

This is the final draft of the article published in July 2018 in *Philosophy and Phenomenological Research* 97 (1): 249–61. doi: 10.1111/phpr.12506. Please cite the published version.

REPLIES TO MY CRITICS

I should like to thank Ram Neta for commissioning this symposium, and Jennifer Hornsby and Timothy Williamson for taking the trouble to write about my work. I have included a brief statement of the views they criticize, at the beginning of each of my replies, so that readers who are not familiar with *Action, Knowledge, and Will* can follow the exchange without reading the relevant chapters first.

Reply to Hornsby

I argue in *Action, Knowledge, and Will* (henceforth *AKW*) that in order to make progress in the philosophy of action, we need to distinguish between four aspects or dimensions of human agency, which were traditionally combined in the idea of a will: the physical dimension, which we think and reason about by means of concepts such as *event* and *cause*; the psychological dimension, where the relevant concepts include *desire* and *intention*; the intellectual dimension, where they include *knowledge* and *reason*; and the ethical dimension, involving *voluntariness* and *choice*.

Jennifer Hornsby says her comments are about ‘what some might think of as a fifth dimension’ of human agency, ‘the metaphysical’. In fact they are concerned with the concepts I discuss under the ‘physical’ heading. The names I give to the dimensions of agency are of course all optional: ‘conative’ would do as well as ‘psychological’, ‘intellectual’ could be replaced by ‘rational’, and so on. But there is no need to postulate a metaphysical dimension of human agency in addition to the physical, even if there are non-physical agents—e.g. souls or, more plausibly, institutions—because the same group or family of concepts would be included under both headings.

These four dimensions of individual human agency are not the only ones that matter to philosophy. On the contrary, the social and political dimensions of agency are just as important as the ones I examine in *AKW*. But the latter are the ones that were combined or amalgamated in the concept of the will—I refer here to the will of an individual: the ‘general’ will is another matter—and the result of amalgamating them was that profoundly different problems about human agency were confused with each other, or assumed to have the same solution.

For example, consider the following two questions. First, what makes a change in my body, such as a movement of a limb, attributable to me personally as the agent? And second, what makes my conduct qualify as voluntary, and therefore potentially

culpable or meritorious?¹ The first question is about the distinction between activity and passivity, the distinction between causing and undergoing change, whereas the second is about an agent's liability to a particular range of responses or judgements, such as gratitude and resentment, or praise and blame. So the first is a physical question whereas the second is an ethical question. And yet they commonly received precisely the same answer: causation by my will.

As I have indicated, Hornsby's comments are about what I call the physical dimension of human agency. Her stated aim is to consider 'whether Hyman's view of acts as causings can accommodate what is said when action verbs are used with imperfect aspect', her view being that it cannot. But in fact this is one of several objections she advances against the position I defend concerning the relationship between the concepts of act, event, and cause. In what follows, I shall briefly summarize the position I defend in *AKW* regarding the relationship between agents, acts and events, before turning to Hornsby's objections to it.

First, I reject the traditional doctrine that some movements of the body qualify as acts, e.g. ones that are caused by desires or by the will. When someone does a simple act, such as raising her arm, her act and the motion of her arm are not one and the same thing.² For to raise something is to *cause* it to rise. So the act of raising her arm is the agent's *causing* of the motion of her arm, and the causing of the motion cannot be identical with the motion caused (*AKW*, pp. 55f). Adopting terminology introduced by Von Wright, we can call the motion of her arm the *result* of her act of raising it, but the motion is not an act, any more than the death of Caesar, as distinct from the act of killing him—i.e. causing his death—was an act.

Now that comparison, and indeed some of the words in which I expressed it, are due to Prichard, who goes on to argue that instead of being identical, the act of raising one's arm and the motion of one's arm are related as cause and effect.³ In her influential 1980 book *Actions*, Hornsby agreed with this proposal, but as she says in her comment, she abandoned it some time ago. In my view, she was right to do so. There are, I believe, compelling logical and epistemological arguments against it (*AKW*, pp. 57-59). But if the act of raising one's arm is neither identical with the motion of one's arm nor related to it as cause and effect, how *are* they related?

The answer I defend in *AKW* is that to do a certain kind of act is to exercise the power or ability to cause a certain kind of change. The change may be a kind of motion, as when someone opens or closes a door or raises an arm, or another kind of change may be involved, as when one burns some toast. The agent is the one that causes the change and the patient is the one that undergoes it—though in

¹ It is debatable whether conduct that is not voluntary can be meritorious. I discuss this question in J. Hyman, 'Voluntariness and Intention', *Jurisprudence* 7 (2017), pp. 692-709.

² The doctrine that some movements of the body qualify as acts is accepted *inter alia* by Wittgenstein, Anscombe, Davidson and Armstrong, and rejected *inter alia* by Von Wright and Kenny. For references, see *AKW*, 3.1.

