

Purity Against Non-naturalism

By
Mohammadjavad Hajialikhani

Submitted to Central European University - Private University
Department of Philosophy

In partial fulfilment of the requirements for the degree of Master of Philosophy

Supervisor: Asya Passinsky

Vienna, Austria
2025

Author's declaration

I, the undersigned, **Mohammadjavad Hajialikhani** candidate for the MA degree in philosophy declare herewith that the present thesis titled Purity Against Non-naturalism is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography.

I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes on any person's or institution's copyright.

I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, 02 June 2025

Mohammadjavad Hajialikhani

Copyright Notice

Copyright ©Mohammadjavad Hajialikhani, 2025. Purity Against Non-naturalism - This work is licensed under [Creative Commons Attribution-NonCommercial-NoDerivatives \(CC BY-NC-ND\) 4.0 International](https://creativecommons.org/licenses/by-nc-nd/4.0/) license.



For bibliographic and reference purposes this thesis/dissertation should be referred to as:
Hajialikhani, Mohammadjavad. 2025. Purity Against Non-naturalism. MA thesis. Department of Philosophy, Central European University, Vienna.

¹ Icon by [Font Awesome](https://fontawesome.com/).

Copyright Notice

some parts of this thesis are revised versions of some parts of my previous term paper, *purity against property dualism*, submitted for the course metaphysics of mind II. The departmental policy allows reusing some of our work up to 4000 words.

Acknowledgements

I would like to thank my supervisor, **Asya**, whose guidance and support enabled me to write—and more importantly, to think through—most of what is presented in this thesis.

I would like to thank many friends, in Vienna and elsewhere, with whom I shared many valuable moments during these years. They made Vienna feel more like home. Their perspectives, conversations, and presence made it possible to think clearly and keep going.

Thank you, **Surin**, for your cozy presence and your constant reminders of what is truly worth living for.

Finally, I would like to thank my family who supported me, in various ways, in every step of my studies.

Table of Contents

Introduction	1
Chapter 1: Grounding, Fundamentality, and Non-naturalism	3
1.1 Fundamentality and Grounding.....	3
1.2 Non-naturalism	6
1.3 Motivations	10
Chapter 2: The Main Argument	13
2.1 The Short Version.....	13
2.2 Fundamental Purity.....	15
2.3 Criteria for the Fundamentality of Objects and Properties.....	17
2.4 Normative Impurity.....	22
2.4 .1 Atomic Normative Facts.....	23
2.4.2 Normative Principles.....	24
Chapter 3: The Purification Strategy	29
3.1 Purifying Atomic Facts	29
3.2 Purifying Principles.....	33
3.3 Conclusion.....	38
Chapter 4: Objections	39
4.1 Overgeneralization	39
4.2 Rejecting Purity.....	41
4.3 Moral Particularism	43
4.4 Other Normative Properties.....	45
Conclusion.....	47

Introduction

Normative facts are those that involve a normative property (such as wrongness, badness, or reasonableness); for example, killing others is wrong; every action that fails to maximize utility is wrong. Non-naturalists claim that there are irreducible, *sui generis* normative properties.

According to one important formulation of non-naturalism, non-naturalism is the thesis that at least some normative facts are not wholly grounded in the physical facts.

In this thesis, I argue against that non-naturalism, understood as the ground-theoretic thesis expressed above, is false, which means normative facts are wholly grounded in non-normative ones, such as physical, mental, or social facts. If God wants to create the world, she only needs to put the non-normative facts in place, and the normative facts will arise automatically. The argument in a nutshell is this: fundamental normative facts violate the principle of purity, according to which fundamental facts contain only fundamental constituents (objects and properties). This is because normative facts always contain a derivative constituent, such as actions, or non-normative properties such as failing to maximize utility. Thus, there are no fundamental normative facts, and the grounding formulation of non-naturalism is false.

The plan of the thesis is as follows. In chapter 1, I start by explaining fundamentality and grounding, which is the metaphysical framework in which the argument is formulated. I will present a version of normative non-naturalism and motivate a grounding formulation of it. In Chapter 2, I then present the outline of the argument: normative facts are impure—which means they contain non-fundamental constituents, but according to the principle of purity, fundamental facts are pure—which means they only contain fundamental constituents. I provide some reasons

for purity. To defend the impurity of normative facts, I provide some criteria for assessing the fundamentality of objects and properties. I discuss two main kinds of normative facts, atomic normative facts and normative principles, and show that they contain derivative constituents. In chapter 3, I discuss the purification strategy, which non-naturalists might employ to make non-naturalism consistent with purity; I argue that various forms of this strategy are theoretically costly. In chapter 4, I discuss some possible objections to the argument, including the objection that this argument overgeneralizes over other domains. Finally, in conclusion, I discuss two ways that non-naturalists might respond to the argument and summarize my argument.

Chapter 1: Grounding, Fundamentality, and Non-naturalism

In this chapter, I will first introduce the grounding framework with some examples. Then I explain the central commitments of normative non-naturalism and why it makes sense to formulate it within the grounding framework—as many non-naturalists have already done. This results in the grounding formulation of non-naturalism, the thesis that I am going to argue against in the next chapter. I then provide some motivations for arguing against the grounding formulation of non-naturalism.

1.1 Fundamentality and Grounding

There has been much discussion about grounding in contemporary metaphysics.² I first try to illustrate this notion with an example. Consider this table. It has some ordinary properties, such as *being brown*. The table is constituted by some atoms that have some physical properties and are arranged in a certain way. It is common to think that the physical facts about these atoms somehow determine or explain ordinary facts about the table: the table is brown *because of* the physical facts about atoms and their arrangement. Once the physical facts about the table are fixed, the ordinary facts such as the table's being brown are automatically fixed—note that the reverse is not true: if we fix the ordinary properties of the table, there are many arrangements and physical properties that might realize those ordinary properties. These observations help us understand the idea that there is some asymmetric explanatory relation between the fact that the table is brown and the fact that its atoms have such and such properties and arrangement.

² See Fine (2001), Schaffer (2009), Correia & Schneider (2012), Audi (2012), and Rosen (2010). See also Bliss & Trogdon (2021) and useful references there.

Grounding is a way of understanding this dependence relation. Grounding is a non-causal, metaphysical in virtue of (or determination, or dependence) relation between entities or facts. If p grounds q, then p metaphysically determines q, or q metaphysically depends on p. Consider some examples. If A is true, then “A or B” is grounded in A: disjunctions are true because one of their disjuncts is true. The fact that the city is crowded is grounded in many facts about the people’s presence in the streets: the city is crowded in virtue of people’s presence and movements. The fact that the set whose only member is Socrates exists is grounded in the fact that Socrates exists. Following our example from the last paragraph, the fact that this table is brown is grounded in physical facts about the microstructure of the table. In short, grounding connects the different *levels* of reality, such as social, psychological, biological, chemical, and physical levels.³

Grounding theorists attribute certain logical properties to grounding. Usually, grounding is taken to be transitive: if A grounds B and B grounds C, then A grounds C. It is irreflexive: no fact grounds itself. It is asymmetric: if A grounds B then B does not ground A (there are no grounding loops).⁴ Sometimes, grounding theorists say that grounding is well-founded, which means there are no infinite chains of grounding.⁵

Fundamentality can be defined straightforwardly in terms of grounding. Fundamental facts are those that are not grounded in any other fact, they are also called *ungrounded* facts.⁶ Derivative facts are those that are grounded in some other fact. If we assume that grounding is well-founded, then all derivative facts are ultimately grounded in the fundamental facts, so

³ See Sider (2020) and Schaffer (2012) on grounding as connecting levels of reality.

⁴ See Fine (2012), Rosen (2010, section 5), and Raven (2013).

⁵ Schaffer (2010) supports this idea. See Cameron (2008) for more discussion.

⁶ For defining fundamentality with grounding, see Schaffer (2009, p. 373).

fundamental facts metaphysically explain (or determine, or constitute) everything else. They are the bedrock upon which everything else is built. To put it in terms of a famous metaphor, if God wants to create the world, she only needs to create the fundamental facts, the rest will follow automatically.⁷

Fundamentality and grounding are not merely theoretical concepts that only grounding theorists are interested in. Often in philosophy, there is a substantive debate over whether facts of a certain kind are fundamental or not, or over *how* they are grounded in more fundamental facts. I will list some examples of debates that grounding can help to formulate. This will also help us to have a better grip on the concept of grounding.

The debate over physicalism about mind can be understood as whether some mental facts are fundamental, or all mental facts are wholly grounded in physical facts.⁸ Many physicalists agree that there are mental states and facts; their disagreement with the dualist lies in that they think mental facts are derivative of, or dependent on, physical facts while dualists think some mental facts are substantive and fundamental facts on their own. As Chalmers puts it: “But the fact that consciousness does not supervene on the physical features shows us that this physical theory is not quite a theory of everything. To bring consciousness within the scope of a fundamental theory, we need to introduce new fundamental properties and laws,” (Chalmers 1996, p. 126)

Questions of grounding normative facts have always been pressing in metaethics. Consider the famous Euthyphro Dilemma, put forward by Plato: Are bad actions bad because God said so, or is it the other way around?⁹ This “because” question might be interpreted as asking how

⁷ See Schaffer (2009: 351) and Tahko (2023, 1.4) on this metaphor.

⁸ See Schaffer (2017) and Schaffer (2023). See Bennett (2011) for a similar idea with a “building” relation instead of grounding. Montero (2001) advocates a fundamentality formulation of dualism. For criticism, see Wilson (2016).

⁹ See Bliss & Trogon (2021, 1.5), and Schaffer (2009, p. 375) for putting forward this example.

normative facts are grounded. As we will see later, the much-discussed relation between natural and normative properties can also be interpreted as a grounding relation.

Grounding can be useful in understanding debates over social construction. Social theorists often claim that social phenomena—such as law, gender, or money—are socially constructed.

However, it is not clear what the metaphysical status of socially constructed objects and facts is.

As Schaffer (2017) observes, social construction can be defined and understood in the grounding framework: to be socially constructed is to be grounded in distinctive, repeated social patterns.

