TRUTH, UNDERSTANDING, AND ELGIN'S CHALLENGE: A NEW CASE FOR VERITISM

Tiegue VIEIRA RODRIGUES, Alexandre ZIANI DE BORBA

ABSTRACT: It is often said that truth is the fundamental value of epistemic norms. Values such as rationality and evidence are valuable insofar as they help to avoid false beliefs and achieve true ones. This view is generally called *veritism*. However, in a recent work, Catherine Elgin (2017) argued that veritism cannot account of the epistemic status of science, since science accepts models and idealizations that are known to be false. Elgin's argument can be reconstructed in the form of an inconsistent triad, from which she concludes that veritism's central prohibition on accepting falsehoods must be rejected. This article offers an alternative defense of veritism. First, we re-evaluate the role of falsehoods in epistemology, drawing on the current debate about *Knowledge from Falsehood* (KFF) to show that the absolute prohibition criticized by Elgin is already under pressure, but for reasons more complex than they appear at first glance. Next, we argue that Elgin's conclusion against veritism is premature. Using precisely the distinction Elgin makes between the attitudes of acceptance and belief, we show that a correctly specified veritism, whose central norm regulates belief exclusively, is not threatened. Finally, we go beyond a purely defensive stance and argue that this framework provides the resources for a robust veritistic account of *understanding*, thus meeting Elgin's deeper challenge.

KEYWORDS: Catherine Elgin, veritism, harmless falsehoods, acceptance, belief, epistemic conscientiousness, knowledge from falsehoods, understanding

1. Introduction

Veritism is a popular, though not uncontroversial, axiological thesis among epistemological theorists (e.g., Goldman 1999; Pritchard 2021). In essence, it is the formal expression of a deeply intuitive idea: that our cognitive life is, in a fundamental sense, oriented towards truth. According to *veritism*, truth is the *supreme* epistemic value, the standard by which all other cognitive goals and practices are measured. Other epistemic goods (such as evidence, critical thinking, or open-mindedness) derive their value from their relation to this primary goal; they are valuable precisely because they are reliable means for achieving the end of obtaining true beliefs and avoiding false ones.[16] This truth-centered structure seems

[16] While Goldman (1999) is a key proponent of this teleological veritism, the view has faced significant criticism. For foundational critiques of veritism's ability to ground all epistemic value, see Kvanvig (2003), who argues for understanding as a distinct and more valuable goal, and Elgin

too grand to be a plausible and elegant basis for the theory of knowledge, but Catherine Z. Elgin (2017) has offered a strong critique of it, arguing that *veritism* is in fundamental contradiction with the practices of our most successful cognitive endeavor: modern science.

Elgin's critique can be reconstructed as an inconsistent triad, a set of three plausible but mutually incompatible claims:

1. There can never be epistemically good reasons to accept a known falsehood.
2. Modern science is epistemically praiseworthy.
3. Modern science accepts models and idealizations that are known not to be true.

Each thesis carries considerable philosophical weight. Statement (1) articulates a central veritistic obligation that can be traced back to the *no false lemmas* principle, which arose from the analysis of Gettier cases.[17] It encompasses the intuition according to which truth cannot be built on a foundation of falsehood. Claim (2) reflects the unanimous opinion that science is a model of epistemic achievements. In fact, the long-standing philosophical project of distinguishing science from pseudoscience is based on the idea that the label 'scientific' confers a special epistemic value. Finally, claim (3) is a sociological and historical observation about scientific practice. It is acknowledged that scientists are comfortable using models (such as the ideal gas law, which treats molecules as dimensionless points with no intermolecular forces) that are known to be literally false but, nonetheless, are indispensable for predictions and explanations.

In face of this inconsistency, one must give up of something. Elgin's solution is to reject thesis (1), which, in her opinion, requires the total rejection of *veritism*. Although most responses[18] to Elgin focus on contesting thesis (3) (for example, by arguing that scientists do not really *accept* falsehood, but rather some true theses about models), our analysis follows Elgin in rejecting (1) but contests that this step is fatal to *veritism*.[19] However, it is important to recognize that, while structurally

(2017), who argues that this formulation of veritism is incompatible with the practice of science. See also Pritchard (2021) for a recent analysis of these challenges.

[17] See Gettier (1963). The *no false lemmas* condition emerged as a primary strategy to save the *Justified True Belief* (JTB) account of knowledge. Classic proponents of this fix include Clark (1963) and Lehrer (1965), who argued that knowledge cannot be inferred from a false belief, as any such inference would make the resulting true belief only accidentally true. Our paper challenges the scope, not the spirit, of this prohibition.

[18] See Nawar (2021); Le Bihan (2021); Rice (2021); Lawler (2021); Frigg and Nguyen (2021); see also Elgin (2021) for replies.

[19] Many commentators on Elgin's *True Enough* (2017) have taken this first path, attempting to show that scientific models are not *felicitous falsehoods* but rather *partial truths* (Le Bihan 2021),

useful, this triad represents a specific line of attack within Elgin's broader, more holistic challenge. Her central claim is not merely that science uses falsehoods, but that the real *goal* of research is often *understanding*, not truth, and that our best epistemic practices reflect this alternative goal. For Elgin, science's reliance on *felicitous falsehoods* is a symptom of a deeper discrepancy between *veritism* and the real *telos* of our most successful cognitive endeavors.

Before we proceed, we also need to clarify the specific goal of our defense. Veritism is not a monolithic doctrine. Consequentialist or teleological versions, which assume that epistemic correctness consists in promoting or maximizing the total number of true beliefs, are notoriously susceptible to objections such as the *swamping problem* (in which the value of a reliable process is overshadowed by the value of the true belief it produces) and the *trivial truths problem* (which suggests that a *veritistic* would devote himself to the pursuit of an infinity of trivial truths) (see Goldman 2002; Pritchard 2021).[20] Our defense is more restricted and, as we claim, more robust. We defend a *normative veritism* that focuses on the central prohibitive norm governing the attitude of believing: one should avoid *believing* a falsehood. This *deontological* restriction is, as we have stated, at the heart of the veritistic's commitment, and by focusing our defense on it, we protect our argument against a series of standard objections that seek to reject consequentialism.

Our argument unfolds in the following stages. We begin in Section 2 by moving beyond a simple reliance on the concept of 'Knowledge from Falsehood'. Instead, we situate the discussion in the more nuanced contemporary debate between the KFF and *Knowledge Despite Falsehood* (KDF) theses. This allows us to establish a critical distinction that underpins our entire project: the difference between reasoning from an unknown false belief and the deliberate scientific practice of *accepting* a *known* falsehood. With this groundwork laid, Section 3 presents our central defense. We argue that the primary *veritistic norm* (to hold only truths) applies exclusively to the attitude of belief, whose essential aim is to represent the world accurately. We then show that strategically *accepting* a known falsehood

---

that they provide access to facts about the model (Nawar 2021), or that their epistemic value can be explained through a *understanding-based realism* (Rice 2021). Our approach concedes more to Elgin's setup – accepting the rejection of (1) – to build a more robust defense of (2).

[20] The *swamping problem*, most prominently discussed by Kvanvig (2003) and Pritchard (2010), holds that the value of a reliable belief-forming process (like perception) is *swamped* by the value of the true belief it produces, making reliability instrumentally valuable but not epistemically valuable in its own right. The *trivial truths* problem, also a staple of these critiques, points out that if truth were the sole epistemic good, an agent could maximize this good by memorizing a phone book, an outcome that seems epistemically worthless. Our deontological focus on the norms of believing sidesteps these objections, as it is not a maximizing or consequentialist thesis.

is not a violation of this norm, but rather an act of *instrumental epistemic rationality* performed in the service of truth. Building on this, Section 4 transitions from a defensive posture to a constructive one. Here, we outline how our framework provides the tools for a robust, truth-based theory of *understanding*, directly addressing the most significant aspect of Elgin's challenge. In order to test the resilience of our account, in Section 5, we anticipate and respond to two key objections: the charge that our view is too permissive and the claim that it devalues the concept of understanding. Finally, Section 6 draws these threads together, offering our concluding remarks on what a more sophisticated and scientifically plausible veritism looks like.

