



Knowledgeably Responding to Reasons

Joseph Cunningham¹

Received: 6 October 2017 / Accepted: 10 July 2018 / Published online: 25 July 2018
© The Author(s) 2018

Abstract

Jennifer Hornsby has defended the Reasons-Knowledge Thesis (RKT): the claim that Φ -ing because p requires knowing that p , where the ‘because’ at issue is a rationalising ‘because’. She defends (RKT) by appeal to the thought that it provides the best explanation of why the subject in a certain sort of Gettier case fails to be in a position to Φ because p . Dustin Locke and, separately, Nick Hughes, present some modified barn-façade cases which (a) seem to constitute counterexamples to (RKT) and (b) undermine Hornsby’s way of motivating it by rendering their alternative Reasons-Explanation Thesis (RET) a better explanation of Hornsby’s datum. This paper defends (RKT) and Hornsby’s argument for it against those objections. First, I point out that their supposedly intuitive verdict about the relevant barn-façade cases is not as intuitive as they think. Second, I point out that even if we share the intuition: we have strong reason to doubt the verdict anyway. And finally, I point out that since (RET) is independently implausible, the two problems can be tackled anyway.

Jennifer Hornsby (2007a, b, 2008) has defended the following *Reasons-Knowledge Thesis*, where the ‘because’ at issue is a rationalising ‘because’:

(RKT) Necessarily, if $S \phi$ s because p then S knows that p

Hornsby argues for (RKT) by appeal to a certain type of Gettier case which is intended to generate the intuition that the agent in that case does not ϕ because p . Since the agent also does not know that p , an abductive argument for (RKT) is made available. Other adherents of (RKT) include Unger (1975), Hyman (1999, 2006, 2010, 2011, 2015), Marcus (2012), McDowell (2013), and Roessler (2014). Of those, all but Roessler argue for (RKT) in the same way Hornsby does.

A seemingly devastating objection to (RKT) has been raised separately by Hughes (2014) and by Locke (2015). In essence, the objection is that there is a class

✉ Joseph Cunningham
joseph.cunningham@jesus.ox.ac.uk

¹ Jesus College, University of Oxford, Turl Street, Oxford OX1 3DW, UK

of Gettier cases, different from the class on which adherents of (RKT) rely, the members of which constitute counterexamples to (RKT) and serve to motivate the following *Reasons-Explanation Thesis*, a competitor to (RKT) which seems to provide a *better* explanation of Hornsby's original intuition than (RKT) itself:

(RET) Necessarily, if S ϕ s because p then S's belief that p is explained by the fact that p

This paper aims to defend (RKT) against the objection. I will argue that the supposedly problematic class of Gettier cases do not generate the intuitions Hughes and Locke want them to, and that even if they do the relevant problems for (RKT) can be avoided anyway. Section 1 introduces (RKT) and Hornsby's argument for it. Section 2 introduces (RET) and the objection to (RKT). Finally, Sect. 3 defends (RKT) against the objection.

1 The Reasons-Knowledge Thesis

Certain facts constitute reasons for us to act in certain ways and maintain certain attitudes. That my football team has won promotion is a reason for me to celebrate in some way. That the exit-poll predicts a Labour victory at the by-election is a reason to believe that Labour will win. That my friend is having his wedding outdoors is a reason for me to hope for good weather. Those facts are *normative reasons* for me to ϕ , where ' ϕ ' is a stand-in for the relevant action or mental-attitude verb.

An agent might be motivated to ϕ by a fact, in a way that manifests their treating the fact in question as a normative reason for them to ϕ . We can record that state of affairs in the following non-psychologistic way:

($\neg\psi$) S ϕ s because p

The 'because' which appears in the above schema is a rationalising 'because': the kind of 'because' which holds only when the agent's ϕ -ing is the manifestation of an exercise of their capacity for rationality, and hence differs from the merely causal 'because' which appears, for example, in the sentence: 'The leaves on the tree are rustling because there is a strong gust of wind'.¹

I will call cases in which the ($\neg\psi$) condition holds *good cases*. We should also acknowledge the existence of *bad cases*. Bad cases are cases in which the agent *thinks* that there is a fact present, ϕ s in a way that manifests their treating that apparent fact as a normative reason for them to ϕ , but the apparent fact in question does not really obtain. In bad cases, the relevant ($\neg\psi$) statement is not true of the agent. Only an instance of the following psychologistic schema, which records the obtaining of a condition that's neutral on whether the agent is ϕ -ing in the light of a fact, is true of them:

¹ From hereon, unless stated otherwise, it should be taken for granted that what I have in mind when using 'because' talk is the rationalising sense of 'because'.

(ψ) S ϕ s because S believes that p.

How the condition specified by (ψ) statements relates to the condition specified by ($\neg\psi$) statements is an issue that will not concern me here. Instead I want to focus on the following thesis about the nature of the good case in particular, Hornsby's Reasons-Knowledge Thesis:

(RKT) Necessarily, if S ϕ s because p then S knows that p.

According to (RKT), it's not just that ϕ -ing in the light of the fact that p requires believing that p: that belief must have a positive epistemic status, and nothing short of the status of *knowledge* will do.

Here is another way to think about (RKT). It's trivial that ϕ -ing in the light of the fact that p requires it to be a fact that p. Also, it's plausible that ϕ -ing in the light of the fact that p requires believing that p. But presumably, there will need to be some connection between the fact that p and the agent's belief that p, in such cases. (RKT) can be seen as one way of cashing out that connection: it consists in the agent's belief that p being an article of knowledge that p.

How might one argue for (RKT)? Hornsby (2007a, b, 2008) defends the thesis by appeal to a certain Gettier case, adapted from Ryle (2000). I will give Locke's (2015) version of the case, which is the same as Hornsby's in all substantive respects,² for it is his discussion of it which will occupy me later on:

Thin Ice... Edmund has just asked his mother whether the ice on the pond is thick enough to skate. His mother tells him that it is too thin, and indeed it *is* too thin. However, unbeknownst to Edmund, his mother believes that the ice is thick enough to skate, but is trying to trick Edmund into thinking that the ice is too thin. On the basis of his mother's testimony, Edmund believes that the ice is too thin. And because he believes this he stays off the ice. (Locke 2015: 222)

Edmund has a true belief that the ice on the pond is too thin to skate, and he refrains from skating on the pond in a way that manifests this belief, as well as manifesting his treating what he believes as a normative reason for him to refrain from skating. His belief is acquired off the back of his normally reliable mother's testimony, and so is justifiably held. Nevertheless, Edmund does not know that the ice is too thin to skate: the case is a Gettier case. Intuitively, Edmund does not refrain from skating because the ice is too thin: he does not refrain in the light of the fact that the ice is too thin. According to Hornsby, that's precisely because he does not *know* that the ice is too thin. His lack of knowledge, says Hornsby, is the best explanation of why he fails to be in a position to refrain from skating *because* the ice is too thin. So it must be that that sort of 'because' can hold between the agent's ϕ -ing and the fact that p only if the agent knows that p: (RKT) is true.