(Schaffer 2017, p. 2454) Grounding definition can make sense of the intuition that socially

constructed “is generated by, dependent upon, and explicable on the basis of social patterns”

because “The grounded is non-fundamental, and is generated by, dependent upon, and explicable

on the basis of the grounds. (Schaffer 2017, p. 2452 - 2454) There is a whole debate on how to

understand social construction via grounding.¹⁰

1.2 Non-naturalism

After seeing the basics of the grounding framework, it is time to see the basics of normative non-naturalism, the view that is the central focus of my argument. Non-naturalism in metaethics is the

thesis that there are irreducible, *sui generis* normative properties.¹¹ I use normative broadly to

refer to all moral and evaluative properties and relations, such as *must-be-doneness*, *goodness*,

being better, *wrongness*, etc. Many non-naturalists focus only on moral properties. This does not

matter for the argument in this thesis, so I use the broader term, ‘normative’.

¹⁰ See Passinsky (forthcoming) for both a good overview and a new way of understanding social construction with grounding.

¹¹ See Ridge (2019).

Non-naturalism is usually formulated in many ways, and it is safe to say that it is a label for a family of philosophical theses. As observed by McPherson and Plunkett (2022, p. 676), some of these theories have lighter metaphysical commitments,¹² but some of them are more robust and metaphysically loaded. In this thesis, I focus on the second kind of non-naturalism, exemplified by Enoch (2011).¹³ Enoch uses the term "robust realism" to refer to this view. I will briefly mention three central commitments of non-naturalists; they are usually shared by all non-naturalists but the way that the third one is developed here is unique to the second kind of non-naturalism.

- (I) Non-naturalists are realists about normative properties and facts. This realism has two components. First, normative sentences are factual or truth-apt: they report some state of affairs in the world, similar to sentences such as "the cat is on the mat" and contrary to expressions such as "hello". This first component of realism is sometimes captured by saying that normative judgements express beliefs, rather than some non-cognitive state. Second, at least some normative sentences are true. These components together entail that normative realists reject expressivist theories and error theories about normativity.¹⁴
- (II) They also reject response-dependent theories of normativity, according to which normative judgements are constitutively dependent on human opinion.¹⁵
- (III) Non-naturalists think that normative properties are not natural. They claim that these normative properties are not identical with natural properties and cannot be wholly

¹² See (1996), Parfit (2011), and Scanlon (2014).

¹³ Also see FitzPatrick (2008), and Bengson & Shafer-Landau & Cuneo (2024), Dunaway (2014), and Leary (2017).

¹⁴ This is usually called "minimal realism". See Enoch (2011, p. 3) and van Roojen (2015, chapter 2).

¹⁵ Enoch (2011, p. 3-4).

explained with or reduced to natural properties. Enoch puts it as “Robust Realism rejects the naturalist claim that— in a sense yet to be precisified— normative facts are nothing over and above natural ones. Normative facts are just too different from natural ones to be a subset thereof.” (Enoch 2011, p. 4) Non-naturalists thus reject naturalistic explanations of normativity in terms of biological or psychological facts. They think that normative properties are of a new kind; they are something on their own.

In a nutshell, non-naturalists are realists who think that morality is objectively out there, and not because of any natural properties or facts, but as an entirely new stuff. No story of the world is complete unless mentioning some normative facts or properties.

Many non-naturalists have formulated non-naturalism in the grounding framework. The idea is that non-naturalism is the claim that normative facts are not fully grounded in natural facts.

Enoch and Weinshtock Saadon, for instance, write that: “we [robust realists] take the basic moral (or normative) facts to be entirely fundamental, that there’s nothing in virtue of which these facts obtain.” (Enoch and Weinshtock Saadon 2022, p. 449). And Enoch (2019) writes:

Let me stipulate that the Robust Realism I’m interested in here is committed to the following: *No moral fact is fully grounded in non-normative, natural facts...* This claim is meant to capture the idea that moral facts are as real, as ontologically respectable, and the most fundamental of them are as fundamental, as any others. When the gods finished creating the natural facts, as it is sometimes said, they still weren’t done—they still had to put in place the moral ones as well.” (Enoch 2019, pp. 2-3)

Several other prominent non-naturalists have also proposed similar formulations.¹⁶ In the following, I will discuss two motivations for formulating non-naturalism with grounding.

First, it is intuitive to think of normative naturalists as engaging in the project of *explaining* normative properties and facts with natural facts—facts about mental states, human biology, conventions, or something similar (Note that usually, in this context, natural fact just means non-normative fact (maybe excluding supernatural facts about God), so it encompasses a very broad range of facts). If this explanation is metaphysical, assuming a very tight connection between metaphysical explanation and grounding, they are trying to ground the normative in the natural. As Rosen puts it “... the ethical naturalist’s key thought is that the normative facts stand to the non-normative facts in precisely this intimate relation [grounding] ... [Ethical naturalists] agree that the normative facts are not further facts, superadded to reality after the non-normative facts are fixed.” (Rosen 2016, p. 157) Non-naturalism can then be formulated as the denial of naturalism (and supernaturalism etc.), so it would be the thesis that some normative facts are not fully grounded in non-normative ones.

Secondly, non-naturalists think that normative properties are *just too different* from natural ones (or supernatural facts, or any other kinds of facts)¹⁷. But if p fully grounds q, then p and q cannot be “just too different” facts because, at least according to one way of theorizing about grounding, it is a very strict metaphysical connection: a fact is *almost* nothing over and above its grounds.¹⁸

For instance, if social facts are ultimately grounded in physical facts, though different from

¹⁶ For more formulations of non-naturalism with grounding or fundamentality along this line, see Rosen (2016, forthcoming), Dunaway (2014, p. 650), and Morton (2020). For dissent, See Leary (2021) and Bengson & Shafer-Landau & Cuneo (2024), and McPhersen and Plunkett (2022). For other examples of understanding non-naturalism with grounding, see Berker (2018) and Bader (2017).

¹⁷ Enoch (2011, part 5.1)

¹⁸ Rosen (2016, p. 9-10).

physical facts, social facts cannot be too different from the physical ones. Of course, social facts are different from the physical facts, but not *too* different—the existence of these social facts is compatible with physicalism. So, according to non-naturalists, at least some normative facts are not fully grounded in any other kinds of facts (like natural facts).¹⁹ Another way to make this point is that the phrase “just too different from natural facts” is obscure and “not being wholly grounded in the natural” is a plausible way of making it clear and understandable.

Assuming that grounding is well-founded, one noteworthy consequence of non-naturalism, formulated as the thesis that some normative facts are not fully grounded in non-normative facts, is that there are fundamental normative facts. To see this, consider a normative fact that is not fully grounded in non-normative ones. If grounding is well-founded, this fact should have some ultimate grounds—fundamental facts that ground it. Since this fact is not fully grounded in non-normative facts, among these ultimate grounds there should be at least one normative fact. So, there is at least one fundamental normative fact.

In this thesis, I will argue against the grounding formulation of non-naturalism—more specifically, I will argue against the claim that there are fundamental normative facts. There are other ways to formulate non-naturalism using contemporary metaphysical tools, and there is a debate on what is the best formulation of non-naturalism. For those who accept the grounding formulation, this argument would be an argument against non-naturalism. For simplicity, usually when I say “non-naturalism” from now on, I mean the grounding formulation of non-naturalism.

1.3 Motivations

¹⁹ Enoch (2019, part 6) suggests this picture.

We saw how normative non-naturalism can be formulated in the grounding framework. The main aim of my argument is to argue against this grounding formulation of non-naturalism. In this part, I will give some motivations to engage in such an argument.

The main motivation is to shed light on the naturalism vs. non-naturalism debate. The argument is against the grounding formulation of non-naturalism. If that is a good formulation of non-naturalism, as I tried to say so in the last section, then my argument would be an argument against non-naturalism. For those who accept the grounding formulation, this will be a good motivation—the argument adds to the rich debate between naturalists and non-naturalists. For those who are still not convinced that grounding is a good tool for formulating non-naturalism, this is not a good motivation—since my argument is only against the grounding formulation, not any formulation, of non-naturalism. Thus, I will provide some other motivations for my argument against fundamental normative facts.

First, that there are fundamental normative facts is an interesting thesis in its own right since it says something about the fundamental layer of reality. It is usually important to know whether facts of some kind are fundamental or not: to know whether reality contains mental, social, or lawlike facts in the most fundamental layer. Some metaphysicians, such as Sider (2011, 4.5), have even posited a special epistemic value for beliefs about fundamental reality. It is interesting if reality has normative facts in its fundamental layer: if God wants to create the world, she needs to directly create some normative facts after creating everything else. The fundamentality of normative facts is a substantive metaphysical doctrine with important implications—for instance, if it is true, then general physicalism, the thesis that everything is wholly grounded in the physical facts, is false. Thus, for anyone who is interested in knowing whether general

physicalism is true, the following argument would be important. Even if the fundamentality of normative has nothing to do with non-naturalism, it is still worth perusing.

The second motivation comes from the debate over how to understand non-naturalism. There is a growing literature on how to best formulate non-naturalism, and the grounding formulation has always been a serious option. Many non-naturalists have engaged with grounding and fundamentality formulations of non-naturalism; some have endorsed these formulations, and some have rejected them.²⁰ Both parties want to formulate a strong version of non-naturalism that cannot be simply refuted (at least, not refuted by a new argument that hasn't been among the traditional arguments against non-naturalism). The following argument, if sound, is a new argument that refutes a (important) grounding formulation of non-naturalism. Therefore, for non-naturalists who want a theory that cannot be simply refuted, it supports the case against grounding formulations: it makes more sense to formulate non-naturalism with, say, essence to avoid this argument.

²⁰ see footnote 15.

Chapter 2: The Main Argument

In the previous chapter, I discussed grounding, non-naturalism, and the grounding formulation of non-naturalism. In this chapter, I will present the main argument. First, I present a short version of the argument which has two premises. Then I defend the premises in turn.

2.1 The Short Version

For this argument, we need to extend the notion of fundamentality and non-fundamentality to objects and properties. An intuitive characterization is this: the solar system, laptops, and the properties of *being a city* and *being red* are derivative; electrons and the property of *having a negative charge* are fundamental. In the same way that the brownness of the table is not fundamental while the physical properties of its fundamental particles are, the table itself is not fundamental while its fundamental particles are fundamental. For now, this intuitive characterization is enough; I will talk more about how to theorize about fundamental objects and properties later.

Let's call a fact *pure* iff all its constituents are fundamental—when it involves only fundamental entities. Then the main argument against the grounding formulation of non-naturalism, formulated succinctly, is:

(Normative Impurity) All normative facts are impure: “killing persons is wrong” and “every action that fails to maximize utility is wrong” involve derivative entities such as persons, acts of killing, or the property of *failing to maximize utility*.

