111).

Importantly, Hall's remarks are first an endorsement of condition (3) of the explanatory principle. When explaining what work the element does, the explanation ought to be able to generate understanding. Second, Hall argues that Lewis's account fails condition (3) because it does not include any such explanatory capacity (similar to Sider's and

Hildebrand's critiques of Tooley's speculative theory). But even if Lewis's proposal falls afoul of the same criticism Lewis lodges against others, the goal is to highlight how this debate in the metaphysics of chance can be understood as centring around the explanatory principle, for which Hall's critique of Lewis is just as well placed as Lewis's critique of Armstrong. Both demand fair play; a chance realist needs to properly explain how the chance-making element makes the chances.

3.3 Mathematical Truth

Departing from areas where Lewis is directly involved, Benacerraf (Benacerraf, 1973) outlines an influential dilemma for accounts of mathematical truth. The clause relevant to the explanatory principle pertains to Benacerraf's second condition which states that an adequate theory of mathematical truth must account for *how* we know the mathematical propositions we do:

> 'Mathematical propositions cannot make it impossible for us to know that they are satisfied. […] the concept of mathematical truth, as explicated, must fit into an over-all account of knowledge in a way that makes it intelligible how we have the mathematical knowledge that we have' (Benacerraf, 1973, p. 667)

Benacerraf defines a key condition that theories of mathematical truth must satisfy: no element can play the role of mathematical truth-maker if there is no way of accounting for how it allows us to know the truth of (at least some) mathematical propositions. As we possess knowledge about the truth of some mathematical propositions, there must be an explanation of how such knowledge could have been obtained.[5]

[5] A generalised version of this condition can be found in (Peacocke, 1999), when he outlines an integration challenge: a requirement to, 'reconcile a plausible account of what is involved in the truth of statements of

To frame this homogenously, consider the Platonist approach to mathematical truth. The Platonist proposes that mathematical propositions are made true by entities confined to the immaterial realm; the proposition 'there is at least one prime number' is made true by the existence of any platonic number that is prime: *#7*, for example. Say the proposition 'there is at least one prime number' is one we know is true then, how we possess such knowledge must be accounted for. This is where the explanatory principle applies as there ought to be a good, non-trivial, explanation for how *#7* can make the truth of 'there is at least one prime number' known. For element *#7* to warrant being assigned the function of making the proposition 'there is at least one prime number' true in this theory of mathematical truth, some features must explain how it plays this role:

a: *#7*

b: Making the proposition 'there is at least one prime number' true

EP: If *a* plays role *b*, there must be an explanation *φ* of how *#7* makes the truth of the proposition 'there is at least one prime number' knowable

φ: ?

Benacerraf's critique is that *φ* is absent (1); that no explanation is available to the Platonist because they 'depict truth conditions in terms of conditions on objects whose nature, as normally conceived, places them beyond the reach of the better understood means of human cognition' (Benacerraf, 1973, p. 667). Similarly, Field raises epistemic questions about 'how we can know anything at all about entities in a platonic realm outside of space-time and causally isolated from ourselves and everything we can perceive' (Field, 1988,

a given kind with a credible account of how we can know those statements, when we do know them' (Peacocke, 1999, p. 1).

p. 75) and also echoes Benacerraf's claim by arguing that the truth of mathematical propositions cannot be explained because they depend 'on facts involving platonic entities that reside in a realm outside of space-time' (Field, 1989, p. 68)

In reply, consider that Plato (Meno 81d) provides a speculative account (not too dissimilar in spirit to Tooley's speculative account of laws) where knowledge is merely a process of remembering the forms that are already constituent to us. Benacerraf considers this idea of *anamnesis* and frames it as a potential answer to the explain how question, '[Plato could appeal to] the concept of anamnesis at least in part to explain how, given the nature of the forms as he depicted them, one could ever have knowledge of them' (Benacerraf, 1973, p. 675). This could feature in place of φ as a potential explanation of how we know the truth of some mathematical propositions (such an explanation is also non-trivial (2)). Few would be willing to accept this obscure conception of knowledge, but criticism here would count as a failure of the less detrimental condition (3); that there's no *good* explanation.

However, expressing an explanation capable of generating understanding is not something out of reach for the Platonist. In Field's argument, an adequate explanation is a causal explanation, and Linnebo contests Field's conflation of explanation with causal explanation (Linnebo, 2006, p. 553). When proceeding with a more accommodating account of explanation, Linnebo argues that Platonists can explain how mathematical propositions are knowable and explicitly frames his argument as an answer to an 'explain how' question (Linnebo, 2006, p. 570).[6]

[6] There are some differences as Field's formulation of the argument is used instead of Benacerraf's, but the relevant dialectical similarities remain.

Ultimately, regardless of whether the Platonist can produce an explanation that generates understanding, the dialectic fits with that of enforcing an explanatory demand. All parties here agree that it would be problematic for the proposal to be devoid of an explanation and the dispute centres on what counts as an adequate explanation; again, everyone accepts the need to play fair.

3.4 Grounding

Another area of metaphysics that faces criticism for positing elements where it is unclear how they play their intended roles is grounding. Wilson (Wilson, 2014), argues that a widely applicable 'big-g' grounding relation, that is supposed to play many of the roles traditionally performed by 'small-g' relations (type identity, mereological part-whole relations, causal composition relations, set membership relations, etc.), is too coarse-grained to properly characterise the appropriate dependence relations. Wilson argues that the 'small-g' relations are capable of capturing 'diverse forms of metaphysical dependence in a genuinely explanatory and illuminating way' (Wilson, 2014, p. 539) that is not applicable to the 'big-g' relations. Her complaint is analogous to that of a failure to adhere to the explanatory principle. What Wilson wants is an explanation of how the 'big-g' grounding relation plays any of the many roles it has been assigned and without it, 'Grounding cannot do the work proponents of Grounding want it to do' (Wilson, 2014, p. 542).

One target of Wilson's argument is Rosen's grounding formulation of naturalism (Rosen, 2010), which can be expressed concisely here as the thesis that normative facts depend upon some arrangement of non-normative facts via the 'big-g' grounding relation. And a central part of Wilson's criticism is that Rosen's approach to naturalism, 'tells us almost nothing about how, exactly, normative […] goings-on stand to naturalistic goings-on'

(Wilson, 2014, p. 545). Naturalists, Wilson argues, are concerned with more than the mere fact that the normative facts depend on non-normative facts e.g., they care about whether the connection is reductive or irreducible (or irreducible but nothing over and above), whether it is objective or subjective, whether it is efficacious and if the means of bringing about the desired outcome is distinct to the normative domain.

To express this argument explicitly as an explanatory failure, consider, for example, a case where the non-normative fact [nN] and the normative fact [N] are connected via the 'big-g' grounding relation '→': [nN] → [N]. Then, for '→' to play the role of dependence relation in the theory of naturalism, there must be further details about how the dependent facts stand with respect to the facts they depend on. This is where the explanatory principle applies as there ought to be a good, non-trivial, explanation for how [Nn] stands with respect to [N]. For the element [nN] → [N] to warrant being assigned the function of making [N] depend upon [nN] in this theory of ground, some features must explain how it plays this role:

a: [nN] → [N]

b: Making [N] depend upon [nN]

EP: If *a* plays role *b*, there must be an explanation *φ* of how [N] depends on [nN]

φ: ?

There is an obvious candidate explanation for *φ* in the 'big-g' grounding relation. As this element expresses the relevant dependence relations between normative and non-normative facts, it explains how the normative facts stand with respect to the non-normative facts. So, no failures on condition (1).

In turn, Wilson's objection could be understood in one of two ways. First, one objection is that the only explanation available on Rosen's account is trivial, and so fails condition (2). Claiming how [nN] stands with respect to [N] is explained by the 'big-g' relation is tantamount to stating that how some facts depend on others is wholly explained by the fact that they stand in some dependence relation or other. More than stipulating some dependence relation is required to non-trivially explain how some facts depend on others. The other way of presenting Wilson's objection is that the only explanation available on Rosen's account is not capable of generating understanding, and so fails condition (3). The explanation tells us almost nothing about how the 'big-g' relation plays the role of connecting normative to non-normative facts. It is not an adequate explanation, nor in any way capable of generating a sufficient degree of understanding about how [nN] stands with respect to [N].

To satisfactorily explain how the normative facts stand with respect to the non-normative facts the relation must become more fine-grained. But any further exposition on the nature of the grounding relation results in a relation that is too specialised to retain its widespread application (to remain 'big'). In place, Wilson lists many examples of the various 'small-g' candidates (Wilson, 2014, pp. 546-547) that explain how the normative facts stand with respect to the non-normative in a way that generates understanding. All are compatible with a general dependence relation, but each possesses different metaphysical features that are used to explain how the normative facts depend on the non-normative facts, for which, 'each of these metaphysically informed views is genuinely illuminating' (Wilson, 2014, p. 547). All the work to be done by a dependence relation is not done by the general 'big-g' grounding relation but by any of a number of 'small-g' grounding

relations because it is possible to present an adequate explanation of how they carry out such work.

Overall, Wilson outlines a condition on dependence relations that restricts their scope. To provide a reasonable account of how some facts depend on others the relation must be understood in a sufficiently piecemeal way. This kind of condition could be generalised beyond grounding. Upholding that any element that remains underdetermined about many of its aspects will fail to adequately explain how it does any of the work it is assigned in any specific theory. An upshot here is that a solution is readily available by simply taking a stance on the underdetermined features of the element before assigning it any roles to play.

Ultimately, the take-home message of this section is that Wilson's argument can be framed around the need to properly explain how elements play their roles. She wants the 'big-g' grounding theorist to play fair by not simply calling the relation a dependence relation before showing how it can do the work of any particular dependence relation. Just as it is for cases of mathematical truth, and the need to explain how the known truths are knowable, for chances and the need to explain why we should set credence correspondingly, or for laws and the need to explain why there can be no exceptions to the regularities. In each of these domains, there are prominent arguments that can be neatly framed as a demand to explain how particular elements play their role; the explanatory principle (or something closely related to it) already features across metaphysics and is a key motivating factor in various arguments.

4.0 META-PHILOSOPHY

4.1 EXPLANATORY METAPHYSICS

So far in this paper, I have shown how a number of arguments across metaphysics can be understood in terms of an appeal to the explanatory principle. The explanatory principle captures not only a viable approach to Lewis's arguments concerning laws and chances, but a pervasive sentiment already harnessed in a range of metaphysical topics. As such, I envision the explanatory principle as a core ideal of a particular class of epistemically guided metaphysics. I will call this programme 'explanatory metaphysics':

> *Explanatory Metaphysics*: metaphysical proposals must adhere to the *explanatory principle*.

Explanatory metaphysics is the metametaphysical view that metaphysics should respect the explanatory principle: it maintains that our metaphysical theories ought to contain elements whose roles are adequately explained. This approach to metaphysics will be appealing to those who want to avoid esoteric theories where the mechanisms by which our elements work are vague; something that I have shown to already be accepted by many.

Explanatory metaphysicians question whether universals can play the role of a law on the basis that nothing explains how the lofty notion of a universal necessitates. They question whether the special chance-making relation can play the role of chance-maker on the basis that the only explanation comes from labelling the special relation as a chance-making one. They question whether platonic entities can play the role of mathematical truth-makers on the basis that there is no explanation for how we can come to know any mathematical truths. They question whether the 'big-g' grounding relation can play any

of the myriad of roles assigned to it on the basis that it is too coarse-grained to obtain a good explanation for how it plays any one particular role. In each of these cases, there is a homogenous structure to the arguments, requesting an explanation of how, and rejecting trivial and bad answers. Each of these arguments presupposes explanatory metaphysics; for a metaphysical proposal to be viable, any elements that are given a role to play must also be accompanied by an explanation of how they play that role.

Although generally standardised, there can be variation when determining what is required of the element to play any particular role; whether laws must necessitate, chances must constrain rational credence, mathematical truth-makers must be knowable, or the dependence relation must not be underdetermined. In the case of chance, for example, it is a further and distinct question if we ought to accept the Principal Principle, but the explanatory challenge presented here is indifferent. Whatever ends up being used to define an adequate theory of chance, there must be an explanation for how the element conforms to it. The same goes for each other case so far discussed; it does not matter if laws must be understood as exhibiting modal force, if mathematical truths must be knowable, or if dependence relations must be fine-grained. What matters is that once one agrees to a particular condition being employed to define the adequacy of a theory, it is explained how any proposed element satisfies that condition.

4.2 Systematic Conflicts

From here, Explanatory metaphysicians should apply the explanatory principle to any and every area; to protest certain views and ratify others. The same methodological discussions are warranted universally, to sanction views that clearly define how their elements do their work and discard those that do not. The principle is topic-neutral: it uses only logical vocabulary along with the notions of theory, theoretical role, and explanation

– and we are capable of presenting theories and explanations within any domain. So, those who accept the explanatory principle in one domain ought to be receptive to its implementation in all. It would be bizarre to accept the need for explanation in one area but remain silent in others.

So, for someone like Lewis – someone who is undeniably an explanatory metaphysician – any concurrently proposed metaphysical theories ought to be compatible with the explanatory principle. Lewis outright describes himself as a systematic philosopher, building a total package of concurrently entwined metaphysical proposals (Lewis, 1983b, p. ix) with all of his metaphysical views intended to be compatible with each other. But Lewis is also receptive to being proven wrong by a 'sceptical reader' (Lewis, 1983b, p. x) who reveals some parts of the overall picture to be incompatible with others. My aim in this section is to defend Lewis from a sceptical reader; in particular, from the scepticism presented by Schaffer (Schaffer, 2016, pp. 580, footnote 2). This section is an expansion of the issue initially raised in §2.1 of this paper, whereby the strong principle conflicts with Lewis's own metaphysical proposals. If this is still the case in the explanatory reading, and areas of Lewis's metaphysics fail to adhere to the explanatory principle, then this would add credence to Schaffer's dismissal of the fair play argument.

First, it should be noted that where I assert that Schaffer is uncharitable to Lewis in interpreting the 'fair play' argument as merely a rejection of primitives, I have gone to lengths to define a charitable interpretation of Lewis's argument. There are still viable questions to unearth here but they arise after an extensive exegesis of Lewis's work (and not from the brief eisegesis provided by Schaffer). Moreover, Schaffer uses the potential conflict as motivation to reject the fair play argument. If the fair play argument undermines another part of Lewis's metaphysics, then the fair play argument must be

incorrect. But there are other perfectly viable options to resolve this kind of inconsistency; notably, by retaining a commitment to the explanatory principle and discarding the metaphysics that it is incompatible with.

I have already briefly touched on one area of contention in Lewis's proposed theory of chance. Hall (Hall, 2004) argues that Lewis's own theory of chance does not meet the standards it asks of others. But it is not a conclusive result, and one could argue that appeal to elements like frequencies, best systems, and symmetries offer far more resources to aptly explain how they do their work over other primitive metaphysical whatnots. In particular, Hall's critique is that although Lewis does attempt an explanation it is not good enough (3) and there is a certain degree of ambiguity surrounding what is a sufficient explanation. Furthermore, when faced with this kind of conflict it is also perfectly feasible to retain the explanatory principle and revise the theory of chance such that it is better suited to explain how the proposed element does its work. This is the key point I intend to argue for in the remainder of this section: *pace* Schaffer, when there is a conflict between specific elements of Lewisian metaphysics and the explanatory principle, it is the former, not the latter that should be dispensed.

To see how Lewis goes to lengths to ensure that his metaphysical proposals do not conflict with the explanatory principle, consider Lewis's stance on the singleton element in set theory. Unlike multi-membered sets, which can be described simply as the formation of many entities into one, a set containing only one member cannot be described in such a way. Instead, defining the singleton set is difficult and can only be understood by an obscure relation to its only member. Lewis is open about singletons being primitive within his system and about the epistemic deficit that this creates: 'We were told nothing about the nature of the singletons, and nothing about the nature of their relation to their

elements' (Lewis, 1991, p. 31). But whereas an explanatory metaphysician takes no issue with the existence of primitive elements, in contrast, if singletons are proposed to play a role in set theory, there must be an explanation of how they play that role. Lewis's approach might seem to lack any account of how the singleton element could play its role without simply being named as playing that role.

Lewis has the foresight to recognise that his stance here might lead to accusations of a double standard; that he is accusing others of not playing fair while, when constructing his preferred theory of sets, also not playing fair. van Inwagen (van Inwagen, 1986) outlines an argument of this kind explicitly against Lewis's fair play objection to magical ersatzism. In response to Lewis's objection that magical ersatzism relies on unfairly named relations between entities in the actual and abstract worlds (Lewis, 1986, p. 182), van Inwagen argues that Lewis relies on an exactly analogous relation when he accepts the notion of a singleton in set theory. This kind of relation, accepted by Lewis, is the same as which Lewis critiques the ersatz theorists for accepting in their theory of modality.

Lewis is aware of van Inwagen's complaint and concedes to neither situation being satisfactory:

> 'It's a nasty predicament to claim that you somehow understand a primitive notion, although you have no idea how you could possibly understand it. That's the predicament I'm in when I accept the notion of singleton, and that's the predicament I claim others are in if they accept the alleged notion of realization of abstract simple possibilities' (Lewis, 1991, p. 36)

The key problem is that for both elements there is no way of understanding how the element could play its role if we understand nothing about the element outside of what role the element is supposed to play. Lewis frames the relevant issue in terms of there being 'no independent grasp on what sort of relation we were dealing with' (Lewis, 1991, p. 37). And the problem here is that this is precisely the same criticism expressed in the fair play arguments. When presented as an explanatory deficit the worry is that there is no explanation for how the element plays its role outside of appealing to the element itself and what role it is supposed to play.

To again frame this homogenously consider, for example, the singleton set: {Socrates}. The role of this set is to bear a membership relation to its only member Socrates. This is where the explanatory principle applies as there ought to be a good, non-trivial, explanation for how {Socrates} relates to Socrates. In this theory of sets, for the element {Socrates} to warrant being assigned the function of making Socrates a member of the singleton set {Socrates} some features must explain how it plays this role:

a: {Socrates}

b: Making Socrates a member of the singleton set {Socrates}

EP: If *a* plays role *b*, there must be an explanation φ of how {Socrates} is made to contain Socrates

φ: ?

Lewis needs to explain how his element of a singleton plays its role in his theory of sets. Yet, even when Lewis expands his discussion on set theory (Lewis, 1993), he concedes that 'because we know so little about the singletons, we are ill-placed even to begin to understand the relation of a thing to its singleton' (Lewis, 1993, p. 217). But crucially,

when considering potential answers of the form: the singleton {Socrates} has Socrates as its member because it has the property of being the singleton of Socrates, Lewis remarks,

> ‘That's just to go in a circle. We've named a property; but all we know about the property that bears this name is that it's the property, we know not what, that distinguishes the singleton of x from all other singletons.’ (Lewis, 1993, p. 217)

Lewis consistently upholds the need to play fair and the impermissibility of simply naming the element as playing that role without providing an explanation for how it plays that role. So, Lewis concedes, when faced with the failure of set theory, that an account of set theory that includes individual singleton set-member relationships is untenable. As a result, he moves to adopt a structuralist, nominalist, account of set theory that has no primitives in either the general relation of set membership or the singleton relation of set membership (Lewis, 1993, p. 221). This resolves the explanatory problem because by adopting a structuralist account of singleton functions, ‘we may bid farewell to all our worries about how we understand 'the' singleton function. That which is not there need not be understood’ (Lewis, 1993, p. 221).

Summing up: Lewis faces a dilemma between his initially proposed theory of sets and the explanatory principle that he resolves by rejecting the initial, unfair, theory in favour of another theory that avoids this difficulty. When faced with the incompatibility between fair play and singletons, Lewis retains the former meta-philosophical standard and jettisons the latter first-order theory from his systematic worldview. There might be other places in the Lewisian metaphysical system where this kind of conflict arises. Lewis’s commitment to perfectly natural properties (Lewis, 1986, p. 63) is one area that comes to mind but even there Lewis does offer some account of how this primitive element is apt

to play at least some of the roles it has been given by appealing to objective similarity (Lewis, 1983a, p. 365).[7] If that explanation turns out to not be sufficient, and if we accept that Lewis would go to the same lengths to revise his theory of natural properties as he was for his theory of sets, then natural properties ought (if possible) to be discarded too.

Luckily, there are many alternative accounts of Lewis's proposals that try to alleviate or discard entirely a commitment to natural properties. For example, (Meacham & Eddon, 2015) outline a number of plausible versions of Lewis's best-system accounts of laws that dispenses with a commitment to natural properties. Likewise, (Cohen & Callender, 2009) outline a similar proposal for altering the best-system account of laws to avoid the questionable commitments resultant of Lewis's primitive metaphysics. That is not to say that Lewis would find these views acceptable. Some might worry that the pragmatic approach makes the laws unacceptably anthropocentric (Jaag & Loew, 2020, pp. 2538-2540) or subjective (Hall, 2015, p. 38). As such one might pursue alternative accounts whereby the feature that explains how laws necessitate is secured without explicitly appealing to factors linked to the epistemology of agents. One route could be to rigidify the objective standards of simplicity and strength (Lewis, 1994a, p. 479), or another could be to demonstrate that the desiderata of simplicity and strength in the canonical best systems analysis are alone sufficient to do the work required (Friend, 2022).

Regardless of the choice one makes here, any account that explains how the elements play their roles will find equal backing from the explanatory principle and when facing these challenges Lewis has a number of options. Likewise, for the metaphysics of chance, if Hall is correct in demonstrating that Lewis's theory of chance rests on equally unfairly

[7] A further discussion of natural properties in relation to explanation can be found in Paper II.

named elements, revisions to the problematic parts of the theory are in order. For frequentist theories of chance, the same kind of revisions appear plausible; Hicks (Hicks, 2017) advances a frequentist position that builds 'accuracy' into the notion of chance, which seems to be a good candidate for meeting the demands of the explanatory principle.[8] This process is warranted throughout Lewis's systemic metaphysics as if one wants to endorse the explanatory principle (something Lewis seems to strongly do) then any concurrently proposed theories must also adhere to it.

4.3 Resolutions

Worries relating to systematic conflicts are not unique to Lewis. Any explanatory metaphysician who makes systemic proposals will face the same kinds of questions. They must be careful to ensure that none of their proposals conflict with the explanatory principle. A useful consequence of this discussion on Lewis is that the potential options available when one faces systematic inconsistencies with the explanatory principle can be clearly articulated.

One available option when faced with unfair metaphysical proposals, as has been shown, is to simply reject the metaphysics. Any first-order metaphysical theory incompatible with the meta-philosophical proposal ought to be replaced with theories that better explain how their elements do their work. Another option for individuals unlike Lewis, who do not embrace the explanatory principle as fervently, is to reject the explanatory principle when facing a conflict between their metaphysical proposals and the principle. Armstrong might be receptive to this move because it allows him to retain his universal account of laws. Or Schaffer might be receptive to it because he believes one need only stipulate

[8] See (Hicks & Wilson, 2021) for a related discussion on the explanatory role of chances.

what the roles are and need not explain their being played. Overall, this option would rest on adopting a distinctly non-explanatory approach to metaphysics:

> *Primitive Metaphysics*: metaphysical proposals need not adhere to the *explanatory principle*.

There is nothing obviously objectionable about primitive metaphysics outside of the fact that it licenses elements to play roles without any explanation for how they do so. I suspect that primitive metaphysics would be appealing to anyone who views the cost of discarding their metaphysical proposals to be greater than the cost of not being an explanatory metaphysician.

However, there is an alternative that does not implicitly rest on a choice between explanatory metaphysics or primitive metaphysics (or on their compatibility with various other metaphysical proposals). We might choose to not enforce the explanatory principle absolutely and allow for certain metaphysical proposals to be exempt. This is akin to adopting a weak version of the meta-philosophical explanatory principle and could be understood as a positive proposal to approach a system of metaphysics holistically; comparing contrastive systems to see which one posits elements that explain more. Not everything needs explaining and the system of theories that manages to explain the most is the preferable one. After all, some theories explain some phenomena better than others. Take, for example, Humean theories of laws being better suited to explain how science identifies laws while non-Humean theories are better suited to explain why the world exhibits regular patterns.

This approach has a *prima facie* problem with assessing the explanatory capabilities of competing systems. Being able to explain some phenomena might be more important for

some but not for others; the Humean might assign a lot of weight to being able to explain how science generates laws while conversely, the non-Humean might assign more weight to explaining why the world exhibits regular patterns. It is a real worry for an explanatory comparison of this sort to be incommensurable; Humeans and non-Humeans explain different facts and there is no way to tell which explanations ought to hold more weight.

While there can still be grey areas where it is not obvious if the proffered explanation is adequate (if the explanation satisfies condition (3)), explanatory metaphysics, whereby metaphysical proposals *must* adhere to the explanatory principle, provides a more robust basis to assess whether a theory is explanatorily adequate. Further discussions are warranted here but the stronger meta-philosophical view, one that allows for no exceptions and requires all theories in any given metaphysical system to possess an explanation for how their elements play their roles, does not face the *prima facie* issue outlined here.[9]

5.0 CONCLUSION

This paper began by arguing against two possible approaches to Lewis's fair play argument: one for being too strong and the other for being too weak. In turn, I developed and argued for an explanatory approach to Lewis's argument. To substantiate the explanatory principle as not merely resulting from the eccentric concerns of a few isolated metaphysicians, but as a prevalent motivation across metaphysics, I have reformulated some influential arguments in terms corresponding to those of the explanatory principle. Then, I addressed a potential area of concern by showing how Lewis strove in his own metaphysics seeking compatibility with the explanatory principle. When faced with a

[9] An expanded, and more general, discussion on these kinds of holistic comparisons features in Paper V, §1.2.

conflict, as an explanatory metaphysician, we must discard or modify the metaphysical theory which violates it. Alternative options are available for non-explanatory metaphysicians and further discussions are warranted on neighbouring meta-philosophical proposals. From here there remains the expansive task of discovering what other metaphysical proposals violate and conform to the explanatory principle, and of exploring the programme of explanatory metaphysics that it opens up.

PAPER II: REALISM AND THE VALUE OF EXPLANATION[10]

ABSTRACT

Properties are cheap, and this is problematic. The difficulty is to distinguish between the properties that are important and noteworthy and those that are insignificant or irrelevant. A familiar solution is to adopt an inegalitarian classification that demarcates a limited number of privileged properties. Yet, Dasgupta (Dasgupta, 2018) poses a serious challenge to realism. He argues that there is no acceptable explanation of why the realist's privileged properties deserve the value realists assign to them and are consequently absent of value. In response, this paper defines and defends an alternative non-explanatory account of normativity that can be employed by realists to justify the attribution of value to their privileged properties. Unlike the comparable realist positions of (Lewis, 1984) and (Sider, 2022), who argue that it is sufficient to defend realism solely on realist terms, I engage with the challenge on unfriendly grounds. First, I demonstrate that anti-realists face a similar challenge to that directed against realism and therefore, realism is no worse off than competing anti-realist accounts. Second, I show that enforcing a connection between explanation and normativity rests upon an unexplained assumption that prevents any universal demand that properties must have their value explained from being imposed. Additionally, I show that this argument translates over to a similar debate in meta-ethics, address the underlying pragmatic motivation for Dasgupta's argument, and explore a nearby view connecting value to phenomenal acquaintance.

OVERVIEW

In §1, background on the problem of the cheapness of properties and the normative responses to this problem made by realists and anti-realists is outlined. In §2, Dasgupta's challenge to realism is presented and the important premise of explanatory normality is clarified. In §3, three arguments are presented. The first defends realism from Dasgupta's challenge on the realists' terms where realist-friendly assumptions are permitted. The second and third arguments pose the same problems for anti-realism and explanatory normativity and therefore do not rest upon assumptions to which only a realist would agree. In §4, with the central arguments in place, I show how my response to Dasgupta's challenge also applies to the relevant discussions in meta-ethics, directly address the underlying pragmatic motivation for explanatory normativity, and define a phenomenal view that separates understanding from explanation.

[10] This chapter is an extended version of a paper (Andrews, 2023) accepted for publication in *The Philosophical Quarterly*, 73, (4), 2023, pp. 1305–1314; [https://doi.org/10.1093/pq/pqad052].

1.0 BACKGROUND

1.1 THE CHEAPNESS OF PROPERTIES

Properties are abundant. As Lewis puts it:

> 'Any class of things, be it ever so gerrymandered and miscellaneous and indescribable in thought and language, and be it ever so superfluous in characterising the world, is nevertheless a property. So there are properties in immense abundance' (Lewis, 1983a, p. 346).

Contrast being square with being a shape with four equal straight sides and four right angles. Or contrast being simultaneous with being at the same time and not being the left half of a trout. Or contrast being an Ibex ram with being composed of oxygen, carbon, hydrogen, nitrogen, calcium, and phosphorus molecules in the configuration of an Ibex ram.[11] In each case, the former is a property that appears relevant and noteworthy while the latter appears arbitrarily invented and inappropriate. When constructing a theory, where theorising is understood as an activity that aims to accurately describe the world, the use of the latter property would not affect the truth of the theory. Yet, we theorise with a preference for the former. Call this the *problem of the cheapness of properties*: when we engage in theorising countless alternative properties could be used but only a limited number of them should be. The challenge is to demonstrate what is so unfitting about those numerous odd properties. This challenge has often been neglected. For example, Loewer notes that 'every other account of laws - both Humean and non-Humean - helps

[11] These examples are adaptations of those given by Dorr (2019). See for further exposition on natural properties along with (Dorr & Hawthorne, 2013) for a discussion on naturalness more generally (and critically). Other cheap properties include disjunctive properties (Hirsch, 1993) and projective properties (Goodman, 1955).

itself to a distinction between properties that are fit and those that are unfit for laws.' (Loewer, 1996, p. 109).

1.2 Eliteness and Value

The solution to the problem of the cheapness of properties is to identify some from the abundance of properties that *ought* to feature in our theories: adopting an inegalitarian classification between properties and selecting a limited number of them as theoretically preferable to the rest. We can separate those properties that ought to feature in our theories and call them *elite*, from those that ought not to feature in our theories and call them *non-elite*. Then if asked what is wrong with those miscellaneous and gerrymandered properties the answer is that they are non-elite and deserve no place in our theories. Summarised into a proposal:

Eliteness Proposal: the elite properties are those we should theorise with.

It is important to recognise that this demarcation is inherently normative. The elite properties can be said to have value in virtue of making our theorising more accurately describe the world; elite properties possess value-theoretic upshots. The eliteness proposal does not deny that some properties exist or that it is possible build theories with non-elite properties. Instead, it suggests that those theories using non-elite properties are *worse* than those that utilise elite properties; the elite properties guide *good* theorising.

What results from the eliteness proposal is the need for an account of what makes a property elite. After accepting that not all properties are equal, it remains to be answered what the elite properties have that the rest lack.[12] The answer principally rests on whether

[12] There are similarities between this question and the explanatory principle as outlined in Paper I of this thesis; an expanded discussion comparing the explanatory notions at play in both papers can be found in the next chapter titled 'BRIDGING PAPERS I, II, & III'.

the distinction between the properties that appear noteworthy and significant and the properties that appear superfluous and disposable is determined by something inherent in the properties, or by their relationship to us as theorisers. In other words, the difference comes down to realism or anti-realism.

1.3 Realism and Anti-Realism About Elite Properties

Realist approaches rely on a metaphysical claim about the world to identify the elite properties while in contrast, anti-realist approaches rely on a theoriser-relative claim to identify the elite properties. One important difference is that for realists there is a singularly objectively correct group of elite properties; realists maintain that there is a distinct and objective set of elite properties *all* theorisers should use. For anti-realists, any one of several groups of elite properties can be correct, as a property might be elite to one group of theorisers but non-elite to another; there is then no objectively correct set of elite properties. To plainly define these positions:

> *Realism about elite properties*: objective posit φ provides an account of what properties are objectively elite for all actual theorisers.
>
> *Anti-realism about elite properties*: mind-dependent posit ψ provides an account of what properties are elite relative to a particular group of actual theorisers.

For example, one prominent realist approach asserts that the world's objective structure makes a property elite. The two main proponents of this view are Lewis (Lewis, 1984) and Sider (Sider, 2011). The central claim is that a property is elite when it demarcates the metaphysically objective structure of reality. As Lewis remarks, 'Only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature.' (Lewis, 1984, p. 227). Sider makes similar claims

about objective structure, ‘Although their beliefs are true, those beliefs do not match the world's structure’ (Sider, 2011, p. 2).

Although Lewis and Sider define distinctive accounts of the world’s structure - Lewis invokes naturalness as a monadic predicate of properties while Sider uses a generalised structural operator – they both make the relevant objective commitments and there are extensive similarities between their positions such that they can be grouped together here. Likewise, For Lewis naturalness plays many roles. Lewis’s conception of *perfectly natural* properties, which is a matter of degrees and entwined with notions of fundamentality, is separable from the categorical role of naturalness determining eliteness (Schaffer, 2004). The categorical role is all that is utilised for realism about naturalness in this paper.

Alternatively, a notable anti-realist view asserts that linguistic history makes a property elite. This view comes from Goodman (Goodman, 1955) who claims that certain properties are elite due to them being entrenched in our linguistic practices; a limited number of properties traditionally feature in our language, and this accounts for why we should theorise with them. Theorisers with different linguistic histories will identify a different set of elite properties as nothing makes our linguistic history objectively special.

In general, realists like Lewis and Sider are substituting objective structure for φ in their version of realism about elite properties while anti-realists like Goodman substitute linguistic history for ψ in his version of anti-realism about elite properties. This kind of connection between the eliteness proposal and various realist and anti-realist positions can be found across metaphysics. Other realist and anti-realist accounts that do not explicitly invoke objective structure or linguistic history also fit this model with the main

difference coming down to whether eliteness is accounted for by an objective or subjective posit. For further realist examples, φ could be the objective property of being a law and the elite properties being those that feature in the laws (Maudlin, 2007), or φ could be an objective grounding relation and the elite properties being those that are ungrounded (Schaffer, 2017), or φ could be a sparse set of universals and the elite properties being those that correspond to the universals (Armstrong, 1978). For further anti-realist examples, ψ could be deductive connexions between predicates in a theory and the elite properties being those that play more central classificatory roles (Taylor, 1994), or ψ could be the biological traits of theorisers and the elite properties being those that are appealing to a particular species (Putnam, 1990), or ψ could be the preferences of a group of theorisers and the elite properties being those favoured to feature in their theories (Dasgupta, 2018).

For simplicity, for the rest of this paper, I will focus on two particular realist and anti-realist views. For realism, the discussion will centre on *realism about naturalness*, whereby the objective natural structure of our world identifies a single objectively elite set of natural properties. For anti-realism, the discussion will centre on *preference anti-realism*, where the preferences of theorisers identify the elite properties relative to a group of theorisers who exhibit such partialities. For both, in the arguments that follow nothing specific hangs on choosing to focus on any particular realist or anti-realist account. The same arguments apply to any version. The choice to focus on naturalness and preferences is done solely because the objection I aim to defend realism from is directed against naturalness and the alternative anti-realist position (which supposedly avoids the faults of the realist account) appeals to preferences.