What we have here is a case in which the agent has a justified true belief that p but does not know that p and where it is an intuitive datum that they do not ϕ because p. The argument for (RKT) is an abductive inference from that intuitive datum to (RKT). As I mentioned in the introduction, most of (RKT)'s supporters

² Hornsby's own version of the case is given, for example, in Hornsby (2008: 251).

argue for the thesis in exactly the same way, by appeal to cases of the same structure. McDowell (2013: 17) relies on Hornsby's own thin ice case. Unger (1975: 209–210) has his own case involving rainfall, Hyman (1999: 447) relies on a tennis example, and Marcus (2012: 37–38) relies on still further examples. The alternative cases here are interesting, but to keep things simple I will ignore them and conduct the discussion solely in relation to *Thin Ice*. It would make little difference were one of Unger, Hyman or Marcus's cases selected instead.³

2 The Reasons-Explanation Thesis

Hughes (2014) and Locke (2015) have separately raised what I think is the same objection to (RKT) and to Hornsby's argument for it. I will focus on Locke's more detailed discussion, taking it for granted that what I say by way of exegesis in the current section and criticism of it in the next applies equally well to Hughes' discussion.

Hornsby focuses on a kind of Gettier case in which, as in Gettier's own cases, there is no explanatory connection of any kind between the fact that *p* and the agent's justified true belief that *p*. Edmund truly believes that the ice is thin, and he believes this because he's been told it by his mother. His mother is normally reliable, but on this occasion she believes the opposite of what she tells him, so that the fact that the ice is thin does not in any way explain why she tells him that it is. Because the fact that the ice is thin is not explanatorily connected to his mother's telling Edmund that the ice is thin, and Edmund's belief is explained by his being told so by his mother, Edmund's belief is itself explanatorily unrelated to the fact to which it corresponds: the fact that *p* is no part of the explanation, either distal or proximate, of why Edmund believes that *p*.

But there are Gettier cases in which there *is* an explanatory connection between the fact that *p* and the agent's belief that *p*: cases in which the agent has a justified true belief that *p*, they believe that *p* in a way that is to be explained by the fact that *p*, but in which they do not know that *p*. Barn façade cases, in which the agent's belief that *p* is causally explained by the fact that *p*, are supposed to fit that description and it is precisely cases with the structure of barn façade cases which Locke thinks cause a problem for Hornsby. Locke presents us with the following variant of *Thin Ice*, modified so that it exemplifies the structure of a standard barn façade case:

Fake-mother Invasion. Edmund has just asked his mother whether the ice on the pond is thick enough to skate. His mother tells him that it is too thin, and indeed it *is* too thin. However, unbeknownst to Edmund, his home has just been invaded by aliens disguised as perfect duplicates of his mother. If Edmund had spoken with any one of these aliens, he would have been told that the ice is perfectly safe. Just by chance, Edmund is talking to his actual

³ It is interesting to note that in Hyman's work, (RKT) forms the basis of a distinctive account of knowledge itself, an account according to which we are to think of knowing that *p* as possessing the ability to ϕ for the reason that *p*.

mother, who tells him that the ice is too thin, and she tells him this because she inspected it this morning and found it to be so. On the basis of his actual mother's testimony, Edmund believes that the ice is too thin. And because he believes this, he stays off the ice. (Locke 2015: 223)

In this variant of the case, it is supposed to be intuitively plausible that Edmund has a justified true belief that p , has a belief which is to be explained by the fact that p , but does not know that p . He does not know that p , this time, not because his mother is out to deceive him, but because, given the presence of the imposters who are disposed to tell him that not- p , he could too easily have believed falsely (relative to the same kind of method). It is also supposed to be intuitively obvious however that he *does* refrain from skating on the pond because the ice is too thin.

This generates two problems for the proponent of (RKT) which it is worth separating. First, we have the problem that we seem to be confronted with a straight-up counterexample to (RKT). *Fake-mother Invasion* is a case which (supposedly) involves an agent who ϕ s because p without knowing that p . Cases of an analogous structure are also supposed to constitute counterexamples.⁴

But there is a second problem generated by the relevant class of cases. To get a grip on this second problem, let's consider Locke's Reasons-Explanation Thesis:

(RET) Necessarily, if S ϕ s because p then S 's belief that p is explained by the fact that p .⁵

(RET) presupposes that if S ϕ s because p , then S believes that p . It then places a restriction on the sorts of belief that can play the relevant sort of role: that the reason why S believes that p is the fact that p . Paradigm cases in which the agent's belief that p is explained by the fact that p include cases in which the agent perceives the state of affairs which corresponds to p and thereby believes that p ; the agent believes that p because they remember that p ; the agent believes that p because someone who knows that p has told them that p ; and cases in which they have come to believe that p by introspection. We should also add barn façade cases, and cases with the same structure, such as *Fake-mother Invasion*, to the list.

As should be apparent from the set of illustrative examples just selected, for the fact that p to explain why the agent believes that p in the relevant sense is not for the agent's belief that p to be *wholly explained* by the fact that p : it suffices that the fact is *part of the explanans* of the relevant explanation. Moreover, there is no requirement that the fact be a part of the *proximate explanation* of the agent's belief. In the case of perception or memory, for example, it might be that the fact causally explains why the agent has an experience or memory, thought of as a state of mind which is constitutively independent of the relevant fact, and the experience or memory so conceived itself proximately explains the agent's belief that p . Finally, we have been focusing our attention on a certain sort of *rationalising* explanation in

⁴ Locke also offers us barn façade-style variants of Unger and Hyman's cases. Hughes (2014: 461) offers a barn façade variant of Hyman's tennis case too, as well as a fresh case involving a hiker's hut. My response to Locke's variant of *Thin Ice* is just as effective against these other cases.

⁵ Locke (2015: 225) himself labels the thesis *The Explanatory-Representation View* and formulates it in a slightly different way, though the differences do not matter for my purposes.

this paper—the sort of rationalising explanation paradigmatically specified using a $(\neg\psi)$ statement—but (RET) doesn't require that the explanatory link between belief and fact be a rationalising link. That would be an extra claim that goes beyond what (RET) says. (RET) simply requires that there be an explanatory link of some kind.

(RET) is, strictly speaking, consistent with (RKT). One could say *both* that knowledge that p is required for ϕ -ing because p *and* having a belief that p that's explained by the fact that p is required as well. However, Locke intends his (RET) to be a competitor to (RKT), and given something I said about the latter above it is indeed natural to read (RET) in that way. What I said above was that ϕ -ing because p plausibly requires a belief that p which is in some way connected to the fact that p , and that (RKT) is one way of cashing that thought out. But it is quite natural to read (RET) as an alternative way of cashing out the relevant thought. Like Locke, then, I will treat (RET) and (RKT) as in competition.

