

Better understanding through falsehood (forthcoming in *Pacific Philosophical Quarterly*)

Benjamin T. Rancourt

December 29, 2014

Abstract

Can understanding be based on false beliefs? In this paper, I argue that it can. In order to defend this answer, I first argue that the best way to understand the question is that it is about whether one can increase one's degree of understanding by coming to have an overall less accurate body of beliefs. To demonstrate this, I identify three sufficient conditions for one body of beliefs to be more accurate than another. Next, I appeal to two methods of comparing degrees of understanding, methods that are used in everyday life, in educational settings, and in psychology. With these methods of measuring degrees of accuracy of beliefs and degrees of understanding, I show that understanding can be gained by acquiring false beliefs. In addition, I show that this conclusion settles the ongoing debate over whether understanding is factive. Understanding is not factive.

1 Introduction

Can false beliefs improve understanding? A common intuition is that they cannot, that false beliefs always hinder understanding. There is an expectation that understanding and the truth go hand in hand. However, in this paper I argue that sometimes false beliefs can increase understanding. In addition, one can lose understanding by gaining true beliefs. The relationship between understanding and the truth is not as simple as one might have thought; it is not simply a matter of false beliefs hindering understanding while true beliefs improve it.

The fact that false beliefs can improve understanding is inimical to reductive accounts of understanding that treat it as a species of knowledge. According to reductive accounts, to understand a topic is just to have the right kind of propositional knowledge about that topic. Since knowledge that p can only exist if p is true, the reductive view requires that understanding can only be based on the truth. The beliefs must constitute knowledge. This paper will provide reason to reject these reductive views.

In addition, this paper demonstrates that by answering the question of whether false beliefs can improve understanding, we also answer the question whether understanding is factive, a question debated in (Kvanvig 2003), (Elgin 2007), (Elgin 2009), (Pritchard 2008), (Riggs 2009), and elsewhere.¹ This debate has largely focused on understanding as an all-or-nothing state, rather than as coming in degrees. I show that if one's degree of understanding can be increased through false beliefs or hindered through true beliefs, then understanding is not factive in the sense used in these debates.

Truth is often taken to be the only, or highest, epistemic good. On the other hand, understanding is often taken to be the goal of inquiry. The argument in this paper undermines the connection between understanding and the truth, thus creating a tension between these two claims. This provides occasion to reassess the roles of truth and understanding in inquiry.

What role does truth play in understanding, then? I argue that understanding means having the ability to find the truth. It does not always require already having true beliefs. In particular, some false beliefs can aid this ability, and some true beliefs can hinder it, but the ability to find the truth is what matters. Thus, understanding retains an important connection to the truth, just not the connection many have suspected.

The argument presented below does not depend on a detailed account of understanding, but instead employs two methods by which we measure and compare degrees of understanding in real life. The first method, used in classrooms, job interviews and psychology experiments, involves finding answers to questions, especially novel questions. If one has a greater ability to answer novel questions about a topic, then one has greater understanding of that topic. The second method, also employed in evaluating understanding, is reasoning ability. If one can reason better about a topic, then one has a better understanding of the topic

¹Note that this is an extension of the original use of the term 'factive'. The debate is not whether that which you understand must be true, but whether understanding of something must be based on true beliefs. See section 5.

(for more detail, see section 2.2).

Using these methods of comparison, I identify two broad classes of cases where false beliefs can improve understanding and true beliefs can hinder understanding. The first class of cases involve approximation and idealization. Approximations and idealizations are false, but believing them to be true can sometimes increase understanding. The second class of cases involves misleading true beliefs. Someone with a misleading true belief can have a lower degree of understanding about the topic than someone with fewer true beliefs, or even false beliefs. Understanding can be increased through the acquisition of false beliefs.

The rest of the paper is structured as follows: in section 2, I clarify the nature of the question about the relationship between understanding and truth. Next, section 3 provides two classes of counterexample to the factivity of understanding. In light of these counterexamples, Section 4 considers what role truth plays in understanding. Section 5 shows how the argument here implies that understanding is not factive, and explains how the present argument avoids objections to past arguments against factivity. Section 6 explains why there is no reason to impose a truth requirement on understanding.

2 Understanding and the truth

Consider these two questions one can ask about the relationship between understanding and truth.

Object Question Must that which one understands be true?

Basing Question Must the beliefs upon which understanding is based be true?²

An example illustrates the difference between the two: Suppose John understands why boats float. The object question is whether John's understanding why boats float implies that boats float.³ On the other hand, the basing question is whether John's understanding why boats float implies the truth of the beliefs that allow

²This is closely related to the question whether understanding is factive in the sense employed by (Kvanvig 2003), and those adopting his use of 'factive', such as (Elgin 2007) and (Riggs 2009). See section 5 for details.

³Technically, the object of understanding in the example is the question why boats float. However, following convention, for the object question we decompose 'understanding why boats float' into the verb *understanding-why* taking the proposition that boats float as its object. For the basing question we need not make this assumption.

him to understand why boats float, such as his beliefs about boats, water, floating, etc.

We can dispense with the object question briefly. On one hand, there is good reason to think that understanding why p implies that p is true. One cannot understand why rocks float in air, because rocks do not float in air. In the same vein, it also appears that understanding that p implies that p is true.⁴ For these kinds of understanding, one can only understand what is true.

However, not all understanding requires the object of understanding to be true. One can understand a false theory. For example, I understand several JFK assassination conspiracy theories that are false in almost every detail. Further, the object question does not make sense when applied to understanding a person, or a painting, or a symphony, since these things are not truth apt. Relatedly, one can understand a subject matter, yet if subject matters are sets of questions as argued in (Lewis 1988), it does not make sense to ask whether a subject matter is true or false. For these kinds of understanding, there is no requirement that what you understand must be true. This covers the object question.

On the other hand, the basing question—whether one can understand based on false beliefs—applies to and understanding theories, subjects, objects, and people, even in cases where the object question makes no sense. Though one can understand a false theory, it is still an open question whether that understanding must be based entirely on true beliefs (perhaps true beliefs about what the theory implies). Even though a person or a symphony cannot be true or false, it is still an open question whether one’s understanding of a person or a symphony must be based entirely on true beliefs (perhaps true beliefs about the person or symphony).

The basing question arises because many beliefs can be implicated in understanding. In understanding why boats float, one must draw on beliefs about buoyancy, density, pressure and so forth. As a first approximation, the basing question is whether these beliefs must be true.⁵ In each case, we can ask whether the beliefs that are implicated in understanding must be true.

