

# Seek the Joints! Avoid the Gruesome!

## Fidelity as an Epistemic Value\*

Peter Finocchiaro

November 18, 2021

### Abstract

A belief is valuable when it “gets it right”. This “getting it right” is often understood solely as a matter of truth. But there is a second sense of “getting it right” worth exploring. According to this second sense, a belief “gets it right” when its concepts accurately match the way the world is objectively organized – that is, when its concepts are joint-carving, or have *fidelity*. In this paper, I explore the relationship between fidelity and epistemic value. While many philosophers (especially metaphysicians) acknowledge fidelity’s value, they overlook just how much it might disrupt our understanding of epistemic value. To tease out this disruption, I draw on the Jamesian balance between seeking the truth and avoiding the false. A similar balance must be struck both within the pursuit of fidelity itself (“seeking the joints” and “avoiding the gruesome”) as well as between the pursuit of fidelity and the pursuit of truth. I then give an argument against the claim that truth is the higher epistemic good.

**Keywords:** Epistemic Values, Fidelity, Fundamentality, Veritism, Hierarchy of Goods, Metameta-physics, William James

## Introduction

Some beliefs are more valuable than others. These differences in value may be traced to a multitude of sources. But I want to focus on just one particular source. Some beliefs are more valuable others insofar as the former are right and the latter are wrong.

Traditionally, this value in “getting it right” has been called epistemic value. Many a fight has started over the precise nature of epistemic value and its relationship to other forms of value. Some philosophers, for instance, contest that epistemic value is entirely reducible to some other value, like moral value. Along those lines, some would say that “getting it right” is valuable only insofar

---

\*I want to thank Andrew Brenner, Ben Cross, Kate Finley, Asher Jiang, Michael Longenecker, Matt Lutz, Zee Perry, Juha Saatsi, Ru Ye, the Australian National University Epistemology reading group, and the audience at my session of the Thinking About Theory Choice workshop for their comments, advice, and encouragement in the development of this paper. Special thanks belong to Liz Jackson, who helped me to see the broader significance of my initial ideas. Without her, this paper would be much more narrow and boring. Finally, thanks to Timothy Perrine for indulging my constant need to talk through the issues I was encountering. Without him, this paper would be much more confused – as would I.

as it contributes to the well-being of the individual or their society. Others philosophers balk at such reductions. They think that “getting it right” is valuable on its own, independent of any connection it might have to the moral domain.

Rather than engage in these long-standing fights, I’d like to start a new one. I’d like to complicate the notion of “getting it right” and its attendant epistemic value.

Many philosophers, myself included, understand “getting it right” as a correspondence between the world and what is believed of the world: to “get it right” is to have an accurate sense of how things are. It’s easy to take this notion of accurate correspondence as equivalent to that which we call truth: a belief “gets it right” just in case it is true.

But I think that doing so overlooks a crucial part of the story. I think that there is a second sense of “getting it right” that is not merely a matter of truth. Consider the following example. While catching up over a pint, Zhang San shows off her new jade bracelet to Li Si. The bracelet is green. In the spirit of Goodman (1955), let’s say that *x* is grue iff either *x* is green and observed before 3000AD or blue and not observed before 3000AD. Because the bracelet has been observed before the year 3000AD, it is also grue. The particular year on which this definition rests is of course entirely arbitrary. Furthermore, as the definition suggests, the truth conditions for beliefs about grue bracelets do not demand the existence of some property, *being grue*. Instead, the beliefs could be true so long as the bracelet instantiates the property *being green* at the appropriate time (or instantiates the property *being blue* at the appropriate time). Nevertheless, it would be better if Li Si came to believe that the bracelet was green rather than believe that the bracelet was grue. The first belief would “get it right” in a way that the second would not. This is so because the concepts that would be employed in believing the former more accurately match the way the world is objectively organized.

This second sense of “getting it right” should sound familiar to many philosophers, especially those who are familiar with contemporary metaphysics. But, as it currently stands, no one seems to fully appreciate its connection to epistemic value. If “getting it right” is partially determined by the concepts one employs, then epistemic value is determined along at least two dimensions. The first dimension (the one that has received plenty of attention) covers value with respect to truth. The second dimension (the one that seems to me underappreciated) covers value with respect to using the right concepts – what I call *fidelity*. If epistemic value is determined along separate

dimensions, then we are confronted with a number of pressing questions, including questions about how we should understand the relationship between the two values as well as questions about how we should balance our pursuit of them.

In this paper, I will flesh out these questions and offer some preliminary answers to them. I begin in section 1 by more precisely detailing this second sense of “getting it right” that I call fidelity. Then, in section 2, I explore the ways in which fidelity might have epistemic value. Supposing that it does, we must consider how to weigh its value, especially in relation to the more familiar value of truth. To that end, in section 3 I motivate the claim that the value of fidelity cannot be explained in terms of the value of truth. Then, in section 4, I draw on the Jamesian balance between seeking the truth and avoiding the false to evaluate some of the ways we may weigh the value of fidelity in the pursuit of “getting it right”.

## 1 Fidelity: A Primer

There are two senses in which a belief can “get it right”. The first sense of “getting it right” is truth. A belief “gets it right” when it is true and does not “get it right” when it is false. The second sense is what I call *fidelity*. A belief “gets it right” when it has fidelity and does not “get it right” when it lacks fidelity.

Fidelity is not some newfangled philosophical posit, though it has at times been presented as such. Philosophers have acknowledged it, in one form or another, for millennia. Plato, in the voice of Socrates, spoke admirably of carving each kind along its natural joints while avoiding any splintering indicative of bad butchery. The medieval scholastics fiercely debated whether the categories (e.g. substance, accident) accurately divide the world or rather reflected conceptual or linguistic biases. More recently, David Lewis stressed the importance of interpreting the words and thoughts of others as tracking the comparatively more natural properties of the world, and Theodore Sider has discussed the importance of employing ideology that matches the world’s metaphysical structure.<sup>1</sup>

These philosophers are all, more or less, articulating the sense of “getting it right” that I call fidelity. (Personally, I favor the nuances behind Sider’s articulation of fidelity. But that’s not

---

<sup>1</sup>For more on these different articulations of fidelity see *Phaedrus*: 265e; Pasnau (2011); Lewis (1983): 370–377; Sider (2011): i, 1–3.

important here.<sup>2</sup>) In each case, the posited correspondence is a form of representing the world. What that representation precisely consists of differs. But, loosely speaking, the representation is organizational. Because the correspondence is to be evaluated by the accuracy of its representation, its success is essentially sensitive to how the world is organized. Furthermore, the way this correspondence “gets it right” is not equivalent to the way truth “gets it right” – not obviously so, anyway. To use an imperfect, but hopefully illuminating, slogan: matters of truth are decided by ‘yes’s and ‘no’s, whereas matters of fidelity are decided by ‘these’s and ‘those’s.