I'm now in a position to bring the second problem for (RKT) into focus. What we have are two claims about the nature of the good case which are in competition: (RKT) and (RET). The former is to be motivated by appeal to an abductive inference, the premise of which is that Edmund in *Thin Ice* does not ϕ because p . However, with *Fake-mother Invasion* on the scene and (RET) introduced, that abductive inference begins to look problematic. (RET) can explain why Edmund fails to ϕ because p in *Thin Ice*: his belief that p is not explained by the fact that p . But it can *also* explain why Edmund *is* in a position to ϕ because p in *Fake-mother Invasion*: his belief *is* explained by the fact that p . (RKT) cannot, however, explain our intuitive verdict about the latter, for the simple reason that it is inconsistent with that verdict. This undermines the original motivation for (RKT) because, in the light of all this, (RET) is revealed to have more explanatory power than (RKT): it can explain our intuitions in a greater range of cases than can (RKT). (RKT), then, cannot be the *best* explanation of why Edmund cannot refrain from skating because the ice is too thin in *Thin Ice*: (RET) is the better explanation of our intuition in response to that case.

To ease the discussion let us label and codify these two objections:

The Counterexample Problem. *Fake-mother Invasion*, and structurally analogous cases, are counterexamples to (RKT).

The Motivation Problem. (RET) provides a better explanation of Hornsby's intuitive datum that the agent doesn't ϕ because p in *Thin Ice*, and in structurally analogous cases, than does (RKT). This renders (RKT) unmotivated.

The next section defends (RKT) against each objection at length.

3 The Reasons-Knowledge Thesis Defended

This section defends (RKT) and Hornsby's argument for it by supplying three separate ways in which the two problems identified can be successfully tackled. The three responses are supplied in order of ascending strength: each successive

response grants more to Locke than the last, whilst still showing that the two problems for (RKT) can be successfully undermined nevertheless. First, Sect. 3.1 argues that Locke's claim that Edmund ϕ s because p in *Fake-mother Invasion* lacks intuitive appeal. This stops the pair of problems so much as getting off the ground. Second, Sect. 3.2 grants to Locke the intuitiveness of the claim that Edmund is in the good case, but goes on to provide (RKT)-independent grounds for doubting that claim, so that even if we accept that it is an intuitive thing to think, we are at best warranted in suspending judgement over its truth. This, again, serves to stop the pair of problems so much as getting off the ground. Finally, Sect. 3.3 grants to Locke the intuitiveness of his claim that Edmund is in the good case *and* grants that there are no (RKT)-independent grounds for doubting it, but goes on to argue that the two problems can be successfully tackled anyway, largely because (RET) is an independently implausible thesis.⁶

Before getting to my response to Locke, however, I want to note that there is a further way of responding to him which I won't be pursuing in any depth here: to argue that Edmund in *Fake-mother Invasion* does in fact *know* that the ice is too thin. I am not altogether unsympathetic to this move.⁷ However, in order to make my main line of response to Locke attractive to as many philosophers as possible, I will not pursue it here: for the purposes of this paper, it is granted that agents in barn façade-style cases lack knowledge.

3.1 Is it Intuitive that Edmund ϕ s Because p?

About *Fake-mother Invasion* and cases with the same structure, Locke finds it intuitive that the agent ϕ s because p: he finds it intuitive that Edmund refrains from skating on the pond in rational response to the fact that the ice is too thin. If it turns out *not* to be intuitive that Edmund keeps off the ice because it is too thin, in the rationalising sense, then the Counterexample and Motivation Problems do not get off the ground: each of those problems relies on the thought that it is *intuitively plausible* that the Edmund of *Fake-mother Invasion* is in the good case of rationally ϕ -ing, after all.

In this sub-section I want to argue that it is *not* intuitive that the Edmund of *Fake-mother Invasion* ϕ s because p. It should be noted that this falls short of demonstrating that it is *false* that Edmund ϕ s because p: from the fact that a claim lacks intuitive support, it does not follow that that claim is false. Still, as I have just pointed out, demonstrating that the claim that the Edmund of *Fake-mother Invasion* ϕ s because p lacks intuitive appeal is enough for my dialectical purposes.

My strategy for proving that it is not intuitive that Edmund ϕ s because p is this. First, I suggest that we acknowledge a distinction between ϕ -ing because p in the sense we're interested in, and ϕ -ing in a way that is caused by a fact which is treated by the agent as a normative reason for them to ϕ , but where their ϕ -ing does not manifest their treating it as such. Second, I suggest that with this distinction made

⁶ For a discussion of these issues consonant with, though independent from, my own, see Hyman (2015: 152–158).

⁷ Compare Hawthorne and Gendler (2005) and Lycan (1977, 2006).

explicit, we'll find that it is not intuitively clear that Edmund ϕ s because p , in the required sense, after all.

Let us begin with the distinction. What is of interest to us are rationalising explanations of the form: *S ϕ s because p* . We can contrast these rationalising explanations with the following three sorts of non-rationalising explanation.

First, we are familiar with cases in which a non-agential, natural event is causally explained by the occurrence of a second non-agential natural event: *the grass is wet because a rainfall occurred last night* provides just such an explanation. Clearly, these are not rationalising explanations, not least because the event to be explained is not agential.

Second, there are merely causal explanations of why an agent performs some action, engages in some omission, or holds some attitude possible, but which don't do so by appealing to a fact treated as a normative reason to ϕ by the agent: *my next door neighbour attends church because he finds it comforting* might provide just such an explanation. Although these explanations purport to explain an agential event, they remain non-rationalising precisely because they do so by appealing to some fact which is not treated by the agent as a normative reason for them to engage in the relevant act, omission, or attitude.

Third, there are non-rationalising explanations which function to explain why the agent acts, omits, or holds some attitude in a merely causal way by appeal to a fact which just so happens to be one which is treated by the agent as a normative reason for them to engage in that very act or omission, or to hold that very attitude. This third sort explanation remains non-rationalising because although the fact appealed to is treated by the agent as a normative reason for them to engage in the ϕ -ing at issue, the ϕ -ing in question does not manifest the agent's treating the fact in question as a normative reason for them to ϕ , and hence the explanation operates without appealing to the thought that the agent's ϕ -ing does so.