2.0 REALISM AND THE ABSENCE OF VALUE

2.1 EXPLANATORY NORMATIVITY

Realists about natural properties face a serious challenge from Dasgupta (Dasgupta, 2018). The challenge principally rests on endorsing an explanatory account of normativity and then demonstrating that there is no explanation for why the realist's natural properties deserve value. In other words, Dasgupta argues that the realist is unable to provide an account of why their elite properties deserve to be called elite when the explanation appeals solely to a metaphysical posit. Ironically, as Dasgupta notes (Dasgupta, 2018, p. 286), his argument is an adaptation of Lewis's objection to primitivism about chance (Lewis, 1994a, p. 484). Realism about naturalness rests upon an analogous kind of primitivism about metaphysical structure and as such, Dasgupta argues that it faces an analogous kind of objection. Lewis objects that the chance primitivist fails to 'play fair' when claiming that objective chances are "real" chances. To play fair they must demonstrate why the property deserves the name "real" without simply affirming that it is objective.[13]

To see how Dasgupta's argument works it is important to clarify his underlying conception of normativity before moving to show how the realist fails to satisfy the demands required to obtain this kind of value. Dasgupta's adaptation of Lewis's argument to normativity and naturalness asks the realist to meet an explanatory requirement: there needs to be an explanation, beyond simply assuming that "natural" properties are elite, for why "natural" properties are elite. Importantly, the explanation must meet a certain standard of non-triviality. The explanation needs to provide an account of why the natural

[13] This closely resembles the non-triviality clause defined in Paper I of this thesis.

properties are elite that does not merely assume, paraphrase, or repeat that natural properties are elite. Any explanation that invokes the explanandum as explanans does not illuminate in any capacity: simply reiterating the thing that is supposed to be illuminated does not make it any brighter. Call this explanatory normativity:

> *Explanatory Normativity*: if a property ought to feature in theories, then there must be an explanation for why it ought to feature in theories.[14]

This principle outlines a criterion for the normativity of elite properties. For elite properties to uphold their privileged status there must be an explanation for why they deserve to be valued. If some properties are proposed as elite but fail to explain this, then they do not qualify for elite status. Importantly, the criterion for non-triviality states that permissible explanations must demonstrate why elite properties are valuable and must not simply assume, paraphrase, or repeat that they are. Trivial explanations are explicitly excluded. As will become apparent, this clause is doing important work and it is worth highlighting that even though a trivial explanation might be on offer, this will be taken as equivalent for present purposes to providing no explanation. Overall, explanatory normativity affirms that if we should theorise with a property there should be good reasons for why we should theorise with it and if there is no explanation for a property's having value, the property is not valuable.

Although explanatory normativity is directly inspired by Dasgupta's work drawing connections between normativity and metaphysics, the underlying explanatory

[14] There is a shift in terminology here with the use of 'why' in contrast to 'how' as used in the previous chapter. Dasgupta uses 'why' when interpreting and applying the fair play argument to normativity. As this chapter focuses on Dasgupta, I will stick with his terminology. While I believe the 'how' locution is the most apt, the contrast between the different explanatory requests made by 'why' and 'how' do not bear substantially on the arguments outlined here. Further discussion on this can again be found in the following chapter, specifically p.86.

motivation is more widespread. In general, something akin to explanatory normativity is broadly accepted and intuitively plausible; it seems obvious that if some properties ought to feature in our theories, we ought to be able to say why this is so. Similar doubts are already posed for Lewis's commitment to an objective natural structure in his theory of reference (Hawthorne, 2007) which is highlighted in the account's derogatory branding as 'reference magnetism' (Hodes, 1984, p. 135). Likewise, Lewis's proposal that the laws of nature need to be formulated with natural predicates faces criticism for similar reasons (Cohen & Callender, 2009); (Loewer, 2007). This cohort of anti-naturalness theorists share similar explanatory ideals and will likely endorse or at least be sympathetic to explanatory normativity. I will call those who accept explanatory normativity, independent of realism or anti-realism, *explanatory metaphysicians*.

Dasgupta falls into this group of explanatory metaphysics and employs explanatory normativity to argue that realists about naturalness cannot explain why natural properties deserve to be called elite. Where the realist might appeal to the connections between natural properties and laws, causation, or induction to explain why natural properties warrant elite status, none of these moves is deemed adequate (Dasgupta provides a detailed discussion on the impermissibility of many potential realist options (Dasgupta, 2018, pp. 295-308) and I do not contest these points). The problem chiefly stems from realists being restricted to 'unfair' accounts that reiterate natural properties as aligning with the world's natural structure or paraphrasing natural properties as being "really" natural but do not in any reasonable capacity explain why they have value. Any claim that natural properties matter because the world has an objective natural structure assumes that theories ought to adhere to the world's metaphysical structure before having shown that such a posit deserves to feature in theories. Realists cannot affirm their metaphysical

posit to have value until they have shown the metaphysical posit to be something worth valuing: 'don't call it "naturalness" until you've shown that it is something that should guide our theorizing' (Dasgupta, 2018, p. 286).

Realists claim that natural properties are elite because of a metaphysical posit about the world's structure, but this is no explanation for why, when theorising, we should care about natural properties. The question realists must answer is: why should a metaphysical posit about the structure of our world have any bearing on what properties should and should not feature in our theorising? Realists might try to defuse such a question with the obvious answer that natural properties matter because the world has an objective structure. But this is to assume realism in a debate about realism – to assume that our theorising ought to adhere to the world's metaphysical structure before showing that such a posit deserves to feature in our theories.

To further substantiate the objection, one could (for comparison) propose that the world actually has an objective *gratural* structure, where gratural properties are those altered in a projective manner like *grue*: if observed before time t being green, or if not observed before time t, being blue (where time t is some point in the future). On the assumption that properties are amply cheap, there are many second-order properties (natural, gratural, etc.), that range over contrasting first-order properties (green, grue, etc.), any of which have equal claim to being valuable if we first posit the world to have a corresponding metaphysical structure.[15] If natural properties have value because they accord with the world's natural structure, then there is just as much reason to think that gratural properties have value because they accord with the world's gratural structure. As Dasgupta writes,

[15] This degree of property cheapness is assumed by Dasgupta (Dasgupta, 2018, p. 289).

‘don’t say that naturalness is itself natural and graturalness is not, for we are in the middle of trying to explain why naturalness matters!’ (Dasgupta, 2018, p. 289). Consequently, realists about naturalness fail to present a satisfactory account of the normativity of elite properties; they fail to provide an illuminating account of why natural properties warrant the value ascribed to them.

To summarise the argument: explaining the value of elite properties is a requirement for a property to have value, realism about naturalness does not satisfy this demand and therefore, realism is absent of value. In turn, Dasgupta goes on to argue for anti-realism, whereby natural properties are just those we *prefer* (Dasgupta, 2018, pp. 308-309). The preferences of theorisers explain why certain properties ought to feature in theories and no objectively correct set of properties possesses value-theoretic upshots.

2.2 Two Contrasting Interpretations

With the problem and some of its implications stated I will now move towards resolving it. Before doing so, a crucial distinction needs clarifying. There are two ways explanation can be understood, and the choice between them drastically alters the demands expressed in explanatory normativity. The notion of explanation is central to the objection and yet it is something about which realists and anti-realists disagree. Dasgupta notes that his objection to realism could be that realists lack an objective explanation for why natural properties are elite or it could be that realists lack an explanation, relative to our preferences, for why natural properties are elite (Dasgupta, 2018, p. 312). The argument runs differently depending on which notion of explanation is in play.

On the one hand, if realism is assumed and explanation construed in distinctly realist terms, where the explanation must appeal to objective facts, then the objection to realism is a *reductio*:

1. Assume realism.
2. Assume explanatory normativity.
3. From 1 and 2, there must be an explanation that meets the realist's objective standards for why only natural properties have value.
4. There is no explanation for why only natural properties have value.

C. Realism is false (and this contradicts the first premise assuming realism).

Call this argument, where realism is assumed, the *internal criticism*. On the other hand, if anti-realism is assumed and explanation construed in distinctly anti-realist terms, where the explanation can appeal to the subjective preferences of theorisers, then the objection to realism runs as follows:

1. Assume anti-realism.
2. Assume explanatory normativity.
3. From 1 and 2, there must be an explanation that meets the anti-realist's non-objective standards for why natural properties have value.
4. Anti-realists explain the value of natural properties in virtue of our preferences.[16]
5. Theorisers with contrasting preferences confer value to contrasting properties (natural, gratural, etc.).
6. 5 contradicts the realist claim that only natural properties have value.

[16] I doubt a realist would accept explanations relative to us and therefore, there is no realist-acceptable explanation, and the argument then continues as in the internal criticism.

C. Realism is false.

Call this argument, where anti-realism is assumed, the *external criticism*. Although both arguments come to the same conclusion, they are established on different assumptions. The external criticism is that without making assumptions only a realist would agree to, the realist is ineffective at convincing the anti-realist (or even a fence-sitter) that their view is correct. Perhaps it is not too surprising that realism faces opposition when non-realist assumptions are made, but vulnerability to the anti-realist version of the objection will dissuade the non-aligned. The internal criticism makes the stronger claim that realism is indefensible even if realism is assumed. It is an exceptionally worrying outcome if realists are unable to make their case even on their own terms. Further, it is worth explicitly stating that although explanation is the key feature whose interpretation changes the argument, the difference extends to assuming the realist and anti-realist positions wholesale, not just their conceptions of explanation.

As I argue at the end of the next section, responding to both criticisms is important to exonerate realism. As such, I will now present three related arguments that together exonerate realism and define a novel account of non-explanatory normativity. These arguments also expose problems for anti-realists and explanatory metaphysicians who endorse explanatory normativity independently of realism or anti-realism. First, in §3.1, I respond to the internal criticism by arguing that when realism is assumed, explanatory normativity may be consistently and coherently rejected, and an alternative non-explanatory normative principle accepted in place. Second, in §3.2, I respond to the external criticism and demonstrate that when anti-realism is assumed, there is no satisfactory explanation for why the elite properties of anti-realism deserve elite status. The situation is analogous to the accusations Dasgupta has logged against realism. Thus,

unexplained value becomes obligatory for both realism and anti-realism. Finally, in §3.3, I expose a problem for explanatory metaphysicians - those who maintain explanatory normativity independent of realism or anti-realism - by demonstrating that explanatory normativity itself rests on a questionable assumption about the connection between explanation and value that is itself unexplained.

3.0 EXPLANATION AND VALUE

3.1 THE INTERNAL CRITICISM

When assuming realism, the demand to explain why natural properties deserve to feature in our theories might well be a fair one, but it is not a request a realist needs to acquiesce to. By not assuming explanatory normativity and denying premise 2 of the internal criticism, no *reductio* obtains. In place of explanatory normativity, an alternative principle affirming that it can be that some properties ought to feature in our theories without further explanation for why they ought to, is needed:

> *Primitive Normativity*: if a property ought to feature in theories, there does not need to be an explanation for why it ought to feature in theories.

Primitive normativity affirms that some properties should feature in theories as a brute fact that lacks any explanation. Realists who accept primitive normativity take it as a basic commitment of their view that natural properties have value-theoretic upshots. Importantly, I am conceding that Dasgupta is correct in arguing that there is no non-trivial explanation for why the realist's natural properties are elite. I do not contest the charge that the realist cannot produce an acceptable explanation; at least not in any sense that will meet the non-triviality standards inherent in explanatory normativity. Hence, the suggestion here is for the realist about naturalness to outright reject explanatory

normativity and declare the elite status of natural properties to obtain without any non-trivial explanation.

Initially, this might seem to be just as much of a concession as it is a reply. Dasgupta argues that realists lack an explanation for the value of natural properties and my suggestion is to accept natural properties as possessing unexplained value. The problem realists face is that they offer no explanation of the value of their whatnots and the response suggested is to take unexplained whatnots as valuable. Some might worry that explaining why the unexplained whatnots are valuable is still explanatorily inadequate. But remember this is a defence of realism from the internal criticism. It is an argument against the charge that realism cannot defend itself on its terms. Realists assert that the world does exhibit an objective natural structure and therefore, natural properties are valuable because theories are about that world! Realists are free to claim that they can see, dimly enough, why objective structure imbues natural properties with value. Just because the value of natural properties cannot be explained does not mean those properties are not valuable. On the contrary, the value of realist natural properties being unexplainable leads not to the absence of value but to accepting value that is unexplained. This contravenes explanatory normativity, but it does not unavoidably falsify realism. There is an obvious way out in accepting a different – realist-friendly – conception of normativity.

To clarify the normative proposal, the value resulting from accepting primitive normativity is best understood as a kind of intrinsic value (understood equivalently to non-instrumental value). Generally, intrinsically valuable properties are valuable in and of themselves with nothing outside of the properties determining why they have value. This kind of value is not instrumental, ethical, or prudential but instead is a kind of

epistemic value – the relevant ideal is *epistemically good theorizing*. As the normative proposal is primitive, analysis of it is going to have its limits. But I have provided a functional account by stating what role the value plays and how it relates to other comparable conceptions of value. Any further analysis of how intrinsic value pertains to natural properties, going into further detail on how science identifies the valuable properties, is beyond the scope of this proposal. Ultimately, it is better to theorise with natural properties and this value is derived from their intrinsic value.

When the realist about naturalness adopts primitive normativity, they are affirming the value of natural properties to be inherent in the natural properties themselves. Nothing outside of the natural properties determines why these properties are elite and any attempt to explain the intrinsic value of natural properties will have to trivially appeal to the value inherent in natural properties. Therefore, the realist ought to reject explanatory normativity in favour of a non-explanatory normative principle that captures the connection between the eliteness proposal and their metaphysical posits. By embracing the intrinsic value of natural properties, realists are adopting a distinctly realist conception of value that obtains without a non-trivial explanation.

Before moving forward, it is worth noting that Dasgupta considers a realist response that denies premise 2 of his argument (i.e., the response rejecting explanatory normativity) and accepts this response as a coherent option: 'To be clear, there is nothing *incoherent* about denying premise 2. It is an option in logical space—indeed I suspect it is the realist's only retreat. But I insist that we recognize it as the radical position that it is' (Dasgupta, 2018, p. 294). It might seem radical to a non-realist, but even so, just being called a radical is no objection. I take Dasgupta's stance here to be more an expression of disapproval than a critique. Putnam criticises Lewis's realist metaphysics for sounding 'medieval'

(Lewis, 1984, p. 229), to which Lewis's response can be neatly adapted here: what's wrong with being a radical? More power to the radical metaphysicians.

The key difference between these approaches to normativity is that Dasgupta seems motivated by the idea that we only ought to attach normative significance to that which we can explain why significance attaches to them. It is a comfortable proposal to suggest that we should only care about the parts of reality that we can explain but that is no reason to assume we live in such a world. To echo the words of Camus, 'it is natural to give a clear view of the world after accepting the idea that it must be clear' (Camus, 1955, p. 40). The realist ought not accept this assumption about the connection between theorising and value. Instead, they should accept a converse connection in the form of primitive normativity. This fits with the general attitude which realists take towards theorising. They accept that we might be incorrect about which properties are natural/intrinsically valuable as merely an inevitable consequence of adopting a robust distinction between reality and (even optimal) theorising. Compare Devitt's response to Putnam's model-theoretic argument, wherein Devitt argues that realists ought to freely accept every proposition we think is true might be false simply as a consequence of believing the world is altogether independent of theorising:

> 'Imagine an intelligent organism like a human but with very inferior sense organs. We have no difficulty in understanding how a theory might be ideal for that organism and yet from our superior perspective be clearly false. What the realist believes is that this might be so for any organism including a human being; for the realist believes that reality is altogether independent of experience.' (Devitt, 1983, p. 297)

It is plausible for realists that, based on our finite capacity as theorisers and the extension of reality beyond those limits, we will not be able to explain everything that has value-theoretic upshots. Analogously, realists about natural properties should accept that every property we think is natural might be unnatural, as a consequence of believing nature's joints exist entirely independently of theorising. The plausibility of such a proposal relies on accepting that there are aspects of reality that we may never be able to uncover and one place one may look in support of such a suggestion is found in the ubiquitous writings of Lewis, 'Why is Humility "ominous"? Who ever promised me that I was capable in principle of knowing everything?' (Lewis, 2009, p. 211).

Importantly, since the normative proposal is primitive, it entails accepting a certain degree of epistemic humility. In this sense, primitive normativity constitutes a concession to the idea that theorising about reality might not be co-extensive with reality; there may be features of reality whose value we are never able to explain but are no less significant, and these features arise in cases alike the value of natural properties. So one way of looking at this aspect of Dasgupta's argument is as an adaptation of Putnam's cartesian credo, whereby the realist commits to the notion that every proposition we think of as true might be false (Button, 2013, pp. 10-11), to apply to normativity and naturalness. The realist commits to the notion that every property we think is natural (and possesses intrinsic value) might not be. Yet, the inability to provide an explanation for why natural properties have value is simply a consequence of believing that what determines whether a property is natural is entirely independent of us. Devitt's argument translates over too: this is not a bug, but a feature of realism.

However, there is one additional question that should be addressed. Adopting primitive normativity leads to a further challenge as outlined by Dreier (Dreier, 2015, p. 179) in the

context of moral realism, concerning the moral realist's ability to explain why it is irrational to be unmotivated by what one believes one ought to do. Even if the moral realist takes it as a brute fact that some facts are normative, this is not in itself a good reason to think that no explanation is required for why it is irrational:

> 'First, we might wonder whether any explanation is needed. Could it just be a brute fact that it is irrational to be unmotivated by one's normative thoughts? In a way it seems like a good candidate: it is very obvious that this is irrational. But that's not a good reason to think it needs no explanation! (We are not asking for a justification, after all.) So if it has no explanation, I think, we would want an account of why it doesn't: a second-order explanation.' (Dreier, 2015, p. 179)

More generally, just because one adopts the stance that some normative fact is unexplained does not support the affirmation that no explanation is needed. Instead, for cases where there is no explanation for why the elite properties are elite, there should be a reason for why there is no explanation. In other words, there should be a second-order explanation that explains why no explanation is correct; if we are plausibly reject explanatory normativity and instead accept primitive normativity then we should be able to provide some further explanation, explaining why no explanation is correct.

For realists, no first-order explanation is the correct approach because intrinsic value cannot be explained in a way that satisfies the standards of explanatory normativity. This has been noted by Dorr: 'it is unclear whether any claim of [intrinsic] value could be explained in a way that would satisfy Dasgupta's criteria for adequate explanation' (Dorr, 2019), and Pérez Carballo also states that when facing the question, 'why is it better, epistemically, to use concepts that correspond to elite properties? […] I think it would not

be surprising if such an answer is not forthcoming.' (Pérez Carballo, 2020, pp. 312-313). As a result, accepting that the reason natural properties ought to feature in our theories comes from the natural properties themselves undermines any attempt at non-trivially explaining why the natural properties ought to feature in our theories. If the value of natural properties is inherent to natural properties, then the only viable first-order explanation for why they have such value must appeal to natural properties and their innate value! Subsequently, by accepting natural properties as having intrinsic value there is no permissible route towards obtaining a first-order explanation of why we should theorise with natural properties in accordance with the standards outlined in explanatory normativity.

In turn, accepting the intrinsic value of natural properties provides a second-order explanation of why there is no first-order explanation. Specifically, the reason there is no non-trivial first-order explanation for why natural properties are valuable is that the value of natural properties is not apt for a non-trivial first-order explanation. The elite status of the natural properties cannot be explained non-trivially because the elite status of these properties is accounted for in their intrinsic value. The realist explains the absence of explanation by appealing to intrinsic value, providing the second-order explanation of why no explanation is correct – an explanation supporting their acceptance of primitive normativity.

My defence of realism shares some structural similarities with Sider's (Sider, 2022) reply to Dasgupta. But there are some substantial differences. First, I explicitly concede to an explanatory deficit while Sider is not so conciliatory. Second, I employ the concept of intrinsic value and use this to answer the second-order challenge while Sider employs an externalist conception of value in his argument, whereby the value of natural properties

is fixed by objective features of our world. Thirdly, Sider believes 'joint-carving' to also be instrumentally valuable (Sider, 2011, p. 61), while I make no so such claim. The positive thesis I present in this paper amounts to the conjunction of realist metaphysics, primitive normativity, and intrinsic value. Not arguing for the instrumental claim is advantageous to realists because it avoids superfluous commitments. As Thomasson argues, the instrumental value of theorising is fully captured 'without positing extra metaphysical facts about 'structure', without invoking epistemic difficulties about how we could know such facts, or about why such metaphysical facts should be theory-guiding' (Thomasson, 2020, p. 451).[17]

But the most significant difference is that *pace* Sider and Lewis, I do not believe that responding to the internal criticism is alone a sufficient response to Dasgupta's critique of realism. To stop here and defend realism solely on the realist's terms is dialectically dubious. Putnam decried Lewis's response to his model-theoretic arguments as 'just more theory' (Putnam, 1985, p. 18), and a similar challenge applies as the realist's value claims are just more unexplained value. Lewis's response to such concerns is to reiterate a strong commitment to realism: 'If I am looking in the right place for a saving constraint, then realism needs realism' (Lewis, 1984, p. 228). Similarly, Sider also deploys realism to defend realism, and further advocates for a defensive stance: 'for it is no part of the mission to convince Dasgupta, or even a fence-sitter, that the outlook is correct.' (Sider, 2022, p. 10).

Critically, however, these arguments (including mine up to this point) rely upon making realist assumptions to save realism. To stop here raises a serious worry as there appears

[17] Christine Korsgaard raises similar problems for dogmatic rationalism in (Korsgaard, 2009, pp. 52-57).

to be no case for accepting realism unless one is already a realist. This kind of argument will not persuade anyone but those already persuaded. Adopting defensive stances leads to philosophical stalemate and debate over metaphysical realism has already arrived at such an impasse.[18] The realist claims their objectors are caricaturing their position if it is not understood through the lens of suitable realist assumptions and their objectors claim that making such assumptions is dogmatic. Both realists and anti-realists alike might find it hard to see common ground and disagreement is bound to arise around what notions one is allowed to accept as foundational. But if we concern ourselves solely with those who already accept our views, we may well spend our time transporting sand to the beach. To provide any useful development to the dialectic one cannot maintain a defensive stance. Accordingly, I will now advance a defence of realism on unfriendly grounds.

3.2 The External Criticism

Forgo realist assumptions for minimal anti-realist ones. Recall that the anti-realist view I will focus on is preference anti-realism: the elite properties are those that are preferred by a group of theorisers (the same argument outlined here applies if one replaces preferences with linguistic history, or species, etc.). On this anti-realist model, there can be many sets of properties that are elite simply because they are preferred by the relevant community; the natural properties, the gratural properties, etc. As only those properties that are preferred are elite, an inegalitarian classification is not established between the natural and unnatural properties nor between the gratural and ungratural properties but instead between the preferred and unpreferred properties. Individually, each community will uphold a different set of properties to be elite, but a ubiquitous feature exhibited by all

[18] See (Khlentzos, 2016) for a summary of the stalemate present in the literature.

elite properties is that they are preferred. Even if these properties are only elite relative to a particular group of theorisers, the preference anti-realist asserts that all elite properties are elite by virtue of being preferred. Therefore, call the elite properties of preference anti-realism - those that have their elite status explained in virtue of preferences - *prefatural* (and recall that properties are cheap, so no qualms should be had about the formation of this property). Now, since the preference anti-realist affirms that prefatural properties are elite, explanatory normativity applies: if prefatural properties are elite there must be some non-trivial explanation for why prefatural properties are elite.

What explanation could be provided? As the anti-realist establishes that the elite properties are those that are preferred, it appears that their answer is restricted to appealing to some feature or other of our preferences; stating that being preferred explains why prefatural properties are elite. However, the challenge I pose here is to explain why, when theorising, we should care about what properties are and are not preferred. The request is to explain why preferred properties matter and appealing to preferences alone tells us nothing as to why prefatural properties should feature in our theories. Likewise, any additional appeal to preferences is equally unilluminating; we cannot say that we prefer green because we prefer natural properties because we prefer prefatural properties. Simply reiterating or paraphrasing that a set of properties is preferred goes no way in explaining why it is the case: the task is to state why preferred properties matter and not simply reiterate that they are "preferred"!

Just like realism requires the value of natural properties to fall out freely from an 'unfair' posit about structure, anti-realism requires the value of prefatural properties to fall out freely from an 'unfair' posit about preferences. Yet, this only works when one assumes that our theories *ought* to utilise the properties we prefer. Simply *being* preferred tells us

nothing as to why these properties ought to feature in theories unless one assumes that the only properties that ought to feature in theories are those that are preferred. After all, just being preferred does not imbue a property with value-theoretic upshots, any more than calling someone 'Goodman' makes them virtuous. Be my guest – posit a set of properties that you prefer if you want. But play fair. Do not call it "prefatural" until you have shown it to be something that should guide our theorising.

What other than preferences could anti-realists appeal to in an attempt to explain the eliteness of prefatural properties? As anti-realism is partially motivated by the idea that they can explain why their elite properties are elite, I suspect there will be resistance to the claim that preferences are all that matter and in fact when asked to explain why prefatural properties matter, the answer will be that preferences *explain* why those properties have value. When asked to explain why prefatural properties are elite, the answer is that preferences provide a non-trivial explanation for why those properties are elite. Green is elite because we prefer it, and because preferences non-trivially explain why green is elite, prefatural properties are elite.

Yet, appealing to explanation does little to resolve the problem and instead simply pushes the burden from preferences to explanation. If the elite properties are only those that have their elite status explained, it turns out that the inegalitarian classification is not established between the preferred and unpreferred properties but instead between those such that there is an explanation for why they are elite and those such that there is not. Individually, each community will uphold a different set of properties to be elite in virtue of varying preferences, but a ubiquitous feature exhibited by all elite properties is that they have their elite status explained. We can call this version of anti-realism *explanatory anti-realism* and call the properties it designates as elite *explatural*: the elite properties

are those that have their elite status explained. Once more, as the explanatory anti-realist affirms that explatural properties are elite, explanatory normativity applies again: if explatural properties are elite there must be a non-trivial explanation for why explatural properties are elite. This condition is not satisfied by merely stating that preferences explain, but demands an answer to the question: what makes the properties with their elite status explained warrant elite status? And any feasible answer will include some appeal to the fact that they are explained; that the elite status of explatural properties falls freely from the fact that they are explained. Claiming that properties with their value explained are elite because they are (putatively) explained is no more acceptable than the realist claiming that properties that follow the world's structure are elite because they follow the world's structure.

From these *tu quoques*, the combination of anti-realism and explanatory normativity becomes untenable as it leads to the following *reductio*:

1. Assume anti-realism.
2. Assume explanatory normativity.
3. From 1 and 2, a non-trivial explanation that meets the anti-realist's non-objective standards is required to account for why the elite properties are elite.
4. There is no non-trivial explanation for why prefatural or explatural properties are elite.

C. Anti-realism is false

Anti-realists face the same dilemma as realists. One solution for committed anti-realists is to simply join realists in rejecting explanatory normativity in favour of primitive normativity and the unexplained – intrinsic – value of preferences. Perhaps this view is

palatable for some, but if the sole motivation for anti-realism had been that it meets the demands of explanatory normativity, anti-realists now lack a positive reason to be anti-realists. More substantially, recall that the external criticism of realism was motivated to avoid the dialectical issues associated with assuming a realist-friendly conception of value to defend realism. As the anti-realist equally relies on the same notions, the dialectical issues dissipate. The realist is no longer precariously resting their view on assumptions only a realist would agree to, as the anti-realist must make similar assumptions. As a result, the conclusion is to recognise primitive normativity not merely as a tool for realists to avoid explanatory objection but as a useful normative concept required by both realists and anti-realists alike.

Alternatively, explanatory metaphysicians who wish to retain explanatory normativity, and who reject both realism and anti-realism, do not appear to have an easy solution. For one, on the assumption that the only viable means of construing explanation are realist or anti-realist, there is now no plausible conception of explanation available to formulate explanatory normativity with. Furthermore, explanatory normativity appears to undermine itself. If all elite properties must have their value explained, then the only elite properties are explatural. Regardless of what does the explaining, there is no further explanation of why being explained matters. This undercuts any attempt at enforcing explanatory normativity. Committed explanatory metaphysicians might try to resist this conclusion, but I will now go on to expand the *tu quoque* and apply it to explanatory normativity more generally.

3.3 The Value of Explanation

Explanatory metaphysicians, those who accept explanatory normativity independent of realism or anti-realism, implicitly assume that the only properties that are valuable are

those with an explanation for why they are valuable. This kind of assumption is plain to see in Dasgupta's argument as he asserts that 'there must be some explanation' (Dasgupta, 2018, p. 289) for why the natural properties are elite; without an explanation those properties are not elite. But it is reasonable to ask why only those properties that have their elite status explained warrant elite status and, in line with explanatory normativity, to expect a good answer; it is a reasonable request to explain why explanation plays this central metaphilosophical role. Just as I highlighted Dreier's suggestion earlier in the paper for a second-order explanatory demand applying to primitive normativity, I suggest that a similar second-order demand must also apply to explanatory normativity. Just because one adopts the stance that normative facts are explained does not support the affirmation that there 'must be some explanation'. Explanatory metaphysicians assume that explanation is a prerequisite for normativity. My complaint is that there is no good reason for this assumption.

What explanation could be provided? To start one could argue that it is obvious why we must explain why the elite properties are elite, and therefore no explanation is needed. Properties that have their value explained seem clearly more deserving of value than those that do not. But this reply entails primitive normativity, as the properties whose value is explained are assigned elite status without an explanation for why this is the case. This is not a response so much as a concession that we must explain even while there is no explanation for why we must explain.

The challenge is met by providing an explanation for why there must be an explanation. Recall that I argued, in response to the internal criticism, that realism meets this challenge by adopting an intrinsic conception of value to explain why primitive normativity is correct. The realist answer is that no explanation is correct because the value of their

natural properties is not apt for explanation. A similar move is needed in defence of explanatory normativity, adopting a conception of value that explains why explanatory normativity is correct. Contra the intrinsic value suggested for realism, instrumental value is the likely candidate. By appealing to instrumental value one can claim that properties with their value explained are valuable because the kind of value they possess only obtains when an explanation is provided. However, if explanatory normativity is concerned exclusively with this kind of value, it does not serve as a universal criterion for determining what properties do and do not have value and therefore, poses no threat to the realist position I have outlined (nor to primitive normativity in general).

The only way explanatory normativity poses a threat to realism (and to primitive normativity) is on the assumption that the only kind of value is that prescribed by explanatory normativity. Upholding such a view rests on further assumptions about the relationship between value and explanation. Effectively, this answer to the second-order challenge of why the elite properties are only those that have their elite status explained is that explanation is just part of what makes a property elite. This would rest on explanatory metaphysicians embracing something akin to Schaffer's (Schaffer, 2016) proposed solution to the inference problem by stipulating that explanation plays the role of determining what properties are valuable; asserting that it is the business of explanation to evaluate. But again, my complaint about such a move is that it rests upon assuming a connection between eliteness and explanation in a similar way to the realist's assumption of a connection between eliteness and reality. Holding fixed the impermissibility of trivial explanations, the challenge to explain why explanation determines what is valuable is not satisfactorily answered by simply affirming that explanation plays this role. Such a move rests upon merely stipulating a connection between explanation and value and not

explaining it, which if acceptable allows the realist to solve the problem by stipulating a connection between objective structure and value.

The only way to make this argument work, I think, is to drop the non-triviality requirement on explanation and affirm that this kind of explanation constitutes a benign sort of circularity. Yet, this rests upon accepting trivial explanations as satisfactory in a way entirely counter to the objection. To reiterate, the objection first posed to realism explicitly includes a clause prohibiting trivial explanations. If explanatory normativity is obtained only by virtue of trivial explanation, then the realist is free to revive any of their trivial explanations and regard them as satisfactory. Explicitly, if the explanatory metaphysician is allowed to claim that properties with their elite status explained are elite because being explained is part of what makes a property elite, the realist is allowed to claim that natural properties are elite because being natural is part of what makes a property elite. Regardless of the plausibility of this line of reasoning, the objection itself is chiefly concerned with the exclusion of trivial explanation and therefore it offers no plausible way out for explanatory metaphysicians.

Additionally, it is also worth noting that the arguments pertaining to explanatory normativity, both in conjunction with anti-realism and when considered independently, are iterations of the following question: if something is explained, is it explained that it is explained? This is akin to Sider's *Purity Principle*: 'fundamental truths involve only fundamental notions' (Sider, 2011, p. 106) or as phrased by Dasgupta, 'if it's a brute fact that the X facts explain the Y facts, don't we still have facts about Y on the ground floor' (Dasgupta, 2017, p. 317 (footnote 21)). Dasgupta responds to some of these concerns (Dasgupta, 2014), but my argument here is that such a response is liable to the same kind of criticism he levies against realism. Most significantly, he remarks that this kind of

question stems from a general worry resulting from doing metaphysics in terms of explanation, and not from anything about normativity nor anti-realism specifically (Dasgupta, 2017, p. 317 (footnote 21)).[19] This line of reasoning does not offer any respite to the explanatory metaphysician as the challenge can be reformulated for those who approach metaphysics in terms of explanation: if metaphysics ought to be done in terms of explanation, there ought to be an explanation for why we should approach metaphysics in terms of explanation. There is no explanation, outside of trivially asserting that metaphysics is an explanatory process, for why we should approach metaphysics in terms of explanation. Those who believe some parts of metaphysics cannot be explained are free to make that contrasting assumption.

Overall, the most succinct way to summarise my argument against explanatory metaphysics is that just calling some property explained does not make it elite any more than calling someone Moore makes them multiply.[20] Be my guest – posit that explanation determines what is and is not elite if you want. But play fair when playing fair. Do not interweave explanation with value until you have shown explanation to be the deciding factor for what is and is not valuable. Explanatory normativity rests upon conferring a special status upon explanation while providing no explanation to support such bestowal. Explanatory metaphysicians may well claim that explatural properties are obviously more deserving of value, or that value is innately tied to explanation, or that metaphysics is

[19] The objection in this paper is logged against normative non-naturalism. I have replaced naturalism with the relevant counterpart, anti-realism.

[20] There are poignant parallels to be drawn between G. E. Moore's meta-ethics of normativity and Dasgupta's objection. Dasgupta notes this in (2018, pp. 294-295). There is also an interesting dialectic similarity in Moore's reply to scepticism and the development of dogmatism by James Pryor (2000), 'For Moore doesn't seem ready or able to offer any considerations at all in favor of the claim that he has a hand—even defeasible, ampliative considerations—without begging the question against a skeptic who refuses at this stage of the dialectic to grant the existence of the external world.' (Pryor, 2000, p. 518).

simply explaining. But these claims are comparable to realist accounts that are deemed unacceptable: natural properties are obviously more deserving of value, or that value is innately tied to the metaphysics of the world, or that metaphysics is simply joint carving. The former set of trivial explanations cannot be permissible if the latter are not.