As I understand it, fidelity is ultimately a feature of concepts. A concept is a mental representation. So, too, is a belief. But these two kinds of representation are not independent. A belief is constituted by concepts and what a belief represents depends on the concepts that constitute that belief. Consequently, the fidelity of a belief is derived from the fidelity of the concepts that constitute it. Li Si’s belief that the bracelet is green is partially constituted by the concept *bracelet* and the concept *green*. The fidelity of her belief therefore depends on the fidelity of those two concepts. In this way, fidelity is like other decomposable features of complex wholes.

Of course, not everyone shares my representation-based understanding of beliefs and concepts. Some philosophers may understand belief in terms of something non-representational, like patterns of behavior or causal relations. There’s no need to fight about the nature of beliefs right now. I am interested in the mental activity that occurs in someone’s mind when, for example, they just remembered that they left the stove on. I will call that mental activity a belief, though not much rests on whether or not it “really is” a belief. This mental activity is partially constituted by other mental activities. Similarly, though, I don’t want to insist at all costs that these smaller activities “really are” concepts. I only assume that there are such representational mental activities.

That’s the basic idea behind fidelity. Beyond that, there are more complex questions to address. For instance: does fidelity come in degrees, and if so, how fine-grained is the measurement? I will assume, without argument, that fidelity does come in degrees and that there are two factors that figure into this gradation. First, beliefs typically employ more than one concept, and so a belief can have greater or lesser fidelity insofar as some of its concepts match the world and others do not. Second, each individual concept might have greater or lesser fidelity insofar as they more accurately match the world than rival concepts. I will also assume that there is some threshold

---

<sup>2</sup>For more on these nuances, see Finocchiaro (2021b); Finocchiaro (2021a).

above which the fidelity of a belief is “close enough”. Every belief above this threshold is a *fidelic*, or *joint-carving*, belief. At the limit, a fidelic belief might just be a belief that has the maximum degree of fidelity. Similarly, I will assume that there is some threshold below which the fidelity of a belief is “too far off”. Every belief below it is a *non-fidelic*, or *gruesome*, belief. Finally, I will assume that these two thresholds are one and the same. As a consequence of this “bivalence” about fidelity, every belief is either fidelic or non-fidelic. These assumptions might very well be wrong. But they will help to streamline my discussion of fidelity’s epistemic value.

In the example above, the belief that the bracelet is green has greater fidelity than the belief that the bracelet is grue. But fidelity need not be restricted to the natural world. The belief that murder is wrong has greater fidelity than the belief that murder is wrong on Wednesdays because morality does not (ordinarily) depend on the day of the week. Though the day-specific moral belief is true, it misrepresents the universality of the moral facts. Furthermore, some beliefs about social categories like gender, race, and sexual orientation might have greater fidelity than others, even if these categories are social constructions. Fidelity is applicable to any belief about any topic insofar as that belief reflects a way that the world is organized. Nevertheless, in what follows I will not rely on any specific assumptions about which of our concepts reflect how the world is organized and which do not.

Of course, plenty of philosophers deny that the world is organized.<sup>3</sup> While such a denial does not automatically constitute a denial of fidelity, it does trivialize it. Consider the analogous situation for truth. Many philosophers hold that truth is a kind of correspondence, a belief being true just in case it “gets it right” with respect to what’s going on in the mind-independent world. Those who deny that there is a mind-independent world do not thereby deny this kind of correspondence. But they do trivialize it and are consequently likely to replace the correspondence with some other relation, perhaps even calling that relation “truth”.

As the analogy to truth shows, I need not fully respond to such deniers before exploring the potential value of fidelity, no more than I would need to fully respond to anti-realists before exploring the potential value of truth. Thus, in this paper I will assume without further argument that belief

---

<sup>3</sup>See, for instance, Goodman (1955); Rorty (2000). It is important to contrast these “full blown” deniers of organization from more moderate deniers like Thomasson (2014). Moderate deniers do not deny that the world is in some ways organized. Rather, they deny that this organization can do all the work that some philosophers attempt to make it do. Many deny, for example, that fidelity is a notion that applies to logical concepts like that of quantification. Moderate deniers still face the epistemological issues raised in this paper.

has this second sense of “getting it right” and that in a good number of cases its success or failure is a non-trivial matter.

That being said, it is worth saying one more thing about the fidelity of belief. The details here will be important in the subsequent exploration of its value.

Often, fidelity is given a linguistic gloss. As Sider sometimes puts it, “a fact is fundamental when it is stated in joint-carving terms” (Sider (2011): i). But this linguistic gloss is inappropriate when it comes to the question at hand. Belief is a mental phenomenon. We sometimes express our beliefs with words, and those words might (or might not) carve the world at its joints. But so far as the fidelity of a belief goes, it is fidelity with respect to our concepts – more specifically, our primitive concepts.

What makes a concept *primitive*? Here, I want to be as accommodating as I can be and avoid needlessly presupposing any specific theory of concepts. So, in lieu of a full account of conceptual primitiveness, I offer a partial characterization that should suffice for the purposes of this paper.

First, a rough definition: a concept is primitive relative to an agent just in case (i) the agent regularly employs the concept, and (ii) the agent predominantly employs the concept to the exclusion of rival concepts. According to this definition, neither the source of the concept nor the time of its acquisition determines whether or not it is primitive (though, to be sure, these factors often shape which concepts agents have primitively and which they do not). The definition also does not explicitly acknowledge what the agent could do with a concept, what they would ideally do, and so on. The definition is about what the agent in fact does with their concepts. Furthermore, by acknowledging the regularity of the concept’s employment, the rough definition allows for change in which concepts are primitive relative to the agent.

Some examples may help to further illuminate what I mean by conceptual primitiveness. Zilin is a mathematician who accepts the successor-based reduction of the natural numbers and the arithmetic functions. This morning, when she calculated the number of eggs to buy for breakfast, she multiplied the number of people eating by the number of eggs each wanted. The cognitive story never involved succession. Assuming that this is what regularly occurs in Zilin’s head, multiplication is a primitive concept for her, even though she could and sometimes does “reduce” it to succession. Adam is an American who now lives in New Zealand. At first, whenever he read the weather forecast, he needed to mentally convert from Celsius into Fahrenheit before he could determine

whether he should wear a jacket. Now, decades after his relocation, he can determine what to wear without converting. At first, only Fahrenheit was conceptually primitive for Adam. But over time it has been replaced by Celsius. Zilin and Adam meet at an academic conference. While describing her background, Zilin shares a story about the time she and her sister visited their uncle in Chengdu. Adam shares a similar story about the time he and his younger sister visited their father's older brother in San Francisco. Of course, both can speak English and (Mandarin) Chinese. Because they are practicing their non-native languages, Zilin shares her story in English, using the English words 'sister' and 'uncle', while Adam shares his story in Chinese, using the Chinese words '妹妹' and '伯伯'. Despite what Zilin says in English, she conceives of her sister as her 妹妹 and her uncle as her 伯伯. Despite what Adam says in Chinese, he conceives of his 妹妹 as his sister and his 伯伯 as his uncle. Their primitive conceptions of family are different. Adam's primitive conceptions are based on the gender of his family members, whereas Zilin's primitive conceptions are based on gender as well as birth order.

