Penultimate draft. Please cite final version.

The Role of Explanation in Understanding

Kareem Khalifa

Department of Philosophy

Middlebury College

Middlebury, VT 05753

Forthcoming in *The British Journal for the Philosophy of Science*

Abstract

Peter Lipton has argued that understanding can exist in the absence of explanation. We argue that this does not denigrate explanation's importance to understanding. Specifically, we show that all of Lipton's examples are consistent with the idea that explanation is the ideal of understanding, i.e. other modes of understanding ought to be assessed by how well they replicate the understanding provided by a good and correct explanation. We defend this idea by showing that for all of Lipton's examples of non-explanatory understanding of why p, there exists a correct and reasonably good explanation that would provide greater understanding of p.

Table of Contents

1. Lipton's framework ***3***

2. Clarifying Lipton's framework ***7***

3. Explanatory idealism ***9***

3.1. Misunderstanding without explanation ***10***

3.2. How to interpret Lipton's examples ***12***

4. Examples of modal understanding 15

5. Critical information and the Galileo example 19

6. Unification via tacit analogy 23

6.1. Explanatory unification 24

6.2. Unification and explanatory idealism 26

7. Tacit understanding of causes 29

7.1. Semantic knowledge 31

7.2. Tacit explanations 35

8. Conclusion 37

If the first decade is any indication, understanding promises to be a lively topic among epistemologists and philosophers of science throughout the 21st century[1]. While consensus is not yet established, many hold that explanation plays an indispensable role in understanding. De Regt ([2009], p. 25) expresses this idea succinctly:

> understanding a phenomenon = having an adequate explanation of the phenomenon.

Similarly, Kvanvig ([2003], p. 192) writes,

> Understanding requires the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information.

By contrast, Peter Lipton's ([2009]) posthumously published 'Understanding Without Explanation' offers several examples wherein explanation appears unnecessary for understanding.

At first blush, Lipton's examples might appear to diminish the centrality of explanation to understanding. However, we shall argue that Lipton's arguments are consistent with explanation still playing a special role as the *ideal* of understanding. More precisely, other modes of understanding ought to be assessed by how well they replicate the understanding provided by knowledge of a good and correct explanation. We call this position *explanatory idealism* about understanding[2].

Section 1 presents Lipton's general framework, while Section 2 clarifies certain aspects of that framework. Section 3 then shows how explanatory idealism clarifies certain lacunae in Lipton's position. Most notably, explanatory idealism entails that for every instance of non-explanatory understanding, an explanation exists that would provide greater understanding. Sections 4 through 7 then apply this thesis to Lipton's various examples of understanding without explanation.

1. Lipton's framework

Lipton's arguments assume two things. First, Lipton claims to be only concerned with cases of understanding *why* something is the case:

> I am sympathetic to the broader conception of understanding that encompasses understanding *how* in addition to understanding *why*, but my present purpose is to show that, even on a narrow conception of understanding as understanding why, we may nevertheless get understanding without actual explanation (Lipton [2009], p. 54).

Like Lipton, we focus on understanding-why, but with one small caveat. Both Lipton and our points also apply to certain uses of 'understanding how.' For instance, 'How do

amoebae reproduce?' is an explanation-seeking question, and thus perfectly relevant to the issue at hand.

However, our discussion does not concern other forms of understanding-how, e.g. understanding how to sew. Lipton is only distancing himself from this 'procedural understanding' in the passage above[3]. Unless otherwise noted, 'understanding-why' and 'understanding' are hereafter umbrella terms for the grasping of information that would answer an explanation-seeking question of any form. As a result, the following examples are relevant:

(1) Darwin understands why species are adapted to their environments;

(2) Darwin understands the adaptation of species to their environments; and

(3) Darwin understands how species evolve.

However, examples such as the following are irrelevant:

(4) Darwin understands Greek; or

(5) Darwin understands how to stuff animals.

Lipton's second assumption is that 'it is more natural to identify understanding with the cognitive benefits that an explanation provides rather than with the explanation itself' ([2009], p. 43). Since this figures prominently below, call this *Lipton's Assumption*. More precisely,

> *Lipton's Assumption:* If knowing that *b* constitutes understanding of *p*, then a correct explanation *e* of *p* provides knowledge that *b*.

Four clarifications are in order. First, *b* is the content of Lipton's cognitive benefits, which he characterizes in terms of 'four kinds of knowledge: of causes, of necessity, of

possibility, and of unification' (Lipton [2009], p. 43). Thus, in different contexts, *b* may be equivalent to one of the following:

(i) *the phenomenon cited by e causes p;*

(ii) *e entails that p is necessary;*

(iii) *e entails that p is possible;* or

(iv) *e unifies p.*

Second, it is important that Lipton's Assumption take understanding to be provided by *correct* explanations. For instance, conspiracy theories and the machinations of invisible demons can well nigh explain anything, but presumably they provide no bona fide understanding because such explanations are so wildly incorrect. For the purposes of this paper, we assume that correct explanations must be approximately true.

Different readers might construe explanatory correctness in terms of (strict) truth, empirical adequacy, justification, fit with background belief, etc. In principle, adopting these different standards of explanatory correctness could be assimilated to the arguments that follow, though we stress that this would require reinterpreting Lipton's cognitive benefits.

For example, in Lipton's framework, correct causal explanations provide knowledge that *A causes B*. However, since knowledge entails truth, a constructive empiricist might interpret correct causal explanations as providing either: (a) knowledge that *the hypothesis that A causes B is empirically adequate,* or (b) empirically justified beliefs that *A causes B*. Whether (a) or (b) better captures the empiricist's notion of causal understanding, and the degree to which constructive empiricism is compatible with explanatory idealism, are issues for another paper. Furthermore, given Lipton's ([2004],

pp. 184-206]) commitment to realism, our interpreting explanatory correctness as approximate truth enables a tidier comparison with his views.

Third, explanatory correctness should be contrasted with explanatory goodness. An incorrect explanation might be simple, powerful, fruitful, and accord with all of the evidence, and hence be good. Below, we discuss explanatory goodness' role in understanding.

Fourth, we follow Lipton in treating explanations as propositional in character. Indeed, the leading alternative—that explanations are things in the world, such as causes—sits uncomfortably with talk of correct and incorrect explanations, as it is odd to speak of correct and incorrect causes. Of course, propositions *about* causes are perfectly acceptable on the view endorsed here.

As should already be obvious, Lipton's Assumption gives explanation a privileged role. We return to this point below. For now, let us observe how this assumption still sits comfortably with his main thesis—that there can be understanding without explanation:

> The switch from identifying understanding with explanation to identifying it with some of the cognitive benefits of an explanation…makes this essay possible. For by distinguishing explanations from the understanding they provide, we make room for the possibility that understanding may also arise in other ways (Lipton [2009], p. 44).