(Fundamental Purity) All fundamental facts are pure: fundamental facts do not involve derivative entities.

Therefore, non-naturalism is false.

Before proceeding, I should mention that Justin Morton has provided a similar argument against normative principles being fundamental (Morton 2020). He is arguing against what he calls “structured non-naturalism”, a version of non-naturalism that is committed to normative facts being grounded in principles. In one step of the argument, he argues against the fundamentality of normative principles with a thesis equivalent to purity.

Here is his argument. He mentions a few examples of normative principles, such as “An act is wrong iff it doesn’t maximize happiness.” (Morton 2020, p. 189). Then he says that if we accept usual normative principles as fundamental, then “happiness and painfulness turn out to be fundamental properties, and beliefs and wills turn out to be fundamental objects. But that seems, at the very least, radically controversial.” (Morton 2020, p. 189) A bit later, he mentions that atomic normative facts are not good candidates for fundamentality since “are acts of destruction fundamental entities?” (Morton 2020, p. 193) Morton then provides some replies to this argument and rejects them, such as the reply that normative principles might be like “when the fundamental particles are arranged thus-and-so, then one ought to Φ ” (Morton 2020, p. 190). He rejects it immediately since “it would be outrageous to hold our normative metaphysics hostage to such principles. That’s just not what the fundamental normative principles look like, on any plausible normative theory.” (Morton 2020, p. 190)

Here is how what I present here differs from Morton’s. First, Morton’s discussion is very brief and is part of a larger argument, while here I try to engage with every step of the argument in detail. Second, he does not provide more support for the claim that actions and painfulness are not fundamental. As we will see, I try to support this claim by providing some systematic criteria

for the fundamentality of objects and properties and applying them to the constituents of normative facts. Third, I discuss different non-naturalist responses in detail. Here I tried to examine some other non-naturalist replies and objections that are not mentioned by Morton. Admittedly, the reply mentioned in the last paragraph is similar to the purification strategy, but Morton does not put any time in saying why this response fails, while I explain in detail why this strategy fails both for atomic normative facts and normative principles. Fourth, the part about non-naturalists' strategy of another fundamental realm in the conclusion is new.

2.2 Fundamental Purity

Purity says that derivative entities cannot be a constituent of fundamental facts. Sider (2011, chapter 7) is one of the first people to formulate and advocate purity. He writes that “*fundamental truths involve only fundamental notions*. When God was creating the world, she was not required to think in terms of nonfundamental notions like city, smile, or candy.” (Sider 2011, p. 106) Below I give three reasons for accepting purity.

The first reason to accept purity is that it is highly intuitive. It's odd that the solar system being such and such, or something being a city, are fundamental facts. There is no mention of cityhood, laptops, or the solar system in the most fundamental layer of reality—all facts about the solar system are grounded in facts about its microstructure. As Sider writes, “This [purity] might seem obvious and uncontroversial.” (Sider 2011, p. 107) Purity is a widely accepted thesis, especially in the literature on how grounding facts are grounded.²¹ Grounding facts are facts of the form of “p grounds q”. They always involve one derivative constituent, which is q, so they are impure.

²¹ See Sider (2011), deRosset (2013), Rosen (2010), Dasgupta (2014) who explicitly endorse purity (though maybe not under the name “purity”. Raven (2015) and Bennett (2011) are also relevant. For dissent, see Barker (2023).

Purity entails that grounding facts cannot be fundamental and are grounded. Hence, there is a question of how grounding facts are grounded.

The second reason to accept purity comes from DeRosset (2013). He introduces LINK:

LINK: “ e_1, \dots, e_n are the entities that ground entity e only if e 's existence and features are all explicable solely by reference to the existence and features of e_1, \dots, e_n ” (DeRosset 2013, p. 4)

LINK plausible principle, connecting the notion of fundamental fact to the notion of fundamental entity. For instance, if the house is grounded in its building blocks, then if we put the blocks in place then everything about the house is fixed—all features of the house in explicable only with reference to the blocks. Thus, every fact about a derivative object is itself grounded in facts about fundamental objects that ground it. DeRosset notices that, from LINK, we can infer that “an entity e is fundamental if e 's existence or its possession of some feature is fundamental,” which is equivalent to purity. So, purity is plausible to the extent that LINK is plausible.

The third reason comes from a plausible way of understanding the notions of fundamental object and fundamental property. One might posit a *definitional* connection between fundamental facts, objects, and properties, which amounts to defining fundamental object and fundamental property as every object or property that is a constituent of a fundamental fact. Alternatively, following Sider (2011, p. 116), one might define fundamental fact as any fact that involves only fundamental objects and properties.²² If one accepts this understanding of these notions, purity would be true by definition.

²² Sider calls them fundamental “notions”.

Sider (2020) notes that, when we want to argue that a fact is not fundamental because it is not pure, we do not need Purity in a general form. My argument would work with weaker, less controversial principles stating that actions of killing, or the property of failing to maximize utility, should not appear in fundamental facts. These principles are individual instances of Purity. They are what Sider calls “one-off” principles.²³ I hope that Purity, or one-off instances of it, is intuitive enough as I will assume Purity in what follows. After I present the argument, I will talk a bit more about Purity in Chapter 4.

2.3 Criteria for the Fundamentality of Objects and Properties

Before defending Normative Impurity, we need to introduce the criteria for assessing the fundamentality of objects and properties, as we need to assess whether objects or properties in normative facts are fundamental or not. With the help of the literature, I will introduce four intuitive criteria that signal that an object or property is non-fundamental.

(I) Being mid-level: An object is mid-level iff (i) it has parts, **and** (ii) it is a part of something—e.g., tables and planets. Being mid-level is a good sign of being derivative since mid-level objects are not fundamental on two main pictures of the priority of parts and wholes. I will discuss these two views in what follows.

If one accepts that wholes are completely grounded in their parts, then mid-level objects are derivative since they have parts. As Schaffer (2010, 1.4) calls it, this is the *pluralistic* picture. Alternatively, following Schaffer (2010), if one thinks that parts should be grounded in the

²³ See Sider (2020, footnote 6).

wholes, then mid-level objects are derivative since they are parts of other objects—this is the *monistic* picture. These two principled ways of thinking about the priority relation between parts and wholes both entail the non-fundamentality of mid-level objects. I find other theories about this priority unmotivated and arbitrary. Even Sara Bernstein, who argues for the fundamentality of mid-level objects, accepts that “[t]raditionally, these two views are thought to be the major options in logical space for proponents of grounding and fundamentality (Bernstein 2021, p. 1068).²⁴ In light of these considerations, I propose that being mid-level is a sign of non-fundamentality, at least for those who hold the monistic or pluralistic view.

(II) Naturalness: Perfectly natural properties are first introduced by Lewis (1983). They are supposed to capture the intuitive difference between the property of *being an electron* and *being an electron or a cow*, or between *being green* and *being grue* (as Goodman (1955) has first formulated, something is grue iff it is green and observed before some future time, or it is blue and not observed before that time). They are the special properties “whose sharing makes for resemblance, and the ones relevant to causal powers.” (Lewis 1983, p. 347). Non-natural properties are gerrymandered and arbitrary, while natural properties carve the reality at its joints. They also have other theoretical utilities, such as formulating physicalism (Lewis 1983, p. 364).

It is intuitive to think of fundamental properties as perfectly natural properties. Sider agrees with that: his notion of structure or fundamentality is an extension of David Lewis’ notion of natural properties and relations: “If the concept of structure is to play this role in metametaphysics, it must be generalized beyond Armstrong’s notion of a universal and Lewis’s notion of natural properties and relations.” (Sider 2011, p. 8) One reason for equating naturalness and

²⁴ Note that the two views that Benrstein refers to accept the well-foundedness of grounding.

fundamentality is that fundamental properties and natural properties have similar theoretical roles. For instance, that “sensing red” is a natural property threatens physicalism just in the same way that “sensing red” is a fundamental property. So, not being perfectly natural is a sign that an entity is non-fundamental.

(III) Having vague boundaries: The term ‘P’ is vague when it has borderline cases, such as ‘bald’. As Rohan Sud observes, it is often assumed that if one rejects metaphysical vagueness, then a predicate being vague means that it has no corresponding perfectly natural property (Sud 2018, p. 243). For instance, Sider (2003, part 3) argues that existence is not vague since existence is a unique natural kind. His reasoning is as follows. First, he assumes a linguistic theory of vagueness, according to which whenever the term ‘P’ is vague, it is because there are multiple candidate meanings that are equally good as a meaning of the term—these are the precisifications of the term. Second, he assumes the reference magnetic theory of meaning (semantic values), according to which the meaning of a term is determined by a balance between use and eligibility, where natural properties are more eligible. Then, when there are multiple precisifications that are equal with respect to use, there is no natural property nearby: if there were a natural property nearby, we would have determinately referred to that natural property because of the reference magnetism—which means the term wouldn’t be vague anymore.

Thus, if ‘P’ is vague then there is no natural property in the vicinity of what we are talking about when we use ‘P’. If we see perfectly natural properties and fundamental properties as being the same, then ‘P’ being vague is a sign that it refers only to non-fundamental properties. This reasoning presumes a linguistic theory of vagueness, but it’s not harmful given the popularity of

linguistic theories of vagueness.²⁵ This line of thought can be extended to objects. A referring expression ‘S’ has vague boundaries iff for some object O it is borderline whether O is part of S—e.g., oceans and galaxies. Then we can say that if ‘S’ is vague then it does not refer to a fundamental object, by extending reference magnetism to referring expressions. This is more controversial but possible.

(IV) Being Explainable: This test is inspired by a plausible principle proposed by Louis DeRosset that I discussed in the last section:

LINK: “ e_1, \dots, e_n are the entities that ground entity e only if e ’s existence and features are all explicable solely by reference to the existence and features of e_1, \dots, e_n ” (DeRosse 2013, p. 4)

The idea is to connect grounding between facts and grounding between entities: if the existence and features of some entity are grounded in the existence and features of other entities, those other entities ground that entity. This means that if the existence and features of some entity are grounded in facts that do not contain that entity, then that entity is non-fundamental.

