Limning Structure as an Epistemic Goal

Allan Hazlett

In the *Phaedrus*, Socreates sympathetically describes the ability "to cut up each kind according to its species along its natural joints, and to try not to splinter any part, as a bad butcher might do." (265e)[1] In contemporary philosophy, Ted Sider (2009, 2011) articulates the same idea:

> The world has an objective structure. [C]ommunities that choose the wrong groupings may get at the truth, but they nevertheless fail badly in their attempt to understand the world. […] There is more to be *discovered* [than the truth], more that is mandatory for inquirers to think about. The world has objective streaks in it; it has structure. […] You can state truths if you don't speak in terms of this structure, but you miss out; you are deficient along one of the main axes of cognitive success. (2009, pp. 398-9)

> [I]t's *better* to think and speak in joint-carving terms. […] The goal of inquiry is not merely to believe truly (or to know). Achieving the goal of inquiry requires that one's belief state reflect the world, which in addition to lack of error requires one to think of the world *in its terms*, to carve the world at the joints. (2011, p. 61)

As I shall put it, Plato and Sider's idea is that limning structure is an epistemic goal. My aim in this paper is to articulate and defend this idea. First, I'll articulate the notion of structure (§1) and the notion of an epistemic goal (§2). Then (§3), I'll argue against some proposals for understanding the idea that limning structure is an epistemic goal: limning structure is neither an aim of belief, nor of inquiry, nor of concept possession. Next (§4), I'll propose a framework for understanding the idea that limning structure is an epistemic goal, and defend that idea. Finally (§5), I'll discuss the possibility that there is no structure.

1 On "epistemic goals"

Epistemologists agree that true belief is an epistemic goal:[2] some argue that true belief is the only fundamental epistemic goal (Foley 1987, Alston 1989, Sosa 1991, Goldman 1999, Treanor forthcoming), and others argue that true belief is not the only fundamental epistemic goal (Zagzebski 1996, Grimm and Ahlstrom forthcoming).[3] What does it mean to say that something is an epistemic goal?

Linda Zagzebski (1996) argues that "all intellectual virtues have a motivational component that aims at cognitive contact with reality." (p. 167) This suggests the view that epistemic goals are species of cognitive contact with reality: true belief, knowledge, understanding, and the like. Stephen Grimm and Christopher Ahlstrom (forthcoming) develop this idea by defending the view that "the fundamental epistemic goal is … *accuracy*." True belief is a species of accuracy, but

[1] Alexander Nehamas and Paul Woodruff's translation, from Plato 1997.

[2] Some say that *truth* is an epistemic goal, but their concern is with *true belief*, rather than with truth, per se.

[3] A goal is **fundamental** when its status as a goal cannot be explained by the existence of other goals. Having reliable faculties, you might argue, is an epistemic goal, but its status as an epistemic goal is explained by the fact that true belief is an epistemic goal.

this formulation leaves open the possibility of other species. I'll assume that this approach to the notion of an epistemic goal is right, i.e. that *x* is an **epistemic goal** iff *x* is a species of accuracy.[4]

I take accuracy to be the same as correctness, as well as the same as "getting it right" (as Grimm and Ahlstrom put it). Essential to these notions is the possibility of being **mistaken** or in **error**. Where there is accuracy, there is the possibility of inaccuracy; where there is correctness, there is the possibility of incorrectness; where there is "getting it right," there is the possibility of "getting it wrong." To make sense of true belief as a species of accuracy, for example, we need to be able to make sense of the idea that false belief is a species of inaccuracy. For every species of accuracy, there is a corresponding species of inaccuracy.[5] The possibility for accuracy brings with it the possibility for inaccuracy, and vice versa. Take the state of having a mass of one kilogram. The only way to understand this as a species of accuracy would be to conceive of some other mass or masses as species of inaccuracy. To think of the state of having a mass of one kilogram as a species of accuracy would require thinking of the state of having a mass of two kilograms (e.g.) as a species of inaccuracy. But we do not think this: to have a mass of two kilograms is not to be mistaken or in error. And, for the same reason, to have a mass of one kilogram is not to be correct or accurate.

2 On "structure"

Recall Borges' Chinese encyclopedia, in which:

> … animals are divided into (a) those that belong to the emperor; (b) embalmed ones; (c) those that are trained; (d) suckling pigs; (e) mermaids; (f) fabulous ones; (g) stray dogs; (h) those that are included in this classification; (i) those that tremble as if they were mad; (j) innumerable ones; (k) those drawn with a very fine camel-hair brush; (l) etcetera; (m) those that have just broken the flower vase; (n) those that at a distance resemble flies. (1999, p. 231)

You might think that the dictionary's categories do not "carve the world at the joints," or, as I shall put it, that the dictionary's categories do not "limn structure."[6]

We can understand the notion of structure in propositional terms. Suppose that the following are both true propositions:

(1) that Noah has been drawn with a very fine camelhair brush.
(2) that Noah is a brown spider monkey.

You might think that knowing (2), but not knowing (1), is an instance of limning structure, on the grounds that the property of being a brown spider monkey (i.e. the property of being a member of the species called *Ateles hybridus*) is part of the world's structure, and that the property of having being drawn with a very fine camelhair brush is not. We could put this by saying that

[4] This avoids the fraught question of why epistemologists speak in terms of "goals," the answer to which is completely obscure (Sosa 2003, Kelly 2003, Grimm 2008, Hazlett 2013).
[5] These two species need not be understood as themselves species of the same genus: we understand false belief as the species of inaccuracy that corresponds to knowledge, but we need not think of false belief and knowledge as species of the same genus (although we may think of them this way).
[6] Which doesn't rule out other (perhaps related) flaws, e.g. that the dictionary's categories are neither exclusive nor exhaustive.

brown spider monkeys form a natural kind, while things that have been drawn with a very fine camelhair brush don't. (More on these assumptions in a moment.) We can give a more rigorous articulation of the idea that some propositions are "part of the world's structure." I assume a property of properties and relations (in the sequel, let "properties" cover both): that of being "structural." Structural properties are also known as **natural properties** (Armstrong 1978, Lewis 1983, 1986, Sider 2009, 2011, Mellor forthcoming) and as **sparse properties** (Lewis 1986, pp. 59-61), by contrast with non-natural properties and abundant properties, respectively. Natural properties are shared by members of a **natural kind**. As David Lewis (1983) puts it, natural properties are those "whose sharing makes for resemblance, and the ones relevant to causal powers." (p. 13) By contrast, non-natural properties are "as gruesomely gerrymandered … as you please. They pay no heed to the qualitative joints, but carve things up every which way." (Lewis 1986, p. 59)[7] The property of **being grue** – i.e. the property of being either green and observed before 1st January 2015 or blue and observed after 1st January 2015 – is intuitively a non-structural property. As Sider would put it, structural properties are part of the world's objective structure, and non-structural properties aren't. We'll adopt the following rough characterization of **structural properties**: structural properties are those whose sharing makes for resemblance and which are relevant to explanation. I leave open the question of whether the notion of structure is primitive (cf. Sider 2011, pp. 15-8), e.g. whether the notion of a structural property can be reduced to some other more primitive notion or notions.

We'll take propositions to be complexes consisting of individuals, properties, and logical operations. Assuming structural and non-structural properties, here's a first stab at the notion of a "structural proposition": if $F_1 \ldots F_n$ are the properties that (partly) constitute the proposition that p, then the proposition that p is structural iff $F_1 \ldots F_n$ are structural.

We could say something similar if we took propositions to be sets of possible worlds. Following Lewis (1986), we could identify properties with the sets of their instances and propositions with properties that are instantiated only by entire possible worlds (p. 53). Then:

> [P]ropositions can be conceived as abundant or sparse, and sets of worlds may accordingly be divided into the more and less natural. This is automatic, given the division of properties plus the identification of propositions with properties of worlds. (p. 62)[8]

We run a risk if we follow Lewis here, for if (2) is structural, then so is:

(3) that Noah has the property of being a brown spider monkey such that 2+2=4.