One belief has a higher degree of fidelity than another to the extent that the primitive concepts employed in that belief better represent the way the world is actually organized. What someone believes can come apart from what they say. So, too, can the fidelity of someone's belief come apart from the fidelity of how they express that belief. Similarly, the fidelity of someone's belief can come apart from the fidelity of the proposition they believe. Some theories of propositions do not distinguish between propositions about younger sisters and propositions about 妹妹. On such theories, the content of Zilin's beliefs may be the same as the content of Adam's beliefs. But clearly something different is going on in their minds. Zilin employs the comparatively more detailed Chinese familial concepts even though she utters English familial words and even if she expresses a comparatively coarse-grained proposition. This conceptual employment, whatever it is, is what is important in the evaluation of a belief's fidelity.<sup>4</sup>

## 2 The Epistemic Value of Fidelity

Once we acknowledge fidelity as a second sense of "getting it right", we face a wide range of questions regarding the precise nature of its epistemic value. In this section, I'll articulate some potential

---

<sup>4</sup>Cf. Dasgupta (2018); McDaniel (2017)

answers to these questions. Though some of these answers seem to me to be more plausible than others, I think they all deserve to be considered more fully.

Truth and fidelity are only two among many epistemic features that are sometimes regarded as epistemically valuable. Some others include: accuracy, adequacy, aptness, coherence, creditability, explanatoriness, generality, intellectual virtue (and all of the more specific intellectual virtues like caution, courage, humility, and perseverance), justification, knowledge, proper functioning, reliability, systematicity, and understanding, as well as social features like trust, epistemic justice, and epistemic autonomy. (I'm sure there are others that I've missed, and for that I apologize.)

Which of these features are actually valuable and why they are valuable is a matter of dispute.<sup>5</sup> But almost everyone agrees that there are *some* connections between them. Let me briefly introduce some terminology to help articulate these connections, and then apply that terminology to the case of fidelity.

Some beliefs we value for their own sake. Philosophers disagree as to what feature makes these beliefs valuable for their own sake. But many think that, for at least some of these beliefs, we value them at least in part because they are true. In this sense, truth has *final epistemic value*.<sup>6</sup> Similarly, some beliefs are epistemically disvaluable because they are false; falsity has final epistemic *disvalue*. But consider an epistemic feature like reliability. If some belief of mine comes from a reliable belief-forming process, then that belief is probably true. Because reliability directs me toward that which has final epistemic value – in this case, toward the truth – reliability itself has a sort of indirect epistemic value. Call this indirect value *instrumental* epistemic value. Many philosophers, though, think that the instrumental value of reliability exhausts its epistemic value; they deny that reliability has final epistemic value. If they're right, we can say that reliability *only* has instrumental epistemic value, or has *mere* instrumental epistemic value.

Many of the above-mentioned epistemic features are indirectly connected to each other in the way that reliability is indirectly connected to truth. But some of them have an even more intimate connection. For instance, part of what it is to know something is to have a belief that is true.

---

<sup>5</sup>Among other disputes is the dispute about the foundations of epistemic normativity. Roughly, the dispute is about whether epistemic normativity is ultimately based on the acquisition of epistemic goods, or, instead, ultimately based on the observation of epistemic duty. While I talk in a way that implicitly endorses the first position, much of what I say should be salvageable on the second. For more on this dispute, see Berker (2013).

<sup>6</sup>Sometimes, this sort of value is called *intrinsic* value. There's no need to fight about the terminology here. I choose the word 'final' instead of 'intrinsic' to leave open the theoretical possibility that epistemic value is extrinsic or relational. See Korsgaard (1983).

Plausibly, we value knowledge for its own sake; it has final epistemic value. But some philosophers (e.g. David (2005)) think that its final value is entirely comprised of the final value of the truth that it includes. For them, knowledge has only *non-basic* final epistemic value, whereas truth has *basic* final epistemic value. Other philosophers (e.g. Greco (2003)) think that some other component of knowledge contributes final epistemic value. And, in theory, a philosopher could argue that knowledge has basic final value that is greater than the basic final value of its elements.<sup>7</sup>

The instrumental/final distinction and the basic/non-basic distinction are importantly different. I won't attempt to give a full account of the difference here, because I do not think that the details matter for what I want to say in this paper. Suffice to say that the connection between the instrumental and the final seems to be something like a causal, correlative, or probabilistic connection and the connection between the basic and the non-basic seems to be something like a constitutive, mereological, or grounding connection. That gloss leaves a lot to be developed, but it should be informative enough to put the distinctions to fruitful use.

To understand the overall epistemic value of fidelity, we need to identify the source of that value. Does fidelity have instrumental value? If so, is that all it has? Or does it also have some sort of final value? If so, is its final value basic or non-basic?

Fidelity might have instrumental value. For instance, because fidelity tracks objective similarities, maybe fidelity promotes more successful inductive inferences. Or maybe fidelic beliefs lead to better theories, which lead to better experimental designs, which lead to more and better discoveries. Maybe not. Maybe, because of our psychological limitations, thinking in fidelic terms actually increases the number of errors we make. (Imagine trying to describe the fall of the Roman empire using the concepts of quantum physics!) Unlike reliability, fidelity lacks analytically-specified instrumental connections. Questions about the instrumental value of fidelity are therefore empirical to a degree that renders me unable to do anything more than speculate. Since such speculation has little value, I leave these questions to future empirical studies. I simply note that the claim that fidelity has instrumental epistemic value is not *obviously* wrong.

Thankfully, I can do a bit more than speculate about fidelity's final value. Many philosophers seem to think that it does have final value. One of the most prominent proponents of this position is Sider. As he forcefully puts it:

---

<sup>7</sup>Jonathan Kvanvig discusses this position in Kvanvig (2003): 140–156.