Thus, say an object S is *explainable* iff all facts about it are grounded in facts about other entities. For instance, every property of this table (color, rough surface, etc.) and its existence is grounded in facts about its microstructure. Being explainable is strong evidence of non-fundamentality. To see why, consider an explainable object, like a table. When every fact about the table is grounded in facts about other entities, it seems redundant and arbitrary to posit the table as fundamental. The same consideration can also be said for explainable properties. We

²⁵ See references in Sorensen (2023), for instance on supervaluationism which is one important linguistic approach to vagueness.

expect that if an object or property is fundamental, there is at least one fundamental fact about that entity. Otherwise, it is not clear what we mean by calling an entity fundamental.

It would be helpful to apply these criteria to some intuitive examples and see whether they work properly in these examples.

(1) Consider an electron; it is plausibly a fundamental object because it is considered a basic building block of reality in physics. It is not mid-level since it does not have any parts—if it has a part, it means that it is not a fundamental particle. It is not explainable since many facts about electrons are fundamental facts—such as their velocity or charge. It does not have vague boundaries in the same sense that a pile of sand has vague boundaries. Maybe it is vague whether it occupies some space, but for every other object, it is not vague whether that object is part of the electron or not. Thus, it does not have any non-fundamentality criteria, which is expected.

(2) Consider the property of being a city; intuitively, it is a derivative property. It is not a perfectly natural property since cityhood does not account for any objective similarity in the world. It has vague boundaries: for many given places with borderline populations, it is hard to tell whether it is a city or not. At least, for places that are not yet labeled as a city or a non-city, it can be vague whether a place is a city or not. It is also explainable. That this place is a city is explained by its population and many other factors about its buildings and what happens inside it.

(3) Consider Vienna; intuitively, it is a derivative object. It has vague boundaries: even though there is a border in the maps that distinguishes Vienna from its surrounding lands, still for many atoms in the border it is borderline whether that atom is part of Vienna or not. It is a mid-level object: Central European University is part of Vienna, and Vienna is part of Austria. It is

explainable: there are many, many facts about Vienna, but all of them are explained based on what is going on in the smaller parts of the city. For instance, that Vienna has clean weather is explained by the distribution of different gases in its air.

2.4 Normative Impurity

The impurity of *familiar* normative facts seems plausible since these facts involve high-level, complex objects and properties that cannot be fundamental. It's hard to believe that an action of killing, or the property of utility maximization, is a fundamental entity (I will argue for this later). But Normative Impurity is an ambitious thesis since it states no normative fact is pure. To reject it, non-naturalists only need to show that a small number of normative facts are pure. They might even introduce new normative facts to do this job. Thus, I defend Normative Impurity by investigating different forms of normative facts that are mentioned in the literature: atomic normative facts and normative principles.²⁶ These correspond to two main forms of non-naturalism: brute and principled non-naturalism. I will examine these candidate normative facts and show that they have derivative constituents, so they are not pure.

For simplicity, I assume that the basic normative property is *wrongness* (W), so I discuss only normative facts that have wrongness as a constituent. Also, I assume that the correct normative theory is that everything that fails to maximize utility is wrong. Nothing hinges on these two assumptions. Occasionally, I consider other normative theories to show that the same problems occur for them.

²⁶ See references in footnote 15.

2.4 .1 Atomic Normative Facts

Each atomic normative fact states that a specific, non-repeatable action token is wrong: John's cheating on yesterday's exam is wrong. Their logical form is $W(g)$, where g is an action token. They can also take other forms, such as $W(g, x, t)$, which says agent x 's doing g in time t is wrong. These differences do not matter for our purposes, as long as these facts involve action tokens or action types. *Brute non-naturalism* is the corresponding view that some atomic normative facts are fundamental, and they ground other normative facts.

Atomic normative facts are not pure because they involve *actions*. Actions are not fundamental.

To see this, let's apply our criteria:

- (I) Actions have smaller spatial and temporal parts: they contain simpler or smaller actions, or body movements. They are parts of other things, such as bigger actions. For instance, John's secretly seeing the other student's paper is part of John's cheating, which is in turn a part of John's trying to impress his friend with good grades. Thus, actions are mid-level.
- (II) Terms referring to actions are vague. It is borderline whether some last nanoseconds of the cheating are included in what 'John's cheating' refers to or not: When exactly did the cheating end? The spatial borders of actions are also vague because the spatial borders of the body are vague. Thus, actions have vague boundaries.
- (III) Facts about actions are explained by instantiations of physical and mental properties. Consider John's cheating on yesterday's exam. Some facts about it include what it has caused, its causes, whether something is a part of it, when it occurred, and its

agent. All of these facts are explainable by facts about the action's simpler constituents, such as the details of John's movements, sensory experiences, and maybe intentions and thoughts. Therefore, actions are explainable.

Non-naturalists might object that we are begging the question in the third criterion: if there are fundamental normative facts about actions, then surely actions are not explainable because this normative fact is a fact about *them*! In response, I would say that to avoid begging the question on both sides, we should put aside both non-naturalism and its denial and assess the fundamentality of actions in a way that is neutral to there being any fundamental normative fact. This way, we will see that all facts about actions are grounded except *possibly* their normative status. This does not entail that actions are explainable, but I think that it is strong evidence that they are explainable: the vast majority of facts about actions are grounded. This is how the explainability criterion is often used: we assess whether the vast majority of facts about an object are grounded or not.

2.4.2 Normative Principles

Brute non-naturalism is not a popular thesis since it is considered at odds with moral or normative practice: whenever an action is wrong, we always search for some features of the action—such as its harming other people—that explain its wrongness. To explain the pattern of normative instantiations, non-naturalists usually introduce normative principles. For instance, Rosen (2016) suggests that every particular normative fact “is grounded in some non-normative fact $\phi(\alpha)$ together with a bridge law connecting F and ϕ ” (p. 163). In general, these principles express a law-like relation between a non-normative property (failing to maximize utility, or N) and a normative property (wrongness, or W). The principle can take many forms, including: (i) A

universal generalization: Necessarily, for all x , whenever x fails to maximize utility, then x is wrong (as suggested by Rosen 2016, p. 164); (ii) a higher-order relation between the non-normative and the normative: $L(N, W)$. L is a relation of higher-order normative necessitation, or normative grounding, normative lawhood. The details of the normative principle do not matter here. The relevant thing is that the normative principle somehow contains both a natural and a non-natural property.

The principle metaphysically explains the normative correlation by unifying the elements of the correlation and showing that each element of the correlation is an instance of the same principle. This can be done by grounding particular facts in the laws. Following Rosen: “Particular facts of the form $F\alpha$ are grounded in the natural features φ of α , together with a general law that somehow bridges the gap between the two.” (Rosen 2016, p. 23) The claim is that whenever X is wrong, $W(X)$ is grounded in $N(X)$ plus the principle that connects N and W . So atomic instantiations of wrongness are grounded in the non-normative property instantiation plus the normative principle, and the normative principle is fundamental. For example, the wrongness of John’s cheating is grounded in the fact that it failed to maximize utility, plus the principle that everything that fails to maximize utility is wrong. We call this *principled non-naturalism*.²⁷ This theory is more popular in the literature than brute non-naturalism.²⁸ Thus, normative principles require more attention and scrutiny.

²⁷ Rosen (2016) introduces these two types of non-naturalism. See also Rosen (2017).

²⁸ See Enoch (2019), Rosen (2016), Rosen (2017), and FitzPatrick (2008).

Normative principles cannot be fundamental since they are not pure. This is because all such principles contain a derivative non-normative property. Let's apply our criteria for the fundamentality of properties to N, the property of failing to maximize utility:

- (I) 'failing to maximize utility' is a vague term because the term 'utility' is vague. Even if utility is restricted to mental states, it is not clear how to weigh different kinds of pleasures, pains, and other kinds of mental states against each other. Utilitarians have tried to precisify the notion of utility, but these precisifications just limit the amount of vagueness, rather than eliminate it. Consider the states of having pleasure from drinking water and having pain from pinching your arm. How these two states are to be compared? There are infinite ways of weighing them based on assigning different real numbers to them. Unless there is some accurate measurement of degrees of pleasures and pains, 'utility' remains a vague term: for many combinations of pleasures and pains, it remains borderline whether that combination increases or decreases the utility of the person involved. The economic concept of subjective value, the amount of money you are willing to pay to have that pleasure or to mitigate that pain, does not help. The subjective value reflects how much you think that pleasure and pain have utility, not how much actually they have utility—you might irrationally pay a lot of money to have a petty pleasure. Moreover, there is another dimension in which 'failing to maximize utility' is vague: who counts? It is not clear how the utility of different non-human animals is to be counted and at what point of a fetus's development its utility is counted. So, in many cases where different animals' utilities are opposed to the utility of humans, it is borderline whether some action fails to maximize utility or not.

- (II) N is explainable. It is a complex property, summing individual utilities of many persons, which means every fact about N is grounded in simpler properties of individual utilities. Even individual utilities are grounded in a complex pattern of the person's phenomenal states, desires, and environmental factors. Every fact about N is also grounded in facts about consequences of actions, which are complex causal facts connecting actions and the environment. This means that N is explainable.
- (III) The property of failing to maximize utility does not seem to be perfectly natural—its instances do not exhibit a high degree of objective similarity. This property is not even an intrinsic property: whether an action fails to maximize utility depends on external factors and other events. For instance, if we add one human to the world, an action done by someone else that fails to maximize utility might now maximize utility. It is implausible that a property so sensitive to external factors accounts for some objective similarity between its instances. Surely, instances of some properties can be to some degree objective due to external factors, such as social properties, but they cannot have the maximum degree of objective similarity required for perfect naturalness.

Other candidate non-normative properties (to appear in the normative principle) do not fare better. Consider the property of *treating others as mere means*. Again, every instantiation of this property is grounded in the intentions of the agent, the relation between the action and other people (which other people are involved), causal relations between these others and the ends of the agent (whether they can be means to her end), whether others have consented (or would consent) to be used in such a way, etc. These properties, in turn, are grounded in countless other more fundamental properties.²⁹ The vagueness of the previous discussion on who counts as

²⁹ SEP Treating as means

“others” reappears here: is treating a highly intelligent animal as mere means wrong? The general point is that there is no fundamental non-normative property that unifies all wrong actions.

Chapter 3: The Purification Strategy

These brute and principled non-naturalists can reshape normative fundamental facts and make them pure. This is done by substituting usual normative facts with equivalent facts that have some other logical form. We can call this *purifying* normative facts. I will examine purifying atomic normative facts and normative principles in turn and argue that this strategy makes non-naturalism implausible.

