

# Knowledge, Action, and Defeasibility

Carlotta Pavese

This is a draft. For *Reasons, Justification, and Defeaters*,  
edited by Jessica Brown and Mona Simion,  
Oxford University Press 2020.

## 1. Introduction

One can intentionally do something only if one knows what one is doing while one is doing it. For example, one can intentionally kill one's neighbor by opening their gas stove overnight only if one knows that the gas is likely to kill the neighbor in their sleep. One can intentionally sabotage the victory of one's rival by putting sleeping drugs in their drink only if one knows that sleeping drugs will harm the rival's performance. And so on. In a slogan: *Intentional action is action guided by knowledge*.

Although prominent philosophers have found this idea platitudinous (Hampshire 1959) and others have been at pains to argue for it (Gibbons 2001, Pavese 2019, Beddor & Pavese 2019), some of the crucial motivations for it are underappreciated in contemporary debates both in epistemology and action theory. This essay reviews some arguments for a 'knowledge-centered psychology' — a psychology where knowledge enters center stage in an explanation of intentional action — and discusses some consequences for the debate on the *defeasibility of know-how*.

The essay is divided into two parts. §2 summarizes some considerations in favor of a knowledge-centered psychology and defends its core idea against a recent attack by Cath (2015); §3 argues that a knowledge centered psychology motivates intellectualism about know-how and in particular it supports the claim that the epistemic profile of know-how is the same as that of knowledge.<sup>1</sup> The focus will be on whether know-how is defeasible in the way knowledge is often supposed to be defeasible in contemporary epistemology. A knowledge centered psychology predicts that the defeasibility of know-how patterns with that of knowledge: know-how and knowledge are both defeated when one's ability to intentionally act is defeated. By replying to a challenge raised in the recent literature by Carter & Navarro (2018), I argue that this prediction is actually borne out.

## 2 Towards a knowledge-centered psychology

### 2.1 Explaining attempts versus explaining intentional successes

A long tradition in the philosophy of mind assigns beliefs a central role to play in folk psychological explanations of intentional behavior (e.g., Stich 1978; Fodor 1987; Lewis 1976; Stalnaker 1984). More or less explicitly, this tradition confines psychological explanations to an explanation of *attempts*. To provide an illustrative example, consider

---

<sup>1</sup> For a more extended discussion of the claim that a knowledge centered psychology supports intellectualism about know-how, see Pavese 2018, 2020. For an extension of the argument to intellectualism about *skills*, see Beddor & Pavese 2019.

the usual example of a psychological explanation, where one's belief that there is water in the fridge and one's desire to drink it together are supposed to explain one's attempt to grab a bottle of water from the fridge. Success happens when the world cooperates — when there is indeed water in the fridge. If one's belief is true, then one will succeed at finding a bottle; if one's belief is false, one will not succeed at finding water. The dominant thought behind a belief-centered psychology is the idea that, as far as the *psychological* explanation of behavior goes, whether the world complies (e.g., whether there is water in the fridge) is irrelevant: what is to be explained is the fact that *one attempted to get water from the fridge*, whether or not one has succeeded. And one's belief that there is water in the fridge, together with one's desire to drink it, suffices to explain one's attempt, whether or not the belief is true.<sup>2</sup>

The assumption that psychological explanation should be confined to explaining attempts, rather than successes, relies on the idea that actions are *decomposable into mental "conditions" and non-mental "conditions"* — into attempts, on one hand, and into bodily movements, causal chains, or successes on the other. On this picture, attempts are supposed to be *narrow* mental conditions. A little terminology, from Williamson 2000, might help. Let us a *condition* be something that obtains or fails to obtain at a case, and let a *case* be a centered world: an ordered set of a possible world, a time, a place, and an agent. A condition C is narrow if and only if for all cases a, b, if a is internally alike b and b is internally like a, then C obtains at a if and only if C obtains at b. Two cases are internally alike if and only if the total internal state of an agent in a is exactly alike that of the agent in b. Another way of saying this is: Narrow mental conditions supervene or are determined by internal physical states: *no difference in whether they obtain without a difference in the internal state of the agent*.