More precisely:

> *Understanding Without Explanation* (UWE)*:* For some instances in which knowing that *b* constitutes understanding of *p*, it is possible for something other than a correct explanation to provide knowledge that *b*.

Thus, just as burning lumber is not the only way to provide heat, so explanation is not the only way to provide understanding. Specifically, Lipton presents the following as examples of understanding without explanation:

- Visual models and manipulations can provide tacit knowledge of causes (44-46)
- Non-explanatory deductive inferences can provide knowledge of necessity (46-49).
- Incorrect or 'merely potential' explanations can provide knowledge of possibilities (49-52).
- Non-explanatory analogies provide tacit knowledge of unification (52-54).

We will examine the details of these examples below. For now, it suffices to observe that they fit UWE's template: in all four cases, we can imagine someone not having a correct explanation of the phenomenon, yet still having causal, unifying, or modal knowledge.

2. Clarifying Lipton's framework

To render these ideas more precise, the sense in which something 'provides' understanding must be unpacked. Unfortunately, Lipton offers few details on this front. We propose the following as a friendly articulation of Lipton's Assumption:

> (6) If knowing that *b* constitutes understanding of *p*, then there exists a correct explanation *e* of *p* such that knowing that *e explains p* entails knowing that *b*.

On this view, explanations provide understanding by being known. This accords with Lipton's four forms of understanding. Clearly, if *e* is a correct causal or unifying

explanation, then if *S* knows that *e explains p*, *S* has knowledge of causes/unification. Furthermore, since both causal and unifying explanations frequently entail propositions about what is necessary or possible, we assume that understanders are deductively competent enough to infer modal knowledge from their explanatory knowledge.

Correspondingly, UWE can be rendered more precise along analogous lines:

(7) It is possible that, for all correct explanations *e* of *p,* a person *S* does not know that *e explains p*, yet *S* still understands *p*.

This sits comfortably with Lipton's examples. A person can achieve the relevant kinds of knowledge by thinking through visual models, manipulations, deductive inferences, merely potential explanations, or analogies without knowing a correct explanation.

Our gloss of Lipton's Assumption (6) might raise two possible worries. First, *knowledge* of correct explanations may be too demanding of a condition. Perhaps merely having true beliefs about a correct explanation, or simply being able to explain correctly, suffices to have understanding with explanation.

In reply, the most debated aspect of understanding's status as knowledge concerns its immunity to certain forms of 'environmental' epistemic luck. Since this epistemological issue need not concern us here, if proponents of this objection are correct (Kvanvig [2003], [2009a]; Pritchard [2010]), then we consider 'knowledge' to be imprecise shorthand for 'true beliefs that are immune to all epistemic luck save environmental luck.' In this case, understanding is still very similar to knowledge, and coincides with it in many instances. However, if replies to this objection are sound (Grimm [2006]; Khalifa [2011]), then even these qualifications are otiose.

Second, knowledge of correct explanations might be insufficient for understanding (de Regt [2009a], [2009b]; Grimm [2010]; Pritchard [2010]). For instance, students often come to know correct explanations through rote memorization without understanding. However, this overlooks the fact that understanding comes in degrees, such that these students have *some* understanding, even if *more* understanding is in the offing. On the Liptonian framework we are adopting, this greater understanding amounts to knowledge of, e.g. further causal or theoretical details, more and better explanations, etc.

3. Explanatory idealism

While Lipton's position is consistent, it raises a question. Given that he countenances understanding without explanation, why does his Assumption also accord explanation a privileged role in understanding? Explanatory idealism answers this question.

Roughly stated, if explanation is the ideal of understanding, i.e. if we measure our understanding by how well it approximates the cognitive benefits provided by knowing a correct and good explanation[4], then Lipton's Assumption would be a natural consequence. We would also expect the non-explanatory modes of understanding that animate Lipton's examples to stand in a different relationship to understanding than explanation does in Lipton's Assumption (§3.1).

Additionally, if explanatory idealism is true, Lipton's examples have specific consequences he does not explore. Section 3.2 presents the general contours of these consequences; the balance of the essay details how those consequences play out in Lipton's specific examples of understanding without explanation.

3.1. Misunderstanding without explanation

If explanatory knowledge is the ideal of understanding, then we should expect a certain asymmetry between the cognitive benefits provided by explanatory knowledge versus those provided by Lipton's non-explanatory modes of understanding. Specifically, since an ideal sets a standard, then (trivially) it also always meets that standard. By contrast, everything else meets that standard with varying degrees of success.

For instance, morally ideal people are moral by definition; the rest of us are, at best, contingently moral. Similarly, Lipton's Assumption entails that knowledge of correct explanations provides understanding by definition. Lipton anticipates this aspect of explanatory idealism when he describes his position as one that 'lets explanations set the standard for what kind of knowledge counts as understanding' (Lipton [2009], p. 54])[5].

However, Lipton does not develop another side of this point. Just like morally imperfect people meet a moral standard contingently, non-explanatory practices provide understanding contingently, i.e.

(8) It is possible that *S* does not understand why *p*, but has:

(a) A merely potential explanation of *p*,

(b) A non-explanatory deductive inference that concludes with *p*,

(c) An analogy involving *p*,

(d) A visual model of *p*, or

(e) A manipulation to which *p* refers,

Consider merely potential explanations. These are propositions that would, if true, explain a phenomenon. By contrast, actual explanations are true potential explanations,

i.e. correct explanations. Restricting Lipton's Assumption to *correct* explanations was precisely because some potential explanations do not provide understanding. To repeat, these are outlandish potential explanations such as conspiracy theories and fairy tales.

However, even when potential explanations are not outlandish, they may not provide understanding. For instance, if a doctor misdiagnoses the cause of a patient's symptoms, then, regardless of how reasonable her diagnosis, she misunderstands why the patient has the symptoms she does.

Similarly, while the following deductive inference is sound, it provides no understanding of why parity is conserved in strong interactions:

(9) *Sound Argument Providing No Understanding*

(a) Either parity is conserved in strong interactions or unicorns exist.

(b) Unicorns do not exist.

(c) ∴ Parity is conserved in strong interactions.

Likewise, many analogies fail to provide understanding, e.g.

(10) FOUR: NUMBER :: D: LETTER.

This analogy is perfectly good, but it does not tell us why four is a number, D is a letter, or why they are similar. Similarly, we do not understand why people are happy as a result of knowing that ☺ is a visual model of a happy person.