³ H.A. Prichard, *Moral Obligation* (Oxford: OUP, 1949), p. 191.

some cases, such as suicide, these are one and the same—and the kind of act done depends on the kind of change caused. If the change is something's burning it is an act of burning something, if the change is something's dying it is an act of killing something, and so on. A particular act is therefore a particular instance of an agent's causing a certain kind of change.

The relationship between an act and its result—the causing of a change, and the change caused—is therefore similar to the relationship between an instance of a relation and one of the relata, for example, a marriage and one of the spouses. Why just 'similar to'? Why not say that an act *is* an instance of a relation, the causal relation between the agent and the act's result? I argue that this would not be quite right, because acts are dynamic whereas relations are static, i.e. ways in which one thing can stand to another thing, or several things can stand to one another (AKW, p. 74). That is why there is no analogue of the relationship between a relation and its converse in the case of acts. For example, 'Brutus was older than Caesar' and 'Caesar was younger than Brutus' ascribe converse relations, whereas 'Brutus killed Caesar' and 'Caesar was killed by Brutus' describe one and the same act, with a single 'direction' or dynamic orientation (AKW, p. 71ff). It is also why acts take time whereas relation-instances do not. For example, David was the father of Absalom if, and only if, David begat Absalom. But begetting Absalom was one thing and being his father was another. The former may have been pleasant whereas the latter seems to have been painful, and we can ask when and where David begat Absalom, and how long it took, but not when and where David was the father of Absalom, or how long *that* took.

Hornsby criticises this theory of action on several grounds. The first, which is about the use of a verb in the imperfective or progressive form, concerns her most, and I shall discuss it at some length. I shall respond to the others more briefly.

(1) Hornsby's first criticism has to do with the use of the imperfective form of the verb to say that an act or event was or is in progress at a given time, without implying that it has been or will be completed: for example, 'Ann was drying a plate'. As I point out in a note in AKW, it is sometimes claimed that one cannot properly be described as being engaged in causing (as opposed to attempting to cause) a result, such as a plate's becoming dry, unless it actually occurs (AKW, p. 34n). (Call this the 'causing' claim.) But if that is right, then 'Ann was drying a plate' cannot be glossed as 'Ann was causing a plate to become dry'.

My comment on this line of thought is rather brief, perhaps too brief. Hornsby implies that I accept the 'causing' claim, but in fact I simply point out that *if* we accept it, then 'we cannot infer from the fact that to do an act of a certain kind is to cause a change of the corresponding kind, that to be doing an act of the same kind is to be causing a change of the corresponding kind.' Now Hornsby is sure that the inference is valid. 'It doesn't seem possible', she writes, 'to agree with Hyman that to dry the plate is to cause a certain change, yet to deny [...] that to be drying the plate is to be causing such a change.' And she may be right. As I see it, this is a dilemma for those who find the 'causing' claim plausible: either give up the claim—or at least interpret it as a cancellable implicature—or give up the inference.

This may seem like a minor point, but Hornsby explains that she is concerned about it because the ‘tense and aspectual behaviour of verbs needs to be brought into account’ when we explain the meaning of causative verbs. As she explains in her 2012 article ‘Acts and Activity’, where she spells out her position in more detail, her own view now is (i) that verbs in the imperfective aspect can be used to attribute *activity* of some kind to an agent, as opposed to a particular *act*; (ii) that an act is ‘comprised from a bit of activity’, much as an object of a certain kind, say a pill, consists of a quantity of stuff, such as aspirin; (iii) that what I call the result of an act—i.e. the change caused by the agent—is ‘comprised from’ the same ‘bit of activity’ as the act itself; and (iv) that the act and its result are one and the same event. So, to return to the question of the meaning of causative verbs, a single event ‘might be described equally as “her raising her arm” or as “her arm’s going up”’.⁴

The only kind of act Hornsby discusses in detail in the 2012 article is a person’s raising her arm, and she does not explain there how her picture applies to non-basic acts. But in her comment on my book, she argues that in this kind of case too an act is the same event as its result. ‘The flag’s rising’, she proposes there, ‘simply was its being raised by John’, and since she agrees that when John raises a flag, the flag’s being raised by John and John’s raising the flag are one and the same act, it follows that according to her view, just as a single event ‘might be described equally as “her raising her arm” or as “her arm’s going up”’, another single event might be described equally as “John’s raising the flag” or as “the flag’s going up”. So this is not Davidson’s theory that every act is a movement of the agent’s body. Hornsby agrees with Davidson that a person’s raising her arm and the motion of her arm are one and the same event. But whereas Davidson identifies John’s act of raising a flag with the motion of his arm, Hornsby identifies it with the motion of the flag.⁵