[I]magine a universe that is entirely full of fluid. A plane divides the universe into two halves, one in which the fluid is uniformly red, the other in which the fluid is uniformly blue. . . Now imagine a group of people who encounter this universe, but accord no special status to the dividing blue-red plane. Instead of thinking of the universe as divided into the red and blue halves, they think of it as being divided in half by a different plane. . . And they do not use predicates for red and blue. Instead, they have a pair of predicates that they apply uniformly within the two regions separated by their dividing plane. These predicates. . . cut across the predicates ‘red’ and ‘blue’. The regions to the left of the [dividing] line they call “bred”; the regions to the right they call “rue”. It is almost irresistible to describe these people as making a mistake. But they’re not making a mistake about where the red and blue regions are, since they make no claims about red or blue. And they make no mistakes when they apply their own concepts. The regions that they call “bred” are indeed bred, and the regions they call “rue” are indeed rue. The problem is that they’ve got the wrong concepts. They’re carving the world up incorrectly. By failing to think in terms of the red/blue dividing plane, they are missing something. Although their beliefs are true, those beliefs do not match the world’s structure (Sider (2011): 2–3).

According to Sider, the mistake these people make is not a merely instrumental one. When they divide the world into bred regions and rue regions, they make a mistake with respect to what they believe now, regardless of their mistaken beliefs’ consequences. It would have been better if they had used concepts that divided the world into blue regions and red regions. If they had done so, their beliefs would have had greater fidelity. In virtue of that greater fidelity, their beliefs would have had greater final epistemic value.

Yet even if red/blue beliefs have greater basic final value than bred/grue beliefs and even if that difference is because the former have a greater degree of fidelity, it does not follow that fidelity itself has basic final value. I see at least two alternative explanations for the difference in value.

First, the red/blue beliefs might constitute a complex epistemic feature that itself has basic final value. Whether or not complex epistemic features have basic final value is a controversy I don’t intend to settle here. But suppose that they do. Arguably, fidelity is a constitutive element of some complex epistemic features. Consider the state of understanding. When an individual comes to understand something, they come to acquire something that goes beyond mere truth. Perhaps in some cases this acquisition involves fidelity.<sup>8</sup> To understand something is, in part, to possess a genuine explanation of that thing. These explanations may be causal – as when a chemist explains the reactive dispositions of the elements – or they may be non-causal – as when a philosopher explains that goodness is what the gods love. These explanations only reflect a

---

<sup>8</sup>Cf. Hazlett (2017): Section 4.3.

genuine understanding when they are stated in joint-carving terminology. In other words, fidelity is a constitutive element of understanding. So, if understanding has basic final value, then the final value of fidelity might be merely non-basic.

Second, the red/blue beliefs might form an *organic unity* that is more epistemically valuable than the organic unity formed by the bred/rue beliefs. As Moore (1903) puts it: “The value of a whole must not be assumed to be the same as the sum of the values of its parts” (79). Moore suggests that the (non-epistemic) value of someone admiring something beautiful is greater than the admiration and the beauty taken on their own. But the value of an organic unity need not be an “amplification” of the value of its parts. The value of a whole may be positive even if one or all of its parts lacks value. Perhaps, for example, musical rests lack (non-epistemic) value even though music is generally more (non-epistemically) valuable with rests than without. We might say something similar about the value of complex epistemic features and their constituents.<sup>9</sup> But the point stands even for beliefs that do not constitute complex epistemic features. Even if a belief that is both true and fidelic is more valuable than a competing belief that is true and non-fidelic, it does not follow that fidelity itself has value, let alone basic final value.

This fact raises a troubling possibility. “Getting it right” is epistemically valuable and true beliefs are more valuable when they are also fidelic. But what if that’s simply because such beliefs are a specific way of getting the truth right? In other words, what if the value of “getting it right” derives entirely from the basic value of truth? To that possibility I now turn.

### 3 Veritism Unmotivated

As I mentioned earlier, many philosophers think that truth is a basic epistemic value. Many of those philosophers go even further and think that truth is the only basic epistemic value. Such philosophers – *veritists* – think that the apparent epistemic value of all other epistemic features can be either explained away or explained in terms of truth. Kristoffer Ahlstrom-Vij, for example, argues that the state of understanding is epistemically valuable only as an instrumental means to further truths; any apparent surplus value derives either from the practical value of having

---

<sup>9</sup>Cf. Zagzebski (2004): 193–196.

systematized beliefs or from the prudential value of feeling a sense of closure and completeness.<sup>10</sup>

At first glance, fidelity seems amenable to a similar treatment. Even if fidelity is a distinct sense of “getting it right”, its apparent value might be explained in terms of truth. Look at how contemporary metaphysicians talk about the importance of fidelity: they talk about the importance of acquiring fundamental truths, the importance of securing inductive inferences, the importance of establishing substantive disputes where one disputant affirms and the other denies the same claim, and so on. In short, their talk looks suspiciously truth-centric.

But looks can be deceiving. Below, I sketch five initially plausible strategies a veritist may use to explain the value of fidelity in terms of the value of truth. Each strategy faces problems without straightforward solutions. A firmly committed veritist might find a way to solve these problems. Yet I think the problems are significant enough to show that the value of fidelity is not straightforwardly reducible to the value of truth. As it stands, then, the veritist’s commitment is unmotivated.

The first strategy takes the value of fidelity to be merely instrumental to the basic value of truth. If someone uses joint-carving concepts, they are more likely to acquire true beliefs and avoid false beliefs. Similarly, if someone uses gruesome concepts, they are less likely to acquire true beliefs and avoid false beliefs. As I mentioned before, there is nothing inherent to the notion of fidelity that analytically guarantees such instrumentality. So in the absence of empirical evidence there is no general reason to think that fidelity is instrumental to the truth. However, there is a connection between fidelity, objective similarities, and inductive inferences. Inductive inferences succeed because the future resembles the past; they succeed because in general objects that are objectively similar remain objectively similar. Thus, the green objects we observe now will be like the green objects in the distant future, whereas the grue objects we observe now might not be like the grue objects in the distant future. Consequently, inductive inferences made on the basis of joint-carving concepts are more likely to lead to true beliefs; similarly, inductive inferences made on the basis of gruesome concepts are liable to lead to false beliefs.

Unfortunately, the connection between fidelity and inductive inferences is not quite that simple.

---

<sup>10</sup>See Ahlstrom-Vij (2013): 31–34. Perrine (2020): 82–84 offers a similar argument against the value of understanding, though he himself is not a veritist. I should note that my gloss on the veritist position is potentially underdeveloped – see Perrine (2021). As far as I can tell, though, these complications do not impact what I say in this paper.