In many of these cases, *some* kind of understanding may still be in play, e.g. understanding *that* parity is conserved or *that* four is similar to D. But this does not amount to understanding *why* these facts are so. As already stated, only the latter concerns us here. If we recall that procedural understanding is not our concern, parallel

considerations apply to manipulations. For example, many people understand *how* to ride bicycles, but far fewer understand the physics that make bicycling possible.

In these examples, merely potential explanations, deduction, analogy, visual models, and manipulation each provide cognitive benefits, but these benefits should *not* be identified with understanding-why. By contrast, Lipton's Assumption expresses the exact opposite when it comes to knowledge of an explanation—its benefits *should* be identified with understanding-why. Thus, just as the morally ideal person sets and meets the moral standard by definition, knowledge of a correct explanation sets and meets the standard of understanding by definition. Just as mere mortals are not guaranteed to be moral, so too these non-explanatory practices are not guaranteed to provide understanding. As a result, explanatory idealism provides a plausible rationale as to why there is a special place for explanatory knowledge in understanding even if it is not the only means of achieving understanding.

3.2. How to interpret Lipton's examples

Explanatory idealism also suggests bolder departures from Lipton's Assumption. Ideals have a way of exceeding the actual. For instance, the ideal society remains forever elusive. (Explanatory idealism is a bit less utopian in this regard: many correct explanations are readily available.) While even our best current explanations can be improved upon, examples of non-explanatory understanding more starkly highlight the degree to which we fall short of our explanatory ideal. This point can be used to provide novel interpretations of Lipton's examples of understanding without explanation.

The general idea is this: wherever there is understanding without explanation, there is always a knowable explanation of the same phenomenon that would provide greater understanding than its non-explanatory counterpart. Call this the *Superior Explanation Thesis* (SET). More precisely:

(SET) For any non-explanatory way *w* to understand *p*, there exists a correct and reasonably good explanation *e* such that the understanding of *p* provided by *w* is a proper subset of the understanding of *p* provided by knowing that *e*.

Just as actual societies are deficient when compared to an ideal society, SET claims that non-explanatory understanding is deficient when compared to certain good and correct explanations.

Six aspects of SET deserve clarification. First, SET and Lipton's Assumption are not a package deal. In particular, one could hold Lipton's Assumption while claiming that an explanation exists that provides *at least as much understanding* as its non-explanatory counterpart, or that certain forms of understanding without explanation have no explanatory counterparts because correct explanations are unknowable. Moreover, Lipton does not even say this much (though we suspect it is a consequence of his view.) In this regard, our position is stronger than Lipton's.

Second, we follow Lipton in countenancing both causal and unificationist explanations. While our position would be unaffected by adopting a more pluralistic account of explanation, it is unnecessary for our argument.

Third, by 'reasonably good' explanations, we mean explanations good enough to be recognized by disinterested readers as neither borderline dreadful nor superhumanly demanding. All of our discussions below rest on such examples.

Fourth, we 'measure' degrees of understanding using the following:

(11) *General Argumentative Strategy*

(a) For each of Lipton's examples of understanding without explanation, there is a non-explanatory way *w* of understanding *p* such that if *S* understands *p* via *w*, then there is some cognitive benefit *b** such that *b** satisfies Lipton's Assumption and *S* does not know that *b**.

(b) There exists a correct and reasonably good explanation *e* such that if *S** knows that *e*, then *S** knows that *b*,* and *S** also knows all cognitive benefits *b* that comprise *S's* understanding of *p* via *w*.

(c) Thus, per Lipton's Assumption, *b** is something about *p* that *S** understands but that *S* does not.

(d) So a correct and reasonably good explanation provides greater understanding of *p* than the understanding provided by the non-explanatory way of understanding *p*.

Fifth, the conclusions of these arguments (11.d) provide some evidence for SET. We will argue that a wide variety of examples of understanding without explanation yield instances of (11.d), lending SET further plausibility. This evidence is not intended to be conclusive. Rather, it only shows that explanatory idealism is defensible, and shifts the burden of proof towards those who would draw stronger conclusions about the dispensability of explanation to understanding from examples akin to Lipton's.

Sixth, we will argue that the cognitive benefits provided by explanations are more varied than the four forms of knowledge that Lipton cites. In particular, we will include true beliefs about the failures of rival explanations, knowledge of inferential structure, and semantic knowledge as cognitive benefits to be considered alongside knowledge of causes, unification, necessity, and possibility.

4. Examples of modal understanding

We will first apply our strategy (11) to Lipton's claim that 'we can gain actual understanding from merely potential explanation' (Lipton [2009], p. 49). On this line of argument, not only actual explanations furnish knowledge of what is possible, as 'a merely possible explanation may also give information about the modal status of an explanation' (Lipton [2009], p. 51). Thus, Lipton is claiming that we sometimes understand ways a phenomenon could have been without knowing its actual cause or the unifying pattern actually underlying it.

Lipton provides two clear rationales for how modal understanding without actual explanation is possible[6]. First, potential explanations 'may show a degree of contingency in the actual explanation' (Lipton [2009], p. 51). For instance, suppose that a firm hires Jones because he had extensive prior experience in the industry. Moreover, the contingency of his hiring is highlighted by the fact that his other credentials were fairly nondescript, such that he would not have been hired had he lacked this experience. Thus, the falsehood of other, potential explanations of Jones' hiring highlights its contingency. Suppose that an agent knows that if other hiring criteria (e.g. education) had been the deciding factor, Jones would not have been hired, but does not know of Jones' superior

experience. In this case, the agent can be said to possess some understanding, although she lacks knowledge of a correct explanation of Jones' hiring.

Second, potential explanations 'may show necessity… by revealing fail-safe overdetermination,' as the following illustrates:

> Suppose that a boxing match between Able and Baker is rigged so that Baker—though in fact the far better boxer—would take a dive in the tenth round. Knowing this helps us to understand why Able won, even if as a matter of fact Able floors Baker with a lucky uppercut in the fifth (Lipton [2009], p. 51).

Suppose that someone knows that if the fight had lasted until the tenth, Baker would have taken the dive, but does not know that Able's fifth round knockout actually caused his victory. Because this person possesses knowledge of salient counterfactual scenarios, she understands without knowing an actual explanation.

As mentioned above, we accept that *some* understanding exists in these cases. We only aim to show that a correct explanation would provide *greater* understanding. To that end, these two examples both presuppose that we gain understanding when we know what would happen had things turned out differently. Specifically, had Jones had fewer years of industry experience, he would not have been hired; had Able not landed the lucky uppercut in the fifth, he still would have won the boxing match.

