

Normalcy and the Contents of Philosophical Judgements

Georgi Gardiner
Rutgers University

Abstract

Thought experiments as counterexamples are a familiar tool in philosophy. Frequently understanding a vignette seems to generate a challenge to a target theory. In this paper I explore the content of the judgements that we have in response to these vignettes. I first introduce several competing proposals for the content of our judgements, and explain why they are inadequate. I then advance an alternative view. I argue that when we hear vignettes we consider the normal instances of the vignette. If the normal instance of the vignette is a counter-instance, the vignette constitutes a challenge to the target theory. I argue that this proposal shows how responses to vignettes are a kind of ordinary, everyday judgement, and I explain how the proposal avoids the problems generated by competing views. Finally, I argue this ‘normalcy proposal’ most naturally accords with our understanding of the method.

Keywords

philosophical judgements, counterexamples, philosophical methodology, intuitions, Timothy Williamson, normalcy judgements.

1. Introduction

Thought experiments as counterexamples are a familiar method in philosophy. Amongst the most paradigmatic include Edmund Gettier’s use of vignettes to argue that knowledge is not justified true belief, Frank Jackson’s Mary vignette against physicalism about conscious experience, and the ‘murderer at the door’ vignette against the view that lying is always all-things-considered impermissible.¹ In typical instances of the ‘vignette as counterexample’ method, we take a theory or claim, such as:

(1) Knowledge is justified true belief.

or

(2) Necessarily, knowledge is justified true belief.

We then present a vignette, short story or anecdote challenges the target theory or claim. An example of a vignette that challenges claims (1) and (2) is,²

ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large

¹ These examples are from Gettier (1963), Jackson (1982), and Constant (1797) respectively.

² This vignette is found in Williamson (2007: 183).

sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf.

If the method is successful, then considering the vignette provides reason and motivation to abandon, revise, reduce confidence in, or reconsider the target theory or claim. The ubiquity and apparent efficacy of thought experiments as counterexamples in philosophy raises several questions, such as (how) does this method work, what kind of judgement do we have in response to these vignettes, and what is the epistemic status of these judgements?³ In this paper I focus on the content of the judgement we have in response to these cases. Since this judgement seems to both motivate and legitimate certain inferences—such as casting doubt on claims (1) and (2)—understanding the content of this judgement is important. And resolving the ‘content question’ illuminates other questions about this central method in philosophy.⁴

2. Two Initial Proposals

To see why this question is tricky, consider some proposals. Consider first the view that the content of the judgement we have in response to the vignette is simply the negation of the target claim. Call this the ‘direct proposal’ since it posits that upon hearing the vignette we directly grasp, unmediated by a bridge proposition, the falsity of the target claim.⁵ In the above example the target claim is:

(2) Necessarily, knowledge is justified true belief.

We can understand this claim as entailing claim (3):

(3) Necessarily, for any subject x and proposition p , x knows p iff x has a justified true belief in p .

To pin down a precise formulation of the target claim, we can formalise (3) as:

(4) $\Box \forall x \forall p (K(x, p) \equiv JTB(x, p))$

³ Note that judgements in response to vignettes are not, by any means, the only kind of philosophical judgement, and counterexamples are not the only role vignettes play in philosophical theorising. See Gardiner (2015) for discussion of various kinds of philosophical judgement.

⁴ Williamson (2007), Malmgren (2011), and Ichikawa and Jarvis (2009; 2013) advance views about how the content question bears on other questions. Williamson (2007) argues that his view allows these judgements to be continuous with other everyday judgements; Malmgren (2011) and Ichikawa and Jarvis (2009; 2013) argue that a virtue of their account is that it renders these judgements apriori. Whilst in this paper I bracket the question of whether the judgements are apriori or aposteriori, I argue that answering the content question helps vindicate these judgements, and so legitimates some uses of ‘thought experiments as counterexample’ in philosophy.

⁵ Sosa (2006) appears sympathetic to the direct proposal, since he suggests that considering a vignette can lead us directly to endorsing claims such as claim (5).

Call claim (4) the JTB view of knowledge (strictly speaking, it is an entailment of the JTB view⁶). The vignette challenges claim (4) by seeming to present us with a counter-instance—a case where someone possesses justified true belief that *p* without knowledge that *p*. Accordingly the direct proposal holds that the judgement has the following content:

(5) Someone could have a justified true belief that *p*, yet not know that *p*.

We can formalise this possibility as:

(6) $\Diamond \exists x \exists p (JTB(x, p) \ \& \ \neg K(x, p))$

The first proposal, then, is that when we hear a vignette, such as the Orlando vignette, we form a judgement with the content articulated by (6). This suggestion has several virtues. Firstly, (6) directly contradicts (4), and so this proposal explains why the judgement in response to these cases constitutes a problem for the target theory. Secondly, it captures the immediacy that the judgement often has: hearing a successful vignette can quickly produce a grasp of the falsity of the target claim.⁷ The challenge is immediately apparent. By contradicting the target claim, the proposed content (6) vindicates this.

There are problems, however, with the direct proposal. Firstly, the proposed judgement is not *about* the vignette presented. But it seems that our response to a vignette is about the vignette and includes content specific to the particular vignette. The possibility judgement (6) seems to arrive via, and be mediated by, some judgement that concerns details of the case: It seems as though our bridge judgement has something to do with Orlando, the justification of testimonial beliefs, whether an alternative book renders Orlando's belief true etc., and from this bridge judgement we infer (6).

Relatedly, the proposed content does not posit difference in content between the judgements we have in response to the Orlando case and an alternative Gettier case, such as:⁸

SHEEP. Driving past a field, Roddy looks at a sheep-shaped, white object in normal daylight across a field. As a naturalist with keen eyesight, Roddy is reliable at recognising sheep. His visual inputs accord with his background belief that there are sometimes sheep grazing in this area. Roddy accordingly forms the belief 'there is a sheep in the field'. As it happens, however, what Roddy saw is not a sheep but a sheep-shaped rock. Unbeknownst to him, behind the rock a sheep is grazing.

This vignette also challenges the target claim. As with the Orlando vignette, the case seems to exhibit the possibility of justified true belief without knowledge. But according to the direct proposal, the judgement in response to the sheep vignette is identical to the judgement in response to the Orlando vignette. This is implausible.⁹

⁶ Focusing on claim (3), formalised as (4), as the target claim is useful because most accounts that aim to understand knowledge as justified true belief will entail (3). It is important to note, however, that claim (3) concerns necessary and sufficient conditions, whereas an account of knowledge is typically a claim about constitution, identity, metaphysical dependence, reduction, conceptual equivalence, or linguistic synonymy. See also Williamson (2007: 183).

⁷ Or, less boldly, we immediately grasp the tension between the vignette and the target claim.

⁸ Pritchard (2012: 251), adapted from Chisholm (1977: 105).

⁹ Malmgren (2011: 284) independently makes a similar point.

Thirdly, the direct proposal does not provide much explanatory power. It does not generate resources to explain how the vignettes precipitate or vindicate judgement (6); it does not explain, in other words, why rejecting or revising the target theory can be the correct response to a vignette. The proposed content also does not explain why some instances of forming the judgement upon hearing a vignette are mistaken or illegitimate. Suppose we hear a vignette that is incoherent, has nothing to do with belief formation, or does not describe an instance of justified belief. In these cases, if a hearer were to judge in response ‘someone could have a justified true belief that *p*, yet not know that *p*’, i.e., claim (6), her response does not seem appropriate. It is only an appropriate conclusion when the vignette warrants the judgement. But this proposed content does not seem to offer resources to explain when the judgement is, and is not, appropriate.¹⁰

An alternative proposal posits that the judgement concerns the vignette itself more directly. The proposal holds that the content is something like,

- (7) Orlando has a justified true belief but does not know that he owns a book that once belonged to Virginia Woolf.

This proposal, unlike the ‘direct proposal’ above, captures that the operative judgement concerns the vignette. The judgement is about Orlando. And this proposal provides resources for explaining the difference between those judgements that are justified, proper, and appropriate responses, and those that are not. Roughly, the judgement is appropriate whenever the vignette describes the protagonist’s having a justified true belief that is not knowledge.

But there are weaknesses. Firstly, it is not clear how the judgement appropriately attributes properties to Orlando; it is common knowledge that Orlando does not exist. Theorists who discuss thought experiments typically realise the vignettes are fictitious, and so cannot judge (7) as a straightforward matter of fact. (Some vignettes deployed as counterexamples are, or purport to be, veridical. In these cases, content (7) seems more plausible. But these cases constitute the small minority of vignettes in philosophy, and the method is successful regardless of whether the vignettes purport to be true.) So if the judgement resembles (7), more must be said about how the content features a non-existent person, and whether or how the acknowledged fictionality of the vignette plays a role in the judgement.¹¹

¹⁰ One response here might be to argue that the response judgement (6) is appropriate only when it is well-grounded, based on sufficient relevant evidence, reliably formed etc. According to this response, judgement (6) sometimes is inappropriate, even though it is always true; whether it is apt depends on connections between the possibility judgement and the vignette. The response maintains, however, that the judgement’s content is simply (6). The success of this proposal demands discussion of how the possibility judgement might relate to the vignette. This essay can be seen as doing this work: I explain how the possibility judgement (6) can be legitimately arrived at via considering the vignette. And so although I do not think (6) is the correct proposal, the insights in this essay can be harnessed by advocates of this proposal, as a way to explain how and when (6) is a legitimate and reasonable response to considering a vignette.

¹¹ In light of this worry, an alternative proposal holds that the content is something like, ‘within the vignette, Orlando has a justified true belief that *p*, but does not know that *p*’. One concern about this proposal is that vignettes can be inconsistent and perhaps even conceptually incoherent. Advocates of this proposal must explain how Orlando’s epistemic states within the vignette bear on which epistemic states are possible. The proposed judgement, in other words, does not itself threaten (4), since it is *prima facie*

3. The Necessity Proposal

Given the problems faced by the first two proposals, it seems worthwhile to examine other possibilities. To better articulate other proposals, I first introduce some terminology.¹²

(8) $V(x, p)$

This terminology expresses that every sentence in the vignette is true of a person and proposition. It can be understood informally as:

(9a) The vignette is true of person x and proposition p .

or

(9b) x stands to p as described by the sentences in the vignette.

or

(9c) x and p satisfy the vignette.

I take claims (9a-c) to provide roughly equivalent glosses of $V(x, p)$.¹³ I understand the vignette to be the set of claims used to describe a case. An instance of a vignette is any situation in which someone satisfies the vignette.

We can employ this apparatus to articulate a third proposal for the content of our operative judgements in response to vignettes. The third proposal posits that the judgement is ‘for any x and p that satisfy the vignette, x has a justified true belief that p , but does not know that p ’. Put slightly differently, ‘anyone who stands to a proposition as described by the vignette has a justified true belief in that proposition, but does not know it’. This content can be expressed:

(10) Necessarily, if the vignette applies to a person and proposition, then that person possesses a justified true belief in that proposition, but the belief is not knowledge.

possible that impossible things can be embedded in a ‘within the vignette’ operator. Thanks to Daniel Rubio for helpful discussion of these issues.

¹² This terminology resembles Williamson’s ‘GC (x, p)’, where ‘GC’ stands for ‘Gettier case’ (Williamson (2007: 184)). I prefer to not use Williamson’s terminology for several reasons. Firstly, it is ambiguous between the mere *sentences* of the vignette being true of a person, and the (more full-bodied, richer) fiction being true of a person. See Ichikawa (2009) and Ichikawa and Jarvis (2009; 2013) for more on this distinction. Secondly, ‘Gettier case’ seems like a success term: something is a Gettier case only if it successfully exhibits justified true belief without knowledge. But the mechanisms underlying the methodology are present whether or not the vignette is auspicious. A third reason is that my account of the method generalises to all thought experiments as counterexamples, and thought experiments in other roles, and so there is reason for the terminology introduced to apply smoothly outwith Gettier cases.