3.1 Purifying Atomic Facts

This is how brute non-naturalists might purify atomic normative facts. Suppose John's cheating on yesterday's exam is ultimately grounded in P_1, P_2, \dots, P_n . Every P_i is a fundamental fact; for instance, they are facts about the fundamental particles comprising John's hands. Non-naturalists then claim that $W(P_1, P_2, \dots, P_n)$ is the fundamental normative fact that expresses the wrongness of John's action. The wrongness property does not attribute wrongness to one specific action; rather, it attributes wrongness to some fundamental facts taken together, i.e. the plurality of them. Wrongness is like a relation over fundamental facts, stating that the *joint instantiation* of its relata is wrong. It's clear that $W(P_1, P_2, \dots, P_n)$ is pure, so it can be fundamental.

The more important issue is that this strategy posits unexplained correlations at the fundamental level. Consider all collections of fundamental facts out there (it's quite a huge number of collections). A small number of these collections instantiates W . Surprisingly, all W -instantiating collections have common features: all of them ground actions, all of them ground the presence of an agent, and all ground actions that fail to maximize utility. This means that facts in one of these collections have a correlation with being an action and failing to maximize utility. Call this the normative correlation. It cannot be a mere coincidence. However, according to the above picture,

there is no further explanation for that: these wrongness facts, or the totality of them, are not grounded in anything else. Since ground and metaphysical explanation have a tight connection, we can conclude that there is no metaphysical explanation for these wrongness facts.³⁰

Admittedly, the bruteness of the normative correlation is an *essential part* of the brute non-naturalist view: once you are a brute non-naturalist, you are already committed to the normative correlation having no explanation. Thus, it seems that this objection begs the question against brute non-naturalists. In response, I would say that my aim is *not* to convince the brute non-naturalist to change her view, but to better understand the costs of brute non-naturalism from my own point of view, or from a neutral point of view. The following discussion will show that brute non-naturalism is costly, given some natural epistemic and metaphysical principles. Surely, this discussion is not effective against someone who has already accepted the costs.

Let's illustrate why the unexplained normative correlation is implausible with an example.

Consider the collections of fundamental facts such that whenever all members of a collection are present, then there is a white object. These collections are correlated with the property of being white. There is a straightforward explanation of this: for each of these collections, its members ground the whiteness of an object, so it is no surprise that a white object exists if those members obtain. Now, to make a surprising correlation, suppose that being in pain is a fundamental property of a bunch of particles. If someone declares that all white-instantiating collections also instantiate being in pain, then this correlation is surprising and cannot be explained easily. In our case, there is a correlation between instantiating W, which is a fundamental property, and jointly grounding actions or failing to maximize utility.

³⁰ Grounding is either identical to metaphysical explanation or backs metaphysical explanation. See Raven (2015).

One might object that unexplained correlations happen all the time. Suppose the spin of electrons (which can be zero and one) is regular over consecutive observations: there is one after zero and zero after one (10101010...). There might be no metaphysical explanation for this pattern—especially if we adopt a Humean picture of laws of nature. So, the objection goes, the same is true for the pattern of wrongness instantiations.

However, the normative correlation is more surprising, or striking, than any old pattern we might see in nature's regularities. The claim is that (1) surprisingness, the quality of something crying out for explanation, comes in degrees, and (2) the normative correlation is more surprising than the spins' regularity.

To see why surprisingness comes in degrees, consider a second kind of pattern: a sequence of zeros and ones that encodes *Hamlet* in binary code (it's like what's written in a file containing Hamlet in your computer). The occurrence of this pattern in nature, without any explanation, is totally unbelievable: imagine your reaction if you find out that the sequence of observations of spins tells you the story of Hamlet. The Hamlet pattern is more surprising, striking, or unbelievable than the zero-after-one pattern: it cries out for explanation more than the zero-after-one pattern.

To argue that the normative correlation is more surprising, we need to know what makes a correlation or pattern surprising. There is a literature on this, and it is hard to adjudicate this matter here. I briefly mention some accounts, extracted from Bhogal (manuscript), and show that all of them entail that W-instantiations are surprising.

(1) A fact is surprising if and only if it is unlikely according to the current theory, and there is some other hypothesis that makes it more likely (Horwich 1982). If so, then the normative

correlation is surprising since it is unlikely according to brute non-naturalism but is far more likely according to many other theories. Brute non-naturalism does not predict the normative correlation, so it does not raise the probability of normative correlation. On the other hand, other metanormative theories predict the normative correlation. For instance, if atomic normative facts are grounded in normative principles or laws, then the correlation is expected—the correlation would be certain. Or if normative facts are grounded in human desires and motivations, then this correlation is more likely because we actually have desires for maximizing the utility of all—maybe for evolutionary reasons.

(2) A pattern is surprising to the degree that it is not simple. For instance, the Hamlet pattern is less simple than the consecutive zeros and ones, so it is more surprising (Bhogal ms, 3.2). It is clear that “being correlated with actions” and “being correlated with utility maximization” do not make a simple pattern, as these properties are very high-level and complex (see next section for more on this).

(3) A fact is surprising when it suggests that there is some authorship or agency causing the fact: when the hypothesis of some agent bringing about the fact provides understanding about that fact (Bhogal ms, part 4). The Hamlet pattern, for instance, is a clear example of this. Note that, as Bhogal emphasizes, “suggesting” authority here does not mean that the actual explanation of the fact is that there is an authority. Following Lipton (2004), he draws a distinction between loveliness and likeliness of an explanation. The authority explanation might be an unlikely one for other reasons—for instance, for the simple reason that there is no God. But the authority explanation is a lovely one: it provides understanding if it were true.

The normative correlation suggests someone behind the curtain creating the correlation.

Somebody cared about us and created normativity in a way that obeying it increases our utility:

normative facts could command murdering and annoying others, but they do not. Also, it is surprising that normativity is about us: instead of actions, chairs and tables might be wrong and right.

Admittedly, the normative correlation is not a problem only for purified strategy: brute non-naturalism, in general, entails that there is an unexplained normative correlation between wrongness instantiations and failing to maximize utility. The specific problem for the purified strategy is the correlation with actions: all wrongness instantiating collections ground actions. To see this, consider a purity-friendly version of brute non-naturalism: actions are fundamental, and some of them are fundamentally wrong. Here there is a correlation between normative instantiations and actions, but it can be better explained since actions are fundamental: the line between actions and non-actions is not an arbitrary one, so it is expected that actions have some fundamental features. However, in the purification strategy, actions are not fundamental and the correlation is not explained.

3.2 Purifying Principles

Suppose the normative principle is $L(N, W)$. Gideon Rosen in personal communication suggested substituting N with a fundamental expression to make the normative principle pure. In the following, I will explain and criticize one way of doing so. Although this might not be the exact strategy that Rosen has in mind, I think it is sufficiently close to it. At least, it is the most powerful purification strategy that I can think of.

Non-naturalists need to substitute N with an expression ‘ M ’ such that (i) all constituents of M are fundamental entities, and (ii) the property M is *sufficiently close* to N : for now, let’s say M should be at least necessarily coextensive with N . In this case, $L(M, W)$ would be the purified

principle that can be fundamental. We call M the *metaphysical definition* of N. This strategy has been adopted elsewhere in metaphysics, especially when it is needed to get rid of some derivative constituents in a fact.³¹ The metaphysical definition of some property is a disjunction of properties that describe, in fundamental terms, each of its possible instances. For example, “being a chair” is defined by “being a duplicate of this bunch of fundamental particles (the first instance of chair) or being a duplicate of that bunch of fundamental particles (the second chair), or...”

Here is how to construct metaphysical definitions in detail. We need to construct *duplicate properties* of objects. Duplicate property D of an object o is the property that (i) is constituted by fundamental entities, and (ii) its instantiation guarantees the instantiation of a duplicate of that object. D can be easily constructed. First, suppose that the existence of o is grounded in $R_1, R_2 \dots R_m$. For each R_i construct the property of instantiation of that fact (P_i); for instance, if R_i is “electron e has spin s”, then its correlative property would be “being an x such that x is electron and x has spin s”— P_i is a property such that its instantiation guarantees the existence a duplicate of the subject of R_i . Now the duplicate property D is a relation over m number of objects and is the conjunction of all P_i s (being $x_1 \dots x_m$ such that x_1 is P_1 and x_2 is $P_2 \dots$). The instantiation of D guarantees that duplicate objects of the subjects of R_i s would exist, and since R_i s ground the existence of o, it guarantees the existence of a duplicate of o. For instance, the instantiation of the duplicate property of the Earth guarantees that an object qualitatively indiscernible to Earth is instantiated.

³¹ For instance, see Sider (2020 p. 759) and Schaffer (2013).

Now we can construct metaphysical definitions. The property N has some actual and possible instances, call them $O_1 \dots O_n$: “being a chair”, for instance, is realized by many chairs in our world and other possible worlds. For each O_i , call its duplicate property M_i . M is the disjunction of all M_i s: being $x_1 \dots x_m$ such that it is a duplicate of O_1 or it is a duplicate of $O_2 \dots$. Since all N s are constituted by fundamental entities, M is also solely constituted by fundamental entities.

Now the non-naturalist strategy would be to remove “failing to maximize utility” from the principle and to put in its place *its metaphysical definition* M: L (M, W) is the new principle; it claims that there is a higher-order normative necessitation or grounding between the natural property M and the normative property W. This new principle is pure, so it can be fundamental and ground individual instances of wrongness. An action X is wrong because the fundamental particles that constitute X satisfy M, and because everything that satisfies M also satisfies W (which is the principle itself, or a consequence of it).

I mention three problems with the purified normative principle. First, the purified normative principle, posited to explain the normative correlation, cannot do any explanatory work because it is not unifying. Unification means showing that different cases are all cases of a single, general pattern. Unification is usually taken to be necessary, or important, for any explanation (Schaffer 2017a, p. 306). To see why the principle is not unifying, note that in its current format, it contains a description of all actions that possibly fail to maximize utility. Thus, instead of cohesively *unifying* all instances of the normative correlation under a single property, the principle merely restates the correlation in one big disjunctive law. These principles are very much similar to “supervenience principles” that connect moral to natural features—principles that are constructed based on how moral facts supervene on natural facts. These principles are taken to be not

explanatory, and contain irrelevant information (see Little, 2000 part 3; Ridge and McKeever 2016, part 3; Dancy 2004, 5.3).