Many of the stock examples of joint-carving concepts and their gruesome counterparts involve temporal “shuffling”. But fidelity can sometimes be a matter of specificity. The concept *green* has greater fidelity than the concept *green or blue*. Yet inductive inferences that move from “This F is green” to “All Fs are green” are less reliable than inductive inferences that move from “This F is green or blue” to “All Fs are green or blue”. The opposite is true when the form of the inference is reversed. Importantly, though, more specific concepts are not always more joint-carving. As I suggested earlier, the concept *wrong* is epistemically better than the concept *wrong-on-Wednesday* because the former better represents the universal, non-specific, nature of morality. All of this points to the conclusion that the value of fidelity cannot be fully explained in terms of its instrumental value in inductive inferences. Perhaps there is some other reason to think that fidelity has instrumental value. But the burden of proof rests with the veritist.

According to a second strategy, fidelity is an indicator for something else that has epistemic value but fidelity itself does not have epistemic value. Consider the psychology of Li Si, who looks at Zhang San’s bracelet and forms the belief that the bracelet is grue. Most likely, she also believes that the bracelet is green. If she doesn’t believe that the bracelet is green, then something has gone wrong. The only reliable way to form the belief that the bracelet is grue is to first observe that the bracelet is green. But if Li Si does not believe that the bracelet is green then she is failing to believe something that her perceptual evidence shows is obviously true. Thus, the second strategy provides a forked explanation. Either someone who has a gruesome belief also has its joint-carving counterpart, or they do not. If they do, then the gruesome belief only has a sort of non-basic value based on the basic value of the truth of a distinct belief. If they do not believe the joint-carving counterpart, then they are failing to base their beliefs on the evidence that they are given. That failure is instrumentally disvaluable insofar as it will prevent someone from acquiring true beliefs and avoiding false beliefs. Thus, fidelity appears to have value because joint-carving beliefs indicate that the believer is *not* employing an instrumentally disvaluable process of belief formation. But the indication is merely a correlation. In this way, the second strategy tries to explain away the apparent value of fidelity.

This explanation makes sense in the specific case of Li Si looking at a jade bracelet, but it fails to generalize. First, Li Si could form the belief that the bracelet is grue on the basis of Zhang San’s testimony that the bracelet is grue. The veritist may respond by saying that even in that

circumstance Li Si should be able to deduce that the bracelet is green, and so if she does not believe that the bracelet is green then there is something wrong with how she is arriving at the truth. Fair enough. But the issue depends on which concepts are primitive for the agent. For normal humans like Li Si, the concept *green* is primitive and the concept *grue* is conceptually defined in terms of *green*. Thus, it would be strange for Li Si to believe that the bracelet was grue but not believe that the bracelet was green. Yet there could be aliens for whom *grue* is primitive. They could even lack the concept *green*. For such aliens, the second strategy's explanation for the merely apparent value of fidelity fails.

And yet we don't need aliens to show that the explanation fails to generalize. Li Si believes that the bracelet is made of jade. But Li Si is not a chemist and does not know the chemical make-up of the bracelet. Thus, she does not believe that the bracelet is made of nephrite, that it contains silicon, and so on. As a matter of fact, everything that is made of jade also contains silicon. Nevertheless, there is nothing unusual with how Li Si forms her belief. Consequently, a veritist who adopts the second strategy would have to deny that Li Si's belief that the bracelet is made of jade is less valuable than the belief that the bracelet is made of nephrite.

The third strategy employs the notion of a fidelity belief. A *fidelity belief* is a belief that explicitly maintains that some concept is joint-carving. The strategy says that every individual believes the fidelity beliefs that support their primitive concepts. Perhaps, even, they must believe them. Plausibly, every individual must be immodest in that they believe their beliefs to be more accurate than anyone else's – if they were not immodest, then they would just adopt the beliefs of those they judged to be more accurate, in which case they would still believe their beliefs to be the most accurate, albeit indirectly.<sup>11</sup> The third strategy extends this need for immodesty to fidelity: every individual must believe that their concepts are more joint-carving than the alternatives – otherwise, they would come to employ those alternative concepts. So, this third strategy concludes, an individual's fidelic belief has a sort of non-basic epistemic value insofar as it entails that they have a true fidelity belief.

But this strategy seems psychologically unrealistic. It seems to me that plenty of people have beliefs such that they have no opinion about the fidelity of the concepts employed in those beliefs. Children, for instance, have beliefs but lack an explicit concept of fidelity. Even those who do

---

<sup>11</sup>Cf. Horowitz (2014); Joyce (2009)

have an explicit concept of fidelity may be ambivalent about the fidelity of some of their beliefs. Recall that an individual can lack full voluntary control over which concepts are primitive to them. When Adam first moved to New Zealand, he could not help but think of temperature in terms of Fahrenheit. The third strategy would say that Adam had a fidelity belief to the effect that the Imperial measurement system is joint-carving. And yet Adam may have viewed the Imperial measurement system that his brain internalized as an unfortunate burden to be shed, rather than as a joint-carving success. I admit that there is some kind of tension between using a concept and regarding that concept as defective. But that tension is not a matter of inconsistency in belief.

A variation of the third strategy fails for a similar reason. Suppose that a fidelic belief entails a true belief about the objective organization of the world and a gruesome belief entails a false belief about the objective organization of the world. Then the apparent value of fidelic beliefs can be explained in terms of the final value of these entailed belief. Yet these entailments do not hold in general. I believe that the jade bracelet is green, I believe that the jade bracelet was observed before 3000AD, and I believe that something is grue iff either it is green and observed before 3000AD or blue and not observed before 30000AD. So I naturally also believe that the jade bracelet is grue. Yet I also explicitly reject that the world is objectively organized in a grue way. So either I hold contradictory beliefs or, more realistically, in general a belief does not entail a belief about objective organization.

The fourth strategy identifies a value discrepancy in true beliefs. Many philosophers (though certainly not all) maintain that some beliefs are more valuable than others. Often, this difference in value is explained by the comparative *significance* of the beliefs.<sup>12</sup> The first strategy, then, claims that fidelity tracks significance. For example, the belief that all emeralds are green is more valuable than the belief that all emeralds are grue because the former is more significant than the latter. Likewise, beliefs about chemistry (e.g. that an oxygen atom has eight protons) are more valuable than beliefs about mixology (e.g. that a sidecar has cognac, orange liqueur, and lemon juice) because the concepts inherent to chemistry have greater fidelity than the concepts inherent to mixology and that greater fidelity indicates greater significance. And by thinking in terms of bred and rue, Sider's imaginary group of people missed an opportunity to attain a "deeper", more significant, truth about their world.