2. Harmless falsehoods and the limits of prohibition

The first assertion of Elgin's triad (that there can never be good epistemic reasons for accepting a known falsehood) is in line with a strong and historically rooted intuition. This intuition is a direct philosophical descendant of the *'No False Lemmas'* principle, a condition for knowledge proposed in the wake of Gettier's famous counterexamples. In its simplest form, the principle states that *inferential* knowledge cannot be validly derived from a false premise or '*lemma*' (Audi 2011). If a person's *reasoning* is based *on* a falsehood, the resulting belief, even if true, is considered epistemically flawed, as its connection to the truth is merely accidental. In this context, this prohibition is applied in its strictest form, and its consequence is that any dependance on a falsehood is, by definition, an *epistemic error*. However, when one looks to contemporary literature there is a much more intricate and controversial relationship between falsehood, justification, and knowledge, suggesting that such a generalized prohibition may be untenable. At the heart of this debate is the lively discussion about the possibility of *knowledge from falsehood*.[21]

2.1 The debate between KFF and KDF

The philosophical landscape on the nature of knowledge has long been marked by a fundamental principle: one cannot obtain knowledge from falsehood. This idea, often summarized in the *No False Lemmas* principle, states that any conclusion derived from a false premise is epistemically falsified and cannot be considered true knowledge. However, this seemingly intuitive principle faces a significant challenge in the literature on *Knowledge from Falsehood* (KFF). Proponents of KFF present compelling cases in which individuals apparently acquire knowledge through a chain of arguments that includes a false belief. Warfield (2005, 412) says "that

[21] For the history of this topic before the publication of Warfield (2005), see De Almeida (2019).

relevant falsehoods sometimes play a central epistemizing role in inference." What is more, in recent years, several epistemologists have defended the possibility of knowledge from falsehood, forcing us to rethink the traditional understanding of how knowledge is acquired.[22]

KFF proponents often cite the *Inaccurate Clock* case as a paradigmatic case for their position. You go to the extremely reliable clock on your kitchen wall. You are used to trust the clock; it has never let you down before. The clock says 2:58 p.m. and you conclude that you won't be late for your appointment at 4 p.m. appointment. Intuitively, it seems that you know you won't be late. But unknown to you, your clock is two minutes slow; in fact, it is 2:56 p.m. As a matter of fact, your belief ('that it is 2:58 p.m.') is false. Nonetheless, the conclusion you were able to draw from it ('that you will not be late for your 4 p.m. appointment') is true and it seems like something you still know. But how is this possible? The KFF theorist would claim that this is a genuine case of knowledge from a falsehood, showing that the no false lemmas principle is too strong.

However, this interpretation is far from widely accepted. The main contradiction comes from the *Knowledge Despite Falsity* (KDF) thesis. KDF theorists argue that in cases such as the "inaccurate clock," the false belief does not actually perform epistemic work.[23] Instead, they claim that falsity is epistemically inert or merely a causal springboard, while knowledge is based on a separate, often unexpressed true belief. In the case of the clock, a KDF theorist might argue, for example, that your knowledge is not based on the false belief that 'it is *exactly* 2:58 p.m.', but rather on an associated true belief, such as 'it is *approximately* 3 p.m.' or 'according to my reliable watch, it is 2:58 p.m.'.

As an illustrative example, let's consider a comparison. Imagine a situation where a detective receives an anonymous tip: that the suspect is hiding in 'flat number 7B' of a certain building. The detective drives to the building, but the flat

[22] Proponents of KFF are numerous and the cases varied. See, for example, Warfield (2005) on deductive reasoning from false premises; Goldberg (2001) on knowledge from false testimony; Bernecker and Grundmann (2019) on knowledge from forgetting (which involves a false belief); and Fitelson (2010), Luzzi (2010, 2014, 2019), and Klein (2008) for general defenses. According to Elgin (2019), scientific knowledge and understanding can emerge from the use of approximations and idealizations, even though these are literally untrue. However, for a belief derived from an idealization to be classified specifically as knowledge from falsehood, two assumptions must be met: the idealizations must be considered complete falsehoods rather than simply approximate truths, and the scientists using them must be unaware that they are literally false. Both assumptions have been challenged by Sorensen (2013).

[23] The possibility of KDF is acknowledged by Goldman (1967, 368), Lehrer (1965, 169–71), Saunders and Champawat (1964, 9), and Swain (1981, 149–50).

numbers are confusing, and he fortuitously enters 'flat number 7D', which is right next door. There, he finds the suspect and arrests him. The original clue was wrong, the suspect was not in 7B. Did the detective obtain knowledge (the suspect's whereabouts) from false information? A KDF theorist would say no. The crucial information was that the suspect was 'in that area of the 7th floor', and the slightly inaccurate information was close enough to the truth to lead the detective to the right place. The wrong information ('7B') was not the actual source of knowledge.

Leading proponents of the KDF view, such as Ian Schnee (2015), have argued emphatically that there is no knowledge from falsehood. Schnee and others suggest that in every alleged case of KFF, one of two things happens: either the subject does not really have knowledge, or there is a hidden true belief that does all the work. According to this view, falsehood does not contribute to knowledge, but the subject knows *despite* its presence. This *resistance strategy*, as it is sometimes called, aims to preserve the traditional view of knowledge by reinterpreting the evidence in these complicated cases.[24]

The debate between KFF and KDF theorists is not just a philosophical curiosity but has a profound impact on our understanding of what knowledge is and how it is acquired. If KFF is possible, our theories of knowledge need to be more flexible and allow for a certain degree of *epistemic luck*, in which we may be wrong about some details but still be right overall. If, on the other hand, we assume that KDF is in fact the correct approach, then we can conclude that the traditional view according to which knowledge is understood as grounded in truth remains valid.

2.2 Distinguishing believing from accepting

Although the debate over *Knowledge from Falsehood* (KFF) represents a fascinating challenge to traditional epistemology, the direct application of its logic to the role of idealizations in science involves a subtle but profound category mistake. The entire KFF/KDF controversy is based on a particular psychological scenario: a subject who forms a false *belief* without recognizing its falsehood. The person in the *Inaccurate Clock* case genuinely *believes* that 'is exactly 2:58 p.m.'; for her, this statement is a representation of reality. She does not consciously *accept* it as a useful fiction or as an approximation that is sufficiently accurate for administrative purposes. Her reasoning stems from an unconscious error. This is the decisive context for the KFF

[24] Besides Schnee (2015), this *resistance strategy* is ably defended by Ball and Blome-Tillmann (2014), who argue that alleged KFF cases are really instances of *knowledge despite falsehood* where the false premise is epistemically superfluous, and the subject possesses a distinct, knowledge-grounding true belief.

debate: it examines whether knowledge can arise accidentally from an unknown falsehood.

Catherine Z. Elgin's defense of scientific idealization, however, concerns a wholly different epistemic practice: the conscious and deliberate *acceptance* of a falsehood in order to achieve a particular cognitive goal. This distinction between *believing* and *accepting* is not a mere philosophical subtlety; it describes two fundamentally different relationships that a subject can have with a proposition. To *believe* a proposition P is to regard it as true. Believing is, in the common philosophical representation, an involuntary mental state that *aims at truth* (Williams 1973).[25] We don't seem to be able to simply choose our beliefs. Instead, they appear to be imposed to us by our encounter with what we consider evidence. When one looks at a blue sky, one does not decide to believe that it is blue; one simply has that belief, it is irresistible for her.