To see why the disjunctive normative principle cannot be explanatory, suppose it is of the form $L(P_1 \text{ or } P_2 \text{ or } \dots \text{ or } P_n, W)$. This law means that whenever each of the facts (or properties) $P_1 \dots P_n$ obtains, then W also obtains. This amounts to saying that for each P_i , if P_i obtains, then W obtains. This last sentence is exactly what $L(P_i, W)$ says. Therefore, positing this law is equivalent to positing many local laws: $L(P_1, W), L(P_2, W), \dots, L(P_n, W)$. This point is more vivid if L is a simple material conditional (or even a strict conditional): $L(P, Q)$ just is “if P , then Q ”. Then $L(P_1 \text{ or } P_2 \text{ or } \dots \text{ or } P_n, W)$ is logically equivalent to $L(P_1, W)$ and $L(P_2, W), \dots$ And $L(P_n, W)$. It is obvious that positing many principles of this local form is ad hoc and not explanatory. For each instance of the normative correlation, we simply posited a new principle to locally explain why the non-normative instantiation is correlated with wrongness instantiation. This huge number of normative principles are equally in need of explanation as atomic normative facts.

Second, the principle contains many disjunctions and conjunctions. However, intuitively, all facts containing disjunction or conjunction are grounded in their conjuncts or disjuncts—such as paradigm cases of grounding, “ $A \& B$ ” and “ $A \text{ or } B$ ”. Grounding theorists widely agree that the truth of conjunctions and disjunctions is grounded in their conjuncts and disjuncts.³² Admittedly, the normative principle differs slightly from “ $A \& B$ ” or “ $A \text{ or } B$ ” because the disjunction appears in the description of one of its properties (if the principle is a higher-order fact), or in the antecedent of a conditional (if the principle is a universal generalization). But still, I take it that

³² See Rosen (2010, p. 117) and Bliss and Trogon (2021, 1.4).

the grounding theorists want *all* facts in which conjunctions and disjunctions appear to be partly grounded in the disjuncts or conjuncts. For example, suppose causation is a fundamental relation between facts. Then facts like “(A or B) cause C” are plausibly not fundamental; they are true either because A causes C or because B causes C. If virtually all facts containing conjunctions or disjunctions are not fundamental, then it is surprising that the normative principle is the only exception.

Third, the principle cannot satisfy the necessary co-extensiveness constraint and be pure at the same time. I borrow this idea from Schaffer (2013, pp. 746-747), where he criticizes metaphysical semantics, the principles connecting fundamental to derivative in Sider’s fundamentality framework. Note that for the new principle to be necessarily co-extensive with the old one, the duplicate property of *all* possible instances of N (failing to maximize utility) should appear in the principle. However, some possible instances of N have impure duplicate properties since they are ultimately grounded in *alien* fundamental properties, which are properties that are not instantiated in the *actual* world but nonetheless are fundamental in some other world. Below is one possible scenario of this kind.

Suppose N is all about reducing pain and increasing pleasure and that our world is a physicalistic world, so all instances of N are grounded ultimately in the brain states of conscious creatures. But there are possible non-physicalistic worlds with non-physical minds and spirits capable of experiencing pain and pleasure. In these worlds, N is ultimately grounded in some non-physical properties that are not fundamental in our world, so possible instances of N have fundamental descriptions that have alien properties as their constituents. These descriptions must appear in the new principle if it wants to be coextensive with the old principle. But the presence of alien properties in the normative principle would either violate purity or simplicity: if these properties

are not fundamental in our world, then purity is violated; if we posit these uninstantiated properties as fundamental, then simplicity is violated—since there are a huge number of properties that could possibly be fundamental. If we remove problematic possible instances from the principle, the principle would be less explanatory: it cannot unify wrong actions in our world with wrong actions in the worlds with alien properties.

3.3 Conclusion

I argued that non-naturalism is untenable because normative facts violate purity: normative facts involve derivative entities, such as actions and non-normative properties such as failing to maximize utility. I also discussed attempts to make normative facts consistent with purity, and none of them were satisfying. Purifying atomic normative facts makes them arbitrary and surprising; purifying normative principles makes them unexplanatory and disjunctive.

Chapter 4: Objections

4.1 Overgeneralization

Non-naturalists might say, “Your argument is against many similar philosophical theses. For instance, consider property dualism: one interpretation of it could be that there are fundamental mental properties but no mental substance. If so, then some mental facts are fundamental, such as the fact that Jones is in pain. For this fact to be pure, Jones should be fundamental. But it is hard to see how Jones is fundamental if there are no mental substances: Jones is probably a person, or a body, or a brain, the existence of which is grounded in many other physical facts. Similarly, psychophysical laws between mental states and their physical basis, which some consider to be fundamental, often involve derivative physical properties. This means all mental facts have derivative constituents (bearers of phenomenal experiences), and so mental facts are not fundamental.

The same can be said against views that claim there are fundamental facts about human agency or will, but there is no new fundamental entity as the “doer” of actions. The same for *social* properties attributed to some ordinary objects or wholes (groups, nations) that are clearly not fundamental: for instance, a possible view (social property dualism) according to which groups of individuals can have a fundamental property of *collectively* intending to do something.³³ Or there cannot be a fundamental law of nature between some derivative properties, like many higher-level physical laws. In general, whenever a theory claims that there are some fundamental facts in some area, it should posit both some fundamental properties and fundamental objects.

³³ This is just a possible view, not something that someone has proposed (as far as I know). There is a view that social entities are fundamental, and a property dualistic parallel of this view is imaginable.

This is too much for your argument to prove. Even if these theories are false, it is bad to eliminate them all at once from the start. At least, the argument loses dialectical strength since it targets many views other than normativism.”

Response: I will respond to this objection by considering different debates one by one. First, let’s discuss the case against property dualism. This worry about property dualism is a real worry that has been brought up in the literature—not framed in terms of purity, but very similar worries (see Zimmerman 2010 and Schaffer 2023). I believe that purity is really a consideration against property dualism, and this doesn’t weaken the above argument. Purity does not outright falsify property dualism because property dualists might purify mental facts, as seen above, with high theoretical costs. Other theoretical advantages of property dualism might allow this. What purity does is that it puts a significant cost on property dualism or non-naturalism. Furthermore, property dualists have other options, such as panpsychism or substance dualism, that are consistent with purity: in panpsychism, the subject of mentality is a fundamental particle; in property dualism, it is a fundamental mental substance. Thus, property dualists can adopt these dualistic views while preserving one main tenet of their view—that there are fundamental mental facts. Non-naturalists don’t have similar options.

For other views, the worry is less significant because the parallels of non-naturalism are not attractive positions from the start. The parallel view for the social facts, social property dualism, is an unmotivated view: there is a view, ontological holism, which posits both social entities and fundamental social facts,³⁴ but as far as I know, there isn’t a view in the literature positing only fundamental social properties. The orthodoxy about fundamental laws of nature holds that they

³⁴ See Epstein (2018) for social holism. He writes that holism is “the claim that social entities are fundamental, independent, or autonomous entities, as opposed to being derived from individuals or non-social entities.” (1.2)

obtain between elite properties—elite properties are natural properties in terms of Lewis (1983), and are usually taken to be fundamental; see Sider (2011), who treats fundamentality as an extension of the notion of naturalness. If a law of nature contains a non-elite property, it is evidence that this law is not fundamental (see Hicks and Schaffer 2017, 6.3).

The fact that the parallel of non-naturalism is not attractive might be explained by an implicit belief in purity: exactly because it is hard to imagine fundamental laws or social facts containing derivative properties or objects, no one has advanced such a theory. In some cases, such as property dualism, the purity-friendly theory (substance dualism) was the default option until it faced some other problems, such as mental causation. This moved theorists to explore other options and reject fundamental mental objects. In the case of non-naturalism, I think no one could imagine what a normative “substance” is, so non-naturalism became the view that there are fundamental normative properties that are instantiated by some derivative entities.

4.2 Rejecting Purity

Maybe purity should be rejected. Purity is a general principle about how grounding (or fundamentality) works. But it is unlikely that such a general consideration immediately falsifies many theories in different areas, including non-naturalism. So, we should reject Purity and adopt a conception of fundamentality that is not governed by purity. This is not very costly. Purity is intuitive, yes, but there is no direct argument for it—it seems more like a methodological principle. If so, then generalization gives us good reasons to abandon Purity so that we can accommodate more theories in the grounding/fundamentality framework.”

Response: Purity is more central to fundamentality and grounding than this objection assumes.

In the last chapter, I provided some reasons for purity. I think that’s the strongest way to support

purity. But for those who are still unconvinced, I provide two theoretical reasons for accepting purity.

First, purity allows us to have fewer primitives among the notions of fundamental fact, fundamental object, and fundamental property. If purity is true, fundamental fact can be *defined* as “a fact that contains fundamental objects and properties”³⁵, or fundamental object and property can be defined as “an object or property that appears in fundamental facts.” Call this the *definitional approach*. A nice feature of the definitional approach is that it frees us from treating one (fundamental facts) or two (fundamental object and property) notions as primitive. But if purity is false, there are counterexamples to these definitions, so we need to reject the definitional approach and have independent notions of fundamental objects, properties, and facts. Therefore, rejecting purity would result in a less economical theory of fundamentality.

Second, if we allow three primitives and reject purity, there would be unexplained regularities. Even if purity is false, it is true in the majority of the cases: most candidate fundamental facts that we are aware of are pure; also, most derivative objects and properties never appear in a fundamental fact. Thus, we need to explain why almost all fundamental facts are pure except these ones—for non-naturalists, the fundamental normative facts. A feature of the definitional approach is that it can easily make sense of these regularities: when we *identify* fundamental facts with facts containing fundamental objects and properties, there is no need for any explanation because purity is true by definition, and the regularities would be explained trivially.

Another nice feature of the definitional approach is making sense of the other direction of purity, which I call *purity reverse*: every pure fact is fundamental. Purity reverse is even more intuitive

³⁵ Sider (2011) has a similar view.

than purity: I cannot think of a fact that is entirely constituted by fundamental entities but is grounded in some other facts. The rejection of purity reverse is also not beneficial for non-naturalists or theorists in other areas—such as those who say that there are fundamental mental facts. The definitional approach vindicates purity reverse: if we define fundamental fact as any fact that is constituted by fundamental objects and properties, then purity reverse is true by definition. If we define fundamental objects and properties with fundamental facts, purity reverse is not true, but we can add it to our theory as an axiom with no harm because we have only *one* primitive notion so far (fundamental fact). On the contrary, treating three notions as fundamental does not vindicate purity reverse and we need to treat it as a primitive fact or reject it on pain of being unintuitive.