But it is widely held that information about these kinds of counterfactuals is necessary for explanations. For instance, Woodward ([2003], p. 221) argues that explanations must answer 'what-if-things-had-been-different' questions, or w-questions for short. If correct, good explanations cannot be restricted to facts in the actual world,

but must already intimate information about possibilities. We can now implement our strategy (11):

(12) *The Strategy Applied to Modal Examples*

(a) For Lipton's examples of understanding possibilities without explanation, there is a potential explanation of *p* such that if *S* understands *p* via this potential explanation, then *S* does not know the actual explanation of *p*, i.e. that which causes/unifies *p*.

(b) There exists a correct and reasonably good explanation *e* such that if *S** knows that *e*, then *S** knows the actual explanation of *p,* and *S** also knows all of the answers to the w-questions that comprise *S's* understanding of *p* via a potential explanation.

(c) Thus, per Lipton's Assumption, an actual explanation is something about *p* that *S** understands but that *S* does not.

(d) So a correct and reasonably good explanation provides greater understanding of *p* than the understanding provided by the non-explanatory way of understanding *p*.

Following Lipton, we take (12.a) to be uncontroversial, so the real question is whether explanations must answer w-questions, as Woodward suggests. Suppose that Charlie knows that Jones was hired *and* had superior experience. However, Charlie has no knowledge of counterfactual scenarios related to Jones' hiring. For instance, Charlie has no idea about what would have happened had Jones lacked experience, or perhaps falsely thinks that Jones would have gotten the job without this experience. It seems doubtful that Charlie is in any position to know that Jones was hired *because* he had superior

experience. Thus, a correct explanation of *p* requires correct answers to w-questions about *p*. In other words, knowledge of a correct explanation provides knowledge of possibilities.

One might argue that (12.b) does not mean that explanatory knowledge entails knowledge of these specific possibilities. While it is certainly true that not all good and correct explanations answer every w-question, SET only requires *one* correct explanation to answer the same w-questions as a merely potential explanation. Furthermore, since the ability to answer more w-questions tracks with explanatory goodness, explanations that function as ideals answer more w-questions[7].

Furthermore, the conclusion (12.d) is intuitive. For instance, suppose that someone knows that Jones' education alone would not have resulted in his hiring, but does not know what actually prompted his hiring. When asked why Jones was hired, the answer will have to be indirect or oblique, e.g., 'Not because of his education.' By contrast, a person with a correct explanation will be able to say that Jones was hired not because of his education *and* because of his experience. Intuitively, the latter person better understands Jones' hiring.

Similarly, a person who only knew that Baker *could have* lost the match because he would have taken a dive in the tenth, but does not know that Baker was *actually* knocked out fair and square in the fifth, doesn't understand why Baker lost the match as well as someone who also knows about the actual course of events. Thus, in both examples, either a correct explanation already provides the counterfactual information, or else far less understanding is to be had.

Finally, this argument can work for any example involving knowledge of possibilities qua understanding without explanation. Our interpretation of Lipton's Assumption implies that knowledge of possibilities only constitutes understanding if the possibilities in question are consequences of propositions about causes or unification. But since these explanations are clearly stronger than their consequences concerning possibilities (e.g. *if e then p is possible*), this means that understanding via actual explanations will always feature claims that cannot be replicated by merely potential explanations.

5. Critical information and the Galileo example

Our strategy also applies to another of Lipton's examples, concerning Galileo's demonstration of why gravitational acceleration is independent of mass. Galileo supposes the contrary, and then considers what would happen if a heavier mass m_1 were attached to a lighter mass m_2. Since, *ex hypothesi*, m_2 falls at a lesser speed than m_1, the two masses should fall more slowly than m_1 alone. However, when considered as one mass, $m_1 + m_2$ is heavier and should thus fall faster than m_1 alone. But since the same thing cannot fall both more quickly and more slowly than m_1, acceleration must be independent of mass (Galilei [1914], pp. 62-63; Lipton [2009], p. 47).

Lipton ([2009], p. 48) contends that Galileo's argument is not an explanation, because:

> Rather than saying directly why acceleration must be independent of mass, the argument works by showing that the contrary assumption would entail a contradiction.

Granting Lipton these points, then, despite lacking an explanation, Galileo knows that the conclusion of his argument is necessarily true. Knowledge of necessity is a cognitive benefit provided by an explanation. So, per Lipton's Assumption, Galileo understands why acceleration must be independent of mass. Thus, we have a genuine case of understanding without explanation.

In making the case for explanatory idealism, it's important to note that Galileo is criticizing a potential explanation—a point Lipton does not acknowledge. Specifically, Galileo is assessing whether or not acceleration *depends* on mass, which is tantamount to considering whether or not an object's being of a certain mass *explains* how it accelerates. So, Galileo treats mass as potentially explaining acceleration. If the mass explanation were correct, then m_1*'s* mass would explain its acceleration. But Galileo then considers what would happen if we attached m_2 to m_1, and his *reductio* is a criticism of the mass explanation; it shows that differences in mass fail to explain differences in speed.

Let *critical information* be true beliefs that potential explanations are incorrect. Then knowing an explanation requires critical information. Specifically, true beliefs falling short of knowledge are generally thought to be lucky, and critical information helps to mitigate that luck. For instance, if someone has no way of ruling out some other underlying cause of a patient's symptoms, at best he has luckily guessed the cause of those symptoms, but he does not know. Similarly, if someone knows the correct explanation of acceleration (e.g. that acceleration is a result of the net forces acting on an object), then she should be able to rule out rival explanations, such as the mass hypothesis, or else she merely has guessed luckily. So if an explanation provides

someone with explanatory knowledge, she also possesses critical information. Presumably, better explanations rule out more rival hypotheses, and thus convey a lot of critical information.

On this interpretation, Galileo's demonstration provides critical information about how objects accelerate. Implementing the strategy from above, we get:

(13) *The Strategy Applied to Galileo*

(a) For Lipton's example of Galileo's understanding without explanation, there is a non-explanatory *reductio* about how objects accelerate such that if Galileo understands how objects accelerate via this *reductio*, Galileo does not know the actual cause of objects' acceleration.

(b) There exists a correct and reasonably good explanation *e* such that if *S** knows that *e*, then *S** knows the actual cause of objects' acceleration, and *S** also knows all the critical information that comprises Galileo's understanding of how objects accelerate via the *reductio*.

(c) Thus, per Lipton's Assumption, knowledge of the actual cause of how objects accelerate is something about how objects accelerate that *S** understands but that Galileo does not.