¹³ Note that this terminology is different from Ichikawa and Jarvis’s ‘truth in fiction’ apparatus. See Ichikawa and Jarvis (2009: esp. 227; 2013, esp. 204–5). This is because some facts might be true in the fiction, yet not articulated in the vignette itself and so irrelevant to whether $V(x, p)$. According to Ichikawa and Jarvis, it is true in the fiction that Orlando has two eyes, for example (Ichikawa and Jarvis (2009: 227)). But the question of eye-count does not arise for determining whether the vignette is true of somebody: the vignette does not mention eye-count. Thus if someone, S , has one eye, then it is possible for $V(x, p)$ to hold for S , but not possible for Ichikawa and Jarvis’s ‘true in the fiction’ relation to hold for S .

Claim (10) can be formalised as:

$$(11) \quad \Box \forall x \forall p [V(x, p) \rightarrow \Box (JTB(x, p) \& \neg K(x, p))]$$

By itself the truth of claim (11) does not threaten the target theory. The two claims are compatible provided that the vignette could not be satisfied; if the vignette is impossible, then $V(x, p)$ could never obtain, and (11) is trivially true. But we can combine (11) with the further premise that the vignette could be true of some x and p :

$$(12) \quad \Diamond \exists x \exists p V(x, p)$$

This proposal, which I call the necessity proposal, has many virtues. The conjunction of claims (11) and (12) conflicts with the target claim (6), and seems to explain several features of the methodology. Upon hearing a vignette the judgement that the vignette challenges the target theory is often quick. As soon as we comprehend the vignette we perceive the problem. We need not consider a particular variation or specification of the case, such as the possible variation where Orlando purchases the other Woolf-owned book over the internet or has possessed it for decades. This speed and effortlessness might be explained by the necessity proposal: if every case that satisfies the vignette is a counterexample to the target theory, the person need not think of a particular instance—any instance will suffice. We need not consider a specific instance of the vignette, and this proposal can explain that phenomenology.¹⁴

Relatedly the proposal might capture the generality that we understand the method to have. When we comprehend that the vignette constitutes a problem for the target theory, we also comprehend that various specific details do not affect its dialectical force. We recognise that it does not matter whether the bookseller has brown or blonde hair, for example, or whether Orlando is married. In this sense, the methodology has a kind of recognised generality.¹⁵ Arguably this proposal captures the generality by recording in the content of the judgement that any instance satisfying the vignette is a problem for the view. Note that it is not simply that there is generality—i.e. that Orlando's marital status is irrelevant—we also can recognise this generality. And the necessity proposal arguably captures this; the proposal holds that this generality is located in the *content* of the judgement.

The central problems with this proposal are that the proposed content, judgement (11), is false, and obviously so. A vignette cannot hope to provide a comprehensive accounting of a case—vignettes are typically between 3 and 12 lines long—and inevitably leaves a great deal unspecified. Many unarticulated facts are consistent with the vignette; $V(x, p)$ obtains in instances where the protagonist has brown hair, or blonde hair, or is bald. And it is not only incidental facts about Orlando's hair colour and marital status: $V(x, p)$ also holds—the vignette is satisfied—by instances where the given vignette is supplemented with epistemically relevant facts. Consider for example the following two vignettes:

¹⁴ I think this first putative advantage is spurious. This is because the thing it putatively explains—the speed of the judgement—is not well explained by the proposed content. The proposed content, if formed responsibly, requires that we make a judgement covering *every* possible version of the case. This process, even if possible, would not be immediate, contra the conjectured advantage.

¹⁵ See Malmgren (2011: especially 281 ff) for discussion of the generality of the method.

GULLIBLE ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. Orlando should have known better, since just the previous evening he watched the local news describe the bookseller as a fraud, liar, and manipulator, who was selling worthless books as historical artefacts. Orlando recognised the man from the exposé, but as he entered the shop Orlando decided to disregard the exposé. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf.

MARRIED ORLANDO. A fraudulent bookseller fakes evidence which appears to show that a particular book once belonged to Virginia Woolf. Orlando, deceived by the forgery, purchases the book for a large sum. He thereby comes to believe that he owns a book that once belonged to Virginia Woolf. As it happens, Orlando does own a book that once belonged to Virginia Woolf, since another book he owns—one he does not associate with her—once belonged to Woolf. He phones his wife to describe the purchase, and she—a competent and rigorous bookseller herself—tells him about the other Woolf-owned book.

These variants are consistent with the original Orlando vignette. They merely augment the vignette with further details. Some instances such that $V(x, p)$ —i.e. some instances in which the original vignette is satisfied—are also instances that satisfy the gullible Orlando or married Orlando vignettes. But the former intersection—situations that satisfy both the Orlando vignette and the gullible Orlando vignette—includes many instances in which Orlando is not justified in his belief. He ignored relevant evidence. The second intersection includes many instances where Orlando possesses knowledge—Orlando’s reliable wife informed him that p . Thus there are instances where $V(x, p)$, but where x lacks JTB or x possesses knowledge. Judgement (11) is false.

There are many possible ways to deviantly satisfy the Orlando vignette, such as instances in which Orlando is obsessed with Woolf, and so is disposed to believe that all kinds of his purchases once belonged to Woolf, or in which he resides in a country where it is common knowledge that booksellers are terrible frauds, or in which Orlando lives one thousand years after Woolf, and so any books once owned by her would be prohibitively expensive. Any of these ‘deviant variations’ render his belief unjustified. There are also many deviant variations that would render his belief knowledge. Call the falsity of claim (11) owing to deviant variations of the vignette the ‘deviant variation’ problem.¹⁶

The falsity of judgement (11) is a problem for the necessity proposal because if employment of counterexamples to challenge target theories rests on a false judgement, then our resultant beliefs are unjustified. We should not, for example, have rejected the JTB view of knowledge based on Gettier vignettes. The necessity proposal undermines, rather than vindicates, the method.

But the *existence* of deviant variations is not the only problem with the method. Perhaps, after all, our inference from hearing a Gettier vignette to rethinking the JTB view of knowledge is a flawed inference, and maybe a clear-eyed understanding of the method would reveal this. If the method is flawed in this way, the falsity of (11) is not a strike against the proposal. This would be a surprising conclusion, but not one we should dismiss outright. The second problem is that (11) is obviously false, yet many careful thinkers are convinced

¹⁶ The problem of deviant variations for the necessity proposal is discussed in Williamson (2005: 6–7; 2007: 185ff); Malmgren (2011: 275 ff); Grundmann and Horvath (2014); Ichikawa and Jarvis (2009: 223–5; 2013: 200–1).

by Gettier cases and other vignettes. It is implausible that so many thinkers commit such an elementary error. Even if we are wrong, we are not foolish.¹⁷

A further problem with the necessity proposal is that when we hear a compelling vignette, we continue to think it poses a challenge to the target theory even after we also realise that there are possible deviant ways to augment the case.¹⁸ We simply acknowledge that these deviant variations, whilst consistent with the text, are not the normal understanding of the text. But according to the necessity proposal, acknowledging these possible deviant variations should cause thinkers to judge the vignette benign; they realise claim (11) is false.

One could attempt to defend the necessity proposal by ruling out deviant variations. A flatfooted strategy posits a (tacit) ‘that’s all folks’ clause at the end of every vignette used in theorising. This clause conveys that no other facts should be imported; the story is complete as told. This defence of the necessity proposal holds that whenever theorists consider a vignette as a counterexample, they implicitly acknowledge that they cannot import any further facts. Thus the deviant cases described above, since they rely on importing “extra” facts, are ruled out.

But this flatfooted proposal will not work. We cannot consider a vignette without importing *some* facts not explicitly appearing in the vignette. These include that Orlando is alive, for example, and that the book is not on fire. If we are limited to explicit claims in the vignette, the case might be literally unimaginable.

A slightly more refined defence argues that we hold the epistemically relevant facts fixed (or, if the target theory concerns ethics or aesthetics, the ethical or aesthetic facts respectively). This proposal allows that we are free to round out the story as needed to comprehend the case, but proscribes importing facts that might bear on whether the protagonist possesses justified true belief or knowledge. But this strategy also lacks promise. Inevitably there are epistemically relevant facts not mentioned in the vignette itself. These include that Orlando is not a fool, that the bookseller is not a well-known fraud, and that Orlando did not later receive testimony about his other Woolf-owned book. Other facts, such as that Orlando does not own every book in the world, and that his recent purchase was not published within the decade, are not explicitly epistemic but are epistemically relevant in this case because they bear on the normative status of Orlando’s doxastic attitudes. We must import these extra facts into our understanding of the story: They are not explicitly stated in the vignette, but are expected. They fill out how we normally understand the case and help constitute typical instances of the vignette. I develop the role of expectations, normal understandings, and typicality judgements in section seven.

4. The Possibility Proposal

The obvious falseness of proposed judgement (11) suggests the necessity proposal does not capture the operative judgement we have in response to vignettes. But perhaps this is unsurprising: judgement (11) is far stronger than required. Rather than requiring every

¹⁷ See also Malmgren (2011: 276).

¹⁸ Malmgren (2011: 276) also makes this point, although she offers a different diagnosis—appealing to the vignette designer’s intentions—to explain why considering deviant variations does not undermine the perceived potency of the vignette.

instance that satisfies the vignette to exhibit justified true belief without knowledge, (plausibly) all that is required is just one possible instance where a person possesses justified true belief without knowledge. This is sufficient to challenge the target theory.¹⁹ Realising this suggests a natural alternative proposal. This fourth proposal holds that the judgement the vignette elicits is that at least one possible instance of the vignette exhibits justified true belief without knowledge. This proposal is advanced by Anna-Sara Malmgren, and can be expressed:²⁰

- (13) It is possible that some x and p satisfy the vignette, and x has a justified belief that p but does not know that p .

This can be formulated as:

- (14) $\Diamond \exists x \exists p [V(x, p) \ \& \ (JTB(x, p) \ \& \ \neg K(x, p))]$

This possibility claim entails the falsity of the target claim. Note this proposed content entails the possibility claim (12) $\Diamond \exists x \exists p V(x, p)$, which states that the vignette is possible, and so claim (12) need not be independently endorsed.

The possibility proposal has virtues. Firstly, it is psychologically plausible; whilst it does not seem that we consider or judge every possible instance that satisfies the vignette, it is plausible that we ‘latch onto’ one instance. Secondly, supposing the right conditions, the judgement is true. The judgement is false if either i.) the target claim is true, and so anyone who has justified true belief possesses knowledge, or if ii.) the vignette is abysmal, such that no augmenting of the vignette can lead to an instance that exhibits knowledge without justified true belief. An example of an abysmal vignette might be:

NO ONE. All the humans and animals have died. Earth was the only planet with life, and the explosion destroyed all living things. Gradually the buildings decay and collapse. A mug filled with rainwater falls through a rotted picnic table and smashes on the ground.

Perhaps no instance of this vignette, no matter how creatively augmented, can exhibit justified true belief without knowledge. This is because any candidates for Gettierisation are stipulated dead. But with either of these disjuncts—the truth of the target theory or an abysmal vignette—we ought *not* judge that the vignette generates a counter-instance to the theory; in both cases it cannot. Thus the possibility proposal, unlike the necessity proposal, seems to generate the right results about when the proposed judgement is true.²¹

The possibility proposal has been criticised for not vindicating the generality of the judgements we have in response to vignettes.²² There are several ways to unpack this

¹⁹ One possible counter-instance is sufficient to challenge a target theory that is, or entails, a necessity claim. One possible counter-instance is unlikely to challenge a target claim concerning what is typical, normal, or characteristic, such as ‘knowledge is normally justified true belief’ or knowledge is characteristically justified true belief’. I return to this issue in section 7.2.

²⁰ Malmgren (2011).

²¹ I argue below that this virtue is spurious. The proposed content is true about vignettes that are not genuine challenges to the target theory, and so is not a plausible candidate for the content of our judgement.