Since they're the same set of worlds. But you might think that the property of being a brown spider money is more structural than the property of being a brown spider monkey such that 2+2=4. For this reason, we'll stick with the conception of propositions as constituted by properties, individuals, and logical operations, inheriting the costs of such a conception of

[7] To remain as neutral as possible I've dropped two assumptions that Lewis makes: that structural properties are necessarily intrinsic and that structural properties are necessary non-disjunctive.

[8] N.b. that Lewis doesn't define the structuralness of propositions in terms of the structuralness of ordinary properties.

propositions; if you prefer the Lewisian conception, this will not make a difference in what follows.[9]

Just as properties can be divided into the structural and the non-structural, you might think that individuals and logical operations can be divided into the structural and the non-structural (cf. Sider 2011, Chapters 6, 9, and 10). Consider Gilbert, intuitively a more structural individual than the mereological fusion of Gilbert and Sullivan, or (even worse) the mereological fusion of everything that is not Gilbert (cf. Mellor forthcoming, §7). So we should amend our definition of a **structural proposition**: a proposition is structural iff the properties, individuals, and logical operations that constitute it are structural, and non-structural otherwise.[10] (However, our focus will be on structural properties in what follow.) Just as we say that it is true that p, when the proposition that p is true, we'll say that **it is structural that p** when the proposition that p is structural.

Some defenders of structure defend a hierarchy of structure, proposing a primitive relation of being more structural than (Lewis 1983, 1986, pp. 60-1). Others treat being structural as an all-or-nothing affair (Sider 2011, pp. 128-36, Mellor forthcoming). If the structuralness of properties (for example) comes in degrees, then so will the structuralness of propositions. In what follows I'll assume that being structural is all-or-nothing, but this will only be for the sake of simplicity.

We might be attracted to the idea of degrees of structuralness from another angle. Consider three true propositions:

(4) that this molecule is composed of carbon.
(5) that the mereological fusion of Gilbert and Sullivan is composed of carbon.
(6) that the mereological fusion of Gilbert and Sullivan has been drawn with a very fine camelhair brush.

Suppose that the property of being composed of carbon is structural, that the property of having been drawn with a very fine camelhair brush is non-structural, and that molecules are structural individuals, but that mereological fusions of human beings aren't. On our definition, (4) is structural, but neither (5) nor (6) is. However, you might think that (5) is *more* structural than (6) – the intuition being that (5) has at least *one* structural constituent (the property of being composed of carbon), whereas (6) has *no* structural constituents. So we may ultimately want to say that propositions admit of degrees of structuralness for this reason as well. But, again, I will assume, for the sake of simplicity, that being structural is all-or-nothing.

My articulation of the notion of a structural proposition is neutral as to where structure is to be found. This is important. Although I assumed, in explaining the notion of structural properties, that the property of being a brown spider monkey is structural and that the property of having been drawn with a very fine camelhair brush is not structural, this is not entailed by my definition of a structural proposition. And, more importantly, this is not a premise in my argument that limning structure is an epistemic goal (§4). Indeed, that argument makes no assumptions about

[9] See also Sider's (2011) formulation: "[A] proposition is joint-carving to the extent that it can be simply expressed using joint-carving concepts, given some appropriate notions of simplicity and proposition." (p. 61)

[10] This would allow us to say that disjunctive propositions are non-structural, on the assuming that disjunction is a non-structural logical operation. This jibes with the idea (which we have left open) that disjunctive properties are non-structural.

where structure is to be found. So my argument is neutral on which, if any, of these propositions are structural:

(7) that negatively charged particles repel other negatively charged particles.
(8) that bronze is partially composed of copper.
(9) that 17 is a prime number.
(10) that *Guernica* is an example of modernism.
(11) that sectarian conflict was among the causes of the Thirty Year's War.

I do not assume that all structure is physical structure, e.g. that only physical properties are structural: I leave open the possibility of, for example, chemical, biological, and psychological structure. I do not assume that all structure is "natural," in the sense of "the natural sciences," as opposed to "the social sciences": I leave open the possibility of, for example, economic and historical structure. I do not assume that all structure is the object of some scientific, as opposed to humanistic, inquiry: I leave open the possibility, for example, of structure described by art criticism. I do not assume that all structure is in "the world," where this refers to the spatiotemporal universe: I leave open the possibility of, for example, logical and mathematical structure. Finally, I do not assume that there is any structure at all: I leave open the possibility that there are no structural properties, individuals, or logical operations (cf. §5).

In line with my neutrality on the question of where structure is to be found, I leave open the question of whether structure should be identified with the fundamental (cf. Sider 2011, Chapter 7), e.g. whether structural properties should be identified with fundamental properties. Perhaps all structure is fundamental structure, but perhaps not: perhaps there is dependent, supervenient structure. The admission of physical structure (e.g. the admission of physical properties as structural) does not preclude the admission of dependent, supervenient biological structure (e.g. the admission of biological properties as structural). (As well, I do not assume that all that is fundamental is structural.)

"Structural" is a term of art, intended to distinguish, for example, those properties whose sharing makes for resemblance and which are relevant for explanation. It is preferable to "natural," which has connotations of "the natural sciences" and "the natural world," which suggests the impossibility of, e.g., economic or art critical structure. For the same reason, talk of "structure" is preferable to talk of "carving nature at the joints" or of "carving the world at the joints." Of course, you might have views about where structural properties are to be found – you might think that all structure is physical structure, for example. You'll have to evaluate my proposal that limning structure is an epistemic goal with those views in mind. But in considering the view that limning structure is an epistemic goal we need not assume any view about where structure is to be found.

3 How not to understand limning structure as an epistemic goal

Below (§4), I'll articulate and defend the view that limning structure is an epistemic goal. In this section I'll criticize some views that might be offered as articulations of the idea that limning structure is an epistemic goal: that limning structure is an aim of belief (§3.1), that limning structure is an aim of inquiry (§3.2), and that limning structure is an aim of concept possession (§3.3). I'll conclude by briefly defending my assumption (§1) that epistemic goals are species of accuracy (§3.4).

3.1 Limning structure an aim of belief

Sider (2011) writes that limning structure "is a constitutive aim of the practice of forming beliefs, as constitutive as the more commonly recognized aim of truth." (p. 61) We'll say that a **belief is structural** just in case its propositional content is structural (§2), just as a **belief is true** just in case its propositional content is true. A **constitutive aim** of Φing is one that is essential to Φing, i.e. an aim that all instances of Φing necessarily have, and that makes them (explains the fact that they are) instances of Φing. Is it then plausible that being structural is a constitutive aim of believing, in the same way that being true is?

"Aim" is ambiguous here, as between (at least) a psychological and a normative reading. On the **psychological** reading, a constitutive aim of Φing is a desire or intention that is necessary for Φing. As Sider admits (ibid.), being structural is not a constitutive psychological aim of believing: it is possible for someone to form beliefs without desiring or intending that they be structural. .

An alternative to the psychological reading is a **normative** reading, on which a constitutive aim of Φing is a standard by which Φing may be evaluated. So Sider writes that being structural is "a standard by which beliefs and believers may be evaluated." (Ibid.) But beliefs and believers may be evaluated relative to all kinds of standards. What is needed, to show that being structural is a *constitutive* normative aim of believing, is some sense in which the standard of being structural flows from the essential nature of belief. Now the reason people have for thinking that the standard of being true flows from the essential nature of belief comes from some observations about conscious belief and about conscious belief formation, or "doxastic deliberation":

> If in full consciousness I could acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e. as something purporting to represent reality. (Williams 1973, p. 148)
>
> In order for a propositional attitude to be an attitude of *belief*, it cannot represent itself as wholly unaccountable to truth or evidence. […] A self-representation of certain of one's attitudes as 'aiming at' truth is *partially constitutive* of belief. (Railton 1997, pp. 57-9)
>
> Truth is not an optional end for first-personal doxastic deliberation. [W]ithin the first-person deliberative perspective … the question *whether to believe that p* seems to collapse into the question *whether p is true*. [O]ne cannot settle on an answer to the question *whether to believe that p* without taking oneself to have answered the question *whether p is true*. (Shah 2003, p. 447)

These phenomena, so the argument goes, reveal something about the essential relationship between belief and being true: that truth is a constitutive aim of believing. But no analogous phenomena exist that might suggest the same relationship between belief and structure. It is possible to acquire, "in full consciousness," a belief irrespective of whether its content is a structural proposition; one can represent one's beliefs to oneself as unaccountable to structure, and being structural *is* an optional end for first-personal doxastic deliberation: I can ask whether to believe that p while explicitly setting aside the question of whether it is structural that p. Therefore, it is not the case that being structural is a constitutive aim of believing, in the way that being true is said to be. To be sure, evaluating beliefs vis-à-vis the standard of being structural is *one* way we may evaluate them. But this standard does not flow from the essential nature of belief.