As a result, to avoid explanatory metaphysicians undermining themselves, explanatory normativity cannot be enforced as a universal criterion on what properties do and do not have value. It can only apply when some degree of primitive normativity is permitted that allows for the unexplained value of explanation. But this licences the permissibility of unexplained value that undercuts any objection to realism on the grounds that their value rests on primitive normativity. Both realists and explanatory metaphysicians are free to make their respective assumptions; the realist is free to assume, without further explanation, that the elite properties are simply those that align with objective structure (or any other objective posit) and the explanatory metaphysician is free to assume, without further explanation, that the elite properties are simply those that have their elite status non-trivially explained. Both views require their respective primitives, and it cannot be objectionable to assume that value trivially arises from objective structure if one is also assuming that value trivially arises from explanation.

Aside from exonerating realism, a conclusion I aim to draw from this discussion is the prevalence of primitive normativity and unexplained value. Realists, anti-realists, and explanatory metaphysicians require something akin to it. Where for the former two, primitive normativity can be accepted without much issue, explanatory metaphysicians need to adapt their core commitments. I have here suggested that they restrict the application of their principle to a particular kind of value that only obtains when it is explained. Perhaps there is a version of explanatory normativity that suitably incorporates

primitive normativity, say by weakening it to allow for the unexplained value of explanation, or perhaps a sufficient distinction can be drawn between first and second-order explanations to avoid explanatory normativity applying to itself.

An approach of this latter sort would rest on imposing a criterion for non-triviality for first-order explanation whereby natural properties must be fairly explained but explatural properties (and why we must explain) can be 'unfairly' explained and allowed to be trivial. Enforcing this distinction would require support, but there might be a case to be made for it. Importantly, this approach must be sufficiently nuanced such that it also rules out the realist making a similar kind of move. If trivial explanations are permitted here, why are they not permitted for the naturalness of natural properties? Likewise, the second-order demands to explain why we either need to explain or why there is no explanation seem to go together. If we no longer need to non-trivially explain why we need to explain, then the realist needs no longer non-trivially explain why they cannot explain.

Other options might be worth pursuing but I can foresee no route that renders permissible triviality or primitivism for explanatory metaphysicians while ruling out a similar move for others; something that fundamentally rebukes the objection to realism, vindicates the realist of any dialectical complaints about defending their view solely on their own terms, and quashes accusations that primitive normativity is implausible.

4.0 FURTHER DISCUSSIONS

4.1 META-ETHICS: NORMATIVE NON-NATURALISM

Before concluding there are three areas that warrant further discussion. The first, which will be examined in the remainder of this section, is the consequences for related arguments in Meta-Ethics. The second, in §4.2, is the pragmatic assumption that underlies

Dasgupta's argument. The third, in §4.3, is the viability of connecting phenomenal acquaintance with value as a response to the absence of explanation.

To begin, a general form of the problem realism faces from explanatory normativity features prominently in areas of meta-ethics. For example, Christine Korsgaard writes about the challenges that come with positing facts that exist beyond any individual's mind about what is a reason:

> 'The difficulty with this account in a way exists right on its surface, for the account invites the question why it is necessary to act in accordance with those reasons, and so seems to leave us in need of a reason to be rational.' (Korsgaard, 2009, p. 53)

Assigning some objective metaphysical posit with the status of being a reason lacks the requisite support for being rational – it fails to account for why we *ought* to act rationally. The realist may well claim that the reason is that objective posit, but this answer does not explain why we must act in accordance with them. This answer only works when one first assumes that the metaphysical posit is action-guiding before showing why it deserves to be action-guiding.

Dasgupta has himself put the same argument forward against normative non-naturalism, the view that the good is both unique and irreducible. In (Dasgupta, 2017) an argument is presented whose structure is a direct counterpart of what is discussed in this paper (without realism about naturalness as the target of the objection). The argument starts by recognising that if something is good then it implicitly commits one to the notion that it should be promoted. Then, it assumes a marginally adapted version of explanatory normativity: that if some feature of our world should be promoted, there must be an

explanation for why it should be promoted (Dasgupta, 2017, p. 305). The non-naturalist maintains that the reason for this normativity, the reason that we should promote goodness, is due to some metaphysical feature of goodness; it should be promoted because of some brute, metaphysically objective, primitive fact. If we call this feature P, then the argument can be presented as follows:

1. Assume normative non-naturalism (that goodness is P)
2. Assume explanatory normativity.
3. From 1 and 2, there must be an explanation for why P should be promoted.
4. There is no explanation for why P should be promoted.

C. Goodness is not P (non-naturalism is false)

Non-naturalists derive their normativity from some inherent metaphysical posit; that posit, P, is the reason why those acts should be promoted. Dasgupta raises the problem that for this feature to be goodness there must be some explanation for why it should be promoted, and simply asserting that the property is goodness provides no reason to promote it:

> 'To paraphrase David Lewis, be my guest—posit all the sui generis whatnots you like. But play fair in naming your whatnots. Don't call any alleged feature of reality "goodness" unless you've already shown that you have something we should promote.' (Dasgupta, 2017, p. 301)

With this version of the argument outlined, the same dialectical division applies. If non-naturalists are permitted to make non-naturalist assumptions, they can simply reject the second premise and not assume explanatory normativity. The replacement is again a primitive form of normativity that allows for goodness to be promoted even if there is no

explanation for why it should be promoted. Non-naturalists might not be able to fairly explain why P should be promoted, but there is no requirement for them to do so. There is an obvious answer to why P should be promoted for non-naturalists: P is Goodness! There is, of course, no explanation for why goodness should be promoted; that is implicit in the approach to normativity I am suggesting. Then following this, the suggestion is to analogously accept P as possessing intrinsic value to explain why adopting primitive normativity is the correct stance. As the factor that makes P normative is constricted to the intrinsic value inherent in P, the absence of explanation is explained because nothing outside of P can be invoked to explain why it should be promoted and therefore, no 'fair' explanation is obtainable. When the non-naturalist is freely permitted friendly assumptions, an appeal to intrinsic value as understood as a kind of primitive normativity - for which there is no explanation - counters the internal criticism.

What remains from here is to defend non-naturalism on non-defensive terms. Once again, the two kinds of responses can be presented. First, it stands that the alternative view, naturalism, also fails to explain why their normative features should be promoted. Second, the demand itself rests on a normative claim about what should and should not be promoted that is itself unexplained.

To start with, naturalists replace the metaphysical posits with features outside of goodness itself to account for why goodness should be promoted. This could be features like happiness or pleasure and they employ these to satisfactorily answer the explanatory challenge: good acts should be promoted because they bring about more happiness. But here we can ask, what makes happiness so special? Run the same argument:

1. Assume normative naturalism (that goodness is happiness)

2. Assume explanatory normativity.
3. From 1 and 2, there must be an explanation for why happiness should be promoted.
4. There is no explanation for why happiness should be promoted.

D. Goodness is not P.

I imagine that just as non-naturalists would decry the original objection by stating something like: 'it is obvious why P should be promoted', the naturalist would decry this argument by stating something like: 'it is obvious why happiness should be promoted'. But this is no explanation. This only works when one assumes that we *ought* to promote actions that bring about happiness. An act simply bringing about happiness tells us nothing as to why these acts ought to be promoted unless one assumes that the only acts that should be promoted are those that bring about happiness.

Then for the meta-argument: if it is not just happiness but the fact that *happiness explains* why certain acts should be promoted, then the burden again sits with explanation. Why should only those acts that explain why they should be promoted, be promoted? What is it about explaining that instils normative force into an act? Ultimately, the naturalist stands on equal footing with the non-naturalist with respect to satisfying the demands of explanatory normativity. Neither can adequately explain; the non-naturalist requires the promotion of P to fall freely from P while the naturalist requires the promotion of happiness to fall freely from happiness. In response, both should accept primitive normativity and the unexplained – intrinsic – value of their respective normative posits. Interestingly, Dasgupta is aware of this kind of worry for non-naturalism:

> ‘I don’t claim that naturalism avoids the […] argument, just that the matter is unclear and deserves scrutiny. But I do claim that the naturalist must play fair. She can’t say, ‘Well, my theory is that goodness = pleasure; hence, since it’s obvious that goodness should be promoted, it follows that pleasure should be promoted’—that’s not playing fair! No, she must explain why pleasure should be promoted without assuming that pleasure = goodness (or show that no explanation is needed).’ (Dasgupta, 2017, pp. 317-318)

My argument does not stop with naturalism but extends to explanatory normativity itself. The notion of playing fair extends to explanation. Dasgupta makes an implicit and unspoken assumption when requesting that we play fair and assuming explanatory normativity: that we ought to play fair! I ask that the normativity claimed to be inherent in playing fair also be subject to the same demand it expresses. Play fair in playing fair; it must be explained why we must explain without assuming that we must explain. Dasgupta asserts that denying the second premise of his argument is an unattractive bullet to bite (Dasgupta, 2017, p. 307); (Dasgupta, 2018, p. 287). In doing so he is explicitly expressing preference for a particular view on normativity; that for it to be an attractive position it must conform to his explanatory demands. But again, this is no explanation. It is merely an assumption that explanation is inherently tied to normativity; that value must be explained. It is all well and good to tell me that this explains that, but my question is: why does that matter? Why does one need to explain anything for value to obtain?

4.2 Pragmatic Motivation

Second, it is worth plainly highlighting that Dasgupta’s position is motivated by pragmatism. In a draft paper, ‘Undoing the Truth Fetish: The Normative Path to Pragmatism’, there is similar opposition to realism, but this time in the domain of

language and thought. One notable difference is that when conceding that the realist might escape that problem by adopting something akin to primitive normativity there is further reasoning that aims to express why this move is not credible. In short, where Button (Button, 2013) outlines the epistemic opposition to realism – that it entails a strange commitment to the fact that everything we understand as true might be false - Dasgupta is extending this complaint by attaching a normative consideration to the strange commitments a realist makes. Even if the realist's primitive properties do exist, there is no reason to care about them.

To see this, a situation is presented considering the truth and reference relations of the words 'birds' and 'pigs'. Where 'pigs' refers to pigs and 'birds' refers to birds; 'birds fly' is true if and only if birds fly and 'pigs grunt' is true if and only if pigs grunt. Then, the realist also instantiates a relation between the word and world that is entirely independent of us and yet still normatively significant; call these metaphysical relations real-reference and real-truth. The problem that arises is that the "real" relations might not satisfy the initial truth and reference relations set out for 'birds' and 'pigs' as these "real" relations are entirely consistent with the possession of some objective metaphysical relation. Therefore, it could well be the case that 'pigs' refers to pigs while it really-refers to birds and in this situation, 'pigs fly' is false but really-true. Dasgupta's stance is that this sort of 'error' should be of no concern to us:

> 'It's that this kind of "error" shouldn't bother us in the slightest. Imagine an oracle informs us that our ideal theory is true but Really False. It tells us that "Pigs fly" is false (obviously) but Really True. Why should we care? Our theory is ideal by our standards and does everything we want of it—it's even true! All we've been

> told is that it lacks some extra primitive property, but so what? That's a minor curiosity at most' (Dasgupta, 2020, p. 26).

On first appraisal, this argument appears persuasive. Yet, there is much to unpack here.

To start with, it is worth explicitly stating that the realist is not affirming the claim that pigs fly is true, only that it could be. No one is suggesting that we scrap all modern science for this sake, but rather that all modern science might in principle turn out to be incorrect. That is not to say that it is, but it is to say that all our theorising may be incorrectly describing the world. A basic tenet of realism is that as the world is entirely independent of us, it may be the case that we are radically incorrect in our understanding of it, but this is no objection. It is merely a reiteration of what realism is. Even if there is a global divergence ultimately, a realist stands their ground on this issue; unlike pragmatists, they do not dilute reality's objectivity to appease our epistemic limitations.[21] They simply accept that we might be wrong. To counter any lingering concerns here, it should also be noted that the realist is free to assume, as a primitive part of their view, that although we might be incorrect, there is some degree of correlation between those properties that are objectively valuable and those that are identified by science.

On the question of why we should care the realist has a modest answer. If you take reality to be altogether independent of us, and theorising is the activity of saying preferably true things about reality, then it certainly warrants care. Here, the pragmatist will disagree as they take inquiry to aim at serving our interests and that the truth of our theories is little more than what this inquiry arrives at. The realist has no reason to adopt the pragmatic approach to theorising and the pragmatic metaphysics it leads to. Instead, they start with

[21] I have discussed the envisioned robust theorising/world distinction in more detail on pp.56-57.

a realist understanding of theorising and accept the realist metaphysics it leads to. This kind of relationship, between an understanding of truth and a conception of metaphysics, has a well-documented history. For example, Coherentist theories of truth face the charge of leading to idealist metaphysics (BonJour, 2000, pp. 146-147).

Lastly, when pressed with the claim that it would be 'madness' to accept the normative significance of these "real" relations and primitive properties (Dasgupta, 2020, p. 26), I ask: on what authority has such a decision been made? Why should normative significance not be assigned to those primitive properties? If it is so obvious that the only notions worth valuing are the ones that we can explain, I ask you to play fair in making such a normative claim. Simply stating that primitive normativity is absurd is no objection and claiming that normativity must be explained is little more than a reiteration of the initial objection; that an explanation is required for why any set of properties that claims to be normative, is normative. That objection has already been addressed.

4.3 Phenomenal Acquaintance

Finally, there is perhaps one strategy that undercuts explanatory normativity by appealing to phenomenal acquaintance. The core claim is that because we, as theorisers, are in a phenomenal association with some properties it is obvious why they deserve value. I bring up this idea because it is one Dasgupta suggests as a possible resolution in the conclusion of his paper (Dasgupta, 2018, pp. 318-319). Consider again, happiness and the question of explaining why it should be promoted. One way to answer this question is to state the obvious: that happiness should be promoted because it is good! For some things, call these *acquaintables*, we understand what they feel like and there is no need to explain why they are valuable. Thus, this claim allows for a second-order explanation for why

there is no first-order explanation; being phenomenally acquainted with happiness reveals why one need not explain why it should be promoted.

Crucially, Dasgupta thinks that this kind of argument is not viable for the realist. Properties like happiness or green, etc. are acquainted with us in a way that natural properties are not. This is not to say that we could not conflate the natural properties with the acquaintables but that the properties proposed by realists, such as that of Lewis's and Sider's objective structure, are not ones with which we are in any way acquainted.[22] This claim could be contested. Lewis's account of natural properties often appeals to objective similarity and similarity between two properties could be something we can be acquainted with. But the problem with this view is that for the realist to make this kind of claim they would need to posit an objective relation of acquaintance and the elite properties being those that stand in this relation. Resting the basis of objective reality on being in a phenomenological relation makes the metaphysical posits dependant on theorisers:

> 'On this view the privileged properties, the ones on which one's metaphysics is built, are not the properties of fundamental physics but rather those revealed in perception. Perception reveals the maximally "elite" properties; other properties like those of physics are derivative upon these' (Dasgupta, 2018, p. 319)

The resulting kind of idealist worldview will not be appealing to many. It would certainly be unappealing to a realist like Lewis who has suggested that fundamental physics reveals the most natural properties (Lewis, 1994a, p. 474). I believe there is a viable nearby view

[22] A case could be made for acquaintance with higher-order conceptions of naturalness (Schaffer, 2004).

whereby the value of elite properties obtains without explanation while avoiding the resultant idealism.

That view shares an important attribute with Dasgupta's: explanation and understanding come apart. These two notions are often conflated, and a simple and popular view is that to understand something is to be able to explain it. But cases like that of phenomenal acquaintance, as indicated by Dasgupta, provide a counterexample by apparently allowing us to understand something without being able to explain it. It seems a given that if you can explain something, you understand it, but the contrary claim, that if you understand something, you can explain it, is less clear-cut. Accepting that explanation and understanding come apart shifts the debate away from the demand to explain value and toward the demand to understand value. Although on the realist view outlined in this paper, I have argued that the realist should accept their inability to explain the value of their properties, there is a nearby realist view that can grasp at understanding if not explanation.

Take those properties with intrinsic value. For the realist, they are those objective properties (this argument would also work for an anti-realist view that accepts their properties to possess intrinsic value). Then, the core claim for this new view is that solely in virtue of those properties possessing intrinsic value, we can understand that value. I attribute this idea to Forrest (Forrest, 1991) whereby we can understand that which is intrinsically valuable simply because it has that value. To substantiate this claim, Forrest explores several cases where understanding exists without explanation and that understanding is derived simply from the fact that the thing in question possesses intrinsic value. By accepting the intrinsic value of their elite properties' realists are also capable

of accepting that we understand why natural properties are valuable simply by virtue of them being valuable; it is just obvious to realists why some properties, those carved at the joints, ought to feature in theories. There is no explanation for this, but that does not entail that we cannot understand why they matter.

To be clear, the parallel drawn with Dasgupta's proposed phenomenal view is invoked here to substantiate separating explanation from understanding and importantly, this realist view entails no idealism. Those objective properties with intrinsic value confer understanding solely by virtue of the fact that they possess that value. They are not valuable because we are phenomenally acquainted with them. On the contrary, they are valuable, and we can understand that simply because they are valuable. There remains ignorance about which properties are the intrinsically valuable ones – but we can understand why they should feature in our theories simply because they possess that value. Realism still requires a leap of faith in accepting that the properties scientists discover correspond to those with intrinsic value. But again, pointing that out remains no objection – that is just what realism is.

Regardless of the plausibility of this kind of view, it does offer an interesting alternative for a suitably minded realist; one who wants to retain the claim that even if they cannot explain why natural properties possess value, they still understand why they are valuable. I will leave any further exposition of the view to any such suitably minded realists. It exists here simply to note that it is a possibility in nearby conceptual space.

5.0 CONCLUSION

To resolve the problem of the cheapness of properties, realists are free to adopt primitive normativity to account for the value of the properties they privilege. Where Dasgupta

assumes that explanatory normativity adjudicates what properties are and are not elite, in contrast, I have argued that explanatory normativity does not determine what properties are and are not elite nor does it decide what kind of value is valuable. To once more paraphrase Lewis in response to Putnam's objection to realism:

> 'We are in the presence of a paradox, as surely as when we meet the man who offers us a proof that there are no people, and in particular that he himself does not exist.[23] It is out of the question to follow the argument where it leads. We know in advance that there is something wrong, and the challenge is to find out where. If the paradox-monger is good at his work, we stand to learn something; and indeed, I think that [Dasgupta's Detonation] affords an important lesson.'[24]

Explanatory normativity teaches us a lesson about what metaphysical theses suit what kind of value and includes acknowledging primitive normativity as part of a novel normative account permitting realists to derive unexplained value from their metaphysical claims.

The dialectical criticisms that typically plague realist arguments have also been circumvented. Anti-realism has been shown to also require primitive normativity and therefore if primitive normativity is problematic both realism and anti-realism face these challenges jointly. As I have likewise argued that explanatory normativity rests upon unexplained assumptions about value, the conclusion is to recognise not only the permissibility of primitive – unexplained – normative connections but also their ubiquity.

---

[23] See, Unger (Unger, 1979).

[24] The original passage concerns Putnam's model-theoretic objection to realism (Lewis, 1984, p. 221). I have replaced 'Putnam's Paradox' with Sider's (Sider, 2022) apt terminology 'Dasgupta's Detonation'. There are similarities between Dasgupta and Putnam, notably in their original steadfast endorsements of realism and subsequent desertion via pragmatism.

The consequences of this extend beyond the discussion on realism and as shown, apply equally to comparable debates in meta-ethics, delineate pragmatic from realist motivation, and allow for potential alternatives that bypass explanation but still obtain understanding. Ultimately, this paper refutes the view that value *must* be explained.

BRIDGING PAPERS I, II, & III

Papers I and II both draw heavy inspiration from Lewis's 'fair play' argument. The first paper of this thesis is chiefly a work of exegesis with the aim of presenting a viable interpretation of the argument; it results in the explanatory principle. A similar account can be found in the second paper. It draws heavily on work done by (Dasgupta, 2018) who, also inspired by Lewis's fair play argument, defines an explanatory demand for normative claims; something I define as explanatory normativity. Due to the similarities between them, the purpose of this bridging section is to briefly compare the two papers, make some clarificatory remarks, and address what might be viewed as a potential conflict between claims I have made in both papers. The discussion ends by noting the overarching dialectical progression that leads to the third paper on meta-ground.

To start, in Paper I §2.1 and §4.2, I argue that Schaffer (Schaffer, 2016, p. 580) is uncharitable to Lewis. To briefly reiterate, Schaffer interprets Lewis's request for fair play as the rejection of primitives and then, due to the presence of primitives in Lewis's own metaphysics, dismisses the argument entirely. In paper II §2.1, I note that Dasgupta's argument (Dasgupta, 2018, p. 286) against (Lewis's) realism is ironically inspired by Lewis's fair play argument. Both arguments are similar with respect to the fact that they interpret the fair play argument in a way that comes into conflict with another claim proposed by Lewis. As such, one might think the accusations of being uncharitable I log towards Schaffer also apply to Dasgupta. But I do not think this is the case. There are substantial differences between their arguments – one as a rejection of primitives and the other as a demand for an explanation – and a conflict is not so obvious in Dasgupta's case. Schaffer's argument conflicts with an explicit passage of Lewis's where he outright

allows for the legitimacy of primitives, while Lewis makes no explicit claim allowing for unexplained features. Further, Dasgupta does not assert that his explanatory criterion is one accepted by Lewis but instead uses the fair play argument as inspiration for an explanatory criterion.

I, however, argue that the most suitable interpretation of the fair play argument is as an explanatory criterion. If explanatory metaphysicians, such as Lewis, are also inclined to accept Dasgupta's demand to explain the value of elite properties, then there seems to be a conflict. My explanatory interpretation of fair play differs from explanatory normativity as explaining *why* natural properties have value is distinct from explaining *how* natural properties do their work. But the two explanatory demands are sufficiently similar such it requires no great effort to translate the claim of one over to the other. Phrased using the corresponding terminology, consider, for example, objective structure demarcating the property of "Green". The role objective structure plays in a normative theory is to make theories that use the property "Green" better than those that use other unfitting properties such as "Grue". Here the explanatory principle applies as there ought to be an account of how objective structure makes those properties valuable. From Dasgupta, we find that no good answer is forthcoming.

As a result, there appears to be a conflict between two of my arguments. In paper I §4.2, I suggest that if there fails to be an explanation for how natural properties play any of their given roles in a theory then they ought to be discarded. Then, in paper II §3.1, I propose that as realists about natural properties cannot explain the value of their properties, they ought to reject explanatory normativity in favour of an alternative non-explanatory account. Although incompatible, these claims are not intended to be understood as conjoined in a systematic metaphysical picture. In paper I, the claim is: *if*

you are an explanatory metaphysician, you can and ought to reject any element that plays a role that is not explained (for which natural properties appear to be one). The claim I make in paper II is: *if* you are a realist, you can and ought to unproblematically reject explanatory normativity. In both cases, I make no further claim about being an explanatory metaphysician or a realist.

It would be problematic if one tried to accept both concurrently; demanding fair play while allowing for the value of realist properties to obtain unfairly. Lewis's proposed systemic metaphysical outlook appears to do just this, and accordingly, it leads to a dilemma. On the one hand, Lewis could retain realism and reject explanatory metaphysics but in doing the motivation for many of his other metaphysical proposals would be lost. On the other hand, he could retain the explanatory principle and reject realism but as realist assumptions are pervasive throughout Lewisian metaphysics vast revisions would be in order. Lewis thought van Inwagen had put him in a difficult position with his *tu quoque* for singletons such that he remarked the choice of outcomes is 'almost entirely between the fires, not between the frying pans' (Lewis, 1991, p. 36). This new dilemma is even worse, and the inferno threatens to engulf extensive and important parts of Lewis's metaphysics; the choice to avoid the multi-dog-headed monster Scylla by sailing via the whirlpool Charybdis must be made when a hungry Cerberus baths in the maelstrom.

Yet, these ashes are fertile space for new metaphysics to grow. There already exist anti-realist versions of parts of Lewis's metaphysics (e.g., (Hall, 2015)) and many other places of Lewis's metaphysics have been revised with the aim of being more compatible with explanation in mind (e.g., (Cohen & Callender, 2009)) Whether Lewis would find these moves appealing is another question. He remarks that if the standards of simplicity and strength used to determine the best system come from us:

> ‘[S]ome ratbag idealist might say that if we don't like the misfortunes that the laws of nature visit upon us, we can change the laws-in fact, we can make them always have been different-just by changing the way we think! (Talk about the power of positive thinking.) It would be very bad if my analysis endorsed such lunacy.’ (Lewis, 1994a, p. 479)

While Lewis does not appear to be receptive to ‘ratbag idealists’, Hall has argued that the anti-realist approach is most compatible with Humean laws (Hall, 2015, p. 38) or maybe the resulting anthropocentrism is unproblematic (Jaag & Loew, 2020). Either way, there exists an easy way out of the dilemma by discarding realist commitments.

The alternative is to either revise the explanatory principle or reject it entirely. In paper I §4.2 I argue that the latter is an unlikely choice for Lewis. However, the former might offer a means of escape. Importantly, the explanatory principle would need to be adapted such that the demand to explain how elements play their roles is maintained but realist elements are made exempt from this demand. This move can be substantiated by the arguments already posed in §3.2 and 3.3 of paper II. There I argue that both anti-realists and explanatory metaphysicians fail to adequately meet the explanatory demand. Generally speaking, the conclusion to draw is that not everything can be explained. Even explanatory metaphysicians must take the role explanation plays as primitive and realists are free to make the same kind of move; taking the role realism plays as unexplained. Once some unexplained notions are permitted, drawing a line between what does and does not require explanation becomes difficult and a realist can look to exploit this. A viable explanatory meta-philosophical view that permits exceptions allows the realist to make “real” elements exempt from the demands.

There might be a sharp boundary between the permissibility of unexplained explanation and the impermissibility of unexplained realism such that it prevents the realist from making this move. This stance will hinge on the provision of a second-order explanation that accounts for the demands of the first; explaining why an explanation is required or why no explanation is required. One might try to put explanations of explanations as uniquely not in need of an explanation because they are already explanations. Regardless, this is just a brief outline of a possible response; however, questions about higher-order explanations inspire the next paper that looks at meta-ground. What explains grounding explanations? While the discussion comes to focus on ground, it has implications not only for metaphysical explanation but for all kinds of explanation.

PAPER III:
UP-AND-DOWN META-GROUND

ABSTRACT

When a grounding connection is established, the question arises: is it grounded? This is the question of meta-ground and while most accept that the answer is yes – grounding connections are grounded – there is little consensus on what facts are responsible for them. This paper argues that the best answer is to ground the grounding connection in both the ground and grounded fact; a view I call Up-and-Down Meta-Ground. In support, I demonstrate how the view preserves the appealing features of various competing accounts while avoiding their pitfalls. Moreover, I argue that Up-and-Down Meta-Ground has two unique virtues, being both the only fully explanatory account of ground and the only view compatible with Humean nomic principles.

OVERVIEW

In §1, background is provided on ground, its connection to fundamentality, and the question of meta-ground. In §2, three already proposed solutions to the question of meta-ground are outlined and the advantages and disadvantages of each are considered. They include assessing whether the view leads to a problematic regress, whether the view is adequately explanatory, and compatibility with Humean or non-Humean nomic principles. In §3, Up-and-Down Meta-Ground is outlined and is shown to avoid the problems facing previously considered and in some cases being uniquely compatible with some desirable features of ground. In §4, some further considerations are explored when making temporal comparisons with grounding. Finally, §5 concludes by mapping out the various solutions to the question of meta-ground and what theoretical virtues they are compatible with. Up-and-Down meta-ground is argued for as the optimal solution; I also highlight its unifying nature.

1.0 BACKGROUND

1.1 WHAT IS GROUND?

It is instructive to begin by providing an overview of what ground is; such background will, however, be brief as I will be assuming a broad familiarity with the concept.[25] With that said, Fine defines ground as a non-causal kind of explanation (Fine, 2012, p. 37), Sider defines ground as a non-modal metaphysical connection between levels (Sider, 2020, p. 747), and Schaffer defines ground as comparable to metaphysical causation. But, where causation connects the world across time, grounding connects the different levels of the world (Schaffer, 2012, p. 122).[26] These definitions put ground in a similar camp to a broad range of metaphysical notions such as supervenience, dependence, mereological parthood, identity, and reduction.

To classify ground independently of these comparable metaphysical posits is controversial and the objection has been raised that there is little more to ground than the metaphysical notions already prevalent in metaphysics. Wilson (2014) argues that there is no theoretical work for grounding over and above the level-connecting relations already utilised in metaphysics.[27] To distinguish ground, and avoid the brunt of this objection, explanation can be invoked alongside the level-connecting aspect of other metaphysical relations; explanation is the key identifying feature of ground that distinguishes it from

[25] Influential papers in the literature include (Fine, 2001), (Schaffer, 2009), and (Rosen, 2010).

[26] For an extensive discussion on levels see (List, 2019), and specific to the ground theoretic hierarchy see (Werne, 2020). Additionally, even though it is generally thought that causation and grounding are distinct, both (Schaffer, 2016) and (Wilson, 2018a) offer detailed accounts of the numerous parallels between them while (Bernstein, 2016) provides opposition to any such conflation.

[27] Moreover, she rightly points out that the story of motivation for grounding as a novel kind of metaphysical relation is ill-founded.

other metaphysical level-connecting relations.[28] Therefore, the understanding of ground I will be working with in this paper is as follows:

Ground: an explanatory dependence relation between different levels of reality.

Before proceeding, some routine assumptions about ground are required. I assume grounding to be an asymmetric, irreflexive, transitive relation between a plurality of facts and another plurality of facts; in any case, where it is not specified otherwise, being grounded will refer to a case of full grounding. Each of these assumptions is widespread. For example, transitivity and irreflexivity axioms are assumed by (Fine, 2010, p. 100) and (Correia, 2010, pp. 261-262). This is not to say that they are uncontentious. Transitivity is disputed by (Schaffer, 2012), irreflexivity is disputed by (Jenkins, 2011), asymmetry is disputed by (Barnes, 2018), and all three are disputed by (Rodriguez-Pereyra, 2015). Additionally, disputes as to whether grounding is solely between facts can be found in both (Schaffer, 2009) and (Jenkins, 2011), and disputes as to whether grounding is best expressed utilising operators or relational predicates to connect the explanation to the explained can be found in (Raven, 2015, pp. 324-325).

1.2 Fundamentality & Ground

Much of the debate about the proper way to characterise ground is superfluous to this paper as this paper is concerned with one of ground's core attractions: characterising the fundamental. Schaffer expresses this plainly when he distinguishes between two classes of entities: the derivative and the fundamental. He defines the former as that which is grounded and the latter as that which is ungrounded (Schaffer, 2009, p. 373).[29] Grounding

[28] See (Dasgupta, 2017), particularly p.75, for a related argument.

[29] Other examples of the link between fundamentality and grounding can be found in (Bennett, 2011, p. 27) and (Sider, 2011, pp. 105-106). Its plausibility is also discussed in (Dasgupta, 2014).

provides us with a means of determining what the fundamental entities are: those without a ground. With this in mind, a further question arises. Are facts about grounding fundamental? If, for example, the fact [*a*] grounds fact [*b*], is the grounding connection [[*a*] grounds [*b*]] grounded? This is commonly understood as the question of meta-ground (Litland, 2020):

> *Question of Meta-Ground*: what, if anything, grounds the grounding connection [[*a*] grounds [*b*]]?

The question of meta-ground is about what grounds grounding, and it puts grounding theorists in a serious predicament. But, broadly speaking, there are two simple answers. Either grounding is grounded and is therefore not fundamental; or grounding is ungrounded and is therefore fundamental.

1.3 Primitivist Meta-Ground

Primitivism is the view that grounding connections are not grounded and are therefore fundamental:

> *Primitivism*: the grounding connection is ungrounded.

The primitivist approach to meta-ground has few advocates and is generally not regarded as a viable option.[30] Briefly, I will note two prominent arguments concerning primitivism.

The first is based on Sider's Purity principle. Sider expresses Purity as, 'fundamental truths involve only fundamental notions' (Sider, 2011, p. 106). On accepting Purity no ungrounded fact can contain any fact that is non-fundamental. Leaving [[*a*] grounds [*b*]]

[30] Nicholas Jones, 'Ungrounded Ground' is cited as defending the view by authors ((Frugé, forthcoming); (Wilson, 2019)), but in correspondence, Jones has stated that he no longer advocates for the view.

ungrounded would also leave a fact that contains a non-fundamental fact [*b*] ungrounded. Anyone who desires their account of meta-ground to be consistent with Purity cannot leave the grounding connection ungrounded. Second, Bennett argues that 'whatever the (contingent) fundamental elements of the world are, they are open to free modal recombination' (Bennett, 2017, p. 190). Her argument principally rests on the claim that there is no reason to reject that the fundamental facts are freely modally recombinable and that this is not the case for non-fundamental facts as being grounded in another fact prevents that fact from being modally recombined.

Neither of these arguments should be taken as delivering a knockout blow to the primitivist. Purity is merely assumed based on intuition and even though she is sympathetic to it, Bennett points out that it is uncomfortably close to question-begging (Bennett, 2017, p. 190). The modal recombination argument too can be resisted and Wang (Wang, 2020) offers arguments for doing so. Further, Bennett's support for the principle consists in demonstrating that its refutation would lead to unpalatable consequences, but it is an open question as to whether the primitivist is willing to stomach them. Regardless of one's response to these arguments against primitivism, this paper is concerned with accounts of meta-ground that ground the grounding connection.

1.4 Anti-Primitivist Accounts of Meta-Ground

Anti-Primitivism is the overarching term for accounts of meta-ground that ground grounding; some fact(s) grounds the connecting fact [[*a*] grounds [*b*]]. Even though this is the foremost approach, there is little consensus regarding what facts ground the grounding connection. Anti-primitivism comes in many forms and this paper will explore the viability of four contrasting, but interrelated, approaches:

Downward Anti-Primitivism (DAP): the grounded fact grounds the grounding connection.

Upward Anti-Primitivism (UAP): the ground fact grounds the grounding connection.

Connectivist Anti-Primitivism (CAP): a general connection between the ground and grounded fact grounds the grounding connection.

Up-and-Down Anti-Primitivism (UDAP): both the ground and grounded fact ground the grounding connection.

Each of these approaches will be assessed primarily against three factors. The first factor relates to whether the resulting hierarchy of grounding facts leads to a problematic regress. The second factor relates to whether the account of meta-ground is explanatorily adequate. The third factor relates to compatibility with both Humean and Non-Humean nomic accounts of ground. Ultimately, I intend to argue for a holistic solution that combines the positives and avoids the pitfalls of the first three views when defining the latter. I will argue that UDAP offers the best solution to the problem of meta-ground.