Let us take some examples of this third sort of non-rationalising explanation. Suppose a chicken sexer comes to believe that the chick is female by dint of sub-personally detecting the presence of a female pheromone being emitted by the chick at that time. Then we could explain why they believe as they do by saying: *the chicken sexer believes that the chick is female because they detect the presence of a female pheromone the chick is emitting*. In providing this explanation, we'd have given a merely causal, or at least non-rationalising, explanation of a certain belief by appeal, we can suppose, to a fact which is treated by the agent as a normative reason for them to adopt that belief. But the belief does not manifest their treating the relevant fact as normative reason; the chicken-sexer's belief is caused by a sub-personal process, after all. Thus, the explanation provided would be non-rationalising. Similarly, suppose one has been invited to do a steep climb up a cliff with one's friends. Knowing that one is not a good climber, one might feel nervous about the climb. That feeling of nervousness might cause one to refuse to climb, in the non-rational manner that emotions can sometimes exert an influence on one's behaviour, when one gets to the bottom of the cliff. We could thereby explain why the agent refuses to climb by saying: *they refuse to climb the cliff because they are feeling nervous*. Feeling nervous is treated as a perfectly good reason to refuse to climb by the agent. But their refusal does not manifest their treating it as a good

reason. Rather, the refusal is non-rationally triggered by the feeling of nervousness. Thus, the explanation provided would count as non-rationalising.⁸

What I want to bring into focus is precisely the distinction between ϕ -ing because p in the rationalising sense of interest to us, and ϕ -ing in a way that's merely caused by some fact which so happens to be one which is treated by the agent as a normative reason for them to ϕ . I want to explicitly distinguish, that is, good cases of rationally ϕ -ing from the third sort of non-rationalising case just described.

With that distinction brought into focus, we are in a position to see the problem for Locke's suggestion that it is *intuitive* that the Edmund of *Fake-mother Invasion* ϕ s because p , in the rationalising sense of interest to us. Let us ask: with respect to *Fake-mother Invasion*, is it intuitively clear that the explanation *Edmund keeps off the ice because the ice is too thin* is really an example of the sort of explanation we're interested in—a ($\neg\psi$) rationalising explanation—as opposed to a merely causal explanation which just so happens to appeal to a fact which is treated by Edmund as a normative reason to refrain as he does? I don't think it is *intuitively* clear that that's so; the case is precisely a case in which there is a causal link between Edmund's belief and the fact, after all.

It should be emphasised that I am not trying to prove that *Fake-mother Invasion* is an instance of the third sort of non-rationalising explanation. That would be to try to demonstrate that it is *false* that Edmund, in *Fake-mother Invasion*, ϕ s because p , which, as I have already made explicit, is not my aim. My point is simply one about the *intuitiveness* of the claim that Edmund's is a good case of rationally ϕ -ing. Let us distinguish good cases of rationally ϕ -ing from the third sort of non-rationalising case described above. And let us now ask: is it *intuitively plausible* that the Edmund of *Fake-mother Invasion* keeps off the ice because the ice is too thin, where this is a rationalising 'because' and hence not the kind of 'because' which holds merely in the third sort of non-rationalising case? I submit that this is not *intuitively* plausible, which is weaker than asserting that *Fake-mother Invasion* is an instance of the third sort of non-rationalising case. Demonstrating that the claim lacks intuitive appeal, however, still suffices to undermine the Counterexample and Motivation Problems.

This completes the argument of this sub-section. But it is worth noting, before we move on, that the problem concerning the intuitiveness of the claim that Edmund is in the good case identified just now is exacerbated for Locke by his own choice of words for describing the intuition he has in response to *Fake-mother Invasion*. Instead of appealing to a ($\neg\psi$) statement, he says: 'Intuitively, Edmund does treat the fact that the ice is too thin as a normative reason' (223). Although Locke intends to capture the phenomenon of ϕ -ing because p , in the rationalising sense of 'because', using such talk (Locke 2015: 219), there is a way of reading it on which it comes out as an intuitively correct thing to say in response to *Fake-mother Invasion*, but doesn't capture the phenomenon of ϕ -ing because p nevertheless. That's because we can read 'S treats the fact that p as a normative reason' *de re*. On this reading, what the relevant sentence says is: 'of the fact that p , S treats *that* as a normative reason'. That might be an intuitive thing to say about the Edmund of

⁸ McDowell's (1996: 163) cyclist case is a further example of this third kind of non-rationalising case.

Fake-mother Invasion, but it would be true merely if a (ψ) statement were true of him and what appears to be a reason to him just so happens to be one: a condition which clearly holds in *Fake-mother Invasion*, but which does not add up to it being true that S ϕ s because p nevertheless.⁹ It would also be true merely if the third sort of non-rationalising explanation canvassed here were to hold of Edmund. Once we're clear that Locke's chosen sentence admits of this *de re* reading which is of no use to him, is it really clear that, read in the way that we require, the sentence is intuitively correct? Again, I don't think so.

3.2 Is it True that Edmund ϕ s Because p?

As it stands, I have suggested that it is not intuitive that the Edmund of *Fake-mother Invasion* is in the good case of rationally ϕ -ing, once attention has been drawn to the distinction between the third sort of non-rationalising explanation and ($\neg\psi$) explanations proper. I have also pointed out that Locke's own formulation exacerbates that worry for him. But let us now *grant* that Locke's claim that Edmund is in the good case has intuitive plausibility. The intuitive plausibility of that claim only supplies *prima facie* warrant for believing it. If there are strong grounds for doubting that Edmund is in the good case, then the warrant to believe that Edmund is in the good case provided by the intuitive appeal of that claim will not warrant that claim *all things considered*; at best, it will leave us in a position where, all things considered, we ought to suspend judgement over whether Edmund is in the good case. And if suspension of judgement is the attitude that's warranted overall, then as long as the strong reasons for doubting that Edmund is in the good case are independent of the truth of (RKT), the Counterexample and Motivation Problems do not get off the ground once again.

In this sub-section, I precisely aim to provide strong (RKT)-independent grounds for doubting that the Edmund of *Fake-mother Invasion* is the good case. My first aim is to make explicit something which has been so far left implicit: that there is a certain kind of *anti-luck* condition on being in the good case of rationally ϕ -ing. I will then argue that there are cases structurally analogous to *Fake-mother Invasion* about which there are strong reasons for thinking that the anti-luck condition is not met. We thereby have strong reason to think that the Edmund of *Fake-mother Invasion* is not in the good case after all, so that the warrant provided for thinking that they are by the (supposed) intuitive pull of that idea is over-ridden.

First, let me try to bring into focus the anti-luck condition on ϕ -ing because p. Let us return to the contrast I drew at the beginning between the (ψ) condition and the ($\neg\psi$) condition: the distinction between the condition that holds just when the agent ϕ s because they believe that p, and the condition that holds just in case the agent ϕ s because p. The former condition is neutral on whether p is a fact, it only requires p appear to the agent to be so. The latter condition, however, requires that p is a fact.