2.2 Limning structure as an aim of inquiry

Sider says that thinking about structure is "mandatory for inquirers" (2009, p. 398) and that limning structure is "the goal of inquiry." (2011, p. 61) Again, we should understand this as a claim about the a constitutive aim (§3.1): in this case, about a constitutive aim of inquiry. On a psychological reading, our claim says that inquiry requires a desire or intention to limn structure. But this is implausible. You can engage in inquiry without seeking to limn structure: my inquiry into whether the mereological fusion of Gilbert and Sullivan ever wore two pairs of grue trousers merely seeks the truth about a certain question, and does not seek to "carve the world at the joints." The goal of limning structure is optional for inquiry, and in a way that the goal of true belief isn't: you can't coherently claim to be inquiring unless you're trying to arrive at a true belief (or beliefs). "Inquiry" that doesn't seek truth is no inquiry at all, but rather pretended or feigned inquiry. Inquiry that doesn't seek structure may be uninteresting or insignificant (cf. §4.2), it may even be foolish or a waste of time. But it does not thereby cease to be inquiry. The psychological aim of *some* instances of inquiry is limning structure, but some instance merely seek true belief.

What about a normative reading of the claim that limning structure is a constitutive aim of inquiry? Perhaps we might say that inquiry that doesn't lead to limning structure is thereby a failure, qua inquiry. But this too is implausible. Consider my inquiry about G+S's trousers (§3.1), where I explicitly set aside the goal of limning structure. My inquiry is successful if I acquire the true answer to my question, and there is no sense in which my inquiry has failed, in virtue of failing to limn structure. No such negative evaluation is mandated by the essence of inquiry. Some inquiries seek to limn structure, but others seek merely the truth about some question, and in the latter case, limning structure is not an appropriate evaluative standard.

3.3 Limning structure as an aim of concept possession

Sider (2011) asks us to imagine a community who think of the word in terms of 'grue'-like predicates, as opposed to 'green'-like predicates, and argues that "it is almost irresistible to describe these people as *making a mistake*." (p. 2) But this is not the mistake of false belief. Their mistake "is that they've got the wrong concepts," and so "[t]hey're carving up the world incorrectly." (Ibid.)

Is this plausible? We require an alternative conception of limning structure (cf. §3.1) to make sense of this idea; let's say that a **concept C limns structure** iff C picks out a structural property, individual, or logical relation (§2). Consider the view that the possession a concept C is correct iff C limns structure, and incorrect otherwise. This is not plausible: we are not mistaken or in error just in virtue of possessing the concept of being grue (§1). What then (if anything) differentiates us from the members of Sider's imagined community? Perhaps this: in addition to possessing the concept of being grue, we also possess the concept of being green, whereas the members of Sider's imagined community possess only the former concept. So perhaps we should say that the set of concepts Σ possessed by a person or community is correct iff a sufficient number of the members of Σ limn structure, and incorrect otherwise. This still is not plausible. Imagine someone such that none of her concepts limn structure, but who is indifferent to this. We tell her that her concepts are all 'grue'-like, and she admits never having had an opinion on the matter of whether her concepts picked out structural properties, individuals, and logical operations. She is a kind of pragmatist, indifferent to structure, whose concepts, by her own lights, are designed for other purposes. There seems no mistake, no incorrectness here. If you desired or intended to possess concepts that limn structure, but didn't, we would count this as a kind of failure. And it would be a mistake to employ a set of concepts that don't limn structure while thinking that they limn structure. It would be a mistake is to think you're limning structure, when you aren't. But this would just be good old-fashioned incorrectness in belief.

3.4 Epistemic goals and accuracy

I have assumed that epistemic goals are species of accuracy (§1). As defenders of truth as an aim of belief remind us, beliefs are correct iff true, and incorrect otherwise. This is why true belief is a species of accuracy. When you believe a proposition that isn't true, you are wrong, mistaken, and in error. No analogous claim can be made when it comes to being structural. It is not the case that beliefs are correct iff structural, and incorrect otherwise. It is not incorrect to believe that this emerald is grue; when you believe that this emerald is grue, you are neither wrong, nor mistaken, nor in error. For this reason, structural belief is not a species of accuracy, and therefore not an epistemic goal.

Below, I'll propose an alternative conception of limning structure (§4). Might we, instead, drop the assumption that epistemic goals are species of accuracy? It seems to me that we must make something like this assumption, to preserve the idea that epistemic goals enjoy some kind of *value*. Assume, again (§1), that the property of being a brown spider monkey is structural and that the property of having been drawn with a very fine camelhair brush is not, and suppose that you believe (1) and I believe (2). Both beliefs, by hypothesis, are true. What's better about my belief, by contrast with yours? You might point out that there is a kind of *similarity* between my belief and the world's structure, that is not present between your belief and the world's structure. My belief employs the concept of being a brown spider monkey, which corresponds, per our assumption, to a structural property. Your belief enjoys no such distinction. But what sort of distinction is this? Imagine that you and I both attend Sarah's birthday party, and both wear baseball caps, but that my baseball cap is of the same type as the baseball cap that Sarah is wearing. There is a kind of similarity between my cap and hers, which your cap does not enjoy. But there is nothing *good* about this. And there is, intuitively, nothing *epistemically* good about this, nothing that resembles the paradigm epistemic goal of true belief. In true belief, there is not mere similarity between mind and world, there is accuracy. (Spelling out the difference between mere similarity and accuracy, of course, is a central and difficult issue in the philosophy of mind.) The situation at Sarah's party might be different if you and I were *trying* to wear caps that matched Sarah's cap. In that case, it might make sense to speak of accuracy and inaccuracy, or at least of success and failure. But recall that we have seen that structural belief is not a species of accuracy; belief does not aim at being structural (§3.1). To preserve the idea that epistemic goals have some kind of value, we should not treat mere *similarity* between mind and world as sufficient for something's being an epistemic goal. We should require *fit* between mind and world, i.e. we should retain the assumption that all epistemic goals are species of accuracy.

4 How to understand limning structure as an epistemic goal

We need to explore an alternative framework, if we are to defend the idea that limning structure is an epistemic goal. In this section I provide such an alternative (§4.1), discuss the relationship between structure and significance (§4.2) and between structure and understanding (§4.3), and consider the idea that limning structure is a fundamental epistemic goal (§4.4).

4.1 Limning structure as a species of theorizing

The problem with our previous attempts to articulate the idea that limning structure is an epistemic goal (§3) was that neither belief, nor inquiry, nor concept possession plausibly aims at being structural. Structural belief, for example, is not a species of accuracy. So what we need is something such that its structural instances are members of a species of accuracy.

To fulfill this need, I'll describe a propositional attitude, distinct from mere belief, which does aim at being structural (in the relevant sense of "aim"). This attitude is nameless in English, but

we can get a grip on it through two characterizations. I'll call this propositional attitude **theorizing that p**, and I shall propose a conception of limning structure as a species of theorizing that p. But it is essential to keep in mind that I am not attempting conceptual analysis on the notion of theorizing. "Theorizing" is here used as a term of art, whose meaning is entirely determined by the following two characterizations.