2.0 GROUNDING GROUND

2.1 DOWNWARD ANTI-PRIMITIVISM

To start with, *Downward Anti-Primitivism (DAP)* is the view that the grounded fact grounds the grounding connection. Expressed with particular facts:

DAP: [*b*] grounds [[*a*] grounds [*b*]]

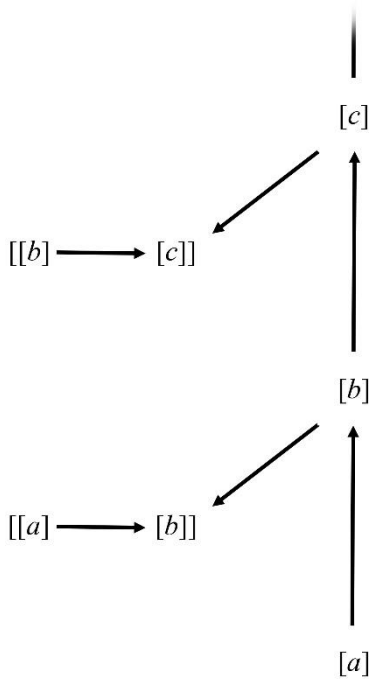

Figure 1 Downward Anti-Primitivism.

To clarify the figures used in this paper: letters represent facts, and solid head arrows represent grounding connections between facts. Where only a single arrow grounds a fact, there is an instance of full grounding and when more than one arrow grounds a fact, there is an instance of partial grounding. For DAP, there are three atomic facts [*a*], [*b*], and [*c*], the connecting facts [[*a*] grounds [*b*]] and [[*b*] grounds [*c*]], and the various grounding connections between them. Notably, [*c*] is grounded in [*b*], and [*c*] grounds [[*b*] grounds [*c*]]. It is the same for [*b*] being grounded in [*a*], and [*b*] grounding [[*a*] grounds [*b*]]. Finally, as [*a*] is ungrounded it is fundamental. (One feature of these diagrams is that the grounding connection between [a] and [b] is represented twice, first on the right with the vertical arrow and second on the left horizontally. The only difference between them (in

the case of DAP) is that a grounding connection between [b] and [[a] grounds [b]] is also represented on the left. The diagram is presented this way to allow for the visualisation of the relevant grounding connections whose depiction would otherwise be obscure).

This position is commonly attributed to (Fine, 2012, p. 76), as he states, 'that it is something about the grounded that explains why it is grounded' (but since Fine is concerned with the logic of essence it is unclear whether this entails a position on the logic of ground). The view has some intuitive plausibility, as one might think a connection between ground and grounded fact is underpinned by the fact that has been grounded. However, beyond intuition, there are few arguments in favour.

Additionally, DAP faces difficulties when dealing with a regress that arises for all anti-primitivists. Each instance of grounding generates a new fact about those facts grounding. For example, if [*c*] grounds [[*a*] grounds [*b*]], the fact [[*c*] grounds [[*a*] grounds [*b*]]] will be generated, and in line with anti-primitivism will need grounding. Then onwards *ad infinitum*, each instance of grounding must have its corresponding grounding connection grounded. While the existence of an infinite regress is not in itself problematic (Sider discusses several possible problems with the grounding regress and reveals many of them to be unproblematic (Sider, 2020, pp. 751-753))[31] a regress of grounding facts is problematic, according to Bennett, if it fails to 'ground out' in a fundamental fact (Bennett, 2011, pp. 30-31). Likewise, Wilson affirms that we ought to be able to identify a downward chain of grounding facts that terminates in a fundamental fact (Wilson, 2019, p. 7). The regress may continue indefinitely but it is only problematic when each iteration cannot identify the fundamental fact that sits at its end. I will present this as the need for

[31] The truth regress is an often-quoted case of harmless regress (Armstrong, 1989, p. 56); (Rabin & Rabern, 2016) discuss in detail the worries around well-founding grounding hierarchies.

the ‘ultimate ground’, the fact that eventually grounds all the facts in the regress, to be a fundamental fact; the regress may acceptably cease in a fundamental fact because fundamental facts are understood as just those without a ground.

Consider, for example, consider the following grounding regress:

IV. …

III. [[*d*] grounds [[*c*] grounds [[*a*] grounds [*b*]]]]

II. [[*c*] grounds [[*a*] grounds [*b*]]]

I. [[*a*] grounds [*b*]]

This is problematic because there is no ultimate ground; no fundamental fact sits at the end of the regress, and it ramifies endlessly without grounding out in some fundamental fact. We must be able to identify what the ultimate ground is and in cases like this, we cannot.

Next, returning specifically to DAP. If stipulating that the grounded fact grounds both the grounding connection and all subsequent facts in the regress that follow from the initial grounding, it generates the following grounding hierarchy:

IV. …

III. [[[*b*] grounds [[*b*] grounds [[*a*] grounds [*b*]]]].

II. [[*b*] grounds [[*a*] grounds [*b*]]]

I. [[*a*] grounds [*b*]]

Each successive fact is grounded by a non-fundamental fact and consequently, the hierarchy never gets acceptably ‘grounded out’. [*b*] is recurringly invoked to ground all other facts in the regress but [b] is not a fundamental fact and therefore, the regress is problematic. Bennett takes this to be one of the main problems with DAP. Perhaps, an

advocate of DAP could resist the problematic aspects of the regress or somehow circumvent it entirely (I argue later in §3.2 that the kind of regress faced here is unproblematic in almost all cases). But absent other endorsements for DAP aside from intuition, I will move on to other approaches.

2.2 Upward Anti-Primitivism

The second version of anti-primitivism is *Upward Anti-Primitivism (UAP)*. This is the view that the ground grounds the grounding connection. Expressed with particular facts:

UAP: [*a*] grounds [[*a*] grounds [*b*]]

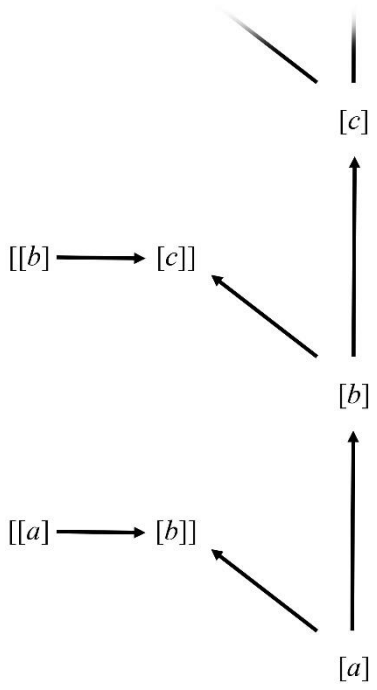

Figure 2 Upward Anti-Primitivism

In Figure 2, the connecting facts are grounded by the grounds: [b] grounds [[b] grounds [c]], [a] grounds [[a] grounds [b]], and [a] is the sole ungrounded (fundamental) fact. UAP is held by Bennett (Bennett, 2017, pp. 187-213), deRosset (deRosset, 2013), and Litland (Litland, 2017). Again, UAP has some intuitive plausibility as Bennett argues 'anything that builds a building fact thereby also builds the built entity involved in that fact' (Bennett, 2017, p. 196). Although 'building' is used by Bennett, since she intends ground to be a type of building relation, the intuitive sentiment remains. It is reasonable to think that the connection between ground and grounded fact is intuitively underpinned by the ground but both Bennett and deRosset are chiefly motivated by the desire to avoid a problematic regress. UAP does this by ensuring that each level of the grounding regress always grounds out in a fundamental fact:

IV. …

III. [[a] grounds [[a] grounds [[a] grounds [b]]]]

II. [[a] grounds [[a] grounds [b]]]

I. [[a] grounds [b]]

A 'bottom-up' hierarchy of grounding facts results. The regress is one that is always grounded out in the fundamental fact [a] and is therefore unproblematic. Regardless of the endless ramification, at each subsequent iteration of the regress the eventual answer to what grounds everything is always [a] and the ultimate ground a fundamental fact.

So, in favour of UAP is some intuitive plausibility and the avoidance of a problematic regress. In opposition, Dasgupta argues that the view does not properly capture the explanatory nature of grounding (Dasgupta, 2014, pp. 571-572). Take for example facts about activities engaged at an event to ground the fact that the event is of a certain kind: facts about people kicking a football, scoring goals, complaining about offside, etc.

ground the fact that the event is a football match. Similarly, facts about people presenting papers, asking questions, entertaining existential dread, etc. ground the fact that the event is a philosophy conference. People engaged in football-like activities grounds the fact that the event is a football match while people engaged in conference-like activities grounds the fact that the event is a philosophy conference. Now, the initial grounding question is:

> Why is this event a conference (instead of a football match)?

And the answer is: that people are engaged in conference-like activities. We ground one set of facts in another, [a] grounds [b]. Then, with this question answered the subsequent meta-ground question arises:

> Why do those activities make the event count as a conference (instead of a football match)?

This question asks what grounds the fact that some fact is grounded, what grounds [[a] grounds [b]]? What is it about people being engaged in certain activities that make an event be of a certain kind? UAP provides the same answer to both questions because what makes this event a conference and what makes those activities count as a conference is simply the fact that people are engaged in conference-like activities. Dasgupta contests UAP's answer to the second question for being explanatorily deficient. Being engaged in conference-like activities tells us nothing as to why those activities constitute a conference; simply asserting what facts constitute a conference does not tell us why those facts constitute a conference. The challenge is to explain why some facts make that event a conference instead of a football match, and simply reiterating that people are engaged in conference-like activities does not. An explanation needs to tell us why conferences are constituted by conference-like activities.

Both Bennett and deRossett offer similar responses to charges of explanatory inadequacy. deRossett's approach is to adopt a deflationary account of explanation that maintains that there 'is nothing interesting or informative to say about the nature or constitution of explanatoriness in general' (deRosset, 2013, p. 23). In effect, the explanation provided by the ground, although minimal, is still adequate; there is little more to a grounding explanation other than that which is expressed by the ground. Bennett's argument explicitly distinguishes between a metaphysical and an epistemic conception of explanation (Bennett, 2017, pp. 202-203). She claims that UAP captures the explanatory nature of grounding in a metaphysical sense as the explanation accounts for why those activities generate, produce, or make the event count as a conference (instead of a football match). What UAP fails to do is to explain in an epistemic sense as it does not shed any light, nor render intelligible, why or how the event occurred.

Even though Dasgupta and advocates of UAP might be interpreting explanation in different senses, the objection is not seamlessly rebuffed by deflating the explanatory force of grounding to mere metaphysical explanation.[32] First, it is not entirely clear that even in a metaphysical sense anything is properly explained by the ground alone. A metaphysical explanation is intended to express why [*a*] produces (makes happen or is responsible for) [*b*], but it is hard to see how it does this without making any mention of [*b*], or the connection between [*a*] and [*b*]. This kind of concern is reflected in (Dasgupta, 2014, pp. 571-573), (Wilson, 2019, p. 6), and (Thompson, 2019, pp. 322-323). Relatedly, Kim calls this 'the problem of explanatory ascent' (Kim, 2005, p. 107): how can one

[32] Dasgupta remarks about the difference in usage, 'It may be that Bennett or deRosset mean something different by "ground" than I do. In particular, if "ground" in their mouths is not constitutively tied to explanation then my remarks here do not engage with their view' (Dasgupta, 2014, p. 573).

account for the explanandum on a higher level when the explanans is limited to a lower level?

Second, renouncing the epistemic explanatory force of grounding leaves it indistinguishable from the myriad of other level-connecting metaphysical relations. Bennett and deRosset may well use supervenience or mereological parthood to express the connection between levels (and in doing so they need to respond to the charge that the grounding relation thus defined is redundant (Wilson, 2014)). Critically, no reply from advocates of UAP attempts to capture an inflated account of explanation. Instead, they have generally retreated to diminishing ground's explanatory force. This is particularly problematic here because one of ground's most attractive features is that it is explanatory. Those unwilling to accept such a concession will find UAP unacceptable. In turn, these concerns have motivated another form of Anti-Primitivism that aims to be suitably explanatory.

2.3 Connectivist Anti-Primitivism

The third version of Anti-Primitivism is *Connectivist Anti-Primitivism (CAP)*. This is the view that a principle relating the ground to the grounded at least in part grounds the grounding connection:

> *CAP*: a general connection between [*a*]-like things and [*b*]-like things grounds [[*a*] grounds [*b*]]

This is the typical approach, but the connection can be defined in many ways. It could be necessities, conceptual truths, or metaphysical laws, that explain why [*a*]-like things are connected to [*b*]-like things (and the ground [*a*] is also often invoked as a partial ground along with the connecting principle).

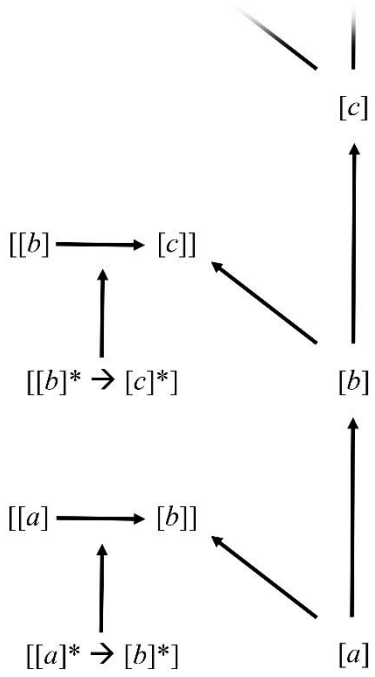

Figure 3 Connectivist Anti-Primitivism.

In Figure 3, there are facts [[a]* → [b]*] and [[b]* → [c]*] that represent a general connection between [a]-like things and [b]-like things. The '*' denotes the fact's likeness (replaceable with essence, necessity, principles, etc) and the '→' denotes a conditional connection: if [a]* things obtain then so do [b]* things. Finally, the connecting fact is grounded in both the ground [a] and the connecting fact [[a]* → [b]*].

Some prominent advocates of CAP include (Dasgupta, 2014), (Wilsch, 2015), (Rosen, 2017), and (Wilson, 2019). What unites CAPs is the motivation to provide an explanatory account of grounding by explicitly mentioning the connection between grounds and grounded. Returning to the previous example of Dasgupta's, the occurrence of conference-like activities makes the event count as a conference (instead of a football

match) because conference-like activities and conferences are connected by some principle linking the two (essence, modal necessity, etc.). Therefore, this view avoids explanatory complaints by explicitly mentioning the connection between [a] and [b] when grounding the fact [[a] grounds [b]].

Yet, while CAP is explanatorily inflated some notable consequences arise after evaluating the resulting grounding regress, where the ground and connecting principle ground all subsequent facts:

IV. …

III. [[[a]* → [b]*], [a] grounds [[[a]* → [b]*], [a] grounds [[a] grounds [b]]]]

II. [[[a]* → [b]*], [a] grounds [[a] grounds [b]]]

I. [[a] grounds [b]]

The connecting principle is reiterated at each stage and therefore, to be unproblematic, must either be grounded or fundamental. First, if the connecting principle is grounded in its instances then it is circular and Bennett argues that this violates the asymmetric nature of grounding, 'the general fact that if something a-like obtains, so too does something b-like is made true by that a-like thing and b-like thing, and that one, and so forth' (Bennett, 2017, p. 208). The principle grounds its instances and the instances ground the principle creating an unacceptable grounding circle. CAPs can avoid Bennett's objection by grounding the linking principle in something other than its instances (although it is unclear what other facts could be called upon without the regress proliferating problematically).[33] The alternative route is to present the connecting fact as fundamental. Then, the grounding regress grounds out in the fundamental facts [a] and [[a]* → [b]*].

[33] (Wilson, 2019, pp. 8-9) defines an account of Connectivism that avoids detrimental circularity.

Most Connectivists are happy to assign fundamental status to their connecting principles and according to Schaffer, some metaphysical laws must be fundamental (Schaffer, 2017, p. 315). However, with this move they take a stance on their nomic status.[34]

In other words, just as in the case of casual explanation, non-causal grounding explanations are divisible into Humean and non-Humean camps. Making sense of this distinction when translated from the laws of nature to the laws of metaphysics comes with some difficulties. For instance, one of the chief motivations for Humeanism is Hume's Dictum (Wilson, 2010) which states, generally, that there are no necessary connections between distinct matters of fact. Translating this over to connections of ground is challenging. First, because ground connections are generally understood to hold necessarily, and second, because Humeans often take higher-level facts to reduce to lower-level facts and therefore any fact grounding in another would not qualify as distinct. There are, as with most features of ground, arguments for (Lovett, 2019) and against (Baron-Schmitt, 2021) the necessity of grounding but overall, it is unclear what a faithful Humean account amounts to.

To plainly distinguish Humean from non-Humean accounts of ground, I will set aside Hume's dictum in favour of another feature that it is often conflated with: Humean Separability. Although, as Maudlin argues, there are questions surrounding the motivation for Separability (Maudlin, 2007, pp. 61-64), it offers a clear way of distinguishing the two contrasting nomic accounts of ground. Mauldin's (Maudlin, 2007, p. 51) definition of Humean Separability is stated in terms of Supervenience but is easily translated in terms of ground:

[34] See (Wilsch, 2015) and (Giannotti, 2022) for nomic accounts of ground.

> *Humean Separability*: The complete physical state of the world is grounded by the intrinsic physical state of each spacetime point and the spatial-temporal relations between those points.

Separability is found in Lewis's doctrine of Humean Supervenience (Lewis, 1987, pp. ix-x): all that exists is an arrangement of fundamental facts and everything else supervenes on it. For grounding to follow analogously, all that can exist at the fundamental level is an arrangement of localised facts and everything else must be grounded in it.

In turn, assigning fundamental status to the non-localised Connectivist principles of the sort [[a]* → [b]*] violates Humean separability. Moreover, this is explicitly part of what motivates Bennett to reject CAP (Bennett, 2017, pp. 212-213) and pushes her towards favouring UAP: 'It is in this sense that I say that upwards anti-primitivism is Humean in a way that connectivism is not. It is Humean in that the "local matters" individually are enough to generate the connections' (Bennett, 2017, p. 213). CAP, understood as positing fundamental nomic facts, violates Humean Separability in a way that UAP does not. For CAP the grounding hierarchy 'grounds out' in two fundamental facts – the localised fact [a] and the non-localised connecting principle [[a]* → [b]*] – while for UAP the grounding hierarchy 'grounds out' solely in the localised fact [a].

These considerations are meant not presented as objections to CAP but to highlight the dialectical connection between the metaphysics of ground and the type of nomic connection. Ultimately, for those sympathetic towards a Humean approach, fundamental status cannot be granted to the linking principles and the connecting fact must be generated from localised facts. Importantly, there is a meta-philosophical trade-off going on in the background of this debate whereby one faces a dilemma in either obtaining an

inflated explanatory account resulting from positing general connecting principles or one obtains a Humean account resulting from positing only localised fundamental facts. I highlight this here as any view that does both would possess a distinct advantage and the view that I will defend in the next section possesses this advantage.

3.0 A UNIFIED APPROACH

3.1 UP-AND-DOWN ANTI-PRIMITIVISM

With the first three approaches outlined, I will now detail the view I intended to endorse. The fourth version of Anti-Primitivism is *Up-and-Down Anti-Primitivism (UDAP)*. This is the view that the ground and grounded fact together ground the grounding connection:

UDAP - [*a*], [*b*] ground [[*a*] grounds [*b*]]

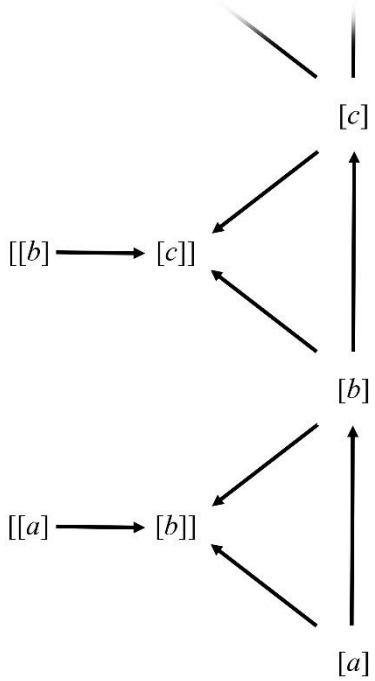

Figure 4 Up-and-Down Anti-Primitivism

In Figure 4 the connecting facts are grounded by both the ground and grounded fact: [b], [c] ground [[b] grounds [c]] and [a], [b] ground [[a] grounds [b]]. Again, [a] remains the sole ungrounded and fundamental fact. This view is briefly gestured to by Audi (Audi, 2020) and Bennett (Bennett, 2020, pp. 786-788) but both dismiss the view's plausibility for reasons I will address in the next section.[35] I to aim develop the view and demonstrate how it can address the issues discussed so far by expressing how the resulting regress is unproblematic, how an adequate explanation can be obtained, and how it is compatible with both Humean and non-Humean nomic accounts.

[35] Something like UDAP is also considered in (Raven, 2009).

3.2 Regress Revisited

To begin with, UDAP when understood as positing the ground and grounded facts to ground both the grounding connection and all subsequent regressive facts results in the following regressive grounding hierarchy:

IV. …

III. [[*a*], [*b*], ground [[*a*], [*b*], ground [[*a*] grounds [*b*]]]]

II. [[*a*], [*b*], ground [[*a*] grounds [*b*]]]

I. [[*a*] grounds [*b*]]

As the non-fundamental fact [*b*] is always reiterated in each stage of the regress it appears *prima facie* problematic. However, I argue that the ultimate ground in which all other facts 'ground out' in remains [*a*]. To see how this is achieved, note that [*b*] is grounded. The first level of the regress grounds [*b*] in [*a*]. So, any complaint that this regress leads to a chain of grounds that does not ground out in the fundamental is false. The regress continuously invokes the ungrounded fact [*a*] and the grounded fact [*b*] in every subsequent fact. Therefore, all the facts in the regress do ground out in something fundamental or are themselves fundamental. Just as UAP avoids the problems associated with the regress by positing only one ultimate ground, one ungrounded fact that grounds out all the facts in the regress, UDAP also only posits one ultimate ground. All facts in the regress ground out in [*a*], just detouring via [*b*]. Each reiteration of the regress has the same answer: [*a*], [*b*], but [*b*]'s reiteration is not problematic because [*a*] has already grounded [*b*]!

The crucial claim made here is that any regress resulting from [*a*]'s grounding [*b*] will only arise once [*b*] is grounded. Therefore, any subsequent reiteration of [*b*] will produce

no regress of facts that are not already aptly grounded out by [*a*]. This resolves the problematic aspect of the regress, but further considerations need to be addressed. Notably, the reason that both Bennett and Audi deny that [*b*] can play a role in grounding the grounding connection is that they believe the connecting fact must be grounded *before* the grounded fact. Otherwise, 'it's as though b makes it be the case that a brings it into existence, which cannot be right' (Bennett, 2017, p. 206). The crucial assumption they are opposed to is clarified by Audi when he states, 'on the assumption that a given dependence relation is always prior to its dependent relatum […] this involves a thing's being prior to itself' (Audi, 2020, p. 691). These considerations and connecting the use of temporally loaded terms in the generally non-temporary case of ground lead to further questions that will be explored in greater detail later in the paper (see §4.1).

For now, in reply, I reject Bennett and Audi's key assumption and instead assert the opposing assumption that the connecting fact must be grounded after the grounded fact: if [a] grounds [b], [b] must be grounded before any fact about [b]'s grounding is grounded. To support this approach, note that questions about what grounds a grounding connection arise only once one fact is grounded in another. If [a] has not 'yet' grounded [b], there is no meta-fact about what grounds [[a] grounds [b]] as no grounding has yet occurred. Therefore, when the UDAPs posit that [b] partly grounds [[a] grounds [b]], [b] must already be grounded by [a]. Otherwise, the fact [[a] grounds [b]] would not exist!

To clarify, there are two ways of presenting this argument. First, if explicitly concerned with the relation the claim is that the *relata* must be prior to the relation. No relation can exist if nothing is being related and therefore, the *relata* [*a*] and [*b*] must exist prior to the relating connection, [[*a*] grounds [*b*]]. The relation is between the two facts and if they do not exist then nothing can be related. Alternatively, if explicitly concerned with the

facts the claim is that the components of a fact must be prior to the fact they compose. No composite fact can exist if its components do not exist. If neither [*a*] nor [*b*] exist, then no meta-fact containing them can either.[36] I have no preference for which formulation to prefer as the consequences of both are the same: [*b*] is grounded prior to [[*a*] grounds [*b*]]. To summarise, the general argument is that the grounding connection can only exist after the facts it concerns have been grounded, otherwise no meta-fact would yet exist. If [*a*] has not yet grounded [*b*], there is no fact about [*a*]'s grounding [*b*].

However, while both ways of ordering the grounded and grounding fact can be argued for, to defend UAP, one need not inevitably accept my or Bennett's and Audi's proposed ordering. It seems in most cases of ground that both assumptions hold; that in most cases (if not all) the grounded fact should not exist without the connecting fact and concurrently, the connecting fact should not exist without the grounded fact. It is odd to have the connection between two facts before having the thing that is being connected, but likewise odd to have the connected fact before having the thing doing the connecting work. For example, consider [Socrates exists] grounds [{Socrates} exists]. It is odd to be in a situation where [Socrates exists] and [[Socrates exists] grounds [{Socrates} exists]] obtain but [{Socrates} exists] does not. Equally unusual is a situation where [Socrates exists] and [{Socrates} exists] obtain but [[Socrates exists] grounds [{Socrates} exists]] does not. The imperative result I seek here is that the priority of fundamentality (the order in which facts are grounded) in grounding relations appears underdetermined by temporally allegorical claims of the sort that [b] must exist before or after [[a] grounds [b]]. As these constraints appear internally conflicting, they should not influence which

[36] I assume this holds on any reasonable view of facts. I leave the onus on an objector to show otherwise.

way we order grounding facts; they do not bind us to a particular stance on the priority of grounding connections.

Clarifying the conflict between temporally metaphorical existence claims and fundamental priority leads towards rejecting all claims of this sort. This includes discarding both the proposal made by Bennett and Audi, that [[a] grounds [*b*]] must exist *before* [*b*], and my proposal that [[a] grounds [*b*]] to exist *after* [*b*]. This strategy might undermine my argument against Bennett and Audi's proposed ordering of grounding facts, but it also undermines their assumption which my argument was designed to oppose. As such, we ought to step back from using temporal metaphors about what facts must exist *before* or *after* others to determine which priority they warrant in a hierarchy of grounding facts. Crucially, following this, simply making an alternative assumption to Bennett and Audi about the order of priority stands perfectly permissible; I am free to claim that the facts are ordered from ground [*a*] to grounded fact [*b*] and then to the grounding connection [[*a*] grounds [*b*]].

Finally, before moving on I will briefly address some other *prima facie* concerns. First, the resulting order of fundamentality might appear unusual. This is particularly evident when exploring the immediate and mediate grounds, where the immediate ground is 'at the next lower level' and 'mediate grounds are always mediated through immediate grounds' (Fine, 2012, pp. 50-51), since [*a*] both mediately and immediately grounds [[*a*] grounds [*b*]]. Yet, so long as partial grounds are at some finite level lower this combination does not seem problematic (Fine notes this is the case of truth (Fine, 2012, p. 51)). Accordingly, I am inclined to adopt this distinction as a feature and not a bug, committing to two distinct (non-identical partial orderings of the same relata) hierarchies, one ordering between the initial facts [*a*], [*b*], [*c*], etc. and one ordering between those

facts and their grounding connections (and the regress of grounding facts). It ought not to be overly surprising that a view that is both upwardly and downwardly directed results in a hierarchy of grounding facts perpendicular to the hierarchy between those facts; a 'sideways' ordering. Ultimately, both still ground out in [a] and therefore avoid the initial regressive troubles the section was designed to appease.

Second, I am not suggesting that [b]'s grounding role is supplanted by [a], but that everything ultimately arises out of [a]. The structure determining how [[a] grounds [b]] arises from [a] requires [b] to play a role. Third, others might worry about [a] grounding [[a] grounds [b]] twice over. Initially when [a] grounds [b] directly and again, via transitivity when [b] grounds [[a] grounds [b]]. But a fact is not overdetermined by ancestors of what determines it; in the classic firing squad case, it is the multiple gunshots that overdetermine and not a single gunshot together with the order to fire.

However, another concern is raised by (Litland, 2018) who argues that involving both the ground and grounded facts together to ground the grounding connection is inconsistent with certain standard principles of the logic of ground. Briefly in response, Litland's discussion centres on the logic of ground and the proposals in this paper focus on the metaphysics of ground.[37] As I am only concerned with the metaphysics, these issues on logical ground are extraneous to my proposal. Additionally, many of the self-referential cases he discusses, when translated over to the metaphysics of ground, would raise no issues beyond those that have already been addressed in this section.

[37] A similar distinction was made earlier in this chapter for the ambiguity of Fine's endorsements of DAP. He is concerned solely with the logic of ground and that does not entail a similar commitment in the metaphysics of ground.

3.3 Partially Deflated Explanation

Moving on to explanatory issues. Dasgupta argues that UAP does not properly capture the explanatory nature of grounding by appeal to the ground alone (Dasgupta, 2014, pp. 571-572). When faced with the question of why certain activities make an event count as a conference (instead of a football match), the explanation needs to tell us why conferences are constituted by conference-like activities. There are two options available for UDAP to approach this explanatory question.

The first parallels UAP's deflationary stance and asserts that by explicitly mentioning both the ground and grounded fact an adequate explanation is provided. To address the explanatory challenge, we must answer the question of why certain activities make the event count as a conference (instead of a football match). UDAP's answer is that this event is a conference, and those events are ongoing. What more is there to say? It will always be the case that what makes certain activities constitute a certain kind of event is that the event is of a certain kind and those activities are occurring. Dasgupta acknowledges that his objection does not apply to this conception of grounding (Dasgupta, 2014, p. 572) and therefore, UDAP avoids objection based on explanatory inadequacy.

This response will not satisfy everyone since explanatory complaints against UAP are based not only on the grounded fact's absence in the explanation but also on the absence of a principle connecting the two. This issue is summed up neatly by Schaffer: 'Explanation has a tripartite structure of source, link, and result […] involving ground (the more fundamental source), principle (metaphysical principles of grounding), and grounded (the less fundamental result)' (Schaffer, 2017a, p. 3). In this sense, a *fully* explanatory account of ground includes all three parts. The approach suggested for UDAP

only meets the explanatory challenge halfway and ground's explanatory capacity is still partially deflated. This will displease those who seek a *fully* explanatory account of ground but constitutes an improvement over the adversely deflated account offered by UAP.

To obtain a linking principle the same two options are available as were for CAP: Humean or non-Humean. The non-Humean approach would entail stipulating nomic principles that play a role grounding the connecting fact alongside [*a*] and [*b*]. We would end up with a slightly different view: *Connectivist Up-and-Down Anti-Primitivism*. I see no obvious problems with such a view, but I struggle to see why those motivated as Connectivists are (i.e. more by the desire to provide an explanatorily adequate view than to adhere to Humean strictures) would migrate their allegiances over to this new view. Instead, the UDAP approach will be more appealing for those who want ground to be explanatorily adequate while also not requiring a commitment to distinctly non-Humean nomic principles. To achieve this, a novel approach to grounding is required; unsurprisingly, a Humean approach.

3.4 Humean Grounding Laws

To develop a Humean account of ground it will be helpful further to draw out the analogy with causation.[38] Humeans generally adopt a best-system approach to the laws of nature, and in particular, Lewis defines the canonical best-system account (Lewis, 1994a, pp. 478-480) whereby laws are nothing more than true generalisations of patterns in the fundamental constituents that capture the best combination of simplicity and strength in that world. The fundamental facts are a mosaic of localised points that exhibit regular

[38] Humean approaches to grounding are briefly mentioned in (Rosen, 2017, pp. 287-288), and (Sider, 2020, pp. 754-755).

patterns. Nomological parameters that aim to be both simple and informative are then generalised from these patterns; the best are the laws.

To articulate the distinction between Humean and non-Humean approaches, consider, for example, a causal regularity: [[*a*] causes [*b*]]. The non-Humean claims that this pattern is exhibited in the mosaic because of a law [if [a] then [*b*]]; the laws are fundamental parts of our world that determine what regularities the mosaic exhibits. By contrast, Humeans claim that this regularity is a basic fact, exhibited in virtue of nothing else. The laws are just helpful ways to classify these regularities in a simple and informative way for agents like us.

From here, one way to translate Humeanism about the laws of nature to laws of metaphysics is to simply extend the systems account so that it includes metaphysical laws. This proposal has been suggested by Marshall (Marshall, 2015, p. 3157) but he implies that the distinction between the two different kinds of laws is that the causal laws are contingent, while metaphysical laws are necessary. I wish to remain neutral on the issue of whether metaphysical laws are necessary (or whether causal laws are contingent) and instead continue with a Humean account based solely on Humean Separability. As such, the proposal I will outline here simply extends the systems account to include those kinds of connections between the levels of reality often ascribed to ground. The proposal is that Humean nomic principles of ground are simply ideal generalisations of regular grounding connections exhibited between localised grounding facts.

One crucial point is that Humean laws of nature require both cause and consequence. Laws are generalised from repeating observations of cause and effect. The same is true of grounding laws. The laws can only be derived from a systematisation of both ground and grounded facts together; laws are generated by generalising repeating occurrences of

[*a*] and [*b*] and [*a*] alone is not sufficient. Yet, to remain true to Humean Separability everything must arise from a localised mosaic of fundamental facts; everything must reduce to [*a*]. These might seem like incompatible theses, particularly for someone endorsing UDAP. However, as argued earlier, [*b*] can play a role in grounding [[*a*] grounds [*b*]] after [*a*] has grounded [*b*]. Similarly, once [*a*] has grounded [*b*], we can then generate metaphysical connections linking them together. The only fundamental fact, in this case, is [*a*], that first grounds [*b*] and then together these facts ground the grounding connection [[*a*] grounds [*b*]]. Subsequently, this grounding regularity is generalised into a law. This provides a translation of Humeanism from causal laws to grounding laws, where priority is first given to the mosaic, then to regularities of grounding facts, and finally to the nomic principles. Most notably, this account only works on the UDAPs model.

To summarise, Humean grounding laws are simply systematisations of grounding regularities. The whole process starts from the fundamental fact [*a*], which grounds [*b*]. Then these two facts [*a*], [*b*] together ground the connecting fact [[*a*] grounds [*b*]]. Finally, from [[*a*] grounds [*b*]] we generalise to a Humean law [if [*a*]-like things obtain then so do [*b*]-like things]. This final move is just a shorter way of stating that which is already present, and it expresses nothing over and above those facts already mentioned. The linking principle is not distinct nor fundamental, but instead simply a shorthand way of expressing those facts already grounded.