My view about these claims is as follows. (i) is incontrovertible: the imperfective form certainly can be used to attribute activity of some kind to an agent, such as knitting or walking, as opposed to a particular act, although it can also be used in a habitual sense, as in ‘Anne is making her own breakfast, while Henry is away’. (ii), the claim that that an act is ‘comprised from a bit of activity’, is plausible in many cases, especially where the act is one that can be explicitly described as an activity with a specific terminus or span, such as knitting a sweater or walking across Brooklyn Bridge, although it should be borne in mind that activities such as knitting or walking consist in turn of specific acts, such as doing a purl stitch or taking a step. It is less plausible in some other cases, such as casting a vote or moving a bishop, where the kind of activity we would normally describe the agent

⁴ J. Hornsby, ‘Actions and Activity’, *Philosophical Issues* 22 (2012), pp. 234f. See also J. Hornsby, ‘Basic Activity’, *PAS*, suppl. vol. 87 (2013), pp. 1-18. The analogy between act/activity and thing/stuff was originally drawn in Alexander P.D. Mourelatos, ‘Events, Processes, and States’, *Linguistics and Philosophy* 2 (1978) pp. 415-434.

⁵ Hornsby’s position is also defended in U. Coope, ‘Aristotle on Action’, *PAS*, suppl. vol. 81 (2007), pp. 109-38, and D. Hillel-Ruben, ‘One-particularism in the theory of action’, *Philosophical Studies* (2017).

as being engaged in is not casting or moving, but, say, political activity or playing chess. For in these cases, the activity consists in acts such as the ones mentioned, rather than the other way around, just as walking consists in taking steps, whereas taking steps does not consist in walking. However, if every basic act consists in (or is 'comprised from a bit of') the activity of moving parts of one's body, then there is a sense in which every non-basic act also 'ultimately' consists in activity of this kind.

That is all I have to say about (i) and (ii). Hornsby says she doubts whether my view of acts as causings 'can accommodate what is said when action verbs are used with imperfect aspect'. Presumably she has points (i) to (iv) in mind. But my position is certainly consistent with (i) and (ii). The points it is *not* consistent with are (iii) and (iv): the claim that result of an act is 'comprised from' the same 'bit of activity' as the act itself, and the claim that an act and its result are one and the same event.

But (iii) and (iv) cannot be right. This should be obvious in the case of non-basic acts—acts whose results are not movements of the agent's body—such as raising a flag or Caesar's murder. The result of Caesar's murder was his death, and his death was not 'comprised from a bit of activity'—some vigorous stabbing by the conspirators, for example—and it cannot be the same event as his murder. For the death was caused by haemorrhage or shock, whereas the murder was not. But it is equally true in the case of basic acts. When I raise my arm, the motion of my arm is not 'comprised from a bit of activity', and cannot be the same event as my act of raising it. For the motion of my arm was caused by a contraction of my biceps, whereas my act of raising it was not. Furthermore, the claim that an act and its result are one and the same event faces the familiar and insurmountable difficulties about the location of acts that were explored in the 1960s and 1970s, and which I discuss in chapter three of *AKW*. Most of the literature on this topic concerns non-basic acts, but the problem is no less acute in the case of basic acts. For as Davidson points out, 'if a man's arm goes up, the event takes place in the space-time zone occupied by the arm; but if a man raises his arm, doesn't the event fill the zone occupied by the whole man?'⁶ Davidson takes the example to show that we do not have adequate criteria for the location of an event, but this is simply an evasion. Acts or events that do not occur in the same place cannot be identical. 'Fill the zone' may not be exactly the right phrase, but if the motion of a limb occurs in a certain place, and the act of moving it does not occur in that place, then they are distinct.

For these reasons, while my view of acts as causings cannot accommodate everything that Hornsby says about 'what is said when action verbs are used with imperfect aspect', I believe it can accommodate the parts of what she says that are true.

(2) *Suppose that John set his plate drying machine to start in five minutes' time. One might think that in this case it was the machine, not John, that dried the plate; and that here John did cause the plate to dry: he did so by setting the machine.*

⁶ D. Davidson, *Essays on Actions and Events* (Oxford: OUP, 1980), p. 124.

In *AKW*, I argue that ‘to raise, move, wet, dry, kill, etc., something is to cause it to rise, move, become wet, become dry, die, etc., regardless of identities of the agent and the patient and regardless of the means by which the act is done’ (*AKW*, pp. 36f). I discuss the kind of objection Hornsby raises on pp. 38f, and I do not have much to add to what I say there. It is true that we tend to reserve the phrase ‘cause . . . to rise (move, die, etc.)’ for the case where something is raised (moved, killed, etc.) more indirectly than usual; and equally that we tend to reserve the causative verbs for cases where causation is relatively direct, although the proximity of the connection suggested by the causative expression is relative to what is normal in the kind of case concerned. But we do not observe this rule consistently. For example, we do not find the statement that Stalin killed millions jarring, although he did not do so personally, or with his own bare hands. Furthermore, if we did observe the rule strictly, it would not follow that to raise, move, kill, wet, dry, etc., something is *not* to cause it to rise, move, die, become wet, become dry, etc. On the contrary, if ‘raise’ meant ‘cause to rise by some relatively direct means’, then to raise something would be to cause it to rise by some relatively direct means, and hence, *a fortiori*, it would be to cause it to rise.