(d) So a correct and reasonably good explanation provides greater understanding of how objects accelerate than the understanding provided by the non-explanatory way of understanding how objects accelerate.

Let us anticipate two objections to this argument. First, one might complain that we have pulled a bait and switch: Lipton claims that knowledge of necessity is a form of

understanding, and we have changed the topic by rendering critical information the relevant form of understanding.

It's unclear that this is an objection at all. After all, Galileo's reasoning is a refutation of Aristotelian explanations of motion. Moreover, Galileo's mastery of critical information appears more fundamental to his understanding than his knowledge of necessity: consider that even if Galileo failed to show that the mass explanation is *necessarily* false, he still would have offered an especially strong criticism.

The shift in emphasis from knowledge of necessity to critical information gains further credence when considering recent arguments denying that Galileo's reasoning amounts to a *reductio* (Gendler [1998]; Schrenk [2004]; Vickers [forthcoming]). In effect, these arguments suggest that Galileo's reasoning only shows the implausibility, but not the impossibility, of the mass explanation. Such an interpretation accords well with our privileging of criticism over necessity.

A second objection is that appealing to critical information only works for understanding obtained via *reductio*, but not for Lipton's more general claim that knowledge of necessity via deduction can furnish examples of understanding without explanation. Absent a compelling example, this concern is difficult to address. On the one hand, some deductive inferences are explanatory, so these cannot be used as examples of understanding without explanation. On the other hand, some knowledge of necessity, e.g. knowledge of tautologies, seems to fall short of understanding. While navigating these two extremes is *prima facie* possible, it behooves the objector to offer an example.

6. Unification via tacit analogy

Thus far, we have seen that critical information and knowledge of possibilities are cognitive benefits that either explanations or their alternatives can provide. However, these are only 'oblique' forms of understanding. By contrast, only explanations have provided 'direct' understanding, i.e. they provide the information we associate with direct answers to why-questions. Consequently, this aspect of explanations has thus far proven essential to highlighting the greater understanding provided by explanations.

However, Lipton has other examples of understanding without explanation wherein this 'direct understanding' is on full display. In these examples, Lipton pries apart understanding from explanation by appealing to tacit knowledge[8]. For instance, he writes:

> Kuhn [offers] an account of … the scientists' ability to select problems and generate and evaluate solutions. And these abilities correspond also to a knowledge that goes beyond the explicit content of the theory. The exemplars provide knowledge of how different phenomena fit together. The unarticulated similarity relations that the exemplars support provide a taxonomy that gives information about the structure of the world. They thus have the effect of unifying the phenomena, and they do this by analogy, not by explanation (Lipton [2009], p. 53).

This suggests that unification without explanation works as follows:

(14) *Unification by Tacit Analogy*

(a) *S* knows the following exemplar: *B* is a solution to problem *A*.

(b) *S* is able to identify that *C* is a problem analogous to *A*, and that C has a solution, D, analogous to B.

(c) So S knows that A, B, C, and D are unified.

Since an explanation could provide knowledge of unification, Lipton's Assumption treats this as a form of understanding. It also appears to forgo explanatory knowledge, and is presumably tacit because the person cannot say how *C* and *D* are analogous to *A* and *B*, respectively. Thus, as before, we have understanding without explanation.

Before using our strategy (11) to show that this example poses no threat to explanatory idealism (§6.2), explanatory unification must be carefully distinguished from 'Kuhnian' unification by tacit analogy (§6.1).

6.1. Explanatory unification

Lipton is not explicit about how explanations provide knowledge of unification. However, applying our strategy requires a clearer account of how unification via explanation operates. To that end, we provide a brief sketch of Philip Kitcher's ([1989]) influential 'unificationist' account of explanation. Essential for our purposes is Kitcher's ([1989], p. 430) claim that 'Science supplies us with explanations whose worth cannot be appreciated by considering them one-by-one but only by seeing how they form part of a systematic picture of the order of nature.'

More precisely, Kitcher holds that we gain knowledge of an explanation *e unifies p* (i.e. we 'see' this 'systematic picture'), when we know that that the derivation of *p* from *e* is an instance of a broader derivation pattern (or schema) that can be applied to phenomena other than *p*. This suggests the following explanatory alternative to (14):

(15) *Explanatory unification*

(a) S knows that B explains A and that D explains C[9].

(b) S knows that B's explaining A and D's explaining C are both instances of a more general pattern or schema G.

(c) So S knows that G unifies A, B, C, and D

Schemas such as *G* contain information such as the following:

External Pathway Explanation Schema:

Explanation target:

Why does a **cell** become defective in a **function**?

Explanation pattern:

The **cell** is destructively affected by **external agents,** such as bacteria, viruses, or autoimmune cells.

These **external agents** operate by means of **pathways** that enable them to invade and disrupt the **cell**.

So the **cell** becomes defective and cannot carry out its **function** (Thagard [2003], pp. 244-245)[10].

Here, the boldfaced letters are variables that are filled in by different values. The unifying power of the schema is directly proportional to the variety of values that these variables can assume while still yielding correct instances of the explanation schema.

Importantly, on Kitcher's view, an explanation's providing knowledge of unification *entails* that it does so by way of fitting our explananda into a more general pattern or schema. Since this idea is paramount in what follows, it is worth stressing that it animates most other accounts of explanatory unification. For instance, Friedman ([1974], p 19) writes that when we explain via unification, 'We replace one phenomenon

with a *more comprehensive* phenomenon, and thereby effect a reduction in the total number of accepted phenomena. We thus genuinely increase our understanding of the world.' Bartelborth's ([2002], p. 91) more recent unificationism claims that 'explanations promote our understanding of the world [by] embedding … our observations, events, and other facts into more general patterns that bring our different observations together in a coherent world view.' Indeed, even advocates of non-unificationist accounts of explanation often favor explanations involving broader or more invariant explanatory generalizations, e.g. (Woodward [2003]).

To recap, Lipton was not explicit about how explanations provide knowledge of unification. According to many sustained reflections on this topic, explanations provide this kind of understanding via general schemas or patterns of reasoning. Hence we assume that if explanations provide knowledge of unification, then they provide knowledge of general schemas. Should other accounts of explanatory unification bypass knowledge of general schemas, we would have to evaluate those accounts, and their ramifications for explanatory idealism.

