

## Partial Truth versus Felicitous Falsehoods

Soazig Le Bihan

Received: date / Accepted: date

**Abstract** Elgin has argued that scientific models that are, strictly speaking, inaccurate representations of the world, are epistemically valuable because the “falsehoods” they contain are “felicitous”. Many, including Elgin herself, have interpreted this claim as offering an alternative to scientific realism and “veritism”. In this paper, I will argue that there is a more felicitous interpretation of Elgin’s work: “felicitous falsehoods” do play a role in the epistemic value of inaccurate models, but that role is of instrumental value. Elgin’s view is not best understood as claiming that falsehoods provide scientific understanding in and of themselves, only that they facilitate epistemic access to the fundamental, even if partial, truths that are contained within models. While falsehoods may be felicitous in that they facilitate exemplification, the epistemic value of inaccurate models ultimately relies on their partial accuracy.

**Keywords** Understanding · Idealization · Models · Factivity · Realism

### 1 Introduction

In her *True Enough*, Elgin sets up a dilemma for philosophers of science and epistemologists: (1) philosophy of science and epistemology have taken truth to be the ultimate epistemic value; (2) modern science as it is practiced is replete with products and practices for which truth is not, and arguably ought not be, the ultimate epistemic value; so (3) either modern science is not epistemically valuable, or philosophy of science and epistemology are lacking, and are in need of reform.<sup>1</sup> Unsurprisingly, Elgin defends the latter option and attempts to articulate an account of the epistemic value of science in which truth is not the ultimate value. In Elgin’s words (Elgin, 2018, 15):

If truth is mandatory, much of our best science turns out to be epistemologically unacceptable and perhaps intellectually dishonest. Our predicament is this: We can retain the truth requirement and construe science either as cognitively defective or as noncognitive, or we can reject, revise, or relax the truth requirement and remain cognitivists about and devotees of science. I take it that science provides an understanding of the natural order. By this I do not mean merely that an ideal science would provide such an understanding or that at the end of inquiry science will provide one, but that much actual science has done so and continues to do so. I take it, then, that much actual science is cognitively reputable – indeed, estimable.

The view according to which truth is a “requirement” is dubbed “veritism” by Elgin. She claims that her account of scientific understanding offers an alternative to veritism, such that modern science, including its falsehoods, can be construed as having epistemic value.

---

University of Montana

<sup>1</sup> When referring to scientific representations, it is, strictly speaking, inappropriate to characterize them as “true” or “false”, as representations are neither true or false, but only accurate or inaccurate representations of their target, where accuracy is assessed according to multiple dimensions that crucially depend on the user and usage of the representation (van Fraassen, 2008). That said, in this article, I will characterize scientific representations as “true” or “false”, for simplicity and also to align myself with Elgin’s own practice.

The starting point of this paper is to note that Elgin’s challenge to philosophy of science and epistemology is both legitimate and important. That said, I will argue in Section 2 that by her own definition of veritism, Elgin’s view *is* veritist, contrary to her claims. In Section 3, I will demonstrate that her arguments against veritism are faulty. Finally, in Section 4, I will characterize how her view importantly and insightfully differs from typical veritist accounts. Overall, the conclusion is that Elgin’s account of the epistemic value of scientific representations containing falsehoods is original and important, but not for the reasons she claims it is.

## 2 Elgin: Veritism and Partial Truth

The main goal of this section is to argue that, despite her claims to the contrary, Elgin’s view is “veritist” by her definition of that term, which will be specified below (Subsection 2.1). On Elgin’s view, I will argue, the epistemic value of scientific products containing blatant falsehoods is grounded in truth. Granted, scientific products are valuable under her account even if they are only *partially* true, but truth is still a “requirement” (Subsection 2.2). In fact, her view aligns itself with many other philosophers of science who take partial truth to be a necessary condition for epistemic value (Subsection 2.3). Overall, the conclusion of the section is that Elgin’s view fails to provide an alternative to veritism and is instead well-aligned with the views of philosophers that count as veritists by Elgin’s definition.

### 2.1 Elgin’s Definition of Veritism

It is important to specify how Elgin defines the notion of veritism for the sake of this paper. At the beginning of her discussion, Elgin defines veritism as the view according to which truth is the sole, necessary and sufficient standard of epistemic value (Elgin, 2018, 9-10):

According to veritism, truth-conduciveness is *the* appropriate standard of assessment for epistemic policies, practices, and their products. [*my emphasis*]

That said, Elgin relaxes this definition in the ensuing discussion (Elgin, 2018, 10-11). If truth is the sole, necessary and sufficient standard of epistemic value, then any truth is epistemically valuable. But if that’s the case, then accumulating trivial truths is epistemically valuable. Since that’s not acceptable, Elgin explains, veritists have amended their views in various ways. She mentions several authors who she takes to have followed that path: Pritchard, Sosa, Treanor, and Hetherington. The common theme is that, in the face of the problem of trivial truths, these authors have distinguished between different degrees of “quality” of truths which provide different degrees of epistemic value. On Sosa’s account, for example, true beliefs are valued only if they are the result of the application of intellectual virtues (Sosa, 2007). Treanor argues to the effect that not all truths “weigh” the same: some truths weigh more because they are more conducive to increasing knowledge and reducing ignorance than others. Just like gold miners look for great veins of gold rather than flakes or dust, epistemic agents look for truths that promise to yield more truths (Treanor, 2013). Neither Sosa nor Treanor thus takes mere truth to be the sole, necessary and sufficient standard of epistemic value. That said, in Elgin’s view, they remain veritist, because, still, in their views, “truth is epistemically central” (Elgin, 2018, 11). I take this to mean that veritism is the view according to which:

- considerations of truth are necessary, even if not always sufficient, when assessing epistemic value;
- when other considerations apply, truth is always of central concern.

In what follows, it will be shown that, under this definition, Elgin’s account of the epistemic value of scientific representations containing blatant falsehoods is veritist.

### 2.2 Elgin’s Veritism

Remember that Elgin’s starting point is that scientific products that contain blatant falsehoods are a challenge for veritism. Now Elgin admits that not all scientific products containing falsehoods will

be problematic. For instance, scientific theories and models that have been articulated in the past but have since been found to be inaccurate representations of the world cause no trouble for veritism. The reason is that, while they do indeed contain falsehoods, they have been rejected precisely for reasons in alignment with veritism. Such instances of scientific falsehoods are not epistemically valuable in themselves. If they are valuable, they are so in a derivative and instrumental way, i.e. as steps toward theories that are relatively closer to the truth. Similarly, anomalies are instances of scientific falsehoods that are not problematic for veritism. Anomalies are typically taken to be an indicator that our scientific representations are inaccurate and hence in need of modification. In other words, anomalies point to falsehoods that scientists hope to see disappear. As such, they are not valued for themselves: if they are valued at all, it is only in a derivative and instrumental way, i.e. as steps toward theories that are relatively closer to the truth. Truth thus remains the standard of epistemic value of scientific practice in these two contexts. Veritism can only be challenged by instances of scientific falsehoods that would be valued for themselves and of which the epistemic value would not be instrumental and derivative vis-à-vis the value of truth.