(3) *If ‘raise X’ meant ‘cause X to rise’, then one would expect ‘John raised the flag slowly’ to be equivalent to ‘John caused the flag to rise slowly’. But the latter sentence conveys that the flag’s rising was slow. It doesn’t convey, as the former does, that John did something slowly.*

True, ‘John caused the flag to rise slowly’ doesn’t convey that John did something slowly. The adverb is in the wrong place, and the way to convey that John acted slowly would be to say ‘John slowly caused the flag to rise’. But it does not follow that ‘raise the flag’ does not mean ‘cause the flag to rise’. Sameness of meaning does not imply substitutability *salva veritate* without any change in word order. For instance, ‘X took out a contract on Y’ means ‘X hired someone to murder Y’. But ‘Michael took out a contract on Fredo regretfully’ does not mean that the hitman was required to commit the murder in a regretful state of mind. The correct paraphrase is ‘Michael regretfully hired someone to murder Fredo’.

(4) *‘Paul melted the butter’ conveys that Paul and the butter stand respectively in the relation that any two things stand if the first melted the second. Here there is no claim that an act in any sense is a relation. The claim is that the two-place predicate ‘melted’ expresses a relation.*

I doubt whether this can be right. ‘Paul melted the butter’ may indeed *convey* that Paul and the butter stand respectively in the relation that any two things stand if the first melted the second, if ‘convey’ means suggest or imply, but this is not what someone who utters the sentence ‘Paul melted the butter’ *says*. What they say is that Paul melted the butter, which is analogous to the statement that David *begat* Absolom (in which the predicate is an action verb), rather than the statement that David *is the father of* Absolom (in which the predicate expresses a relation). If we want a predicate analogous to ‘x is the father of y’, which expresses the relation that any two things stand in if the first melted the second, then we shall have to use ‘x is the one who melted y’, or something of that kind. But whereas *this* predicate

certainly does express a relation, if an act *is* not a relation—as Hornsby appears to concede—then the predicates ‘x melted y’ and ‘x begat y’ do not *express* relations.

Hornsby says that ‘no-one really thinks that acts themselves are included in the category of relations. Hyman’s real opponent here, I think, is someone who wants to allow that relational predications may be made with action verbs.’ In fact, Kenny showed that the claim that acts are relations had been implicit in the treatment of action verbs by logicians since Russell, when he (Kenny) made philosophers aware that it was problematic.⁷ Davidson then reasserted the claim explicitly in several places.⁸ And where Davidson led, many others followed. But the substantial point is that Hornsby draws a distinction without a difference. For acts *are* relations if, and only if, the predicates we use to report them *express* relations. Perhaps the word ‘convey’ is the culprit, since we can *imply* that the relational predicate ‘x is the father of y’ applies to David and Absalom by *saying* that David begat Absalom, but it does not follow that ‘x begat y’ also expresses a relation, which—since begetting is not a relation—it does not.

Reply to Williamson

In the last three chapters of *AKW*, I defend the theory that knowledge is the ability to be guided by the facts. To know a fact, I argue, is to have the ability to be guided by it, to respond to it rationally, in what one thinks, or feels, or does. The core of the argument is quite simple.

First, the traditional assumption that knowledge is a species of belief is mistaken. Both can be described as mental states, in a broad sense of the term. But knowledge and belief have different objects: the object of belief is a proposition, whereas the object of knowledge is a fact (*AKW*, p. 163).⁹ Moreover, knowledge is an ability, whereas belief is a disposition. Belief, like many other attitudinal dispositions, such as love or trust, can be foolish, passionate, or whole-hearted; whereas knowledge, like other mental skills and abilities, such as the ability to calculate or reason, cannot be any of these things. We ask *why* rather than *how* someone believes something, or trusts someone, wanting a justification; whereas we ask *how* rather than *why* someone knows something, wanting to be told the means by which their knowledge was acquired (*AKW*, p. 164).

But if knowledge of a fact is an ability, what is it an ability to do? In other words, what kinds of thought, feeling, or behaviour express a person’s knowledge of a fact? It should be obvious that thought, feeling and behaviour do not express a person’s knowledge of fact whenever they are explained by that fact. For example, someone with diabetes might feel despondent because her blood sugar is low without *knowing* that her blood sugar is low. Feeling despondent is simply an effect low

⁷ A.J.P. Kenny, *Action, Emotion and Will* (London: Routledge & Kegan Paul, 1963), ch. 7.