First characterization of theorizing that p: just as belief aims at truth, theorizing aims at truth and structure. The nature of the "aim" of belief is controversial (Wedgwood 2002, Shah 2003, Steglich-Petersen 2006, Sosa 2009), but it is agreed that your belief that p is correct iff it is true that p, and incorrect otherwise. So we can say that:

> Your **theorizing that p is correct** iff it is both true and structural that p, and incorrect otherwise.

Just as being true is a constitutive standard of correctness for belief (§3.1), being true and structural is a constitutive standard of correctness for theorizing.[11] Suppose you theorize that Noah is a brown spider monkey (§1). Given the assumption that the property of being a brown spider money is structural, your theorizing is correct. Suppose now that you theorize that Noah has been drawn with a very fine camelhair brush. Given the assumption that the property of having been drawn with a very fine camelhair brush is not structural, your theorizing is incorrect, because it is not structural. Finally, suppose that you theorize that Noah is a great blue heron. Given the assumption that the property of being a great blue heron is structural, your theorizing is incorrect, just because it is false.

Second characterization: just as knowledge is the achievement of truth, such that we can think of false belief as failed knowledge, understanding is the achievement of structure, such that we can think of false or non-structural theorizing as failed understanding. It is instructive to consider the idea that **knowledge is apt belief** (Sosa 2007, Greco 2010; cf. Zagzebski 1996). On this view, true belief is understood as a success, and knowledge, a species of true belief, is the corresponding achievement. Just as an archer can succeed in hitting her target, and can be credited with an achievement when such success manifests her ability in the domain of archery, a believer can succeed in believing the truth, and can be credited with an achievement – knowledge – when such success manifests her ability in the epistemic domain, and, in particular, her ability vis-à-vis the goal of true belief. Analogously, we can think of **understanding as apt theorizing**. True and structural theorizing is a success, and understanding, a species of true and structural theorizing, is the corresponding achievement. Understanding, in this sense, is success that manifests someone's ability vis-à-vis the goal of true and structural theorizing. How does this help us understand the notion of theorizing? Just as we could understand belief as that which (metaphorically) aspires to knowledge (cf. Williamson 2000), we should understand theorizing as that which (metaphorically) aspires to understanding. In something like the way that knowledge is "factive," understanding is "factive": both knowledge and understanding preclude inaccuracy, incorrectness, mistake, and error (cf. §4.3). Belief is the "non-factive" correlate of knowledge, and theorizing is the "non-factive" correlate of understanding. Belief is what both knowledge and falsehood have in common; theorizing is what both understanding and misunderstanding have in common.[12] There is more that needs to be said about understanding, in the sense articulated (§4.3). But we shall move on towards an articulation of the notion of limning structure.

---

[11] Just as true belief is belief that p where it is true that p (§3.1), **structural theorizing** is theorizing where it is structural that p.

[12] If understanding belief in terms of knowledge is an instance of "knowledge-first epistemology," then the present approach is an instance of "understanding-first epistemology."

What is the relationship between theorizing and believing? We could think of theorizing as a species of believing. Where mere believing aims at truth, theorizing aims at truth *and* structure. We could offer a reductive account of theorizing in terms of belief, on which S theorizes that p iff S believes that p and S believes that it is structural that p. But we need not think of theorizing as involving any such belief about structure. The unsophisticated beliefs of children and non-human animals suggest that believing that p does not require believing that it is true that p; likewise we might not want to require believing that it is structural that p, for theorizing that p. Essential to the notion of theorizing is the idea that, just as belief represents its content as true, theorizing represents its content as structural. But belief that p does not represent its content as true in virtue of the believer believing *that it is true that p*. The believer simply believes that p. Likewise, so the argument goes, theorizing that p does not represent its content as structural in virtue of the theorizer believing *that it is structural that p*. Just as it is intrinsic to belief that its content is represented as true, it is intrinsic to theorizing that its content is represented as structural. However, this is not material to our understanding of the notion of theorizing; we might opt for the reductive account of theorizing in terms of belief suggested above. If theorizing is understood as a species of belief, then although being structural is not a constitutive aim of belief (§3.1), we might say that being structural is a constitutive aim of a species of belief, namely, theorizing. Alternatively, we might not want to think of theorizing as a species of belief. My defense of limning structure as an epistemic goal is neutral on the question of the relationship between theorizing and believing.

With the notion of theorizing articulated, by appeal to the two characterizations provided above, we can articulate the notion of limning structure:

> **Definition:** An instance of **theorizing that p limns structure** iff it is both true and structural that p.

Given this definition, the conditions for an instance of theorizing being structure-limning and the conditions for an instance of theorizing being correct are the same. Therefore, structure-limning theorizing is a species of accuracy. Given the assumption that epistemic goals are species of accuracy (§1), it follows that structure-limning theorizing is an epistemic goal. Just as true belief is an epistemic goal, in virtue of being a species of accuracy, structure-limning theorizing is an epistemic goal, in virtue of being a species of accuracy.[13]

How can this conclusion be resisted? Below (§4.3), I'll consider some objections to what I said above about understanding. I think the most important worry that must be addressed, in defense of structure-limning theorizing as an epistemic goal, is the worry that the propositional attitude of theorizing, in some sense, does not really exist. Now this raises some interesting and difficult questions about what it is for a propositional attitude to really exist. We should avoid these debates if we can. My purposes will be served if the case for the real existence of theorizing is as good as the case for the real existence of belief. So what convinces us that the propositional attitude of belief really exists, in the relevant sense?

We have, I shall argue, at least two reasons for thinking that belief really exists. First, we sometimes inquire about whether p, i.e. we set out to discover the truth about whether p, and we are familiar with belief as the state we find ourselves in when we think we have answered our question, when we think that our inquiry has come to an end. We know what it is to be curious about whether p, and we know what it is to have that curiosity satisfied, i.e. to form a belief

[13] True and structural theorizing is a species of limning structure, but this leaves open the possibility of other species of limning structure, e.g. the non-propositional representation of structure.

about whether p. Second, we are familiar with cases of knowledge, but we are also familiar with cases that we can think of as cases of failed knowledge. The virtuous inquirer with misleading evidence, for example, does not know that p, because it is not true that p, and so she merely believes, perhaps with justification, that p. Belief is real because knowledge is real; once we are familiar with the reality of knowledge, we can easily recognize the reality of belief.

There are analogous reasons for thinking that theorizing really exists. First, we sometimes inquire not merely with the aim of discovering the truth, but also with the aim of discovering structure. And we are familiar with the state we find ourselves in when we think we have succeeded. Imagine a team of chemists investigating a mysterious substance. "What is it?," they wonder. There are many true answers to their question: that it is a mysterious substance being investigated by some chemists, that it is a substance such that 2+2=4, that it is the fourth substance they have investigated since lunchtime. But none of these is the answer they seek. They seek an answer such as: that it is a compound of carbon, hydrogen, and fluoride. They seek an answer that is both true *and* structural. So we know what it is to be curious about structure, and we know what it is to have that curiosity satisfied – that's what it is to theorize that p.

Second, we are familiar with cases of understanding, but we are also familiar with cases that we can think of as failed understanding or misunderstanding. Consider Stephen Grimm's (2010) case of the spilled coffee: a woman sitting in a café spills her coffee; I think it's because her leg brushed against the table, which caused the cup to tip over, while you think it's because a shaman cast a spell on the cup, causing it to tip over. At most one of us understands why the coffee spilled; our explanations are incompatible and at most one is correct. At most one of us has understanding, the other has failed understanding or misunderstanding. I have introduced the jargon "theorizing" to capture the state that a person is in when she fails to understand or misunderstands. We can see that theorizing is real, because understanding is real; once we are familiar with the reality of understanding, we can recognize the reality of theorizing.