Highly idealized models are promising candidates to challenge veritism because many of them contain falsehoods that are not taken to be defects but rather valuable assets.<sup>2</sup> Some of Elgin's favorite examples of this kind of model are also classic examples in the literature: the ideal gas model and the Hardy-Weinberg model in population genetics. The former represents gases as made of dimensionless molecules interacting only by colliding perfectly elastically with one another and the walls of the container. The latter represents the genetic variation of a population that is isolated from all evolutionary forces, is infinite, and features random mating. These kinds of idealizations are clearly blatantly false, or, as Elgin puts it, they "are true of nothing real" (Elgin, 2018, 15). Dimensionless molecules with no intermolecular attraction do not exist. Populations that are infinite and mate randomly do not exist either. And yet, both models are considered as epistemically valuable: they are taken to provide us with some understanding of gas behavior and population genetics. The question is how.

Elgin's answer to that question is that scientific models need not be completely true to be epistemically valuable, all they need to be is "true enough". In other words, epistemic value is a threshold issue, not an all-or-nothing issue as the veritist, according to Elgin, would require. Now what does it take to be "true enough"? According to Elgin, scientific claims are true enough if (Elgin, 2018, 15-28):

1. they constitute close approximations; or
2. they constitute not-so-close approximations that are practically useful – in contrast to some intractable truth; or
3. they isolate features of target systems that we deem important.

What is interesting is that none of this stands in conflict with veritism as defined by Elgin. Remember that, according to Elgin's definition, veritism is the view that truth is a necessary condition for epistemic value, and a consideration of central concern even when other conditions apply. By that definition, my contention is that Elgin's account of why the products and practices described above are epistemically valuable is consistent with veritism.

Let us begin with close approximations. Close approximations are approximations to the truth, and are assessed on the basis of how close to the truth they are. Approximations that would not closely approximate truth would not count as epistemically valuable in this context. Truth is thus a necessary condition and central concern in this account of the epistemic value of close approximations, which is consistent with veritism.

Second, not-so-close approximations that are practically useful are accepted, according to Elgin, because their deviations from truth is not too detrimental to the possibility of using the models for prediction (and possibly explanation). So, such approximations are still monitored and assessed with an eye towards truth. They can easily be described as a "second best" option: getting to the truth would be preferred, but, given that the truth is "intractable", perhaps because of the complexity of the target system, one

<sup>2</sup> Elgin also takes thought experiments to be promising candidates to challenge veritism. She argues that the veritist cannot account for the epistemic value of thought experiments because they are "morally, practically, and physically unfeasible" (Elgin, 2018, 15). It is unclear how exactly the moral, practical, or physical feasibility of an experiment bears on the truth of the scientific claims that are made in relation to the thought experiment. I will not discuss this issue any further in this paper: suffice it to say, Elgin's claim is at least controversial, as many would argue that thought experiments are typically valuable because they afford knowledge, which implies that their epistemic value is related to truth (e.g. Norton (2004)).

settles for a not-so-close-but-tractable approximation instead. The approximation would not be valuable, however, if it was not allowing for true predictions, i.e. for a true description of the target system's behavior under certain conditions. So a minimal requirement of truth is thus still in operation, and truth is still of central concern even if other considerations, in this case of tractability, apply. Again, Elgin's account is consistent with veritism.

Finally, the notion that models isolate features of target systems that we deem important implies that models accurately represent these important features. In other words, such models are partially accurate or partially "true" representations of the world. They are epistemically valuable, according to Elgin, because they get something right about the world. Hence, here again, truth is a necessary condition and central concern for epistemic value, and Elgin's account is consistent with veritism.

That Elgin has failed to cut ties with veritism is clear in the following quote, which concludes her discussion of the epistemic value of scientific models containing falsehoods.(Elgin, 2018, 28):

The foregoing examples show that in some cognitive contexts we accept claims that we do not consider true. But we do not indiscriminately accept falsehoods either. The question then is, what makes a claim acceptable? Evidently, to accept a claim is not to take it to be true, but to take it that the claim's divergence from truth, if any, is negligible. As Strevens (2008) puts it, the divergence from truth is not a difference maker. The divergence need not be small, but whatever its magnitude, it can be safely neglected. We accept a contention, I suggest, when we consider it true enough. The success of our cognitive endeavors indicates that we are often right to do so. If so, a contention is acceptable when its divergence from truth is negligible. In that case it is true enough.

So, scientific claims are epistemically valuable only if they are true enough and they are true enough only if the divergence from truth is negligible. In other words, truth is still a necessary condition and central concern when assessing epistemic value. Given that it is most often than not somewhat intractable, however, one needs to relax their standard, and accept scientific claims which are partially false whenever the falsehoods included are "negligible". Importantly, counting as negligible is determined by the standard of truth. Again, Elgin's account is perfectly consistent with the veritist ideal of the truth being a necessary condition and central concern for epistemic value. In fact, I will argue below, Elgin's account is consistent with many other accounts of the epistemic value of scientific representations containing blatant falsehoods that would qualify as "veritist" under her definition. What all these accounts, including Elgin's, have in common, is that these representations possess epistemic value essentially because they are partially true.

### 2.3 Veritism and Partial Truth

The main problem with Elgin's argument to the effect that her account stands in conflict with veritism is that, while she explicitly defines veritism as the view that takes truth to be a necessary condition and central concern for epistemic value, she then argues as if it involved a much stronger commitment to truth. Her argument relies on the assumption that "abject fidelity to truth [be] required" (Elgin, 2018, 27) by the veritist, such that no scientific models containing any degree of falsehoods could be deemed valuable. But of course that is not the case. One can take truth to be necessary condition and central concern for epistemic value and value scientific models containing falsehoods as long as they also meet the truth requirement to some extent. The analysis of Elgin's own account of the epistemic value of scientific models containing falsehoods is a case in point. The vast majority of philosophers of science who would agree to be called "veritist" under her definition would not agree to commit to a requirement of "abject fidelity to truth". They would instead consider models that are partially true to be epistemically valuable. In fact, they would do so for similar reasons as Elgin does.

The fact that Elgin refers to Strevens in the quote above is telling, as he is one of the clear examples of philosopher of science whose views align themselves with the veritist ideal as defined by Elgin, but without a requirement of complete truth. According to Strevens (Strevens, 2008):

- scientific understanding is nothing more than grasping scientific explanations,
- scientific explanations are nothing but causal explanations, and

- causal explanations are those that provide a minimal and sufficient set of true difference-makers for the target phenomena.

Idealized models, on Strevens' view, are epistemically valuable if and only if:

- they accurately represent a minimal and sufficient set of difference-makers; and
- they only inaccurately represent non-difference-makers.

In other words, idealized models are epistemically valuable on Strevens' view only because of the truths they contain. Scientific falsehoods are acceptable if and only if none of them are about the difference-makers, i.e. none of these falsehoods are of the important processes at play in the production of the phenomena. On Strevens' account, truth is thus a necessary condition and a central concern for epistemic value.