Acceptance, on the other hand, is understood as a voluntary and context-dependent act. As Jonathan Cohen (1992) argued, *accepting* a proposition means treating it as true for a particular purpose and in a particular context, regardless of whether one believes it. In this sense, *acceptance* can be seen as a *pragmatic* stance, not necessarily a *doxastic* one. Consider a defense lawyer in court. She might say to the jury, 'Suppose my client was near the scene of the crime'. She definitely does not believe in this premise, and her entire case may be based on the fact that she is wrong. But she temporarily *accepts* it as a strategic move in order to show that even if it were true, it would not be enough to prove her client's guilt. Her acceptance, in this situation, is a voluntary decision, guided by a practical goal (persuasion) and immediately discarded as soon as the context of that specific argument passes.

This seems to exactly be the stance of a scientist who tends to use an idealization. For instance, a physicist who uses the ideal gas law, which models gas molecules as dimensionless points that exert no intermolecular forces, does not seem to *believe* for a moment that gas molecules behave or are really like that. She knows very well that this model is strictly flawed. However, she *accepts* it because this consciously simplified model is incredibly useful within a certain range of temperature and pressure. It makes calculations manageable, isolates the central relationship between pressure, volume, and temperature, and provides a powerful

[25] Williams's (1973) argument that *belief aims at truth* is foundational, but this claim is itself a complex philosophical debate. The *aim* has been interpreted constitutively (that aiming at truth is what makes a state a belief, as Williams suggests), normatively (that a belief is correct if and only if it is true, e.g., Velleman 2000), or instrumentally. For our purposes, the crucial feature is the phenomenological and involuntary link between the state of believing *p* and taking *p* to be true, which distinguishes it from accepting *p*.

framework for making accurate predictions and promoting scientific understanding (Elgin 2017). The falsehood is not an error in his reasoning process, but an intentional feature of his modelling strategy.

Therefore, the psychological and epistemological scenarios are very different. The KFF debate asks: 'Can an epistemic agent who stumbles into a false belief still manage to acquire knowledge from it?' The subject's false belief is an *unforeseen obstacle* that he has overcome or thought through carefully. In contrast, the use of idealizations raises another question: 'How can an epistemic agent who consciously and strategically uses a falsehood use it as a tool to improve their understanding of the world?' Here, the falsehood is a *consciously chosen tool.*

To claim that KFF theories justify the use of idealizations is like arguing that it is rational to intentionally consume poison for nutritional purposes because a person who accidentally ingests a small, non-lethal dose of poison survives. In the first case, it is knowledge despite a cognitive error; in the second, it is the cognitive benefit of a *conscious strategy.* Consequently, evidence that knowledge can (probably) be derived from an unconsciously held false belief does not entitle us to draw direct conclusions about the rationality or epistemic legitimacy of consciously and voluntarily adopting a falsehood as a premise in scientific research. These are two different philosophical problems, requiring different analytical tools.

2.3 The role of harmless falsehoods

The complex debate between advocates of KFF and KDF is more than an esoteric thought experiment; it forces us to refine our understanding of how truth and falsehood work in our cognitive lives. Rather than trying to determine a definitive winner in this particular dispute, we can draw a more fundamental and universal lesson: any plausible cognitive theory must be sophisticated enough to recognize that not all falsehoods are equal. The literature addressing these puzzles has largely agreed on a crucial distinction between *malicious* and *harmless* falsehoods, a distinction that is essential for understanding the role of idealizations in science (Fitelson 2017).[26]

A *malicious falsehood* is the classic villain of epistemology. It is a statement that actively sabotages the path from evidence to conclusion, making any resulting true belief a matter of pure and simple luck. If we look at the heart of Gettier-type cases, these are the kinds of falsehoods that we are going to find. For example, consider a typical picture: one looks at a field and sees an animal that looks just like a sheep. Based on her perception, she comes to the conclusion: 'There is a sheep in

[26] See Klein (2008) and De Almeida (2017).

this field'. However, her premise is based on seeing a cleverly camouflaged wolf. Unbeknownst to her, *there is* a real sheep in the field, hidden behind a hill. Her conclusion ('There is a sheep in this field') is true, but her reasoning is wrong. The wrong premise ('The animal I see is a sheep') is *malicious* because it completely separates the justificatory connection between her evidence and the fact that makes her belief true. Her belief is true due to chance, a *felicitous* accident. A *malicious* falsehood functions just like a defective compass that, by a one-in-a-million chance, points north at the exact moment she observes it. She has arrived at the right place, but her guide was useless.

A *harmless falsehood*, on the other hand, works very differently. Although it is literally false, it serves as a reliable and sometimes even necessary guide to a true conclusion. This is because the aspects in which the statement is false are irrelevant to the conclusion in question. Its falsehood is, in a sense, epistemically inert noise that does not interfere with the signal of the conclusion. Consider again the case of the *Inaccurate Clock*: its premise is 'It is exactly 2:58 p.m.', while the actual time is 2:56 p.m. This premise is false. However, it reliably leads you to the correct conclusion: 'I will not be late for my 7 p.m. appointment.' The falsity is harmless because the conclusion does not depend on *the accuracy* of the time. The part of your belief that really matters (that 'it is approximately 3 p.m. and you still have time') is true enough. The two-minute error is a deviation from the exact truth, but it is a deviation against which the conclusion is entirely robust. The falsehood is a harmless rounding error, not a fundamental error in reasoning.

This is precisely where the true strength of this concept lies, which has evolved from simple thought experiments to the core of our most successful cognitive endeavors. Catherine Elgin's *felicitous falsehoods* in science are paradigmatic examples of *harmless* falsehoods that are used systematically and productively. Think, for example, of the calculations done by physicists of centuries past to model the Earth's orbit around the Sun. They represented the two bodies as perfect spheres with all of their mass located at a single point, ignoring the fact that we had and have a terrestrial planet with a rugged surface because of mountains, oceans, and an atmosphere, as well as an uneven density. These assumptions were known to be falsehoods. Yet they were not only arguably good enough approximations that they could be used to make extremely accurate predictions of the trajectory of said planet, but they were very convenient in terms of the mathematical and computational effort that would have been required if those *technical* details had to be modeled rigorously. The Earth's deviation from a perfect sphere was irrelevant to the large-scale mechanics of its orbit. By omitting these chaotic and causally insignificant details, the model allowed for a manageable

calculation and a clearer understanding of the central gravitational relationship (Strevens 2008).[27] Falsehoods are not only harmless, but they are also useful, as they allow cognitive access to a complex reality.

This poses a profound challenge to any simplistic form of *veritism*: the view that truth is the only definitive epistemic good. If our goal is merely to accumulate true beliefs, then these scientific models based on falsehoods must be rejected. However, this would mean discarding precisely the tools that have enabled scientific progress over centuries. The motivation to question a rigid principle such as *no false premises* does not come from the controversial KFF thesis, but from the indisputable epistemic utility of harmless falsehoods. The task of the *veritistic*, therefore, is not to abandon the confession of truth, but to formulate it more carefully. The challenge is to explain how our cognitive economy can so fruitfully integrate the role of harmless falsehoods. The solution seems to lie in recognizing that we can adopt different attitudes towards the statements we use. The way we deal with a known and strategically useful falsehoods is fundamentally different from the way we deal with a statement we consider to be the literal truth, and so it should be.