Now, attempts are supposed to be 'narrow' mental conditions: whether they obtain or not at a case depends entirely on the physical internal state of their agent. As narrow mental conditions, an attempt does not encompass those external aspects of the world that make for an agent's success. Hence, because narrow, they are not 'factive' in the sense that an attempt can be successful or fail. So an attempt to F does not entail successful F-ing.

Suppose intentional action does arise when narrow mental conditions combine with non-mental conditions — i.e., the subject's bodily movements and whatever else the external world contributes to the success of the agent. A psychological explanation of intentional behavior, understood as attempts, does not need to appeal to anything more than to narrow mental states — i.e., nonfactive mental states such as mere beliefs — for they seem like good candidates to explain and ground narrow mental conditions. If when we are explaining behavior, all we are trying to explain is an attempt, such as the attempt at finding a bottle of water in the fridge, rather than the success at finding it, then all we need is a psychological theory that encompasses non-factive attitudes such as belief.

---

<sup>2</sup>That a belief-centered psychology aims to explain attempts rather than successes is made explicit in Stich 1998.

Suppose however that intentional actions were *not* decomposable into a mental narrow component and a nonmental component. In particular, suppose actions could be intentional — and hence mental — beyond the attempts' contribution. Then this residual 'mentality' of actions too would call for an explanation and presumably one explanation that would reduce the intentionality of the action to some mental state of the agent. Now, for this sort of explanation, belief alone can not suffice.<sup>3</sup> For, if the mental condition to be explained is not narrow (like attempts) but broad (like intentional successes), then its explanation calls for a non-narrow condition — i.e., or broad condition — one that is world-encompassing.

Indeed, there are good reasons to think that actions are not decomposable in mental and nonmental components, and even if attempts were in some sense components of actions, the mentality of actions would not be exhausted by the mentality of attempts. Here is an argument for this conclusion. If attempts exhausted the mentality of actions, then provided that one attempted to perform an action  $\phi$ , one's eventual success at  $\phi$ -ing would have to be intentional. For on this picture, the intentionality, and hence the mentality, of actions would be exhausted by their attempt. However, there are a variety of cases in which one attempts at  $\phi$ -ing, succeeds, and yet their action fails to be intentional. That suggests that the intentionality of actions cannot be reduced to the intentionality of attempts; and intentionality being one mark of the mental, that suggests that the mentality of actions cannot be reduced to the mentality of attempts.

A well-known example in action theory illustrating how attempts and intentionality can come apart is **Mary the Bomber** (cf. Mele and Moser 1994; Gibbons 2001; Beddor & Pavese 2019). Mary intends to kill her uncle by a bomb in his house and then, after moving a safe distance away, pressing the large red button on the remote control device. She does not know much about how these things work and thinks that pressing the button will cause the bomb to detonate but has no idea about the details of this process. Her belief is true and, we can suppose, justified. But here is what happens. A satellite, launched by the National Security Agency and designed to prevent bombings of just this kind, intercepts Mary's transmission; this causes the satellite to send a warning to the intended victim; but, because of an unfortunate choice of frequency, this causes the bomb to detonate. Mary killed his uncle and caused the bomb to detonate and did intend both things. But she did not do either of these things intentionally.

In this example, Mary attempts to perform an action (killing her uncle) and succeeds but intuitively her action is not intentional. Hence, the intentionality of this action is not reducible to the intentionality and mentality of the attempt.<sup>4</sup> Examples like **Mary the Bomber** suggest that intentional actions cannot be decomposed into a 'narrow' mental component and a non-mental component, where the 'narrow' mental component

---

<sup>3</sup>Cf. also Levy (2013) and Williamson (2015). Gibbons (2001) also argues that psychological explanations should explain successes and not merely attempts.

<sup>4</sup> Cath (2015) objects to Gibbons' example being a case of intentional action without knowledge. But see Pavese (2018) for a reply.

is mere attempt or trying and where the combination of the narrow mental component and the non-mental component somehow issues intentional successes: the intentionality of actions is a broad, rather than a narrow, condition. When it comes to explaining a broad, rather than narrow conditions, a kind of psychology where a broad mental condition such as knowledge enters center stage is more promising than a kind of psychology centered on a narrow mental condition such as belief: *narrow can be explained with narrow. But broad ought to be explained with broad.*