So Strevens' account is consistent with veritism as defined by Elgin and also consistent with models containing falsehoods being epistemically valuable. This in turn implies that, if Elgin takes her account to be compatible with Strevens, then she must take it to be compatible with veritism, and she must take veritism to be compatible with partially true models being epistemically valuable.

Other examples of authors who would happily be called "veritist" by Elgin's definition and still value partially true models include Mizrahi, who takes models that provide close approximations or not-so-close-but-useful approximations as epistemically valuable on the basis of truth-related considerations. Close approximations are close approximations to true representations: the closer to truth, the better. Truth is the standard. He also takes not-so-close-but-useful approximations as epistemically valuable: the ideal gas law, for example, is valuable because, Mizrahi claims, it gives accurate predictions regarding the behavior of the target systems under certain conditions (Mizrahi, 2012, 242):

The behavior of gases usually agrees with the predictions of the Ideal Gas Law when the temperature of a gas is not close to its liquefaction point and when pressure is approximately under 10 atm.

Mizrahi takes this view to be perfectly compatible with a "partial-truth" account of the epistemic value of scientific models (Mizrahi, 2012, 245):

The Ideal Gas Law doesn't give us the truth, the whole truth, and nothing but the truth about the behavior of gases. Rather, it tells a partial, true story about the behavior of certain gases under specified conditions. Admittedly, there are deviations from ideal-gas behavior. But scientists know that and, presumably, the law is not meant to have universal scope.

Elgin articulates two types of objection to accounts à la Mizrahi (Elgin, 2018, 28-31). The first is to claim that not all idealizations are close approximations to the truth: in the case of the ideal gas law, the idea of dimensionless molecules is no where near the truth. The second is to claim that not all approximations are meant to be improved or not all idealizations discarded over time: many are taken to be indispensable and valuable.

Regarding the first kind of objection, Mizrahi's response is along the following lines. The idea gas law is true of gas behavior under particular conditions. Ideal gases exist (are "real") as they are defined as those gases that obey the law under certain conditions. There is no need, Mizrahi explains, to commit ontologically to the content of the hypotheses made for the theoretical derivation of the law. In other words, Mizrahi focuses on the fact that the model provide a partially true account of gas behavior.

Note that this is compatible with Elgin's own view of not-so-close-but-useful approximations. Elgin herself distinguishes between different ways in which approximate models can be valuable, as explained in Subsection 2.2. Some approximations are close to the truth, and valued as such, some others are not close to the truth, but valued as "second best" if the truth is somehow intractable and the approximate models make adequate predictions regarding the system's behavior. Mizrahi would certainly agree.

Regarding the second kind of objection, i.e. that there are approximations that scientists neither mean nor want to improve or discard, one has to shift back to accounts in which scientific models containing falsehoods are epistemically valuable because they get something right about what matters the most. Strevens' account presented above is definitely of this kind. There is value, in Strevens' account, on focusing solely on minimal sets of difference-makers and neglecting those factors that do not make a difference. This is because what is core to the value of models is their explanatory power, and good explanations are precisely these explanations laying out minimal sets of difference makers for the target

phenomena. Adding extra, true, but unimportant factors as part of the explanation does not improve the quality of an explanation, and hence does not increase the epistemic value of the model.

Potochnik’s account would be another example of an account of this kind. She argues that highly idealized scientific models are representations of causal patterns in the world (Potochnik, 2017). These causal patterns are accurately represented by the models and the epistemic value of the models stems from the fact that they get something right about these causal patterns. So, partial truth is a necessary condition and central concern. Keeping track of causal patterns is important because they “are key to human insight into and influence over our world” (Potochnik, 2017, 33). On her view, idealizations are indeed not always meant to be improved or discarded over time, quite the contrary. This is because (1) the world is incredibly complex, and (2) human cognition is limited. So, idealizations are core to the value of idealized models, in that they allow the models to reveal these causal patterns in which we are most interested from a practical standpoint.

Note that both Strevens’ and Potochnik’s views are compatible with Elgin’s own account according to which scientific representations containing blatant falsehoods are valuable because they isolate features of target systems that we deem important. So, overall, Elgin’s three-pronged account of the epistemic value of highly idealized models is compatible with accounts of various veritists. Mizrahi, Potochnik, Strevens, and others would agree that these models are epistemically valuable because:

1. they constitute close approximations; or
2. they constitute not-so-close approximations that are practically useful – in contrast to some intractable truth; or
3. they isolate features of target systems that we deem important.

Now Elgin develops some arguments *against* the notion that veritism is acceptable. In what follows, I argue that Elgin’s arguments against veritism miss their targets.

### 3 Faulty Anti-Veritism Arguments

In Section 2, I argued that Elgin’s account is in agreement with the main tenets of veritism as she defines it. That said, Elgin claims to demonstrate that veritism is an inappropriate view of science. This section examines her three main arguments to that effect. Subsections 3.1 and 3.2 each focuses on one of two lines of arguments that veritism stands in conflict with holistic confirmation, which Elgin takes to be characteristic of the sciences. Subsection 3.3 is devoted to her argument that veritism is committed to scientific understanding being factive, which she takes to be unacceptable. It will be argued that Elgin’s arguments are unsuccessful. That said, the analysis will reveal which version of veritism she rejects, which will in turn help us make sense of her claims that she is not veritist.

#### 3.1 Granularity of Justification

Let us begin with Elgin’s first main argument against veritism as conflicting with holistic confirmation. Consider a scientific account  $S$  constituted of several propositions. Veritism, Elgin explains, requires that each individual propositions in  $S$  be true for  $S$  to be true. Call this the “granularity” of the truth of  $S$ .

Granularity of Truth ( $G_T$ ): A set of proposition  $S = \{S_1, \dots, S_n\}$  is true iff  $\forall i \in \{1, \dots, n\}, S_i$  is true.

This is unproblematic. In fact, it is so unproblematic that it is unclear how it could be taken to be false, i.e. how anyone could claim that there is a sense in which a set of proposition could be “true” without each of the propositions contained in the set being true.  $G_T$  amounts to either a category mistake, as sets are not the kind of things that are “true”, or a simple definition of what is to be meant by a “true set of propositions”. So, the idea of granularity of truth is most charitably characterized as a simple definition for “true set of propositions”.

Now Elgin claims that it follows from granularity of truth that justification is also granular (Elgin, 2018, 12):

If this [notion of granularity of truth] is right, the sort of justification of interest in epistemology is in the first instance the justification of individual propositions.

It is not clear what exactly Elgin means by granularity of justification: it is useful to look at her argument to get a better sense of it. As mentioned above, the first step of the argument is to claim that, because the veritist commits to  $G_T$ , she also commits to granularity of justification. Elgin then contrasts the veritist view with her *holistic* view (Elgin, 2018, 12 sq.). On this view, one comes to accept an “account” (i.e. “a constellation of mutually supportive commitments that bear on a topic”) as a whole. An account includes many kinds of claims, including factual claims, normative rules, and meta-level constraints. The important point for us regards justification: an account has *epistemic* priority over its components. In other words, the acceptability of individual claims is derivative on the acceptability of the account. There is no way, on Elgin’s view, to secure individual statements independently. The next piece of the argument consists in displaying how science functions under a (at least partially) holistic form of justification. In the sciences, Elgin argues, justification is *de facto* holistic and not granular. Scientific accounts are made of “mutually supportive commitments” (Elgin, 2018, 36). The conclusion then follows that because of the commitment of veritism to granularity of truth, and hence of justification, and because of the “unavailability of sentence-by-sentence verification or justification” (Elgin, 2018, 37) in the sciences, veritism is an inadequate view for science.