3. A defense of veritism: belief, acceptance, and epistemic rationality

Our central thesis is that, contrary to Elgin, the rejection of statement (1) is entirely compatible with veritism. The conclusion that the usefulness of falsehoods in science leads to the failure of veritism is precluded precisely by the distinction that Elgin highlights between the attitudes of *belief* and *acceptance*.

3.1 The right place for veritistic norms

The apparent conflict between our commitment to truth and the successful use of falsehoods in science generates significant philosophical tension. However, the solution does not require us to abandon the *veritistic* principles that underpin our cognitive life. Instead, it requires a more precise understanding of where exactly these principles apply. By relying on the crucial distinction between the mental states of *belief* and *acceptance*, we can correctly map the epistemic terrain and show that the norms of truth govern their actual domain with full authority (Cohen 1992; Elgin 2017). This is not a clever way of sidestepping the problem, but a necessary clarification of the different roles these attitudes play in our cognitive economy.

[27] Strevens (2008) calls this the *kairetic account of explanation*, where idealizations function by isolating the *difference-makers* for an event while omitting causally irrelevant details. This view is central to a broader literature on how models, by virtue of being false, perform their explanatory function. See also Frigg and Nguyen (2019) on scientific representation and Rice (2021) on how idealizations can provide understanding of real-world phenomena.

Let us consider *belief* first. To believe a statement is to regard it as true, to take it as an accurate representation of the world. This attitude is fundamentally characterized by its *involuntariness*. As Bernard Williams (1973) has accurately argued, *one cannot* simply *decide* to believe something, just as one can decide to raise one's arm. Belief is triggered by encountering what one considers to be evidence. If someone offered you a million pounds to truly believe that the sky is green, you could not simply decide to believe it. You could say the words and play the part, but you would lack the inner conviction, as it cannot be directly controlled by the will. This is because the *constitutive goal* of belief is truth. A belief that does not aim at truth is not belief, but another mental state that masquerades as such. As its identity is linked to the accurate representation of reality, belief is also generally *context-independent*. If one takes that whales are mammals, one considers this to be true irrespective of whether one is in a biology class, participating in a whale watching tour, or reading *Moby Dick.* The context of your activities does not alter the truth value you assign to the proposition.

Acceptance, on the other hand, is a completely different type of cognitive attitude. Accepting a proposition means considering it true for a particular purpose and in a particular context. It is, first and foremost, a *voluntary* act; a decision or commitment made to serve a particular purpose (Bratman 1992).[28] In this sense, the key purpose of acceptance lies not in the truth, but in *usefulness*. There seems to be uncontroversial that we accept a premise because it is useful for calculations, to simplify a problem, to win an argument, or to investigate a hypothetical scenario. Consequently, acceptance is radically *context-dependent.* A physicist accepts the idealization of a frictionless plane to calculate the basic trajectory of an object, but immediately rejects this acceptance when the goal changes to the construction of a braking system, in which friction is the most important factor. Their acceptance is a temporary tool, adopted for one task and set aside for another. Above all, acceptance *does not require sincerity.* It is possible to accept a proposition coherently and, at the same time, believe in its negation. The scientist accepts that gas molecules are dimensionless points for the purposes of her model, while believing that they are, in fact, physical objects with volume and intermolecular forces. This psychological state, which would be a sign of profound irrationality for belief, is perfectly consistent with acceptance and often very productive.

[28] The distinction is famously articulated by L. Jonathan Cohen (1992), who argues that belief is a disposition (a feeling or state), whereas acceptance is a policy or action (a decision to treat p as true in a line of reasoning). Bratman (1992) develops this by arguing acceptance is voluntary and context-dependent, noting one might *accept* a premise for the sake of a practical project even while not believing it. Elgin (2017) deploys this distinction as a central plank of her argument.

With this distinction, we can take the decisive step: we can recognize that the central *veristic norm* (that is, the requirement to avoid falsehood and defend only true attitudes) applies directly and exclusively to the attitude of belief. This is not an *ad hoc* determination that serves to save a theory. There is a fundamental philosophical reason for this restriction. Belief is the only doxastic attitude whose *constitutive goal* is truth. To be in a state of belief simply *means* to take something as fact. Therefore, belief is the real and only bearer of fundamental *veristic value*. Subjecting belief to a primary norm other than truth (such as utility or convenience) would mean fundamentally misunderstanding the purpose of the attitude of belief. It would be like evaluating a news report according to the standards of epic poetry; the evaluation would miss the essential purpose of the activity.[29]

The norms governing acceptance, on the other hand, are instrumental and strategic. The question we ask ourselves when accepting something is not 'Is this true?', but 'Is this *useful* for my current purpose?'. However, that is not the same as understanding that acceptance just occurs in a vacuum without truth. Considering the case of epistemic contexts, the usefulness we attribute to the acceptance of a statement is normally measured by how well it serves the long-term ultimate goal of obtaining meaningful and relevant true beliefs or a deeper understanding of a subject. The scientist accepts the *falsehood* of a perfect vacuum in order to gain a *true* understanding of the law of universal gravitation. Falsehood is a calculated and strategic deviation, a structure used to construct a building of true knowledge and then dismantled.

Thus, *veritism* is not abandoned, but rather correctly classified. The direct and prohibitive norm 'Do not take a wrong stance' controls belief. The indirect and instrumental norm 'Accept what serves your ultimate epistemic goals' controls acceptance. By identifying belief as the right place for our fundamental confession of truth, we do not eliminate the problem of *felicitous falsehoods*. We recognize that our cognitive toolkit contains more than just a hammer for hammering in truths. It also contains a range of special tools for examining, modelling, and simplifying; tools that we use strategically and voluntarily on our way to a truer and richer understanding of the world.

[29] This does not imply that belief is not subject to other norms (e.g., of evidence, justification, or rationality). However, under a veritistic framework, these other norms are themselves instrumental in service of the constitutive norm of truth. The norm 'believe only what is supported by evidence' is a good norm because following it is the most reliable way to satisfy the primary norm 'believe only what is true'.

3.2 Acceptance as instrumental epistemic rationality

By correctly identifying belief as the sole goal of direct *veritistic norms*, we can recognize that the practice of accepting falsehoods in science is not the epistemic scandal it may seem at first glance. From a *veritistic* point of view, *believing in* a known falsehood is clearly an epistemic error. This would be a failure in the most basic cognitive duty: to reconcile one's conception of the world with the world itself. However, *accepting* a known falsehood is not necessarily an error. When a scientist accepts an idealization such as the ideal gas law, she is not committing an error, for the simple reason that she is not being deceived. She does not believe that the model is true; on the contrary, her professional competence consists, in part, in knowing exactly where and how it differs from reality. Since that she does not really believe that the given falsehood is true, she cannot be judged to have committed the most serious *veritistic sin*. She is not like someone who has been deceived by a very convincing illusion, she is more like a spy who consciously uses a false passport – not because she is actually confused about her own identity, but because she understands that the false document is a necessary tool for accomplishing an important mission.

Furthermore, if the accepted falsehood is *harmless* (a falsehood whose inaccuracies are irrelevant to the conclusion at hand) a veritistic has every reason to endorse this action. This step can redefine the acceptance of *harmless* falsehoods – not as a reluctant concession or a violation of our principles – but as a sophisticated and admirable expression of our commitment to the truth. We can argue that accepting a known and harmless falsehood is an act of what James Montmarquet (1993) calls *epistemic conscientiousness*.[30] This virtue is not about rigid and dogmatic adherence to a rule, such as 'never tell a lie'. It is about having a genuine and effective desire to achieve truth and avoid error. A truly conscientious epistemic agent is pragmatic and resourceful. They understand that the path to knowledge is often winding and that the most effective route to deep truth sometimes involves a strategic and temporary detour through a simplifying untruth. The scientist who uses an idealization is not intellectually lazy, but intellectually intelligent, in choosing the best tool for the cognitive task. Their actions, therefore, *serve* truthful objectives and do not contradict them.