⁴A potential problem here is that one can say to a colleague, “I understand that you got a promotion,” as a way of expressing that one has reason to think she got a promotion but that one is not sure and is seeking confirmation. However, it is plausible that this is a different use of ‘understanding that’ that does not affect the use that concerns us here.

⁵There is an analogous debate regarding knowledge, but this debate does not occur under the heading ‘Is knowledge factive?’ but rather under the heading ‘Can knowledge be based on false lemmas?’ A defense of the ‘No false lemma’ thesis can be found in (Harman 1973), and objections can be found in (Klein 2008).

Which beliefs are implicated in understanding? Without taking a stand on any specific account of understanding, we can say that if understanding is constituted entirely by beliefs as reductive theories claim, then the beliefs implicated in understanding are simply the beliefs that constitute understanding. If, on the other hand, understanding is a kind of ability, then the beliefs implicated are the beliefs that contribute to the ability. Some cognitive abilities depend on the right basis of beliefs from which to draw appropriate inferences. For example, the ability to calculate a projectile's path depends on beliefs about Newton's laws, or similar beliefs. The beliefs are implicated in understanding to the degree that they contribute to the cognitive ability that is understanding.

Given the above, the simplest way to clarify the basing question is to ask whether the following strict truth condition is met:

Strict Truth Condition If S understands a topic, then every belief implicated in S's understanding that topic is true.

I use 'topic' as a general term for the different objects that understanding can take as complement, including propositions, why-questions, persons, works of art, subject matters, and theories. According to the Strict Truth Condition, a single false belief relating to X precludes understanding X.

However, this way of defining the issue is too strict to be interesting; understanding almost trivially fails the Strict Truth Condition. To illustrate, suppose Sally has a single false belief about a minor aspect of the circulatory system of bats. Other than this, Sally has only true beliefs and is an expert on the bat circulatory system. She can explain in detail how every part of it works (except for the one minor error), and she can successfully experiment on it and manipulate it. By any sane standard, Sally understands the bat circulatory system despite her one false belief. Thus, the Strict Truth Condition is false.

To locate a more interesting, less strict way of addressing the basing question, it helps to recall that understanding comes in degrees. We are concerned with whether one can acquire understanding by gaining false beliefs. Since understanding is not just all-or-nothing, we should not restrict ourselves to all-or-nothing truth requirements. What matters is the degree to which the beliefs implicated in understanding are connected to the truth, in terms of how many true beliefs versus false beliefs one has, and how closely one's false beliefs approximate the truth. Let us call the degree to which a body of beliefs approaches the truth the body's *degree of veracity*.

Now, the motivation for asking the basing question is to figure out whether increasing understanding requires approaching the truth by increasing the degree

of veracity of one's beliefs. Does greater understanding of a topic imply believing more truths about the topic? Does less veracity among one's beliefs lead to a lower degree of understanding? In short, can false beliefs bring about greater understanding?

To summarize, understanding is based on background beliefs and (perhaps) cognitive capacities. If we hold capacities fixed, the question is whether false beliefs are capable of increasing understanding. The issue can be encapsulated in the following condition:

Covariance Condition The degree to which one understands X varies in the same direction as the degree of veracity of the beliefs that are implicated in understanding X, holding cognitive capacities fixed.

The Covariance Condition is motivated by the idea that the success involved in understanding is having true beliefs, so that degrees of veracity and degrees of understanding always vary in the same direction. If one can gain understanding by lowering the degree of veracity of one's beliefs, then the answer to the basing question is 'No, understanding need not be based entirely on truth.' Thus, to answer the basing question, we just need to be able to compare degrees of veracity and degrees of understanding to determine whether the covariance condition holds.

2.1 Comparing degrees of veracity

In order to find counterexamples to covariance, we do not need a full account of degrees of veracity, we only need sufficient conditions for when a body of beliefs A has greater veracity than another body of beliefs B. Three such sufficient conditions are:

Body of beliefs A has a higher degree of veracity than body of beliefs B if

1. A and B contain the same beliefs, except for additional true beliefs in A,
2. A and B contain the same beliefs, except for additional false beliefs in B,
3. A and B only differ in that A contains beliefs that are closer approximations to the truth while B contains inferior approximations.

Whatever other complications are present in determining degrees of veracity, adding true beliefs and nothing else increases the degree of veracity. Similar considerations apply to the other two conditions. If we deal only with cases involving the above three conditions, we avoid the need to deal with trade-offs between more false beliefs in one area compensated by more true beliefs in another, and other trade-offs. We also avoid the problem of determining how to count beliefs. These sufficient conditions for greater veracity will enable us to show that one can achieve a greater degree of understanding with a lower degree of veracity, in violation of the covariance condition.

2.2 Comparing degrees of understanding

Assessment of putative counterexamples to the Covariance Condition also requires a way of comparing degrees of understanding, in addition to a way of comparing degrees of veracity. We have already identified three sufficient conditions for comparing degrees of veracity of beliefs. In the literature, previous descriptions of degrees of understanding are in terms of degrees of “breadth and depth” (Kvanvig 2003), or the degree of one’s ability to “connect, synthesize, and grasp a body of information,” or “unify a body of information,” (Elgin 2007, p. 41). While these are ways of describing degrees of understanding, it would be preferable to have measures that are more specific. This section identifies two such methods of comparing degrees of understanding.

In everyday situations as well as in more formal endeavors such as psychology and education, we often need to make comparisons of degrees of understanding. Two methods frequently used are to compare:

1. The ability to answer questions correctly, including novel questions, and
2. Reasoning ability.

The greater one’s ability to answer questions about a topic, the greater one’s understanding of that topic. The greater one’s reasoning ability on a topic, the greater one’s understanding of that topic. Both of these allow us to make comparative judgments about degrees of understanding sufficient to identify counterexamples to Covariance.

The first method of comparison is to compare the ability to answer questions regarding a topic and to answer them correctly. This is the method that teachers use to test students, employer’s use to test job candidates, and psychologists use to test subjects. Novel questions are important to such tests, because verifying that

someone understands a subject requires more than giving them a series of questions or problems that they have already encountered, to which they could have simply memorized the answers. To test whether someone understands quantum mechanics, one asks novel questions about quantum mechanics. To test whether someone understands why solar eclipses occur, one asks novel questions regarding why solar eclipses occur. The wider the range of questions one can answer, the greater one's understanding.⁶ In general, to test for understanding, we present subjects with a range of questions from the familiar to completely novel. This requires them to go beyond what they may have memorized and demonstrate their understanding.⁷

When using this measure, it is important to be clear about which topic is under consideration. For example, when we ask, "How well can Susan answer questions about morality?" we might be asking how well she can answer questions about what is right and wrong. On the other hand, we might be asking how well Susan can answer questions about why actions are right or wrong. One can be able to answer questions about what is right or wrong without being able to answer any questions about why actions are right or wrong. Suppose utilitarianism is true, but that one can better determine what is right and wrong in specific cases using Kantian ethics. Then one who believes Kantian ethics is correct will, according to this measure, understand what is right and wrong better than one who believes utilitarianism is correct. On the other hand, she has worse understanding of *why* actions are right and wrong.⁸ Clarity about the topic is important.