⁸ E.g., D. Davidson, ‘Adverbs of Action’, in B. Vermazen & M.B. Hintikka (eds), *Essays on Davidson* (Oxford: OUP, 1985), p. 232.

⁹ I discuss this point at greater length in J. Hyman, ‘Knowledge and Belief’, *PAS*, suppl. vol. 91 (2017), pp. 267-289.

blood sugar sometimes has. But suppose she worries about hypos (episodes of low blood sugar), checks her blood sugar level, discovers that it is low, and feels despondent for this reason. In *this* case, the statement that she feels despondent because her blood sugar is low is a different kind of explanation. It is the same kind of explanation we would give of an intentional act, e.g. if we explained that she ate a biscuit because her blood sugar was low. It presents the fact that her blood sugar is low as *her* reason—not merely *the* reason—why she feels despondent, in other words, it presents it as a fact she responds to rationally, a fact she is guided by, or takes it consideration or account.

In *AKW* I defend the theory that a person's knowledge of a fact is her ability to be guided by it, in what she thinks or feels or does. One might think that there is no difference between knowledge and belief in *this* respect—that they explain a person's thoughts, feelings and behaviour in exactly the same way. But I argue that this is a mistake (*AKW*, 6.4). It is true that knowledge and belief tend to have the same mental and physical effects. For example, a man who has never travelled from Athens to Larissa, and does not know the way, but has the right opinion about the route purely by chance, will take the same road as a man who does know the way, and has made the journey many times. But one cannot be guided by a fact one does not know, any more than one can be guided by a sign one cannot see, so explanations that involve knowledge are different from ones that involve belief. 'He took the left fork because he knew that it leads to Larissa' refers to a feature of his situation he was aware of and took into account, whereas 'He took the left fork because he believed that it leads to Larissa' merely refers to his state of mind.

In his interesting and resourceful comments, Timothy Williamson disputes the claim that knowing a fact can be equated with having the ability to be guided by it, or, as he expresses the idea, with having the ability to 'do things *because p* (on the relevant reading of "because")'. His discussion falls into two parts. In the first part, he concedes for the sake of argument that being able to do things *because p* is both necessary and sufficient for knowing that *p*. Then, in the second part, he contests this proposition by means of three examples. The first example is meant to show that knowing that *p* is not necessary for having the ability to do things *because p*, the second and third ones are meant to show that it is not sufficient. I shall begin with the examples.

(1) *Right now, with my eyes open, I have the ability to move cautiously because my eyes are shut. It is just that, to exercise that ability, I have to shut my eyes. [...] Thus knowing that my eyes are shut is not necessary for being able to do things because my eyes are shut. It is only necessary for doing things because my eyes are shut.*

Is Chris Froome able to ride a bike in his sleep? Probably not. Somnambulism is well documented, bike-riding while asleep less so. But Chris does not lose his ability to ride a bike when he falls asleep, and then recover it again when he wakes up. His ability to ride a bike is undiminished, although he cannot exercise it in his sleep. In other words, he has the ability to ride-a-bike while asleep, but he doesn't have the ability to ride-a-bike-while-asleep. Williamson thinks that the ability to be guided by a fact is similar: 'Right now, with my eyes open, I have the *ability* to move cautiously because my eyes are shut. It is just that, to exercise that ability, I have to

shut my eyes.’ But can that be right? Consider a different case. Suppose I do not have a bike. Does Chris have the ability to ride *my* bike? The answer is surely that if I acquired a bike, he would have the ability to ride it, but as things stand there is no such ability since there is no such bike. It would be wrong to say: ‘Right now, Chris has the *ability* to ride my bike. It is just that, for him to exercise that ability, I would have to acquire one.’ Similarly, if my eyes are open, there is no such fact as the fact that my eyes are shut, so the question whether I am able to be guided by it does not arise. As soon as I shut my eyes, the question *does* arise, and the answer may be either yes or no, depending on the circumstances. For example, I may open and close my eyes while I am asleep, but I cannot be guided by the fact that my eyes are closed in these circumstances, unless they are also circumstance in which I can know that my eyes are closed. Hence, the example fails to show that knowledge of fact is not necessary for the ability to be guided by it, because in the example there is no such fact. So the objection fails.