²² Malmgren (2011: 281–5).

criticism. One way is to note that it does not matter for the methodology whether the protagonist has two children, blonde hair etc. But the proposal says we fix on some particular instance, and so focus on an instance in which hair colour, etc. is determined. We should be able to recognise that the judgement is legitimate about a range of variations; it does not matter what hair colour Orlando has, for example. The implicit generality, it is argued, is important to our understanding of the method and so should be reflected in the judgement. But the judgement posited by the possibility proposal does not contain or allow any generality. Another way to unpack this criticism is to note that if a theorist thinks that the Orlando vignette is a challenge to the target theory when Orlando is brunette, but not when he is blonde, she has misunderstood the method. As Malmgren describes the criticism,²³

[If I hold these two beliefs, i.e. that Orlando has justified true belief without knowledge when he is brunette, but not when he is blonde] I am being confused or inconsistent: there is a problematic discord between my earlier and later judgement [...] Now, the natural way to account for this—the natural diagnosis of my apparent rationality failure—is to say that I am here *contradicting* my original intuitive judgement (or that I am contradicting something that obviously follows from it).

But suppose that the content of our judgement is, as Malmgren suggests, the possibility proposal. Then the person who judges that blonde Orlando knows, but brunette Orlando does not, has not judged anything inconsistent. The contents of the judgements concern different particular possible instances of the vignette. Since there is no generality in the judgement, the criticism contends, the proposal cannot explain the fault with this pair of judgements.

Malmgren responds to the generality objection. She distinguishes between generality in the content of a judgement and generality in the grounds for a judgement. She argues that the objection requires that the generality—and associated tension—be located in the content of the judgement. Malmgren concedes that if tension between the judgements—that brunette Orlando has justified true belief without knowledge and blonde Orlando does not—is in the content of the judgement, she cannot explain this. Her proposal cannot explain why these two judgements are *contradictions*.²⁴ But, Malmgren argues, the tension is actually in the *grounds* for holding the judgement: when someone judges that brunette Orlando has justified true belief without knowledge, this rationally commits her to the same judgement when something irrelevant, such as hair colour, is altered.²⁵ I will not adjudicate whether Malmgren's response is persuasive. We might note the dialectical force of the putative objection as a conditional: *if* the person who judges that 'brunette Orlando, but not blonde Orlando, has justified true belief without knowledge' has thereby contradicted herself, this is a serious problem for the possibility view. If she has simply violated rational demands, but not contradicted herself, this is not a problem for the possibility proposal.

The possibility proposal suffers from other problems. Consider someone, Ann, who believes there can be justified true belief that is not knowledge.²⁶ Ann antecedently denies the target claim. Perhaps she is convinced by thinking about lottery cases, kinds of testimony, or barn façade cases, or perhaps other reasoning has led her to think that

²³ Malmgren (2011: 285, emphasis in original).

²⁴ Malmgren (2011: 291).

²⁵ Malmgren (2011: 290–4).

²⁶ This objection is also found in Ichikawa and Jarvis (2013: 203).

knowledge differs from justified true belief. Now suppose Ann hears the Orlando vignette, but does not have the ‘Gettier intuition’; she does not judge that the vignette as told describes a case of justified true belief without knowledge. Given Ann antecedently believes that justified true belief and knowledge can diverge, she should agree that *some* instance that satisfies the vignette exhibits justified true belief without knowledge. All Ann must do is provide additional details that transform the vignette into a lottery case, for example, and judge that the variation exhibits justified true belief without knowledge. Whilst she does not think the case presented is compelling, Ann does agree that:

(13) It is possible that some x and p satisfy the vignette, and x has a justified belief that p but does not know that p .

(14) $\Diamond \exists x \exists p [V(x, p) \ \& \ (JTB(x, p) \ \& \ \neg K(x, p))]$

Ann will, in other words, upon hearing the Orlando vignette, agree to the content proposed by the possibility proposal. This suggests the possibility proposal cannot capture disagreement about putative counterexamples. Consider someone else, Bob, who judges that the Orlando vignette is a successful Gettier case. Bob believes that the natural understanding of the vignette describes someone who has justified true belief but not knowledge. Surely Ann and Bob disagree in their judgement about the case. But they agree about claim (14). If claim (14) represents the judgement we have about the case when we take it to challenge the target claim, this renders Ann and Bob’s agreement about (14) confusing. An account of the judgement in response to vignettes should represent Ann and Bob as having contrary responses to the vignette: one straightforwardly judges the vignette to exhibit justified true belief without knowledge, whilst the other does not. But, according to the possibility proposal, Ann assents to Bob’s judgement. This is a problem for the possibility proposal.²⁷ One diagnosis of this problem suggests that the operative judgement makes reference to whether knowledge obtains in specific versions of the vignette, not simply ‘some possible variation of the vignette’: Ann and Bob do disagree about more specific versions of the vignette.

There is a further problem for the possibility proposal. Consider this vignette,

ANEMIC VIGNETTE. Ted looks at a tree; he then forms a belief.

Note that the Anemic Vignette, like the Orlando vignette, can be filled out with supplementary details. We can augment the case with details about what Ted believes and how. Supposing that knowledge diverges from justified true belief, a hearer can provide supplementary details so that Ted possesses a justified true belief that is not knowledge.

²⁷ One response holds that whilst Ann and Bob agree on claim (14), their grounds for assenting differ. Bob’s grounds for (14) is reflection on the case as presented. Ann’s grounds are an antecedent commitment to the possibility of justified true belief without knowledge combined with the judgement that nothing in the presented vignette precludes augmentation into such a case. (Contrast the ‘no one’ vignette, above, which cannot be transformed into a case of justified true belief without knowledge.) But this response is unlikely to work: Ann and Bob disagree about the vignette itself and whether the presented case is a problem for the target theory. Very plausibly they have different judgements about the case, and this should be reflected in the proposed content. The possibility proposal cannot do this.

Since such augmentation is possible, and the vignette is possible, claim (14) is true of the anemic vignette. A hearer can endorse judgement (14), in other words, about the Ted vignette. But Ted provides no evidential weight as a counterexample to the target theory of knowledge. Before we are moved by the anemic vignette, we must learn more about how Ted's belief is justified, whether it is true, and what might impede Ted's knowledge. The existing vignette is not persuasive. The mere fact that we can (assuming knowledge is not justified true belief) endorse claim (14) about the vignette is neither here nor there in epistemological theorising.

The same problem applies to botched vignettes. A botched vignette is when the speaker, who has a compelling vignette in mind, bungles important details in the delivery, so that the normal interpretation of the vignette is not one which exhibits justified true belief without knowledge.²⁸ Since botched vignettes can often be augmented into successful counterexamples, the possibility judgment (14) is usually true of botched vignettes, but this does not seem like the right result. Botched vignettes do not present a challenge to the target theory.

This might seem like a strange criticism of the possibility proposal—after all, anyone who thinks the anemic or botched vignette as it stands constitutes a problem for the target theory is mistaken. But I think it illuminates a problem with the possibility proposal: the proposal cannot vindicate the evidential significance of vignettes as counterexamples. Vignettes generate potent, compelling evidential weight. One explanation of their effectiveness is that when we hear a vignette we thereby *confront* an instance of, for example, justified true belief that fails to be knowledge. The vignette does not merely make us realise that such as instance is (merely) possible. The compelling vignette itself presents us with an instance that is counter to the target theory.²⁹ But the possibility proposal does not paint this picture. The possibility proposal only says that we judge that some version or other of the vignette is justified true belief without knowledge. The proposal does not quite capture the way that we are *faced* with a counter-instance.

Note too that vignettes play a social role: theorists agree, disagree, and persuade each other about vignettes. The possibility proposal might render it mysterious how vignettes fulfil this social role since, according to the possibility proposal, the counter-instance itself is not communicated by the vignette. Only the judgement that 'a counter instance based on or consistent with this vignette is possible' is evinced by the vignette. But it seems that when we recount vignettes we convey more than that a possible variation of the vignette is a case of justified true belief without knowledge. Instead we seem to convey something more like particular illustrations of justified true belief without knowledge.

To illuminate this social role, suppose a theorist told an anemic or botched vignette, and a second theorist, on that basis, judged the target theory false. This would be a concern. It does not seem like good philosophical communication or a legitimate basis on which to reject the target theory. If we learn that the intermediary judgement was merely (13) 'it is possible that some x and p satisfy the vignette, and x has a justified belief that p but does not know that p', this will not allay our concerns. To judge that the theorists were working well,

²⁸ An example of a botched vignette is, 'Roddy was looking at a field, and he saw lots of sheep. One the things he thought was a sheep was actually a rock. Roddy formed the belief "there are sheep in the field".'

²⁹ If the Orlando vignette merely describes something that *could* be a case of justified true belief without knowledge, as the possibility proposal suggests, this provides significantly less evidential force against the target theory.

we want to know about the specific variant of the vignette each had in mind, to see whether it was a lottery case, a Gettier case, or some other case. But the possibility proposal does not generate resources to explain why this interaction is ineffective. According to the possibility proposal, this ‘possibility’ content is all we convey, enable or evince when we employ vignettes as counterexamples.

To summarise this strand of objection: Insofar as (14) captures the content of our judgement, then to that extent the vignette is not dialectically compelling. A vignette that generates a stronger judgement would be more compelling. But vignettes in philosophy are remarkably compelling; given how compelling vignettes are, (14) cannot adequately capture our reaction to them. (Indeed perhaps less compelling counterexamples do generate judgement (14). Perhaps sometimes we hear a case and judge ‘there’s a counterexample lurking in there’, or ‘some version of this case is a counterexample’. But such counterexamples do not typically legitimate jettisoning established theories. Instead they are an invitation to construct a better vignette.) Given the method’s effectiveness, the content of our response to the vignettes must be stronger than the possibility proposal.

A general difficulty with the possibility proposal is that judgement (14) is too weak. Given that vignettes can be augmented, and knowledge and justified true belief diverge, the judgement is often true of a vignette. It is true of *all* possible vignettes other than those that stipulate either i.) that the person knows p, ii.) that the person does not have a justified true belief that p, or iii.) there are no potential candidates for Gettierisation, such as in the ‘no one’ vignette.³⁰ Unless one of these conditions holds, the vignette can be transformed into a case of justified true belief without knowledge, such as a Gettier case, barn façade case or lottery case. Accordingly claim (14) is true of almost every vignette. This suggests the judgement we have in response to cases is stronger.³¹

5. Williamson’s Counterfactual Proposal

The fact that we are *struck* by vignettes as counterexamples suggests the judgement elicited by the vignettes concerns a more specific instance of the vignette. We are confronted by a problem instance, and this should be reflected in the content of our judgement. A fifth proposal for the content of our operative judgement is advanced by Timothy Williamson. He holds that the content of the judgement is the subjunctive:³²

- (15) Were the vignette to be satisfied by an x and p, then she would have justified true belief, but not knowledge, that p.

³⁰ Perhaps more precisely the first two conditions should read i.) that every person knows everything and ii.) that no person has a justified true belief. These conditions entirely preclude judgement (14) being true of an augmented variation of the vignette.

³¹ As I explain in section 7.2, an account of the content of our judgements in response to vignettes should explain other roles that vignettes play in philosophy, such as motivating claims, conveying ideas, illustrating distinctions, and homing in a topic or debate. I argue below that the possibility proposal is ill-equipped to do this.

³² Formalising the precise content of Williamson’s counterfactual proposal is complicated by the problem of donkey anaphora for counterfactuals. For discussion see Williamson (2001: 195–9 and appendix two), Ichikawa and Jarvis (2009: 225, n.4; 2013: 201, n.6), and Malmgren (2011: 277; 306).