The rationale behind this method is that the ability to answer questions, especially novel questions, expresses the underlying state of understanding. It captures the fact that understanding contributes to our grasp of reality through allowing us to find the truth. Genuine understanding is not idle. Either understanding is itself an ability that contributes to answering questions, or it is constituted by beliefs that are readily applied to allow one to answer questions. Because understanding plays this role, the ability to answer questions reveals degree of understanding. If there are no interfering factors, such as intoxication or uncharacteristic perfor-

⁶Like any method of measurement, interfering factors can prevent the use of this method from giving an accurate result. This practical concern about application will not affect the counterexamples to covariance given below.

⁷This is nearly universally accepted in the education literature and in psychology. For examples from education, see (Skemp 2006) and (Wiggins and McTighe 2005). For examples from psychology, see research on the illusion of understanding, such as (Rozenblit and Keil 2002) and (Fernbach, Rogers, Fox, and Sloman 2013).

⁸I thank an anonymous reviewer for this example.

mance errors, then the test of correctly answering questions reveals the degree of understanding. While individual tests may not perfectly implement the strategy, the principle that greater understanding produces greater ability to answer questions is all we need here. The basing question, in these terms, is whether one can increase one's ability to answer novel questions solely through adopting an overall less true set of beliefs. If one can do this, then the Covariance Condition fails.

The second method of comparing degrees of understanding is by level of reasoning ability. When someone understands a topic, they are able to reason well about that topic. Greater reasoning ability implies greater understanding. Either understanding is an ability that takes part in reasoning, or it is constituted by beliefs that enable better reasoning. There are various aspects of reasoning ability regarding a topic: the ability to distinguish the relevant from the irrelevant, the ability to bring the right facts to mind at the right time, and the ability to evaluate how evidence relates to conclusions within the topic. For example, if one can tell which equation to apply to a physics problem, one understands physics better than someone with the same mathematical abilities but who struggles to figure out which laws to appeal to. (As with the previous measure, clarity about the exact topic is important.)

This method can be used to assess understanding when one is not in a position to ask the subject questions. Suppose we wish to assess whether John or Larry understands Susan better by studying essays they wrote about her in the newspaper. Suppose each essay only contains true statements, but John relates facts and anecdotes that capture the central features of Susan's character while Larry relates facts and anecdotes that give a completely misleading picture of her character. We are not in a position to ask them any questions about Susan, but we can still compare their degrees of understanding. Assuming it was not Larry's goal to be misleading, he does not understand Susan as well as John does, for he displays worse reasoning abilities about Susan. John can properly assess what is relevant or important about her, while Larry cannot.⁹ Thus, reasoning skills provide another method upon which we have an independent grasp, and which can provide information about degrees of understanding.

Both measures are intended to measure understanding, and thus they should agree on cases. There might be concern that they do not in fact measure identical states. Rather than arguing that they do measure the same underlying state,

⁹Reasoning ability is closely related to the ability to answer questions, since reasoning ability is what allows one to answer questions. However, they are different methods. One can assess novel questions without directly addressing reasoning skills, and one can assess reasoning ability without asking questions.

consider the consequences if we assume they do not. What would be missing from the ability to answer questions that prevents it from being a complete measure of understanding? The missing component is *how* one answers questions. Some ways of answering questions, such as memorizing a list, are effective but not enough for understanding. Understanding is not just answering questions correctly, but answering them in the right way. In other words, one must answer them using proper reasoning ability. This suggests, then, that increasing both ability to answer questions and reasoning ability increases understanding, and likewise for decreases. When the measures agree, that is sufficient.

On the other hand, what is missing from reasoning ability that prevents it from being a complete measure of understanding? Here, the missing component is proof that the reasoning in question is not just idle, mental wheel-spinning. Considering reasoning alone could lead to the conclusion that understanding is merely a useless sense of satisfaction with how one can reason about a topic (see (Trout 2002) for more about this objection). One must be able to apply this reasoning ability for it to be understanding. Thus, the potential deficiency of this measure is addressed by the first measure. Even if the the ability to answer questions and reasoning ability do not coincide, together they are sufficient to measure changes in understanding. If both measurements give the same verdict, as in the cases below, they are adequate to identify counterexamples to Covariance.

The above measures of understanding provide sufficiently precise measurements of degrees of understanding. The counterexamples to Covariance below involve agents for which the differences in degrees of understanding are larger than any uncertainty in the applications of the measures.

3 Cases of understanding through falsehood

We now have enough information to compare degrees of veracity among bodies of beliefs and to compare degrees of understanding. In this section, we employ these methods to examine a series of cases in which an agent with greater veracity among the beliefs implicated in understanding nevertheless has a lesser degree of understanding.

3.1 Approximation and idealization

The first counterexample to Covariance is approximation.¹⁰ Approximation is the use of a value, equation or statement that is not exactly correct but that is nevertheless close to the truth.¹¹ Approximations are strictly speaking false. If it is exactly 45.2 degrees Celsius, 45 degrees Celsius can often serve as a good approximation.¹² The temperature is not 45 degrees, but it is approximately 45 degrees.

Like approximation, idealization is the use of falsehoods in problem solving. An idealization is a simplified version of a concrete situation that abstracts away from some features for the sake of clarity.¹³ Some idealizations are approximations, but some are not. As an illustration of this, in formulating linguistic theories, linguists employ the idealization that speakers have infinite capacities to perform certain kinds of operations (e.g. embedding), abstracting away from human limitations. However, humans do not even approximately have infinite capacity.

While approximations and idealizations are strictly false, they are employed in reasoning as if they are true. One makes derivations from them as one does from any belief. While we often use approximations with the explicit knowledge that they are approximations, suppose one believes that an approximation is literally, exactly true. Then one has a belief that is strictly speaking false. If this false belief were removed and replaced with the exact truth (or a better approximation), the overall body of beliefs would have greater veracity. There are cases where believing a closer approximation to the truth leads to lesser understanding.