It might help to clarify the proposal by considering a simple case of metaphysical regularity; say, [*tables*] being grounded in [*atoms arranged table-wise*]. On recognising this regularity, that whenever we observe [*atoms arranged table-wise*] we also observe [*tables*], instead of listing each occurrence of atoms arranged table-wise grounding tables,

we produce the simple yet informative nomic principle: if [*atoms arranged table-wise*] obtains then so does [*tables*]. Likewise for all grounding connections. The non-localised facts become little more than mundane generalisations that nonetheless play an important role in explaining the link between the ground and grounded facts. It is worth plainly stating that the connection between the linking principle and grounding facts is not a metaphysically robust one. It plays a role in the explanation but does not ground. Likewise, Humean causal laws do not metaphysically connect local instances to particular patterns but are still used to provide an explanation of that pattern. Further discussion on this distinction is beyond the scope of this paper, but in effect, the point being made here is that the debate around the explanatory nature of Humean causal laws will translate over to Humean grounding laws for which there are viable options.[39]

Next, it is worth briefly addressing some opposition to Humean grounding. Notably, Rosen (Rosen, 2017) argues that some domains of facts appear to be grounded by general principles, notably legal facts. To support this, Rosen presents the following example. Consider legislation stipulating that the buildings on Main Street cannot be taller than six stories. Then, we can ask what makes a seven-story building on Main Street illegal, and the answer is the legislation. The law, being a general principle that governs how tall buildings on Main Street can be, both precedes any of the instances and serves a non-Humean role governing what makes a building legal. Therefore, the general principle grounds particular facts about the legality of buildings on Main Street and not the instances. If this is true, then it conflicts with the UDAP Humean approach stipulating that no grounding connections hold between general nomic principles and the facts

[39] See (Lange, 2023) for a great summary of the relevant debate, (Loewer, 2012, pp. 131-132) for the claim that the explanatory force of Humean laws of nature can be separated into metaphysical and scientific explanations, and (Hicks, 2021) for a Humean approach similar to that which I would endorse.

subsumed under it as they are just generalisations. Interestingly, Bennett concedes that 'Grounded Fact Principle-ism is true of legal facts' (Bennett, 2019, p. 515). I am not so receptive to such a concession.

The reply is simple. The law is not a general principle. It is just the ground. The confusion, at least in this case, seems to be coming from the use of 'laws' as both a term denoting general nomic principles and as facts that ground the legality of other facts. Contra to the suggestion in the previous paragraph, legislation is not a principle that governs but is instead simply a fact that makes other facts legal (or illegal). Notably, when the gavel knocks in the legislature, a new fact is fashioned that grounds facts about the legality of the buildings on Main Street. That 'law' is in the place of [*a*], and it grounds the fact that buildings are either legal or illegal [*b*]. Then, the grounding connection is simply: [[legislation stipulating that buildings taller than six stories are illegal] grounds [that fact that a seven-story building on Main Street is illegal]]. Anyone is free to call the ground a 'law', but that does not make it a nomic principle any more than being called 'Mane' Street makes the asphalt sprout hair.

This fits the UDAP Humean model, as on recognising this connection to be a regularity - that when we pass laws about the building regulations it grounds the legality of buildings - we can systematise it into a nomic principle: if [a law stating that a building on Main Street cannot be taller than six stories] obtains, then so does [the illegality of any building of seven stories or more]. First, the law makes the building illegal, and what makes this the case is both that this building is illegal and that it is a law. Then, we systematise these facts into a nomic principle: if laws obtain then so do facts about what is and is not illegal.

To conclude this discussion on Humean grounding, I believe it is worth acknowledging that accounts of Humean grounding are undeveloped in the literature at present. The version I have outlined here is open for debate, but it should stand to present one potential way of developing the position. Relatedly, those who desire an account of ground that is fully explanatory (as is the chief motivation for CAP) should be satisfied with the version of UDAP I have outlined here. A tripartite explanatory structure is attained because the grounded and ground facts are explicitly mentioned, while the linking principle arises by systematising those facts into laws.

Finally, I have already explicitly shown that UAP fails to be *fully* explanatory because of the lack of grounded fact and nomic principle. DAP also fails due to inversely lacking the ground fact and nomic principle. CAP also fails for, although containing an appeal to the nomic principle (and often ground), it lacks the grounded fact. Overall, those who want their account of meta-ground to be 'fully' explanatory - one that explicitly contains all three elements of a tripartite explanation - will find UDAP to be the only viable (Humean or non-Humean) option.

4.0 TEMPORAL COMPARISONS

4.1 THE ITERATIVE CONCEPTION OF GROUND[40]

In §3.2, temporal terminology is used to argue for ordering grounding facts in particular ways. In this section, I aim to clarify the temporal analogue with grounding and to offer a novel interpretation of Bennett's claim that the world unfolds upward (Bennett, 2017, p. 206). While I have already argued against the idea that unfolding needs to be upwardly

[40] Both Litland (Litland, 2017) and Kovacs (Kovacs, 2020) utilise 'iterative' in reference to the question of meta-ground; once a fact is grounded, we can then ask what grounds that grounding connection. To avoid confusion: the iterative conception of ground outlined in this section is entirely unrelated to their usage of the term.

directed as we can understand relata before relation and relation before relata without much issue, the act of unfolding itself remains ambiguous. So, the chief questions I aim to answer in this section are: what is it for the world to unfold? And must we accept the idea that the world does in fact unfold? To answer these questions, I will distinguish between two contrasting conceptions of ground. First, call the view that grounding facts unfold one after the other dynamic grounding:

Dynamic Grounding: the hierarchy of grounding facts unfolds iteratively.

To construct a framework for understanding how to implement a dynamic structure of grounding facts, it is worth considering a similar approach to another metaphysical notion: sets. The iterative conception of set theory is a tensed account of set theory, whereby sets are formed at different stages (Boolos, 1971). Potter defines iterative set theory as consisting of a sequence of successively generated levels (Potter, 2004). Those levels in iterative set theory can be interchanged with the levels of metaphysics with an analogous sequence being drawn iteratively from fundamental to derivative facts via grounding relations.

In this worldview, at the start, the dynamic world consists solely of fundamental facts. The zeroth stage is merely an unordered collection of ungrounded facts whose existence does not depend on any other fact. Then, at stage one, all facts that bear a grounding relation to the fundamental facts in stage zero are grounded. Next, at stage two, all facts bearing a grounding relation to the facts grounded in stage one are grounded. And so on. The facts unfold, one level after another, each level consisting of the facts that bear a grounding relation to the facts in the previous stage. The dynamic world unfolds from the most basic facts to the most derivative in an ordered fashion. This allows for a precise

interpretation of what Bennett is suggesting when remarking that [*a*] exists *before* [*b*]: simply that [*a*] is positioned at an earlier stage in the dynamic order of grounding facts than [*b*]. The more fundamental facts are prior to the derivative and the grounded arise after their grounds.

This might work for full grounding, but some nuance is required to account for partial grounding. The general idea remains the same, but if any fact is grounded in facts that do not belong to the same level, then the dynamic story of each level consisting of that grounded in the previous level does not work. More exposition is required to account for this kind of partially grounded fact. The answer I will suggest is that such a fact is grounded at the stage directly following the stage where *all* its partial grounds have been grounded. Whichever of the partial grounds is grounded last immediately precedes the level at which, together with all other partial grounds that are situated at some preceding level, ground that next fact. This allows for partial grounds to come from any level already grounded to play a role in grounding a subsequent level. For example, if [*c*] is grounded in both [*a*] and [*b*], and [*a*] belongs to the zeroth stage while [*b*] belongs to the first stage, [*c*] is grounded in the second stage because it is the stage immediately following the stage where all its partial grounds have been grounded.

There is also a metaphysical issue here concerning how seriously the dynamic unfolding of facts should be understood; whether the more fundamental facts are *literally* building the less fundamental one after another or if the unfolding is merely a *metaphorical* means of understanding the fundamentality hierarchy. Both approaches appear tenable, but the iterative conception of set is often accompanied by a caution against taking the iterative metaphor too seriously. Such warnings are present throughout (Studd, 2013), and Linnebo

remarks that a 'similar attitude lives on today in Boolos (1989)'s admonition against taking the iterative conception's talk of "set formation" too literally' (Linnebo, 2013, pp. 207, fn. 5). Perhaps these concerns translate over to ground, but I see nothing overtly objectionable about adopting the alternative stance. There exists the conceptual space for both a stronger view, whereby the more fundamental facts literally build the less fundamental facts and they do not exist until the iterative unfolding has reached their level. There also seems space for a weaker view, whereby we just take the unfolding as a symbolic means of ordering the grounding hierarchy. I will remain neutral on this distinction and instead focus on the unfolding aspect, without committing to a position on whether it is literal. Importantly, while the weaker dynamic view might only be dynamic in name, it appears different from the entirely non-dynamic approach whereby there is no unfolding, and the facts instead appear 'all-at-once':

Static Grounding: all facts in the grounding hierarchy arise at the same 'time'.

On the static view, all the facts in a grounding hierarchy are not in any sense before or after any other. Some are more fundamental or less fundamental than others, but this bears no translation to temporal analogues. All facts appear together. The world is neither 'unfolded' nor 'folded' and at no point does it make sense to compare the more fundamental fact [*a*] arising *before* the less fundamental fact [*b*].

Those who hold robust views on the atemporal nature of grounding will likely be drawn to static grounding and there is nothing obviously incorrect about it. I suspect opinions on dynamic/static grounding will raise no crises for those with particular opinions on meta-ground. For instance, Wilson affirms that he cannot see anything preventing the grounded facts from appearing all at once out of the ungrounded facts (Wilson, 2019, p.

9). The intention in this section is not to argue for or against dynamic or static grounding, but to instead outline the two views and their differences. Some may have compelling intuitions for one approach or the other. I do not as they both appear to offer tenable variations of ground.

4.2 Compatible Dynamic-Static Meta-Ground

With exposition on the notion of 'unfolding' in place, I will now turn to interpreting the directions in which that unfolding can occur. Bennett has strong intuitions supporting the idea that the grounded fact cannot precede the grounding connection (as discussed in §3.2), but there is just as much reason to think the opposite. One alternative way of interpreting these conflicting arguments is as endorsements for Staticism about grounding. If the priority of fundamentality in grounding relations is underdetermined by temporal claims such as '[*b*] must exist before [[a] grounds [*b*]]' (or vice versa), then why bother with using the temporal notion at all? So, the following is only for those who feel an independent pull towards dynamic grounding, and with the intention of clarifying the dynamic ordering I proposed for UDAP.

To begin, consider one interpretation of the UAP view whereby the ground builds the connecting fact and grounded fact at the same 'time'. As depicted in Figure 5:

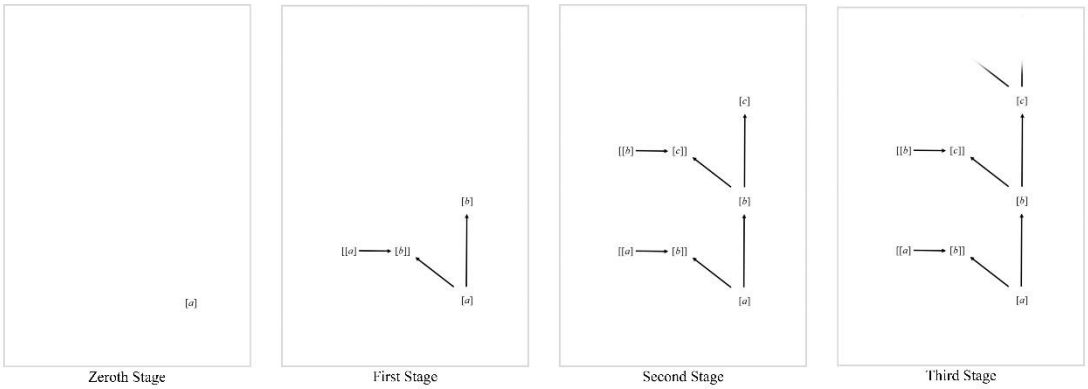

Figure 5 Dynamic UAP

Figure 5 represents the priority of fundamentality (the order in which the facts unfold). Starting with the most basic fact [*a*], then both the grounding connection [[*a*] grounds [*b*]] and the grounded fact [*b*] arise at the next stage. The unfolding then continues from [*b*] to the grounding connection [[*b*] grounds [*c*]] and the grounded fact [*c*] at the next stage. This is the ordering Bennett seems to have in mind, but it might not be correct. One could put the grounding connection in the stage before the grounded fact and not alter the view's 'upward' intentions.

I view the upward ordering as just one of many plausible options. The ordering I suggest for UDAP involves an additional caveat moving upwards and then 'sideways'. As depicted in Figure 6:

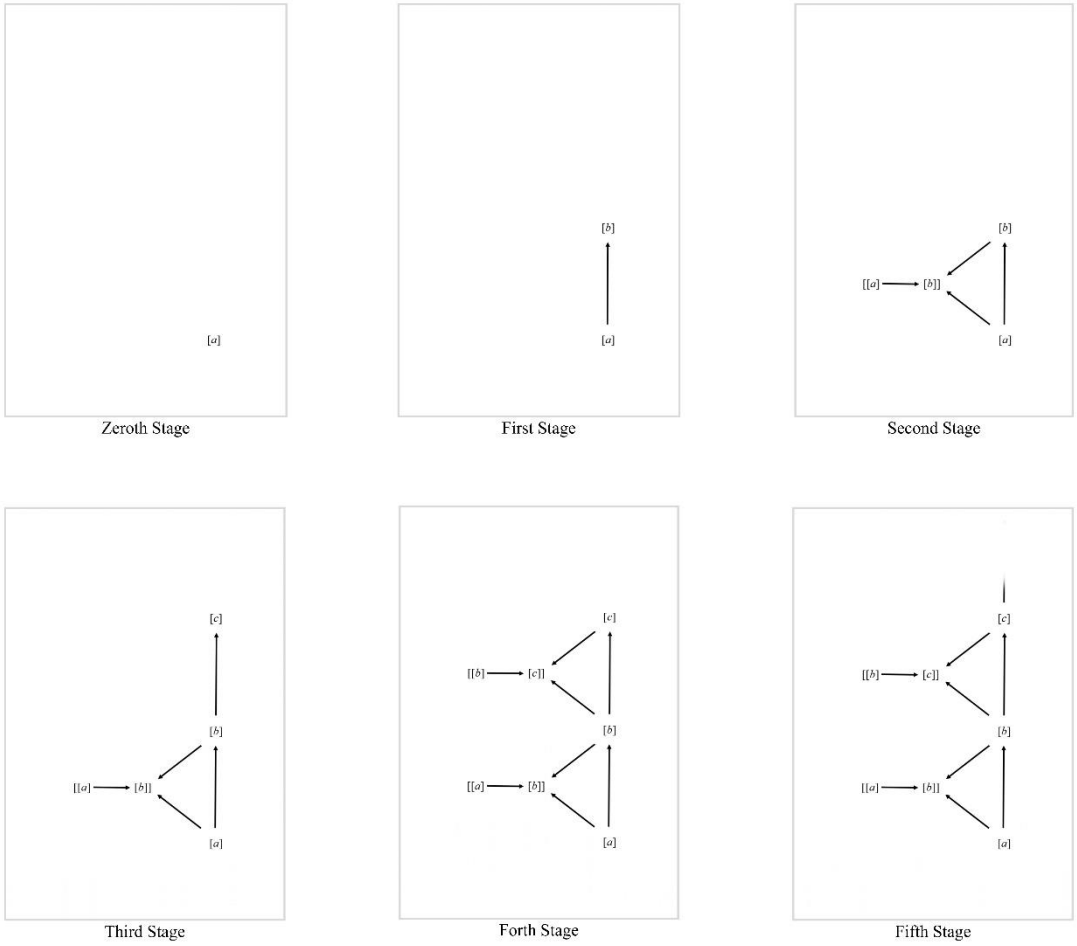

Figure 6 Dynamic UDAP

Figure 6 represents a distinct direction of fundamentality. Starting with the most basic fact [*a*], then with the grounded fact [*b*] arising at the next stage. This is followed by the grounding connection [[*a*] grounds [*b*]]. The unfolding then continues from [*b*] to [*c*] and afterwards to the grounding connection [[*b*] grounds [*c*]]. The facts remain the same, and so too do the grounding relations exhibited between the facts. All that changes is the order in which the grounding hierarchy unfolds. Overall, this section is not meant to include an argument for either UAP's or UDAP's dynamic ordering but is intended to provide clarification on the differences. Either route for ordering the dynamic unfolding, or even

options entirely distinct from these two, appear tenable. The aim here is merely to clarify the different ways one might choose to order the hierarchy.

4.3 Temporal Position of Nomic Principles

In the same passage where Bennett remarks on her temporally loaded account of ground, she also defines an argument against CAP. This argument relies on a more traditional understanding of temporally loaded notions of 'before' and 'after'. To be clear, in this section, one will need to set aside the previous discussion of temporal ground and instead accept these terms to merely refer to something occurring at an earlier or later time.

To begin, Bennett's argument is directed against the nomic principles posited by CAP (Bennett, 2017, pp. 208-211). She argues that they must be understood as generalisations and that generalisations must be temporally posterior to their instances as they can only be established once there are at least some instances to generalise from. In other words, the facts [a], [b] must exist before the fact that [[a]-like things ground [b]-like things]. To support this claim, Bennett asks us to consider some grounding facts that do not exist at the present but will at some future point: say the first instance of [p] grounding [q] occurs at the future time t. Before t it cannot be that there are [p]-like things, nor [q]-like things (assuming that the nomic principles can only be established via generalisations of the facts they connect). Therefore, in the first instance of [p] and its grounding [q] at time t, the grounding of [[p] grounds [q]] occurs before the existence of [p]-like things and [q]-like things. The instances of grounding precede nomic principles; as the facts are grounded before establishing any principles about those facts' grounding, nomic principles cannot play any role in the grounding.

So long as we accept nomic principles to be nothing more than generalisations of their instances the argument is persuasive. Someone with a suitable non-Humean approach to nomic principles would be best placed to reject the claim that the nomic principles are mere generalisations. If the nomic principles of ground are truly analogous to the nomic principles of causation, then they would be nothing like generalisations. Non-Humeans posit metaphysical primitives (be they primitive laws, essences, or dispositions) that exist before any of the instances they govern because that is what explains why the world exhibits the patterns it does. Analogously for grounding, if the nomic principles are metaphysical primitives that govern the grounding hierarchy, they too must exist before the facts they ground. All in all, Bennett's argument relies on adopting a weaker understanding of nomic principles and an advocate of CAP can avoid the force of her objection by strengthening their nomic principles; by making an explicit commitment to non-Humeanism.

Inversely, understanding nomic principles as mere generalisations is more like the Humean approach – one that I have outlined in §3.4. In this case, the argument runs into the inverse problem. Someone with a suitable Humean approach to nomic principles would not posit a robust connection between the generalisation and the instances. So, while Bennett is correct to assert that the generalisations must be temporally posterior to the instances, this is nothing a Humean would disagree with. Nor does it pose any issues for the Humean account of grounding outlined for UDAP that I have proposed.

As such, I believe the best way of interpreting Bennett's argument is as intending to force the issue on the status of nomic principles. If one wants them to play a role in grounding the grounding connection, they must be understood not as a generalisation but in a non-

Humean metaphysically robust sense. Perhaps advocates of CAP are fine with making this kind of commitment, but advocates of UDAP need not make this choice. Both a Humean and non-Humean approach is compatible with UDAP. The nomic principle can ground the grounding connection alongside the ground and grounded fact or it can serve a weaker – non-grounding – explanatory role. Overall, the only viable version of CAP appears to be non-Humean while this restriction does not apply to UDAP.

5.0 CONCLUSION

5.1 COMPARING ACCOUNTS OF META-GROUND

To begin drawing this paper to a close, I will now provide a comprehensive overview of the various versions of meta-ground and their standings with respect to the arguments presented in this paper. The intent here is to clearly show what view is compatible with what and ultimately that UDAP represents the most harmonious approach. The views explicitly discussed already include DAP, UAP, CAP and UDAP. There are also variations on these – each may have a Humean and non-Humean variation or a dynamic and static variation, etc. But many of these have been shown to be untenable positions – I have argued that the only viable Humean view is UDAP. Overall, to help spell out the views and their various compatibility, I will first define the aspects with respect to which accounts of grounds are being compared.

First, each account of meta-ground is assessed relative to whether the regress is problematic. UAP had no problems there, and I have argued that UDAP (and consequently DAP) does not lead to a problematic regress. CAP might face issues here, but it depends on the status of their nomic principles and whether they are grounded in their instances. So, there are only some possible issues for CAP on this front. Second,

the explanatory capacity of each view is assessed. Broadly, there are three different levels of explanation the accounts of meta-ground can obtain: A *full explanation* view consists of all three elements in the tripartite structure, a *partial explanation* view consists of two out of three elements, and lastly a *deflated explanation* view consists of only one out of three elements. DAP and UAP both obtain deflated accounts, CAP obtains a partial account, and UDAP a full account. Third, compatibility with Humean or non-Humean nomic principles is assessed. I define a distinctly Humean approach in terms of Humean Separability that only UDAP is compatible with. Additionally, because of Bennett's temporal argument in §4.3, CAP appears to only be compatible with non-Humean nomic principles while UDAP is compatible with either. Neither UAP nor DAP posit nomic principles (perhaps this makes these views more Humean – as Bennett holds – but they still lack any fleshed-out nomic account). The last argument concerns the distinction between dynamic and static grounding. While some views might be incompatible with ways of ordering the dynamic unfolding, when permitted a free choice of how the grounding hierarchy unfolds, all views appear capable of defining a dynamic and static variation respectively. As such, I will set this consideration aside.

Drawing all this together, the following comparisons can be made:

	Unproblematic Regress	*Deflated Explanation*	*Partial Explanation*	*Full Explanation*	*Non-Humean Nomic Principles*	*Humean Nomic Principles*
DAP	✓	✓	X	X	X	X
UAP	✓	✓	X	X	X	X
CAP	✓	✓	✓	X	✓	X
UDAP	✓	✓	✓	✓	✓	✓

Figure 7 Comparison between theories of Meta-ground

While this kind of rudimentary comparison lacks nuance and making comparisons between metaphysical theories is not as simple as merely listing arguments, Figure 7 depicts UDAP as avoiding some of the main problems facing competing accounts.

5.2 The Unifying Account of Meta-Ground

To conclude, it is worth drawing out another one of UDAP's chief virtues; it is unifying. UDAP incorporates positive aspects from all three other anti-primitivist positions by being concurrently upwardly directed, downwardly directed, and Connectivist. Further, the motivation underscoring each of these views is also incorporated into UDAP. Those who feel an intuitive pull towards DAP ought to be satisfied that the grounded fact can play a role in grounding the grounding connection. Those who are motivated to avoid a problematic regress ought to be satisfied that UDAP's regress grounds out in the fundamental. Those who are motivated to provide an explanatory account of ground

ought to be satisfied with UDAP's (either Humean or non-Humean) tripartite explanation. Overall, all of the desirable features are retained while the problematic aspects are avoided.

PAPER IV: ULTIMATE-HUMEANISM

ABSTRACT

Humeans (Esfeld & Deckert, 2017) argue that the most parsimonious ontology of the natural world compatible with our best physical theories consists exclusively of particles and the distance relations between them. Prominent criticisms of Super-Humeanism question the view's coherency (Matarese, 2020), explanatory capabilities (Wilson, 2018b), empirical adequacy (Lazarovici, 2018), and plausible temporal persistence (Simpson, 2021). These objections all stem from the common thought that Super-Humean reduction is excessive. This paper argues by contrast that Super-Humean reduction goes insufficiently far, by showing there to be a more parsimonious ontology compatible with physics: Ultimate-Humeanism. This novel view posits an ontology consisting solely of the particles and distance relations required for the existence of a single brain. Super-Humeans impose conditions on what counts as an ontology of the natural world to avoid their view slipping into this kind of solipsism, but these conditions are arbitrarily imposed and once this is exposed, Super-Humeans face a dilemma. Either they can embrace Ultimate-Humeanism as the minimal ontology of the natural world compatible with physics, or (more likely) they can rethink the methodology that got them there. Overall, this paper argues that Super-Humeanism currently lacks principled motivation, outlines a framework for naturalistic ontological reductions, and exposes the consequences of unchecked adherence to a simplicity-driven methodology. More generally, this paper explores the process of deriving a minimal, physically adequate, ontology of the natural world and highlights the consequences of pursuing excessive ontological reductions.

OVERVIEW

First, in §1 the methodological background employed by Super-Humeans is outlined and the sole guiding principle of minimal sufficiency is introduced. Next, in §2 ontological reduction is applied to non-Humeanism and Humeanism, resulting in Super-Humeanism. Then, in §3 Ultimate-Humeanism is described and argued for as the minimal ontology of the natural world that meets the standards of sufficiency outlined in the Super-Humean methodology. Ultimate-Humeanism is also defended against two potential objections. First, on grounds that fails to remain a realist project that squares with the definition of the natural world set out by Super-Humeans and second, that it relies upon a radical reinterpretation of physics. Finally, in §4 the conclusion is drawn out: either Ultimate-Humeanism is embraced as the new minimal ontology of the natural world or the methodological principle that implies it is rejected. In both cases, Super-Humeanism currently lacks principled motivation either because the methodology of prizing minimality above all else is abandoned, or because it is not the minimal ontology.

1.0 METHODOLOGICAL BACKGROUND

1.1 THE CANBERRA PLAN FOR PHYSICS

Inspired by the Canberra plan outlined by Lewis (Lewis, 1986) and Frank Jackson (Jackson, 1994), Super-Humeans conceive of metaphysics as an ontological enterprise where once the fundamental parts of the ontology are explicated, all other parts of the ontology will be accounted for in terms of the fundamental. In this sense, the metaphysical goal Super-Humeans aim for is a classification of the fundamental entities in terms of which the rest of reality can be understood. To achieve this aim, a central distinction is established between two classes of fundamental and non-fundamental entities. For the former, Super-Humeans interpret physics as providing an account of reality's most basic constituents. For the latter, as many aspects of reality do not feature in the fundamental story, they are either eliminated from the ontological list or located elsewhere in it.

Location is the task of accounting for the implicit existence of the non-fundamental in terms of the explicit fundamental ontology. The Super-Humean methodological approach to location diverges somewhat from the traditional understanding of metaphysics as ontology. Notably, Michael Esfeld (Esfeld, 2020a, p. 126) conceives of location in terms of identity where the metaphysical task is to show how some features of reality are identical to others. Esfeld also makes the stronger claim that 'taking recent developments in metaphysics into account that have moved away from the notion of supervenience, "location" is in the following conceived in terms of identity' (Esfeld, 2020, p. 126). Whether location as identity is widespread in contemporary metaphysics does not matter for this paper; what is important is that Super-Humeans conceive of location in this way.

After establishing an identity relation between two distinct features, the ontology need only explicitly include one of them to implicitly include the other. Consider for example

chairs and *atoms arranged chair-wise*. If both exist, then these are either identical or not identical. If chairs are identical with atoms arranged chair-wise, then if chairs exist, the existence of atoms arranged chair-wise is nothing new and the ontological list does not expand. Inversely, if chairs are not identical with atoms arranged chair-wise, then if chairs exist, the existence of atoms arranged chair-wise is something new and will expand the ontology. This is how Esfeld conceives of location, but such adaptation does not depart significantly from the original proposal of Lewis and Jackson as similar ontological reductions can be achieved without positing identity as the locating relation.[41]

The Super-Humean methodology also pays closer attention to physics, and prizes minimality to a greater degree than Lewis and Jackson proposed. Physics is not merely taken into consideration, but the ontology ought to be 'in the position to tell a story how the established physical theories can be construed on the basis of such a minimalist ontology' (Esfeld, 2020b, p. 1891). For minimality, they build on Jackson's suggestion of not, 'letting a thousand flowers bloom but rather that of making do with as meagre a diet as possible' (Jackson, 1994, p. 25). These two features outline the sole principle guiding the Super-Humean project: minimal sufficiency. Where minimality is achieved by positing an ontology consisting of the fewest fundamental elements while employing location (or elimination), as described in the previous paragraph, to account for the non-fundamental; the minimal ontology posits the fewest fundamental entities and then locates the rest of reality's features in the fundamental. Sufficiency is achieved by the ontology being compatible with our most prominent physical theories. This includes a

[41] Other relations could be used in place, e.g., Grounding or Supervenience. Esfeld ascribed Supervenience as Lewis's and Jackson's locating relation. Lewis believes there is a role for supervenience in certain domains (e.g. reduction of mind (Lewis, 1994b)) but in general might be better understood as an identity theorist for his stance on mereology (Lewis, 1991).

connection between our empirical observations and ontology that allows for our physical evidence to be understood in terms of the ontology, but further discussion of this qualification comes later in the paper. Overall, Super-Humeans interpret naturalistic metaphysics as targeting the most parsimonious physically adequate ontology. It is, for them, an enterprise designed to answer the following question: 'What is a minimal set of entities that form an ontology of the natural world, given what science tells us about the natural world?' (Esfeld, 2020a, p. 126).

1.2 Whose Minimum?

As Super-Humeans place high importance on minimality, it is imperative to clarify what counts as the minimal ontology. Most notably, reduction is not a process of building an ontology with fewer elements but instead one of taking a given ontology and eliminating or locating elements to increase simplicity. Super-Humeans are not concerned with building the simplest ontology of the natural world but are instead concerned with deconstructing parts of a natural ontology into its simplest, but still sufficient, form. As a result of the deconstructive nature of reduction, contrasting starting ontologies lead to contrasting minimums.

This distinction, between local and global minimums, is drawn by Alastair Wilson (Wilson, 2018c). The former is a preference for simpler theories within the bounds of any given ontology while the latter makes the stronger claim that no other ontological system is objectively simpler than it is: 'that no system entails the existence of objectively less stuff than it does' (Wilson, 2018c, p. 6). Consider, for example, a hiker trying to find their way down from a mountain range where they can only decide the direction to travel by assessing the gradient in the local vicinity. Where the range consists of many peaks and troughs, the hiker will most likely find themselves in a mountain lake or nearby valley,

which is the lowest *nearby* point, but not the *absolute* lowest point. Depending on where one starts, the decent to the local minimum will lead to contrasting places at contrasting heights. The same goes for ontological reductions; the local minimum is relative to the bounds of the starting ontology and will lead to the nearest mountain lake while the global minimum is relative to all ontologies and will lead to the Mariana Trench.[42]

It is difficult to establish a clear criterion for parsimonious evaluation between contrasting ontological systems; to say which local minimum is truly the most parsimonious. So, although it might seem that Super-Humeans seek a minimum in comparison to any and all given ontologies of the natural world, they are at most targeting a local minimum.[43] It remains an open question whether other systems entail the existence of objectively less stuff. This kind of comparison between global minimums would consider different fundamental elements, say for example, between Super-Humeanism's particles and Wallace and Timpson's Spacetime States (Wallace & Timpson, 2010). And when comparing the two ontologies there is no common standard to evaluate which is simpler. Ambiguities can also arise when evaluating parsimony between minimums within the same ontological system as different elements in the ontology can be reduced. This results in ontologies consisting of contrastive entities, both of which are simpler than the one they are a reduction of, but it being unclear (in comparison between them) which is simplest. In general, specifying a straightforward means of parsimonious evaluation between contrasting entities is prohibitively challenging.[44]

[42] This is known as 'Gradient Decent' in mathematics.

[43] Esfeld and Deckert have confirmed that are only concerned with a local minimum (Wilson, 2018c, p. 7)

[44] Sober (Sober, 1988, pp. 37-70) objects to the plausibility of evaluating the simplicity of theories independent of a given domain. See also (Woodward, 2003, pp. 288-295) for similar criticism of Lewis's vague use of simplicity.

As such, one uncontentious way to establish that an ontology is simpler is if the reduced ontology is a strict subset of the original ontology. This is to say if at least one entity in the original ontology does not exist in the reductive ontology, then it is uncontroversially simpler. This method allows for straightforward evaluations of simplicity between ontologies. For example, if ontology *a* contains entities φ, ψ, and, ω: $a = \{\varphi, \psi, \omega\}$, and reductionists locate the element φ in ψ and ω to form the strict subset of *a*: $b = \{\psi, \omega\}$, we can state that *b* is more parsimonious than *a* as there is at least one element of *b* that is not in *a* while nothing additional has been added: $b \subsetneq a$. Likewise, if the reductionist goes a step further and locates ψ in ω to form the strict subset of *b*: $c = \{\omega\}$, we can state that *c* is more parsimonious than *b*: $c \subsetneq b$. As *c* is also a strict subset of *a*, it too is more parsimonious than *a*: $c \subsetneq a$. Here a hierarchy of simplicity emerges from the most proliferated, *a*, to the simpler *b*, to the simplest *c*, where parsimonious evaluations of the sort, *c* is simpler than *a*, hold without quarrel. This understanding of minimality will be utilised in this paper when assessing the ontological standards of minimality attained by Super-Humeanism. This application of the principle raises no problems for Super-Humeanism and, although limiting the extent to which their ontology can claim to be the universal ontological minimum, it provides a robust and practical basis to carry out evaluations of simplicity.

1.3 End of the Line

The final part of the methodological background required is to establish a metametaphysical principle on the limitations of ontological reduction. For any reduction of a given ontology, there will be certain parts of the ontology that are denoted as indispensable; held fixed as bedrock commitments vital to the ontology. Yet, the indispensable parts ought to have a good reason for why they need to remain. Otherwise,

when one embraces the search for minimality there is little reason to stop the reduction anywhere but at the absolute minimal reduction. This is of primary concern for individuals who aim to reduce an ontology that already possesses the virtue of being simpler than others; those that push the boundary for minimality further. In these cases, there is disagreement about where the reduction ought to stop and what parts of the ontology are indispensable.

A wide range of considerations might be drawn upon to enforce a particular stopping point, but the stopping point should not be arbitrary. When searching for minimality there ought to be a well-motivated end to the reduction. Otherwise, *ceteris paribus*, if it is permissible to assign ad-hoc bedrock status to any part of the ontology, all reductions could be avoided by affirming any elements up for location as bedrock commitments. When there are competing reductive ontologies with different stopping points, if the less parsimonious one does not have a well-motivated stopping point, then there is nothing to prevent its collapse into the more parsimonious one. Conversely, if the more parsimonious one does not have a well-motivated stopping point, then there is nothing preventing similarly arbitrary conditions from being imposed on the entities it aspires to reduce.