---

<sup>12</sup>Cf. Alston (2005); Hazlett (2017); Treanor (2014)

This strategy supplies incorrect value judgments. Significance amplifies value: a true significant belief has greater epistemic value than its insignificant counterpart, while a false significant belief has greater epistemic disvalue than its counterpart. But fidelity does not amplify value. It adds value. A true belief that is also fidelic has greater epistemic value than a true belief that is non-fidelic. But a false belief that is fidelic has *less* disvalue than a belief that is both false and non-fidelic.

Think about specific comparisons. Probably most people have the intuition that the belief that the bracelet is green has more value than the belief that the bracelet is grue. Speaking personally, I also have the intuition that the belief that the bracelet is not green has less disvalue than the belief that the bracelet is not grue. But I might be unusual in that respect. So let me address two motivations a veritist may give for the opposite intuition and then offer a motivation for my intuition.<sup>13</sup>

First, the veritist may say that it's worse to believe that the bracelet is not green because it is more blameworthy – epistemically blameworthy – to form that type of belief. In most circumstances, it's obvious that the bracelet is green, and so someone who believed that it was not green must have made a careless mistake. In contrast, *being grue* is a complicated property. It is more excusable for someone to make a mistake about which things are grue.<sup>14</sup>

This motivation assumes an implausible connection between epistemic value and epistemic deontology. To be sure, the extent to which someone is criticizable for their belief has something to do with the disvalue of that belief. But the fact that one belief is less criticizable than another is insufficient to show that the former is less disvaluable than the latter. Recall the aliens with the grue-ified conceptual scheme. For them, *green* is comparatively more complicated, and thus it is more excusable for them to make a mistake about which things are green. But the value (or disvalue) of the belief is unchanged.<sup>15</sup>

A second potential motivation is based on a more sophisticated understanding of truth. Many philosophers talk as if the amount of truth contained within a belief is readily discernible from the surface grammar of the sentence that expresses that belief. But that has to be a simplifying

---

<sup>13</sup>Many thanks to an anonymous reviewer for not letting me get away with intuition-mongering here!

<sup>14</sup>Thanks to an anonymous reviewer for providing this interesting case.

<sup>15</sup>Here, I assume that the epistemic value of a belief is not relative. There are plausible reasons to reject this assumption. But those reasons tend to focus on complex epistemic features and are thus contrary to the veritist position that I'm arguing against here.

assumption. Arguably, for instance, the belief that Socrates is dead is a logically complex belief that contains the belief that Socrates was alive and the belief that Socrates is not alive. In this particular case, ‘green’ and ‘grue’ are single words. But perhaps the concept *green* is more informationally rich than the concept *grue*. If so, then the belief that the bracelet is green would contain more truth than the belief that the bracelet is grue. Similarly, the belief that the bracelet is not green would contain more falsehood than the belief that the bracelet is not grue. Consequently, the true joint-carving belief would be more valuable than the true gruesome belief, and the false joint-carving belief would be more disvaluable than the false gruesome belief<sup>16</sup>

Decisively discrediting this motivation would require a fully developed account for measuring truth; no one, myself included, has one of those. Still, I find this motivation suspicious. Why should we think that *green* is more informationally rich than *grue*? We often describe *grue* as a disjunctive concept. But being disjunctive doesn’t necessarily entail being less informative. In fact, it seems to me that saying something is grue communicates the same amount of information communicated by saying something is green. The only difference is that in the case of grue the information is mixed together in a way that we find somewhat confusing. Think of it this way. There are four possible states for the bracelet to be in: (1) green and observed before 3000AD, (2) green and not observed before 3000AD, (3) blue and observed before 3000AD, and (4) blue and not observed before 3000AD.<sup>17</sup> All of the beliefs under consideration rule out exactly two of these possibilities. The belief that the bracelet is green rules out (3) and (4); the belief that the bracelet is not green rules out (1) and (2). The belief that the bracelet is grue rules out (2) and (3); the belief that the bracelet is not true rules out (1) and (4). Thus, on further examination it seems that the beliefs contain the same amount of information, and therefore would contain the same amount of truth or falsehood.

Finally, here’s how I would motivate my intuition that the belief that the bracelet is not green has less disvalue than the belief that the bracelet is not grue. Generally speaking, we want to improve the value of our beliefs. But if improving the fidelity of our beliefs sometimes worsens their value, then we will often find ourselves in highly counterintuitive situations. Suppose Zilin believes that her current scientific theory is false. (Perhaps she is motivated by the pessimistic meta-

---

<sup>16</sup>Cf. Treanor (2014): 554–558

<sup>17</sup>In truth, there are more possible states corresponding to colors other than green and blue. But that is a complication that does not impact what I am saying and could have been eliminated by defining ‘grue’ differently.

induction.) She knows that, for whatever reason, the true scientific theory is beyond her reach. My intuition is that in this situation Zilin should stick with her current theory. But someone who accepts the opposite intuition would have to say that Zilin would be better off grue-ifying her theory by replacing its concepts with less fidelic counterparts. For the same reason, teaching philosophy would become a risky business. By teaching my students that the wrongness of murder is not relative to the day of the week, I risk the possibility that some students will use the right moral concepts for the wrong moral beliefs. From the perspective of epistemic value, then, my lectures could cause serious harm! The strangeness of these situations suggests my intuition is the right one. It is better to have a false belief that is joint-carving, even if that belief is still overall disvaluable.

The fifth and final strategy I'll consider attempts to explain away the apparent epistemic value of fidelity. While fidelity is valuable, it would be a mistake to take that value to be epistemic. Perhaps thinking about the world with joint-carving concepts is practically valuable. Consider, for example, the progress made in theories of physics, where old concepts are replaced by new, seemingly more joint-carving, concepts. This advancement in our concepts has enabled us to develop technologies like lasers, atomic bombs, and smart phones. Or perhaps thinking about the world with joint-carving concepts creates a sense of "getting it right" that is pleasurable or otherwise prudentially valuable. Finally, perhaps it is valuable (either practically or prudentially) for someone to be "on the same page" as other people when it comes to the concepts that they employ. Fidelity may then serve as a convenient means of coordinating our use of concepts.

I do think that fidelity has non-epistemic value in some of these ways. But I doubt that they can adequately explain away all of its apparent value. These non-epistemic explanations leave large swaths of philosophy inexplicable. Metaphysicians fiercely disagree about the fundamental nature of reality. Yet many of them go to great lengths to show that their disagreements lack practical consequences. Certainly, such metaphysicians often take pleasure at "getting it right". But their pleasure is not merely a matter of fidelity. They also take pleasure in believing the truth! Thus, in response to the fourth strategy, I offer the following *tu quoque*: if the non-epistemic value of fidelity should lead us to doubt that fidelity has epistemic value, then the non-epistemic value of truth should also lead us to doubt that truth has epistemic value.