Consider physics, a field where approximations are ubiquitous. Today, we know that Newtonian Mechanics (NM) is not the exact truth. It is an approximation that works well in many of the situations where we need to use physics.

¹⁰Elgin also discusses approximations and outdated scientific theories in her argument against factivity in (Elgin 2007). However, she does not use them as her counterexamples to factivity; she allows that they might provide only honorary understanding, not genuine understanding. Idealization is the only example she defends as a counterexample to factivity. Thus, by identifying approximation as a counterexample, the argument in this paper goes further. For more details about Elgin's arguments and where they differ from mine, see section 5.1.

¹¹The notion of "close enough" is difficult to define exactly, but for our purposes we can rely on a working sense of approximation. Approximation has been explored for example in (Elgin 2004).

¹²The statement "The temperature is between 45 and 46 degrees," is not an approximation in the same situation, since it is a strictly true statement about the temperature. The fact that it does not give the exact temperature does not make it an approximation.

¹³There is some controversy regarding just how accurate this account is, for example see (Cartwright 1995) and (Kowalenko 2009), but it is good enough to give a working sense of what idealization is.

Strictly speaking, it is false. General Relativity (GR), while probably not exactly true either, is a better approximation that gives accurate predictions in a much wider range of circumstances. Thus, believing NM rather than GR is a decrease in degree of veracity; a shift from a closer approximation to a worse approximation. Let us assume for the sake of argument that GR is exactly true. Consider a student who has been taught GR, and has enough background in math to grasp the meaning of the claims. Consider also a second student who has the same background but has learned Newtonian Mechanics and believes it to be the exact truth. The second student has fewer true beliefs (by lacking belief in GR), and more false beliefs (by believing in the exact truth of NM). Otherwise, they are the same. Thus, the first student's beliefs have a greater degree of veracity.

Suppose both students consider the subject of the motions of heavenly bodies. The student who only knows GR spends hours fruitlessly trying to solve the field equations. She does not come close to describing the motion of any heavenly body in any form, no matter how approximate.¹⁴ On the other hand, the student using NM with the same level of mathematical skill is able to derive equations for orbits that are approximately correct for several bodies. As she acquires more true beliefs about celestial motion, she continues to use NM fruitfully in answering further questions. With this in mind, let us turn to their degrees of understanding.

First, compare their degrees of understanding using the ability to answer questions. Given a series of novel questions about celestial motion, the student who believes in NM, and whose beliefs have lower veracity, is able to answer many correctly. On the other hand, due to the difficulty of solving Einstein's Field Equations, the student who knows GR cannot answer any. Finding solutions to Einstein's Field Equations is so difficult that solutions are named after their discoverers (e.g. the Schwarzschild solution). Since the subject is celestial motion and not the universal laws governing such motion, by the measure of answering questions, the student whose beliefs have less veracity has greater understanding, which implies that lower veracity among beliefs can provide greater understanding, in violation of the Covariance Condition.

Next, we compare degrees of understanding by comparing reasoning ability. Return to our two students before they learn physics, whether NM or GR. They each contemplate the motions of heavenly bodies, but neither one has the ability to reason through the topic of planetary motion. They have no sense of what keeps everything moving in the observed orbits, or how hypothetical bodies would

¹⁴Anyone who has learned GR should recognize the plausibility of this happening to a student who has more than enough mathematical skill to master NM.

move. Then one student learns NM, taking it to be the exact truth, and the other learns GR. The first student can identify which information is relevant, which factors play a significant role in determining celestial motion and which do not. When confronted with features of the motion in the heavens, she can often explain what is going on; she can successfully navigate the topic. However, the student who learns GR is not in the same position. Relying on GR, the student knows the equations governing the universe, but cannot use them to work out even general constraints on how heavenly bodies move. Though possessing new true beliefs, the student is still confused. If this student comes to believe that GR is false and that NM is true, she will thereby increase her ability to reason about celestial motion and thus she will have greater understanding. Both measures give the same verdict, and the Covariance Condition fails.¹⁵

One might object that the above does not show that falsehood can increase understanding; there are true beliefs in the vicinity of the false belief (belief in NM), and these true beliefs do the work in understanding. The student who believes that NM is exactly true presumably also believes that NM is accurate to within practical limits, which is true. Perhaps it is only this true belief, and not her false beliefs, that are implicated in giving her understanding of celestial motion.

However, this objection fails to undermine the case. First of all, in realistic cases, she reasons from belief in NM, not from a hypothetical belief that NM is true enough. The fact that we can find true beliefs she has that she could have used is not relevant. In addition, one has greater understanding by the first measure when one has a greater underlying ability to answer questions. Understanding contributes to the ability to apply one's beliefs, whether understanding is constituted by beliefs or is itself an ability. In this case it is NM that contributes to the student's abilities. Once the student gains her false belief that NM is true, she has the ability to answer novel questions about celestial motion. She does not need to infer anything weaker before she actually has the ability. While she could infer that NM is true within error tolerance, she need not do so. Even if she uses NM to build up a large body of true beliefs about celestial motion, the false NM beliefs are still implicated in her understanding. It is not a ladder she has thrown away; it continues to play an active part in the success of her reasoning.

One has greater understanding by the second measure when one has greater reasoning ability. Again, once she comes to hold the false belief in the truth

¹⁵The same analysis can also be applied to idealizations that are believed to be true. The example of ideal gases is used in (Elgin 2007) to argue that understanding is not factive, though her argument has a weaker conclusion, as shown in section 5.1.

of NM, she then has the greater ability. She need not infer weaker beliefs to increase her skills. Furthermore, when actually making derivations, predictions and explanations, the actual belief she draws upon is false. Her derivations do not start from the premise that NM is true within certain tolerances; they start from Newton's laws. The belief about error tolerance need play no role in her reasoning. The false beliefs are central and are implicated in her understanding. They continue to play a role in her reasoning. Thus, approximation does provide an example of gaining understanding through falsehood.

Thus, by both measures of understanding—novel questions and reasoning ability—approximation is a counterexample to the factivity of understanding. One can gain understanding by lowering degree of veracity.