2.0 THE SEARCH FOR MINIMALITY

2.1 NON-HUMEANISM

To begin exploring the Super-Humean ontological reduction we must first identify the ontology for which they aspire to derive a minimally sufficient version. That ontology is materialistic and based on fundamental physics and can be expressed as consisting of three distinct elements. First, there are localised particles that instantiate properties. It is

worth noting right away that competing physical ontologies might posit fields (in addition to or in place of particles) as part of the fundamental ontology but Super-Humeans believe these elements reduce to particles (e.g., they endorse a particle interpretation of quantum field theory (Deckert, et al., 2019)). As this debate can be side-stepped here, for clarity I will assume Super-Humeans are successful in this reduction and that particles are all that is required.[45]

Second, the particles instantiate properties. Some of these properties are intrinsic to the particles and physics provides some prominent candidates for these e.g., mass, charge, etc. Other properties are external and denote spatial-temporal relations between particles that arrange them on a space-time manifold. The resulting configuration of particles exhibits certain patterns of behaviour, or regularities e.g., charged particles repelling, particles with mass attracting, etc. Physicists then use these regularities as the basis for defining the laws of nature. The laws are the third and final part of the starting ontology that contains the following elements:

non-Humeanism (NH) = {Space-time Particle Manifold, Properties, Laws}

The significant aspect of this view is that the laws of nature are fundamental.[46] Laws are given such status by non-Humeans for various reasons, but they all share the common theme affirming that the modal facts arise before the regularities; that the laws govern the

[45] Conversely, (Caulton, 2018) argues that Super-Humeans are unsuccessful in this regard. (Bhogal & Perry, 2017) provide an alternative Humean method of dealing with the non-localised posits of quantum mechanics while also discussing the Super-Humean approach (and raising some problems).

[46] For dispositionalists (e.g. (Bird, 2007)), who are anti-Humean but take the laws to stem from the dispositional essence of properties, the story here would be slightly different. The elements that are reduced would be the dispositional essences instead of laws. This makes little difference to the dialectic here and a Humean reduction is still applicable. For the sake of clarity, I will set the disposition view aside.

regularities and explain why the particle configuration exhibits such patterns.[47] non-Humeans believe that the laws of nature are bedrock commitments indispensable to the ontology but regardless, their ontology merely serves here as the starting point for Super-Humean reductionism.

2.2 HUMEANISM

Before Super-Humeanism lies regular Humeanism. Humeans deny that laws of nature are part of the fundamental ontology. To achieve this reduction a systems approach to laws is adopted: notably, the best systems account endorsed by Lewis (Lewis, 1994a, pp. 478-480) whereby laws are located in true generalisations of regularities that capture the best combination of simplicity and strength in that world. Laws, for Humeans, are nothing more than nomological parameters employed by us to systematise the regularities. They play no role in governing patterns in the property instantiating particle manifold and arise after, not prior to the regularities.

Importantly, Humeanism achieves greater simplicity by locating one element of non-Humeanism, resulting in an ontology consisting of the following elements:

Humeanism (H) = {Space-time Particle Manifold, Properties}

As $H \subsetneq NH$, together with the definition of minimality outlined previously in this paper affirming that strict subsets of ontologies are unquestionably simpler, we can conclude that Humeanism is more parsimonious than non-Humeanism. Then, recalling the guiding maxim of minimal sufficiency, if it is possible to construct our best physical theories from this ontology, it ought to be preferred over non-Humeanism. As will be shown in the next

[47] Some views invoke god, others structure imposing primitives, and others dispositional essences; see (Hildebrand, 2020) for a summary of these views.

section, our best physical theories can be constructed from an ontology that is a strict subset of Humeanism and therefore, it has all the ontological resources (and more) available to meet the standards of sufficiency.

2.3 Super-Humeanism

Next, Super-Humeans (Esfeld & Deckert, 2017) go a step further than Humeans and locate some properties. Specifically, they locate the intrinsic properties into the property of standing in an external relation to all other particles. The intrinsic properties are understood as nothing more than arrangements of particles where the external relation of being in a particular place relative to all other particles is identical to the particle's instantiating a property. For example, where a Humean asserts that the property of mass is instantiated by a particle, a Super-Humean asserts that mass is identical with a particular configuration of particles. As a result, their ontology only consists of particles and the external relations between them. To express the gain in ontological simplicity, it is first worth individuating the plurality of properties in the Humean ontology:

$$H = \{\text{Space-time Particle Manifold, Property}_1, \text{Property}_2, \ldots, \text{Property}_n\}$$

Then, the Super-Humean locates all other properties into one property of being in an external relation to all other particles:

$$\textit{Flat-Humeanism (FH)} = \{\text{Space-time Particle Manifold, External Relations}\}$$

Yet, I have named this view Flat-Humeanism because this is not exactly the Super-Humean position. Their reduction goes even further, applying not only to some properties but also to some particles in the space-time particle manifold.

While there are some Humean approaches according to which laws are underdetermined before the end of time, leaving the future open (Beebee & Mele, 2002), Humeans

generally posit a four-dimensional configuration of particles that consists of many three-dimensional layers.[48] This commitment is clear to see in Lewis's work as he asserts that it is the totality of manifest matters of particular fact that form the basis that generalised laws are based on. The mosaic contains all past, present, and future space-time points:

> 'the neo-Humean thesis that every contingent truth about a world—law, dependency hypothesis, or what you will—holds somehow in virtue of that world's total history of manifest matters of particular fact. Same history, same everything.' (Lewis, 1981, p. 20)

The laws that result are those that provide the best systematisation of this totality and capture the simplest and most informative representation of the regularities that can be found throughout all times. In contrast, Super-Humeans only posit a single spatial configuration of particles that changes. Or in other words, Humeanism is Eternalist and Super-Humeanism is Presentist. Briefly, Eternalism is the view that everything exists regardless of its temporal location; entities in the past, present, and future all exist simpliciter. Presentism is commonly formulated as the view that only present entities exist; there are no non-present basic constituents.

While the Flat view appears tenable, the Super-Humean doctrine is formulated solely with spatial relations and only one spatial configuration of particles instead of spatiotemporal relations and many temporally distinct particle configurations (Esfeld, 2020b, p. 1894). Super-Humeanism is expressed in this manner because they claim it is ontologically more parsimonious; each time slice contains more particles and therefore is an addition to the ontology. As they write, 'Presentism, thus conceived, is the most simple and

[48] Non-Eternalist views lead to questions about how the laws can supervene on non-existent basic constituents (Hüttemann, 2014).

parsimonious ontology, since only one configuration exists' (Esfeld & Deckert, 2017, pp. 151-152). As Simpson notes, presentism appears to be Esfeld's view: 'This suggests super-Humeans should adopt something like the A-theory of time, in which there is only one configuration of particles that exists at the present moment. (This appears to be Esfeld's position.)' (Simpson, 2021, p. 903). Similar remarks are made by Super-Humeans in other papers (Esfeld, 2020b, p. 1895).

To once more precisify the gain in simplicity, it is worth individuating the space-time particle manifold into individual particle configurations indexed to different times:

> *FH* = {$\text{Configuration}_{t1}$, $\text{Configuration}_{t2}$, …, $\text{Configuration}_{tn}$, External (Spatial-Temporal) relations}

Then, if $\text{Configuration}_{t1}$is taken to be the present particle configuration, the proper Super-Humean position can be expressed as follows:

> *Super-Humeanism (SH)* = {$\text{Configuration}_{t1}$, External (Spatial-Distance) Relations}

To specify exactly as Super-Humeans do, their view consists of only two axioms:

1. There are distance relations that individuate simple objects, namely matter points.
2. The matter points are permanent, with the distances between them changing. (Esfeld & Deckert, 2017, p. 4)

The first axiom specifies that the ontology contains particles arranged into a spatial configuration and the second, specifies they are enduring entities whose sole feature is the distance they stand to all other particles. I see no serious issues with the first axiom;

with adopting a particle ontology and differentiating them solely by distance relations. In this regard, I will raise no issue with whether Super-Humeans are successful with their reduction and accept that particles and distance relations are identical with intrinsic properties e.g., that they can give an account of mass in terms of a particular pattern in the changing particle configuration.

However, the second axiom complicates matters of minimality as change was not included in the Humean ontology and therefore, the ontology of Super-Humeanism is not a strict subset of the Humean ontology. Change has been introduced as a distinct element in Super-Humeanism that was not an element in Humeanism (or any other previous ontology). Perhaps there is ground for objection on the basis that there is no clear way of determining whether one changing spatial configuration is more parsimonious than many static ones. They are different types of things, and no clear methodology is given for how to evaluate these discrepancies. There are also differences with the reduction of past and future spatial configurations as they are not located in the present configuration but are instead eliminated as redundant ontological commitments, superfluous to meeting the criteria of sufficiency. I will not engage with these worries here.[49] Instead, I will discount primitive change as an element in the ontological set, focusing on the gain in simplicity achieved because there are fewer configurations, while also highlighting that Super-Humean reductive methodology is not restricted to location but also includes the elimination of everything beyond that required for sufficiency.

As an additional aside here, it is worth plainly articulating the various options in conceptual space when aspiring to reduce an Eternalist Humean ontology:

[49] Marmodoro (Marmodoro, 2018) engages worries like these, particularly the notion of change.

Humeanism = { $\text{Configuration}_{t1}$, $\text{Configuration}_{t2}$, …, $\text{Configuration}_{tn}$,
Property_1, Property_2, …, Property_n}

The first option I have already outlined is to remove the properties but keep all of the spatial configurations. Effectively, by only reducing the intrinsic properties but retaining an ontology that is temporally extended, we end up with a 'flat' ontology:

Flat-Humeanism = { $\text{Configuration}_{t1}$, $\text{Configuration}_{t2}$, …, $\text{Configuration}_{tn}$,
External (Spatial-Temporal) relations}

Another option is to do the opposite and remove the non-present particle configurations but keep the intrinsic properties. In this case, we end up with a temporally 'thin' ontology:

Thin-Humeanism = { Configuration_1,
Property_1, Property_2, …, Property_n}

Lastly, there is the option to do both; to remove both non-present spatial configurations and intrinsic properties. That is the Super-Humean view:

Super-Humeanism = {Configuration_1,
External (Spatial-Distance) Relations}

While I see each of these contrasting ontologies as worthy of examination, I am not advocating for any of them. The aim of this aside, in depicting these various alternatives, is to help clarify the ontology proposed by Super-Humeans and articulate precisely how their reductionism is 'super'.

That said, in the relevant comparison between Super-Humeanism and Humeanism, parsimony has increased in two respects: fewer properties and fewer particles. As an ontology is simpler when it reduces or eliminates elements from another ontology while adding nothing new, and Super-Humeanism reduces intrinsic properties and eliminates

non-present particles, it is simpler. In other words, the Super-Humean ontology is a strict subset of the Humean ontology: $SH \subsetneq H$. Following the methodology, as Super-Humeanism achieves greater parsimony, if it is possible to construct our best physical theories from it, it ought to be preferred.[50]

Esfeld and Deckert proceed with a rigorous examination of prominent contemporary theories of fundamental physics and demonstrate that they can be accounted for in terms of their particle-only ontology. For example, they endorse a conception of quantum field theory compatible with permanent matter points and relational space (Esfeld & Deckert, 2017, pp. 99-130) and similar investigations are conducted for quantum mechanics and relativistic physics (Esfeld & Deckert, 2017, pp. 59-98 & 131-166). Super-Humeans demonstrate how our best physical theories are constructed from their ontology, building our physical theories from a presently changing particle configuration. Chiefly, Super-Humeans are capable of deriving (in a metaphysical sense) a complete account of our best contemporary physical theories from the complete Super-Humean particle mosaic. To explore how they achieve this it is first worth detailing the considerations in the background of their epistemology and their precise formulation of sufficiency.

2.4 Sufficiency

The Super-Humean methodology affirms that an ontology is sufficient if we can construct our best physical theories from it, but the specific conditions for sufficiency are more nuanced. Notably, a prevalent assumption that underscores the Super-Humean particle-

[50] Avoiding the 'baroque' metaphysics in Lewis's adoption of quiddities is another advantage Esfeld and Deckert claim over Humeanism (Esfeld & Deckert, 2017, p. 46). Yet, this is more an expression of disapproval than a critique. Putnam criticises Lewis's realist metaphysics for sounding 'medieval' (Lewis, 1984, p. 229), to which Lewis's response can be neatly adapted here: 'what's wrong with sounding baroque? More power to the baroque metaphysicians.' Also see, (Darby, 2018) for further discussion on this supposed advantage.

based ontology is the belief that all experimental evidence is of particles and the changing distance relations between them: 'To put it in a nutshell, particle evidence is best explained in terms of particle ontology' (Esfeld & Deckert, 2017, p. 33). Initially, this stipulation raises some circularity worries as the evidential support for the Super-Humean ontology is derived from interpreting physics in a way that only arises if one first assumes a preferential ontology: particle ontology is adopted in part because our physical evidence is about particles, while part of the reason for conceiving of our evidence as particles is because it is best explained in terms of a particle ontology.[51] It is equally plausible for proponents of contrasting ontological systems to make analogous claims: a field ontology is adopted in part because our physical evidence is about fields and part of the reason for conceiving of our evidence this way is because it is best explained in terms of a field ontology.

So, to support their suggested interpretation of our evidence, Esfeld and Deckert expand on sufficiency by suggesting that the ontology must be connected to observable facts (where observable facts are, for example, dots displayed on a screen as the results of the double slit experiment in quantum mechanics). An explanatory chain is required that connects our empirical observations to the ontology; one that allows us to interpret our evidence as really being about what sits at the end of this chain. Start by explaining the observable facts in virtue of neuroscientific facts, and on to facts about molecular biology, chemistry, and statistical physics of large ensembles of elementary particles, before finally reaching the ontological facts: fundamental particle physics (understood as nothing more than particles and distance relations).[52] Their core claim is that the *best*

[51] (Wilson, 2018c, p. 585) also raises a similar circularity objection.
[52] This compositional chain is suggested by Esfeld and Deckert (Esfeld & Deckert, 2017, p. 32).

explanation of our observable facts is found by interpreting our evidence to ultimately be about particles and their spatial configuration. It is noteworthy to underline that what counts as the 'best' explanation is being presented in a Super-Humean light as simplicity partly determines the best explanation: 'The simplicity and parsimony of this proposal are part of the case for its being the best explanation' (Esfeld & Deckert, 2017, p. 33).

To summarise, in addition to being able to interpret physics in terms of the proposed ontology, an explanatory connection must be established between the proposed ontology and observable facts. This connection allows for our evidence to be conceived of in terms of the proposed ontology. Thus, when Esfeld and Deckert state that an ontology must be empirically adequate (Esfeld & Deckert, 2017, p. 3), what they have in mind is the establishment of a causal chain connecting the epistemologically evident facts to the ontology via the best explanatory causal chain, for which 'best' is determined in part by simplicity.

3.0 ULTIMATE-HUMEANISM

3.1 THE FINAL REDUCTION

Having established both a clear means of assessing ontological minimality and fleshing out the Super-Humean sufficiency requirement, it stands to reason that if an alternative ontology achieves greater simplicity while also being sufficient, it takes precedence over Super-Humeanism. To approach this goal, a further reduction of elements in the Super-Humean ontological set is required. This can be achieved by separating the present particle configuration (Configuration_1) into individual particles indexed to a particular spatial location relative to all other present particles:

$$SH = \{\text{Particle}_1, \text{Particle}_2, \ldots, \text{Particle}_n, \text{Spatial Distance Relations}\}$$

From here, one obvious means of attaining greater minimality is to eliminate all but one particle to form the following ontology:

Absolute Humeanism = {Particle_1}

Such a view is arguably the simplest ontology attainable but whether such ontology is sufficient is doubtful. One would need to assume a kind of ontological monism (e.g., (Schaffer, 2010)) and then connect the singular entity to the observable facts. Yet, in the Super-Humean case, it is the plurality of individual particles in a changing configuration that explains the observable facts. Some adaptation might be plausible by treating the configuration as an indivisible entity (as is discussed and rejected by (Esfeld & Deckert, 2017, p. 28)), but this would either result in a contrasting non-local minimum ontology or lack a plausible explanation of the observable facts. Instead, a more plausible ontology consists of the particles required for the existence of a single human brain:

Ultimate-Humeanism (UH) = {Particle_1, Particle_2, …, $\text{Particle}_{\text{<enough for a brain>}}$, Spatial Distance Relations}

This view consists of some particles in a changing configuration such that they constitute the macroscopic stability of a human brain and the manifestation of a phenomenology that includes conscious experience (assuming consciousness supervenes on the brain and the falsity of the extended mind thesis).[53] Everything else, all laws, all intrinsic properties, all non-present particles, and all non-brain present particles are nothing more than nomological parameters that feature in the explanation of the motion of the brain-particles but are themselves not part of the ontology. So, in addition to the Super-Humean

[53] This is a mainstream and uncontroversial assumption that while worth explicitly noting is not something I will argue for. Relatedly, Lewis makes similar arguments for the supervenience of mind on the brain in (Lewis, 1994b).

reduction, Ultimate-Humeans assert that particles outside the brain are identical with a particular configuration of changing brain-particles, and are therefore no addition to the ontology. Where a Super-Humean thinks the property of charge is nothing more than a fictitious entity, merely a collection of particles in a changing configuration, the Ultimate-Humean thinks that all particles outside of a brain are nothing more than fictitious entities, merely collections of brain-particles in a changing configuration. Just as the Super-Humean asserts that admitting mass and charge as intrinsic properties of particles introduces surplus structure to the ontology, the Ultimate-Humean asserts that admitting all present particles also introduces surplus structure to the ontology.

As the goal here is to show that proper adherence to the Super-Humean methodology results in Ultimate-Humeanism, it is important to address any potential disanalogies between the Ultimate-Humean reduction and the Super-Humean reduction. To begin, concern with the reductive analogy might arise as the Ultimate-Humean reduction applies to particles and not properties. Yet, Ultimate-Humeanism constitutes a similar kind of reduction that is just applied to a different element in the ontological set. Where Super-Humeans go from all properties to one, Ultimate Humeans go from all particles to some. Additionally, Super-Humeans already reduce non-present particles in their adoption of presentism and if comparisons of minimality are to hold between Eternalist Humeanism and Presentist Super-Humeanism, the Super-Humean reduction already applies to particles.

There might also be resistance to Ultimate-Humeanism's discriminative privileging of some particles, the brain-particles, over others, with the consequence that the reduction is mistakenly dissimilar from that which Super-Humeans endorse. Yet, once more Super-Humeans do the same thing with properties by discriminatively privileging one property,

the distance relations, over others. The same is true for non-present particle configurations as Super-Humeans are selecting only some of a certain element to serve in their ontology: only the present particle configuration. Overall, Super-Humeans are selective about some temporally distinct particles (and one type of external property) but not others and Ultimate-Humeans are selective about some spatially distinct particles but not others.

Lastly, as Ultimate-Humeanism does not attain greater simplicity by locating a different type of element but only by locating a token quantity of the same type, disanalogy may again arise. Super-Humeanism goes from many different types of properties to only one. Ultimate Humeanism goes from one type of thing, a spatial totality of particles, to a smaller totality of that type, only some of those particles. If a gain in simplicity only matters when there are fewer kinds of entities and not when there is less of the same kind of entity, the reductive analogy fractures. Yet again, problems with this reasoning arise. When Super-Humeans adopt presentism, they commit to a smaller number of particles (only the present ones) and claim it to be preferable because of the gain in simplicity. Further, any claim that present particles are of a different type to non-present particles will allow Ultimate-Humeans the same thought: brain-particles are of a different type to non-brain-particles. Either way, it is preferable to have fewer of the same type, or brain-particles are a different type.

These *prima facie* issues aside, the Ultimate-Humean ontology is easily assessable as more parsimonious than Super-Humeanism since it dispenses with elements from the Super-Humean ontology while adding nothing new: UH $\subsetneq$ SH. Therefore, it follows from the Super-Humean guiding principle of minimal sufficiency that if Ultimate-Humeanism satisfies the criteria of sufficiency, it ought to be preferred over Super-Humeanism.

To begin demonstrating Ultimate-Humeanism's sufficiency, no drastic alterations are required to the Super-Humean account. We begin by interpreting physics as being about particles and distance relations. Our best physical theories are constructed in the same way from relative particle movement. Nothing changes in this respect and as such, compatibility with our best physical theories translates over from Super-Humeanism. The sole change in the ontology is that the number of particles is reduced but the same principles employed by Super-Humeans when accounting for physics in terms of particles and distance relations still apply. They build the basic principles of physics from a totality of particles and relative distance, and likewise, for Ultimate-Humeanism, there is still a totality of particles that stand in relative distance to each other such that allows them to build the basic principles of physics in the same way. Only, in this case, the totality of particles stops at the edge of the brain (further discussion of Ultimate-Humeanism compatibility with physics can be found in §3.3).

From here, the observable facts need to be connected to the ontology such that we can interpret our evidence as ultimately being about the Ultimate-Humean ontology. This causal chain starts with the observable facts and runs down through neuroscience, molecular biology, chemistry, statistical physics of large ensembles of elementary particles, and finally to fundamental particle physics (understood as nothing more than brain-particles and distance relations). Ultimate-Humeanism commits to a causal chain connecting the observable facts solely to brain-particles such that fundamental particle physics is interpreted as being about nothing more than brain-particle configurations. This is a similar level of support to that which Super-Humeans give for interpreting physics to be about nothing more than particle movement. So, if Super-Humeans can claim that their causal chain explains why our evidence is all about particles, then the Ultimate-Humean

can claim that their causal chain explains why our evidence is all about brain-particles. Ultimate-Humeanism is capable of both constructing our best physical theories from their ontology and explaining how the observable facts connect to the ontology.

Further, in the Ultimate-Humean model, the causal chain never extends beyond the brain to explain the observable facts; the dots displayed as the results of the double-slit experiment are explained as being about a particular group of neurons that constitute the experience of those dots. We need not go beyond the brain by further stipulating that facts about particles in a vacuum chamber cause a particular configuration of light to be emitted from the display screen, which in turn causes photons to reach the back of an eye, such that they then cause the neurons to cause experiences. The Ultimate-Humean causal chain is far simpler (and the ontology overall) and as a result, is a better explanation of the observable facts. This correspondingly justifies the initial ontological assumption for interpreting our evidence in a way preferential to Ultimate-Humeanism; for interpreting our evidence to only be about brain-particles. In short: the reason our evidence is only evidence of brain-particles and not of all particles is that when our evidence is of brain-particles, the explanation is simpler.

To summarise, Ultimate-Humeanism is minimally sufficient. It is a simpler ontology whereby all elements that are not brain-particles figure in the explanation of the motion of brain-particles, but they are not themselves part of the ontology. All experimental evidence in fundamental physics is conceived of as evidence of relative brain-particle positions and the change of these positions. Ultimate-Humeanism is, true to its name, the implementation of the Humean reductive doctrine to its sufficient maximum, maintaining scientific realism without including the representational instruments employed by our physical theories in the ontology.

3.2 Realism & The Natural World

With the positive proposal outlined, I will now consider some prominent concerns. The first is that Esfeld and Deckert seek a minimal ontology of the natural world, and 'by the natural world, we mean the physical, spatially extended world' (Esfeld & Deckert, 2017, p. 8). The scope of their project appears to be confined to deriving a minimal ontology after assuming realism about the external world – to deriving the most parsimonious ontology where other brains, dogs, houses, etc. physically exist. Ultimate-Humeanism might seem to conflict with this realist project as those physical objects become nothing more than brain-particles and their movement. But this is not the same as asserting that they do not physically exist. Super-Humeans and Ultimate-Humeans both start from the realist assumption that the physical world exists, that entities alike other brains, dogs, houses, etc. are real, and both also reduce them down to something more fundamental without eliminating them from the physical world. Ultimate-Humeans simply push the reduction further, but this does not diverge from the Super-Humean's realist starting point.

The scope of Esfeld and Deckert's realist project is explicitly confined to what they define as the 'natural world', which is principally based on Descartes's Cartesian criteria for matter that includes extension and motion: *res extensa*. Again, it may seem *prima facie* that Ultimate-Humeanism fails to provide an ontology of the natural world; a world that includes extension and motion. To address both charges, first on extension, Ultimate-Humeanism is a reduction to the particles in a brain and brains are spatially extended. The reduction is not to minds, to matter *res cogitans*, but to a strictly smaller area of matter *res extensa*. Second, on motion, there is just as much moving, even if not as much stuff that moves, in Ultimate-Humeanism as there is in Super-Humeanism. The ontology

consists of particles in a configuration of changing distance relations just as Super-Humean does. Ultimate Humeanism squares with the criteria set out for the natural world just as it does for realism more generally.

Moreover, even though Ultimate-Humeanism is compatible with the Cartesian constraint, there are questions to be asked of the constraint itself. It is a principle that uniquely identifies extension and motion as the parts of our evidence that serve as a guide to ontology. Yet, these criteria are selected based solely on an *a priori* principle that connects them to the natural world without an independent justification. Other parts of our evidence could be elevated in this fashion and serve a corresponding role as a guide to the ontology of the natural world. For example, colour forms a ubiquitous part of our empirical evidence and could be privileged such that it plays a critical role in determining what counts as an ontology of the natural world: matter as *res pigmentum*.[54] Are macroscopic objects not part of our empirical evidence? One could equally favour the content of our evidence of tables and require their inclusion in the definition of the natural world: matter as *res mensa*. Overall, reductionism is not acceptably constrained to a natural world defined by cherry-picking parts of our evidence to privilege.

Consequently, the natural world, even when understood within the scope of the realist Super-Humean project, does not include a request for all present particles. From here, continuing to stipulate that an ontology must extend across all of space is an unjustified claim about what elements of the ontology are indispensable. If this move is permitted, then there is just as much reason for the Humeans, or non-Humeans to assign bedrock status to the parts of the ontology Super-Humeans aspire to reduce. The question here in

[54] To support colour's ubiquity, even agents with achromatopsia can perceive the colours black and white and differentiate between shades of grey in a way analogous to those with colour vision.

its most succinct form is as follows: why must we keep all matter? And Super-Humeans have no good answer.

3.3 Limitations on Reinterpreting Physics

Absent Descartes' authority, Super-Humeans lack a principled reason to differentiate between the elements they deem nothing more than nomic bookkeeping devices and those they deem part of the ontological fabric. An appeal to physics, resistance to Ultimate-Humeanism's pervasive reinterpretation of physics as nothing more than a guide to the present particle configuration in a brain, might underscore a response that would permit Super-Humeans the distinction they seek. Yet, if Ultimate-Humeanism results in a malformed physics that has been reinterpreted to a greater extent than acceptable, the Super-Humean reinterpretation of physics as nothing more than a guide to the behaviour of present particles and their distance relations falls afoul of similar criticism.[55] Most prominently, there ought to be principled reasons affirming when reinterpretations of physics are excessive, and a demonstration that Ultimate-Humeanism infringes them while Super-Humeanism does not. Such (non-arbitrary) distinction is not easily forthcoming and without it, the Ultimate-Humean reinterpretation is not freely dismissed.

To clarify, it is worth explicitly distinguishing between simplicity in ontology and simplicity in representation. One could argue that the simplest way to codify particle configurations is to assign them with properties: the simplest explanation for why some particles exhibit repeating patterns of changing distance relations is because they have mass or charge. The same line of reasoning applies to Humean thinking too: the simplest

[55] Super-Humeans appear capable of reinterpreting physics such that physics could never rule out their ontology; (Matarese, 2021) argues that these worries are particularly damming for a view prised for its naturalism.

explanation for why the property configuration exhibits patterns is because of the laws. Humeans and Super-Humeans are willing to trade off simplicity in the representation of the ontology for simplicity in the ontology. If this becomes problematic when pursued to the extent Ultimate-Humeans have, a non-arbitrary standard is required concerning how complicated the representation can get before it becomes problematic.

Ultimate-Humeanism pushes complexity in representation further, but not so far such that no interpretation of physics is viable. The brain contains enough particles such that the general principles of quantum mechanics can be derived in the same way as Super-Humeans do.[56] And it is possible to construct physics from the brain facts since it includes all the different inputs including testimony that the brain actually receives; this is enough to recover our actual physics. Yet, a huge number of additional higher-level laws will be required to account for the idiosyncratic movement of particles on the peripheries of the brain; those parts that interact with the fictitious non-brain-particles. Each different kind of phenomenal experience requires a distinct higher-level law explaining why those particles behave as such (and why it repeats every time the same conscious experience is exhibited). Actual physics is retained but our interpretation of it becomes much more complicated because of these additional higher-level laws. However, this increase in interpretational complexity is traded off against a gain in ontological simplicity.

Some may contest that the peculiar content and drastic increase in the number of physical principles required is excessive. It might seem obvious that the reason brain-particles

[56] As an adult human brain weighs on average 1.3 kg; imprecisely assume the number of grams in a mole of brain matter to be somewhere between oxygen at 15.999 grams and carbon12 at 12 grams, say 13 grams. With Avogadro's constant, $6.02214076 \times 10^{23}$, expressing the number of particles in a mole, and there being 100 moles in the average adult brain, we can estimate that there are roughly 6×10^{25} brain-particles in the Ultimate-Humean ontology. This is more than enough to derive an interpretation of fundamental particle physics analogous to the means employed by Super-Humeans with all present particles.

behave in certain ways, such that they constitute certain kinds of phenomenal experience, is because of the particles outside of the brain! This is certainly a simpler representation, with fewer laws, and less revisionary physics, but pursuing this argument has troublesome implications. Super-Humeans already complicate physics by making it all about particles; the simplest way of explaining quantum field theory is to posit fields! This is a similar argument employed by those who reject interpreting quantum field theory in terms of particles. The issue is that by reinterpreting physics one struggles to provide a proper account of it. For example, (Matarese, 2021, p. 801) argues when Super-Humeans demonstrate that their particle-only ontology can be applied to Quantum Field Theory by adopting the Dirac sea model, they lack an account of how the language used by physicists reduces to particle movement. Likewise, (Oldofredi & Ottinger, 2021, pp. 30-32) raise several issues of a similar sort. If the simplest interpretation of physics matters, and outweighs the gain in ontological simplicity, Super-Humeans undercut themselves as their methodology is based on trading in a complication of physics for a gain in ontological simplicity. In this sense, Ultimate-Humeanism merely constitutes the proper employment of the Super-Humean methodology by trading in the most complicated representation for the simplest (sufficient) ontology.

One final point on physics is that physics itself is a good place to look for a reason to privilege all of space and consequently, for a principled reason to prevent an ontological collapse into Ultimate-Humeanism. Space is special according to physics.[57] Our best physical theories are defined on a spatial manifold e.g., all four fundamental constants are a function of distance, quantum field theory is expressed on a space(-time) manifold, etc.

[57] Arguably it is space-time that is special to physics but given that this would already cause problems for the Super-Humean position, the argument here grants them the assumption that space is special to physics.

Spatial extension is a ubiquitous feature of our physical theories and adopting space substantivalism based on physics would provide a well-motivated reason to assign bedrock status to space. Unfortunately for Super-Humeans, this kind of reasoning does not translate into their interpretation of physics. Super-Humean space is entirely relational, consisting solely of particles and distance relations and nothing about this understanding privileges space. The saving constraint would be that space is special because of physics, but according to Super-Humeans, space is fully accounted for by nothing more than a changing configuration of particles. By arguing that physics is ultimately all about distance relations, Super-Humeans deprive themselves of a reason for keeping all of space. Ironically, the foundation on which Super-Humean naturalistic metaphysics rests is undermined by their gutting of physics.

4.0 CONCLUSION

4.1 PARSIMONY PROCURED

After accepting Super-Humeanism's methodology, the minimal set of entities that form an ontology of the natural world, given what science tells us about the natural world, is Ultimate-Humeanism. The ontology is simpler while being in a position to construct our best physical theories (given that they are allowed to interpret physics with the same kind of leniency as Super-Humeans do). Two serious objections - that Ultimate-Humeanism is not an ontology of the natural world, and that it requires a radical reinterpretation of physics have also been addressed to further substantiate the view. To press the crux of this conclusion, those who uphold parsimony to be the crucial factor in ontological matters have been provided with a new minimal ontology. It turns out that Super-Humeanism is not minimal: it is an ontology bloated by excessive particles. To make do

with as meagre a diet as possible, those with methodological sympathies for minimality ought to embrace Ultimate-Humeanism instead.

Further, it is worth noting that although Ultimate-Humeanism is simpler than the supposed minimal ontology, it is not the only option that would leave Super-Humeanism without principled motivation. For example, if it turns out that there are good reasons for an ontology to include all spatially extended elements that are represented by observable facts, the minimally sufficient ontology would be more akin to the following:

> *Berkeley-Humeanism* = {Particle_1, Particle_2, …, $\text{Particle}_{\text{<enough for spatially extended perception>}}$, Spatial Distance Relations}.

Super-Humeans face the same problems with this ontology as it is simpler than theirs but does not posit all present particles. Likewise, without a principled reason to retain all present spatial particles, even just one less element is going to raise the same issues:

> *Super-Lite-Humeanism* = {Particle_1, Particle_2, …, Particle_{n-1}, Spatial Distance Relations}

Although any of these ontologies could play the role of Ultimate Humeanism in the larger dialectic of this paper, I have argued that the limits of reduction are found by positing the particles and distance relations for a brain. Interestingly, this reductivist limitation is shared by another philosopher who Esfeld and Deckert quote as inspiring their project: ‘Super-Humeanism can therefore be conceived as being inspired by the following quotation from Russell (1927): There are many possible ways of turning some things hitherto regarded as “real” into mere laws concerning the other things. Obviously there must be a limit to this process, or else all the things in the world will merely be each other’s washing’ (Esfeld & Deckert, 2017, p. 49).

For Super-Humeans that limit is particles and their changing distance to each other. But for Russell and Ultimate-Humeans, that limit is just a brain. The very next line following this quote, which is conveniently absent from the Super-Humean extract, states as much: 'But the only obvious final limit is that set by phenomenalism perhaps one ought to say, rather, that set by solipsism' (Russell, 1927, p. 325). I follow Bertrand Russell, in accepting there are many ways of converting ontological elements into little more than laws about other things and this process must come to an end, but that limit is set by positing just a brain. For me, I mean to imply no necessary connection by using 'Solipsism' to describe Ultimate-Humeanism. It is simply a contingent fact that all that exists in the minimal ontology of the natural world are particles arranged brain-wise.

4.2 The Anti-Thesis

For those with less ardent sympathies for ontological minimality, I present Ultimate-Humeanism as a contingent thesis. I have demonstrated that accepting minimal sufficiency as the sole factor in determining ontology leads to an ontology consisting of nothing more than brain-particles and their relative distance relations. Most will find Ultimate-Humeanism's brain-particle only ontology an unacceptable resting place, instead affirming Ultimate-Humeanism to be nothing more than a consequence of ontological reduction pursued too far. For those so inclined, a different argument can be drawn out:

1. Minimal sufficiency is the sole principle guiding ontology.
2. If minimal sufficiency is the sole principle guiding ontology, then Ultimate-Humeanism is the ontology of the natural world.
3. Ultimate-Humeanism is not the ontology of the natural world.

   Ergo: Minimal sufficiency is not the sole principle guiding ontology.