This connection between utility and the pursuit of truth can be strengthened by considering the acceptance of a harmless falsehood as a paradigmatic act of

[30] See Montmarquet (1993: viii). Montmarquet's virtue epistemology is a *motivationalist* or *internalist*, where epistemic virtue lies in the desire to attain truth and avoid error. This contrasts with *reliabilist* virtue epistemologies (e.g., Sosa, Greco) which define virtues as reliable cognitive faculties. We find Montmarquet's concept more apt here, as the scientist's conscious desire for truth is what motivates the instrumentally rational policy of *accepting* a falsehood.

instrumental epistemic rationality. In the most general sense, instrumental rationality consists of using the most effective means to achieve self-selected goals. If your goal is to get to the city center, it is instrumentally rational to take the express bus rather than walk. In the epistemic domain, the ultimate goal of a veritistic is to acquire a rich and meaningful set of true beliefs and to avoid errors. The *means* to this end are the various cognitive strategies, methods, and procedures we use.

The decisive conclusion here is that the rationality of accepting an idealization should not be evaluated as an isolated action, but as part of *a* broader *strategy*. The strategy goes something like this: 'In context C, I accept model M to achieve goal G'. The rational test that we can apply to this strategy is not to consider whether model M is true, but whether the strategy itself can provid a proper contribution to truth in the long run. Is it likely that an agent who applies this strategy will end up in a better epistemic position (with more meaningful true beliefs) than an agent who refuses to apply it? The answer is clearly 'yes'. Consider an investor whose goal is to maximize their wealth. A single asset in their portfolio may lose value (a *falsehood*). But the *policy* of diversification, which involves maintaining a mix of assets, is instrumentally rational because it is the most reliable long-term strategy for achieving the goal. Accepting a useful falsehood is like holding a strategic asset that, while not performing exceptionally well on its own, grows the entire knowledge portfolio.

The veritistic, therefore, is not trapped and forced to defend the veracity of *the* accepted *model*. That is a losing battle that he does not need to fight. Instead, he only needs to defend the veracity of *the policy* of accepting the model in a given context. And there is ample evidence of this. Accepting the *falsehood* of Newton's laws of motion allows engineers to make a series of *sufficiently true* predictions to land a probe on Mars. Acceptance of the 'falsehood' of a frictionless plane allows a student to gain a real and *true* understanding of the fundamental relationship between force, mass and acceleration and to extract a basic principle from the noise of chaotic reality (Strevens 2008).[31] These idealizations function as an indispensable cognitive framework that allows us to build a stock of knowledge and understanding that would otherwise be inaccessible.

[31] This defense of instrumentalism is compatible with a form of scientific realism. The policy of accepting Newtonian mechanics is rational because it reliably generates truths, and it does so because (we have reason to believe) it is a *veridical idealization* (Sorensen 2013) or a successful approximation of a deeper, true theory (e.g., General Relativity). The veritist is not committed to the truth of the accepted model, but is committed to a truth-based explanation for why the policy of accepting the model is so successful.

This provides a fundamental reason for a veritistic to be intensely concerned with *norms of acceptance* (they must be guidelines that promote truth) without mistakenly applying the direct norm *believe only truths* to the attitude of acceptance itself. *Veritism* is therefore entirely compatible with the rejection of the statement that there can never be epistemically good reasons to accept a known falsehood. It is not weakened by this conclusion but reveals itself to be a more mature and realistic philosophical position, capable of explaining the complex and effective strategies that constitute our most successful cognitive endeavors.

4. Veritism and the value of understanding

The arguments in the previous sections should show that a differentiated realism, which distinguishes between belief and acceptance, can successfully account for the widespread and productive use of untruths in science. However, this has largely been a defensive measure to protect the fundamental principles of a truth-centered epistemology against a strong objection. A truly robust philosophical position, however, needs to do more than just refute objections; it needs to offer a compelling and positive account of our most valuable cognitive achievements. The deeper and more fundamental challenge to *veritism*, presented by revisionist epistemologists – such as Catherine Z. Elgin (2017) and Jon Kvanvig (2003) – is that its basic orientation is misguided. The ultimate gain of research, they argue, is not a scattered collection of individual true beliefs, but a holistic and integrated cognitive state: *understanding*. In their view, understanding is often achieved *through* the use of simplified, non-factual models, suggesting that its value is different from, and perhaps even contradictory to, the pursuit of truth.[32] A successful defense of *veritism* must therefore do more than just make room for useful falsehoods; it must show that it can fully account for the deep epistemic value of understanding itself.

We argue that the framework developed here, which focuses on the *instrumental rationality of acceptance* in the service of veritistic goals, provides precisely the tools necessary for a powerful *veritistic theory of understanding*. This is a very active and fruitful area of current epistemological research, in which many theorists investigate how the value of understanding can be explained without postulating it as *sui generis* epistemic good, separate from truth (see Grimm 2012;

[32] Kvanvig (2003) is perhaps the most forceful proponent of this view, arguing that understanding, not knowledge, is the primary epistemic good, in part because its value is not susceptible to Gettier-style luck. Elgin (2017) argues that understanding, achieved via felicitous falsehoods, is a distinct epistemic achievement often preferable to a collection of *mottled, fragmented, isolated truths*. See also Zagzebski (2001) and de Regt (2017) for defenses of understanding as a *sui generis* epistemic value.

Pritchard 2010; Hills 2016). The immense value we attach to understanding is by no means an alternative to truth but can be explained completely and elegantly in veritistic terms.

First, we need to clarify what *understanding* is. It is proven to be more than just possessing true beliefs; however numerous they may be. Consider a person who knows vast amounts of car facts: the cylinder firing order, optimal tire pressure, specific oil viscosity, and exact torque for every screw. This person has a lot of true beliefs about cars, but if she doesn't understand how they are all related, she doesn't know how an internal combustion engine works. Now consider an experienced mechanic. she may not have memorized all these individual facts, but they understand the *network of explanatory and dependency relationships* that structure the system. In this case, she understands the entire flow of energy transfer from cylinder combustion to the pistons, to the crankshaft, and out to the drive wheels. She can articulate that a plugged fuel injector will cause a misfire and can correctly evaluate the consequences of a slipped timing belt. Her cognitive state is not valuable because of the number of isolated facts it contains, but because of her understanding of the causal and logical architecture of the system.

This *understanding* is a dynamic cognitive ability.[33] It is the ability to recognize connections, provide explanations, make counterfactual considerations ('What if...?'), and predict the behavior of the system under new conditions. It is this characteristic that Elgin attempts to capture in her description. For Elgin, the resulting understanding is not factual, since the scientific models that convey this understanding are themselves idealizations: *felicitous falsehoods*. For example, the behavior of gases is understood with the help of the ideal gas law, a known falsehood. The veritist, however, can tell a different and more nuanced story.

The instrumentally rational acceptance of idealizations is a crucial, perhaps even indispensable, heuristic tool for achieving this understanding of relationships. Reality in all its uniqueness is an extremely complex and noisy system. A totally accurate and faithful representation of a phenomenon would be as confusing as the phenomenon itself, just as a map the same size as the area it represents would be totally useless. A subway map is a powerful tool for understanding a public transport system, precisely *because* it is a felicitous falsehood. It distorts geography,

[33] This *grasping* or *dependency-seeing* model of understanding is common among veritistic theorists. Hills (2016), for example, argues that understanding why *p* consists in the ability to grasp explanatory connections and answer *what-if-things-had-been-different* questions. Similarly, Grimm (2012) argues that understanding is the *grasp of a thing's structure*, which he links directly to truth-conduciveness. This contrasts with Kvanvig's (2003) more holistic view of understanding as a maximal set of coherent beliefs.

straightening lines and balancing distances to perfectly highlight the underlying topological network of stations and connections, the structure that matters to the traveler.