There might be ways for veritists to respond to my arguments against these five strategies. There also might be other strategies for explaining the apparent epistemic value of fidelity that I

haven't considered. Still, my arguments should have shifted the burden of proof onto the veritists. In the meantime, we can consider the consequences of accepting that fidelity has basic final value.

## 4 The Balance of Epistemic Value

By now, I hope to have motivated the claim that the value of “getting it right” is determined along two distinct dimensions: the value of truth and the value of fidelity. If so, then it's worth asking how fidelity is best pursued as well as how its pursuit should be balanced with the pursuit of truth. We want to “get it right”. But, if “getting it right” means pursuing both truth and fidelity, then what do we do when the two come into conflict? Should we prioritize truth over fidelity? If so, how much truth is worth an off-setting lack of fidelity?

These questions may look familiar. They are the same sorts of questions we ask when considering the pursuit of the truth. As William James famously declared, no one simply “pursues the truth”. Such a description obscures an important complexity. In reality, there are two separable pursuits: seeking true beliefs and avoiding false beliefs. How someone balances these pursuits will shape their overall strategy for belief formation. If they privilege seeking true beliefs, then they should just believe everything. Much of what they would come to believe would be false, but they would be guaranteed to believe everything that is true. Similarly, if they privilege avoiding false beliefs, then they should maintain total agnosticism. They would entirely lack true beliefs, but they would also succeed in their goal to never form a false belief.<sup>18</sup>

To give this dual-natured pursuit a less misleading label, call it the pursuit of *alethic value*. Now, some philosophers think that we lack direct control over our beliefs. If they're right, then the “optimal” strategies suggested above are not so easily implemented. Nevertheless, even those philosophers would agree that there is a wide range of available strategies for pursuing alethic value. These strategies influence how skeptical or dogmatic someone is, how often they seek further evidence, and so on. Philosophers disagree as to which of these strategies are rational, and some even argue that there is only one rational strategy for the pursuit of alethic value.

In fact, there is an analogous complexity in the pursuit of fidelity. A joint-carving belief has positive epistemic value and a gruesome belief has negative epistemic value. The pursuit of *fidelic*

---

<sup>18</sup>See James (1979): 17–18.

*value* includes seeking joint-carving beliefs and avoiding gruesome beliefs. Their “optimal” strategies might also be beyond our ability to implement. Some philosophers think that, in one way or another, we lack full control over what concepts we employ.<sup>19</sup> If they’re right, then in some situations we are unable to acquire joint-carving beliefs or avoid having gruesome beliefs. And yet there is still a wide range of available strategies that influence how readily someone adopts new concepts, how hesitant they are to abandon their old ones, and so on.

Philosophers may disagree as to which fidelic balancing strategies are rational. But because they have not explicitly acknowledge the separability of the two pursuits, they in fact have had little to say about the balance between them. As a consequence, many of the recent disputes in metaphysics over ideology – that is, disputes over metaphysical terminology – can be traced to unarticulated differences about the proper balance between seeking the joints and avoiding the gruesome. Here’s one example. In a discussion of the comparative fidelity of logical ideology, Sider advocates for an “egalitarianism” according to which “both  $\exists$  and  $\forall$  carve at the joints” (Sider (2011): 258). He does not seem overly bothered by the possibility that he is carving too aggressively. Others, however, seem more cautious. To avoid arbitrariness in his choice of ideology, Cowling (2013) offers a kind-based principle of ideological parsimony; to avoid both arbitrariness and redundancy, McSweeney (2019) motivates a sort of agnosticism about the fidelity of our logical concepts. Their reasoning can be seen as motivated by an underlying aversion to bad butchery. It’s a move they judge to be worth making, even if they thereby risk missing one of reality’s joints. If there is no uniquely rational balance between seeking fidelic beliefs and avoiding non-fidelic beliefs, then there might be no fact of the matter about which of these strategies is best.

So “getting it right” is a multi-faceted affair. The pursuit of alethic value requires a balance between seeking true beliefs and avoiding false beliefs. The pursuit of fidelic value requires a balance between seeking fidelic beliefs and avoiding non-fidelic beliefs. And the pursuit of the value of “getting it right” requires a balance between the pursuit of alethic value and the pursuit of fidelic value. I don’t mean to suggest that the analogies are as perfect as I’ve drawn them here. I am sure that there are subtle differences worth exploring. But the point stands: the Jamesian framework neatly characterizes the complexity inherent to “getting it right”.

---

<sup>19</sup>For example, linguistic relativists may say that facts about an individual’s first language constrain what concepts they are capable of employing. For some relevant discussions, see Gumperz and Levinson (1996).

And we are once again confronted with a new version of the same question: which strategies for balancing the pursuit of alethic value and the pursuit of fidelic value are rational? I don't have a detailed answer to that question. But I do have an answer to which strategies are not rational. Adapting some useful terminology from Arrhenius and Rabinowicz (2015), we can distinguish between five separate principles regarding the relationship between alethic value and fidelic value:

**Strong Alethic Superiority:** Any amount of alethic value is better than any amount of fidelic value.

**Weak Alethic Superiority:** A sufficient amount of alethic value is better than any amount of fidelic value.

**Archimedeanism:** For any amount of alethic value there is some amount of fidelic value which is at least as good and for any amount of fidelic value there is some amount of alethic value which is at least as good.

**Weak Fidelic Superiority:** A sufficient amount of fidelic value is better than any amount of alethic value.

**Strong Fidelic Superiority:** Any amount of fidelic value is better than any amount of alethic value.

Informally, the first two principles take alethic value to be a "higher epistemic good" than fidelic value. According to them, even though more fidelity is always better than less, there's no amount of fidelity that outvalues a (certain level of) truth. So, for example, the addition of a single truth to someone's body of beliefs would be more valuable than the replacement of those beliefs with perfectly joint-carving ones, even if the replacement did not detract from the alethic value of the beliefs. The same is true for the last two principles, but in reverse: avoiding a million false beliefs wouldn't be worth the smallest decrease in fidelity. In contrast, according to the third Archimedean principle the values have the same upper bound (which might be infinite) and consequently there will always be a trade-off worth making.

I suspect that most philosophers will find both Weak Fidelic Superiority and Strong Fidelic Superiority implausible. But I also suspect that many philosophers may be inclined to privilege the pursuit of alethic value over the pursuit of fidelic value by endorsing either Strong Alethic Superiority or Weak Alethic Superiority. I think that would be a mistake.

Consider Putnam’s model-theoretic argument.<sup>20</sup> Any theory, before we can say whether the theory is true or false, needs to be interpreted. But, on one way of construing Putnam’s argument, almost any extensional theory can be interpreted in a way that makes the theory flawlessly true. So far as the truth of our theories is concerned, it doesn’t matter what the world is really like. The same can be said for interpretations of what we believe. This result threatens the supposed value of “getting it right”.