As the methodological principle of minimal sufficiency implies Ultimate-Humeanism, if Ultimate-Humeanism is incorrect, then we can conclude *modus tollens* that the methodological principle is also incorrect. When the principle is consistently endorsed it leads to an undesirable ontology and on rejecting this ontology, the methodological principle that entails it is also rejected.

It is worth highlighting here that regardless of which conclusion we accept, Super-Humeanism lacks principled motivation. Either a methodological commitment to minimal sufficiency is retained and Ultimate-Humeanism obtains in place of Super-Humeanism. Or, the methodological commitment to minimal sufficiency is rejected; and as both views rely on this premise, both Super-Humeanism and Ultimate-Humeanism lack principled motivation.

4.3 Final Remarks

Minimal sufficiency, as defined by Super-Humeans, fails to secure the ontological consequences Super-Humeans aim for. One means of resolving the issue is to strengthen sufficiency. Instead of relying solely on empirical adequacy and the capability of an ontology to construct our best physical theories, other methodological virtues could be introduced to inflate the minimally sufficient ontology. With some fine-tuning, Super-Humeans might be able to capture a specific combination of theoretical virtues that leads to their reductive endeavour resulting in their desired ontology. Some notable candidates include strengthening the explanatory requirement, perhaps to the degree that includes not only a connection between fundamental physics and ontology but also connects to higher-level sciences (Wilson, 2018b). Another option is to further increase the connection between observable facts and ontology; strengthening how the ‘manifest image’ of the world is explained in terms of the ‘scientific image’ ((Sellars, 1962) and

(Allori, 2013)). But, by taking this approach they risk strengthening sufficiency too much, and the resulting ontology being more expansive than desired. Likewise, if they do not strengthen it enough, Ultimate-Humeanism or other more minimal ontologies will obtain in place.

This renewed balancing act is a consequence of the Super-Humean's devaluing of theoretical virtues other than minimality. Consequently, this paper might be best understood as a cautionary tale about the limits of reduction, which exposes the slippery slope Super-Humeans have created. After having greased their skis, Super-Humeans incorrectly expect their descent to cease before arriving at the incline's lowest point. Ultimate-Humeanism is what sits at the bottom of this ontological hill and as most (Super-Humeans in particular) will find Ultimate-Humeanism repugnant, to avoid it methodological revisions are required. Whether reform can be enacted without comprehensively undermining the motivation for Super-Humeanism is unclear, but the prospects seem poor.

Finally, it is worth highlighting that the general dialectic outlined in this paper applies to all comparable reductivist programs. Any who accept a reductive methodology with a preference for minimality will find themselves facing similar questions. Individuals inclined as such need to articulate how their ontology is simpler while detailing which virtues simplicity is to be balanced against such that the desired ontology results (without resorting to arbitrary constraints). Naturalistic approaches need to explain what counts as empirically adequate and to what extent revisions to physics are permissible along with an account of how the ontology satisfies such criteria. Each of these themes features prominently in the derivation of Ultimate-Humeanism; a fruitful consequence of detailing such ontology, with its solipsism and unorthodox physics, is that it highlights the costs of

an unchecked attachment to ontological simplicity. Having exemplified the ontological minimum, we are now better informed on what the maximum reduction looks like and are better placed to assess whether it is worth the price.

PAPER V: (EXPLICITLY) EPISTEMICALLY MOTIVATED METAPHYSICS

ABSTRACT

This paper links together the previous chapters of the thesis to draw out a unified epistemic methodology for metaphysics which I call *epistemic metaphysics*. The core tenet of this is for metaphysical theories to be accompanied by epistemic principles that provide reasons to accept the corresponding metaphysical theories. Uniquely, I construct epistemic principles via a close examination of contemporary metaphysical debates where epistemic factors play prominent roles. As a result, the epistemic constraints are derived from a generalisation of already imposed epistemic factors and are not imposed externally. In turn, when doing metaphysics, both factors should be considered together as considerations for one will have direct implications for the other. The ultimate metaphilosophical moral is for investigations into metaphysics and epistemology to be conducted in unison.

OVERVIEW

In §1 the structure of the proposal is outlined and two contrastive interpretations for the strength of the motivating connection between epistemology and metaphysics are considered. In §2, examples of epistemic metaphysics are drawn out from the other papers in this thesis before the proposal is applied to itself. Along the way, various aspects unique to each individual paper are used to further clarify parts of the overall proposal. This includes processes relating to combining epistemic principles, the symbiotic evaluation of metaphysics and epistemology, defining (in the relevant respect) what counts as epistemic, and clarifying the proposal's 'explicit' demands. In §3, the paper discusses the uniformity requested of epistemic motivation, chiefly in relation to those who make systematic metaphysical proposals. This is followed by a discussion on the broadly applicable nature of my proposal that draws parallels with (Bennett, 2009) before showing how comparable metaphilosophical accounts of (Ladyman & Ross, 2007) and (Bealer & Strawson, 1992) might be incorporated into my account of epistemic metaphysics. Finally, the thesis concludes by outlining the methodological shift which results from adopting epistemic metaphysics.

1.0 AN EPISTEMIC METHODOLOGY

1.1 STRUCTURE

The goal of this paper is to plainly define and motivate what I call *epistemic metaphysics* which is, in short, a proposal for metaphysics to be more clearly connected to epistemology. To begin, it will be useful to outline the structure of the proposal, its various parts, and how they connect to one another. There are four elements to the proposal: metaphysics, epistemology, the connection between the two, and a generalised meta-philosophical maxim. I will take each of these in turn and discuss what they collectively entail.

First, metaphysical theories aspire to, among other things, describe what exists, what is true, or the relationships between different parts of our world. They are proposals describing the nature of a given metaphysical phenomenon. For example, (Lewis, 1986) outlines an account of the metaphysical notion of modality in terms of concrete possible worlds, (Bourne, 2006) defines the metaphysics of non-present truth in terms of abstract entities, and (Barnes, 2018) suggests that metaphysical dependence relations are best understood in terms of non-symmetry. Each proposal describes the nature of a given metaphysical matter. While it might be difficult to precisely pin down the boundaries of metaphysics (problems have been raised with defining the subject matter of metaphysics (Bennett, 2016); (Hofweber, 2016)), this is no particular issue for my proposal. All that is required here is to accept that there are theories and they aspire to describe the nature of some given, broadly, metaphysical matters.[58]

[58] Even if it turns out that there is no clear metaphysical subject matter, my proposal is readily adaptable for theories *tout court*. One might call arguments about what is right or wrong ethics, or arguments about the most viable interpretation of quantum mechanics (philosophy of) physics. But they are theories about 'what the world is like' and outlining corresponding epistemic motivation still applies. This project would

Second, epistemic principles can be defined in terms that appeal to any epistemic feature of agents like us such as our capacity to know, explain, perceive, etc. From the faculties we possess that allow us to understand the world, principles can be constructed that put a specific demand on our grasp of the world. In this sense, one could, for example, define epistemic principles that require our theoretical proposals to be knowable, explainable, or empirically verifiable. Conversely, one could define epistemic principles that require our theoretical proposals to contain no posits that are unknowable, unexplainable, or empirically unverifiable. These principles are tightly connected to our metaphysical theories and should be sufficiently fine-grained such that it is clear to see if and how the clauses contained within the principle are satisfied by our theoretical proposals.

Third, regardless of what epistemic features are chosen to define the principle, an additional and important condition is that the principle must be formulated in a way that allows for a *motivating* relationship to exist between principle and theory. Epistemic principles must influence the acceptance or rejection of metaphysical theories. This connection is methodological and should capture the impetus one has for taking a stance in favour or against a respective metaphysical proposal. To ensure transparency, all that is required for the principle to be motivating is for it to provide *a reason to accept (or reject)* a metaphysical theory.

Then, the first three elements and their appropriate connections are generalised to mark out a particular branch of metaphilosophy. The goal here is to generate all-purpose epistemic rules for metaphysics from within metaphysics. By acknowledging particular instances where epistemic principles motivate metaphysical theories, these connections

go beyond metaphysics and epistemology and would instead aim for a truly universal, explicitly epistemic, approach to inquiry.

can be generalised into a maxim that can then be employed to motivate other, distinct, metaphysical theories. For example, if it turns out that particular metaphysical theories of modality and non-present truths are motivated by an explanatory principle, one can then generalise the motivating connecting into an explanatory metaphilosophical maxim that is then employed in another area of metaphysics, say of metaphysical dependence, to motivate a particular theory of dependence. These rules come from looking at what motivates metaphysics in one area and then generalising that connection to apply universally to all areas of metaphysical inquiry. It is the process of taking the connections already posited between some specific epistemic principle and some specific metaphysical theories and codifying it into a general epistemic guideline that motivates a particular class of metaphysical theories.

A key presumption of this metaphilosophical proposal is for the epistemic motivation of metaphysical theories to be uniform. This factor might be best expressed in terms of a normative constraint for metaphysical theorisers i.e., if someone accepts an epistemic principle is motivating in one domain of metaphysical inquiry, then they should also accept the generalised maxim based on the principle to be motivating in all domains of metaphysical inquiry. When someone accepts an epistemic principle to motivate a particular metaphysical theory, they are implicitly committing to a particular way of doing metaphysics. The motivation is not constrained to any particular metaphysical domain but outlines the kind of factors that an individual takes to be pertinent epistemic motivation for metaphysics in general. For example, if someone upholds that truths ought to be reconcilable with an account of how we could know them (Peacocke, 1999), then this applies to modal truths, non-present truths, fundamental truths, etc. If the principle holds, then all metaphysical domains positing truths must be consistent with it and

adopting a piecemeal approach, whereby one picks and chooses what metaphysical domains the principle applies in, is unacceptable. I assert this to be the default position.

I leave the impetus on others to contest this default assumption and account for why one is liable to accept an epistemic principle as motivating in one domain but not in others. Anyone who wants to adopt this kind of fragmentary approach to metaphilosophy ought to account for why such non-uniformity is permissible. It seems odd for someone to accept a particular epistemic principle in one domain but reject it in another. The proposed process of metaphilosophical generalisation builds on this underlying uniformity by solidifying the all-purpose nature of any given epistemic motivation – if you find a certain epistemic principle motivating, then you ought to accept *all* of the metaphysical theories that it motivates.

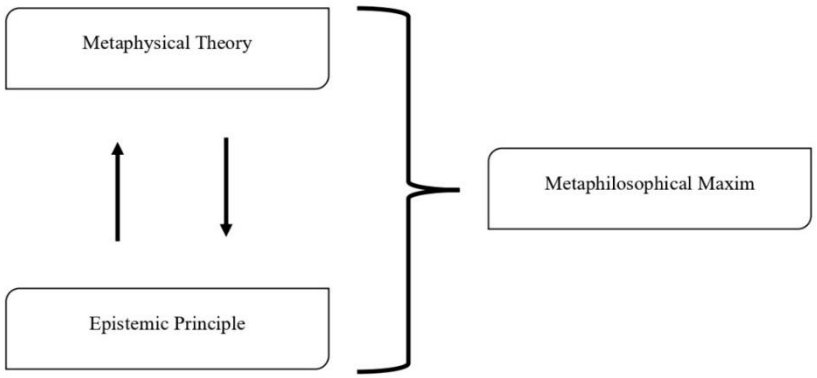

Figure 8 The Structure of Epistemic Metaphysics

Although no addition to the four elements already stated, the epistemic principle and how it motivates the respective metaphysical theory ought to be specified outright. Included in my proposal for epistemic metaphysics is a general methodological claim that one

should not assert something to be the case without demonstrating it to be the case. It is no good to keep quiet about what motivates metaphysics nor is it any better to say that some epistemic principle motivates one's metaphysics without showing this to actually be the case. The connections, and precisely how they are connected, must be explicit. One may well state that their theory is motivated by an epistemic principle but merely stating is not sufficient, instead one must provide a detailed account of how the epistemic principle motivates the view before one ought to accept that to really be the case.

Overall, the four elements outlined in this section, in combination, capture the proposal I propose in this paper. Define a metaphysical theory, a motivating epistemic principle, and allow for the generalisation of this connection into a meta-philosophical rule and you have an (explicitly) epistemic approach to metaphysics.

1.2 Strength of the Motivating Connection

Arguably, the most important part of epistemic metaphysics is the motivating connection between epistemic principle and metaphysical theory. As such, it is worth expanding the discussion to ensure the connection is clear. In particular, reflection is warranted on the strength of the connection. At a first pass, I suggest that the motivating connection between epistemic principle and metaphysical theory is absolute, such that any theories that fail to adhere to a respective epistemic principle are rejected outright. There is no middle ground; consistency with the demands expressed in a principle either does or does not make a metaphysical theory unacceptable. There is, however, the alternative of understanding consistency or inconsistency with the demands of an epistemic principle to simply be one among several merits or demerits of a theory. If a theory meets the standards of an epistemic principle, then that is an advantage of that theory. Likewise, a

theory that fails to meet the epistemic principle's standards is worse off than those that do. Call the two approaches:

> *Strong Motivational Approach*: epistemic principles ratify certain metaphysical theories and rule out others.
>
> *Weak Motivational Approach*: epistemic principles establish advantages or disadvantages of certain metaphysical theories.

The difference between the two approaches comes down to asking: when doing epistemic metaphysics, should we *only accept* those theories consistent with our recognised epistemic principles or should we merely *prefer* those theories? *Prima facie*, either approach appears plausible. In fact, I imagine the most pragmatic approach will allow free choice in any given instance. For some principles, consistency with them is non-negotiable while for others it is a minor shortcoming. What matters here is that this variation is accounted for and that the strength of the connection is clearly expressed.

Much of what I will discuss in this paper concerns the strong motivational approach; places where metaphysicians have found incompatibility with a particular principle to be intolerable. This is mainly due to these areas of dispute being the most useful for exploring the project of epistemic metaphysics. However, I also have some concerns that the weak motivational approach will not lead to a fruitful account of epistemic metaphysics. To express these worries, I will first expand on an account of epistemic metaphysics that adopts the weak motivational approach exclusively.

In the weak motivational approach, no metaphysical theories are ruled out outright on epistemic grounds. Some are just better than others with respect to meeting the demands of certain epistemic principles. They are weighed up against various epistemic factors:

theory *x* is good because it adheres to principle *y* but is bad because it does not adhere to principle *z*. In this sense, the theory is evaluated holistically by assessing the various epistemic principles the theory is consistent with against those that is it inconsistent with. In the end, we could look to discern the 'best' theory to be that which, on balance, comes out on top.

Take, as a brief example, Armstrong's (Armstrong, 1983) proposal for laws to be understood as a second-order necessitation universal that links two first-order universals and the explanatory objection as described in Paper I.[59] Just because the elements in this theory possess no explanations for how they play their roles, this does not mean it should be immediately consigned to the scrapheap, even by those who accept something akin to an explanatory principle. There might be other epistemic principles that the theory does adhere to and when all of them are taken into consideration, it might turn out to be the most viable option.

This kind of situation is conceivable and might be viewed favourably by some. However, I can foresee one *prima facie* problem with the weak motivational approach: how does one conduct this normative evaluation? How is a failure to meet the standards of one epistemic principle comparable to meeting the standards of another? Even for closely related principles I can see there being an issue of incommensurable comparisons. To show this, consider another simplified example, comparing Humean and non-Humean theories. One virtue of the Humean approach is that, on the whole, they are capable of meshing their account of laws with those used by scientists. Conversely, one virtue of the non-Humean approach is that, on the whole, they are capable of explaining why the world

[59] A detailed discussion fitting the content of Paper I into the structure of epistemic metaphysics comes later in §2.1.

exhibits modal regularities. Neither side possesses the virtue of the other to the same degree. Humean laws fail to explain why the world exhibits robust modal patterns and non-Humean laws are often seen to be incompatible with the laws actually used by scientists.[60]

The question then is: which matters more? Humeans tend to think adhering to one epistemic principle is more important while non-Humeans tend to think the opposite. To settle this impasse, we would need to be able to determine how much of an advantage explaining the regularities of the world is, how much of a disadvantage failing to exposit laws like those used by scientists is, etc. The problem is that there exists no common standard to answer these questions. Worse still is that this kind of partisan divergence on the epistemic motivation employed for competing theories is not exclusive to this Humean vs. non-Humean debate. In fact, I speculate that there is no consensus with regards to most metaphysics debates when trying to establish how important each view's motivating epistemic principles are.[61] As a result, it appears prohibitively challenging to make any kind of holistic comparison that would determine which is the overall 'best' theory on epistemic grounds.

In contrast, when adopting the strong motivational approach, epistemic evaluations are not a matter of degrees and consequently do not face the same problem. In every instance where an epistemic principle is used to motivate accepting a metaphysical theory no evaluation of how important the principle is required. Instead, the metaphysical theory

[60] For the explanatory problems with Humeanism see (Armstrong, 1988); (Lange, 2023) and for problems with non-Humeanism's fit with science (Hildebrand, 2020, pp. 8, §3.3) outlines a number of arguments against various forms of non-Humeanism and their struggle to account for the laws used in science.
[61] This might be because we are not asked to be explicit about such things.

either adheres to the epistemic principle or does not. The assessment is absolute and as a result, incommensurability is not an issue.

Evidently, the claim here is not that all problems dissipate. It remains unclear whether we should in fact accept any given epistemic principle or whether the metaphysical theory is in fact compatible with it. However, difficulties surrounding quantitative evaluations, establishing how important a principle is, dissipate. The positive result is that instead of intellectual exertion being put into arguing for or against assigning a particular weighting to a principle (which bears no comparison to any other standard), our investigative efforts can be put into arguing for or against a particular principle, which will have a tangible impact making metaphysical theories either acceptable or unacceptable.

Overall, I do not think this discussion provides substantive reasons to move towards the strong motivational approach or to conversely discard the weak motivational approach. As this paper continues, and particularly in the next section, while I will often talk of epistemic principles as providing absolute reasons, one could reformulate this into weaker talk and my arguments carry over to the weak motivational approach. Both approaches are likely to play a role and it is important to allow for metaphysicians to choose how strongly motivating they find any given principle to be.

2.0 EXAMPLES OF EPISTEMIC METAPHYSICS

2.1 PAPER I

To substantiate the proposed framework, I will now consider several cases that fit into the epistemic metaphysics model. Examples of an epistemic principle motivating a metaphysical theory can be found throughout this thesis and I will detail this connection

in the next four sub-sections focusing on Paper I, Paper IV, Paper II, and Paper V respectively. In some places, it is quite easy to unearth this connection and the terminology is already analogous to that used here. But in others, some further analysis is required. Regardless, there remain numerous cases where it is possible to distil a metaphysical proposal, epistemic principle, a motivating connection between the two, and a resulting meta-philosophical rule.

The first paper in this thesis is probably the easiest place to find the elements and connections of epistemic metaphysics. It is a work that has the intentional goal of defining an epistemic (explanatory) connection with various metaphysical theories. The explanatory principle asks any element proposed to play some role in a theory to be accompanied by an explanation for how it plays that role. The motivational force of the principle comes in the fact that the principle provides both a reason to reject certain metaphysical proposals for being unable to meet this explanatory standard and a reason to accept those that do. This connection is shown to already exist across a wide range of metaphysical domains. The explanatory principle motivates rejecting metaphysical theories of laws and chance that employ universals, theories of mathematical truth that employ platonic entities, theories of ground that employ a widespread dependence relation, and theories of sets that employ primitive singletons. While the paper spends less time on the positive connection, the explanatory principle also motives accepting metaphysical theories of laws that employ best systems, theories of chance that employ frequencies, and theories of ground that employ narrow dependence relations.

The paper also briefly touches on the resulting metaphilosophical outcome. Exploring specific arguments where the explanatory principle motivates metaphysical theories allows us to generalise these connections into a unified maxim for metaphysical theory

choice. That maxim serves as a constraint, ruling out certain theories and demarcating between those that do and do not adhere to the demands expressed in the first-order epistemic principle:

> *Explanatory Metaphysics*: when doing metaphysics, acceptable theories adhere to the explanatory principle.

Explanatory metaphysics ratifies *any* metaphysical theories that are suitably explanatory while opposing *every* metaphysical theory that is unexplanatory. The metaphilosophical maxim is generated by generalising the demand expressed in the first-order explanatory principle such that it applies to any metaphysical theories. It is simply the codification of arguments that utilise the explanatory principle to motivate individual metaphysical theories into a universal rule that serves as a guide to theory choice in *all* areas of metaphysics.

As I mention in §1, this process of generating universally applicable metaphilosophical maxims rests on assuming the epistemic motivation of metaphysical theories applies uniformly. I take it to be the default position for any theorist who accepts an epistemic principle to be motivating in one domain of metaphysics will also accept the principle to be motivating in every domain of metaphysics. In this case, someone who accepts that the explanatory principle motivates a particular account of laws containing elements whose roles are explained should also accept an account of chance that contains elements whose roles are explained. Equally so, the standard position for this explanatorily inclined metaphysical is to reject any theory, be that of laws, chance, dependence, etc. that contains elements whose roles are not explained. If someone is motivated by explanatory

considerations in one area of metaphysical inquiry, they ought to be motivated by explanatory considerations in all areas of metaphysical inquiry.

While explanatory metaphysics applies uniformly, this is not to say that everyone ought to adopt it. There exists the option for rejecting the metaphilosophical approach wholesale and embracing an entirely non-explanatory way of doing metaphysics. This move is also briefly touched on at the end of the first paper and is identified as a possible route for those who have staunch commitments to metaphysical theories containing elements lacking any explanation for how they play their roles. This route consists of suitably epistemically *unmotivated* metaphysicians asserting a *meta-epistemic principle* that applies to the first-order epistemic principle; affirming that no such epistemic principle holds. Explanatorily *unmotivated* metaphysicians make a *meta-epistemic* judgment rejecting the explanatory principle and allowing for the permissibility of theories that contain elements whose roles are not explained:

> *Non-Explanatory Meta-Principle*: how an element plays its role in a theory need not be explained.

This meta-epistemic principle affirms, conversely to the first-order explanatory principle, that the explanatory principle does not motivate certain metaphysical theories; that being able to explain how an element plays its role in a theory does not provide a reason to accept that theory and concurrently, that being unable to explain how an element plays its role in a theory does not provide a reason to reject that theory. The meta-principle is much weaker than the first-order principle as instead of demarcating a set of acceptable and unacceptable theories, it is simply ruling against such demarcation. It is a rejection of the reasons initially given for accepting or rejecting metaphysical theories.

Meta-epistemic principles can also be generalised into metaphilosophical maxims. The default result for an individual who takes themselves to be unbound by the explanatory principle in any given area of metaphysics is that they ought to find the principle unmotivating in every area of metaphysics. Call this non-explanatory metaphilosophical maxim primitive metaphysics:

> *Primitive Metaphysics*: when doing metaphysics, acceptable theories adhere to the non-explanatory meta-principle.

This way of doing metaphysics renders every proposal eligible; primitive metaphysics expresses no constraint on metaphysical theorising except in contrast to the first-order explanatory principle. If, when doing metaphysics, I do not need to explain how the elements in my theories play their roles, then nothing is made off limits. Given this, expressing such a rule might appear pointless. But its job is to serve as a normative constraint for metaphysical theorisers: if someone does not find the explanatory principle motivating in one area of metaphysics, then then they ought not to be motivated by it in every area. By expressing the rule that results from generalising the non-explanatory meta-principle, it rules out situations whereby a metaphysician might try to claim that there is no need to explain how an element plays its role in one domain while concurrently claiming there is a need to explain how an element plays its role in another. Enforcing uniformity across one's epistemic motivation is the chief goal of these metaphilosophical rules and that applies equally to first-order epistemic principles as to meta-epistemic principles.

2.2 Paper IV

In the fourth paper, it remains relatively easy to explicitly present a motivating relationship between epistemology and metaphysics. Super-Huemans invoke a methodological commitment to minimal sufficiency and use this as the sole motivation to accept their Super-Humean ontology. In the paper, I argue that, when scrutinised, the motivating connection they proposed does not hold and instead leads to a more reductive ontology. If that result is found to be unacceptable then revisions are needed to the motivating methodology.

Overall, the connections seamlessly translate over in terms defined here. First, Super-Huemans view metaphysics as exhausted by ontology and so, any proposed ontology is comparable to a metaphysical theory. I will continue to talk about ontologies so that the claims made by Super-Humeans are not perceived to be about metaphysics understood more broadly. Second, what I had previously called 'methodological commitments' are epistemic principles. Super-Humeans seem to accept three distinct principles and conjoin them under the umbrella of minimal sufficiency. To break them down, the first is a principle asking ontologies to be simple:

> *Parsimonious Principle*: ontologies must contain as few elements as possible.

The second is a principle demanding that for any given ontology, it must be possible to recover compatibility with a reinterpreted version of contemporary physics:

> *Physics Principle*: ontologies must be compatible with our best physical theories.

The third is a principle demanding and empirical basis for ontologies:

> *Empirical Adequacy Principle*: ontologies must be causally connected to our empirical observations.

Take all three together and you have the epistemological basis Super-Humeans employ to motivate their ontology. Notably, by themselves, none of these principles outlines the reason Super-Humeans accept their metaphysics. It is the combination of all three that produces the motivational connection; they seek to accept the ontology that makes the least commitment while also being compatible with physics and empirically adequate. So, they obtain a single principle that fits into the program of epistemic metaphysics:

> *Minimal Sufficiency Principle*: ontologies must contain as few elements as possible while being compatible with our best physical theories and causally connected to our empirical observations.

A combinational approach to outlining one's motivating epistemology, as is presented here, is entirely feasible but requires some nuance. First, as shown above, it is important to clearly articulate the united principle that exhibits the motivational connection. Second, any combination cannot include principles that are in conflict. For example, one cannot be coherently motivated to accept a metaphysics that aspires to be both the simplest and most convoluted. Principles can, however, weigh against each other if they pull in different directions. This is what we see in the Super-Humean case. The parsimony principle weighs against the other two by asking for the inclusion of no more elements than that required to be compatible with physics and empirically connected. The goal then is the ideal compromise between two opposing factors.

This kind of balancing must be built into the overall motivating epistemic principle. There might be other ways of presenting this relationship, but here, in the minimal sufficiency principle, I separate the two sides with: 'as possible, while'. That which comes first is taken as a matter of degrees: to be as simple as possible. That which follows is absolute:

while being compatible with physics. For the former, the demand is to obtain to the greatest extent possible while for the latter, the demand is to obtain entirely. While this way of presenting a balancing principle might appear particularly pertinent while invoking some kind of parsimonious principle, it holds more broadly. For example, one principle could ask metaphysical theories to explain as many elements as possible, while being empirically adequate.

Next, as a motivating connection between epistemic principle and metaphysical theory is established, the default position is to generalise the connection into a metaphilosophical rule:

> *Minimally Sufficient Metaphysics*: when doing metaphysics, acceptable theories (ontologies) should adhere to the minimal sufficiency principle.

This maxim is almost entirely redundant as Super-Huemans take ontology to encompass metaphysics and therefore, there are no other parts of metaphysics. All it does here is to clearly articulate the normative constraint on metaphysical theorisers. If someone accepts the minimal sufficiency principle in one area of metaphysics, then they also ought to accept the principle to apply in every area of metaphysics; if they accept the principle, then they should also accept the metaphilosophical rule.

The last matter worth discussing with respect to the fourth paper is the symbiotic relationship between metaphysics and epistemology. The central argument I make in the paper is that Super-Humeans are mistaken in thinking the minimal sufficiency principle motivates their ontology. In fact, the principle motivates a more parsimonious ontology that is still compatible with physics and empirically adequate: Ultimate-Humeanism. Yet, I leave the conclusion open as a disjunctive choice between accepting this new more

minimal ontology or revising the epistemic principle. Super-Humeans face a choice between changing their metaphysics or changing their epistemology.

While this kind of association is implicit in many other places in the thesis, it features as a core part of the argument in paper four. This is a place where metaphysics and epistemology must be done together. Epistemic metaphysics is not a process of adopting an epistemology and then accepting or rejecting a metaphysics'. Neither is it a process of adopting a metaphysics' and then invoking a suitably accommodating epistemology. Epistemic metaphysics is a process of doing both together: evaluating if you are willing to retain a particular epistemic principle when it comes with the cost of losing one (or more) of your metaphysical theories, and *vice versa*, evaluating if you are willing to retain a particular metaphysical theory when it comes with the cost of losing one (or more) of your epistemic principles. In some cases, revisions to each might be in order but in every case, each must be given consideration.

2.3 Paper II

In the previous two sections, the task of defining and connecting metaphysical theory to epistemological principle was largely already completed in the relevant chapter, and this made a translation into the terms of this paper relatively uncomplicated. For this section, I will focus on a connection that is more subtle, as identifying an epistemic principle used to motivate a metaphysical theory and then generalising this connection requires some scrutiny. While much could be said about the influence in paper II of a normative explanatory principle akin to that in the first paper, instead of repeating myself I will focus on another part of the dialectic that plays a less perceptible but still important role.

In the second paper, after presenting an argument defending realism from the internal criticism in §3.1, I question whether it is sufficient. Both my and comparable realist arguments rely on assumptions only a realist would accept; these arguments only have sway over those who are already realists. I call this 'defensive metaphysics' and note that a problem with making arguments that require accommodating assumptions is that it leads to a philosophical stalemate. Each side ends up talking past the other and in the case of realism, non-realists accuse realists of dogmatism in making defensive assumptions while realists accuse non-realists of question-begging for not allowing them to present their view in a viable form. Those who do not see a need for inquiry to progress will find any resulting principle unmotivating, and so while I could offer further arguments in favour of accepting philosophical stalemates to be problematic, they are not needed to outline the relevant epistemic-metaphysical relationship. Instead, any resulting principle is solely for those who accept investigative progress to be motivating. Such individuals uphold arguments must be able to persuade more than those that are already persuaded.

First, it is first important to address a discrepancy as here I am criticising *arguments*, and the format of an epistemic principle is framed in terms of accepting or rejecting *theories*. There are two main ways to accommodate this difference. First, it could be that epistemic principles could be reformulated as motivation for accepting or rejecting arguments. Just as first-order epistemic principles give reasons to accept or reject certain theories, they also give reasons to accept or reject certain arguments.[62] While this is a plausible route, the alternative is to build the argument into the theory and apply an epistemic principle

[62] As arguments are in of themselves motivating – they provide reasons to accept or reject theories – the resulting principle might be best understood as a meta-epistemic principle by constraining the acceptable content of a motivating principle.

to them jointly. In this case, the result would be a first-order epistemic principle that provides reason to reject theories whose tenability rests upon assumptions only accepted by those already accommodating the theory. I will explore the latter going forward (but both options seem open).

The resulting epistemic principle hinges on first acknowledging that a certain argument poses a threat to any given metaphysical theory and a response is required. When responding one cannot construct said reply by making assumptions no one else will accept. The resulting principle can then be expressed as follows:

> *Unentrenched Principle*: the plausibility of a metaphysical theory must not rely upon assumptions only accepted by those who are already advocates of that theory.

Philosophers who find themselves motivated by wanting to avoid metaphysical debates descending into impasses will be receptive to theories that are 'unentrenched'. While there might not be complete concurrence on what arguments require a response for a theory to be tenable, the principle only applies *if* a response is needed. I assert this to be the case for realism in response to Dasgupta's argument of absent value. Moreover, I argue that the realist should not defend their view by making assumptions that are only accepted by realists. The issue here is that realists face the significant challenge of absent value and realism's tenability relies on responding to Dasgupta's argument. I base my reply on something closely akin to the unentrenched principle by arguing that any response must not make assumptions only a realist would accept. In other words, the motivation for realism, the reason to accept it, should not be dependent on 'defensive assumptions'.

From here, the first-order principle can be generalised into a metaphilosophical maxim. As the unentrenched principle is based on a desire to ensure debates do not cease in a draw, call the metaphilosophical rule:

> *Progressive Metaphysics*: when doing metaphysics, acceptable theories adhere to the Unentrenched Principle.

As a default, those who accept the unentrenched principle should be sceptical of any theory that relies on defensive assumptions. For any theory that faces objection (that is such that it warrants a response), the reply ought not to be constructed based on assumptions no one other than those who already advocate for the theory will accept. Again, there should be no cherry-picking across metaphysical domains; if one accepts the unentrenched principle as a motivating force in one area of metaphysics, one also ought to accept it as a motivating force in all areas.

Both the unentrenched principle and maxim have their counterparts. As noted at the end of §3.1 in paper II, both Lewis and Sider favour 'defensive metaphysics' and would reject the unentrenched principle as outlined here. In comparable terms, they would be advocates of the meta-epistemic principle that results from rejecting the unentrenched principle:

> *Entrenched Meta-Principle*: the plausibility of a metaphysical theory can rely upon assumptions only accepted by those who are already advocates of that theory.

Metaphysicians like Lewis and Sider are unmotivated by the first-order principle; they believe that the unentrenched principle does not hold. They do not believe that a theory only plausible on its own terms is a reason to reject that theory. Again, the meta-principle

is much weaker than the first-order principle. The first-order principle demarcates a set of unacceptable theories that only stand on their own terms while the meta-principle merely contravenes this demarcation. The meta-principle constitutes a rejection of the reasons contained in the first-order principle about accepting or rejecting metaphysical theories.

Briefly, while it does not matter for the aims of this paper (to outline and corroborate epistemic metaphysics), in paper two I argue that it is pointless to fight a war with weapons only tangible to your allies. Alternatively, when one concerns themselves solely with those already hunkered down in the trenches with them, there is little need for anything else. But it is worth noting that there exists the alternative option of moving on the offensive and instead of resisting the need to reply in a non-defensive manner, one could work to show how those assumptions are in fact not *only* made by those who already accept the theory – which is the task I embark on in the remainder of paper two.

Regardless, refusing to sally forth is still a viable approach and the project of epistemic metaphysics as outlined here allows for the possibility of such moves. Again, the entrenched meta-principle can be generalised into a metaphilosophical maxim. The default result for an individual who takes themselves to be unbound by the unentrenched principle in any given area of metaphysics is that they ought to find the principle unmotivating in every area of metaphysics. Call this non-progressive metaphilosophical maxim:

> *Stale Metaphysics*: when doing metaphysics, acceptable theories adhere to the Entrenched Principle.

As with the previous case of generalising the non-explanatory meta-principle (§2.1), this way of doing metaphysics renders every proposal eligible and is no constraint on metaphysical theorising. If, when arguing in defence of a given metaphysical theory, I may rest my arguments upon assumptions only accepted by those who are already advocates of that theory, then nothing is ruled off limits. The purpose of expressing this maxim is to make overt the normative constraint that applies to metaphysical theorisers: if someone does not find the unentrenched principle motivating in one area of metaphysics, then they ought not to be motivated by it in every area. This excludes any situation whereby a metaphysician might try to claim that one theory must be defended without friendly assumptions while another theory's defence freely permitted claims only accepted by those who already advocate for that theory. This process results in uniformity across one's epistemic motivation.