This is exactly how scientific models work. The Lotka-Volterra equations, which model the dynamics between predators and prey, are based on a series of obviously false assumptions: that predators have an insatiable appetite, that prey have unlimited food reserves, that the environment is completely stable, and so on. No ecologist *believes* these assumptions. Although, in strategically move, by *accepting them*, she can disregard the (random and confusing) details of a real ecosystem and thus *capture* the fundamental feedback loop that drives population cycles.[34] Such a model illustrates the fundamental relationship of dependence, in which the more prey there is, the more predators can thrive; the more predators thrive, the fewer prey survive; the fewer prey there are, the more predators go hungry; the more predators go hungry, the more prey can recover. The model's *falsehood* is the price to pay for this clear view of the system's core dynamics. Acceptance of the model is the instrumentally rational path to the cognitive state of understanding.

Now, we come to the heart of the matter: why is this state of understanding so valuable from a veritistic perspective? The answer is that understanding, in its relation to truth, is profoundly and perhaps uniquely *instrumentally valuable*. Duncan Pritchard (2010, 74) sums this up perfectly when he describes understanding as a "source of meaningful truths". An agent who understands a system is capable of generating an extensive and powerful network of true beliefs about it. Their understanding of the structure of the system is a truth-generating engine.

Let us return to the mechanic. Her understanding of the engine allows her to produce an unlimited stream of meaningful truths when needed. She can generate new true beliefs ('If the engine overheats and the coolant level is full, then the thermostat is probably stuck'). She can generate true counterfactual beliefs ('If the distance between the spark plugs were greater, the engine would run unevenly'). She can generate true explanatory beliefs ('The reason the car won't start is that the faulty alternator didn't charge the battery'). This generating power is the hallmark of understanding. The value of your cognitive state is not in the static proposition it

[34] This is a classic example of a *toy model* in science. Such models are intentionally simple and highly idealized, not to approximate a specific target system, but to provide "understanding of a more general phenomenon" (Rice 2021, 4104). Philosophers of science like Weisberg (2013) and Cartwright (1983) have argued that this kind of extreme idealization is essential for isolating causal mechanisms that would be obscured in a more complex, *realistic* model.

currently holds, but in its dynamic capacity to produce an abundance of new and useful propositions. The tree is not valuable because of its current leaves, but because of its enduring capacity to produce new leaves every spring.

Moreover, the truths generated by understanding are not trivial. They are *meaningful* truths about the causal, modal, and explanatory structure of the world. A defensible realism must surely deal with the quality and significance of the truths we obtain, not just their quantity. It is better to have a small number of true beliefs about the fundamental laws of physics than a large number of true beliefs about the number of grains of sand on a beach. *Understanding* is the primary capacity by which we have access to these deeper and more significant truths.[35]

This step transforms our project from a narrow, defensive exercise into a constructive one. It shows that *veritism*, when properly understood, not only has the resources to deal with the complex practices of science, but also to provide a convincing, truth-based representation of our greatest epistemic achievements. By distinguishing between the norms of belief and the rational instrumental guidelines of acceptance, *veritism* can explain the entire epistemic arc. It shows why scientists rationally *accept* known falsehoods in their quest for *understanding* and explains why this understanding is so valuable in itself: because it is the most powerful engine we have for the systematic production of meaningful *truths*. The challenge from Elgin and others is therefore not a refutation, but an invitation; an invitation for *veritism* to go beyond simply *counting* true propositions and become a robust theory that celebrates understanding as the crown jewel of our efforts in pursuit of truth.

5. Objections and Replies

No philosophical thesis is complete before it has withstood serious objections. The veritistic framework defended here (which attempts to consider the epistemic role of falsehoods, distinguishing norms of belief from rational instrumental guidelines for acceptance) is no exception. While this approach allows *veritism* to explain scientific practice and the value of understanding, it also invites critical examination.

In this section, we will examine two of the most pressing objections: first, that our representation of acceptance is too *permissive* and risks allowing irrationality; and second, that our explanation of understanding *devalues* it, treating it as a mere

[35] This successfully counters the *trivial truths* objection. A mature veritism is not committed to the equal value of all truths. The challenge, of course, is to provide a non-circular account of what makes a truth *meaningful* or *significant*. A plausible veritistic answer is that "meaningful" truths are those that, once possessed, are instrumentally most valuable in generating other true beliefs (e.g., truths about laws, causes, and dependencies). See Pritchard (2010) for a discussion of this qualitative, as opposed to purely quantitative, dimension of epistemic value.

instrument for the acquisition of true beliefs. In addressing these challenges, we can demonstrate the resilience and explanatory power of a sophisticated veritistic epistemological theory.

5.1 The charge of permissiveness

The first and perhaps most troubling objection is that our criterion for *rational acceptance* is dangerously *permissive*. As we have argued, it can be considered *rational* to *accept* a known falsehood, provided that this approach is beneficial to the truth in the long run. A detractor might argue that this view is too vague, too distant, and ultimately too permissive. If the only check on acceptance is an anticipated gain in long-term truths, what prevents a subject from justifying the acceptance of any belief system that suits them? Couldn't a proponent of a conspiracy theory or an astrological view of the world argue that accepting these falsehoods as working hypotheses helps them *recognize patterns* and ultimately leads them to a deeper truth? In short, the objection is that by replacing the immediate and rigid norm of truth with a diffuse and future-oriented norm of instrumental utility, we open the door to epistemic anarchy. We seem to allow any cognitive practice, however unfounded, as long as it comes with a promise of future truths.[36]

This objection is serious, but it is based on a misunderstanding of what instrumental epistemic rationality requires. A rational policy of acceptance is not a blank cheque, but a cognitive strategy whose legitimacy is subject to rigorous epistemic evaluation. The policy 'In context C, I accept falsehood F to achieve goal G' is not rational simply because an agent declares it to be truth-promoting. Instead, the agent must have *good reasons* or *strong evidence* to believe that the policy is indeed truth-promoting. This requirement for evidence provides the firm and immediate restriction that the objection rightly demands.

We can recognize this constraint by using the example of a scientific idealization compared to a pseudoscientific premise. We have overwhelming evidence, collected over centuries, that the policy of accepting Newtonian mechanics as a model for medium-sized objects at low speeds is an incredibly powerful engine for creating true predictions and promoting genuine understanding. This policy has a proven track record of success. Furthermore, we

[36] This objection echoes the *pragmatist* critique of truth, famously articulated by William James ([1897] 1951) in *The Will to Believe*, where he argues that in certain "forced, living, and momentous" situations, it can be rational to *will* a belief in advance of the evidence. The veritistic framework here is explicitly anti-Jamesian in this regard: the policy of acceptance is not a *will to believe*, but a strategic move whose rationality is itself hostage to publicly available evidence of its truth-conduciveness.

have a deep theoretical understanding of *why* it works: we understand it as a limiting case of a more general theory, the *Theory of General Relativity*. On the other hand, consider a policy that accepts the premise: 'The orientation of Saturn influences my financial prospects'. There seems to be no reliable evidence whatsoever to the effect that this policy promotes truth, and in fact we have plenty of evidence that it does not.[37] A diligent truth-seeker, who is guided by the higher desire to achieve truth and avoid error, would consider the first policy instrumentally rational and the second instrumentally irrational.