6.2. Unification and explanatory idealism

Thus, we now have clear accounts of how both tacit analogy (14) and explanation (15) provide unification. The crucial difference is that only the latter provides knowledge of a general schema *G*. Explanatory idealists can use this feature of explanatory unification to pose a dilemma to their critics. Either the tacit analogies in question contain the information encoded in schemas or they do not. If they do, then Lipton's example is not a genuine case of understanding *without explanation*, as this would amount to the curious claim that understanding without explanation amounts to tacit knowledge of an

explanatory schema. (Indeed, unificationists such as Kitcher, Friedman, and Bartelborth do not deny that these schemas can be tacit.) Alternatively, if the tacit analogies do not contain the schemas' information, then our strategy readily applies, for having knowledge of these schemas is cognitively beneficial. In either case, explanatory idealism is unthreatened.

Focusing on the second horn of our dilemma, reasoning schemas provide knowledge of inferential structure. In tacit unification, we only know that similar problems admit of similar solutions, but we cannot provide reasons for why this is so. General schemas provide those reasons, and inferential knowledge of this sort is cognitively beneficial. This suggests the following:

(16) *The Strategy Applied to Tacit Unification*

(a) For each of Lipton's examples of understanding of unification without explanation, there is a tacit, analogical way of understanding *p* such that if *S* understands *p* via this tacit analogy, *S* does not know that specific explanations can be inferred from a more general schema.

(b) There exists a correct and reasonably good explanation *e* such that if *S** knows that *e*, then *S** knows that specific explanations can be inferred from a general schema, and *S** also has the knowledge that comprises *S's* understanding of *p* via tacit analogy, viz. that *p* is unified with other propositions.

(c) Thus, per Lipton's Assumption, knowledge that specific explanations can be inferred from a general schema is something about *p* that *S** understands but that *S* does not.

(d) So a correct and reasonably good explanation provides greater understanding of *p* than the understanding provided by the non-explanatory way of understanding *p*.

One might resist (16.b) by claiming that knowledge of an explanation and knowledge of a schema ought to be kept distinct, but then we are owed some account of how explanations provide knowledge of unification. Moreover, generalizations of some sort (laws, invariant causal generalizations, models that can accommodate other phenomena, etc.) clearly are essential to many explanations, and these allow us to explain multiple phenomena using the same conceptual framework. Thus regardless of whether *schemas* are the best way to capture this intuition, explanations provide knowledge of unification by fitting an explanandum into a more general framework, and knowledge of this framework differentiates explanatory unification from Lipton's Kuhnian alternative.

In addition to our strategy, there are other reasons to think that explanatory unification is superior to tacit unification. Duncan Pritchard ([2010], p. 82) has argued that 'there is a strong *prima facie* case for thinking that all understanding involves a cognitive achievement,' and explanatory unification is a greater achievement than its tacit counterpart.

Specifically, explanatory unification involves far more demanding criteria of adequacy than the tacit criteria typical of analogies; and overcoming these demands signals a greater achievement. To get a sense of this, Kitcher's account of explanatory unification requires all reasoning schemas to be deductive, all of the premises of any instantiation of such a schema to be acceptable relative to the corpus of background knowledge, and the generation of as many acceptable conclusions from as few and as

stringent schemas as possible. It is hard to find any correlates to these constraints in an analogical route to unification. Moreover, these features of explanatory unification track with other epistemic desiderata that we might associate with better understanding, e.g. simplicity, power, precision, testability, and availability to criticism.

Thus, explanations provide unification through knowledge of general patterns. Lipton's example of unification via tacit analogy then suffers a dilemma. Either tacit analogies rest on these patterns, and hence do not provide understanding without explanation, or these analogies proceed without those patterns, in which case explanations are superior to analogies precisely because they provide knowledge of these general patterns. Additionally, unification via explanation is a greater cognitive achievement than unification via tacit analogy. So insofar as we place greater value on greater achievements, explanatory unification also surpasses analogical unification along this dimension. All of this bodes well for explanatory idealism.

7. Tacit understanding of causes

Finally, Lipton uses tacit knowledge to provide two further examples of understanding without explanation. First, Lipton describes a person who gains tacit understanding of the causes of retrograde motion through a visual model of the solar system (an orrery):

> These visual devices convey causal information without recourse to an explanation. And people who gain understanding in this way may not be left in a position to formulate an explanation that captures the same information. Yet their understanding is real (Lipton [2009], p. 45).

In addition to visual models, Lipton takes manipulation as another site of tacit understanding without explanation:

> […] a scientist may gain a sophisticated understanding of the behavior of a complicated piece of machinery by becoming an expert at using it, and that understanding consists in part in the acquisition of causal information that the scientist may be in no position to articulate (Lipton [2009], p. 45).

Clearly, in both examples, the inquirers possess causal knowledge, and per Lipton's Assumption, they thereby possess understanding. So, why exactly do they lack an explanation? According to Lipton, this is because explanation

> requires that the information be given an explicit representation. In short, there is such a thing as tacit understanding, but not tacit explanation, and this provides the space we are looking for, where there can be understanding without explanation (Lipton [2009], p. 45).

Thus, Lipton's argument runs as follows:

(17) *Lipton's Tacit Understanding Argument*

- (a) When explanations provide causal knowledge, that knowledge is explicit.
- (b) Some visual models and manipulations provide causal knowledge that is not explicit.
- (c) Causal knowledge is a kind of understanding.
- (d) So some visual models and manipulations provide understanding without being explanations.

As we have throughout, we grant that this provides legitimate cases of understanding without explanation, and only seek to show that an explanation provides superior understanding.

This requires a more precise account of explicit and tacit knowledge:

(18) If *S* explicitly knows that *A causes B,* then:

(a) *S* has the ability to reliably identify *A* as a cause of *B*, and

(b) *S* has the ability to communicate that *A causes B* verbally.

Tacit knowledge differs from explicit knowledge with respect to the second condition (18.b). In other words, tacit knowers lack the ability to communicate their knowledge verbally. While Lipton does not present many details about how he distinguishes explicit and tacit understanding, his few remarks are congenial to this gloss. For instance, he describes his tacit knowers as being 'inarticulate' and unable to 'say something' (Lipton [2009], pp. 45-46).

As before, explanatory idealists can inflict a dilemma upon their interlocutors. On the one hand, if knowing an explanation involves verbal abilities, then our strategy applies: explainers have semantic knowledge that their counterparts lack (§7.1). Alternatively, if knowing an explanation does not require verbal abilities, then tacit explanations exist, so these aren't examples of understanding without explanation (§7.2).

7.1. Semantic knowledge

Suppose that the key difference between the agents in Lipton's examples of tacit understanding and their explanatory counterparts is captured by the latter's ability to verbally express otherwise identical causal knowledge. Then, quite clearly, explainers know something that tacit knowers do not: namely that certain verbal expressions refer to certain causal structures. Call this *semantic knowledge*, viz.