But let us now draw a slightly different contrast: the contrast between the agent ϕ -ing because they believe that p *where p is true*, and the agent ϕ -ing because p. Clearly, these two conditions are distinct: an agent might ϕ because they believe

⁹ For more on this, see the next sub-section.

that p , where p is true, without counting as ϕ -ing because p . Consider, for example, the case of an agent who believes that p in an ill-formed manner: an agent who believes that p out of prejudice or wishful thinking, say, but where p just so happens to be true and they ϕ because they believe that p . We would not want to say of this agent that they ϕ because p .

What this brings out is that the contrast between ϕ -ing because one believes that p and ϕ -ing because p does not merely turn on the fact that the latter, but not the former, requires p to be true. Rather, it also turns on the fact that the latter, but not the former, is subject to a certain kind of *anti-luck* condition: ϕ -ing because p requires that there is a *non-accidental fit* of a certain kind between the agent's ϕ -ing because they believe that p , and the fact that p .

Now, I take it that one lesson to draw from reflection on cases such as *Thin Ice* is that it is not sufficient for ruling out the kind of luck inimical to ϕ -ing because p to claim that in addition to believing that p , and to p being true, the agent must also be *justified* in believing that p . That's because the Edmund of *Thin Ice* does not ϕ because p , even though he has a justified true belief that p , and this, intuitively, is because there remains a merely accidental fit between his ϕ -ing because he believes that p , and the status of p as a fact.

There is thus an anti-luck condition on being in the good case of rationally ϕ -ing which cannot be accounted for merely by appeal to the idea that having a justified true belief that p is required for being in the good case.¹⁰ With this point in tow, let us now return to the main issue I wish to address in this sub-section: whether there are strong grounds for doubting that the Edmund of *Fake-mother Invasion* is in the good case. That the ice is too thin to skate is a fact, Edmund is justified in believing it, and Edmund refrains from skating because he believes it. But is there a suitable kind of non-accidental fit between Edmund refraining from skating because he believes the ice to be too thin and the fact that the ice is too thin? If we can establish that there is strong reason to give a negative answer to that question, then we will have strong reason to think that Edmund is not in the good case.

Grounds for thinking that the relevant kind of fit is accidental can be brought out by reflection on the following case which more precisely matches the structure of standard barn façade scenarios than does *Fake-mother Invasion*. The case in question is a variant on Bernard Williams's (1981) well-known glass-of-petrol case:

Fake Gin & Tonic Party. Gordon wants a gin and tonic. He is currently at a party and lined-up in a row on the bar are a dozen-or-so glasses each filled to the rim with a clear, fizzy liquid. Gordon knows how to identify gin and tonic by sight, and the stuff in the glasses looks just like it. However, unbeknownst to Gordon, only one of the glasses contains a *bone fide* gin and tonic; the rest contain a combination of liquids designed to look just like gin and tonic, even to the most discerning observer. Attending to the glass which just so happens to contain real gin and tonic he forms the belief that that glass contains a gin and tonic, reaches for it, and takes a sip, feeling satisfied.

¹⁰ We can think of (RKT) and (RET) as embodying two distinct ways of accounting for the anti-luck condition on ϕ -ing because p . Locke himself shows sensitivity to this point: Locke (2015: 221).

Gordon's justified true belief that the glass contains gin and tonic is (partly and indirectly) explained by the fact that the glass contains gin and tonic. Does Gordon take a sip *because* the glass contains gin and tonic? Is Gordon in the good case of rationally ϕ -ing? If we want to say that Edmund, in *Fake-mother Invasion*, refrains from skating because the ice is too thin, then we will, I think, have to say that Gordon takes a sip because the glass contains gin and tonic; if Edmund is in the good case, then so is Gordon. There is no salient difference between the two cases with respect to this issue, it seems to me.¹¹

I now want to argue, however, that we have grounds to doubt that Gordon's ϕ -ing because he believes that the glass contains gin and tonic displays the kind of non-accidental fit with the fact that the glass contains gin and tonic required for him to count as being in the good case. To see why, we should note certain features of *Fake Gin & Tonic Party*. First, a modal feature: very easily, Gordon could've reasoned in exactly the same way as he does, and thereby ended up performing the very same action, motivated by the same beliefs, even were it false that the glass contains gin and tonic. Second, a claim relating to Gordon's abilities: he is unable to discriminate, relative to the circumstances in which he finds himself, the content of the glass he selected from the content of the other glasses in his peripheral vision of which it is not a fact that they contain gin and tonic. Each of these features of the case give us a sense in which it is merely *an accident* that Gordon ends up sipping in light of a *fact* he treats as a normative reason; that it is merely an accident that Gordon's action, motivated by his belief that the glass contains gin and tonic, is aligned with the fact that the glass contains gin and tonic.

These considerations, it seems to me, supply us with strong, though *prima facie*, reason to doubt that Gordon sips *because* the glass contains gin and tonic. The features of the case just cited indicate that his sipping is only accidentally coordinated with the relevant fact, after all. So: we have strong *prima facie* reason to doubt that the Edmund of *Fake-mother Invasion* is in the good case, too. *Fake-mother Invasion* seems to be analogous to *Fake Gin & Tonic Party* in all salient respects, after all. We can therefore conclude that, intuitive though Locke's claim that Edmund is in the good case might be, all things considered we should not accept it anyway. And if that's so, then the Counterexample and Motivation Problems do not get off the ground.

The only way for Locke to successfully respond to this, it seems to me, is to doubt that the modal and/or ability-based features of *Fake Gin & Tonic Party* suffice to preclude Gordon's action from displaying the kind of non-accidental fit with the fact that's required for being in the good case. That *may* be a plausible move to make, but it comes at the price of making it correspondingly more difficult to defend the claim that Gordon (and Edmund) lack *knowledge* that the glass contains gin and tonic (and that the ice is too thin). After all, the features of Gordon's case that I have cited are, one or the other, pointed to in order to justify the denial of knowledge to the agent in cases with the sort of structure at issue. If it's accepted that those

¹¹ In fact, I think it's clear in the case of Gordon that we lack any determinate intuition about whether he is in the good case, in which case we should also hesitate to ascribe intuitive weight to the analogous claim about Edmund. Thus, *Fake Gin & Tonic Party* buttresses the case I made in Sect. 3.1 against the intuitiveness of Locke's claim that Edmund is in the good case.

features are consistent with the agent's ϕ -ing displaying the kind of non-accidental fit with the presence of the relevant fact required for the good case of rationally ϕ -ing, it becomes correspondingly more difficult to justify the claim that such features preclude the agent's belief that p displaying the kind of non-accidental fit with the fact that p required for knowledge. But, as an opponent of (RKT), Locke would precisely want to deny that Edmund and Gordon have knowledge of their reason.