(2) *A criminal being led out to execution might know the conjunction that the money is hidden in the cellar and he will never do anything even partly because the money is hidden in the cellar; such a conjunction is clearly knowable in principle. [But] it is impossible for him to do anything for that conjunctive reason. Consequently, he lacks the ability to do something because the conjunction holds.*

The first sentence may be true. But the claim that knowledge of a fact is the ability to be guided by it is not about the ability guided by it *purely in what one does*, for one can be guided by a fact one knows in thought or feeling too, and as I emphasize in *AKW*, this counts as an exercise or expression of knowledge, just as much as being guided by a fact one knows in what one does (p. 167). The convict in Williamson’s example may know that he won’t *do* anything, anything significant at least, because the money is in the cellar—for example, he won’t tell a friend where to find it—since the time for action has evidently passed when the place of execution comes into view. But he may still *realize* that the butler will be the person who finds the money, and if he likes the butler, he may *feel pleased* that he didn’t leave it in the attic, as a result. Indeed, even the fact that he will never do anything even partly because the money is hidden in the cellar (as opposed to the simple fact that the money is hidden in the cellar) might prompt him to think or feel something, such as regret. After all, it is a poignant fact, from his point of view. At what point does it become certain that he will longer have any thoughts or feelings of these kinds? How close to death does he need to be? That depends on how long it takes to make an inference. Suppose the answer is 100 ms. Can he know that he will never do or think or feel anything even partly because the money is hidden in the cellar after that moment has passed, assuming he knows that it has passed, i.e. that he has less than 100 ms to go, before his ability to think is finally extinguished? I think not. For if he *did* know that he will never do or think or feel anything etc., this would have to be as a result of making an inference from the fact that it is now too late to make an inference. So I doubt whether this is in fact something he can know.

In sum, Williamson’s conjunction does seem to be knowable in principle, and available to the criminal as a reason for thinking X or feeling Y. Whereas the conjunction that the money is hidden in the cellar and he will never *do or think or*

feel anything even partly because the money is hidden in the cellar is neither knowable by him nor available to him as a reason. So the example does not disprove the equivalence between knowing a fact and being able to be guided by it. On the contrary, it confirms it.

(3) Let s be the sentence "Everything is as it actually is", on the reading formalized in modal logic as $\forall p (p \leftrightarrow @p)$ with quantification into sentence position and the rigidifying "actually" operator @. As uttered in any possible world w , s expresses a proposition p_w true in w and in no other world. [...] Consider a possible world w in which someone knows p_w under the guise of s but never acts on the knowledge. [...] In any other world, p_w is false and so is not a reason for which she can do anything. Thus she does not do anything for the reason p_w in any possible world. [...] Since one has the ability to do something only if it is metaphysically possible for one to do it, even in w she lacks the ability to do things for the reason p_w . Nevertheless, in w , she knows p_w .

It would be easy to offer an accommodating reply to this objection, something to the effect that knowledge of a tautology, or of a tautology of a certain kind, is a limiting or special case. But I don't think we should concede this much to the objection. Roughly, I think the objection shows that once we have introduced the '@' operator, or learned to use the word 'actually' in this way, the principle on which the objection relies, that

(P) one has the ability to do something only if it is metaphysically possible for one to do it,

no longer holds without restriction. Or alternatively, the availability of the '@' operator shows that the principle does not hold without restriction. The reason for saying this is that the exercise of the ability to actually V (where ' V ' stands for a verb-phrase) is excluded in non-actual worlds *a priori*. So if the ability is not exercised in the actual world, it is not exercised in any possible world at all. Hence, if (P) were true without restriction, the ability to actually V would not exist unless it was exercised. But it is a general rule (with a definable range of exceptions) that an ability does *not* need to be exercised to exist. So (P) cannot be true without restriction. I shall explain this in more detail now.

As a rule, an ability does not need to be exercised in order to exist, unless doing the thing is the only way to acquire the ability, as it is for speaking English (despite the apocryphal story about Carlisle), and for walking, in the case of human beings. As Williamson points out, it is easy to define an ability that does not need to be exercised in order to exist by combination, such as the ability to recite a limerick while standing on one leg. But Williamson's argument shows that if we accept (P), we can construct a whole raft of exceptions to this rule. In fact, for any ability to V , the ability to actually V is an exception, so in effect the rule becomes the exception and the exception becomes the rule. Take the ability to actually climb Everest. Evidently, this ability can only be exercised in the actual world, so if a climber doesn't exercise it in the actual world, she doesn't exercise it in any possible world at all. Hence, if (P) were true without restriction, it would follow that

(C) a climber who never actually climbs Everest does not have the ability to actually climb Everest.

And the same goes for every other ability to actually V, including the ability to actually recite a limerick while standing on one leg.

Evidently, (P) contradicts the fundamental idea that an ability, like a disposition, but unlike a tendency, does not need to be exercised to exist. It is true that @ possibly p does not imply possibly @ p , but an ability is not a species of possibility. As I point out in *AKW*, it is an actual property of the substance that possesses it, and not merely the possibility that the kind of act or event that manifests it should occur (p. 180). And as a matter of fact, many climbers who never actually climb Everest *do* have the ability to actually climb Everest. So (P) has to go.