3.2 Misleading truths

Misleading true beliefs are another class of counterexamples to Covariance. A misleading true belief is a belief that inclines one to draw false conclusions. By adopting such a belief, one's overall body of beliefs gains a higher degree of veracity: either one has added one true belief and changed nothing else, or one has removed one false belief and replaced it with a true one. And yet, as I will now show, by both measures, one's degree of understanding decreases.¹⁶

As an example, consider someone who starts out knowing the basic facts of the JFK assassination, such as where it happened, the number of shots, kind of rifle and so on. Now suppose she learns that JFK's brain went missing in the 1960's. This is a new true belief. However, without information about the context, it suggests a nefarious plot that is not there. There is strong evidence that Robert Kennedy took his brother's brain so it could be buried rather than stored as a curiosity in the national archives. The missing brain is not evidence of a cover-up, but in isolation, the fact that the brain is missing can give rise to suspicions that can lead one off track. It is a misleading fact about the assassination.

First compare degrees of understanding before and after learning the misleading truth in terms of the ability to answer novel questions. A misleading truth just is a true proposition that leads one through good inferences to false conclusions. This is the sense in which the truth misleads: when one treats it as one treats other truths, one is led astray. It is the phenomenon summed up in the famous, "A little learning is a dangerous thing."

¹⁶The significance of misleading truths is widely discussed outside the literature on understanding. For example, it figures prominently in the debate about the epistemology of disagreement.

To return to our example, from the fact that JFK's brain went missing, one can correctly infer that it did not fail to go missing, but this is not much help. On the other hand, the false conclusions about conspiracies are more consequential. Given an array of questions about the Kennedy assassination, one would do worse in answering them as a result of the true belief (except those questions relating directly to the location of the brain). The primary effect is doubt, suspicion and inability to draw conclusions. This hurts one's ability to answer questions about the assassination, questions that one could answer if one lacked the misleading true belief. While this will eventually lead one to false beliefs, the hindrance in one's ability to answer questions occurs before false beliefs are inferred. Thus, a true belief can decrease understanding, as measured by the ability to answer questions, in violation of Covariance.

In terms of reasoning ability, misleading true beliefs can decrease reasoning ability in a topic, thereby lowering understanding. If one understands a subject then later adds a new true belief that seems to imply much of what one believes about the subject is wrong, one's abilities to reason go down. With the misleading truth, the effect can be to leave one unsure of what to do, cognitively speaking, perhaps even giving one the false sense that one is doing well when in fact one is confused.

In the Kennedy example, learning about the missing brain leads someone who had a basic sense of how to reason about the assassination to harbor vague suspicions of conspiracies. Worse, the true belief itself provides no clear idea where the suspicions lead or what to do with them. The belief can make irrelevant details seem relevant, as if they are part of a brain-concealing plot. The person does not know what inferences to draw, since the obvious inferences from the misleading belief conflict with the implications of other beliefs. The result is someone who is stuck, confused, and lacking understanding. This effect occurs even before one forms false beliefs as a result of the misleading truth. Measured by reasoning ability, adding this true belief can lower understanding. Both measures show the increase in veracity decreases understanding, so the Covariance Condition fails.

Any misleading truth one cares to mention can give rise to similar counterexamples. As long as the misleading conclusions one is inclined to draw are central to understanding a topic, misleading truths hinder understanding.

The objection raised against approximation and idealization—that the difference in understanding is due to the veracity of nearby beliefs—does not apply here. Suppose John has a misleading true belief. Susan is otherwise in the same position but lacks that misleading true belief. We cannot attribute the difference in understanding to true beliefs that Susan has but John lacks. The difference in the

degrees of veracity of their beliefs goes precisely in the opposite direction: John has more true beliefs. The objection fails. Misleading true beliefs stand as another counterexample to Covariance.

3.3 Summary

The above counterexamples show that the Covariance Condition does not apply to understanding. One can gain understanding by adopting an overall less veridical body of beliefs, and one can decrease understanding by adopting an overall more veridical body of beliefs.

Furthermore, the above examples are not simply outliers that one can set aside. Cases such as these are prevalent. For one, the history of science contains a long succession of people believing approximations or idealizations to be the exact truth, for the limits of the approximation were not yet discovered. This alone shows that counterexamples are not insignificant outliers. Further, misleading truths crop up frequently: they are deliberately promulgated in politics and advertising, and they are a common result of our cognitive limitations. It often happens that we think some fact points to a conclusion, but it turns out not to. The failure of Covariance is not a mere logical trick; it is a central feature of understanding.

Yet surely there must be some negative effects due to having lower veracity of beliefs. Truth must have some role to play in understanding. The next section considers in more detail what effect the degree of veracity of beliefs has on understanding.

4 The role of truth in understanding

I have shown above that understanding can be gained through false beliefs. However, the lesson is not that truth plays no role in understanding.¹⁷ Instead, the lesson is that understanding is not entirely a matter of already having the truth, but is also a matter of *being able to find* important truths. Beliefs take part in understanding via how they contribute to skill in discovery.¹⁸ In order for a belief to

¹⁷Indeed, if one accepts the argument in (Klein 2008) that false beliefs can play a role in producing knowledge, it does not prevent understanding based on false beliefs from playing a role in producing knowledge.

¹⁸This does not imply that before a belief can contribute to understanding one must actually use it to discover information. Understanding is an ability, and one can have an ability that one does not use.

increase understanding by either measure used above, it need only have the right inferential profile. That is what allows it to contribute to reasoning skills and the ability to answer questions correctly.

That said, having a true belief that p at least allows one to correctly answer whether p is true. If the question whether p is true is part of the topic under consideration, then the true belief will help understanding. Further, a true belief has no false deductive consequences, whereas all false beliefs do. Thus, there must always be limits on how much a false belief can increase understanding. Consider the physics students described above. The student using NM understands celestial motion better than the student using GR because given their level of mathematical competence, neither student can use GR. They are not capable of dealing with those questions where GR gives correct results and NM does not. However, with greater mathematical acumen, the student using GR can do better than the student using NM. The student using GR may even be able to derive NM as an approximation, at which point she will have a complete advantage, by being able to use NM while recognizing that it is but an approximation. In the examples above, only cognitive limitations prevent the truth of GR from producing greater understanding to the student who believes it is true.

There are some falsehoods that we do not have the capacity to use badly. Given our state of technological and scientific development, some falsehoods will only have observably false consequences in situations we will rarely or never experience (consider NM in the eighteenth century). In cases such as these, even though agents' beliefs are false, it is not the falsehood of the claims that plays a role in limiting degrees of understanding. The limitation in the degrees of understanding is due to limitations of skill, know-how and technology available. With these limitations, one's understanding is rarely hindered by the fact that the claims are false. When these limitations are removed, the truth once again makes a difference to understanding.