3.3 Is (RET) a Plausible Competitor to (RKT)?

So far, I have tried to establish that Locke's claim that the Edmund of *Fake-mother Invasion* ϕ s because p lacks intuitive appeal (Sect. 3.1) and that, even if it had intuitive appeal, that intuitive appeal wouldn't add up to a decisive reason to endorse the claim, for there are (RKT)-independent grounds for doubting it (Sect. 3.2). Now I want to grant to Locke even more. I want to grant that Edmund being in the good case is intuitively plausible *and* that there are no grounds for doubting it independently of the truth of (RKT). Still, I will argue now, the Motivation Problem can be tackled because (RET) faces a number of counterexamples the upshot of which is that (RKT) is restored to its status as the best explanation of Hornsby's datum (Sect. 3.3.1). With the motivation for (RKT) thus successfully restored, I will then suggest that the Counterexample Problem can be solved because there is a way of explaining away the intuition that Edmund is in the good case available to the proponent of (RKT) (Sect. 3.3.2).

3.3.1 The Motivation Problem Tackled

We have two competing claims on the table: (RKT) and (RET). The Motivation Problem is that (RET) is more explanatorily powerful than (RKT) due to its ability to explain both Hornsby's intuition concerning *Thin Ice* and Locke's intuition concerning *Fake-mother Invasion*, and so (RET) is itself a better explanation of Hornsby's intuition concerning *Thin Ice* than is (RKT). In order to tackle the Motivation Problem I am going to present a set of three counterexamples to (RET) and argue that there is no obvious way for Locke to modify his theory in order to avoid them. I will then argue that, given the availability of these counterexamples, (RKT) is restored to its status as the best explanation of Hornsby's datum that the Edmund of *Thin Ice* fails to ϕ because p. I hope to achieve this, to repeat, against the background of granting the intuitiveness of Locke's datum that the Edmund of *Fake-mother Invasion* is in the good case and granting that there are no (RKT)-independent grounds for doubting that thought.

The first counterexample to (RET) is generated by the possibility of ϕ -ing in response to a universally quantified truth. Consider the following case:

Coffee on Main Street. On his day off, Billy decides to go and read his novel in a coffee house somewhere in town. He knows that all the coffee houses are on Main Street. So: he catches the bus to Main Street.

It is a fact that all the coffee houses are on Main Street, and given that he wants to go to a coffee house, Billy treats that fact as a reason to go to Main Street.

Moreover, he knows that all the coffee houses are on Main Street, and he goes there in a way that manifests his knowledge of that fact, as well as his treating it as a normative reason. Thus, we can correctly say of him that he ϕ s because p: he catches the bus to Main Street because all the coffee houses are there. But this is not a case in which the fact that all the coffee houses are on Main Street explains why he believes that all the coffee houses are on Main Street. What explains why he believes this, we can take it, is an abductive inference (perhaps non-consciously carried out) from a sufficient number of sightings of individual coffee houses on Main Street, plus the fact that, although he knows the city well, he has not spotted any anywhere else. The quantificational fact that all the coffee houses are on Main Street does not itself explain why he believes that very proposition—indeed, it is not obvious what it would be for such a fact to explain why Billy believes as he does.¹² (RKT) can handle this case: Billy *knows* that all the coffee houses are on Main Street, after all.

Now let's consider the second sort of counterexample to (RET), knowledge of the future:

Early Morning Flight. Upon his return from one of Main Street's many coffee houses, Billy recalls that he is due to catch a flight to Edinburgh tomorrow morning. He quickly sets about packing his case in preparation.

That Billy has a flight tomorrow morning is a fact about Billy's future which Billy treats as a reason for him to pack his case now. Moreover, he knows the relevant fact and packs in a way that manifests such knowledge, as well as his treating the known fact as a reason. Thus, he ϕ s because p: he packs his case because he has a flight tomorrow. But the fact that Billy has a flight tomorrow is not itself part of the explanation why Billy holds his belief. He believes it because he remembers the fact that he successfully booked the flight on-line several months ago. That he successfully booked the flight might be causally sufficient (*ceteris paribus*) for him to have a flight tomorrow, but it is not identical to that very fact. Again, (RKT) can straightforwardly handle this case.

Finally, let's consider the third sort of counterexample to (RET), knowledge of probabilistic truths:

Probable Storm. On his way to the airport, Billy notices that the clouds are looking very grey in the vicinity of the airport. Recognising that this makes it probable that there will be a storm, he thereby wonders whether this will delay his flight.

That there will probably be a storm is a fact that is treated by Billy as a reason to wonder whether his flight will be delayed. Billy knows that it is probable that there will be a storm, and wonders about his flight in a way that manifests this knowledge, as well as his treating the known fact as a reason. Thus, he ϕ s because p: he wonders whether his flight will be delayed because there will probably be a storm. But the fact that it is probable that there will be a storm is not part of the explanation why

¹² Notice that it would make no difference were the inference deductive. What generates the counterexample is that the known proposition is universally quantified.

Billy holds that belief. He believes it because he has seen a state of affairs which (*ceteris paribus*) makes it probable that there will be a storm. Once again, (RKT) can handle this.

How might the proponent of (RET) propose to tackle these counterexamples? There are two suggestions which present themselves, both of which involve modifying (RET) in certain respects. I take each in turn, and will argue that neither is successful. After that I will examine two further suggestions that can be extracted from Locke's discussion itself, and argue that neither of them is successful either.

The first option is to modify (RET) so that it requires merely that the agent's belief that p be (perhaps partly and indirectly) explained not by the fact that p , but by a (set of) fact(s) which make p probable. Thus, we end up with the following account:

(RET_{prob}) Necessarily, if S ϕ s because p then S 's belief that p is explained by facts $\{q, r, s, \dots\}$, where given $\{q, r, s, \dots\}$ it is probable that p .

(RET_{prob}) covers *Fake-mother Invasion* because S 's belief that p being explained by the fact that p is a limiting case of S 's belief that p being explained by a (set of) fact(s) which make p probable, and in *Fake-mother Invasion* Edmund's belief that p is explained by the fact that p . What about the trio of problem cases just presented? In *Coffee on Main Street* Billy's belief that all the coffee houses are on Main Street is explained by a set of observed instances of coffee houses in that location, plus the fact that even though he knows the city well, he has not seen any coffee houses elsewhere. Plausibly, these facts together makes it probable that all coffee houses are on Main Street. Similarly with *Early Morning Flight*: that Billy successfully booked a flight to Edinburgh several months ago makes it probable that he is due to catch the flight tomorrow. And in *Probable Storm* the fact that the clouds are grey makes it probable that there will probably be a storm simply because it (nomologically) suffices for the truth of the claim that there will probably be a storm. (RET_{prob}) thus gives us the correct verdicts about the problem cases presented here as well as Locke's *Fake-mother Invasion*.