Williamson uses a Lewisian possible worlds understanding of counterfactuals to interpret claim (15) as something like ‘in the nearest possible worlds where the vignette is satisfied, *x* has a justified true belief that *p*, but does not know *p*’.³³ Since this proposed counterfactual content, unlike the possibility content explored above, does not entail the possibility of the vignettes being satisfied, this proposal requires that we also judge:

$$(12) \Diamond \exists x \exists p V(x, p)$$

which expresses that the vignette could be satisfied by some *x* and *p*. The judgement posited by Williamson’s counterfactual proposal is less demanding than the judgement expressed by (11), the necessity proposal. This is because it does not require that every instance satisfying the vignette manifests justified true belief without knowledge. It thus avoids the ‘deviant variations’ problem faced by the necessity proposal. The proposed judgement is more demanding, on the other hand, than that expressed by (14), the possibility proposal. This is because the truth of the counterfactual (15) requires more than that one possible instance of justified true belief satisfies the vignette and is not knowledge; it requires that those possible instances would obtain, were the vignette true. This allows the proposal to avoid the ‘anemic vignette’ and ‘botched vignette’ problems facing the possibility proposal and better explain the disagreement between Ann and Ben. Were the ‘anemic vignette’ to occur, it almost certainly would not be a Gettier case—since most beliefs are not Gettiered, it would be a remarkable coincidence if Ted’s belief were a Gettier case. And Ann would not endorse judgement (15). So Williamson’s approach posits a claim about which Ann and Ben disagree. Perhaps then, this fifth proposal about the content of the judgement is ‘just right’, laying between judgements that are too weak and too strong.

Judgement (15) seems true about paradigmatic successful counterexamples: were the Orlando vignette to occur, it would be an instance of justified true belief without knowledge. Intuitively, given the way the world is, the variations on the Orlando case that do not exhibit justified true belief without knowledge are likely more modally distant possibilities than variations in which the Orlando case is a Gettier case. Williamson’s counterfactual proposal avoids the ‘deviant variations’ objection; the mere possibility of deviant variations are irrelevant to the counterfactual judgement, and so do not threaten the proposal.

This proposal also plausibly accords with the phenomenology of forming a judgement upon hearing a vignette. Plausibly we consider, ‘were that to occur, would it be justified true belief without knowledge?’ And the proposal seems well-equipped to explain the speed of judgements about successful cases: we need only think about a relatively straightforward counterfactual question.

According to Williamson’s proposal, the world provides the unasserted details.³⁴ Details such as whether Orlando’s wife informs him of the other book or whether the bookseller was featured in the previous day’s news are furnished by whether, were the vignette to occur, these features obtain. These counterfactual facts depend on facts about the actual world. This feature seems like a virtue of the view: it is plausible that (in at least some ways) the world determines unarticulated details of cases. Suppose that in the actual world booksellers were often fraudulent, and this was common knowledge. Suppose bookselling

³³ Lewis (1973).

³⁴ Williamson (2007: 186).

were a devious profession with lax regulations, ubiquitous corruption, and inadequate methods of verifying the authenticity of books. In this scenario, bookselling resembles the practices of fortune tellers, relic touters, and email spammers. If this were the circumstance in the actual world, it seems that were the Orlando story true, Orlando would not have a justified belief. And were this the actual world, if the original Orlando vignette were told hearers would likely not think it causes a problem for the target view, because Orlando's belief would be unjustified. Thus Williamson's counterfactual approach captures that the judgements depend on contingent facts about the world.³⁵

This proposal might also explain our ambivalence about cases such as TrueTemp and the Reliable Clairvoyant.³⁶ These cases are employed to test theories in epistemology, but many people simply cannot adjudicate whether they qualify as cases of knowledge (or whether there is a fact of the matter), and so whether they constitute a challenge to the target theory. Perhaps this uncertainty stems from difficulties forming the relevant counterfactual judgement about the vignette: it is too hard to determine what would be known in a world so epistemologically alien that reliable clairvoyance, or brain scans that determine belief contents, existed. Similarly if a vignette involves many intricacies or implausible psychology it might be harder to determine the relevant counterfactual, and so the vignettes are not perceived as challenges to the target theory. It would be a substantial theoretical virtue if Williamson's proposal not only predicts which cases are problem cases for a theory, but also explains why with some vignettes it is hard to determine whether they challenge the target theory.

Although the proposal has many virtues, there are also problems. The first worry is whether the world really can supply the relevant background details. It seems implausible that the world fills in facts about Orlando and the bookseller; they are fictional characters (or fictional characters *at best*, since perhaps they do not enjoy enough facets to qualify as genuine characters). Given that Orlando is merely fictional, it might be indeterminate whether the nearer worlds are ones where Orlando watches an exposé about the charlatan, is married to another bookseller, or is dreadfully gullible.

We might suppose that most people are not, and so Orlando is likely not, married to a bookseller or terribly gullible. But Orlando is not an actual person, so it is unclear whether this induction is legitimate; it is unclear how the actual world can provide counterfactual facts for non-existent, fictional entities. If I tell an untrue story about Stephen Fry, by contrast, but omit explicit facts about his psychology, the world might fill in that he is not foolishly gullible in my story, because he is not so in real life. Were my story true, it would be true that Fry would not be gullible. But it is not clear how a similar mechanism can apply to fictional characters.

It seems more plausible that the actual world helps determine what the hearer takes for granted and what defaults can be supposed, and so helps determine what the vignette-teller should mention and what she can leave unarticulated. The actual world affects, in other words, how we understand the case. In the actual world, booksellers are not typically fiendish frauds, and so hearers can assume Orlando has default licence to trust the seller without the vignette-teller needing to make this explicit. And most people are not hopelessly gullible, so the fact Orlando is not hopelessly gullible need not be mentioned. But most

³⁵ As I explain in section 7.1, Williamson's counterfactual view overstates the ways that the truth of the judgement, and the success of a counterexample vignette, depends on contingent facts about the world.

³⁶ For the TrueTemp case, see Lehrer (1990: 163–4). For the Reliable Clairvoyant case, see Bonjour (1980).

people do not fortuitously possess books that once belonged to Virginia Woolf, so the teller must explicitly mention this feature. Hearers would not otherwise assume it.

Facts about the actual world inform, in other words, our storytelling norms: if Orlando is gullible and disregarded a televised exposé, this is relevant to the story and is an atypical and abnormal feature, and so the storyteller would mention it. It would not be an expected part of a backstory, and so must be articulated. If something would be abnormal, but is not articulated in the vignette, this licenses hearers to assume the more normal version. Since the vignette-teller did not say Orlando was unusually gullible, for example, the listener can suppose that he is not. But this mechanism is not the actual world directly determining which Orlando-backstory is a closer possibility, as Williamson's counterfactual view holds. Instead narrative norms—which are informed by the nature of the actual world—help determine what hearers should infer and how they understand the case.

A second worry about the proposal is that the counterfactual judgement might be too demanding.³⁷ Suppose that, unbeknownst to the vignette-teller, in the actual world the Orlando vignette is in fact satisfied, but it is satisfied in a deviant way, such that it is not justified true belief or is knowledge. In this eventuality, the actual world is one where something like 'gullible Orlando' or 'married Orlando' obtains.³⁸ This renders the judgement expressed by (15) false. But, the worry continues, this does not seem like the right result. It seems that even if the actual world is, surprisingly, a 'married Orlando' world, and so the putatively fictional Orlando case is actually deviantly satisfied, the Orlando vignette remains a successful challenge to the target theory.

Williamson argues that this objection is mistaken. He argues that if the nearest instance that satisfies the vignette is not a case of justified true belief without knowledge—if, in other words, judgement (15) is false—then the vignette is *not* a successful counterexample, even if no one ever realises this counterfactual. In these cases, he says, our judgement that the vignette challenges the target theory is mistaken. Williamson notes that thinkers are often reluctant to admit error, and claims that if we insist on the potency of a putative counterexample after learning that the nearest satisfaction of the vignette is deviant then we manifest this stubbornness.³⁹

Williamson's response does not sound correct to me. The methodology of using vignettes to illustrate problems with theories requires that those vignettes be sensitive to acknowledged or systematic features of the world. The vignettes are less compelling to the extent that they depart from our ordinary understanding of how the world is, or violate scientific (including psychological) theories. But the method is not hostage to unexpected quirks of the actual world. If, unbeknownst to the vignette-teller, it happens that the Orlando vignette is satisfied by someone who knows that *p*, this should not impugn their using the vignette as evidence against the target theory. A bizarre anomaly in the actual world does not affect whether the Orlando vignette succeeds in challenging the target theory. Whilst the world fills in the gaps—by informing storytelling norms and shaping our expectations about systematic acknowledged features, such as whether clairvoyance is

³⁷ See Williamson (2007: 200); Malmgren (2011: 279); Grundmann and Horvath (2014); Ichikawa and Jarvis (2009; 2013); and Ichikawa (2009) for discussion of the deviant realisations problem.

³⁸ The same would happen if the Orlando vignette is not satisfied in this world, but in the nearest worlds where it were satisfied, a deviant realisation obtains. This also renders judgement (15) false.

³⁹ Williamson (2007: 200–1).

reliable and whether booksellers are often fraudulent—the specific vignettes should not be held hostage to unusual quirks and isolated contingent features of the actual world.⁴⁰

Note that the kind of anomalies Williamson has in mind—the ones that can, according to Williamson, annul a vignette—are truly unusual. They are not systematic fixtures of the world such as the truth of clairvoyance or whether booksellers are typically charlatans. And they are not widely known contingent features, such as whether Virginia Woolf owned books. Instead they are one-off anomalies, including ones that are unknown and unknowable.⁴¹

Suppose in philosophy class a professor tells the sheep vignette to challenge the view that justified true belief entails knowledge. A student raises his hand and comments that this occurrence happened to him once, but explains that ‘in my case there were some other sheep in the field I saw too; it’s just that the first one I thought I saw was actually a rock with a sheep behind it’. In this case the class plausibly learns that the sheep vignette has a deviant actual realisation.⁴² It seems that many students in the class will, despite the interjection, see how the vignette challenges the target theory. They perceive that justified true belief and knowledge can diverge.

According to Williamson’s counterfactual view, these students would be mistaken. But it seems like a perfectly fine piece of reasoning. In fact, if a classmate dwells on the deviant case, and so does not see how the initial vignette challenges the target theory, it seems she manifests less supple and sophisticated thought. The teacher might draw attention to the natural, normal understanding of the vignette and explain how, given the content of the vignette, the most natural interpretation is one where no other sheep were seen.⁴³ If the classmate still cannot see any problem for the target theory, because of the deviant realisation described—but would have perceived the problem had the student never spoken up—it seems that she does not understand the method. The vignette can and should have dialectical force, despite acknowledging a deviant realisation; its success is not hostage to abnormal isolated occurrences in the world.⁴⁴

⁴⁰ As Ichikawa and Jarvis note, if we reject the JTB account of knowledge based on a vignette, and judgement (11), but the vignette in fact has a deviant, non-Gettier realisation, the rejection is itself a Gettier case (Ichikawa and Jarvis (2009: 226; 2013: 203)).

⁴¹ Unpacking the distinction between isolated contingent anomalies and systematic features will take some work. But I think we have a good enough understanding of the distinction to cautiously proceed.

⁴² There are difficulties: it is not clear whether a vignette that is intended to feature wholly fictional characters can be satisfied by real people. As Kripke argues, perhaps fictional protagonists are essentially fictional (Kripke 1980; 2000). As Williamson notes, and Ichikawa and Jarvis develop, we might understand the names in these vignettes as simply placeholders for variables. See Williamson (2007: 184); Ichikawa and Jarvis (2009: 229, n. 13; 2013: 207).

⁴³ Suppose the teacher adds ‘no other sheep were seen’ after the student pipes up. Plausibly according to Williamson’s view the teacher has now presented a different case—the vignette picks out different possible worlds. But a more natural understanding of this is that the teacher has simply further described or clarified the original case. It is the same case, but with more description.

⁴⁴ Ichikawa notes that if we learn about a deviant instance of a vignette, we might change, expound on, or further describe the vignette, so that it precludes this deviance. Whilst this might appear to support Williamson’s counterfactual proposal, because deviant actual realisations cause us to adjust the vignette, Ichikawa notes that we also do this with *merely mentioned* possible deviant instances. This suggests we amend or further describe the cases to be clearer or more convincing in our presentation, rather than to generate the right truth conditions for the counterfactual judgement. See Ichikawa (2009: 439).