Note that the **Mary the Bomber** above can be accounted for by a knowledge-centered psychology. Mary does not really know that she can provoke the explosion by implementing her plan. That is why her success is too coincidental to count as intentional. More generally, the prediction of a knowledge-centered psychology is that if one's belief is Gettiered, then one cannot act intentionally on that belief. This prediction is borne out. To illustrate, consider an adaptation of the fake barns case (Goldman 1976): Daniel is instructed to stop at a barn that he finds on the road to Larissa and wait for further instructions there. The road to Larissa goes through fake barn county. Daniel passes the first barn-looking construction but does not stop for he thinks "There is not enough shadow for me to wait on a car." At the second barn looking construction, Daniel stops and parks. As it turns out, only the second barn-looking construction was a real barn. So Daniel ends up stopping at a barn and intended to do so. However, he did not intentionally stop at a barn: in fact, he would have easily stopped at a fake barn, had the shadow only been present there.<sup>5</sup>

In both cases, their success is too coincidental to count as intentional. More generally, if one possesses knowledge, then one's belief cannot be lucky (Sosa 1999; Williamson 2000).<sup>6</sup> Hence, a knowledge-centered psychology can explain why luck can undermine the intentionality of Mary's and Daniel's successes: by undermining their knowledge. And it is independently plausible that an action based on knowledge is sufficiently under one's control to count as intentional.

## 2.2. *Intentional Action does require knowledge*

And yet some disagree. In a recent paper, Cath (2015, 11) argues that one can have intentional action without knowledge, upon considering cases like the following:

**Bob the Pilot.** Bob wants to learn how to fly in a flight simulator. He is instructed by Henry. Unknown to Bob, Henry is a malicious imposter who has inserted a randomising device in the simulator's controls and intends to give all kinds of incorrect advice. Fortunately, by sheer chance the randomising device causes exactly the same results in the simulator as would have occurred without it, and by

---

<sup>5</sup> This argument is developed in much more detail in Pavese (2020) and Beddor & Pavese (2019). Notice that although the subject intentionally stops and stops at a barn, he does not intentionally stop *at a barn*. There might be a *de re* reading of the intentionality ascription which is true. But the *de dicto* reading is definitely false.

<sup>6</sup> Some object to a modal requirement on knowledge. See Beddor & Pavese (2018) for a recent defense.

incompetence Henry gives exactly the same advice as a proper instructor would have done. Bob passes the course with flying colors. He has still not flown a real plane. Bob has a justified true belief about how to fly, but that justified true belief does not amount to knowledge.

Cath uses this example to argue that one can have know-how even though one's relevant belief is Gettiered, precisely because, he contends, Bob can intentionally fly, even though he does not know the instructions for flying.<sup>7</sup> And for some action F, if one can intentionally F, then one knows how to F (Williamson & Stanley 2001, Hawley 2003, Setiya 2011, Cath 2015, Pavese 2018, 2020).

Cath (2015, 11) thinks that this diagnosis is supported by comparing the case of Bob with the case of Joe, who is a near perfect counterpart of Bob except that he non-controversially knows how to fly: his simulator operated correctly and did so non-accidentally; his instructor intentionally gave him the correct advice, etc. If Joe were to try to fly a plane in normal circumstances he would typically succeed in so doing and his successful actions would be unquestionably intentional actions. And it is an implicit stipulation of the flight simulator case that if Bob were to try to fly a plane in normal circumstances then he would be just as likely to succeed as Joe. Cath contends that not only would Bob succeed as often as Joe but, like Joe, his actions would appear to have all the standard kind of properties that are thought to distinguish merely successful actions from intentional actions. Hence, according to Cath, it is very plausible that Bob's successful actions of flying, like Joe's, would be perfectly under his control or guidance as he performs them.

In order to assess Cath's (2015) argument, let us just focus on Bob's very first successful attempt at flying a plane (because of course Bob might gain more evidence as he keeps practicing, and so his Gettier belief might turn into knowledge after sufficiently many successful attempts at flying). Was that success intentional in the sense of being under Bob's control? Both intuitive and theoretical considerations suggest that the answer to this question ought to be "No".

First come the intuitive considerations. Compare Joe and Bob. Note that, strictly speaking, *Joe knows what he is doing while he flies the plane*: he knows that he is following instructions that are conducive to successful flying. By contrast, by assumption, Bob does not know that. Hence, he does not know what he is doing while he flies the plane. But consider how unintuitive it is to let one's action to be intentional when one lacks knowledge of what one is doing:

---

<sup>7</sup> This example is initially due to Stanley & Williamson (2001) to illustrate that know-how is incompatible with epistemic luck.