But can (RET_{prob}) handle *Thin Ice*? On the face of it, it cannot. For what explains why the Edmund of *Thin Ice* believes that the ice is too thin to skate is that his mother tells him, it's just that she believes that it isn't too thin, but for her own reasons wants Edmund to believe that it is too thin anyway. However, the fact that his mother told him that the ice is too thin makes it probable that the ice is too thin: his mother is normally reliable, after all. So (RET_{prob}) cannot explain why the Edmund of *Thin Ice* is in no position to keep off in response to the fact that the ice is too thin. The same goes for at least some other kinds of Gettier case in which there is no explanatory relation between the belief that p and the fact that p . In at least some of those scenarios, the agent's belief might well be explained by facts that make it probable that p , even though the agent is in no position to ϕ because p . Chisholm's (1989) sheep-in-the-field case, for example, fits the bill: what explains why the agent in that case believes that there is a sheep in the field is that it looks to them to be that way, and it looking to them to be that way makes it probable that it is so. And yet, that agent is in no position (for example) to take a photograph of the field because there is a sheep in it, in the rationalising sense of 'because'.

This brings us to the second option, which is to modify (RET) so that it takes on the following disjunctive form:

- (RET_{dis}) Necessarily, if S ϕ s because p then either (i) S's belief that p is explained by the fact that p or (ii) there exists some set of propositions {q, r, s...} such that S believes that p because S believes each of {q, r, s...} (in the rationalising sense of 'because'), and for each of {q, r, s...} what explains why S believes it is the corresponding fact that q, r, s....

According to (RET_{dis}), we are to carve good cases of rationally ϕ -ing into two fundamentally distinct kinds: those of the standard (RET) variety, in which S ϕ s on the basis of a belief that p which is explained by the fact that p, of which *Fake-mother Invasion* is an example, and those of a distinct variety, in which the agent's belief that p is explained in the rationalising manner by a set of further beliefs which are in turn to be explained (perhaps partly and indirectly) by their corresponding facts. This second proposal has no difficulty handling the three problem cases I've been discussing: the second disjunct succeeds in covering them. In *Coffee on Main Street*, the set of beliefs which explain in the rationalising manner why the agent believes that p are the relevant observational beliefs. In *Early Morning Flight* the set of beliefs in question has one member: Billy's belief that he successfully booked his flight several months ago. And in *Probable Storm* the relevant set of beliefs also has one member: the belief that the clouds are grey. So (RET_{dis}) can handle *Fake-mother Invasion* and the set of three problem cases I've introduced.

But can (RET_{dis}) handle *Thin Ice*? Again, on the face of it, it cannot. That's because the Edmund of *Thin Ice* could plausibly be construed as satisfying (RET_{dis})'s second disjunct. That would be so if Edmund's belief that the ice is too thin were rationally based on a belief he has that his mother has told him that the ice is too thin. If we build this feature into the case, it is still plausible that Edmund doesn't know that the ice is too thin: even though he rationally bases his belief that the ice is too thin on this further belief, the warrant provided by this further belief is plausibly defeated by the fact that his mother is, on this one occasion, not supplying him with trustworthy testimony. But since the belief that his mother told him the ice is too thin is explained by the fact that his mother told him the ice is too thin, were Edmund to rationally base his belief that the ice is too thin on the prior belief that his mother told him, he would satisfy (RET_{dis})'s second disjunct and hence the proponent of that theory would not be able to explain why he is unable to keep to the edge in response to the fact that the ice is too thin. Of course, it might plausibly be doubted that knowing that p by testimony generally involves a belief that p based on a prior belief to the effect that one has been told that p by the relevant source. But even if we were to reject that conception of testimony, still: we could write it into the case that Edmund bases his belief in the way suggested and we'd still get the intuition that he isn't in a position to ϕ because p. As long as that's so, I have my objection to (RET_{dis}).

Locke himself considers a counterexample to (RET) and offers two responses to it. Perhaps those responses enable him to avoid the trio of counterexamples presented here. Let's first look at Locke's own counterexample:

...consider a case where the fact that p and the agent's representation of the fact that p have some *common cause*. Suppose, for example, that Tom wakes up and sees that his bedside clock reads "8:00 a.m." Tom lives in a windowless apartment, so he cannot see that the sun is up. Still, he believes that the sun is up, and, acting on this belief, he gets out of bed. In this case, it might seem that Tom treats the fact that the sun is up as a normative reason, and yet the fact that the sun is up plays no part in the explanatory history of Tom's believing that the sun is up. (Locke 2015: 226–227)

Locke offers us two responses to this case.¹³ His primary response is to say that it isn't really a case of an agent ϕ -ing because p : Tom does not get up because the sun is up; we should stand by the verdicts of (RET) in these cases. Leaving aside the case of Tom, however, I don't think we should simply declare that universally quantified facts, facts about the future, and probabilistic facts are not facts the agent can ϕ in the light of; that would be a quite radical thesis, certainly one in need of argument. So Locke's primary response to his own counterexample to (RET) is not effective against the counterexamples I offer.

What about Locke's second response to his own counterexample? Locke's second response¹⁴ is that there is a way for the proponent of (RET) to allow that Tom's is a case of an agent ϕ -ing because p . We are to replace talk of S's belief that p being explained by the fact that p in (RET) with talk of S's belief that p standing in relation G^* to the fact that p :

Definition of G^* . [S's belief that p]...stands in G^* to the fact that p if and only if S's representation was formed via...a perception that p , a memory that p , or a correctly reconstructed explanatory chain linking a perception/memory of S's to the fact that p . (*ibid.*: 229)

Were (RET) to be developed in this way Locke thinks it would enable the proponent of (RET) to allow that Tom ϕ s because p , for the final disjunct of the definition of G^* is one which is satisfied in the case of Tom. Tom's belief that the sun is up is based on the following reasoning: *my alarm is going off; my alarm is programmed to go off only when the sun is up; so: the sun is now up*. This counts as a correct reconstruction of an explanatory chain linking an observation with the fact that the sun is now up.¹⁵

Does this final modification to (RET) save it from the trio counterexamples I have offered in this section? Perhaps it does. But I don't think it can handle another case I considered above: *Fake Gin & Tonic Party*. That's because Gordon's belief

¹³ Locke also considers a case involving knowledge of a mathematical proposition but rightly notes that it doesn't clearly constitute a counterexample to (RET), and so I will not consider it here.

¹⁴ The second response falls out of his response to a quite separate problem for (RET) not discussed here: the problem of deviant explanatory chains.