This also illustrates the importance of clarifying the topic at issue. While cognitive limitations prevent the students in the example from using GR to answer questions about celestial motion, the student who believes in GR can still answer many questions about the theory itself. While NM did not produce detectably false predictions (aside from a couple of puzzles) in the eighteenth century, it was still a false theory. Thus, if the topic is the general laws of nature rather than simple motion, then the person using NM no longer has the advantage over the person using GR, even if NM is usually more useful.

The problem with false beliefs is that they cannot be entirely isolated. Even useful falsehoods can only have a restricted usefulness in finding information

about the world. A measure of understanding is the ability to answer questions correctly, and any false belief will suggest some false answers. For some false beliefs, the restrictions might be insignificant—like NM in the 18th century—but there are restrictions nonetheless. If one continues to use an approximation or idealization in a situation where greater precision is needed, then failure is the expected outcome.

For this reason, one has an advantage if one realizes that an approximation is not strictly speaking true, or that an idealization is merely an idealization. To summarize the lessons of these examples and how they relate to each measure of understanding, the following list ranks states of belief and skill in terms of their degree of understanding, holding cognitive capacities fixed. The first has the greatest understanding, the last has the least.

1. One has only true beliefs and the ability to use them effectively to solve problems;
2. One has true beliefs that one cannot effectively reason from, but one also has access to useful falsehoods that one can use to solve problems, with the knowledge that they are falsehoods;
3. One has false beliefs that one can use to solve problems.
4. One has true beliefs that one cannot use to solve problems.
5. One has false beliefs that are not useful.

State 1 occurs when one has true beliefs about a topic and can reason from those beliefs directly. State 2 occurs when one has true beliefs about a topic, but cannot reason from them directly (for example being unable to use the equations of GR), yet one knows of an approximation or idealization that one can use to solve problems and answer questions. This is a lesser degree of understanding than 1, since with 1, in addition to the ability to use approximations, one has the ability to answer further questions about how to derive solutions to problems from first principles and one has greater skill in reasoning from these beliefs. With 2 one cannot answer these questions, relying as one does on useful falsehoods to make derivations rather than being able to derive from first principles, and lacks the skill to reason from the beliefs.

States 3 and 4 are the two states that are at issue in the counterexamples to Covariance. State 3 is a state of greater understanding than 4 due to the possession of useful false beliefs. In these terms, the counterexamples in the previous

section show that 3 is correctly placed above 4, which demonstrates the failure of Covariance.

Finally, state 5 is a state of pure error. At least in state 4, one can correctly answer the questions about the true beliefs one has, though one cannot infer anything important from those beliefs. One in state 5 lacks even this ability since the beliefs are false.

The hierarchy shows that we do not need to deny the importance of truth in order to deny the Covariance Condition. The highest degree of understanding results from true beliefs paired with the ability to reason from those beliefs. Truth is important to understanding, but there can be trade-offs due to our finite cognitive capacity. Sometimes we can trade off accuracy for efficacy.

5 Covariance and factivity

The Covariance Condition is not only interesting in its own right. This section shows that if the Covariance Condition fails, then understanding is not factive, in the sense of ‘factive’ used in recent debates in (Kvanvig 2003), (de Regt and Dieks 2005), (Elgin 2007), (Pritchard 2008), (Riggs 2009), and elsewhere. In this debate, the relaxed version of the Strict Truth Condition is a condition known as ‘factivity’—first proposed in (Kvanvig 2003):

Factivity Understanding X is *factive* just in case understanding X implies that all of one’s central beliefs about X and most of one’s peripheral beliefs about X are true.

Factivity avoids the problem of over-strictness that afflicts the Strict Truth Condition. Sally, with her single false bat belief, can understand the bat circulatory system without being a counterexample to factivity. Unlike the Covariance Condition, however, factivity treats understanding as an all-or-nothing matter. It does not address the fact that understanding comes in degrees. Nevertheless, with suitable assumptions the failure of the Covariance Condition implies the failure of factivity.

First, factivity allows that one can understand despite some minor false beliefs, a concession made to avoid being excessively strict. However, it is against the spirit of factivity for false beliefs to benefit understanding. As Kvanvig puts it in his defense of factivity, “When the falsehoods are peripheral, we can ascribe understanding based on the rest of the information grasped that is true and contains no falsehoods. In such a case, the false beliefs are not part of the understanding

the person has,” (Kvanvig 2003, p. 201). According to factivity, false beliefs are deviations from the ideal; they are analogous to the microscopic deviations from ideal flatness of a flat table. These deviations do not contribute to the flatness of the table; rather, they are imperfections that we can ignore. Likewise, one can understand despite minor false beliefs, for they are imperfections that we can ignore. However, the false beliefs are imperfections that cannot contribute to understanding any more than microscopic bumps contribute to a table’s flatness. Therefore, if false beliefs that are implicated in understanding can increase understanding, then factivity falls. In this way the Covariance Condition is an extension of factivity that accounts for degrees of understanding.

We can also show that failure of Covariance implies failure of factivity by relating degrees of understanding to all-or-nothing understanding. The Covariance Condition relates changes in degrees of understanding with changes in degrees of veracity: they must vary in the same direction. One understands in an all-or-nothing sense when one’s degree of understanding is above some threshold. If Covariance does not hold, problems emerge at the threshold.

Given that there is a threshold, it is possible for there to be an agent who does not understand but has a degree of understanding exactly at the threshold. (This assumes that understanding requires being above the threshold. An analogous argument can be given if we assume that one at the threshold understands. In that case, assume the agent gains a true belief that removes her understanding. The same conclusions follow.) If Covariance does not hold, the addition of a false belief that is implicated in her understanding can increase her degree of understanding, putting her above the threshold. The only way this could be avoided would be if every violation of Covariance occurred away from the threshold. While this is not logically impossible, the widespread counterexamples above render it sufficiently unlikely to discount the possibility. Furthermore, if, as thresholds for flatness and tallness suggest, the threshold for understanding is context-dependent, the hypothesis that Covariance violations never occur at the threshold is even less plausible. In other words, merely acquiring a false belief that is implicated in understanding can take her from a state of not understanding to a state of understanding. This violates factivity.