If one still has doubts about purity, I can only point to one thing. Let's grant that purity can be rejected. But still purity is an important theoretical choice in fundamentality: some *might* opt for it—it still has some plausibility. If non-naturalism can be adopted only by rejecting it, I find it costly for non-naturalists: non-naturalists need to take a stand on a highly theoretical issue on fundamentality. Thus, non-naturalism loses *some* plausibility, at least to the extent that purity is plausible.

4.3 Moral Particularism

According to some strands of moral particularism, there is no general principle unifying all moral facts. Suppose that the normative counterpart of this theory, normative particularism, is true. Then there is no non-normative property that its instantiation is sufficient for wrongness or rightness of an action: cases of wrongness cannot be explained by a general non-normative feature such as failing to maximize utility. If so, then our points about the surprisingness of the

normative correlation would be called into question because there is no normative correlation to start with.

Response: The properties that are required to explain the normative domain, such as failing to maximize utility, are **not** the only properties that wrongness instantiations share. Note that “grounding an action” and “grounding the presence of an agent” are also properties that all W-instantiations share: if $W (P_1 \dots P_n)$, then $P_1 \dots P_n$ are grounds of an action. I believe this is surprising enough for my argument to work.

More generally, what our argument needs is a property that all W-instantiating collections have in common, not a property that all these collections are an instance of it: we need necessary conditions for being a collection, not sufficient conditions for it. Grounding an action is a necessary condition for being a W-instantiating collection—if P is a W-instantiation collection, then it grounds an action. But “failing to maximize utility” purports to be both a necessary and sufficient condition for such collections: something is a W-instantiation collection if and only if it grounds something that fails to maximize utility. The stronger sufficiency condition is not needed for having a surprising correlation, while normative particularism only disputes the sufficiency condition—at least as long as we understand normative principles in the form “if something is X, then it is wrong”. For example, moral particularism is consistent with the thesis that all wrong actions fail to maximize utility, but there are many instances of failing to maximize utility that are not wrong—so one cannot reason from failing to maximize utility to the wrongness of an action.

4.4 Other Normative Properties

It might be objected that it was wrong to focus on wrongness: other normative properties, such as goodness or being a reason, are better candidates to deliver fundamental normative facts. At least, it is not clear why all the above considerations would apply equally to fundamental normative facts that involve, say, goodness.

Response: Nothing hinges on choosing *wrongness* as the fundamental normative property.

Normative facts with other normative properties involve similar entities, and trying to purify them faces similar problems. I briefly discuss other normative properties here.

(1) Consider the normative relation of *being a reason for* that is between an agent, a fact, and an action—some fact is a reason for the agent to do an action. This relation involves actions and agents. Furthermore, the facts that are involved in it, facts that are reasons to do something, are usually not fundamental—unless we assume that only fundamental facts give reasons for actions. Principles for the reason relationship would contain a derivative non-normative property—for instance, facts about utility maximization give reasons.

(2) Consider the property of *goodness*. It can be defined on many things, including actions, facts, ordinary objects, etc. Some of these domains clearly contain derivative entities—actions, ordinary objects—so the resulting normative facts would be impure. The only real contender would be to define goodness over facts. This version avoids my argument only if goodness is defined only over fundamental facts, and then the rest of normative reality is grounded by these goodness facts. But this is highly improbable. This amounts to saying that only fundamental facts about fundamental particles are fundamentally good, and the goodness of an action is grounded in the goodness of its fundamental constituents. I have no idea how this can be done. What does

it mean to say that an electron's spin, say, is good, and how can we explain that an action is good from the goodness of fundamental particles? I conclude that facts containing other normative properties do not fare better when it comes to being pure.

Conclusion

I will first discuss how non-naturalism can still have a normative theory that respects purity. I will then discuss how non-naturalists might reject the grounding formulation from the start and formulate non-naturalism with essence, in a way that sidesteps the purity problem.

Non-naturalists can accept that actions, properties of actions, or failing to maximize utility are not fundamental, but claim that these are not constituents of fundamental normative facts.

Rather, fundamental normative facts involve entities from some other fundamental realm that contain fundamental objects and properties independently of normative considerations. One version of this strategy is to construct normative facts using fundamental mental facts and properties. I think that if non-naturalists want to formulate non-naturalism with grounding, this is their best strategy, but it still comes with a high cost. Let me first explain the strategy and then discuss its plausibility.

A version of this strategy is as follows. First, *badness* (denoted by B), not wrongness, is a basic normative property because we cannot attribute wrongness to mental states. Badness is a property of mental states; for instance, it can be a property of phenomenal experiences or intentions. Second, we need some fundamental normative facts. There are two ways to do this, parallel to different versions of non-naturalism we've seen:

- (1) Brute non-naturalism: Some phenomenal experiences (like most of the pains) are bad, and it is a fundamental fact that an experience is bad. Or the same with some intentions. To respect purity, brute non-naturalism should declare phenomenal experiences or intentions as fundamental. Thus, the fact that an experience E is bad is pure since E and badness are both fundamental.

- (2) Principled non-naturalism: There is a fundamental normative principle that states a relation between a non-normative property of phenomenal experiences N_2 , such as *being painful*, and badness: $L(N_2, B)$. N_2 can also be a property of intentions, such as not *intending the maxim to be a universal law*. To respect purity, principled non-naturalism should declare a natural mental property as fundamental. Thus, the fact that being painful normatively grounds badness is pure since being painful and badness are both fundamental.

These theories are more defensible than what we discussed so far because mental properties or states are more likely to be fundamental than actions or properties of actions. Being painful has much more chance of being fundamental than failing to maximize utility. If non-naturalists accept some form of mind-body dualism, for instance, they can reasonably declare some mental properties as fundamental.

I admit that this strategy can successfully generate pure normative facts, thereby rejecting the first premise of my argument. But it comes with a cost: it makes non-naturalism committed to other substantive philosophical claims, such as dualism, and so the success of non-naturalism would be dependent on the truth of these additional claims. This makes non-naturalism harder to argue for: all physicalists about the mind, for example, would reject non-naturalism. Non-naturalists would not be happy with this. Non-naturalists, for instance, have never said that they have any stance on the mind-body problem. But our discussion showed that they should really take a stance about the nature of mind: plausible non-naturalism implies dualism. Even if not a *cost* for non-naturalists, this claim is substantial enough to be the conclusion of our discussion.

This strategy introduces further complications for non-naturalists. Suppose they declare the facts about the badness of phenomenal states, and not intentions, as fundamental. If they want to explain why some action is wrong, they have no other way than to say that wrongness is grounded in failing to minimize badness in the world. They cannot accommodate a Kantian normative theory, since it is hard to see how *treating others as mere means* can be grounded in facts about the badness of phenomenal states. To accommodate a Kantian theory, they should declare intentions as fundamental. And it is hard to see how the wrongness property of other normative theories, such as contractualism, can be grounded in normative facts about mental states. But this shows that non-naturalist first-order theory is dependent on their commitments about which mental property is fundamental (the fundamentality of phenomenal experiences might be more plausible than the fundamentality of intentions). This is a loss of generality for them: they don't have a unified framework to accommodate all normative theories.

Apart from mental entities, non-naturalists could also use other domains with fundamental entities. For instance, suppose that God fundamentally exists and there are fundamental facts about him, including the fact that he is good or the fact that his will is good. Then, these facts about God can be fundamental normative facts that ground the goodness of other things. For instance, that pleasure is good grounded in God has willed us to have it, which in turn is grounded in the fundamental principle that everything that God wills is good. This version of non-naturalism can be consistent with purity, but I guess that it is unpopular. It faces the same problems as mental normativity: non-naturalism would incur heavy commitments to God's existence and properties.

I mentioned that this is the best non-naturalist strategy. But why is this better than rejecting purity? Here non-naturalists are committed to dualism, which is a thesis in another domain; by

rejecting purity they are committed in the same way. I think the important difference is this: dualism or the existence of God seems like a thesis that can have some consequences in metaethics, but purity seems like a highly theoretical constraint that should not have any consequence for the metaethical theory. If dualism is true, then there are fundamental mental states and properties. It can be imagined that these states have a fundamental value based on how they are felt. Also, God is usually taken to be the source of all goodness. Thus, having a metaethical view that is packaged with some commitments in philosophy of mind or religion is not very far-fetched. On the other hand, purity is a very general principle that does not mention anything close to metaethical theorizing. Having a metaethical theory packaged with rejecting purity seems more ad-hoc and unprincipled.

Non-naturalists can reject grounding as a proper way to understand non-naturalism. Especially, they have provided formulations of non-naturalism in terms of essence. The main idea of essence formulations is that non-naturalism is the thesis that the essence of some normative properties cannot be stated or specified in wholly natural terms.

For instance, consider Leary's (2021) formulation. Leary first defines what it is to be an *essentially non-normative* property. The essence of a property is a set of propositions that directly state the property's nature. (Fine 1994, as cited in Leary (2021, p. 805)) Leary defines that a property F is *involved* in the essence of G iff F is a constituent of a proposition in the essence of G. Then a property is natural iff (i) it is a paradigm case scientific or supernatural property, or (ii) its essence involves only non-normative properties, or (iii) its essence contains sufficient conditions for its instantiation that are stated in wholly non-normative terms.

Essentially non-normative properties are those whose essences *ultimately* involve only paradigm scientific or supernatural (or other non-normative) properties. (Leary 2021, p. 806) She then

defines non-naturalism as the thesis that some normative properties are not essentially non-normative: their essences cannot be specified in wholly non-normative terms, and it does not contain sufficient non-normative conditions for their instantiation.³⁶ (Leary 2021, p. 807)

Essence formulations can be compatible with purity because they are compatible with grounded normativity—normative facts being fully grounded in natural facts, which means there are no fundamental normative facts. The reason is that essence formulations leave different possibilities for grounding normative facts.