Finally, I will not say much here in response to objections that are based on the mistaken assumption that my articulation of the notion of theorizing was an exercise in conceptual analysis. Here's an example: "A child can have understanding, but children don't theorize. Therefore, understanding is not a species of theorizing." But this objection assumes an ordinary sense of "theorize," on which it is obvious that children do no theorize. My argument, in defense of limning structure as an epistemic goal, uses "theorizing" as a term of art. Given my articulation of the notion of theorizing, the question of whether children theorize is a question about (i) whether children have propositional attitudes that are correct iff true and structural, and (ii) whether children have understanding. The notion of theorizing was articulated, above, in terms of its correctness condition and in terms of its relationship with understanding. It seems to me that children *do* theorize, in my sense, because they are curious about structure, and do sometimes come to acquire understanding, which is a species of theorizing. But, in any event, I will have nothing more to say about objections that fail to treat "theorizing" as a term of art.

4.2 Structure and significance

I have argued that limning structure is an epistemic goal, by arguing that structure-limning theorizing is an epistemic goal (§4.1). This conclusion jibes with the idea, defended by many epistemologists, that "important truth," rather than mere truth, is an epistemic goal (Haack 1993, Goldman 1999, Zagzebski 2003, Alston 2005, David 2005). The reason for this is that being structural constitutes a species of importance. In this section I defend that idea.

Consider the truth about the exact number of grains of sand on some randomly selected portion of my floor. There seems to be no worth or value when it comes to believing this "trivial" truth (Sosa 2003, Grimm 2008), by contrast with "important" truths, like (so you might think) the truth about whether the Standard Model is correct. But what does it mean to say that some truth is "important"?

It is not the case that there is one correct answer to this question, because "important" is ambiguous in this context. There are several kinds of true proposition that we might want to pick out by speaking of "important" truths; in other words, as I shall put it, there are several species of importance (cf. Baril 2010, p. 222). It would therefore be foolish to attempt to evaluate competing "accounts" of importance. But we can examine the several species of importance.

First, consider the notion of the interesting (Goldman 1999), where the proposition that p is **interesting to S** iff S wants to know whether p, i.e. iff S is curious about whether p. Interestingness is essentially subject-relative. We might speak here of importance *from the perspective* of some would-be believer.

Second, consider the notion of the eudaimonically significant (Zagzebski 2004, Baril 2010, Grimm 2011), where the proposition that p is **eudaimonically significant** iff believing that p would (or could) contribute, causally or constitutively, to one's living well, or to one's living a good life, or to one's flourishing. We might speak here of importance *for the sake of* some would-be believer. (Although eudaimonic significance is not essentially subject-relative, the possibility of developing a relativistic conception of eudaimonic significance is left open; on such a view, a proposition might be eudaimonically significant for me, but not for you, in virtue of the fact that believing that proposition would contribute to my living well, but not to your living well.) Anne Baril (2010) writes that some beliefs "are finally valuable, valuable for their own sakes, *as part of* a (non-derivatively) valuable whole," namely, the good human life of a particular person. Some propositions are more suitable than others to serve as contents for such "worthy" beliefs – and in this sense they have more eudaimonic significance. Alternatively, eudaimonic significance can be understood in terms of instrumental value. Stephen Grimm (2011) argues that knowledge of significant truths may "help to guide our … project of living well." (p. 532) The value of such knowledge derives from the fact that "we are primarily interested in … figuring out how to live well," and "if we can predict and possibly control how the world will unfold then we can try to steer it in a way that promotes our well-being." (p. 534)

Third, consider the notion of **generality**: W.D. Ross (1930) writes that "[k]nowledge of mere matters of fact . . . seems to be worth much less than the knowledge of general principles, or of facts depending on general principles." (p. 139) Here we might speak of the *epistemic* importance of general truths (cf. Treanor forthcoming). Knowledge of generalities may be finally or instrumentaly valuable vis-à-vis the good life, but such knowledge also seems valuable apart from any contribution it makes to the knower's living well.[14]

However, I think there remains a species of importance, indeed a species of epistemic importance, that we have not considered so far. I have in mind, of course, the notion of **being structural** (§2). As with generality, it may be that knowledge of structural truths (cf. §4.3) is

[14] This idea could be challenged. Zagzebski (2004) writes that "[t]here is no independent domain of epistemic value," (p. 353) and Grimm (2011) argues that "it begins to look like our 'purely epistemic interests' [may] in fact be practical at heart," and that the distinction between the epistemic and the non-epistemic "is more complicated than has often been supposed." (p. 535) For more on this issue, see Hazlett 2012 and Hazlett 2013, §9.3.

finally or instrumentally valuable vis-à-vis the good life, but such knowledge also seems valuable apart from that.[15] The (epistemic) importance of structural truths grounds the (epistemic) importance of certain questions, topics, and areas of inquiry. Structural truths enjoy a species of importance that non-structural truths do not enjoy.

This idea provides an appealing diagnosis of our original examples of "trivial" and "important" truths. The proposition that the Standard Model is correct, we assume, is a structural proposition: the Standard Model employs concepts that are at least meant to pick out structural properties. By contrast, the proposition that there are *n* grains of sand on such-and-such randomly selected portion of my floor does not seem like a structural proposition: the property of being on such-and-such randomly selected portion of my floor is not a structural property. These two things explain the relative importance of the former proposition and the relative triviality of the latter. Being structural, in other words, is a species of importance.

Objection: The outcome of tonight's match between Arsenal and Chelsea is very important for me. But the proposition that Arsenal wins the match (for example) is not structural.

Reply: There are several species of importance, and being structural is only one of them. That being structural is a species of importance does not imply that non-structural properties are not important in every sense, only that they are not important in one sense. The importance of the proposition that Arsenal wins the match may be down to the fact that said proposition is interesting to you, or to its eudaimonic significance. (Recall that it is possible to develop a relativistic conception of eudaimonic significance.)

Objection: There are important truths to be found outside the domain of physics. Consider historical, sociological, economic, or aesthetic knowledge. True propositions in these domains are important, but not structural.

Reply: My argument is neutral about where structure is to be found (§2). I have not assumed that all structure is physical structure, nor that the truths of history, sociology, economics, and aesthetics are not structural. My argument implies only that there is a species of importance that such truths enjoy only if they are structural.

Objection: Knowledge of non-structural propositions can be important as a means to (epistemically) valuable ends. Consider the proposition that Noah is the Emperor's favourite animal, which you implied (§2) was non-structural. If this proposition is true, it might be important, in virtue of the fact that knowing this truth is essential to our understanding of history or politics – that Noah was the emperor's favourite animal explains why the Emperor is abdicating the throne, for example.

First Reply: We should distinguish between the structuralness, and therefore the importance, of the following two propositions:

(12) that Noah is the Emperor's favourite animal.
(13) that the Emperor is abdicating the throne because Noah is his favourite animal.

We can concede that (13) is structural, and therefore important, while denying that (12) is structural, on the assumption that causation or explanation is a structural relation. Better (cf. §2):

[15] Sider (2011) speaks of the "*epistemic* value of joint-carving inquiry" (p. 62, my emphasis, and see the title of his §4.5, p. 61) and Nick Treanor (forthcoming) of "*epistemic* contact with the world." (my emphasis)

we can concede, on this assumption, that (13) is more structural than (12), and thus it might be the case that (13) is relatively important, and (12) relatively unimportant. We can respect the intuition that (13) is (relatively) important, without conceding the (relative) importance of any (relatively) non-structural proposition, e.g. (12).

Second Reply: The same objection applies to the other species of importance, such as interestingness and generality. Knowledge of an uninteresting truth might be a means to knowledge of some interesting truth, and knowledge of some particular matter of fact might be a means to knowledge of some generality. Think here of the knowledge of particular items of data in scientific research, which if all goes well are means to scientific knowledge. Are such propositions – e.g. that such-and-such plant grew *n* millimetres on such-and-such a day – important, in the relevant sense, in virtue of their instrumental relation to valuable ends? They are not, in any event, made interesting or general in virtue of their instrumental relation to knowledge of interesting or general truths. It seems to me that what this shows, if anything, is that there is another species of importance (e.g. instrumental value vis-à-vis knowledge of important truths), not that the species of importance canvassed are not really species of importance.