From this summary, one can gather the following definition of granularity of justification:

Granularity of Justification ( $G_J$ ): A set of proposition  $S = \{S_1, \dots, S_n\}$  is justified granularly iff  $\forall i \in \{1, \dots, n\}, S_i$  is justified independently.

One can also give a clear expression of Elgin’s argument:

$$\begin{aligned} \text{Veritism} &\rightarrow G_T \\ G_T &\rightarrow G_J \\ &\neg G_J \\ &\neg \text{Veritism} \end{aligned}$$

The main problem with Elgin’s reasoning is the second premise, i.e.  $G_T \rightarrow G_J$ . From a strictly logical point of view, and whatever granularity of justification is taken to mean, it is unclear how it could follow from granularity of truth as defined above. This is because, if indeed the notion of granularity of truth is a simple definition, then it is unclear how anything could be entailed by it regarding justification, since the definition does not mention justification at all. *Prima facie*, then, the claim that granularity of truth implies granularity of justification is highly dubious.

As a matter of fact, many accounts of justification could be committed to  $G_T$  without being committed to  $G_J$ . Let us see this by way of a simple example. Let  $S = \{S_1, S_2, S_3\}$ . Let us assume  $G_T$  to be true, i.e.  $S$  is true if and only if the conjunction  $S_1 \wedge S_2 \wedge S_3$  is true. Now, we are interested in the justification of  $S_1 \wedge S_2 \wedge S_3$  and its relation to the justification of  $S_1$ ,  $S_2$ , and  $S_3$ , individually. The question of interest is whether  $G_J$  can be violated. Let us focus on justification based on evidence. Consider a piece of evidence  $e$  that we know is the case. Our question becomes: would it be inconsistent to have  $e$  supporting the conjunction  $S_1 \wedge S_2 \wedge S_3$ , but supporting neither  $S_1$  nor  $S_2$  nor  $S_3$  individually (with  $G_T$  still holding true). If some accounts of justification are both compatible with  $G_T$  and such a scenario, then the premise  $G_T \rightarrow G_J$  is not true, and Elgin’s argument is undermined.

There are such accounts. In fact, I content that any Bayesian-style account of justification is friendly to a scenario like the one above, while being also fully compatible with  $G_T$ . Bayesian-style accounts of justification will typically commit to the following probabilistic rule of evidential justification (Williamson, 2002, 186):

Evidential Justification ( $E_J$ ): Evidence  $e$  supports the hypothesis  $h$  if only if  $P(h|e) \geq P(h)$ .

Applying this rule to our scenario,  $e$  supports  $S_1 \wedge S_2 \wedge S_3$  iff:

$$P(S_1 \wedge S_2 \wedge S_3|e) \geq P(S_1 \wedge S_2 \wedge S_3).$$

but at the same time, for this to be non-granular justification, one requires that:

$$\forall i \in \{1, 2, 3\}, P(S_i|e) = P(S_i).$$

Such a scenario is perfectly consistent. All that is required for  $P(S_1 \wedge S_2 \wedge S_3|e) \geq P(S_1 \wedge S_2 \wedge S_3)$  to be the case is that (applying Bayes' Theorem):

$$\frac{P(S_1 \wedge S_2 \wedge S_3).P(e|S_1 \wedge S_2 \wedge S_3)}{P(e)} \geq P(S_1 \wedge S_2 \wedge S_3),$$

which simplifies to:

$$\frac{P(e|S_1 \wedge S_2 \wedge S_3)}{P(e)} \geq 1.$$

So, any situation in which  $\forall i \in \{1, 2, 3\}, P(S_i|e) = P(S_i)$  and  $\frac{P(e|S_1 \wedge S_2 \wedge S_3)}{P(e)} \geq 1$  will qualify as an instance of non-granular evidential justification.

A concrete example might help. Consider the set of three propositions:  $\{S_1$ : 'Special Relativity Theory is true',  $S_2$ : 'The relative velocity between two clocks, Clock 1 and Clock 2, is 0.995c' and  $S_3$ : 'Clock 1 ticks every second relative to its rest frame.'}.<sup>3</sup> Now consider the piece of evidence  $e$ : 'Clock 2 ticks every 10 seconds relative to Clock 1'. The evidence  $e$  that Clock 2 ticks every 10 seconds relative to clock one does support the conjunction  $S_1 \wedge S_2 \wedge S_3$  that Special Theory of Relativity is true *and* that the relative velocity between Clock 1 and Clock 2 is 0.995c *and* that Clock 1 ticks every second relative to its rest frame. This is because  $S_1 \wedge S_2 \wedge S_3 \rightarrow e$ , such that:

$$P(e|S_1 \wedge S_2 \wedge S_3) = 1.$$

Applying Bayes' Theorem:

$$P(S_1 \wedge S_2 \wedge S_3|e) = \frac{P(e|S_1 \wedge S_2 \wedge S_3).P(S_1 \wedge S_2 \wedge S_3)}{P(e)} = \frac{1}{P(e)}.P(S_1 \wedge S_2 \wedge S_3).$$

Given that  $P(e) \leq 1$ ,  $E_J$  holds and  $e$  provides evidential justification to the conjunction  $S_1 \wedge S_2 \wedge S_3$ . Since  $P(e|S_1 \wedge S_2 \wedge S_3) = 1$ , this also implies that one of the conditions for non-granular evidential justification also holds:

$$\frac{P(e|S_1 \wedge S_2 \wedge S_3)}{P(e)} \geq 1$$

However, the evidence  $e$  that Clock 2 is ticking every 10 seconds relative to Clock 1 is probabilistically independent from each of the propositions  $S_1$ : 'Special Relativity Theory is true',  $S_2$ : 'The relative velocity between two clocks, Clock 1 and Clock 2, is 0.995c' and  $S_3$ : 'Clock 1 ticks every second relative to its rest frame'. So  $e$  does not provide evidential justification to either  $S_1$ ,  $S_2$ , or  $S_3$ , taken independently, that is:

$$\forall i \in \{1, 2, 3\}, P(S_i|e) = P(S_i).$$

So, our scenario is a case of non-granular evidential justification.

To conclude, Bayesian-style accounts of evidential justifications are quite common e.g. (Williamson, 2002), (Howson and Urbach, 2006), and (Easwaran, 2011), are fully compatible with  $G_T$ , while also consistent with situations in which  $G_J$  is violated. In other words, the core premise of Elgin's argument is undermined and her conclusion remains unsupported. Elgin's granularity argument against veritism is misguided: just because one takes truth to be epistemically central and truth to be granular, it does not follow that one has to take justification to be granular and reject all forms of holistic justification.