The criterion, therefore, is not just *promotion of long-term truth* but *promotion of evidence-based and verifiable truth*. The policy of acceptance must be part of what Alvin Goldman (1986) would call a *reliable* process**.** By insisting that the instrumental rationality of the policy itself is an evidence-based matter, we avoid a slide into permissiveness and ensure that the acceptance of any falsehood is accountable to the same epistemic court of appeal as any other cognitive practice.

5.2 The devaluation of understanding

A second, more subtle objection is directed against our veritistic representation of understanding. We argue that the deep value of understanding is instrumental: it is a *source of truths*, a cognitive state whose value derives from its powerful capacity to generate a huge network of meaningful true beliefs. An inspired critic, Elgin (2017) or Kvanvig (2003), would argue that this has exactly the opposite effect. *Understanding*, they would say, is a *definitive epistemic good*, not an instrumental one. We try to understand the world for its own sake, not just as a means of accumulating a greater number of true facts. According to this view, veritistic representation devalues our greatest cognitive achievement and reduces the holistic and structural understanding of a phenomenon to a mere factory for producing propositional results. It confuses the real goal of research with a tool used to achieve that goal.[38]

---

[37] This reliance on an evidence-based track record for a policy allows the veritist to engage with the *demarcation problem* between science and pseudoscience. While Popper (1963) famously argued for falsifiability as the criterion, a more modern consensus, following Laudan (1983), suggests that *scientific* practices are those that constitute a *progressive research tradition* – i.e., one with a proven, evidence-based track record of solving problems and generating truths. Astrological *policies of acceptance* fail this test. See Pigliucci and Boudry (2013).

[38] This distinction between final (or intrinsic) and instrumental value is at the heart of axiology. While Kvanvig (2003) and Elgin (2017) position understanding as a final good and truth as (merely) instrumental to it, our account reverses this. We hold that meaningful truth is the final epistemic good, and understanding is the cognitive state that realizes or embodies the possession of such truths. Thus, understanding may be considered constitutively valuable, not just instrumentally, as

This objection impressively expresses the phenomenological understanding that *understanding* is an end in itself. However, it gains strength by implicitly presenting *veritism* as a crude form of *epistemic counting*, in which the goal is only to maximize the number of individual true propositions in the mind. However, mature *veritism* does not advocate such a simplistic view. The veritistic can and should insist that not all truths are equal. The true belief that 'the Battle of Hastings took place in 1066' has a certain value, but it is overshadowed by the value of a true belief about the law of universal gravitation. A differentiated *veritism* deals with the acquisition of *meaningful truths*, especially truths that concern the deep underlying structures of reality: causal relations, explanatory networks, and modal dependencies.

Once we take this step, the objection loses its force. *Understanding* is not just an instrument for generating random truths, but the cognitive state in which *these meaningful truths are grasped in their interconnected structure*. *Understanding* a phenomenon means possessing the most valuable and meaningful truths about it. Understanding is therefore not merely a *means* to the veritistic's end, but the *realization and embodiment* of that end in its highest form.

Let's look at an analogy. An investor's goal is to accumulate wealth. One person could do this by accumulating a large number of individual coins. Another could do this by acquiring a portfolio of interconnected, high-growth assets that form a functional economic engine. Both have *wealth*, but the second person's wealth is more robust, powerful and meaningful. The simple veritistic is like the coin collector. The sophisticated veritistic is like the portfolio owner. *Understanding* is not the factory that produces the coins, but the ownership of the entire high-value portfolio. Its value is not purely instrumental, for by acquiring understanding, one has already achieved the most important part of the veritistic goal: the acquisition of a powerful, coherent, and meaningful stock of truths. By formulating the truth-seeking project in this way, we can consider the intuition that understanding is the ultimate price of research, without abandoning the principle that our cognitive life is and must be fundamentally oriented towards truth.

In dealing with these objections, our veritistic framework is strengthened. It is not overly permissive, as the guidelines it advocates must be based on proof of their reliability. And it does not devalue understanding but celebrates it as the ultimate expression of the search for meaningful truth.

---

it is the form that the highest veritistic good takes.

6. Conclusion

We began with a fundamental tension at the heart of our cognitive life: the deep commitment to truth that anchors our philosophical conception of knowledge, in contrast to the productive and ubiquitous reliance on known falsehoods in science, our most successful endeavor in the pursuit of truth. The challenge was to reconcile the two without abandoning our veristic principles or the rational authority of scientific practice. The framework developed here has shown that this apparent conflict is illusory. Sophisticated *veritism*, based on a fundamental distinction between our cognitive attitudes, is not only resistant to this challenge, but also proves to be a more effective and explanatory theory of knowledge.

The central thesis of our argument was that truth norms are not an inflexible instrument that can be applied uniformly to all our cognitive states. By carefully distinguishing between the involuntary, truth-oriented attitude of *belief* and the voluntary, pragmatic policy of *acceptance*, we find the right place for our core veristic norms. The requirement to defend only truths applies fully to faith, whose goal is to represent the world accurately. Acceptance, however, is subject to another criterion: that of instrumental rationality. We argue that accepting a known and harmless falsehood (such as, for example, a scientific idealization) is not an epistemic failure, but can be a masterpiece of epistemic conscientiousness. It is a rational policy if, and only if, it is part of a reliable, evidence-based strategy for achieving the ultimate veristic goal: the acquisition of a rich and meaningful set of true beliefs.

After establishing this perimeter of defense, we move on to a more constructive project and show how this framework provides a convincing veritistic explanation of *understanding*. Understanding is by no means a competitive *epistemic good*, whose value is independent of truth, but exactly what a mature veritistic seeks. We define *understanding* as the cognitive grasp of the network of causal and explanatory relations that structure a phenomenon. The acceptance of *felicitous falsehoods* is an important, perhaps even indispensable, heuristic for achieving this understanding, as simplified models omit confusing details to highlight the central structure of a system. The immense value of this achievement was, in turn, explained in veritistic terms. Understanding is a source of meaningful truths; a generator that allows an actor to create an extensive and powerful network of predictions, counterfactuals, and innovative explanations.

Projects like this one often present the false issue of having to choose between a truth-centered epistemology and a contextual one, more closely resembling the non-linear, model-based, and fragmented scientific practice we are used to. The epistemic distinction between belief and acceptance allows us to align the normative dimension of epistemology with the practical dimension of success in scientific

practice. It allows for a form of scientific realism that is compatible with the obvious anti-literalism of many of our best theories.[39] We can affirm that science seeks truth while recognizing that its path is often paved with strategically employed falsehoods. Truth, therefore, remains the ultimate goal of our cognitive journey. However, a mature theory of knowledge must recognize that the most direct path is not always the most effective and that the strategic use of a simplified and incorrect map can often be the safest guide. *Veritism* is not threatened by this reality but confirmed by its ability to explain it.[40]

References

Arnold, Alexander. 2013. "Some Evidence Is False." *Australasian Journal of Philosophy* 91 (1): 165–172.

Audi, Robert. 2011. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. 3rd ed. London: Routledge.

Balcerak Jackson, Magdalena, and Brendan Balcerak Jackson. 2013. "Reasoning as a Source of Justification." *Philosophical Studies* 164: 113–126.

Ball, Brian, and Michael Blome-Tillmann. 2014. "Counter-Closure and Knowledge despite Falsehood." *Philosophical Quarterly* 64 (257): 552–568.

Bernecker, Sven, and Thomas Grundmann. 2019. "Knowledge from Forgetting." *Philosophy and Phenomenological Research* 98: 525–540.

Borges, Rodrigo, Claudio de Almeida, and Peter D. Klein, eds. 2017. *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford: Oxford University Press.

Bratman, Michael E. 1992. "Practical Thought and Acceptance in Context." *Mind* 101 (401): 1–15.

Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.

[39] This framework thus offers a path for a *modest* scientific realism. It avoids the anti-realist conclusion that the use of idealizations proves our theories are false (e.g., Cartwright 1983). At the same time, it avoids a *naive* realism that would require our models to be literal, one-to-one representations of reality. By focusing on the policy of acceptance, we can be realists about the truth-conduciveness of our scientific methods without being literalists about the content of every model we employ.

[40] A key avenue route for future research remains the development of a more formal account of *epistemic conscientiousness* as it relates to acceptance. While Montmarquet (1993) provides the virtue-theoretic groundwork, a more detailed analysis is needed to specify the conditions under which an agent has sufficient evidence to deem an acceptance-policy *instrumentally rational*, especially in novel cases of inquiry where a long-term *track record* (as discussed in Section 5.1) is not yet available.

Clark, Michael. 1963. "Knowledge and Grounds: A Comment on Mr. Gettier's Paper." *Analysis* 24 (2): 46–48.

Coffman, E. J. 2008. "Warrant without Truth?" *Synthese* 162: 173–194.

Cohen, L. Jonathan. 1992. *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.

de Almeida, Claudio. 2017. "Knowledge, Benign Falsehoods, and the Gettier Problem." In *Explaining Knowledge: New Essays on the Gettier Problem*, edited by Rodrigo Borges, Claudio de Almeida, and Peter D. Klein, 292–311. Oxford: Oxford University Press.

de Regt, Henk W. 2017. *Understanding Scientific Understanding*. Oxford: Oxford University Press.

Elgin, Catherine Z. 2017. *True Enough*. Cambridge, MA: MIT Press.

Elgin, Catherine Z. 2019. "Epistemically Useful Falsehoods." In *Themes from Klein: Knowledge, Scepticism, and Justification*, edited by Branden Fitelson, Rodrigo Borges, and Cherie Braden, 25–39. Cham: Springer.

Elgin, Catherine Z. 2021. "Replies." *Synthese* 199: 1577–1597.

Fitelson, Branden. 2010. "Strengthening the Case for Knowledge from Falsehood." *Analysis* 70: 666–669.

Fitelson, Branden. 2017. "Closure, Counter-Closure, and Inferential Knowledge." In *Explaining Knowledge: New Essays on the Gettier Problem*, edited by Rodrigo Borges, Claudio de Almeida, and Peter D. Klein, 312–323. Oxford: Oxford University Press.

Frigg, Roman, and James Nguyen. 2019. "Scientific Representation." In *The Stanford Encyclopedia of Philosophy* (Winter 2019 edition), edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

Frigg, Roman, and James Nguyen. 2021. "Mirrors without Warnings." *Synthese* 198: 2427–2447.

Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123.

Goldberg, Sanford C. 2001. "Testimonially Based Knowledge from False Testimony." *Philosophical Quarterly* 51: 512–526.

Goldman, Alvin I. 1967. "A Causal Theory of Knowing." *Journal of Philosophy* 64: 357–372.

Goldman, Alvin I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Goldman, Alvin I. 2002. *Pathways to Knowledge: Private and Public*. Oxford: Oxford University Press.

Grimm, Stephen R. 2012. "The Value of Understanding." *Philosophy Compass* 7 (2): 103–117.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hills, Alison. 2016. "Understanding Why." In *Explaining Understanding: New Perspectives from the Theory of Knowledge and Philosophy of Science*, edited by Stephen Grimm, Christoph Baumberger, and Sabine Ammon, 154–173. New York: Routledge.

Hilpinen, Risto. 1988. "Knowledge and Conditionals." *Philosophical Perspectives* 2: 157–182.

James, William. (1897) 1951. *The Will to Believe and Other Essays on Popular Philosophy*. New York: Dover Publications.

Klein, Peter D. 2008. "Useful False Beliefs." In *Epistemology: New Essays*, edited by Quentin Smith, 25–61. Oxford: Oxford University Press.

Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Laudan, Larry. 1983. "The End of the Demarcation Problem." In *Physics, Philosophy, and Psychoanalysis*, edited by Robert S. Cohen and Larry Laudan, 111–127. Dordrecht: Springer.

Lawler, Insa. 2021. "Scientific Understanding and Felicitous Legitimate Falsehoods." *Synthese* 198 (7): 6859–6887.

Le Bihan, Soazig. 2021. "Partial Truth versus Felicitous Falsehoods." *Synthese* 198 (6): 5415–5436.

Lehrer, Keith. 1965. "Knowledge, Truth and Evidence." *Analysis* 25: 168–175.

Luzzi, Federico. 2010. "Counter-Closure." *Australasian Journal of Philosophy* 88: 673–683.

Luzzi, Federico. 2014. "What Does Knowledge-Yielding Deduction Require from Its Premises?" *Episteme* 11: 261–275.

Luzzi, Federico. 2019. *Knowledge from Non-Knowledge: Inference, Testimony and Memory*. Cambridge: Cambridge University Press.

Montmarquet, James A. 1993. *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman & Littlefield.

Murphy, Peter. 2013. "Another Blow to Knowledge from Knowledge." *Logos & Episteme* 4: 311–317.

Nawar, Tamer. 2021. "Veritism Refuted? Understanding, Idealization, and the Facts." *Synthese* 198 (5): 4295–4313.

Pigliucci, Massimo, and Maarten Boudry, eds. 2013. *Philosophy of Pseudoscience: Rethinking the Problem of Demarcation*. Chicago: University of Chicago Press.

Popper, Karl R. 1963. *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge.

Pritchard, Duncan. 2010. *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.

Pritchard, Duncan. 2021. "Veritism and the Goal of Research." *Episteme* 18 (4): 584–598.

Rice, Collin. 2021. "Understanding Realism." *Synthese* 198: 4097–4121.

Saunders, John Turk, and Narayan Champawat. 1964. "Mr. Clark's Definition of 'Knowledge.'" *Analysis* 25: 8–9.

Schnee, Ian. 2015. "There Is No Knowledge from Falsehood." *Episteme* 12 (1): 53–74.

Sorensen, Roy. 2013. "Veridical Idealizations." In *Thought Experiments in Science, Philosophy, and the Arts*, edited by Mélanie Frappier, Letitia Meynell, and James Robert Brown, 30–52. New York: Routledge.

Sorensen, Roy. 2016. "Fugu for Logicians." *Philosophy and Phenomenological Research* 92: 131–144.

Strevens, Michael. 2008. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.

Swain, Marshall. 1981. *Reasons and Knowledge*. Ithaca, NY: Cornell University Press.

Turri, John. 2011. "Manifest Failure: The Gettier Problem Solved." *Philosophers' Imprint* 11 (8): 1–11.

Turri, John. 2012. "In Gettier's Wake." In *Epistemology: The Key Thinkers*, edited by Stephen Hetherington, 214–229. London: Continuum.

Turri, John. 2019. "Knowledge from Falsehood: An Experimental Study." *Thought* 8: 167–178.

Turri, John, Peter Blouw, and Wesley Buckwalter. 2017. "Gettier Cases: A Taxonomy." In *Explaining Knowledge: New Essays on the Gettier Problem*, edited by Rodrigo Borges, Claudio de Almeida, and Peter D. Klein, 242–252. Oxford: Oxford University Press.

Velleman, J. David. 2000. "The Aim of Belief." In *The Possibility of Practical Reason*, 244–281. Oxford: Oxford University Press.

Warfield, Ted A. 2005. "Knowledge from Falsehood." *Philosophical Perspectives* 19: 405–426.

Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.

Williams, Bernard. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.

Zagzebski, Linda T. 2001. "Recovering Understanding." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 235–251. Oxford: Oxford University Press.