Note too that if another student subsequently contributes, ‘this happened to me too, but I did not see any other sheep. I just continued to believe the rock was a sheep. I discovered the truth months later’ and thereby describes a non-deviant instance of the vignette, plausibly judgement (15) is nonetheless false. This illustrates a further problem with the proposed content: the semantics of counterfactuals plausibly render (15) false if there are some actual deviant realisations and some actual non-deviant realisations. But this seems like a bad result for the counterfactual proposal: If the class hears a successful-sounding vignette, *and* hears compelling evidence that some actual realisations of the vignette are justified true belief without knowledge, this presents a hefty challenge to the target theory. But in this case judgement (15) is false. Since judgement (15) is false whenever actual or modally nearby instances are divided between deviant and non-deviant realisations, proponents of the counterfactual view must provide an explanation of why these vignettes constitute a challenge to the target theory.

The fact that when reasoning with vignettes we can ignore deviant actual instances, as the students in the above class should, provides evidence that discussing vignettes does not commit theorists to transient, quirky, unusual, isolated or unexpected features of the world. Instead it seems we can focus on a normal understanding or typical instance of the vignette. The teacher might even say of the student’s own (deviant) sheep-rock experience ‘that isn’t the *typical* case’, even though it is an actual instance. The actual might not be typical, and the typical might not be actual.

Jonathan Ichikawa and Benjamin Jarvis have argued that a consequence of Williamson’s counterfactual proposal is that we can never know whether a vignette is a successful challenge to the target theory.⁴⁵ This is because, they suggest, the relevant counterfactual is too difficult to know. Ichikawa writes,⁴⁶

The worry is that in too many cases, it is not plausible that we know the relevant counterfactuals, because it is not plausible that we know whether we are situated deviantly in modal space. The world is a big place, and we should not be at all surprised if [a deviant instance obtains]. And I see few prospects for learning whether [the counterfactual judgement is true].

I am not convinced: we do in general know many true counterfactuals, even when we cannot rule out their being false (just as in general we know many truths, even when we cannot rule out their being false). And so perhaps we do know that, for example, were a person deceived by an unscrupulous bookseller in the way described by the vignette, but happened to own a book of that description, they would have a justified true belief that is not knowledge. The facts that would make the protagonist’s belief knowledge or unjustified are unlikely, and so we can know the counterfactual. We are not in a position to rule out that a deviant instance would obtain, but it is sufficiently unlikely, so we can nonetheless possess the relevant knowledge. Thus I am not convinced that Williamson’s counterfactual proposal posits content that cannot be known.

But I have a related worry: The proposal makes the wrong kind of factors epistemically relevant. Whether we are correct that the vignette constitutes a challenge to the target theory should depend on whether the vignette exhibits poor understanding of natural laws or faulty science, is conceptually confused, or relies on implausible claims. According to

⁴⁵ Ichikawa and Jarvis (2009: 226; 2013: 202); Ichikawa (2009: 440).

⁴⁶ Ichikawa (2009: 440).

Williamson's proposal, by contrast, whether a successful-seeming vignette challenges a target theory depends on isolated quirks of the actual world. This does not seem correct: if a sophisticated thinker has devised a vignette, such as the sheep or Orlando vignette, and they have not made conceptual errors, or violated our understanding of the world, then they are in a position to know the falsity of the target theory, regardless of strange and unlikely happenstance. Even if the nearest world happens to be deviant, we still know that the vignette constitutes a challenge to the target theory. This is not to say that we can never be wrong about when a vignette constitutes a successful challenge; but Williamson's view generates the wrong sources of error.

Another weakness of Williamson's counterfactual proposal stems from the 'anemic' and 'botched' vignettes. Recall the anemic vignette,

ANEMIC VIGNETTE. Ted looks at a tree; he then forms a belief.

Claim (15) might well be true of the anemic vignette. It might be that given facts about this world, were Ted to believe something it would be justified true belief without knowledge. And so according to Williamson's counterfactual proposal it is possible—if the world were thus situated in modal space—that the anemic vignette presents a challenge to the target theory. But this cannot be right. The anemic vignette is not, and could not be, a challenge to the target theory. The vignette should not be a successful vignette, in other words, but according to the counterfactual view it might be. As I shall argue, the anemic vignette is not a challenge to the target theory because the most natural understanding of the vignette does not provide an instance of justified true belief without knowledge.

Recall that a 'botched' vignette fails to describe a counter-instance. Suppose that a botched vignette happens to describe what, in the nearest world where it obtains, is in fact a Gettier case. Williamson's proposal says the botched vignette constitutes a challenge to the target theory. But that does not sound right.⁴⁷ Whilst the deviant realisations problem suggests that the counterfactual proposal is too demanding, the anemic vignette and botched vignette problems suggest it is too permissive: Judgement (11) is true in cases where the vignette should not be understood as a challenge to the target theory.

6. The Emendation Approach

The problems with the counterfactual proposal have lead theorists to suggest a different approach to understanding the vignette method. Grundmann and Horvath argue that the necessity proposal can be defended by amending the vignettes. Grundmann and Horvath focus specifically on Gettier cases. They argue that since deviant instances of Gettier-style vignettes can be sorted into two kinds—those where the protagonist fails to have justification, and those where the protagonist knows via some other route—there is a simple and systematic way to amend the vignettes.⁴⁸ They suggest theorists stipulate that the belief is justified. They argue that since theorists already usually stipulate that the proposition is

⁴⁷ Proponents of the counterfactual view do have some resources available. They might argue that whilst the anemic or botched vignette is in fact a successful counterexample to the target theory, theorists lack grounds for thinking that it is. In the absence of good grounds, a theorist who (correctly) takes the anemic or botched vignette to challenge the target theory is making an error.

⁴⁸ Grundmann and Horvath (2014: 428).

believed and true, this amendment is benign. Then we stipulate that if the belief is known it is known only in virtue of the ways described by the vignette. This precludes independent sources of knowledge, such as that described by the ‘married Orlando’ augmentation. Thomas Grundmann and Joachim Horvath illustrate this proposal with an example of a typical vignette in epistemology.⁴⁹

FORD. Smith believes that Jones owns a Ford, on the basis of seeing Jones drive a Ford to work and remembering that Jones always drove a Ford in the past. From this, Smith infers that someone in his office owns a Ford. In fact, someone in Smith’s office does own a Ford—but it is not Jones, it is Brown. (Jones sold his car and now drives a rented Ford.)

Grundmann and Horvath provide a version amended according to their proposal. Following their convention, I have underlined their modifications.⁵⁰

AUGMENTED FORD. Smith justifiedly believes that Jones owns a Ford, on the basis of seeing Jones drive a Ford to work and remembering that Jones always drove a Ford in the past. From this belief alone, Smith logically infers, at time t, to the justified belief that someone in his office owns a Ford, which provides his only justification for that belief at t. In fact, someone in Smith’s office does own a Ford, so that Smith’s latter belief is true—but it is not Jones, it is Brown, and so Smith’s initial belief was false. (Jones sold his car and now drives a rented Ford.) Also, if Smith knows at t someone in his office owns a Ford, then he knows this at t only in virtue of the facts described.

But the ‘emendation proposal’ has problems. Firstly, it is not how vignettes are actually presented, and so the proposal is perhaps best understood as guide to revising philosophical practice, rather than a proposal for understanding the existing method. Some of Grundmann and Horvath’s comments suggest their intention is prescriptive rather than descriptive. They say, for example,⁵¹

Improved case descriptions might nevertheless help prevent philosophical laypersons from misinterpreting the relevant descriptions. In this way, an improved case description like [augmented Ford] might help to avoid potential confounds in experimental studies that test the Gettier judgments of philosophical laypersons.

If intervention, improvement, or prescription is the aim of the emendation proposal, it responds to a different question from the other proposals: Grundmann and Horvath are not asking how the method works, instead they ask how it could be improved. Accordingly the proposal is not a genuine alternative to the others.⁵² But Grundmann and Horvath suggest that ‘augmented Ford’ is how experts implicitly understand the vignettes, and so their proposal is descriptive rather than revisionary. They write,⁵³

something along the lines of [augmented Ford] seems like a good reconstruction of how the relevant experts implicitly interpret the original case description (the Ford vignette) when they perform the thought

⁴⁹ Grundmann and Horvath (2014: 525).

⁵⁰ Grundmann and Horvath (2014: 529–30). For ease of presentation, I collate the underlining from pages 529 and 530.

⁵¹ Grundmann and Horvath (2014: 532).

⁵² Ichikawa and Jarvis (2009: 224) makes a related point.

⁵³ Grundmann and Horvath (2014: 531–2). I have adjusted the way the Ford vignette is referred to, to better match my nomenclature.

experiment. Arguably, no professional epistemologist would be seriously tempted to interpret [the Ford vignette] in any of the deviant ways that are left open by this description. Moreover, the way in which we amended the initial description (the Ford vignette) should be readily available and transparent to every professional epistemologist. It therefore seems quite plausible that professional epistemologists tacitly interpret [the Ford vignette] in a way that roughly corresponds to our improved description (augmented Ford).

Experts do not interpret vignettes as deviant variations. But is this because they respond to an emended vignette, as Grundmann and Horvath suggest, or because they simply interpret the given vignette in a more narrow and less literal way? There is reason to think the latter: Firstly, Gettier-style vignettes are almost never described in the emended way. Given how experts in philosophy are keen to make the implicit explicit and articulate each assumption carefully, this suggests that the amendments to the vignette are not required to successfully elicit the Gettier judgement, and that emending the vignette as described by Grundmann and Horvath is not part of the method. Instead it seems that theorists respond to the vignette *as presented* (not the emended version), but do so in a more sophisticated way than simply quantifying over the instances where $V(x, p)$.

Secondly, laypeople judge in response to vignettes. These responses seem akin to expert responses. But it is unlikely that laypeople respond to the given vignette by ‘translating’ it into the amended vignette, even implicitly. Laypeople rarely use the expression ‘justified belief’, for example, and ‘in virtue of’ is a relatively technical notion.⁵⁴ The amount and content of emendation required varies by vignette in relatively sophisticated ways, furthermore. The above example uses time-indexed beliefs, ‘logically infers’, and other notions that different vignettes do not require. Explicit emendation is challenging.

This suggests that the emendation proposal does not capture what theorists and lay people do. It seems that rather than translating the vignette into the amended vignette, and responding to the amended vignette, people perform something different. In the next section I argue that they respond to the vignette itself, but they do so by interpreting it with good judgement by considering the normal version of the vignette.

Perhaps one of the most powerful and compelling features of the vignette method is that a simple, mundane story elicits a powerful response. The vignette contains few technical or explicitly epistemological details. Yet the response often is a distinct and confident epistemological judgement, one compelling enough to unseat established theories. One weakness of Grundmann and Horvath’s proposal is that it generates a more confusing and complicated vignette to respond to, and so arguably erodes the potency of our response. Grundmann and Horvath argue that since theorists often already stipulate that the proposition is true and believed, it is not damaging to also stipulate that the belief is justified. I demur. One reason the vignettes are so successful is that those very same (largely not explicitly epistemic) features that hearers recognise as making the belief justified nonetheless fail to make the belief known.⁵⁵ That hearers actively realise that the vignette describes a justified belief—rather than being passively told within the vignette—might well contribute to the efficacy of the method.

Finally, as Grundmann and Horvath acknowledge, the method does not readily generalise. Since the augmentation proposal is specific to Gettier-style cases, it does not

⁵⁴ Caton (m/s).

⁵⁵ Many thanks to Lisa Miracchi for emphasising this point.

generalise to vignettes as counterexamples to other necessity claims. It also does not illuminate the other roles that vignettes play in philosophical reasoning. As I argue in section 7.2, my proposal has the advantage of illuminating and explaining various roles of vignettes in philosophy.

7. The Normalcy Proposal

Rather than positing a simple quantification over the situations in which the vignette is true, or considering what is described in the most literal sense by a vignette, I propose that upon hearing a vignette, we instead think of what constitutes a normal, typical, generic, or natural instance of the vignette.