(19) *The Strategy Applied to the Tacit Understanding Argument*

(a) For some of Lipton's examples of understanding without explanation, there is a tacit way of understanding *p* such that if *S* understands *p* tacitly, then *S* does not know that certain verbal expressions refer to certain causes of *p*.

(b) There exists a correct and reasonably good explanation *e* such that if *S** knows that *e*, then *S** knows that certain verbal expressions refer to certain causes of *p,* and *S** also knows the causes that comprise *S's* tacit understanding of *p*.

(c) Thus, per Lipton's Assumption, knowledge that certain verbal expressions refer to certain causes of *p* is something about *p* that *S** understands but that *S* does not.

(d) So a correct and reasonably good explanation provides greater understanding of *p* than the understanding provided by the non-explanatory way of understanding *p*.

We shall address three potential objections to this argument. First, one might object that (19.b) does insufficient justice to tacit knowledge, as tacit knowers often appear to have knowledge that is not easily captured in words. However, SET only requires that at least one explanation can furnish all of the understanding of its counterpart, and these examples are easy enough to construct, e.g.

> *First Example:* Using an orrery, Andrew tacitly understands why retrograde motion occurs; Belle can explain why retrograde motion occurs by using an orrery.

> *Second Example*: By manipulating a complex device, Dan tacitly understands why the device behaves as it does; Emily can explain the device's behavior by manipulating it.

We might think of Belle and Emily as capable of providing 'lecture-demonstrations,' i.e. they can do everything that Andrew and Dan can do, and they can provide running commentary on what they're doing. Hence, there is a straightforward recipe for establishing the second premise of our argument (19.b): we simply imagine tacit knowers with the ability to say what they know.

Second, *pace* (19.c), one might claim that semantic knowledge is not cognitively beneficial. Lipton certainly does not acknowledge it as such. However, it is unclear what resources he has for rejecting it. After all, Lipton's Assumption identifies understanding with the cognitive benefits provided by explanations; all of Lipton's examples of these benefits are kinds of knowledge; and Lipton also requires explanations to be explicit. The only contentious leap we can see here is that the explicitness of explanation might not entail semantic knowledge, so we provide a concrete but simple illustration of the Second Example—involving manipulation—to motivate this connection.

Suppose that Dan knows that he can start and stop an oscillator by flipping a switch, though perhaps he doesn't have words for 'oscillator' or 'switch.' Nevertheless, something like the following is within Dan's ken:

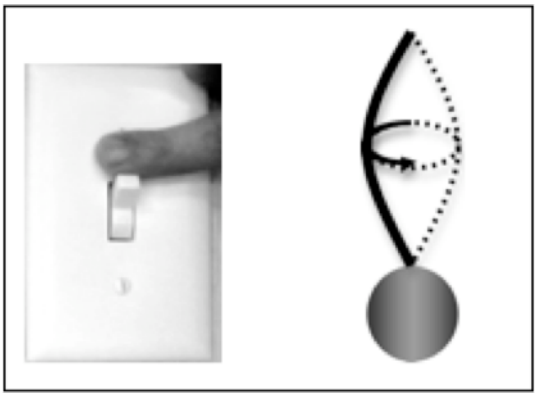

Figure 1: Tacit Representation of Causal Knowledge

Here, Dan needn't have a word for 'causes,' but clearly must have the concept *causes* in order to have causal knowledge. Emily can represent this knowledge as Dan does, plus she can express the following in English:

(20) Flipping the switch causes the object to oscillate.

It is hard to see how Emily could be a competent language-user in this context without having the semantic knowledge that *the English sentence 'flipping the switch causes objects to oscillate' refers to the fact that flipping the switch causes the object to oscillate*. Thus, both she and Dan have propositional knowledge of (20), but Emily also knows that a sentence refers to a causal fact[11]. Hence, if linguistic competence is a mark of explicit knowledge, then the latter implies semantic knowledge.

However, a final objection remains: one might grant that semantic knowledge is cognitively beneficial, but not in a way that promotes *understanding*. We offer two examples to suggest otherwise. First, throughout the history of science, physical interpretations of technical languages are often required in order to advance scientific understanding, and, when correct, such interpretations are instances of semantic knowledge. Second, in an example that is probably familiar to many readers, consider students who object to their grades on the grounds that their writing does not reflect the

depth of their understanding. These complaints frequently ring hollow precisely because verbal expression and understanding cannot be so neatly divorced. In other words, it is natural to think of semantic knowledge as a dimension of understanding.

7.2. Tacit explanations

However, suppose that semantic knowledge isn't cognitively beneficial. Explanatory idealists can still pursue the second horn of our dilemma: *pace* Lipton, explanations in the absence of verbal communication are possible. For instance, Dan's inability to speak or write a sentence like (20) doesn't preclude him from an *explanation* of oscillation. As a result, the tacit-explicit distinction should be handled with greater caution when trying to construct examples of understanding without explanation.

To see this, imagine a third person, Fred, with the same causal knowledge of the oscillator shared by Emily and Dan. However, Fred is mute, and expresses (20) in sign language. Fred lacks verbal abilities, but he's clearly offering an explanation[12]. In other words, our choice to use verbal symbols over other (e.g. gestural or pictorial) symbols is entirely a matter of convention, and no deep facts about explanation hang on this choice alone. For instance, Fred's expression in sign language translates into a readily identifiable explanation in English (20), yet it didn't undergo some magical transformation in the process of translation—it was already expressing an explanation in sign language[13].

This leaves far less room for non-explanatory tacit understanding. Indeed, any symbol system—verbal or otherwise—with interpretive conventions adequate for the purposes of communicating causal information appears expressively adequate for the purposes of explanation. For instance, suppose that Dan flips the switch to demonstrate

how the oscillator goes on and off with the intention of communicating the causal structures at work. If there is enough common knowledge about the meaning of Dan's gestures, then this becomes an impromptu sign language for conveying understanding to an audience.

Consequently, sympathizers of tacit understanding owe us some account of how Fred's explanation in a sign language steeped in tradition differs from Dan's allegedly non-explanatory understanding expressed in an improvised sign language. If no principled distinction exists, then Dan has *explained* oscillation using only (or mostly) gestures. Consequently, this is an example of understanding *with* explanation. Similarly, if we start granting Dan and his audience even simple conventions about ostension and basic verbal communication (e.g. 'If you do this, then that happens'), then it is hard to see why Fred's explaining in a sign language is much different than Dan's explaining in a language composed of words and gestures.