<sup>3</sup> The example is inspired by the chapter on "Rods and Clocks" of Norton's open-source course "Einstein for Everyone", available at [https://www.pitt.edu/~jdnorton/teaching/HPS\\_0410/chapters/Special\\_relativity\\_clocks\\_rods/index.html](https://www.pitt.edu/~jdnorton/teaching/HPS_0410/chapters/Special_relativity_clocks_rods/index.html).



### 3.2 Maximal Inter-Dependence

Elgin deploys a second line of argument in support of her claim that veritism conflicts with holistic confirmation. This second line of argument is rather weak on several fronts, some of which will be discussed below. What is interesting however, is that the argument sheds some light on Elgin's own approach and how exactly her account differs from other forms of (partial) veritism.

The second line of argument that Elgin deploys relies on a notion of "contamination" of an entire scientific account by a few falsehoods (Elgin, 2018, 33 sq.). The conclusion she wants to draw is that scientific accounts cannot be taken to provide knowledge. The argument is not always clear, but a rough idea would be something along the following lines. Any scientific account contains "falsehoods". Taking holism seriously implies that one admit that these falsehoods cannot be "excised" from the account as a whole, and so, the argument goes, the account cannot be regarded even *partially* true. So, any notion that scientific accounts can provide (even partial) knowledge is, according to Elgin, undermined by holism. Now of course, this seems to be a mistake: a falsehood would make it impossible for the account to be true *simpliciter*, but not *partially* so. Now, while Elgin's conclusion outstrips her arguments, these arguments are still interesting because they point to the difference between her account and other versions of veritism.

The most obscure aspect of this line of argument is the notion that falsehoods cannot be "excised" and so the whole account turns out false. This seems to presuppose that any falsehood will have a role to play in any attempt at justifying the truth of each of the other components of the account. As a consequence, the argument would go, any falsehood in an account will somehow epistemically compromise each of the other components of the account. If this generalizes, then the assumption would be that taking holism seriously implies that each proposition (true or false) of the account must be taken to play a role in the justification of any other proposition. Assuming this kind of maximal inter-propositional justificatory dependence seems to be the only way to bring about the conclusion that any falsehood within the account compromises the whole account. It is unclear how such an assumption would follow from holism.

Indeed, on Elgin's account, holism is the view that evidence provides support to the whole account rather than the particular propositions (Elgin, 2018, 35):

Even if the focus of attention is a particular claim, evidence always bears on an account as a whole. Evidence for the claim that a given process is adiabatic is evidence for an entire account of heat transfer, which is tested along with the claim. Evidence that the grooming behavior displays reciprocal altruism is evidence for an entire account of primate behavior, which is tested along with the claim.

From this notion, however, it does not follow that any proposition of the account does play a role in the justification of all of the other propositions. For example, the proposition that grooming behavior is an instance of reciprocal altruism does not necessarily play a role in the justification of the proposition that particular forms of sexual encounters are instances of mechanisms for dominance hierarchies to be implemented, another aspect of some theories of primate behavior. One could argue that the claim regarding dominance hierarchies has been largely falsified, along with some significant portions of our account of primate behavior, while the claim regarding reciprocal altruism remained well-confirmed. Again, from Elgin's notion of holistic justification, it does not follow that each and every proposition in the account plays a role in justifying each and every of the other propositions in the account.

For the purpose of this paper, there is no need to investigate any further the logic and plausibility of the argument from maximal inter-dependence. More interesting is to understand how this assumption points to an important aspect of Elgin's view, which distinguishes her view from other (partial) veritists, and has to do with the question of whether or not falsehoods contribute to the epistemic value of representations.

Elgin's line of argument stands in contradiction with some veritists' claim that falsehoods have little or no role to play in the justification of a scientific account. Strevens, for example, would simply claim that the false components of an account do not matter, as they are not of difference-makers. The components of the account that are true and epistemically valuable are the components representing difference makers. Everything else, including falsehoods regarding non-difference-makers, is of no epistemic value and plays

no significant role in the assessment of the epistemic acceptability/value of the account. Falsehoods, the argument goes, are not conducive of epistemic value.

Elgin would respectfully disagree: she wants to argue for a more encompassing version of holism, in which falsehoods have a positive role to play into the epistemic acceptability and epistemic value of scientific accounts. She wants to defend the claim that falsehoods contribute to the accounts in important ways. In other words, Elgin is mistaken about her target: she is a veritist who holds that partial truth is central to epistemic acceptability and value, but she rejects what could be called “segregationist veritism”. The segregationist veritist holds that partial falsehoods are to be kept separated from the partial truths. They may not have any positive part to play regarding the interdependent support of partial truths. They may not have any positive contribution to play regarding the epistemic acceptability of the account. They add no extra epistemic value to it. Elgin’s view does stand in direct contradiction with these claims.

To conclude, then, Elgin’s argument against veritism, according to which no partial-truth account can be compatible with the holistic nature of justification in the sciences, fails. It relies on the false premise that holism implies some form of maximal-interdependence for the justification of propositions within scientific accounts. That said, Elgin’s argument is revealing of an important difference between her version of veritism and others. Section 4 will be devoted to an explanation of Elgin’s view on the positive contribution that falsehoods can have to the epistemic value of scientific accounts. Before I turn to that topic, however, there is one more argument against veritism that needs to be critically assessed. There again, I will argue that her argument fails and her conclusions remain unsupported.

### 3.3 Factivity

The last line of argument against veritism that I would like to address goes along the following lines:

- (P1) Veritism is committed to knowledge and understanding being factive;
- (P2) Scientific understanding is not factive;
- So veritism is not an adequate view of science.

In what follows, I will argue that Elgin’s argument in support of (P2) fails. Note that it does not follow that scientific understanding is best conceived as factive. Kvanvig (2003), for example, would disagree. Setting that question aside, the point I want to make is that the particular argument that Elgin advances in support of (P2) fails, so that her conclusion that veritism is not an acceptable view of science remains unsupported.

The question of the factivity of understanding is debated in the literature (Grimm, Baumberger and Ammon, 2017, 7-11). Elgin has been defending the view that understanding is not factive for quite a while (See for example her (2009) as well as Chapter 3 of her (2018)). As a reminder, knowledge is generally taken to be factive, i.e. one cannot know that  $p$  unless  $p$  is true. Her argument against the factivity of understanding in her (2018) proceeds from the following assumptions:

- (a) science makes progress,
- (b) scientific accounts may have various proportions of true and false components, and
- (c) people learn scientific accounts at various degrees of precision and depth.

Elgin argues that (a), (b), and (c) imply that the kind of epistemic value afforded by scientific accounts *comes in degrees*. She takes this to imply that the kind of epistemic value afforded by scientific accounts cannot be factive. This is because, according to Elgin, factivity requires truth and truth is an all-or-nothing business. Factivity would thus leave no room for degrees of epistemic value.

According to Elgin, since knowledge is generally taken to be factive, this implies that knowledge is not a good candidate for the epistemic value of scientific accounts. By contrast, she argues, understanding can come in degrees, so it is a good candidate. And indeed (a), (b), and (c) can be interpreted in terms of understanding. Scientific progress can be measured in terms of increased understanding. Scientific accounts may afford various degrees of understanding depending on the proportions of true and false components. Finally, one can learn scientific accounts to various degrees of precision and depth.