4.3 Structure and understanding

In my articulation of the notion of theorizing (§4.1), I appealed to the idea that understanding is apt theorizing. This section further elaborates and defends that idea. More precisely, I'll argue that apt theorizing is a species of understanding. This takes up Sider's (2009) idea that communities that fail to limn structure "fail badly in their attempt to understand the world." (p. 398)

The word "understanding" is highly ambiguous, and (so it seems to me) it would be a mistake to attempt to unify all that we pick out with "understanding" with one "account" of understanding. Consider:

- Linguistic understanding, as when you understand the meaning of a word.
- Knowing the truth conditions of a proposition, as in "Before I studied physics, I didn't understand the proposition that the wave function collapses."
- Explanatory understanding, as when you understand why something is the case.
- Acknowledgement or acceptance, as in "I understand that the jury has reached a decision" or "She understands that she has the right to an attorney."
- General knowledge, thus Ross (1930) writes that "knowledge of general principles [is] what we might call insight or understanding as opposed to mere knowledge." (p. 139)
- Understanding of some part or aspect of the world, e.g. the movement of the planets, the nature of matter, the origin of species, etc.
- Understanding of a topic, as in "She understands chemistry."
- Understanding of a theory, as in "She understands the theory of special relativity."
- Understanding of a person, as "I understand Janice."
- Understanding of some kind of thing, activity, or domain, as in "She understands cars" or "She understands basketball."
- "Understanding-how," as in "I understand how cars work" or "I understand how species evolved."

This makes clear that there are several kinds of cognition that we pick out by speaking of "understanding"; in other words, as I shall put it, there are several species of understanding. I'll

argue that apt theorizing is a species of understanding. Does this suffice as a defense of the idea, to which I appealed above (§4.1), that understanding is apt theorizing? Yes, for I did not there need to appeal to the claim that every instance of every species of understanding is an instance of apt theorizing. Compare the analogous idea that knowledge is apt belief. We do not understand that idea as the claim that every instance of every species of knowledge is an instance of apt belief. There are species of knowledge – practical knowledge ("know-how"), acquaintance – which we understand as distinct from the species identified with apt belief. The same, mutatis mutandis, when it comes to the idea that apt theorizing is a species of understanding.

Let's take a closer look at apt theorizing. I'll then defend the view that apt theorizing is a species of understanding. Your believing that p is apt when it is true that p and the fact that your belief is true manifests your ability vis-à-vis the goal of true belief. Analogously, your theorizing that p is apt when it is both true and structural that p and the fact that your theorizing is both true and structural manifests your ability vis-à-vis the goal of true and structural theorizing. Just as apt belief requires **truth-directed intellectual virtues**, understood as reliable abilities or faculties vis-à-vis the goal of true belief, apt theorizing requires **truth-and-structure-directed intellectual virtues**, understood as reliable abilities or faculties vis-à-vis the goal of true and structural theorizing. The apt believer must manifest her responsiveness or sensitivity to truth; the apt theorizer must manifest her responsiveness or sensitivity to truth and structure.

Recall the analogy of the archer (§4.1), in connection with apt belief. In connection with apt theorizing, we may usefully imagine an archery contest involving both stable and moving targets. To strike a stationary target earns some acclaim, but to strike a moving target earns more. A successful hit is the analogue of true belief, a successful hit on a moving target is the analogue of true and structural theorizing. The analogue of apt theorizing, then, is a successful hit on a moving target that manifests the ability to hit such a target. Such ability requires more than the ability to hit a mere target (i.e. the ability required for apt belief), for someone might excel at hitting stationary targets but be inept at hitting moving targets. The ability required to hit a moving target includes the ability required to hit mere targets, and more. Likewise, the intellectual virtues required for apt theorizing include the virtues required for apt belief, and more.

Apt theorizing is a propositional attitude. Moreover, it is a "factive" propositional attitude (cf. §4.1). But it is "factive" in two ways. First, it is "factive" in the familiar way: someone aptly theorizes that p only if it is true that p. Second, it is "factive" in the sense that someone aptly theorizes that p only if it is structural that p. Consider the chemist investigating a mysterious substance (§4.1). In as much as she is attempting to limn structure, her attempt may fail even if she arrives at a true account of the substance, because her account is not structural. There may be a true division of material things into earth, air, fire, and water, and someone might truly theorize that this mysterious substance is made of earth. But if she is a chemist seeking to limn structure, then she has failed unless she arrives at a conclusion that is not only true but also structural, e.g. that this mysterious substance is a compound of carbon, hydrogen, and fluoride. (As elsewhere, I here make an inessential assumption about where structure is to be found.) Another way to put all of this is that apt theorizing is a species of limning structure.[16]

[16] It is sometimes argued that understanding, unlike knowledge, is not "factive" (Elgin 1996, pp. 122-9, 2006, 2009, Zagzebski 2001; cf. Kvanvig 2003, Riggs 2003, Grimm 2010, Pritchard 2010). I do not think there is any good argument for this view. The case of someone who understands astrology is sometimes offered as a case of "non-factive" understanding, since astrological theory is false. But here we must distinguish between two cases. On the one hand, the sociologist of astrology theorizes about astrological theory, but there is nothing "non-factive" about her

On the assumption that knowledge is apt belief, and if theorizing is a species of belief (§4.1), then apt theorizing is a species of knowledge. For apt theorizing must be apt in two ways: with respect of the truth of the relevant proposition, and with respect to the structuralness of the relevant proposition.

So I maintain that apt theorizing is a species of understanding. You might ask: "What species of understanding is it?" It seems to me, however, that this question is misplaced. Of course, when it comes to the idea that apt belief is a species of knowledge, we can clarify: the view maintains that apt belief is propositional knowledge. We cannot quite say that apt theorizing is propositional understanding, for "propositional understanding" is ambiguous. "Propositional understanding" could refer to acknowledgement or acceptance (above), or to knowing the truth conditions of a proposition (above). Apt theorizing is to be identified with neither of these species of understanding. But it is not the case that, for every species of X, we can say *which* species of X said species is. Propositional knowledge is an example. "Which species of knowledge is it?" This question is misplaced. Once we have identified propositional knowledge, as a species of knowledge, there is no further question of *which* species it is. I maintain the same, mutatis mutandis, about apt theorizing. It is a species of understanding, but there is no further question of *which* species it is. But this raises an obvious question: why think of apt theorizing as a species of understanding? Here are four reasons.

First, the idea that apt theorizing is a species of understanding exploits an independently plausible idea from metaphysics – structural properties – and applies it in epistemology. The idea of structural properties is "put to work" in the theory of understanding. The idea that apt theorizing is a species of understanding, therefore, contributes to the overall simplicity and elegance of our philosophical picture.

Second, the idea that apt theorizing is a species of understanding can vindicate the Aristotelian idea that knowledge of causes and explanations is a species of understanding (cf. Grimm 2006, forthcoming), on the assumption that causal and explanatory relations are structural. Suppose that you theorize that p because q. On the assumption that the relation picked out by "because" is structural, and on the assumption that it is true that p because q, and on the assumption that your theorizing is apt, you have understanding. If the assumption that causal and explanatory relations are structural is plausible, we can treat explanatory understanding, i.e. understanding why p, as a species of apt theorizing.

Third, the idea that apt theorizing is a species of understanding explains why mere knowledge of true and structural propositions can fall short of understanding. There are cases in which someone aptly believes that p, but does not aptly theorize that p. In one kind of case, the person aptly believes that p, but does not *theorize* that p. Consider someone who memorizes truths for the purposes of winning the pub quiz: she learns from a reliable source that Aston Villa won the European Cup in 1982, that Daniel Craig is married to Rachel Weisz, and that $F = ma$. Her concern is just with getting the answers right; she has no concern for limning structure. It is not

---

understanding of astrology, since her thoughts about astrological theory are perfectly accurate thoughts. On the other hand, the astrologer theorizes about the relationship between the positions of the planets and human affairs, but she does not enjoy "non-factive" understanding, because she does not understand the relationship between the positions of the planets and human affairs, in virtue of the fact that her thoughts about said relationship are quite inaccurate. The case of scientific understanding through idealizations and models is sometimes offered as a case of "non-factive" understanding, since said idealizations and models are, strictly speaking, false. But if this speaks against the idea that understanding is "factive," then also it speaks against the idea that knowledge is "factive."

plausible to attribute the propositional attitude of theorizing to her. She may well aptly believe, and thereby know, that $F = ma$, which is a true and structural proposition, but she does not aptly theorize, and therefore does not understand that proposition.