When we reason, we often use the notions of typicality and normality. It is very natural to think things like ‘lying is usually wrong’, ‘we normally meet on Wednesdays’, and ‘Scandinavians typically buy their first home in their early thirties’. It is, I think, less typical to think in terms of strict necessity or mere possibility. Thinking in terms of normalcy and typicality are important and ubiquitous parts of both ordinary and scientific thought, and such thoughts are not easily translated or reduced to claims of probability, prevalence, counterfactuals, or other forms of quantification.

My proposal is that when we hear a vignette we consider the normal understanding of that vignette. If the normal versions of the vignette exhibit justified true belief without knowledge, then the vignette constitutes a challenge to the target theory. We judge something like,⁵⁶

- (16) Normally, whenever a person is related to a proposition such that the vignette obtains, then that person will have JTB without knowledge.

This is formalised below, where NORM denotes the operator ‘normally’:

$$(17) \text{ NORM } \forall x \forall p [V(x, p) \rightarrow (JTB(x, p) \& \neg K(x, p))]$$
⁵⁷

What we judge to be a normal instance of the vignette is influenced our background beliefs about science, laws, politics etc.⁵⁸ It is also influenced by storytelling norms. As described in section five, if something surprising or abnormal happens, it would be mentioned by the vignette-teller. The omission from the vignette thus allows the hearer to infer that the more normal and expected version occurred. Given that it is a vignette about belief-formation, if the protagonist has another route to knowledge, for example, it would be mentioned by the

⁵⁶ Note too that the judgement that the vignette is possible must also be adjusted from ‘the vignette could be true of someone’, ($\Diamond \exists x \exists p V(x, p)$), to ‘normally, the vignette could be true of someone’ or ‘a normal version of the vignette could be true of someone’.

⁵⁷ To emphasise the departure from the usual approaches to vignette satisfaction, we might also proceed using the apparatus ‘NORM (x, p)’ which conveys ‘x and p satisfy a/the normal understanding of the vignette’. We can then posit the judgement, ‘necessarily, anyone who stands to a proposition p as in the normal vignette has a justified true belief that p but does not know that p’. We can formalise this as ‘ $\Box \forall x \forall p [NORM(x, p) \rightarrow \Box (JTB(x, p) \& \neg K(x, p))]$ ’.

⁵⁸ Given the importance of background beliefs in judgements concerning the normal or typical understanding of a vignette, there are parallels with modelling non-monotonic logics, default logics, and other attempts to model natural thought and ceteris paribus clauses.

vignette-teller. But this inference occurs without the hearer explicitly thinking about things like knowledge or storytelling norms; it is simply how we in fact understand anecdotes, explanations, and stories.

The normal understanding of a vignette does not simply interpret the vignette as though it were a fiction, as suggested by Ichikawa and Jarvis. Ichikawa and Jarvis suggest that when we consider a successful vignette our response is:⁵⁹

$$(18) \quad \Box (p \rightarrow \Box \exists x (x \text{ has justified true belief without knowledge}))$$

where p stands for the proposition ‘the members of *story* are true’. *Story* is the set of sentences that are true in the fiction. This is an infinite set, and is based on the given vignette and our interpretation of it. The interpretation provides something richer than the mere vignette, and is constructed using our ability to recognise truth in fiction. Ichikawa and Jarvis provide an example: it is not mentioned in the vignette that Orlando has two eyes, for example, but it is true in the fiction that he does.⁶⁰ So if someone, S , has one eye, then it is possible for $V(x, p)$ to hold for S , but not possible for Ichikawa and Jarvis’s ‘true in the fiction’ relation to hold for S .

Whilst our proposals share some features, I have several concerns about the truth-in-fiction proposal: Firstly, it is hard to see how this view captures agreement and disagreement about vignettes. This is because every hearer will round out the vignette slightly differently. Since they create different stories from the vignette, perhaps their judgements (18) will not have the same content. When hearers seem to agree and disagree about vignettes, in other words, they might have different propositions in mind, which would render agreement or disagreement illusory.⁶¹ The normal version of the vignette, by contrast, supplies fewer details and is therefore interpersonally available: people can agree on the normal version of a vignette, even if they do not agree on the fictional version.

Secondly, it is not clear how we produce the set *story*.⁶² We do not seem to generate a genuine full-blooded fiction from hearing vignettes: when we hear a trolley case, for example, we do not imagine details concerning what kind of train it is, where it is going, or how the driver is feeling (unless we deem these things relevant to either the normal version of the vignette or the philosophical question at stake).⁶³ The normalcy approach does not require that we fill out these kinds of details; according to the normalcy approach we supply some extra details in addition to the vignette, but the case remains largely schematic. The view seems to posit a relatively complicated operative judgement, furthermore, since it appears to invoke the notion of sets and use a meta-language.

Thirdly, the proposal seems too narrow: some vignettes describe acknowledged non-fiction, such as anecdotes or historical events. It seems strange that we would interpret these using our capacity to understand fictions, as the truth-in-fiction proposal suggests. The

⁵⁹ Ichikawa and Jarvis (2013: 207; 2009: 229).

⁶⁰ Ichikawa and Jarvis (2009: 227).

⁶¹ See also Williamson (2009: 468). The same concern arises for Ichikawa and Jarvis’s ‘demonstrative’ suggestion; the concern is that theorists might be picking out different things by ‘that’, rendering agreement illusory (Ichikawa and Jarvis (2009: 232)).

⁶² See also Williamson (2009: 467–8).

⁶³ Williamson argues that the truth-in-fiction proposal is psychologically implausible because it suggests that we imagine the vignette, via daydreaming and imagery (Williamson (2009: 467)).

normalcy approach, by contrast, captures how we interpret both fictional and non-fictional vignettes.

Fourthly, the proposal seems too broad: fictions, and our interpretation of fictions, can be inconsistent or conceptually incoherent. Ichikawa and Jarvis suggest that philosophical vignettes are a special genre of fiction, one that ensures no inconsistent or incoherent fillings-out obtain.⁶⁴ They thereby narrow down the content of the judgement. But by adding these kinds of restrictions, the ‘truth-in-fiction’ operator begins to resemble simply *understanding* the vignette, rather than a specific kind of fictional interpretation as such.

Thus whilst my view shares features with Ichikawa and Jarvis’s proposal it is less wedded to the notion of truth-in-fiction, and accordingly avoids their proposal’s weaknesses. The supplementary details generated by the normal understanding of the vignette ‘fill out’ the vignette—and provide more detail than simply $V(x,p)$ —but the normal understanding does not demand as much filling out as the truth-in-fiction proposal.⁶⁵

Judgements about normalcy are related to counterfactual judgements. Indeed, our understanding of what is normal, typical, or generic affects our judgements about counterfactuals, and is plausibly explanatorily prior to them. We can judge counterfactuals in virtue of judgements about normalcy. The normal understanding of a vignette will often approximate what would obtain were the vignette true. The judgements differ only when either the unexpected or abnormal is a modally nearby possibility, such as when—surprisingly—the actual world turns out to be a gullible Orlando world, rather than a world where were the Orlando vignette satisfied, the normal version would obtain. Or where our judgement of what is typical and ordinary is in error, such as when our scientific theories are mistaken.⁶⁶

7.1 Virtues of the Normalcy Proposal

Since whether something constitutes the normal instance of a vignette is not subject to unexpected contingent anomalies in the actual world, the normalcy approach avoids the ‘deviant realisations’ problem faced by Williamson’s counterfactual approach. Even after we have learnt of a deviant actual or modally nearby realisation, we might well disregard it as an abnormal version of the vignette. The abnormal can be actual. This is illustrated by the class discussing the sheep vignette and its variants. What is normal—and our judgements concerning what is normal—has stability. It does not fluctuate as a result of quirky one-off features of the world.⁶⁷ And so the normalcy proposal captures how the potency of a given vignette is not hostage to random quirks of the world.

What qualifies as a normal version of a vignette is somewhat stable. It is not contingent on fluky isolated facts about the world, as suggested by Williamson’s counterfactual approach. But it will be influenced by, for example, changes in systematic and

⁶⁴ Ichikawa and Jarvis (2013: 211, n. 25).

⁶⁵ For further criticisms of the truth-in-fiction proposal, see Williamson (2009) and Malmgren (2011).

⁶⁶ In section 7.4 I return to the question of to what extent the truth conditions of the ‘normally’ operator are constrained by reality, and to what extent they are influenced by our beliefs, including false beliefs.

⁶⁷ This is not to say that we can never be wrong about which vignettes challenge the target theory. (Williamson (2009) emphasises that we must countenance the fact that we can be wrong about this kind of philosophical reasoning, and accuses Ichikawa and Jarvis’s view of having the result that we are too often correct.) We might think that a normal version of a vignette exhibits justified true belief without knowledge, but be mistaken.

familiar features of the world. What is normal can change over time. Suppose in the future there is rampant bookseller fraud, for example, perfect replication technology allows fakers to create undetectable forgeries, there is inadequate regulation in the market, and much demand for old books. Were the Orlando vignette told in this future, the normal version of the Orlando vignette would not be one where his belief is justified. In this unfortunate future, the Orlando vignette as stated does not provide compelling evidence against the target theory. (The vignette-teller would at least have to augment the vignette to motivate the thought that Orlando was justified in believing the bookseller. Perhaps, for example, by adding that this bookseller was widely regarded as honest.) In this respect the account captures both the modal stability and the contingency of the vignette method.⁶⁸

Unlike the necessity proposal, the normalcy proposal does not quantify over all instances that satisfy the vignette. It narrows down to only those instances that are normal versions of the vignette. And, as I explain in section 7.4, variations of the normalcy proposal narrow down the quantification even further, to just some normal variations, a normal variation, or the most normal variations. Thus the proposal avoids the ‘deviant variations’ problem that plagued the necessity proposal. The deviant variations all involve abnormal understandings of the vignette. Whenever the vignette constitutes a compelling problem for the target theory, the normal version (or most normal versions) of the vignette exhibit a counter-instance to the target theory.

The normalcy approach also avoids the ‘anemic vignette’ and ‘botched vignette’ problems. These problems afflict the possibility and counterfactual proposals because, according to these views, patently inadequate vignettes could qualify as challenges to the target theory. The normalcy proposal avoids this problem because if the vignette presented fails to be one where the normal version of that vignette is justified true belief without knowledge, then judgement (16) is false about that vignette.⁶⁹ The proposal thus explains why these vignettes could not constitute a challenge to the target theory.

Unlike the possibility proposal, the normalcy proposal does not require only that some variation of the vignette is a counter-instance. It must be a normal instance. Since we are typically first struck by the normal understanding of the vignette, this explains why we are struck by the problem instance, and so why the method is so compelling. Unlike the possibility proposal, which faces trouble explaining the forcefulness of the method, the

⁶⁸ Ichikawa and Jarvis (2009: esp. 223) argue for, and I think overstate, the non-contingency of the success of a given vignette. They write,

It is not right to think of Williamson's project as a defense of a traditional understanding of philosophical methodology, for the judgments with which he identifies “intuitions” are scarcely recognizable as the things traditional philosophers had in mind. In particular, on Williamson’s view, the propositional contents of the intuitions in question are contingent. This is problematic, from a traditional standpoint [...] tradition has it that intuitions like the Gettier intuition have necessarily true contents; Williamson’s counterfactuals, however, will vary in truth value between possible worlds.

See also Ichikawa and Jarvis (2013: chapter nine). But surely the judgement should depend on contingent facts—consider the Orlando vignette told in a world where forgery is widespread. Williamson (2007; 2009) argues for, and I think misattributes, the contingency. He thinks quirky, isolated, unknown features can render a vignette impotent. I think the correct view is that facts about the actual world should only affect the success of a vignette insofar as they influence the truth of claims about normalcy, such as claim (16).

⁶⁹ Or ‘a’ or ‘any’ normal version—I return to the question of quantification in section 7.4.

normalcy proposal can explain why the vignettes are potent: the normal understanding of them presents us with a counter-instance.