Now if understanding was factive, the argument goes, it would be an all-or-nothing business, again along the argument that anything factive requires truth and that truth is an all-or-nothing business. So understanding could not come in degrees, and none of (a), (b), and (c) could be true. Overall, then, the argument goes, it is appropriate to claim that scientific accounts are epistemically valuable because they afford understanding, conceived as being non-factive.

The argument needs unpacking. First, we need to specify the target of analysis. Elgin seems to have at least two possible targets in mind in the argument above: (a) and (b) focus on the various degrees of understanding that scientific accounts afford of their target phenomena and (c) focuses on the various degrees of learning that agents may have of scientific accounts. The debate over whether or not scientific understanding is factive has focused much more on the former and much less on the latter, if at all. So let us concentrate for now on the question of whether or not the kind of understanding that scientific accounts afford of their target phenomena has to be non-factive in order to come in degrees.

There are some contexts in which some of Elgin's assertions are true. If our focus is on "knowledge that  $p$ ", where  $p$  is a single proposition, then factivity imposes that one knows that  $p$  only if  $p$  is true. Whether or not a single proposition  $p$  is true is an all-or-nothing business. So, knowledge that  $p$ , because it is factive, cannot come in degrees. Similarly, consider "understanding that  $p$ ": factivity requires again that  $p$  be true for understanding that  $p$  to arise. Given that  $p$  is strictly either true or false, understanding that  $p$ , if factive, must be an all-or-nothing business. In other words, understanding that  $p$ , if interpreted as factive, cannot come in degrees.

However, the debate on whether or not understanding is factive is not about "understanding that  $p$ ", but rather on "understanding why  $p$ ". In this context, it becomes less clear whether the argument advanced by Elgin, i.e. that understanding cannot possibly both be interpreted as factive and as admitting of degrees is successful. For simplicity, let us take understanding why  $p$  to be equivalent to understanding that ( $p$  because of  $q$ ) (Grimm, 2006; Pritchard, 2010; Hills, 2016). In other words, let us define understanding why as directly deriving from explanatory knowledge. What does it mean for understanding why  $p$  thus conceived to be factive?

There is first a rather trivial sense in which "understanding why  $p$ " can be interpreted as factive. It seems reasonable to claim that one cannot understand why  $p$  – i.e. understand that ( $p$  because of  $q$ ) – unless one understand that  $p$  in the first place. Now, as mentioned before, understanding that  $p$  cannot come in degrees if interpreted as factive. So, Elgin's argument applies at least to this rather uninteresting component of understanding why  $p$ .

Now there is another, more interesting component of understanding why  $p$ , more precisely  $q$  and its relation to  $p$ . Factivity in this context is not about the truth of  $p$  but rather about the truth of  $q$  as an explanation for  $p$ . This is more interesting because explanations are typically not simple. If explanations were simple, then there would be a case to make for the factivity of understanding why  $p$  to be incompatible with its admitting of degrees. For example, let  $q$  be a non-compound proposition. In this case, if understanding why  $p$  is nothing else than understanding that ( $p$  because of  $q$ ), then understanding why  $p$  is factive is the requirement that one understands why  $p$  only if ' $p$  because of  $q$ ' is true, which in turns requires that  $q$  be true. Now whether  $q$  is true or not is an all-or-nothing business and does not admit of degrees. So, if  $q$  is a non-compound proposition, then understanding why  $p$  cannot both be factive and admit of degrees (assuming, for simplicity's sake, that whether  $q$  relates to  $p$  in such a way that it counts as a proper explanation for  $p$  is also an all-or-nothing business). Now explanations, and especially scientific explanations such as the ones given by scientific accounts in which Elgin is interested, are complex. In other words,  $q$  will typically consists in a multitude of propositions entertaining a variety of logical relationships. The question is thus whether or not Elgin's argument, according to which scientific understanding cannot be interpreted as factive because it cannot admit of degrees if it is, works in the case in which  $q$  is complex.

It is possible to have a view of scientific understanding that is factive and also comes in degrees. While the truth of a single non-compound proposition is indeed an all-or-nothing business, the truth of an account is not. A set of propositions  $Q = \{q_1, \dots, q_n\}$  might be considered as more or less "true" depending on how many of the  $q_i$  are true.<sup>4</sup> This implies that a scientific account of  $p$ , in which the explanation for  $p$  consists of multiple propositions  $Q = \{q_1, \dots, q_n\}$ , may afford various degrees of explanatory understanding of its

<sup>4</sup> Note that my argument does not rely on the claim that scientific accounts are best conceived of as sets of propositions. That question is very much open. That said, that scientific accounts, including models, can be (ought to be) conceived as collections of propositions has been defended in the literature (Thomson-Jones, 2012). Note that authors who propose

target domain, depending on how many of the  $q_i$  are true. Interpreted this way, scientific understanding afforded by scientific accounts may be consistently conceived as both factive and admitting of degrees. The degrees of explanatory understanding afforded by a scientific account would indeed be directly related to the degree of truth contained in its explanation.

Let us see how this could work with an example. Let us consider a scientific account of phenomena  $p$ , according to which ( $p$  because of  $Q$ ). Let us now take a causal view of explanation, so that understanding why  $p$  is now interpreted as directly deriving from the “knowledge of the causes for  $p$ ”. Finally, let us take a Strevensian view of causal explanation, so that knowledge of the causes for  $p$  is just knowledge of difference-makers  $q_i$ . Difference makers are just these factors which pass some kind of “but for” test: the phenomena of concern would not occur but for the influence of the difference maker. There is no a priori reason to believe that a scientific account could not provide some but not all difference makers for some phenomena of concern, i.e. some but not all of the  $q_i$  are true difference makers for  $p$ . If that’s correct, then it is not logically inconsistent to take that the account can afford various degrees of understanding of  $p$  as directly depending on how many difference makers for  $p$  it accurately represents (how many of the  $q_i$  are true). In a view like this, understanding is factive – one understands why  $p$  only if some causes  $q_i$  are true causes for the phenomena – and still comes in degrees.

There might other ways to articulate a notion of factive-but-possibly-partial knowledge and hence factive-but-possibly-partial understanding afforded by scientific accounts of their target domain, but the example above suffices for the purposes of my argument. It shows that it is possible to consistently conceive of scientific understanding as both factive and admitting of degrees. Hence, when we focus on ‘understanding why’ afforded by scientific accounts of their target phenomena, Elgin’s argument according to which understanding cannot possibly be factive and admitting of degrees is undermined. If this is correct, then (P2) in her overall argument against veritism is false, and her conclusion that veritism cannot be an acceptable view of science remains unsupported.