There are also cases in which someone aptly believes that p, and theorizes that p, but does not *aptly* theorize that p. These are cases in which someone has (and manifests) the relevant truth-directed intellectual virtues, but either does not have or does not manifest the relevant truth-and-structure-directed intellectual virtues. Consider someone comes to know a structural proposition on the basis of testimony, as in the case of a child who comes to know, on the basis of testimony, that her house burned down because of faulty wiring (Pritchard 2010, p. 82). She does not yet understand why the house burned down, for while her testimonial uptake may have been impeccable vis-à-vis the goal of true belief, it may not manifest ability vis-à-vis the goal of true and structural theorizing. Compare the fire chief from whom she got the information: *she* understands why the house burned down, because her knowledge that the house burned down because of faulty wiring manifests her ability vis-à-vis the goal of true and structural theorizing. Her understanding manifests her reliability when it comes to limning structure, her ability to discover the real causal structure behind the fire, by seeing how oxygen, electricity, etc., brought the fire about. The child, although she manifests her reliability when it comes to believing on the basis of testimony, lacks any reliability in the domain of limning structure.[17]

Fourth, the idea that apt theorizing is a species of understanding can vindicate the idea that understanding requires a non-doxastic "grasp" of structure (Zagzebski 2001, Riggs 2003, Grimm 2010, forthcoming). First, apt theorizing is a species of theorizing that p, rather than mere belief that p, and theorizing that p requires the representation of the proposition that p as structural (§4.1). As I have urged, someone can believe that p, where it is true and structural that p, but fail to represent the proposition that p as structural. Understanding, on the present proposal, requires the representation of p as structural. This provides one sense in which understanding requires a non-doxastic "grasp" of structure. Second, apt theorizing requires the manifestation of truth-and-structure-directed intellectual virtues (see above). Thus, on the present proposal, the person who understands must not merely come to theorize that p, where it is true and structural that p, but must also be responsive or sensitive to structure. This provides another sense in which understanding requires a non-doxastic "grasp" of structure.

To sum up: understanding, in the present sense, is distinctive both in virtue of its content (structural propositions) and in virtue of the demand for the manifestation of truth-and-structure-directed intellectual virtues. This captures both the idea (from Aristotle) that understanding is metaphysically distinctive, as well as the idea (more on which below) that understanding is psychologically distinctive.

In this section I have defended the idea that apt theorizing is a species of understanding. Does the plausibility of my thesis, that limning structure is an epistemic goal (§4.1), depend on the plausibility of the idea that apt theorizing is a species of understanding? The status of true belief as an epistemic goal doesn't depend on any claim about knowledge, nor on any particular theory of knowledge. Consider other species of knowledge, such as practical knowledge ("know-how") and acquaintance, which are not special cases of true belief. The status of true belief as an epistemic goal does not depend on any connection between true belief and any particular species of knowledge. So I am skeptical that the status of limning structure as an epistemic goal depends on any connection between limning structure and any particular species of understanding.

[17] What if she had? In that case, she could acquire understanding on the basis of testimony, as when (for example) an experienced colleague of the fire chief comes to understand why the house burned down, on the basis of the fire chief's testimony.

4.4 Limning structure as a fundamental epistemic goal

Is limning structure a *fundamental* epistemic goal (§1)? If so, true belief is not the only fundamental epistemic goal. Imagine that I believe propositions (1) and (2). My beliefs are equals when it comes to the truth of their contents, and so they are equals when it comes to satisfying the goal of true belief. Suppose I know both (1) and (2). Despite this parity, it seems that there is some sense in which knowing (2) is a greater achievement than knowing (2), some sense in which the former instance of knowledge is better than the latter. If so, we shall need an explanation of this: the fact that true belief is an epistemic goal cannot do the job.[18] And if this is right, true belief is not the only fundamental epistemic goal. It seems to me that this argument succeeds, but I won't defend it here. In this paper I have defended the idea that limning structure is an epistemic goal, not that it's a fundamental epistemic goal.[19]

5 What if there is no structure?

You might object that there is no structure, or that there is no real distinction between propositions in terms of structure, because all true propositions are structural. Realists like Plato, Sider, and Lewis respond to Borges' encyclopedia (§2) with the intuition that some categories are better than others: some "carve the world at the joints" and others don't. You might draw the opposite conclusion, as Foucault does in *The Order of Things* (1970): that what Borges' encyclopedia shows is that there is no structure, that all categories are socially constructed, that classification is always relative to a conceptual scheme, where no scheme can claim more objective validity than any other.

Metaphysicians have offered various reasons for believing in structure. Realists about structure appeal to the "Moorean fact" that some things are of the same type as other things (Armstong 1980, Lewis 1983), to the need to objectively ground the distinctness of individual things (Mellor forthcoming), to the utility of the notion of structural properties in articulating other metaphysical notions that are in need of articulation (Lewis 1983, Lewis 1986, p. 63), to the fact that the existence of structure explains the platitudinous values of scientific discovery, causal explanation, and objectivity (Sider 2011, pp. 62-5), or to the intuition that some categories obviously get things right, and others obviously get things wrong (ibid. pp. 1-2). These arguments are controversial and could be resisted. So let us suppose the anti-realists are right, and that there is no structure.

This does not threaten my thesis that limning structure is an epistemic goal. I have argued that theorizing that p is correct only if the proposition that p is structural (§4.1), and that apt theorizing, which is a species of correct theorizing, is a species of understanding (§4.3). If no proposition is structural, then theorizing is never correct, and (this species of) understanding is impossible.[20] But this is compatible with my thesis. For to say that limning structure is an epistemic goal is to say that limning structure is a species of accuracy (§1). If there is no

[18] Note, as well, that we cannot say that seeking structural beliefs is a means to acquiring true beliefs (in the way that we might say that seeking justified or reliable beliefs is a means to acquiring true beliefs), because there is no reliable connection between being structural and being true. For every structural truth, there is a structural falsehood, namely, its negation. (As well, for every structural truth there is a non-structural truth in the neighborhood: everything that is an electron is also either an electron or drawn with a very fine camelhair brush.)

[19] On a related and relevant question, see Treanor forthcoming.

[20] This conclusion follows for any species of understanding that requires the "grasp" of structure. So even if you reject the idea, for example, that understanding is a propositional attitude (§4.3), you may have reason to agree with this conclusion.

structure, then limning structure is impossible, and a certain species of accuracy is impossible. This does not undermine the status of limning structure as an epistemic goal. Compare truth: if no proposition is true, then a certain species of accuracy, namely, true belief, is impossible, and knowledge, a species of true belief, is therefore impossible. But true belief remains a species of accuracy, and therefore an epistemic goal.

The argument that the non-existence of structure precludes the possibility of understanding could be resisted by positing a "non-factive" conception of understanding (§4.3). Similarly, one could resist the point (above) about truth and knowledge by positing a "non-factive" conception of knowledge. The ancients knew that the earth was flat, in the sense that, according to their knowledge, the earth was flat. Likewise, the ancients understood the solar system, in the sense that they had an understanding of the solar system, on which the sun revolves around the earth. Alternatively, the argument could be resisted by adopting a view of structure on which it is, in some sense, "subjective." Suppose there is socially constructed structure, e.g. structure that exists as a result of our way of thinking about the world. One might then be able to agree with the idea that all categories are socially constructed, but retain the possibility of understanding. Compare the view that truth is socially constructed, and the corresponding conception of knowledge (Rorty 1985, Sosa 1993). However, it seems to me that there is a *kind* of knowledge that requires the existence of "objective" truth, and a *kind* of understanding that requires the existence of "objective" structure. My articulation of the notion of apt theorizing (§3.1, §4.1) was meant to correspond to this kind of understanding.