One standard response to the argument is to claim that not all interpretations are created equal. Some interpretations are better than others in virtue of the fact that the former pick up on objective samenesses and differences. Here’s how David Lewis puts it:

Among all the countless things and classes that there are, most are miscellaneous, gerrymandered, ill-demarcated. Only an elite minority are carved at the joints, so that their boundaries are established by objective sameness and difference in nature. Only these elite things and classes are eligible to serve as referents. The world – any world – has the makings of many interpretations that satisfy many theories; but most of these interpretations are disqualified because they employ ineligible referents. When we limit ourselves to the eligible interpretations, the ones that respect the objective joints in nature, there is no longer any guarantee that (almost) any world can satisfy (almost) any theory. It becomes once again a worthy goal to discover a theory that will come true on an eligible interpretation, and it becomes a daring and risky hope that we are well on the way toward accomplishing this (Lewis (1984): 227).

Lewis’s use of the phrase “a worthy goal” is suggestive. Putnam’s model-theoretic argument shows how we can secure maximal alethic value by deprioritizing our pursuit of fidelic value. If either Strong Alethic Superiority or Weak Alethic Superiority were true, then we should hail the argument as our epistemic salvation – how fortuitous it is that the highest good is acquired at so little cost! But we shouldn’t. The worthy goal of “getting it right” depends on the non-infinitesimal value of fidelity. Strong Alethic Superiority and Weak Alethic Superiority are false and the treatment of alethic value as a higher epistemic good is irrational.

## Conclusion

In this paper, I’ve attempted to shed light on some neglected issues inherent to “getting it right”. Once we accept that fidelity is an important epistemic feature, we face a multitude of pressing

---

<sup>20</sup>Cf. Putnam (1977). I am suppressing most of the technical details, but I hope I have remained faithful to the overall thrust of the original presentation.

questions regarding its value in relation to the value of other epistemic features.

Some of these questions I answered. Fidelity is epistemically valuable. Yet the epistemic value of fidelity is actually a complex value constituted by the value of joint-carving beliefs and the disvalue of gruesome beliefs. This complex fidelic value cannot be entirely explained in terms of alethic value. Furthermore, alethic value is not a higher epistemic good in comparison to fidelic value; the value of “getting it right” depends on the combination of both truth and fidelity.

Other questions I have not attempted to answer. One especially pressing question is whether there is more than one rational way to balance the pursuit of these values. Many philosophers respect the differences between the epistemically courageous and the epistemically cautious. Does such respect extend to fidelity? If so, does that suggest new versions of permissivism, especially as it bears on metaphysical disputes?

## References

- Ahlstrom-Vij, Kristoffer. 2013. “In Defense of Veritistic Value Monism.” *Pacific Philosophical Quarterly* 94:19–40. doi:10.1111/j.1468-0114.2012.01442.x.
- Alston, William. 2005. *Beyond “Justification”: Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Arrhenius, Gustaf and Rabinowicz, Wlodek. 2015. “Value Superiority.” In Iwao Hirose and Jonas Olson (eds.), *The Oxford Handbook of Value Theory*, 225–248. Oxford: Oxford University Press.
- Berker, Selim. 2013. “Epistemic Teleology and the Separateness of Propositions.” *Philosophical Review* 122:337–393. doi:10.1215/00318108-2087645.
- Cowling, Sam. 2013. “Ideological Parsimony.” *Synthese* 190:3889–3908.
- Dasgupta, Shamik. 2018. “Realism and the Absence of Value.” *Philosophical Review* 127:279–322.
- David, Marian. 2005. “Truth as the Primary Epistemic Goal: A Working Hypothesis.” In Matthias Steup and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 296–312. Oxford: Blackwell.
- Finocchiaro, Peter. 2021a. “High-Fidelity Metaphysics: Ideological Parsimony in Theory Choice.” *Pacific Philosophical Quarterly* Early Online View.
- . 2021b. “Ideology and Its Role in Metaphysics.” *Synthese* 198:957–983.
- Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Greco, John. 2003. “Knowledge as Credit for True Belief.” In Michael DePaul and Linda Zagzebski (eds.), *Intellectual Virtue: Perspectives From Ethics and Epistemology*, 111–134. Clarendon Press.

- Gumperz, John J. and Levinson, Stephen C. (eds.). 1996. *Rethinking Linguistic Relativism*. Cambridge University Press.
- Hazlett, Allan. 2017. "Understanding and Structure." In Stephen R. Grimm (ed.), *Making Sense of the World: New Essays on the Philosophy of Understanding*, 135–158. Oxford University Press.
- Horowitz, Sophie. 2014. "Immoderately Rational." *Philosophical Studies* 167:41–56.
- James, William. 1979. *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge: Harvard University Press. First published in 1897.
- Joyce, James. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of Belief*, 263–297. Synthese.
- Korsgaard, Christine. 1983. "Two Distinctions in Goodness." *Philosophical Review* 92:169–195.
- Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61:343–377.
- . 1984. "Putnam's Paradox." *Australasian Journal of Philosophy* 62:221–236.
- McDaniel, Kris. 2017. "Normative Accounts of Fundamentality." *Philosophical Issues* 27:167–183.
- McSweeney, Michaela Markham. 2019. "Following Logical Realism Where It Leads." *Philosophical Studies* 176:117–139.
- Moore, G.E. 1903. *Principia Ethica*. Dover Publications.
- Nehamas, Alexander and Woodruff, Paul. 1995. *Plato: Phaedrus*. Indianapolis: Hackett. Translated with introduction and notes.
- Pasnau, Robert. 2011. *Metaphysical Themes: 1274–1671*. Oxford University Press.
- Perrine, Timothy. 2020. "On Some Arguments for Epistemic Value Pluralism." *Logos and Episteme* 11:77–96. doi:10.5840/logos-episteme20201115.
- . 2021. "A Puzzle About Epistemic Value and Steps Towards a Solution." *Synthese* 1–17. Online Early Access.
- Putnam, Hilary. 1977. "Realism and Reason." *Proceedings of the American Philosophical Association* 50:483–498.
- Rorty, Richard. 2000. *Philosophy and Social Hope*. New York: Penguin.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Thomasson, Amie L. 2014. *Ontology Made Easy*. Oxford University Press.
- Treanor, Nick. 2014. "Trivial Truths and the Aim of Inquiry." *Philosophy and Phenomenological Research* 89:552–559.
- Zagzebski, Linda. 2004. "Epistemic Value Monism." In John Greco (ed.), *Ernest Sosa and His Critics*. Oxford: Blackwell.