We have tried to make this argument independently of how one draws the tacit-explicit distinction, since little consensus on this distinction exists. If (18) is a fair gloss of the tacit-explicit distinction and there is no substantive difference between Fred and Dan's expressive capacities, then tacit explanations are possible. If (18) disagrees with one's preferred tacit-explicit distinction and there is still no substantive difference between Fred and Dan's expressive capacities, the challenge remains: why is this an example of understanding *without* explanation? Finally, if Fred has superior expressive capacities to Dan, then we can simply rehearse the arguments in Section 7.1: expressive capacities are cognitively beneficial, providing, *inter alia,* semantic knowledge. In any of

these scenarios, nothing about the tacit-explicit distinction undermines explanatory idealism.

8. Conclusion

Thus, while Lipton has provided compelling evidence that explanation is not necessary for understanding, thinking through his work provides important lessons about the importance of explanation to understanding. Where understanding is achieved without explanation, there is always room for explanatory improvement. This is because explanations are cohesive and communicable bundles of direct, critical, and modal information, while other forms of understanding seem to be more piecemeal in delivering these goods. Moreover, we have also seen that the central assumption animating Lipton's critiques of explanatory understanding sits comfortably with our idea that explanation is the ideal of understanding. Given that few have advocated for the importance of explanation to inquiry more eloquently and ably than Lipton, we hope he would not have disagreed with this conclusion.

Department of Philosophy

Middlebury College

Middlebury, VT 05753

USA

Center for Philosophy of Science

817 Cathedral of Learning

University of Pittsburgh

Pittsburgh, PA 15260

USA

kkhalifa@middlebury.edu

Acknowledgements

Much of this paper was written while on sabbatical from Middlebury College and as a Visiting Fellow at the University of Pittsburgh's Center for Philosophy of Science. I thank both institutions for their support. José Díez, Heather Douglas, Stephen Grimm, Jim Lennox, Bert Leuridan, P.D. Magnus, Elisabeth Nemeth, Laura Perini, Jonah Schupbach, Mark Tschaepe, Pete Vickers, and John Worrall all deserve thanks for their feedback on ideas in this paper. Finally, while I never had the opportunity to meet Peter Lipton, thinking through his work has profoundly influenced my own philosophical positions. I dedicate this paper to his memory.

References

Bartelborth, T. [2002]: 'Explanatory Unification', *Synthese*, **130**, pp. 91-107.

de Regt, H. W. [2009]: 'The Epistemic Value of Understanding', *Philosophy of Science*, **76**, pp. 585-597.

de Regt, H.W. and Dieks, D. [2005]: 'A Contextual Approach to Scientific Understanding', *Synthese*, **144**, pp. 137-170.

Elgin, C. [2007]: 'Understanding and the Facts', *Philosophical Studies*, **132**, pp. 33-42.

Friedman, M. [1974]: 'Explanation and Scientific Understanding', *Journal of Philosophy*, **71**, pp. 5-19.

Gendler, T. S. [1998]: 'Galileo and the Indispensability of Scientific Thought Experiment', *The British Journal for the philosophy of science*, **49**, pp. 397-424.

Grimm, S. R. [2006]: 'Is Understanding a Species of Knowledge?', *British Journal for the philosophy of science*, **57**, pp. 515-535.

Grimm, S. R. [2010]: 'The Goal of Understanding', *Studies in the history and philosophy of science*, **41**, pp. 337-344.

Khalifa, K. [2011]: 'Understanding, Knowledge, and Scientific Antirealism', *Grazer Philosophische Studien*, **83**, pp. 93-112.

Khalifa, K. [forthcoming]: 'Is Understanding Explanatory or Objectual?', *Synthese*, **doi 10.1007/s11229-011-9886-8**.

Longworth, G. [2008]: 'Linguistic Understanding and Knowledge', *Nous*, **42**, pp. 50-79.

Pettit, D. [2002]: 'Why Knowledge Is Unnecessary for Understanding Language', *Mind*, **111**, pp. 519-550.

Schrenk, M. A. [2004]: 'Galileo Vs. Aristotle on Free Falling Bodies', *Logical analysis and the history of philosophy*, **7**.

Schurz, G and Lambert, K. [1994]: 'Outline of a Theory of Scientific Understanding', *Synthese*, **101**, pp. 65-120.

Thagard, P. [2003]: 'Pathways to Biomedical Discovery', *Philosophy of Science*, **70**, pp. 235-254.

Trout, J. D. [2002]: 'Scientific Explanation and the Sense of Understanding', *Philosophy of Science*, **69**, pp. 212-233.

[1] For example, (de Regt [2009]; de Regt and Dieks [2005]; Elgin [2007]; Grimm [2006], [2010]; Khalifa [forthcoming]; Trout [2002]). Following these authors, our focus is

neither on understanding as a special form of social cognition (*Verstehen*), nor as fluency with a language, e.g. (Longworth [2008]; Pettit. [2002]).

[2] Of course, the words 'ideal' and 'idealism' carry many more connotations. Our uses are restricted to the narrow meanings presented in the text.

[3] More precisely, Lipton writes: 'For my present purposes, what is important about abilities is not the *sui generis* forms of understanding they provide [i.e. procedural understanding], but rather the conventional forms of understanding they support, as conventional as knowledge of causes and unification.'

[4] Here, we are not describing how people actually evaluate understanding. Like political ideals, other concerns (e.g. prudential considerations) can override the ideals of understanding.

[5] In the same passage, Lipton describes this approach as 'narrow.' That turns on our earlier observation that procedural understanding is not his concern (nor is it ours).

[6] Lipton lists several other examples under this heading, but their rationale is not as explicit as the ones discussed here. We submit that they can be assimilated to the strategy presented above, or else they are not genuine examples of understanding without explanation.

[7] Parallel objections and replies could be rehearsed in each of the following sections. We spare the reader the redundancy.

[8] Lipton does not reference any works on tacit knowledge, e.g. . It is unclear whether he intends anything precise.

[9] Many unificationists reduce the explanatory relationship to an inferential relationship, and often more narrowly a deductive relationship (Friedman [1974]; Schurz and Lambert

[1994]). Changing the word ‘explains’ to ‘entails’ would not affect the points about understanding we make below.

[10] Kitcher deploys schemas to depict the progressive unification of evolutionary biology and chemistry. Thagard’s example illustrates Kitcher’s idea more compactly.

[11] Here, we follow common conventions in the philosophy of language: sentences are strings of visual or audible symbols that express propositions; propositions are the abstract entities that carry the meaning of the sentences; facts are concrete things in the world, and unlike sentences or propositions, are not capable of bearing truth or falsity.

[12] Indeed, American Sign Language has an expression for ‘explain.’

[13] This is just to say that explanations are fundamentally propositional, rather than sentential creatures, in which case, their syntactical expressions are largely incidental.