I said that limning structure is an epistemic goal, even if limning structure is impossible. Can we really speak usefully in terms of a "goal," when achieving said "goal" is impossible? In some cases, we obviously can: I cannot touch my toes, but my doctor prescribes trying to touch my toes, as a way of stretching my muscles. And a goal may *be* impossible, but not recognized as such by those who seek it. The view that there is no structure does not obviously conflict with the fact that many people seek to limn structure. Even if there is no structure, it seems like what many inquirers, both scientific and non-scientific, are *trying* to do is "carve the world at the joints." Even if they are doomed to fail, this is still their goal. So I do not think it is misleading to speak of limning structure as an epistemic "goal," even if limning structure is impossible.

We have been considering the idea that there is no structure. But the same thoughts apply, mutatis mutandis, when it comes to local anti-realism about structure. Where there is structure to limn, then (this species of) understanding will be possible, but where there is none, (this species of) understanding is ruled out. Anti-realism about biological structure, for example, rules out the possibility of biological understanding (of this species). We can see now that the epistemological question of the scope of human (this specie of) understanding depends crucially on the metaphysical question of where structure exists, whether in physics, in chemistry, in psychology, in economics, or in art history.[21]

Bibliography:

Alston, W. (1989), *Epistemic Justification* (Cornell University Press).

----- (2005), *Beyond "Justification": Dimensions of Epistemic Evaluation* (Cornell University Press).

[21] Thanks to Ted Sider, Nick Treanor, Stephen Grimm, John Greco, and Anne Baril for conversations on this material, and to audiences at the Bled Philosophical Conference on *Knowledge, Understanding, and Wisdom* in 2011, at the Serious Metaphysics Group at the University of Cambridge in 2012.

Armstrong, D.M. (1978), *Universals and Scientific Realism* (Cambridge University Press).

----- (1980), "Against 'Ostrich' Nominalism: A Reply to Michael Devitt," *Pacific Philosophical Quarterly* 61, pp. 440-9.

Baril, A. (2010), "A Eudaimonist Approach to the Problem of Significance," *Acta Analytica* 25:2, pp. 215-41.

Borges, J.L. (1999), "John Wilkins' Analytical Language," in his *Selected Non-Fictions* (Viking, 1999), pp. 229-32.

Brady, M.S. (2009), "Curiosity and the Value of Truth," in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford University Press), pp. 265-83.

David, M. (2001), "Truth as the Epistemic Goal," in Steup, M. (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford University Press), pp. 151-69.

----- (2005), "On 'Truth is Good'," *Philosophical Books* 46:4, pp. 292-301.

Elgin, C. (1996), *Considered Judgment* (Princeton University Press).

----- (2006), "From Knowledge to Understanding," in S. Hetherington (ed.), *Epistemology Futures* (Oxford University Press), pp. 199-215.

----- (2009), "Is Understanding Factive?," in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford University Press), pp. 322-9.

Feldman, R. (1988), "Epistemic Obligations," *Philosophical Perspectives* 2, pp. 235-56.

Foley, R. (1987), *The Theory of Epistemic Rationality* (Harvard University Press).

Foucault, M. (1970), *The Order of Things* (Random House).

Goldman, A. (1999), *Knowledge in a Social World* (Oxford University Press).

Greco, J. (2010), *Achieving Knowledge* (Oxford University Press).

Grimm. S. (2006), "Is Understanding a Species of Knowledge?," *British Journal for the Philosophy of Science* 57, pp. 515-35.

----- (2008), "Epistemic Goals and Epistemic Values," *Philosophy and Phenomenological Research* 77:3, pp. 725-44.

----- (2009), "Epistemic Normativity," in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford University Press), pp. 243-64.

----- (2010), "Understanding," in S. Bernecker and D. Pritchard (eds.), *The Routledge Companion to Epistemology* (Routledge), pp. 84-94.

---- (2011), "What is Interesting?," *Logos & Episteme* 2:4, pp. 515-42.

----- (forthcoming), "Understanding as Knowledge of Causes," *Synthese*.

Haack, S. (1993), *Evidence and Inquiry: Towards Reconstruction in Epistemology* (Blackwell).

Hazlett, A. (2012), "Non-Moral Evil," *Midwest Studies in Philosophy* 36, pp. 18-34.

----- (2013), *A Luxury of the Understanding: On the Value of True Belief* (Oxford University Press).

Johnston, M. (1996), "Is the External World Invisible?," *Philosophical Issues* 7, pp. 185-98.

Kelly, T. (2003), "Epistemic Rationality as Instrumental Rationality: A Critique," *Philosophy and Phenomenological Research* 66:3, pp. 612-40.

Langton, R. (1998), *Kantian Humility: Our Ignorance of Things in Themselves* (Oxford University Press).

Lewis, D. (1983), "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61, pp. 343-377.

----- (1986), *On the Plurality of Worlds* (Blackwell).

----- (2008), "Ramseyan Humility," in Braddon-Mitchell, D., and Nola, R. (eds.), *Conceptual Analysis and Philosophical Naturalism* (MIT Press), pp. 203-21.

Lynch, M.P. (2004), *True to Life: Why Truth Matters* (MIT Press).

----- (2009), "The Value of Truth and the Truth of Values," in A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value* (Oxford University Press), pp. 225-42.

Mellor, D.H. (forthcoming), "Nature's Joints: A Realistic Defence of Natural Properties," in D. Oderberg (ed.), *Classifying Reality* (Blackwell).

Pritchard, D. (2010), "Understanding," in D. Pritchard, A. Millar, and A. Haddock, *The Nature and Value of Knowledge: Three Investigations* (Oxford University Press), pp. 67-89.

Railton, P. (1997), "On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action," in G. Cullity and B. Gaut (eds.), *Ethics and Practical Reason* (Oxford University Press), pp. 53-79.

Rorty, R. (1985), "Solidarity or Objectivity?," in R. Rorty, *Objectivity, Relativism, and Truth: Philosophical Papers, Volume 1* (Cambridge University Press), pp. 21-34.

Shah, N. (2003), "How Truth Governs Belief," *Philosophical Review* 112, pp. 447-82.

Sosa, E. (1991), *Knowledge in Perspective* (Cambridge University Press).

----- (1993), "Epistemology, Realism, and Truth," *Philosophical Perspectives* 7, pp. 1-16.

----- (2003), "The Place of Truth in Epistemology," in M. DePaul and L. Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford University Press), pp. 155-79.

----- (2007), *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (Oxford University Press).

----- (2009), "Knowing Full Well: The Normativity of Beliefs as Performances," *Philosophical Studies* 142, pp. 5-15.

Sider, T. (2009), "Ontological Realism," in D. Chalmers, D. Manley, and R. Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology* (Oxford University Press), pp. 384-423.

----- (2011), *Writing the Book of the World* (Oxford University Press).

Steglich-Petersen, A. (2006), "No Norm Needed: On the Aim of Belief," *Philosophical Quarterly* 56, pp. 499-516.

Treanor, N. (forthcoming), "Trivial Truths and the Aim of Inquiry," *Philosophy and Phenomenological Research*.

Velleman, J.D. (2000), *The Possibility of Practical Reason* (Oxford University Press).

Wedgwood, R. (2002), "The Aim of Belief," *Philosophical Perspectives* 16, pp. 265-97.

Williams, B. (1973), "Deciding to Believe," in his *Problems of the Self* (Cambridge University Press), pp. 136-51.

Williamson, T. (2000), *Knowledge and Its Limits* (Oxford University Press).

Zagzebski, L. (1996), *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowing* (Cambridge University Press).

----- (2001), "Recovering Understanding," in M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford University Press), pp. 237-51.

----- (2003), "The Search for the Source of Epistemic Good," *Metaphilosophy* 34:1/2, pp. 12-28.

----- (2004), "Epistemic Value and the Primacy of What We Care About," *Philosophical Papers* 33:3, pp. 353-77.