Williamson notes that one virtue of his counterfactual view is the continuity between thought experiment judgements and normal everyday judgements.⁷⁰ His view renders thought experiment judgements non-mysterious, because it understands them as a kind of ordinary everyday widely-recognised judgement. The normalcy proposal also enjoys this advantage. Judgements about normalcy are ubiquitous and mundane. They are more familiar to us than judgements about what is a possible or necessary interpretation of a vignette. Normalcy judgements might be even more basic than counterfactual judgements, since plausibly we require the former in order to make the latter.

The normalcy proposal has some similarities with the emendation proposal. The emendation proposal aims to reduce the kinds of variations of the vignette that are relevant to the judgement, and does so by amending the vignette with additional content. The normalcy proposal also restricts which variations of the vignette are relevant, by saying that only the normal variations are relevant. The normalcy proposal, unlike the emendation proposal, is not revisionary about the method. Instead, I posit, it simply describes what people actually do when they hear a vignette. The content is simpler: There was a worry that the experts do not consider the vignettes as described by Grundmann and Horvath. If they did, they would have made the amendments explicit. And there was a worry that the folk do not either: they do not possess the specific conceptual resources to easily think of the vignettes as proposed by Grundmann and Horvath. But laypeople use the vignette method with ease. The content posited by the normalcy proposal is simple. Indeed, considering the normal version of something is a hallmark of normal everyday thinking.⁷¹

7.2 A Unified Account

Unlike Grundmann and Horvath's proposal, the normalcy proposal readily generalises to other vignettes as counterexamples, and other roles that vignettes play in philosophical theorising. Firstly, the emendation suggestion was particular to Gettier cases, and so does not readily generalise to vignettes as counterexamples against other necessity claims. The normalcy proposal readily generalises to other vignettes as counterexamples. But additionally, the normalcy proposal offers a unified account of other roles that vignettes play. Vignettes can motivate claims and principles, illustrate distinctions, communicate ideas, help pin down the phenomena of interest, and frame research questions. And they provide evidential weight against claims that are not, or do not entail, necessity claims. An account of the judgement is better if it unifies these roles. In what follows I illustrate ways the normalcy account can better capture the range of roles that vignettes play.

⁷⁰ Williamson (2007: 192).

⁷¹ Another virtue of the normalcy proposal is that it might explain some cross-cultural variation in response to vignettes. Weinberg, Nichols and Stich (2011) claim, for example, that Asian respondents judge that the protagonist knows in the 'American car' Gettier vignette. Perhaps Asian respondents fill in as a normal background condition that someone who once drove an American car will continue to do so, and so can be known to drive an American car. This background assumption affects their judgement of the normalcy claim. My proposal thus accords with this kind of explanation of cross-cultural variation. See also Sosa (2009).

It is hard to see how the possibility proposal could explain how vignettes help to build a theory, motivate a principle, or illustrate a distinction. If a theorist says that ‘F’s are G’s’, for example, and provides a vignette, it does not much motivate the theorist’s claim if hearers thereby judge only that some possible version of the vignette exhibits an F that is G. Providing a vignette and judging that in *normal* instances of the vignette F’s are G’s is more compelling. It is a good starting place for developing a theory.

Similarly it is hard to build a theory or begin to motivate claims if vignettes, to be potent, must be such that every variant is an instance of an F that is a G. The necessity proposal, by requiring this, renders vignette construction too demanding to play theory-building roles.

Whilst the discussion thus far focuses on vignettes as a challenge to necessity claims, vignettes are also employed to challenge claims about what is typical, normal, characteristic, generic, or usual. A theorist might claim, for example, ‘typically F’s are G’s’, and we might put pressure on that claim by providing an everyday, mundane vignette in which F’s are not G’s. The more everyday and normal the vignette, the better. Suppose a theorist claims that ‘typically testimonial beliefs from strangers are unjustified’. An interlocutor replies with a vignette about a man, Fred, visiting Manhattan who politely asks for directions to the Flatiron Building. The woman he asks, Jenny, is a local and confidently provides directions. This challenge to the target claim is successful, in large part because it is an ordinary, mundane circumstance, and the ordinary, typical instances of the vignette exhibit justified belief. A vignette about an unusual occurrence would be less successful against the target claim. If the vignette stipulated that the woman was Manhattan’s chief cartographer and lives in the Flatiron building, the vignette would be less successful against the target claim. And if the initial vignette elicits the judgement only that some possible variation of the vignette — such as the cartographer augmentation—exhibits justified belief, this is also less potent against the target typicality claim.

I submit that this kind of dialectic—where a vignette challenges a claim about what is generic, characteristic, normal, or typical—is common in philosophy. But the existential proposal has trouble vindicating it: Some possible variant of a vignette’s being a counter-instance would not harm the target claim. In order to challenge a claim about what is typical, generic, normal, or characteristic, the vignette-based judgement must concern a normal instance of a normal vignette.

There is some debate in philosophy about whether philosophical inquiry concerns conceptual facts, metaphysical facts, some other kind of fact, or a combination. This meta-philosophical debate concerns, for example, whether we are investigating what knowledge is or the concept of knowledge. Williamson notes that his favoured proposal, the counterfactual proposal, is ill-suited where philosophy claims concern conceptual, rather than metaphysical, possibility.⁷² He subsequently argues that the metaphysical question is of interest to philosophers, and the conceptual question is not. One virtue of the normalcy approach is its applicability to both families of questions. If the target claim (3) concerns the metaphysical nature of knowledge, we can investigate it by thinking of the typical, normal versions of vignettes. If the target claim concerns the concept *knowledge*, we can also investigate it by thinking of normal versions of the vignette, since what we take to be normal

⁷² Williamson (2007: 205–7).

is influenced by our conceptual scheme; our judgements about normalcy provide a window into our concepts.

7.3 Resistance to the Proposal

There has been resistance to this kind of proposal. Williamson discusses the problem of deviant variations that plagues the necessity proposal. He says in a footnote,⁷³

Merely adding the stipulation that *x* and *p* constitute a normal instance of the Gettier case is unlikely to solve the problem, for the relevant notion of normality is an epistemological one that violates the supposed neutrality of the initial description of the case.

By ‘neutrality’ he means that in telling the vignette we do not bias the judgement with regard to the target claim. In the case of Gettier vignettes, for example, it is not made explicit whether the protagonist knows or is justified. Williamson writes,⁷⁴

Suppose that we fix on a particular Gettier-style story (the one about Orlando would do), henceforth ‘the Gettier case’, told in neutral terms, without prejudice to the target analysis. For instance, it is not explicitly part of the story that Orlando does not know that he owns a book which once belonged to Virginia Woolf.

Malmgren uses the term ‘neutrality’ in a similar way. She writes,⁷⁵

[A vignette] exhibits a certain neutrality—it does not specify that the subject does, or does not, know the relevant proposition; nor does it specify that she does, or does not, have a justified true belief in it. (But note that it *does* specify that she truly believes it.) More generally, a canonical description of a problem case is neutral with respect to the distribution of the ‘test properties’—the properties whose modal connection is at issue in the thought experiment: it does not explicitly stipulate that the test properties are or that they are not instantiated in the given case (nor does it stipulate anything that transparently entails that they are/are not instantiated).

It is not clear why asserting that it is a normal case of the vignette violates neutrality understood in this way. Calling it a normal instance of a vignette does not in general suggest whether the person has a justified belief or knowledge. It is still left to the hearer to determine these facts.⁷⁶

Malmgren also considers and dismisses a view similar to the normalcy proposal. Like Williamson, she reflects on the problem of deviant variations for the necessity proposal. She considers weakening the content of the necessity proposal, whilst retaining the necessity operator, by specifying that the relevant subset of cases is narrower than the set of cases where $V(x, p)$. Malmgren writes,⁷⁷

There are different ways to elaborate on this suggestion, but the basic thought is just to weaken [the judgement proposed by the necessity proposal, (11)] by restricting the scope of the necessity operator to

⁷³ Williamson (2007: 204, n. 22).

⁷⁴ Williamson (2007: 184). This is all Williamson says about neutrality in Williamson (2007).

⁷⁵ Malmgren (2011: 274–5, emphasis in original).

⁷⁶ Note that Grundmann and Horvath’s proposal does violate neutrality with regard to whether the belief is justified. Sosa (2006: 642) suggests neutrality may not be an important feature of the vignette method.

⁷⁷ Malmgren (2011: 285, emphasis in original).

some specific subset of possible worlds that satisfy the given case description [...] Let us use ‘the *intended* Gettier case’ as a placeholder name for that subset of worlds. We can then express the suggestion by saying that the content of the Gettier judgement is the claim that,

- (18) Necessarily, anyone who stands to a proposition *p* as in the intended Gettier case has a justified true belief that *p* but does not know that *p*

Malmgren then criticises this idea,⁷⁸

First, (18) needs to be spelled out in more detail—in particular, the placeholder (‘the intended case’) must be eliminated—and this turns out to be extremely difficult [...] As stated, (18) is way too schematic; we cannot properly evaluate it [...] We need to be provided with a more informative characterization of the set of worlds that makes up the intended Gettier case [...] a characterization that captures all and only non-deviant realizations of the Gettier case, but that does not itself contain terms like ‘deviance’ and ‘intended’ [...] But it turns out to be very difficult to provide a characterization that meets these constraints.

Malmgren argues that formulation (18) must be explained more. She argues that advocates of (18) must translate ‘intended’ into new terms, and may not use terms such as ‘intended’ or ‘deviant’. Presumably ‘normalcy’, ‘typicality’, ‘generically’ and similar terms are also prohibited. Malmgren argues that this constitutes a too difficult task, and resembles the problem of *ceteris paribus* clauses.⁷⁹

I agree the formulation demands elaboration. An advocate of the view would need to say more about the notion ‘intended case’, such as what kinds of details are included in the intended case that are absent from the articulated vignette itself, and whether the intended case can include inconsistencies or conceptual incoherence. But I disagree that this elaboration requires an elimination of the ‘intended case’ placeholder where the illumination cannot include terms such as ‘deviant’, ‘normal’, ‘intended’ etc.

Even if we were committed to thinking those terms cannot be used in descriptions of natural phenomena, such as in chemical and physical theories, the content question concerns how we *think*.⁸⁰ Thinking often uses ideas such as deviance, intended, normally, typically, *ceteris paribus*, generically, and characteristically. We do not need to reduce or eliminate these notions in order to think well. These ideas are successfully employed in thought; understanding things as intended, normal or deviant instances of a kind is an important and apparently irreducible aspect of thought. It would be misguided to attempt to capture the content of thought in general—and the responses to vignettes in particular—without employing these notions.

I do not think ‘intended’ is the right modifier, and so do not endorse (18) as the correct content. This is because I do not think speaker intentions are dispositive about how

⁷⁸ Malmgren (2011: 286–7).

⁷⁹ Malmgren (2011: 289).

⁸⁰ I expect that it is also mistaken to try to understand many other phenomena without using these notions. Arguably we cannot understand causation without these kinds of notions. Perhaps too we cannot understand many natural phenomena, such as biological and psychological kinds, without these notions. Thus we should not be averse to them within science, let alone in an account of human thought; if we should allow these kinds of notions within scientific theories, they should certainly be permitted in theories about the content of responses to vignettes.

we should understand the vignettes.⁸¹ But I think that Malmgren's reasons for rejecting the view are unconvincing.

Similarly Ichikawa and Jarvis express reservations about this line of thought. They write,⁸²

Richard Heck has proposed to us in conversation another move on behalf of Williamson. Instead of a counterfactual, one could render [the judgement] as a generic like 'Fire engines are red,' which is true not because all fire engines are red (we've seen some green ones), but because fire engines typically are red. [The judgement] then becomes: in cases that satisfy the Gettier text, people (typically) have justified true belief without knowledge. The conclusion we arrive at, then, is that typical cases that satisfy the Gettier text are ones in which people have justified true belief without knowledge. We find this approach plausible, although we're hesitant to rest too much on such generics, as their interpretation is a vexed issue [...] It may be that this approach is ultimately similar to the one we will go on to develop.