This invites the question why (P) seems plausible. The answer, I suggest, is that if someone who has the ability to do something attempts to do it in propitious circumstances, she will normally succeed. This is a fundamental feature of our concept of an ability, no less so than the rule, with which it is perfectly consistent, that an ability does not need to be exercised to exist. But it is tempting to infer, first, that if there are *no* circumstances in which someone would succeed in doing something, then she cannot have the ability to do it, and then to infer from *this*, that

(P) one has the ability to do something only if it is metaphysically possible for one to do it,

which is *not* consistent with rule that an ability does not need to be exercised to exist. Which of these two inferences should be rejected depends on the interpretation of the clause 'there are no circumstances in which someone would succeed in doing something', but evidently they cannot both be accepted.

In summary, as long as we ignore the '@' operator, (P) can seem plausible, but once '@' has been introduced, it is clear that (P) cannot be true without restriction, because the exercise of the ability to actually V, unlike the exercise of the ability to V, is restricted *a priori* to the actual world. That is not to say that no modal principle of the kind Williamson relies on is true. It is plausible that one has the ability to actually do something only if it is metaphysically possible for one to do it, for instance, one has the ability to actually climb Everest only if it is metaphysically possible for one to climb Everest. But it does not follow from the proposition that it is metaphysically impossible for someone to do something for the reason p_w , that in w she lacks the ability to do things for that reason. So the objection fails.

It is encouraging to find that none of the alleged counterexamples disproves the proposition that a person knows that p if, and only if, she has the ability to be guided by the fact that p , to respond to it rationally, or take it into account, despite their ingenuity and sophistication. But as I have indicated, Williamson also argues that it is a mistake to equate a person's knowledge of a fact with her ability to be guided by it even if the bi-conditional is true. First, the equation 'makes the state of not knowing that p look too disunified to well explain the inability to do something

because A'. Second, 'not all differences in the ability to do things because p are differences in the knowledge that p , so the ability is distinct from the knowledge.' I shall comment on these arguments in turn.

I found Williamson's first argument puzzling when I first saw it, in a laconic footnote in *Knowledge and its Limits*. Williamson concedes there that someone who believes truly that p without knowing that p cannot do X because p , cannot do Y because p , and so on.

But a single failure to know explains all these incapacities. If the incapacities constituted the failure to know, the correlation between the incapacities would be an unexplained coincidence.¹⁰

In reply, I pointed out that if the incapacities constituted the failure to know, it does *not* follow that the correlation between the them would be an unexplained coincidence, since it would be explained by whatever explains the person's failure to know, such as the fact that she did not see this morning's paper. At the same time, I agreed with Williamson's remark that 'a single failure to know explains all these incapacities':

It explains them by including them, in the way that a person's inability to ride a bicycle explains why she cannot cycle from X to Y, from Y to Z, from Z to W, and so on. (AKW, p. 182)

Williamson remains unconvinced. He points out that if someone is unable to ride a bicycle because she has lost her legs, the fact that she has lost her legs provides a 'unified' explanation of her inability to ride from X to Y, from Y to Z, and from Z to W. But if we imagine 'a random list of a thousand miscellaneous things' I am unable to do (compose a sonata, speak Yiddish, climb Everest, etc.), 'my general inability to do any of those things implies my particular inability to do any given one of them, but it does not well *explain* that particular inability, because the general inability is too disunified.'

The ability account of knowledge [he concludes] makes the state of not knowing that p look too disunified to well explain the inability to do something because p .

This argument seems to me muddled in two ways, first because Williamson confuses the idea of a unified explanation with the idea of a unified ability, and second because he fails to distinguish between a general ability and a number of specific ones.

Presumably Williamson does not object to citing abilities or incapacities as explanatory factors as such. For although the fact that someone has lost her legs may explain why she cannot cycle from X to Y, from Y to Z, or from Z to W, the fact that she is unable to ride a bike can explain this too, regardless of *why* she is unable to ride a bike. And if this *is* the reason, then the fact that she cannot cycle from X to

¹⁰ T. Williamson, *Knowledge and its Limits* (Oxford: OUP, 2002), p. 64.

Y, from Y to Z, or from Z to W, *does* have a unified explanation, in other words, it has a *single* explanation. By contrast, if she is a perfectly competent cyclist—people do cycle with prosthetic legs—but the road from X to Y is impassable, and the roads from Y to Z and from Z to W are too steep, then the fact that she cannot cycle from X to Y, from Y to Z, or from Z to W, *does not* have a unified (i.e. single) explanation. Thus, the fact that someone does not have a general ability, such as the ability to ride a bike, is perfectly capable of explaining why she is unable to exercise this ability in a specific way, and it is also perfectly capable of providing a unified explanation of why she is unable to exercise it in a variety of specific ways.