What about the other claim that Elgin makes, (c) that scientific accounts can be learned to various degrees? It would appear that this is no barrier to understanding being factive. For a scientific account to be a vehicle of understanding, one presumably has to learn something non-trivial about the account. That said, one presumably doesn’t have to learn everything about an account for it to be a vehicle of understanding. E.g. anyone that has done well in an introductory physics class will have learned enough about Newton’s theory to model blocks sliding down inclined planes. On at least some theories of factive theories of understanding, like Strevens (2008), these students will have some degree of understanding of blocks sliding down inclined planes insofar as they have grasped some of the difference makers for the corresponding behavior, e.g. mass of the block, angle of inclination, etc. Now, not everyone can use Newton’s theory as a vehicle of understanding for everything Newton’s theory can be a vehicle of understanding for. E.g. introductory students will have no idea how to model fluid dynamics using Newton’s theory. When someone hasn’t learned a theory well enough for it to be a vehicle of understanding for something, then the issue of factivity of understanding just doesn’t arise. It would be a mistake to say that it fails. If this is correct, then (c) is no obstacle to scientific understanding to be both factive and admitting of degrees. Thus again (P2) is undermined, and Elgin’s conclusion remains unsupported.

To conclude, Elgin’s view is perfectly compatible with the notion that scientific accounts contain truths and that these truths are core to their epistemic value. It is also compatible with the claim that understanding is factive and that science provides partial knowledge from which we may derive partial understanding. That said, Elgin also wants to steer away from the notion that scientific falsehoods in scientific accounts are acceptable *only* because of their negligible character. On Elgin’s view, scientific falsehoods are acceptable not only because we can ignore them, but also because they may have a positive role to play. They are, famously, “felicitous falsehoods”. What role exactly felicitous falsehoods are taken to play within Elgin’s account is the focus of Section 4.

#### 4 Felicitous Falsehoods

In Section 2, I have argued that, on Elgin’s view, partial truth is central to the epistemic value of scientific representations that contain falsehoods. In Section 3, I made the case that, while Elgin’s arguments

---

alternative views such as the semantic view offer alternative ways to measure of degrees of truth and corresponding epistemic value (da Costa and French, 2003).

against veritism and factivity of understanding failed, these arguments suggest some clues regarding how Elgin's view differs from the more common versions of veritism. This section is devoted to articulating that difference more precisely. The following quote is helpful to set the stage (Elgin, 2018, 2):

Although models, idealizations, and thought experiments are inaccurate – sometimes wildly inaccurate – they exemplify features they share with their targets.... They are true enough.

Note first that this quote illustrates my previous point that, on Elgin's view, partially inaccurate scientific representations are epistemically valuable only if they "share" some features with their target systems. The features that are "shared" between the representation and the target are accurately represented. Partial truth is thus core to epistemic value.

Now another crucial element of Elgin's account is also illustrated in the quote above: the concept of *exemplification*. Exemplification is the core concept that makes her account of the epistemic value of scientific representations strikingly different from other versions of veritism. As I suggested in the previous section, Elgin's account differs from other versions of veritism in that falsehoods are not only *acceptable* because we can ignore them, but they are also *valuable* in that they have a positive role to play. Exemplification is the process in which falsehoods may be positively involved, that is, in which they may be "felicitous".

Let us start with a reminder. A good foil to Elgin's view on that matter is Strevens' view. On Strevens' view, epistemic value is rooted only in partially accurate representation of difference-makers for the target phenomena, while inaccurate representation of non-difference-makers is possibly necessary, albeit unfortunate, and (at best) neutral with regards to epistemic value. In other words, falsehoods about non-difference makers do not play any positive role in the representational role of scientific representations and do not contribute to their epistemic value.

Elgin agrees that to distort features of the target system is acceptable only if these features can be considered negligible and the scientific representation remains "true enough" despite these distortions. However, she begs to differ with the claim that these "falsehoods" are epistemically neutral. Falsehoods, according to Elgin, play a positive and important role; they contribute to the epistemic value of scientific representations, in that they help *reveal and highlight* those features that are accurately represented ("shared" with the target systems) (Elgin, 2018, 258):

... scientific models exemplify properties and afford epistemic access to them. By omitting or downplaying the significance of confounding factors (... intermolecular attraction in the ideal gas, friction in the model pendulum), they constitute a cognitive environment where certain aspects of their subjects stand out. They thereby facilitate recognition of those aspects and appreciation of their significance. They thus give us reason to take those aspects seriously elsewhere.

So, as mentioned before, the starting point is that scientific representations are epistemically valuable only if they accurately represent some features of their targets. That said, real-world target systems are typically too complex and too multi-faceted for us to fully comprehend. It follows from this that a perfect, complete replica of a target system, one which would share all the features with that system, would be, in many ways, cognitively useless. That is because the replica would be just as difficult to comprehend as the target system itself. No understanding would be afforded by the replica. But, according to Elgin, what makes scientific representations epistemically valuable is precisely that they afford scientific understanding of their target phenomena. For that reason, scientific representations *must* focus on *selected* features of the system, those we deem important in the particular epistemic or practical context in which we want to use these scientific representations. This further implies that they must misrepresent some other features of the system, perhaps a majority of them. So, while the epistemic value of scientific representations crucially stems from the truths that they contain, it also stems from the fact that they are not fully true, but only partially so. Containing some mixture of truths and falsehoods is thus necessary for epistemic value.

If this is true, Elgin argues, then there is a sense in which falsehoods positively contribute to the epistemic value of scientific representations. They do so, Elgin explains, by giving *epistemic access* to those truths contained in scientific representations. In other words, scientific representations are epistemically valuable not only because they "share" some features with their target systems but also because they make these features salient. To make the truth salient is the positive role played by falsehoods. It is in that sense that "falsehoods" can be "felicitous", i.e. whenever they provide epistemic access to features that otherwise

would be lost, blurred, or obscured. In other words, that scientific representations include falsehoods is a feature, not a bug.

The property that scientific representations have by which they make salient those features of the target system they accurately depict is what Elgin refers to as “exemplification”.<sup>5</sup> Elgin has articulated this notion in numerous publications (e.g. Elgin (2004; 2006; 2007; 2008; 2010)). It is truly the most original and most interesting part of Elgin’s work. It is also what makes her account depart from more traditional views à la Strevens, in which falsehoods at best do not contribute and most often are detrimental to the epistemic value of scientific representations. By contrast, on Elgin’s account, some falsehoods positively contribute to the epistemic value of representations by *facilitating exemplification*.

One could object at this point that, if the analysis above is correct, then Elgin’s account may not be fully veritist. The objection would go as follows. Falsehoods are taken to be epistemically valuable in her account. But falsehoods cannot be possibly intrinsically valuable under a veritist account. So, if Elgin’s account is veritist, falsehoods may be epistemically valuable only if they are instrumentally so. However, according to the analysis above, falsehoods are epistemically valuable because they increase the understandability of scientific representations. As a consequence, if understandability is interpreted within Elgin’s account as independent of truth, then it would seem that Elgin’s account is not veritist after all: felicitous falsehoods do not gain epistemic value via any truths.

The main premise of the objection is that understandability may be interpreted as independent of truth. The premise seems plausible. Several authors have recently articulated accounts of scientific understanding that interpret understanding in a way that is relatively independent from considerations of truth. This includes the views according to which understanding is either an ability to foresee how models apply in particular situations without needing to go through any explicit computation (de Regt and Dieks, 2005), or is related to how “compressed” information is for easy mental use (Wilkenfeld, 2018), or an ability to navigate the space of representations for the target systems (Le Bihan, 2017).