I am sceptical that the fact that generics are a vexed issue provides much reason to avoid using them in an account of the content of our judgements. Firstly, we are able to reason well using generics, *ceteris paribus* clauses etc., despite lacking a full account of them. And the project is to provide an account of *reasoning*. We do not, in general, need an account of something for it to play a successful role in thought. Secondly, people are not sceptical about the existence of something akin to generics. I can understand avoiding positing something where there is some contention about its existence. But in this case theorists typically assume that generics exist but debate their semantics. Thirdly, often in philosophy we make use of ideas in our theories without fully having an account of them. We commonly use notions such as luck, belief, causation, and truth, despite their natures being controversial. Provided we understand them well enough for our purposes, and are not accidentally using an incoherent or misleading notion, this is unproblematic.

Thus whilst there has been some resistance to this kind of suggestion, I do not think this resistance is well-grounded.

7.4 Decision Points

The normalcy proposal as stated generates some choice points—places where differing theories can be developed from the basic idea. Perhaps the most obvious concerns the relevant quantification. The normalcy proposal demotes the importance of which situations are literally satisfied by the vignette, and instead focuses on how we *understand* vignettes. We understand them via their normal and typical instances. When I introduced the normalcy proposal I used universal quantification, and so offered an account that—in that regard—resembles the necessity proposal. I characterised the normalcy proposal as:

- (16) Normally, whenever a person is related to a proposition such that the vignette obtains, then that person will have JTB without knowledge.

⁸¹ There are many reasons to think speaker intentions should not play such a large role. One is that if the speaker botches recounting a vignette, so that the vignette does not describe justified true belief without knowledge, the recounted vignette should not qualify as successful just because the vignette that the speaker intended to tell would have exhibited justified true belief without knowledge.

⁸² Ichikawa and Jarvis (2013: 204, n. 12).

$$(17) \text{ NORM } \forall x \forall p [V(x, p) \rightarrow (JTB(x, p) \& \neg K(x, p))]$$

But universal quantification is not essential to the view. The view could be understood using something akin to the existential quantification of the possibility approach, by positing that the judgement is ‘some normal variations of the vignette exhibit justified true belief without knowledge’. According to this proposal, ‘normally’ is understood as described above, but the judgement is true even if some normal variations of the vignette are not counter-instances to the target claim.

A third alternative uses the formulation ‘most normal variations of the vignette exhibit justified true belief without knowledge’, where ‘most’ is understood quantitatively, i.e. meaning ‘more than half’ or some high proportion. This is akin to formulations that employ the possible worlds framework and say ‘most nearby possible worlds’. We might instead posit that the judgement is ‘The most normal variations of the vignette exhibit justified true belief without knowledge’, where ‘the most normal’ is understood qualitatively. According to this variation of the view, versions of the vignette can be ordered, with some being more normal than others, and the ones relevant to the judgement are the most normal ones. Or we might posit that the judgement is simply ‘normal variations of the vignette exhibit justified true belief without knowledge’, where the advocate proposes that we have a good grasp of what this means independently from a detailed spelling out.

Perhaps too advocates of the counterfactual proposal can adopt a subjunctive version of the normalcy approach. The proposed content would be something like, ‘were the vignette satisfied normally, x would have a justified true belief, but not know, that p’. Employing Lewis’s possible worlds framework for interpreting counterfactuals yields, ‘in the nearest possible worlds in which the vignette is satisfied in a normal way, x has a justified true belief that p, but does not know p’. This proposed content uses the structure of the counterfactual proposal, but content from the normalcy proposal.

The appropriate quantification depends partly on how we understand ‘normal’. If ‘normal version of a vignette’ denotes a narrow range of instances, then ‘every normal variation’ will be a more appropriate quantification. Conversely if ‘normal version of a vignette’ allows wider variation, then something like the quantitative ‘most normal variations’ or the qualitative ‘the most normal variations’ might be more appropriate.

One strategy to adjudicate amongst these quantification options is to consider the unsuccessful vignettes. If a vignette does not provide a challenge to the target theory simply because some normal variations of the vignette are not counter-instances, this suggests that the universal quantification ‘all normal variations’ plays a role in our judgement. If the vignette only fails when no normal understanding of the vignette provides a counter-instance, this suggests that the relevant judgement is ‘some normal variations’. And if the vignette fails when some normal variations, but not the most normal variations, constitute a counter-instance, this suggests the judgement is something like ‘the most normal variations’, which picks out the qualitatively most normal.

A second decision point concerns what to say about atypical vignettes. If a vignette is too weird, plausibly there is no normal version of the vignette.⁸³ This raises the question of how judgment (16) applies to ‘far out’ vignettes. One route maintains that the most normal versions are relevant, even if they are still somewhat abnormal. This is akin to when modally

⁸³ Thanks to Bob Beddor, David Black, Will Fleisher, and Ernest Sosa for discussion of these issues.

distant possibilities are raised in the antecedent of subjunctives, and so we must consider very distant possibilities to evaluate the subjunctive.⁸⁴ We might evaluate what would normally happen if a person were a reliable but unreflective clairvoyant, for example, even whilst recognising that this is a deeply abnormal situation. An alternative route maintains that if there is no normal understanding of the vignette, the vignette cannot present a challenge to the target theory. Advocates of this approach might argue that weird vignettes cannot be well understood, and so lack dialectical force. Alternatively one could advance some combination of strategies: the judgement concerns the most normal version of the vignette, even if somewhat abnormal, until a threshold after which, since there is no sufficiently normal version of the case, judgement (16) cannot be true of the vignette and so the vignette does not pose a challenge to the target theory.

This decision will be influenced, in part, by whether employing weird cases should be vindicated by an account of the method. If we independently think weird vignettes are methodologically suspect, then the normalcy proposal provides resources for explaining why.

Another decision point concerns to what extent, and in what ways, ‘normalcy’ is influenced by our epistemic position, and how much it is simply a function of how the world in fact is. To illustrate this distinction, consider this case,

Sarah is teaching an epistemology class. She intends to explain how knowledge and justified true belief can diverge and tells a vignette. She says,

‘Joey walked into a room. He saw something lush and verdant in a pot and so formed a belief that there was a plant on the table. In fact what he looked at was an abstract sculpture that resembles a plant. Unbeknownst to him, however, there was a mushroom growing in a pot behind the pile of books on the table. And so by chance, Joey’s belief was true.’

Sarah’s class are impressed. They all judge that knowledge and justified true belief can indeed diverge, and proceed to have an intelligent and thoughtful conversation about different ways to modify an account of knowledge to overcome the problem posed by the vignette.

Here is the rub: Sarah—the vignette-teller—is mistaken. Mushrooms are not plants.⁸⁵ But no one in the class knew; few people know mushrooms are not plants. (We can vary the vignette: Sarah could have made a mistake that *nobody*, scientists included, realises is mistaken.) So, is Sarah’s vignette a successful Gettier case?

If the answer is no—Sarah’s case does not provide a successful challenge to the target theory of knowledge—this suggests that what is normal is an objective, mind-independent feature of the world. There is no normal variation of the vignette where Joey has a justified true belief without knowledge, because there is no normal version of the vignette where his belief is true. (‘How can it be normal if it is biologically impossible?’, the thought goes.)

If the answer is yes, this suggests that a normal version of a vignette can be in part determined by what we *take* to be normal. (Perhaps ‘normal’ is understood as something epistemic, such as ‘normal understanding’.) The thought might continue that it is a

⁸⁴ Lewis (1973).

⁸⁵ Phylogenetically, mushrooms are closer to humans than to plants.

persuasive vignette, the students are not being irrational, and if everyone thinks that mushrooms are plants that is sufficient for legitimating the dialectic.⁸⁶

Perhaps both views are partially correct: perhaps Sarah's vignette is problematic in some ways, and would be improved if it reflects biological fact. But the vignette is not valueless; it has some dialectical force. Perhaps a correct understanding of 'normalcy', reflecting both objective and subjective factors, can capture this.⁸⁷

Thus there is some flexibility in the view advanced. I hope to have motivated the view that the normalcy proposal merits further investigation and development. I have explained how the view avoids the problems faced by competing proposals, and have articulated its virtues. The normalcy proposal provides a unified account of the role that vignettes play in philosophical theorising, and captures how we normally think. I suggest the normalcy view best captures the content of the operative judgements we have in response to vignettes.⁸⁸

References

- BonJour, Laurence (1980). 'Externalist Theories of Empirical Knowledge' *Midwest Studies in Philosophy* 5: 53–73.
- Caton, Jacob (m/s). 'Is "Justification" an Ordinary Term?'
- Chisholm, Roderick (1977). *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice-Hall
- Constant, Benjamin (1797). 'On Political Reactions' *France* 5 (1).
- France*, for 1797, Part VI, No. i, page 123, in an article entitled "On Political Reactions" by Constant,
- Gardiner, Georgi (2015). 'Teleologies and the Methodology of Epistemology' in *Epistemic Evaluation: Purposeful Epistemology*, eds. J. Greco and D. Henderson. Oxford: Oxford University Press.
- Gettier, Edmund (1963). 'Is Justified True Belief Knowledge?' *Analysis* 23: 121–3.
- Grundmann, Thomas and Joachim Horvath (2014). 'Thought Experiments and the Problem of Deviant Realizations' *Philosophical Studies* 170: 525–33.
- Ichikawa, Jonathan (2009). 'Knowing the Intuition and Knowing the Counterexample' *Philosophical Studies* 145: 435–43.
- Ichikawa, Jonathan and Benjamin Jarvis (2009). 'Thought-experiment Intuitions and Truth

⁸⁶ If what is normal is partly a function of expectations, this raises the question of how socially determined the expectations and normalcy operator are, versus how private and individual. It should not be wholly private, I think. A person narrating an abysmal or botched vignette to himself should not qualify as a challenge to a target theory just because he takes normal versions of it to exhibit a counter-instance. Even if 'normalcy' is partly a function of expectations and beliefs, there should be some objective or social constraints on what qualifies as a normal instance of a vignette.

⁸⁷ A related decision point is to what extent, and in what ways, the normal version of a vignette is influenced by speaker intention. It is an open question in the philosophy of fiction whether, and how, the intentions of the author affect what is true in a given fiction. See also footnote 81.

⁸⁸ Many thanks to Bob Beddor, David Black, Will Fleisher, Jon Garthoff, Jonathan Schaffer, Susanna Schellenberg, and Ernest Sosa for comments on an earlier draft of this paper. Thanks also to Frances Egan, Thony Gillies, John Greco, Lisa Miracchi, and Daniel Rubio for helpful discussion, and to audiences at Rutgers University and at the Central European University.

- in Fiction' *Philosophical Studies* 142: 221–46.
- (2013). *The Rules of Thought*, Oxford University Press.
- Jackson, Frank (1982). 'Epiphenomenal Qualia' *Philosophical Quarterly* 32: 127–36.
- Kitcher, Philip (2000). 'A Priori Knowledge Revisited' in *New Essays on the A Priori*, eds. P. Boghossian and C. Peacock. Clarendon Press.
- Kripke, Saul (1980). *Naming and Necessity*, Harvard University Press.
- Lehrer, Keith (1990). *Theory of Knowledge*, Westview Press.
- Lewis, David (1973). *Counterfactuals*, Harvard University Press.
- Malmgren, Anna-Sara (2011). 'Rationalism and the Content of Intuitive Judgements' *Mind* 120: 263–327.
- Pritchard, Duncan (2012). 'Anti-Luck Virtue Epistemology' *Journal of Philosophy* 109: 247–79.
- Sosa, David (2006). 'Scepticism About Intuitions' *Philosophy* 81: 633–48.
- Sosa, Ernest (2009). 'A Defense of the Use of Intuitions in Philosophy' in *Stich and His Critics*, eds. D. Murphy and M. Bishop. Wiley-Blackwell.
- Weinberg, Jonathan, Shaun Nichols and Stephen Stich (2001). 'Normativity and Epistemic Intuitions' *Philosophical Topics* 29: 429–60.
- Williamson, Timothy (2005). 'Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking' *Proceedings of The Aristotelian Society* 105: 1–23.
- (2007). *The Philosophy of Philosophy*, Blackwell.
- (2009). 'Replies to Ichikawa, Martin and Weinberg' *Philosophical Studies* 145: 465–76.