So much for unified explanations. But if Williamson's objection is that the ability account of knowledge makes not knowing a fact look like the inability to do the things on a miscellaneous list, the answer is that it does nothing of the kind. It is true of course that my inability to do the things on a miscellaneous list that includes speaking Yiddish does not explain *why* I am unable to speak Yiddish. But the reason for this is not that 'my general inability to do any of those things ... is too disunified.' It is not a general inability at all: it is a conjunction of specific incapacities, one of which is precisely what we are seeking to explain. Whereas the inability to be guided by a fact, to respond to it rationally, *is* a general inability, because the *ability* to be guided by a fact is a general ability, which can be exercised in vastly many different ways in different circumstances, rather like the ability to use a concept or (if we compare the ways in which knowledge can be expressed to journeys) the ability to ride a bike. It is not a conjunction of abilities, let alone a conjunction of miscellaneous abilities.

Admittedly, there is an important difference between the ability to be guided by a fact and the ability to ride a bike, because the ability to be guided by a fact is not the ability to perform a specific kind of act or to engage in a specific activity or range of activities. It is, roughly, the ability to do (or think, or feel) things *for a specific reason*, or *in the light of a specific fact*. So what it is an ability to do is captured by an adverb, rather than a verb. But this does not prevent either the ability or the corresponding inability from playing an explanatory role. For comparison, I point out in *AKW* that unpunctuality is a tendency to do things later than the appointed time. So what unpunctuality is a tendency to do is captured by an adverb too. But this does not mean that someone's unpunctuality is incapable of explaining some particular kind of behaviour to which they are prone, such as failing to pay bills when they fall due, or that it is incapable of providing a unified explanation of several kinds of behaviour.

Williamson's second reason for rejecting the equation between knowledge and the ability to be guided by the facts is that 'not all differences in the ability to do things because A are differences in the knowledge that A, so the ability is distinct from the knowledge.' For example, suppose Tom can run and cycle, but never learned to swim. If he knows that exercise is beneficial, he can go for a run or go for a bike ride for this reason, but he cannot go for a swim for this reason, because he cannot swim. Now suppose he learns to swim. According to Williamson, his ability to be guided by the fact that exercise is beneficial has changed, since it now encompasses going for a swim as well as going for a run or going for a bike-ride, whereas his knowledge of the fact that exercise is beneficial has *not* changed: it remains exactly

what it was. But, as Williamson points out, if 'his ability has changed while his knowledge has stayed the same', the ability and the knowledge cannot be one and the same state.

This argument is simpler than the first, and it has a simpler answer. Tom's ability to be guided by the fact that exercise is beneficial, to respond to it rationally, or take it into account, does *not* change in these circumstances, since (as I put it earlier) the ability to be guided by a fact is not the ability to engage in a specific activity or range of activities. What a person's knowledge of a fact is an ability to do is captured by an adverb, rather than a verb. Hence, Tom's ability to be guided by the fact that exercise is beneficial is one thing, and the range of activities in which he is able to display it is another. The latter changes, but the former remains the same. Dispositions and tendencies exhibit this distinction too. For example, if an unpunctual man starts earning enough to file a tax return, and tends to file it late, his unpunctuality does not change, but he begins to display it in a new (though far from novel) way.

However, the ability theory of knowledge does not imply that a person's ability to be guided by the facts *cannot* change. Abilities change or develop in two main ways: they improve (or decline) and they broaden (or narrow). As Williamson points out, most human abilities can be improved. This is true at least of those that are classified as skills. In fact it should be obvious that abilities we acquire by practice can be improved, whether they are (or depend on) motor skills or intellectual skills, since we improve them as we acquire them, and we can generally improve them further, once we have reached the threshold that allows us to be described as possessing them without qualification. Think of walking or talking, for example. Many abilities can also be broadened or extended. This means different things in different cases, but the general idea is that an ability has a potential field of operation, and someone broadens an ability when she can exercise it across a more extensive field. For example, a pianist can broaden his ability to play the piano by mastering a more varied repertoire, and as we saw earlier someone who believes that exercise is beneficial can broaden his ability to take exercise by learning to swim.

If knowledge is the ability to be guided by the facts, we can easily understand how a person's knowledge in general, and her knowledge within a specific domain in particular, can be broadened. For as her knowledge becomes more extensive, she becomes able to be guided by a more extensive range of facts. As for the improvement of knowledge, knowing a subject better seems to involve both knowing more about it and understanding it better, that is, acquiring a more systematic knowledge of it, and an improved ability to connect the different parts of it and grasp it as a coherent whole. Admittedly, it is hard to see how a person's knowledge of a specific fact can be improved or broadened, although we can improve our *understanding* of a fact, if it has significant ramifications. But the same applies to other intellectual abilities, such as the ability to translate from one language into another, or the ability to do sums. For example, a student can improve her ability to translate from French into English, that is, she can learn to produce better translations, and she can also learn to translate a wider variety of

texts. But it is hard to see how she can improve or broaden her ability to translate 'onze', 'lune', or 'Angleterre'.