Finally, for the discussion in this section, it is worth briefly clarifying what counts as an 'epistemic' principle. In other sections, notions like explanation, empirical adequacy, and compatibility with physics, neatly fit into what is commonly understood as epistemic. That does not quite hold for the discussion around the entrenched principle and progressive metaphysics. The principles and maxims here are formulated by appealing to assumptions, theory plausibility, and theory advocates. These are not explicitly epistemic notions. However, their role is to ensure metaphysical debates can progress and our understanding can advance; these notions are certainly epistemic. The central point is to recognise what factor is the motivating one. In this case, the epistemic reason to accept metaphysical theories concerns the desire to avoid philosophical stalemates. Understood sufficiently broadly (as intended), the principles outlined here fit into the criteria outlined for being a motivating epistemic principle.

2.4 Paper V

In this section, I aim to explore how the kinds of considerations exhibited by epistemic metaphysics apply to itself. It is, after all, a work of metaphilosophy about philosophy and this leads to reflexive questions. While my aim here is to simply outline a way of doing metaphysics – defining its structure, presenting practical examples, etc. – one could make the further claim that metaphysics ought to be done this way. That metaphysics done without epistemology is unacceptable. In turn, such a proposal exhibits a motivating connection. It provides a reason to accept metaphysics accompanied by a suitably motivating epistemology and reject that which is not.

However, as epistemic metaphysics focuses on broader connections between metaphysical theories and epistemic principles and not any specific first-order theories or principles, some adaption is required. Notably, the proposal for epistemic metaphysics is a generalisation of the many connections exhibited between the other first-order epistemic principles and metaphysical theories explored in this thesis. So, it starts with a metaphilosophical rule stating:

> *Epistemic Metaphysics*: when doing metaphysics, metaphysical theories must be proposed alongside a motivating epistemological principle.

On accepting epistemic metaphysics, no metaphysical theory should be accepted without adequate endorsement from a respective epistemic principle. While the other metaphilosophical maxims discussed in this paper are framed in terms of respecting a first-order principle, the presentation here is slightly adapted, but it still expresses a similar demand: when doing metaphysics, theories must adhere to *an* epistemic principle. The difference is that while in other cases the demand is a specific motivating connection

between an epistemic principle and a metaphysical theory, this one only asks for some connection of this sort to exist. Further, while I myself find epistemic metaphysics appealing, the claim here remains conditional. This is not an argument for adopting the epistemic principle; that would require appealing to considerations about the benefits of epistemically informed theories, or the futility of metaphysical theories entirely detached from our abilities to understand them. Instead, this is merely an exploration of the consequences for any who do find epistemic metaphysics motivating.

Additionally, there might be some ambiguity in the proposal concerning what it is for a metaphysical theory to be 'proposed alongside'. One plain way of understanding this is simply a request to express the epistemic principle that provides the reason to accept the theory. This is generally what I have in mind. The key feature is the need to be *explicit* about the epistemic reasons we have for (or for not) accepting any given theory. However, one could map out a stronger approach that requires a detailed accounting of the principle and how it provides the motivation for accepting the metaphysical theory.[63] In practical terms, for any manuscript containing a metaphysical theory, if there is not a section dedicated to extensively accounting for how it is motivated by some relevant epistemic principle, then it is unacceptable. Alternatively, one could do the opposite and advance a weaker approach that does not even ask for the motivating connection to be explicit. Instead, such a connection is only required to be implicit, or even to be possible. Then in practice, for any manuscript containing a metaphysical theory, there only needs to exist the capability of connecting it to a motivating epistemic principle. Perhaps arguments could be made for moving in either direction, but I fall in the middle. Epistemic

[63] Something akin to the explanatory requirement outlined in Paper I: one must provide an explanation of the mechanisms by which this holds.

metaphysics asks that you state what you think motivates your theory. Justify it further if you want but that can be a separate discussion (either contested by others or reinforced by oneself).

Finally, the process of fully outlining epistemic metaphysics behoves us to again consider the result for those entirely unconvinced by the need to conform to epistemic metaphysics. They would be drawn to accept the converse metaphilosophical maxim:

> *Mystical Metaphysics*[64]: when doing metaphysics, metaphysical theories need not be proposed alongside a motivating epistemological principle.

Once more, this is a diluted maxim that allows for metaphysics of all sorts and puts no constraint on any particular way of doing metaphysics. If, when doing metaphysics, I do not need to propose suitably motivating epistemic principles alongside my metaphysical theories, then nothing is made off limits. Instead, all this maxim does is allow for proposals to contravene the constraint expressed in epistemic metaphysics; that such a connection *must* be in place. The purpose of this is again to outline the all-purpose normative constraints that apply to metaphysical theorisers: if someone does not find epistemic metaphysics motivating in one area of metaphysics, then they ought not to find epistemic metaphysics motivating in any area of metaphysics. This establishes the impermissibility of cases whereby a metaphysician claims our metaphysics ought to be epistemically informed in one domain, while concurrently claiming in another domain that no such informed relationship is required. The resulting generality of these rules

[64] As with the use previous use of 'stale' to define a converse metaphilosophical maxim, 'mystical' is a somewhat loaded term. If desired, both could be presented with more neutral antonyms: '*Non-Progressive Metaphysics*' and '*Non-Epistemic Metaphysics*' respectively.

ensures that any given theorises has consistent epistemic motivation that applies equally to all of their metaphysical proposals.

Overall, mystical metaphysical permits metaphysics to float entirely free from epistemology. I suspect that certain *a priori* metaphysicians might find themselves receptive to this approach as, on the surface, many of their theories appear entirely disconnected from any corresponding epistemology. When faced with a choice between an epistemically informed approach to metaphysics and the invalidation of their metaphysical theories, they opt for rejecting the methodological proposal and retaining their metaphysics. This is a perfectly valid approach that in no way undermines my proposal. If anything, it again highlights the symbiotic nature of epistemic metaphysical investigation. However, I will also briefly note that if permitted the option of taking logic as evidence (and in turn using of logic to construct principles that provide reasons to accept or reject theories) epistemic principles could be formulated to fit with the structure outlined here. I leave any further development of these approaches to those inclined as such.

3.0 FINAL SUMMARY

3.1 SYSTEMIC METAPHYSICS

Having completed the bulk of this paper and scrutinized the notion of epistemic metaphysics, I will now begin concluding. First, one area I focus on throughout is the occurrence of potential motivational conflicts between given areas of metaphysical investigation. This is a particularly pressing matter for metaphysicians who make systematic proposals. For a systematic metaphysician, someone concurrently committing to more than one metaphysical theory across domains, it is important to ensure consistent

motivation. Otherwise, one risks undermining the motivation for the complete worldview. In Lewis's systematic metaphysics, for example, if he motivated his theory of laws by appeal to the explanatory principle, he could not then accept a theory of chance that violates it.[65] The epistemic principles are included in the worldview and cannot provide motivation for one part without affecting the whole proposal. Violating this renders the metaphysical liable to undermining arguments in the form of a *tu quoque* (as I explored – and attempt to resolve – for Lewis's (Lewis, 1991, p. 31) notion of a singleton in set theory and van Inwagen's (van Inwagen, 1986) objection).

While some might want to contest this request outright (I suspect the route to do so would involve weakening the motivational connection as explored in §1.2) I see it as a beneficial result. Demarcating various compatible and incompatible theories might undermine some individuals' systematic proposals. But better systems can take their place and those that do will have solid epistemic foundations. Many systematic metaphysical proposals might already appear to be built upon an appropriate set of epistemic principles (e.g., Humeanism and Empirical Principles) but I suspect that homogeneity is not as strong a constraint as it is often perceived to be. There remains a substantive task of corroborating such suspicions and with the structure outlined here such a task is possible.

In many cases, it might just be that there is a mismatch between what one thinks is motivating their view and what actually is. I make this case at the end of paper IV when I suggest Super-Humeans would find an escape route from my criticism if they are capable of identifying a plausible motivating principle that requires there to exist more than just a brain but less than intrinsic properties. In this sense, epistemic principles are

[65] Lewis describes himself as this sort of systemic philosopher in (Lewis, 1983b, p. ix).

quite flexible and for any given systematic proposal, the advocates of which are free to combine them as they please. The complaint I have made here is against conflicting motivations, against saying you find a certain principle motivating in one domain but not in another: or worse still, committing to metaphysics theories in one domain that are incompatible with an epistemic principle that is already accepted another. For any given systematic metaphysics to have robust epistemic foundations, the motivating principles ought to apply uniformly.

Ultimately, one could avoid these kinds of problems entirely by making no systematic proposals at all. Instead of being a systematic philosopher, one could be a *piecemeal metaphysician*, making isolated contributions to particular debates that are not intended to be upheld in conjunction with one another. We can still explore the compatibility between each metaphysical theory and the corresponding epistemic principles used to motivate them. The differences would then be the absence of any grand unified metaphysical theses and the avoidance of forceful concerns that stem from irregular epistemic motivation. Perhaps that trade-off will be preferable for some.

3.2 Metaphilosophy within Metaphysics

Secondly, epistemic metaphysics has been constructed from within metaphysics rather than being externally imposed upon it. There is a contrasting approach whereby one builds an epistemic principle separately from metaphysics and then uses it to constrain what is or is not an acceptable way of doing metaphysics. I have not done this. The epistemic principles presented here have been drawn from various debates in metaphysics. I have engaged in extensive exegesis in many places throughout the thesis to ensure the kinds of epistemic factors used by metaphysicians are incorporated into the wider metaphilosophical proposal of epistemic metaphysics. Building the

metaphilosophy that applies to metaphysics from within metaphysics is important as, after all, how could one know what epistemic factors impact metaphysical disputes if one did not first engage in metaphysical disputes where epistemic factors play central roles?

Additionally, constraints that are externally imposed on metaphysics often fail to be applicable on any piecemeal basis. In such cases, there is still work to be done in exploring what theories meet the standards of the all-purpose constraint. For example, anyone can suggest that our theories must be 'explainable', but that is of no practical help if it is not clear what specific criteria our theories must meet to be explainable, or if it is not clear what actual metaphysical proposals meet that standard. In this sense, I follow in the footsteps of Bennett who asserts that 'we need to give substantive consideration to specific disputes in order to decide whether or not they are one of the problematic ones. We need to do metaphysics in order to do metametaphysics' (Bennett, 2009, p. 43). Both Bennett and I do the extra work of showing how our metaphilosophical proposals constrain actual metaphysical proposals: Bennett's metaphilosophy of epistemicism, the proposal for metaphysical disputes to be dismissed when we are unable to justify either side in the relevant dispute (Bennett, 2016, pp. 22, fn. 1), is followed by specific discussions about how the constraint applies to metaphysics disputes in constitution and composition (Bennett, 2009, pp. 44-45). The goal here is to remove any overshadowing obscurity cast when one only speaks in generalities.

Relatedly, epistemic metaphysics is far more accommodating than some other epistemically motivated metaphilosophical accounts. For example, the *neo-positivistic metaphysics* of Ladyman & Ross (Ladyman & Ross, 2007) rules out large swathes of metaphysics not done squarely within the purview of science. Or as neatly summed up by

Bryant: ‘Ladyman and Ross rail against what they call neo-scholastic metaphysics — metaphysics that floats entirely free of science’ (Bryant, 2020, p. 1867). By contrast, I do not rail against any way of doing metaphysics and have consistently left space for alternative approaches. For any of those explored throughout the thesis, I have noted the viability of taking a different approach: that of being a primitive metaphysician instead of an explanatory one, being a stale metaphysician instead of a progressive one, or even being a mystical metaphysician instead of an epistemic one. I rule none of these out, instead simply noting these alternatives.

Moreover, even the hard-line positivist metaphysics of Ladyman and Ross can be incorporated into epistemic metaphysics, as theirs is a project that can be conceived of as imposing epistemic constraints on metaphysics. The result would be a *scientism principle* and would motivate a particular class of scientifically informed metaphysical theories. The same is true for their opposites in hard-line *a priori* metaphysicians; those who believe science can only be done after one has fixed their metaphysics (Bealer & Strawson, 1992). The result would be an *a priori principle* and it would motivate a particular class of metaphysical theories. Overall, if successful in completing the substantive task of clearly articulating epistemic principles and properly connecting them to suitable metaphysical theories, both approaches would constitute epistemically constrained accounts of metaphysics.

3.3 Epistemic Aftermath

Leading on from the inclusive nature of epistemic metaphysics, one might begin to question the impact of adopting it. If everyone can, present principles and connect them to theories, no matter how eccentric their epistemology, why bother with this process at all? First, until one has done such a task, it is altogether unclear whether one can do it. It

might be that there is just no way of presenting a particular motivating principle or of adequately connecting it to a metaphysical theory. Second, and most importantly, there is still much work to be done in arguing for or against accepting any given epistemic principle and the metaphysical theories it motivates. Just because this is not the task embarked upon here does not mean it is a task not in need of carrying out. In fact, it is arguably the most important consequence of this thesis.

Arguments should be offered in favour or against any given epistemic principle. It is a central question to ask why, for example, should we explain, or why should theories be empirically adequate, or why should they be simple, and so on. Successfully answering any of these questions will lead to tangible results for their respectively motivated metaphysics. This is where the bulk of the novel work should be conducted by those who accept epistemic metaphysics. The project clearly connects metaphysics to epistemology in a way that allows for arguments in epistemology to seamlessly influence metaphysics and *vice versa*. Further to that, there is also the expansive undertaking of defining new principles and theories and of creating various compatible systems that include some principles and metaphysics but exclude others. Ultimately, epistemic metaphysics will allow for the development of novel theoretical outcomes by shifting our investigative focus toward the epistemic principles used to motivate metaphysics.

REFERENCES

Allori, V., 2013. On the Metaphysics of Quantum Mechanics. In: L. Soazig, ed. *Precis de la Philosophie de la Physique.* Paris: Vuibert.

Andrews, S., 2023. Realism and the Value of Explanation. *The Philosophical Quarterly,* 73(4), pp. 1305-1314.

Armstrong, D., 1978. *Universals and Scientific Realism.* Cambridge: Cambridge University Press.

Armstrong, D., 1980. Identity Through Time. In: P. van Inwagen, ed. *Time and Cause: Essays Presented to Richard Taylor.* Dordrecht: Reidel, pp. 67-78.

Armstrong, D., 1983. *What is a Law of Nature?.* Cambridge: Cambridge University Press.

Armstrong, D., 1988. Reply to Van Fraassen. *Australasian Journal of Philosophy,* 66(2), pp. 224-229.

Armstrong, D., 1989. *Universals: An Opinionated Introduction.* s.l.:Westview Press.

Audi, P., 2020. Bennett on building. *Inquiry,* 63(7), pp. 677-692.

Barnes, E., 2018. Symmetric Dependence. In: R. L. Bliss & G. Priest, eds. *Reality and Its Structure.* Oxford: Oxford University Press, pp. 50-69.

Baron-Schmitt, N., 2021. Contingent grounding. *Synthese,* 199(1-2), pp. 4561-4580.

Bealer, G. & Strawson, P. F., 1992. The Incoherence of Empiricism. *Proceedings of the Aristotelian Society Supplementary Volumes,* Volume 66, pp. 99-143.

Beebee, H. & Mele, A., 2002. Humean Compatibilism. *Mind,* Volume 111, pp. 201-223.

Benacerraf, P., 1973. Mathematical Truth. *The Journal of Philosophy*, 70(19), pp. 661-679.

Bennett, K., 2009. Composition, colocation, and metaontology. In: D. Chalmers, D. Manley & R. Wasserman, eds. *Metametaphysics: New Essays on the Foundations of Ontology.* Oxford: Oxford University Press, pp. 38-76.

Bennett, K., 2011. By Our Bootstraps. *Philosophical Perspectives,* 25(Metaphysics), pp. 27-41.

Bennett, K., 2016. There is no special problem with metaphysics. *Philosophical Studies,* Volume 171, pp. 21-37.

Bennett, K., 2017. *Making Things Up.* Oxford: Oxford University Press.

Bennett, K., 2019. Replies to Cameron, Dasgupta, and Wilson. *Philosophy and Phenomenological Research,* 98(2), pp. 507-521.

Bennett, K., 2020. Reply to Audi, Bliss, Rosen, Schaffer, and Wang. *Inquiry,* 63(7), pp. 758-794.

Bernstein, S., 2016. Grounding Is Not Causation. *Philosophical Perspectives,* 30(Metaphysics), pp. 21-38.

Bhogal, H. & Perry, Z., 2017. What the Humean Should Say About Entanglement. *Noûs,* 51(1), pp. 74-94.

Bird, A., 2007. *Nature's Metaphysics: Laws and Properties.* Oxford: Oxford University Press .

BonJour, L., 2000. The Elements of Choerentism. In: F. Dretske & S. Bernecker, eds. *Knowledge: Readings in Contemporary Epistemology..* Oxford: Oxford University Press, pp. 128-148.

Boolos, G., 1971. The iterative conception of set. *Journal of Philosophy,* 68(8), pp. 215-231.

Bourne, C., 2006. *A future for presentism.* Oxford: Clarendon Press.

Bryant, A., 2020. Keep the chickens cooped: the epistemic inadequacy of free range metaphysics. *Synthese ,* Volume 197, pp. 1867-1887.

Button, T., 2013. *The Limits of Realism.* Oxford: Oxford University Press.

Camus, A., 1955. *The Myth of Sisyphus.* New York: Vintage International .

Caulton, A., 2018. A persistent particle ontology for Quantum field theory. *Metascience,* Volume 27, pp. 439-441.

Coates, A., forthcoming. The primitivist response to the inference problem. *Dialectica.*

Cohen, J. & Callender, C., 2009. A better best system account of lawhood. *Philosophical Studies,* Volume 145, pp. 1-34.

Correia, F., 2010. Grounding and truth-functions. *Logique et Analyse,* 53(Special Issue "Explanation in Metaphysics" 211), pp. 251-279.

Darby, G., 2018. A minimalist Humeanism?. *Metascience,* Volume 27, pp. 433-437.

Dasgupta, S., 2014. The Possibility of Physicalism. *Journal of Philosophy,* 111(9-10), pp. 557-592.

Dasgupta, S., 2017. Constitutive Explanation. *Philosophical Issues, 27, Metaphysics, 2017 ,* Volume 27, pp. 74-97.

Dasgupta, S., 2017. Normative Non-Naturalism and the Problem of Authority. *Proceedings of the Aristotelian Society,* 117(3), pp. 297-319.

Dasgupta, S., 2018. Realism and the Absence of Value. *Philosophical Review,* 127(3), pp. 279-322.

Dasgupta, S., 2020. *Undoing the Truth Fetish: The Normative Path to Pragmatism.* [Online]
Available at:
http://shamik.net/papers/dasgupta%20undoing%20the%20truth%20fetish.pdf
[Accessed June 2020].

Deckert, D.-A., Esfeld, M. & Andrea, O., 2019. A persistent particle ontology for QFT in terms of the Dirac sea. *British Journal for the Philosophy of Science* , Volume 70, p. 747–770.

deRosset, L., 2013. Grounding Explanations. *Philosophers' Imprint ,* Volume 13, pp. 1-26.

Devitt, M., 1983. Realism and the Renegade Putnam: A Critical Study of Meaning and the Moral Sciences. *Noûs,* 17(2), pp. 291-301.

Dorr, C., 2019. *Natural Properties.* [Online]
Available at: https://plato.stanford.edu/archives/fall2019/entries/natural-properties/

Dorr, C. & Hawthorne, J., 2013. Naturalness. In: K. Bennett & D. Zimmerman, eds. *Oxford Studies in Metaphysics: Volume 8.* Oxford: Oxford University Press, p. 3–77.

Dreier, J., 2015. Can Reasons Fundamentalism Answer the Normative Question?. In: G. Björnsson, et al. eds. *Motivational Internalism.* Oxford: Oxford Scholarship Online.

Esfeld, M., 2020a. Super-Humeanism: The Canberra Plan for Physics. In: D. Glick, G. Darby & A. Marmodoro, eds. *The Foundation of Reality: Fundamentality, Space, and Time.* Oxford: Oxford University Press, pp. 125-138.

Esfeld, M., 2020b. A proposal for a minimalist ontology. *Synthese,* Volume 197, p. 1889–1905.

Esfeld, M. & Deckert, D.-A., 2017. *A Minimalist Ontology of the Natural World.* New York: Routledge.

Field, H., 1988. Realism, Mathematics and Modality. *Philosophical Topics,* 16(1 Metaphysics), pp. 57-107.

Field, H., 1989. *Realism, Mathematics, and Modality.* Oxford: Blackwell.

Fine, K., 2001. The Question of Realism. *Philosophers' Imprint,* 1(1), pp. 1-30.

Fine, K., 2010. Some Puzzles of Ground. *Notre Dame Journal of Formal Logic,* 51(1), pp. 97-118.

Fine, K., 2012. Guide to Ground. In: F. Correia & B. Schnieder, eds. *Metaphysical Grounding.* Cambridge: Cambridge University Press, pp. 37-80.

Forrest, P., 1991. Aesthetic Understanding. *Philosophy and Phenomenological Research,* 51(3), pp. 525-540.

Friend, T., 2022. The Humean pragmatic turn and the case for revisionary best systems accounts. *European Journal for Philosophy of Science,* 12(11), pp. 1-26.

Frugé, C., forthcoming. Janus-Faced Grounding. *Ergo: An Open Access Journal of Philosophy.*

Giannotti, J., 2022. Fundamentality and Minimalist Grounding Laws. *Philosophical Studies,* Volume 179, pp. 2993-3017.

Goodman, N., 1955. *Fact, Fiction, and Forecast.* Cambridge: Harvard University Press.

Hall, N., 2004. Two Mistakes About Credence and Chance. *Australasian Journal of Philosophy,* 81(1), pp. 93-111.

Hall, N., 2015. Humean Reductionism About Laws of Nature. In: *A Companion to David Lewis.* s.l.:Wiley, pp. 262-277.

Hawthorne, J., 2007. Craziness and Metasemantics. *Philosophical Review,* 116(3), pp. 427-440 .

Hicks, M. T., 2017. Making Fit Fit. *Philosophy of Science,* Volume 84, pp. 931-943.

Hicks, M. T., 2021. Breaking the explanatory circle. *Philosophical Studies,* Volume 178, pp. 533-557.

Hicks, M. T. & Wilson, A., 2021. How chance explains. *Noûs,* pp. 1-26.

Hildebrand, T., 2013. Tooley's account of the necessary connection between law and regularity. *Philosophical Studies,* 166(1), pp. 33-43.

Hildebrand, T., 2020. Non-Humean theories of natural necessity. *Philosophy Compass,* 15(5), pp. 1-15.

Hirsch, E., 1993. *Dividing Reality.* Oxford: Oxford University Press.

Hodes, H., 1984. Logicism and the Ontological Commitments of Arithmetic. *The Journal of Philosophy,* 81(3), pp. 123-149.

Hofweber, T., 2016. How metaphysics is special: comments on Bennett. *Philosophical Studies,* Volume 173, pp. 39-48.

Hüttemann, A., 2014. Scientific Practice and Necessary Connections. *Theoria. An International Journal for Theory, History and Foundations of Science,* 29(1), pp. 29-39.

Jaag, S. & Loew, C., 2020. Making best systems best for us. *Synthese,* Volume 197, pp. 2525-2550.

Jackson, F., 1994. Armchair metaphysics. In: J. O'Leary-Hawthorne & M. Michael, eds. *Philosophy in Mind.* s.l.:Kluwer Academic Publishers, pp. 23-42.

Jenkins, C. S., 2011. Is Metaphysical Dependence Irreflexive?. *The Monist,* 94(2), pp. 267-276.

Khlentzos, D., 2016. *Challenges to Metaphysical Realism.* [Online]
Available at: https://plato.stanford.edu/archives/win2016/entries/realism-sem-challenge/

Kim, J., 2005. *Physicalism, or something near enough.* Princeton: Princeton University Press.

Korsgaard, C. M., 2009. The Normativity of Instrumental Reason. In: *The Constitution of Agency: Essays on Practical Reason and Moral Psychology.* Oxford: Oxford Scholarship Online, pp. 27-68.

Kovacs, D. M., 2020. Four Questions of Iterated Grounding. *Philosophy and Phenomenological Research,* 101(2), pp. 341-364.

Ladyman, J. & Ross, D., 2007. *Every Thing Must Go: Metaphysics Naturalized.* Oxford: Oxford University Press.

Lange, M., 2023. A new circularity in explanations by Humean laws of nature. *Philosophical Studies,* Volume 180, pp. 1001-1016.

Lazarovici, D., 2018. Super-Humeanism: A starving ontology. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics,* Volume 64, pp. 79-86.

Lewis, D., 1981. Causal decision theory. *Australasian Journal of Philosophy,* 59(1), pp. 5-30.

Lewis, D., 1983a. New work for a theory of universals. *Australasian Journal of Philosophy,* 61(4), pp. 343-377.

Lewis, D., 1983b. *Philosophical Papers Volume I.* New York: Oxford University Press.

Lewis, D., 1984. Putnam's paradox. *Australasian Journal of Philosophy,* 62(3), pp. 221-236.

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

Lewis, D., 1987. *Philosophical Papers Volume II.* Oxford: Oxford University Press.

Lewis, D., 1991. *Parts of Classes.* Oxford: Blackwell.

Lewis, D., 1993. Mathematics is megethology. *Philosophia Mathematica,* 1(1), pp. 3-23.

Lewis, D., 1994a. Humean Supervenience Debugged. *Mind,* 103(412), pp. 473-490.

Lewis, D., 1994b. Reduction of mind. In: S. Guttenplan, ed. *Companion to the Philosophy of Mind.* s.l.:Blackwell, pp. 412-431.

Lewis, D., 2009. Ramseyan humility. In: D. Braddon-Mitchell & R. Nola , eds. *Conceptual Analysis and Philosophical Naturalism.* Cambridge: MIT Press, pp. 203-222.

Linnebo, Ø., 2006. Epistemological challenges to mathematical platonism. *Philosophical Studies,* Volume 129, pp. 545-574.

Linnebo, Ø., 2013. The potential hierarchy of sets. *Review of Symbolic Logic,* 6(2), pp. 205-228.

List, C., 2019. Levels: Descriptive, Explanatory, and Ontological. *Nous,* 53(4), p. 852–883.

Litland, J. E., 2017. Grounding Ground. In: K. Bennett & D. Zimmerman, eds. *Oxford Studies in Metaphysics: Volume 10.* Oxford: Oxford University Press, pp. 279-316.

Litland, J. E., 2018. Could the grounds's grounding the grounded ground the grounded?. *Analysis,* 78(1), pp. 56-65.

Litland, J. E., 2020. Meta-ground. In: M. Raven, ed. *The Routledge Handbook of Metaphysical Grounding.* New York: Routledge, pp. 133-147.

Loewer, B., 1996. Humean Supervenience. *Philosophical Topics,* 2(1), pp. 101-127.

Loewer, B., 2007. Laws and Natural Properties. *Philosophical Topics,* 35(1/2), pp. 313-328.

Loewer, B., 2012. Two accounts of laws and time. *Philosophical Studies,* Volume 160, pp. 115-137.

Lovett, A., 2019. A Simple Proof of Grounding Internality. *Thought: A Journal of Philosophy,* 8(3), pp. 154-166.

Marmodoro, A., 2018. Atomism, holism and structuralism: costs and benefits of a minimalist ontology of the world. *Metascience,* Volume 27, pp. 421-425.

Marshall, D., 2015. Humean laws and explanation. *Philosophical Studies,* Volume 172, p. 3145–3165.

Matarese, V., 2020. A challenge for Super-Humeanism: the problem of immanent comparisons. *Synthese,* Volume 197, p. 4001–4020.

Matarese, V., 2021. Super-Humeanism and physics: A merry relationship?. *Synthese,* Volume 199, pp. 791-813.

Maudlin, T., 2007. *The Metaphysics Within Physics.* Oxford: Oxford University Press.

Meacham, C. & Eddon, M., 2015. No Work For a Theory of Universals. In: J. Schaffer & B. Loewer, eds. *A Companion to David Lewis.* s.l.:Wiley-Blackwell, pp. 116-137.

Oldofredi, A. & Ottinger, H. C., 2021. The dissipative approach to quantum field theory: conceptual foundations and ontological implications. *European Journal for Philosophy of Science,* 11(18).

Pagès, J., 2002. The Dretske–Tooley–Armstrong theory natural laws and the inference problem. *International Studies in the Philosophy of Science,* 16(3), pp. 227-243.

Peacocke, C., 1999. *Being Known.* Oxford: Oxford University Press.

Pérez Carballo, A., 2020. Conceptual Evaluation. In: A. Burgess, H. Cappelen & D. Plunkett, eds. *Conceptual Engineering and Conceptual Ethics.* Oxford: Oxford University Press, pp. 305-332.

Potter, M., 2004. *Set theory and its philosophy: A critical introduction.* Oxford: Oxford University Press.

Pryor, J., 2000. The Skeptic and the Dogmatist. *NOÛS,* 34(4), p. 517–549.

Putnam, H., 1985. Models and Reality. In: *Philosophical Papers: Volume 3, Realism and Reason.* Cambridge: Cambridge University Press, pp. 1-25.

Putnam, H., 1990. *Realism with a Human Face.* Cambridge: Harvard University Press.

Rabin, G. O. & Rabern , B., 2016. Well Founding Grounding Grounding. *Journal of Philosophical Logic,* Volume 45, pp. 349-379.

Raven, M., 2009. *Ontology, from a Fundamentalist Point of View. (PhD thesis, New York University).*. s.l.:PhD thesis, New York University.

Raven, M., 2015. Ground. *Philosophy Compass,* 10(5), pp. 322-333.

Rodriguez-Pereyra, G., 2015. Grounding is not a strict order. *Journal of the American Philosophical Association,* 1(3), pp. 517-534.

Rosen, G., 2010. Metaphysical Dependence: Grounding and Reduction. In: B. Hale & A. Hoffmann, eds. *Modality: Metaphysics, Logic, and Epistemology.* Oxford: Oxford University Press, pp. 109-136.

Rosen, G., 2017. Ground by law. *Philosophical Issues,* Issue 27, pp. 279-301.

Russell, B., 1927. *The Analysis of Matter.* London: Routledge.

Schaffer, J., 2003. Principled Chances. *The British Journal for the Philosophy of Science,* 54(1), pp. 27-41.

Schaffer, J., 2004. Two conceptions of sparse properties. *Pacific Philosophical Quarterly,* 85(1), p. 92–102.

Schaffer, J., 2007. Deterministic Chance?. *British Journal for the Philosophy of Science,* 58(2), pp. 113-40.

Schaffer, J., 2009. On what grounds what. In: D. Manley, D. J. Chalmers & R. Wasserman, eds. *Metametaphysics: New Essays on the Foundations of Ontology.* Oxford: Oxford University Press, pp. 347-383.

Schaffer, J., 2010. Monism: the priority of the whole. *Philosophical Review,* 119(1), pp. 31-76.

Schaffer, J., 2012. Grounding, transitivity, and contrastivity. In: F. Correia & B. Schnieder, eds. *Metaphysical Grounding: Understanding the Structure of Reality.* Cambridge: Cambridge University Press, pp. 122-138.

Schaffer, J., 2016. Grounding in the image of causation. *Philosophical Studies,* Volume 173, pp. 49-100.

Schaffer, J., 2016. It is the Business of Laws to Govern. *Dialectica,* 70(4), pp. 577-588.

Schaffer, J., 2017a. The Ground between the Gaps. *Philosophers' Imprint,* Volume 17, p. 1–26.

Schaffer, J., 2017. Laws for Metaphysical Explanation. *Philosophical Issues,* 27(Metaphysics), pp. 302-321.

Sellars, W., 1962. Philosophy and the Scientific Image of Man. In: R. Colodny, ed. *Frontiers of Science and Philosophy.* Pittsburgh: University of Pittsburgh Press, pp. 35-78.

Sider, T., 1992. Tooley's Solution to the Inference Problem. *Philosophical Studies,* 67(3), pp. 261-275.

Sider, T., 2011. *Writing the Book of the World.* Oxford: Oxford University Press.

Sider, T., 2020. Ground grounded. *Philosophical Studies,* Volume 177, pp. 747-767.

Sider, T., 2022. Dasgupta's Detonation. *Philosophical Perspectives,* 36(1), pp. 292-304.

Simpson, W. M. R., 2021. What's the Matter with Super-Humeanism?. *British Journal for the Philosophy of Science,* 72(3), pp. 893-911.

Sober, E., 1988. *Reconstructing the Past: Parsimony, Evolution and Inference.* Cambridge, MA: MIT Press.

Studd, J., 2013. The Iterative Conception of Set. *Journal of Philosophical Logic,* Volume 42, pp. 697-725.

Taylor, B., 1994. Natural Properties in Metaphysics. *Mind,* Volume 102, pp. 81-100.

Thomasson, A., 2020. A Pragmatic Method for Normative Conceptual Work. In: A. Burgess, H. Cappelen & D. Plunkett, eds. *Conceptual Engineering and Conceptual Ethics.* Oxford: Oxford University Press, pp. 435-458.

Thompson, N., 2019. Is Building Built?. *Analysis,* 79(2), pp. 315-327.

Tooley, M., 1988. *Causation: A Realist Approach.* s.l.:Clarendon Press.

Unger, P., 1979. Why There Are No People. *Midwest Studies in Philosophy,* Volume 4, pp. 177-222.

van Fraassen, B., 1989. *Laws and Symmetry.* Oxford: Clarendon Press.

van Inwagen, P., 1986. Two Concepts of Possible Worlds. *Midwest Studies in Philosophy,* 11(1), pp. 185-213.

Wallace, D. & Timpson, C. G., 2010. Quantum Mechanics on Spacetime I: Spacetime State Realism. *British Journal for the Philosophy of Science,* 61(4), pp. 697-727.

Wang, J., 2020. Building and modal recombination. *Inquiry,* 63(7), p. 745–757.

Werne, J., 2020. A grounding-based measure of relative fundamentality. *Synthese,* pp. 1-17.

Wilsch, T., 2015. The Nomological Account of Ground. *Philosophical Studies,* Volume 172, p. 3293–3312.

Wilson, A., 2018a. Metaphysical Causation. *Noûs,* 52(4), pp. 723-751.

Wilson, A., 2018b. Super-Humeanism: insufficiently naturalistic and insufficiently explanatory. *Metascience,* Volume 27, p. 427–431.

Wilson, A., 2018c. Review of Esfeld and Deckert (2018). *Dialectica,* 74(3), pp. 583-589.

Wilson, A., 2019. Making Things Up, by Karen Bennett. *Mind,* 128(510), pp. 588-600.

Wilson, J., 2010. What is Hume's Dictum, and why believe it?. *Philosophy and Phenomenological Research,* 80(3), pp. 595-637.

Wilson, J., 2014. No Work for a Theory of Grounding. *Inquiry,* 57(5-6), pp. 535-579.

Woodward, J., 2003. *Making Things Happen: A Theory of Causal Explanation.* Oxford: Oxford University Press.