In short, factivity is the claim that understanding must be based on true beliefs. While this deals with understanding as an all-or-nothing state, Covariance is the natural extension of factivity to understanding as a graded state. Further, due to the relationship between thresholds and degrees of understanding, violations of Covariance will produce violations of factivity. Thus, the arguments presented above not only show that understanding violates Covariance, they show that un-

derstanding is not factive.¹⁹

Finally, while failure of Covariance implies failure of factivity, the converse is not true. For example, (Riggs 2009) claims that one can understand a topic even if some of one's central beliefs about the topic are false. Thus, if he is correct, factivity is false. However, he still accepts Covariance, claiming that the threshold for understanding can be low enough to violate factivity, but veracity and understanding vary together. The denial of Covariance is thus a stronger claim than the denial of factivity: there are theories of understanding that violate factivity but not Covariance, while the argument above shows that no plausible theory of understanding can violate Covariance without violating factivity.

5.1 Elgin's argument against factivity

The most detailed argument against the factivity of understanding so far is given by Catherine Elgin in (Elgin 2007) and (Elgin 2009). Unfortunately, her argument is subject to serious objections. I will now describe her argument and show that it faces objections that the argument presented in this paper avoids.

Elgin's counterexample to factivity is the use of idealization in science. While Elgin discusses examples similar to those I present in Section 3.1, examples she refers to as "felicitous falsehoods", she does not employ these as her counterexamples to factivity. She allows for the possibility that the false beliefs in cases of approximation only contribute to honorary understanding, not genuine understanding, and they do so only in virtue of the fact that belief in approximations was part of a process that eventually led to accurate beliefs. We had to go through stages of false beliefs to achieve true understanding (e.g. we had to first discover NM before we could get to GR), but only the final state is genuine understanding (Elgin 2007, p. 37). The argument she presents is thus neutral on whether approximations or past false theories are counterexamples to factivity.

¹⁹There may be a response Kvanvig could give to the above argument. His primary concern in (Kvanvig 2003) is to find a state that is like knowledge but that is not subject to the value problem, the problem that knowledge is not more valuable than its constituents. Thus, he could claim that what he calls 'understanding'—the state of having a tightly interconnected body of true beliefs—is such a state, though it is not what we normally refer to as 'understanding.' Such a response may not succeed, however. It is not clear that interconnectedness in itself provides a distinctive epistemic value, unless the interconnectedness allows for better ability to apply those beliefs. Application is what the measures appealed to in my argument above capture, which suggests that understanding of a nonfactive variety is the epistemically valuable state that Kvanvig seeks. There is not space to fully adjudicate this matter here, but these considerations are suggestive.

Her counterexample is idealization. She claims idealizations, though false, function as exemplars of relevant features of the system under study, “affording epistemic access to the features and making their significance manifest,” (p. 41). Some theories are so complex we cannot recognize some of their significant features, while the simplifications of idealizations can bring these features out. By exemplifying relevant features, idealizations allow us to better grasp and unify information about a theory and, if the theory is true, to better grasp and unify information about the world.

For Elgin, this does not yet show that idealizations provide understanding, for her argument depends on showing that they are not just felicitous falsehoods. To demonstrate that they are not, she appeals to a second premise: that the falsehood of idealizations is not a defect. In defense of this, she argues, “The ideal gas is a fiction ...” (p. 40). “Fictive sentences neither are nor purport to be true. They function in other ways. It is no defect in ideal gas descriptions that there are no gases that instantiate them” (p. 41). Thus, idealizations are false, but are a central part of improving understanding, in violation of factivity.

This argument faces objections. First of all, as Elgin herself says, the use of idealization need not involve false beliefs. In fact, Elgin emphasizes that in her example idealizations are knowingly used as fictions to illuminate a true theory. Thus, her example of idealization in science is not an example of false beliefs contributing to understanding—scientists don’t believe that the idealizations are true—and thus is not a counterexample to factivity.²⁰

Similarly, Elgin’s argument depends on the claim that idealizations work by exemplification. (Strevens 2013) objects to this, claiming that the use of idealizations is governed by different standards of correctness. When something is treated as an idealization, it is understood that we need a “translation manual” to determine what the idealization actually says about the world. Further, he claims, successful sciences uses idealizations whose translations are true. That is when they produce understanding. Thus, if he is correct, Elgin’s argument depends on a mistaken account of idealizations. The use of idealizations does not involve false beliefs, because when properly understood idealizations are not false claims about the world. Idealizations are no longer counterexamples to factivity.

Finally, (Mizrahi 2012) argues that scientists continually improve their idealizations, rejecting old ones in favor of new ones that work better. This suggests that, just like “felicitous falsehoods,” idealizations are intermediate steps of falsehood that allow us to work eventually to true theories. If Elgin is neutral about

²⁰A similar point is made in (Mizrahi 2012, p. 247).

felicitous falsehoods, by parallel reasoning she should be neutral about whether idealizations are a counterexample to factivity. It may only provide honorary understanding in virtue of leading us toward more accurate theories.

The argument in this paper avoids these objections. The counterexamples I provide do not rely on controversial claims about how idealization works in science. The counterexamples are not restricted to scientists who know that they are using idealizations as fictions, and thus lack false beliefs about the topic; they involve agents with straightforwardly false beliefs. The measures of understanding show that some felicitous falsehoods can improve understanding. Thus, my argument is not subject to the objections Elgin faces, and factivity remains overturned.²¹

5.2 Riggs' argument against factivity

(Riggs 2009) argues against factivity based on a case that he claims elicits the intuition that factivity is too strict. The case is this: Wayne's wife is afraid of water, and he understands her fear of water based on the belief that she was in a boating accident at the age of three. However, she was never in a boating accident. Since his false belief about the supposed accident is almost all he believes about her fear of water, it is a central false belief about the topic of her fear of water. Riggs claims that intuitively, despite the falsehood, Wayne still understands his wife's fear of water. The false belief allows him to understand the nature of the fear, if not its origin.

This argument relies on intuitions about the case; Riggs does not present a metric by which we can measure whether it is a case of understanding. Unfortunately, the intuition that forms the basis of the counterexample is controversial. It is not clear that Wayne understands his wife's fear of water based on completely false beliefs about a childhood accident. Since she is his wife, it is likely that he has witnessed the effects of her fear and spoken to her about it, so these true beliefs, rather than the false belief about an accident, may be central. If the example were changed to exclude additional true beliefs, it is less plausible that he understands her fear. It is still plausible to conclude that he does not understand her fear of water, though there is some intuitive pull in the other direction.

Further, it is not clear whether the false belief is implicated in understanding her fear of water. Unlike NM, which is used in many calculations of celestial

²¹This is not to say that my argument is in conflict with Elgin's, only that my argument goes further in demonstrating counterexamples to factivity.

motion and thus clearly plays a role in understanding, it is plausible that Wayne's belief about a boating accident plays a role only by allowing him to infer that his wife is deeply afraid of water. This latter belief is true. The conclusions that require inference from the false belief will largely be false, such as the conclusion that a boating accident was the precipitating event. Without a more precise way of evaluating the case, it thus appears inconclusive.