For instance, consider Leary (2017) essential formulation of non-naturalism. In her view, normative facts are fully grounded in natural facts, and this is explained by the essences of hybrid properties that contain both sufficient natural conditions for their own instantiation and sufficient conditions for the instantiation of pure normative properties. For instance, being a promise might be a hybrid property where “one might claim that it’s part of the essence of *being a promise* that if certain natural conditions C obtain, then A promised B to do x, and that it’s also part of the essence of *being a promise* that if A promised B to do x, then A has a reason to do x.” (Leary 2017, p. 99) This essence fact can explain (i) why having a reason to do x is grounded in giving a promise to x, and (ii) giving a promise to do x is grounded in natural conditions C. Conditions (i) and (ii) together, with the transitivity of grounding, entail that having a reason to do x is grounded in some natural conditions C. However, the essence of *having a reason* does not mention sufficient conditions for its instantiation in terms of the natural condition C—the work is done by the hybrid property of giving a promise. Thus, *having a reason* can be not an essentially non-normative property, satisfying the essential formulation of non-naturalism.

³⁶ For another essence formulation, see Bengson and Schaffer-Landau and Cuneo (2024).

In conclusion, I have argued that the grounding formulation of non-naturalism either violates purity or incurs high theoretical costs—such as accepting unexplained regularities or fundamental disjunctive facts. Putting the theoretically costly options aside, my argument entails that one of the following claims should be rejected:

- (I) Non-naturalism
- (II) Non-naturalism is best formulated with grounding
- (III) Purity

I gave some reasons for why I think purity is true, so one of (I) or (II) should be rejected. My own preference is to keep (II)—because of the motivations discussed in chapter 1—and reject non-naturalism. Alternatively, if one strongly believes in non-naturalism, she should reject the grounding formulation and provide another formulation of non-naturalism. Although I think rejecting purity is implausible, it is a theoretical option: If one is a non-naturalist who thinks that grounding is the best way to formulate non-naturalism, then she should reject purity and provide a conception of grounding that renders purity false. Thus, the above argument can be read as a metaethical puzzle that needs to be resolved in a way by rejecting one of the premises, based on one's prior theoretical commitments. In this way, I think non-naturalists and their opponents can both benefit from this argument. For non-naturalists, this argument helps them to better formulate their theory. For opponents, especially those who think that the grounding formulation is motivated, this argument refutes one important formulation of non-naturalism.

Bibliography

- Audi, Paul (2012). Grounding: Toward a Theory of the In-Virtue-Of Relation. *Journal of Philosophy* 109 (12):685-711.
- Bader, Ralf M. (2017). The grounding argument against non-reductive moral realism. *Oxford Studies in Metaethics* 12.
- Barker, Jonathan (2023). Against Purity. *Ergo: An Open Access Journal of Philosophy* 9.
- Bengson, John & Shafer-Landau, Russ & Cuneo, Terence (2024). *The Moral Universe*. Oxford University Press.
- Bennett, Karen (2011). Construction area (no hard hat required). *Philosophical Studies* 154 (1):79-104.
- Berker, Selim (2018). The Explanatory Ambitions of Moral Principles. *Noûs* 53 (4):904-936.
- Bernstein, Sara (2021). Could a middle level be the most fundamental? *Philosophical Studies* 178 (4):1065-1078.
- Bhagal, Harjit (manuscript), Strikingness: When Facts Call Out for Explanation
- Bliss, Ricki & Trogon, Kelly (2021). *Metaphysical grounding*. Stanford Encyclopedia of Philosophy.
- Cameron, Ross P. (2008). Turtles all the way down: Regress, priority and fundamentality. *Philosophical Quarterly* 58 (230):1-14.
- Chalmers, David John (1996). *The conscious mind: in search of a fundamental theory*. New York: Oxford University Press.
- Correia, Fabrice & Schnieder, Benjamin (2012). Grounding: an opinionated introduction. In Fabrice Correia & Benjamin Schnieder (eds.), *Metaphysical grounding: understanding the structure of reality*. Cambridge: Cambridge University Press. pp. 1-36.
- Dancy, Jonathan (2004). *Ethics without principles*. New York: Oxford University Press.
- Dasgupta, Shamik (2014). The Possibility of Physicalism. *Journal of Philosophy* 111 (9-10):557-592.
- deRosset, Louis (2013). *Grounding Explanations*. Philosophers' Imprint 13.
- Dunaway, Billy (2014). Supervenience Arguments and Normative Non-naturalism. *Philosophy and Phenomenological Research* 91 (3):627-655.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford, GB: Oxford University Press UK.
- Enoch, David (2019). How Principles Ground. *Oxford Studies in Metaethics* 14:1-22.
- Enoch, David & Weinshtock Saadon, Itamar (2023). Oh, All the Wrongs I Could Have Performed! Or: Why Care about Morality, Robustly Realistically Understood. In Paul Bloomfield & David Copp (eds.), *Oxford Handbook of Moral Realism*. New York, NY: Oxford University Press. pp. 434-462.
- Epstein, Brian (2018). *Social Ontology*. Stanford Encyclopedia of Philosophy.
- Fine, Kit (2001). The question of realism. *Philosophers' Imprint* 1:1-30.
- Fine, Kit (2012). Guide to Ground. In Fabrice Correia & Benjamin Schnieder (eds.), *Metaphysical grounding: understanding the structure of reality*. Cambridge: Cambridge University Press. pp. 37--80.
- FitzPatrick, William Joseph (2008). Robust ethical realism, non-naturalism, and normativity. *Oxford Studies in Metaethics* 3:159-205.
- Goodman, Nelson (1955). Fact, Fiction, and Forecast. *Philosophy* 31 (118):268-269.
- Hicks, Michael Townsen & Schaffer, Jonathan (2017). Derivative Properties in Fundamental Laws. *British Journal for the Philosophy of Science* 68 (2).
- Horwich, Paul (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.
- Lange, Marc (2013). Grounding, scientific explanation, and Humean laws. *Philosophical Studies* 164 (1):255-261.
- Leary, Stephanie (2017). Non-naturalism and Normative Necessities. *Oxford Studies in Metaethics* 12.
- Leary, Stephanie (2021). What Is Non-Naturalism? *Ergo: An Open Access Journal of Philosophy* 8.
- Lewis, David K. (1983). New work for a theory of universals. *Australasian Journal of Philosophy* 61 (4):343-377.

- Lipton, Peter (2004). *Inference to the best explanation*. New York: Routledge/Taylor and Francis Group.
- Little, Margaret Olivia (2000). *Moral Generalities Revisited*. In Brad Hooker & Margaret Olivia Little, *Moral particularism*. New York: Oxford University Press.
- Loewer, Barry (2012). Two accounts of laws and time. *Philosophical Studies* 160 (1):115-127.
- McPherson, Tristram & Plunkett, David (2022). Ground, Essence, and the Metaphysics of Metanormative Non-Naturalism. *Ergo: An Open Access Journal of Philosophy* 9 (26):674-701.
- Montero, Barbara (2001). Post-physicalism. *Journal of Consciousness Studies* 8 (2):61-80.
- Morton, Justin (2020). Grounding the normative: a problem for structured non-naturalism. *Philosophical Studies* 177 (1):173-196.
- Parfit, Derek (2011). *On What Matters: Volume Two*. Oxford, UK:
- Passinsky, Asya (forthcoming). social construction and meta-ground, forthcoming in *Philosophy and Phenomenological Research*
- Raven, Michael J. (2015). Ground. *Philosophy Compass* 10 (5):322-333.
- Ridge, Michael, "Moral Non-Naturalism", *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition)
- Ridge, Michael & McKeever, Sean (2016). Moral Particularism and Moral Generalism. *Stanford Encyclopedia of Philosophy*.
- Rosen, Gideon (2010). Metaphysical Dependence: Grounding and Reduction. In Bob Hale & Aviv Hoffmann (eds.), *Modality: metaphysics, logic, and epistemology*. New York: Oxford University Press. pp. 109-135.
- Rosen, Gideon (2016). Metaphysical Relations in Metaethics. In Tristram Colin McPherson & David Plunkett (eds.), *The Routledge Handbook of Metaethics*. New York: Routledge. pp. 151-169.
- Raven, Michael (2013). Is ground a strict partial order? *American Philosophical Quarterly* 50 (2):191-199.
- Rosen, Gideon (2017). What is a Moral Law? *Oxford Studies in Metaethics* 12.
- Rosen, Gideon (forthcoming). The explanatory role of moral principles. in P. Boghossian and C. Peacocke, (eds.), *Essays on Normative Realism*.
- Scanlon, Thomas (2014). *Being Realistic About Reasons*. Oxford: Oxford University Press.
- Schaffer, Jonathan (2009). On what grounds what. In Ryan Wasserman, David Manley & David Chalmers (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford, England: Oxford University Press. pp. 347-383.
- Schaffer, Jonathan (2010). Monism: The Priority of the Whole. *Philosophical Review* 119 (1):31-76.
- Schaffer, Jonathan (2013). Metaphysical Semantics Meets Multiple Realizability. *Analysis* 73 (4):736-751.
- Schaffer, Jonathan (2017a). Laws for Metaphysical Explanation. *Philosophical Issues* 27 (1):302-321.
- Schaffer, Jonathan (2017b). The Ground Between the Gaps. *Philosophers' Imprint* 17.
- Schaffer, Jonathan (manuscript). Naturalistic Dualism and the Problem of the Physical Correlate
- Sider, Theodore (2003). Against vague existence. *Philosophical Studies* 114 (1-2):135-146.
- Sider, Theodore (2011). *Writing the Book of the World*. Oxford, England: Oxford University Press.
- Sider, Theodore (2020). Ground grounded. *Philosophical Studies* 177 (3):747-767.
- Sorensen, R. (2023). Vagueness in *Stanford Encyclopedia of Philosophy*.
- Sud, Rohan (2018). Vague Naturalness as Ersatz Metaphysical Vagueness. *Oxford Studies in Metaphysics* 11:243--277.
- Tahko, Tuomas E. (2023). Fundamentality. *Stanford Encyclopedia of Philosophy*.
- Van Roojen, Mark Steven (2015). *Metaethics: A Contemporary Introduction*. New York: Routledge.
- Wilson, Jessica M. (2016). Grounding-based formulations of physicalism. *Topoi* 37 (3):495-512.
- Zimmerman, Dean (2010). I---Dean Zimmerman: From Property Dualism to Substance Dualism. *Aristotelian Society Supplementary Volume* 84 (1):119-150.