¹⁵ Locke borrows the notion of correctly reconstructing explanatory chains from Goldman (1967).

that the glass contains gin and tonic can plausibly be construed as standing in relation G^* to the fact that the glass contains gin and tonic. In turn, that's because there is nothing to stop us construing the case as one in which Gordon implicitly relies on the following line of reasoning in forming his belief that the glass contains gin and tonic: *that glass looks to contain gin and tonic; that's because it does contain gin and tonic; so: that glass contains gin and tonic*. But on that interpretation, Gordon would count as believing in a way that's based on a correct reconstruction of an explanatory chain linking his sensory experience of the glass with the fact that it contains gin and tonic. Thus, by the lights of (RET) modified in the way presently at issue, Gordon displays the sort of contact with the fact that the glass contains gin and tonic which is enough to put him in a position to act or hold an attitude in the light of it. And yet, as I have already argued, that is not a plausible thing to say.

I conclude, then, that (RET) faces a trio of counterexamples and that the four ways of responding to those counterexamples I have canvassed do not succeed.¹⁶ Given that we are granting to Locke here the intuitive plausibility of the thought that the Edmund of *Fake-mother Invasion* is in the good case, and that there are no grounds for doubting it independently of the truth of (RKT), the current state of dialectical play can be presented in the following way. We have a datum to explain: that the Edmund of *Thin Ice* fails to be in a position to ϕ because p . We have two options to explain that datum: (RKT) and (RET). There are three kinds of cases which constitute *prima facie* counterexamples to Locke's (RET): cases of knowledge of universally quantified propositions, knowledge of the future, and knowledge of probabilistic propositions. There is one kind of case which constitutes a *prima facie* counterexample to (RKT): cases with a barn façade-like structure. Now, if we have only two ways of explaining a certain datum, each of those ways has problems, but one has fewer problems than the other, then all things considered we ought to opt for the one which has fewer problems in explaining the datum: that will be the *best* explanation of the datum available. So: (RKT) is the best explanation of Hornsby's datum about *Thin Ice* after all.

3.3.2 The Counterexample Problem Tackled

We are assuming that it is intuitive that the Edmund of *Fake-mother Invasion* is in the good case of rationally ϕ -ing. We are also assuming that there are no (RKT)-independent grounds to doubt that claim. But, even against this background, the Motivation Problem has just been tackled: (RKT) remains the best explanation of Hornsby's datum that the Edmund of *Thin Ice* isn't in the good case. How should we tackle the Counterexample Problem: the simple problem that *Fake-mother Invasion*, and cases with an analogous structure, are counterexamples to (RKT)?

If one has otherwise undefeated reason to think a certain theory true, but one has an intuition the content of which is inconsistent with the truth of that theory, one can discount the intuition if one can come up with a way of explaining away the

¹⁶ Indeed, as one of my reviewers points out, the sheer complexity of the three ways of modifying (RET) in reply to the counterexamples gives us additional reason to be suspicious of those analyses.

intuition. To explain away an intuition is to provide an explanation of why one might have the intuition, in a way that is consistent with the falsehood of that intuition. If one's theory is well-motivated, and one is able to do this, then one is licensed in treating that intuition as an illusion and thus in ignoring it for the purpose of philosophical theorising about the relevant phenomenon. With the Motivation Problem solved, (RKT) is a well-motivated theory. The question is, then: can the proponent of (RKT) explain away the intuition that Edmund, in *Fake-mother Invasion*, keeps off the ice because the ice is too thin? If they can, then the Counterexample Problem is solved too.

Well, I think the proponent of (RKT) has the materials to enable them to do so. They can start by noting two facts about the Edmund of *Fake-mother Invasion*: (i) he refrains from skating because the ice is too thin *where that 'because' is a merely causal 'because'* and (ii) he refrains from skating because he believes that the ice is too thin *where that 'because' is a rationalising 'because'*. When conditions (i) and (ii) hold and the relevant fact, in this case: the fact that the ice is too thin, is a normative reason for the agent to engage in the relevant ϕ -ing, in this case: refraining from skating, we are liable to conflate (i) and (ii) so that it ends up seeming to us as if the agent ϕ s because p in the rationalising sense of 'because'.

So even if it is granted that it is indeed intuitive that Edmund keeps off in light of the fact that the ice is too thin, there is a way of explaining away this intuition available to the proponent of (RKT). And since their theory is antecedently well-motivated, this gives them all that they need to tackle the prospective counterexample.

Acknowledgements I'd like to thank my three anonymous reviewers at Erkenntnis for helpful comments which helped to significantly improve the paper. Particular thanks also go to Matthew Soteriou, Barnaby Walker, Maria Alvarez, and Guy Longworth.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Chisholm, R. M. (1989). *Theory of knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Goldman, A. I. (1967). A causal theory of knowing. *The Journal of Philosophy*, 64(12), 357–372.
- Hawthorne, J., & Gendler, T. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124(3), 331–352.
- Hornsby, J. (2007a). Knowledge, belief and reasons for acting. In C. Penco, M. Beany, & M. Vignolo (Eds.), *Explaining the mental* (pp. 88–105). Newcastle: Cambridge Scholars Publishing.
- Hornsby, J. (2007b). Knowledge in action. In A. Leist (Ed.), *Action in context* (pp. 285–302). Berlin: Walter de Gruyter.

- Hornsby, J. (2008). A disjunctive conception of acting for a reason. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge* (pp. 262–279). Oxford: Oxford University Press.
- Hughes, N. (2014). Is knowledge the ability to ϕ for the reason that p ? *Episteme*, 11(4), 457–462.
- Hyman, J. (1999). How knowledge works. *Philosophical Quarterly*, 50(197), 433–451.
- Hyman, J. (2006). Knowledge and evidence. *Mind*, 115(460), 891–916.
- Hyman, J. (2010). The road to Larissa. *Ratio*, 23(4), 393–414.
- Hyman, J. (2011). Acting for reasons: Reply to Dancy. *Frontiers of Philosophy in China*, 6(3), 358–368.
- Hyman, J. (2015). *Action, knowledge, and will*. Oxford: Oxford University Press.
- Locke, D. (2015). Knowledge, explanation, and motivating reasons. *American Philosophical Quarterly*, 52(3), 215–232.
- Lycan, W. G. (1977). Evidence one does not possess. *Australasian Journal of Philosophy*, 55(2), 114–126.
- Lycan, W. G. (2006). On the Gettier problem problem. In S. Hetherington (Ed.), *Epistemology futures* (pp. 148–168). Oxford: Oxford University Press.
- Marcus, E. (2012). *Rational causation*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1996). *Mind and world*. Cambridge, MA: Harvard University Press.
- McDowell, J. (2013). Acting in the light of a fact. In D. Bankhurst, B. Hooker, & M. O. Little (Eds.), *Thinking about reasons* (pp. 13–28). Oxford: Oxford University Press.
- Roessler, J. (2014). Reason explanation and the second-person perspective. *Philosophical Explorations*, 3, 1–12.
- Ryle, G. (2000). *The concept of mind*. Baltimore: Penguin.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Clarendon Press.
- Williams, B. (1981). *Moral luck: Philosophical papers 1973–1980*. Cambridge: Cambridge University Press.