The problem is that Elgin’s notion of understanding does not fall into that category. She never articulates the notion of understanding in any other way than as a cognitive state that is achieved by having epistemic access to these features that are truly depicted within a scientific representation. On this account, understanding is both directly and indirectly related to truth. It is directly related to truth as truth is the core vehicle of understanding. As Elgin likes to put it, one does not understand combustion via phlogiston theory (Elgin, 2018, 61). It is indirectly related to truth because whatever aspects of understanding are not directly related to truth are interpreted as having epistemic value only to the extent that they allow for truth to be revealed. Thus, when it is claimed that falsehoods have epistemic value, it is indeed compatible with veritism in an instrumental way, as their value is derivative vis-à-vis the disclosure of truth. If the analysis above is correct, then it is not the case that Elgin’s notion of understanding is independent of truth, and the objection fails.

An additional way to address the objection above is to note that the notion of veritism it assumes is not consistent with Elgin’s own definition. The objection assumes that veritism is the view according to which truth is the sole, necessary and sufficient standard of epistemic value. I have argued in Section 2, however, that this is not the definition that Elgin adopts. On her account, veritism is the view according to which:

1. considerations of truth are necessary, even if not always sufficient, when assessing epistemic value;
2. when other considerations apply, truth is always of central concern.

This notion of veritism is consistent with Elgin’s view according to which:

1. the epistemic value of scientific representations stems from their truly depicting *and* making salient – i.e. exemplifying – some features of their target systems, and
2. felicitous falsehoods facilitate exemplification.

On Elgin’s account, that representations *truly* depict some features of their target systems is necessary, albeit not sufficient. Other considerations apply, indeed that these features be made salient, but then truth still remains of central concern. Veritism holds.

---

<sup>5</sup> The extent to which Elgin’s usage of the term “exemplification” is related to Goodman’s views will not be addressed in this paper.

Elgin is thus a veritist, but a veritist who has unveiled an interesting and important aspect of scientific representation, i.e. that some falsehoods positively contribute to the epistemic value of scientific representations by playing the important role to facilitate exemplification.

## 5 Conclusion

In this paper, I argued the following. First, by Elgin’s definition of veritism, i.e. that truth is taken to be a necessary condition and central concern for epistemic value, Elgin’s own account is veritist, despite her claims to the contrary. Second, Elgin’s arguments against veritism and the factivity of understanding are faulty, mainly because they rely on the implausible assumption that veritists/factivists have to be absolutists about truth and reject partial truth as epistemically valuable. Third, what makes Elgin’s account differ from more traditional version of veritism is not the notion that truths are not important, but rather that falsehoods can be important too. “Felicitous Falsehoods” are epistemically valuable because they provide epistemic access to the features that are accurately depicted in scientific representations. They thus facilitate scientific understanding and are instrumental to the epistemic value of scientific representations (Elgin, 2018, 16):

My thesis is that some representations that figure ineliminably in tenable accounts make no pretense of being true, but are not defective on that account. Indeed, I will argue, their deviations from truth are epistemically valuable – often more valuable than the unvarnished truth about the phenomena would be.

That falsehoods can positively contribute to the epistemic value of scientific representations is possibly the most important insight of Elgin’s work.

## 6 Acknowledgments

I would like to thank Armond Duwell, Kareem Khalifa, Henk de Regt, Stephen Grimm, and of course Kate Elgin for multiple illuminating and friendly conversations on scientific understanding. I would also like to recognize various institutions for the support I received from them while developing my views on scientific understanding: the University of Pittsburgh Center for Philosophy of Science, the IHPST in Paris, SND (FRE 3593 Paris-Sorbonne), and Labex Transfer (Ecole Normale Supérieure, Paris).

## References

- da Costa, N. and French, S. (2003). *Science and Partial Truth: A Unitary Approach to Models and Scientific Reasoning*, Oxford University Press, Oxford.
- de Regt, H. W. and Dieks, D. (2005). A contextual approach to scientific understanding, *Synthese* **144**: 137–170.
- Easwaran, K. (2011). Bayesianism I: Introduction and arguments in favor, *Philosophy Compass* **6**(5): 312–320.
- Elgin, C. (2006). From knowledge to understanding, *Epistemology futures* pp. 199–215.
- Elgin, C. (2010). Telling instances, in R. Frigg and M. Hunter (eds), *Beyond Mimesis and Convention*, Vol. 1 (262) of *Boston Studies in the Philosophy and History of Science*, Springer, pp. 1–17.
- Elgin, C. Z. (2004). True enough, *Philosophical issues* **14**: 113–131.
- Elgin, C. Z. (2007). Understanding and the facts, *Philosophical Studies* **132**: 33–42.
- Elgin, C. Z. (2008). Exemplification, idealization, and scientific understanding, *Fictions in Science*, Routledge, pp. 85–98.
- Elgin, C. Z. (2009). Is understanding factive?, in A. Haddock, A. Millar and D. Pritchard (eds), *Epistemic Value*, Oxford University Press, pp. 322–30.
- Elgin, C. Z. (2018). *True Enough*, MIT Press.
- Grimm, S. (2006). Is understanding a species of knowledge?, *British Journal for the Philosophy of Science* **57**: 515–36.
- Grimm, S., Baumberger, C. and Ammon, S. (2017). *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, Routledge.
- Hills, A. (2016). Understanding why, *Nous* **50**(4): 661–688.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12092>
- Howson, C. and Urbach, P. (2006). *Scientific reasoning: the Bayesian approach*, Open Court Publishing.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press.
- Le Bihan, S. (2017). Enlightening falsehoods: A modal view of scientific understanding, in S. Grimm, B. C. and A. S. (eds), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, Routledge, pp. 111–35.
- Mizrahi, M. (2012). Idealizations and scientific understanding, *Philosophical Studies* **160**(2): 237–252.
- Norton, J. (2004). On thought experiments: Is there more to the argument?, *Proceedings of the 2002 Biennial Meeting of the Philosophy of Science Association, Philosophy of Science* **71**: 1139–1151.

- 
- Potochnik, A. (2017). *Idealization and the Aims of Science*, University of Chicago Press.
- Pritchard, D. (2010). Knowledge and understanding, in D. Pritchard, A. Millar and A. Haddock (eds), *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press, chapter 1-4.
- Sosa, E. (2007). *A Virtue Epistemology*, Oxford University Press.
- Strevens, M. (2008). *Depth*, Harvard University Press.
- Thomson-Jones, M. (2012). Modeling without mathematics, *Philosophy of Science* **79**(5): 761–772.
- Treanor, N. (2013). The measure of knowledge, *Noûs* **47**: 577–601.
- van Fraassen, B. (2008). *Scientific Representation: Paradoxes of Perspective*, Oxford University Press.
- Wilkenfeld, D. A. (2018). Understanding as compression, *Philosophical Studies* pp. 1–25.
- Williamson, T. (2002). *Knowledge and its Limits*, Oxford University Press on Demand.