On the other hand, the argument in this paper avoids appealing to controversial intuitions about specific cases. The argument above is based on measures with which to test cases, and the cases are not isolated but are broad classes of counterexamples. The specific measures of understanding presented above are independently motivated. They are the measures that we actually use when we need to test for understanding in the world. The argument above avoids the limitations of Riggs' argument.

The argument in this paper not only refutes Covariance, but shows that factivity fails using arguments that avoid the pitfalls of existing arguments.

6 Falsehood is here to stay

One person's modus ponens is another's modus tollens. One could take the above arguments to show that the two measures of understanding, while often good enough for many purposes, are inaccurate in some cases. In particular, the cases above may be due to a breakdown in these methods of measuring degrees of understanding. In that case, the above arguments may just reveal places where these measures go systematically wrong; they do not reveal the failure of Covariance or factivity. In this section, I resist this line of reasoning and argue that there is good reason not to impose the Covariance Condition.

The key observation is that understanding can be applied. One who understands has the ability to produce explanations, or the ability to apply knowledge to cases, or the ability to sort the relevant from the irrelevant. This is what makes the measures compelling; it is why we use them to test students, job applicants, and psychology test subjects. It is anathema to the very idea of understanding that it could simply be idle. One could not at the same time understand physics while lacking the ability to do anything with physics. That is what the measures capture, the fact that understanding translates into ability, either by being an ability itself or by being the kind of state that is conducive to producing an ability given our cognitive make-up.²² The need to apply understanding motivates the measures

²²For arguments that understanding is an ability, see (de Regt and Dieks 2005), (de Regt 2009),

used in this paper, and the measures show that the core features of understanding that give it its value do not imply Covariance.

Furthermore, the conclusion above already allows for the importance of the truth for understanding; one must be able to find correct answers. Having true beliefs contributes to that, securing a role for truth in understanding. An additional Covariance or factivity requirement would be ad hoc.

In addition, we should expect violations of Covariance. The close tie between understanding and the ability to find the truth suggests the possibility of trade-offs. One ought to expect the possibility of trade-offs, where a decrease in truth can be compensated by benefits elsewhere to produce an increase in the ability. While knowledge requires truth because knowledge means having the truth, understanding gives the ability to find answers and reason well, which does not imply already having the answers. This explains why there is no parallel truth requirement for understanding. The conclusion that one can gain understanding via false beliefs does not overturn a defining characteristic of understanding, as it would for knowledge. There is no need to resist the conclusion that Covariance fails.

Even if understanding is not an ability, the fact that understanding comes in degrees opens the possibility that there can be trade-offs between truth and other factors conducive to understanding. Some have claimed that beliefs constitute understanding when there is a web of strong epistemic connections among the beliefs, as in (Kvanvig 2003). Such an account leaves open the possibility that a loss in veracity can be compensated by an increase in connections (and these connections improve one's ability to find correct answers).

Finally, the denial of Covariance does not eliminate truth from playing an important role in understanding, for it is only at intermediate steps of reasoning to a conclusion that false beliefs can increase understanding. The ability to answer questions correctly means finding true beliefs, for that is how understanding puts us in contact with reality. Though one can draw upon falsehoods to do this, truth must remain in the picture. In short, there is no rationale for imposing the Covariance Condition.

7 Conclusion

Understanding can be increased by the acquisition of false beliefs; understanding is not factive. This undermines claims that understanding is simply a kind of

(Grimm 2010), (Wilkenfeld 2013), among others.

knowledge that. Insofar as understanding is taken as an essential epistemic goal, it casts doubt on the claim that truth is the only epistemic good. While true belief still plays a role in understanding, sometimes the right false belief is better for understanding than the truth.

References

- Cartwright, N. (1995). False idealization: A philosophical threat to the scientific method. *Philosophical Studies* 77(2), 339–352.
- de Regt, H. (2009). The epistemic value of understanding. *Philosophy of Science* 76, 585–597.
- de Regt, H. and D. Dieks (2005). A contextual approach to scientific understanding. *Synthese* 144, 137–170.
- Elgin, C. (2004). True enough. *Philosophical Issues* 14, 113–131.
- Elgin, C. (2007). Understanding and the facts. *Philosophical Studies* 132(1), 33–42.
- Elgin, C. (2009). Is understanding factive? In A. Haddock, A. Millar, and D. H. Pritchard (Eds.), *Epistemic Value*, pp. 322–330. Oxford: Oxford University Press.
- Fernbach, P. M., T. Rogers, C. R. Fox, and S. A. Sloman (2013, June). Political extremism is supported by an illusion of understanding. *Psychological Science* 24(6), 939–946.
- Grimm, S. (2010). The goal of explanation. *Studies in History and Philosophy of Science* 41, 337–344.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Klein, P. (2008). Useful false beliefs. In Q. Smith (Ed.), *Epistemology: New Essays*, pp. 25–62. Oxford: Oxford University Press.
- Kowalenko, R. (2009). How (not) to think about idealisation and *ceteris paribus*-laws. *Synthese* 167, 183–201.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lewis, D. (1988). Relevant implication. *Theoria* 54, 162–174.

- Mizrahi, M. (2012). Idealizations and scientific understanding. *Philosophical Studies* 160, 237–252.
- Pritchard, D. (2008). Knowing the answer, understanding and epistemic value. *Grazer Philosophische Studien* 1(1), 20–50.
- Riggs, W. (2009). Understanding, knowledge and the meno requirement. In A. Haddock, A. Millar, and D. H. Pritchard (Eds.), *Epistemic Value*, pp. 405–428. Oxford: Oxford University Press.
- Rozenblit, L. and F. Keil (2002). The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science* 26(5), 521–562.
- Skemp, R. (2006). Relational understanding and instrumental understanding. *Mathematics teaching in the middle school* 12(2), 88–95.
- Strevens, M. (2013). No understanding without explanation. *Studies in History and Philosophy of Science* 44, 510–515.
- Trout, J. D. (2002). Scientific explanation and the sense of understanding. *Philosophy of Science* 69, 212–233.
- Wiggins, G. and J. McTighe (2005). *Understanding by Design, Expanded 2nd Edition*. San Francisco: Pearson.
- Wilkenfeld, D. A. (2013). Understanding as representation manipulability. *Synthese* 190, 997–1016.