**Aweful** Bob intentionally landed the plane but he did not know that he was landing the plane by following the given instructions.

If one denies that intentional action is guided by knowledge, one is committed to **Aweful** being possibly true. It is simply not true that intuitions stand on Cath's side here.

No important philosophical claim should be motivated merely by intuitions, for they are too unstable and possibly theoretically driven to be conclusive. So here I'd like to put forward and explore a novel theoretical argument to the effect that, *contra* Cath, one's actions cannot be intentional unless they are guided by knowledge.<sup>8</sup>

Cath assumes intentional action ought to be under one's control and assumes that Bob's success are under his control because the success are under "his guidance." Sure, Bob causes the landing and is, to some extent, responsible for it. In this sense, his success is *under their guidance*. But like in **Mary the Bomber**, the belief in question is Gettiered and so happens to be true by luck. If so, both Bob and Mary could have the beliefs they have even if they were false. In this case, an attempt of theirs based on those beliefs would fail. Hence, it could easily have happened that their attempts would be unsuccessful. If an action could have easily failed then it is lucky. And how can a lucky action be under one's control?

More precisely, here is the argument step by step:

- (a) Intentional action is action under one's control. (Premise)
- (b) If an action is lucky then it is not under one's control. (Premise)
- (c) An action is lucky if it could easily not have happened. (Premise)
- (d) A Gettiered belief could easily have been false (Premise).
- (e) An action based on a Gettiered belief could easily not have happened (in the close worlds where the belief is false) (Premise).
- (f) Hence an action based on a Gettier belief is lucky. (c, e modus ponens).
- (g) Hence, an action based on a Gettier belief is not under one's control (b,f, modus ponens).
- (h) Hence, an action based on a Gettier belief cannot be intentional (g and a, modus tollens).

As rarely it happens in philosophy, we have here a deductively valid argument to the effect that intentional actions ought to be based on knowledge. Is the argument sound?

Premise (a) relies on rather minimal assumptions about the nature of action, widely endorsed in action theory (e.g., Mele & Moser 1994, Gibbons 2001), according to

---

<sup>8</sup> Greco (2016) considers a different sort of argument for thinking that knowledge is necessary for explaining action, one that relies on the nature of explanation. See also Pavese (2018).

which intentional action is an action that is under the agent's control. (This premise is actually granted by Cath 2015, 10-11).

Premise (b) is the rather uncontroversial claim that if a certain success is lucky, then it is not under the control of the agent. This is in fact a widespread assumption in the debate on moral responsibility and moral luck (Nagel 1979, Williams 1981, 1993).<sup>9</sup> It is of course a difficult, and a context-sensitive matter, what counts as “under one’s control.” But for our purposes, we do not need to settle this question, the idea being that if a certain success counts as lucky, it is not *sufficiently* under the control of the agent *to count as intentional*. Premise (c) employs a modal definition of lucky event that is rather standard in epistemology (e.g., Pritchard 2005, Sosa 2007), according to which an event is (too) lucky if it could (too) easily not have happened.

The most controversial premise of the entire argument is, I take it, Premise (d) — i.e., the claim that if a belief is Gettiered, then it is unsafe. According to this premise, safety is sufficient for a belief not to be Gettiered. Although a full defense of this claim would bring me too far afield here, let me refer elsewhere for a systematic defense of the sufficiency of safety (Beddor & Pavese 2018:8-12). The main challenges come from putative counterexamples to the sufficiency of safety for knowledge, such as Pritchard’s (2012, 260) Temp case, where the subject forms their beliefs about the temperature in the room by consulting a broken thermometer, one that is manipulated by a hidden agent in the room who is in control of the thermostat and whose job it is to ensure that every time Temp consults the thermometer, its reading corresponds to the temperature in the room. Temp’s belief is allegedly justified, true and safe, but not knowledge. Beddor & Pavese (2018, 8-12) argue that Temp’s belief is unsafe because it is false in many of the close worlds where there is a broken thermometer but no angelic aid and those worlds are relevant for assessing the belief’s safety. More generally, we argue that if the sphere of relevantly close worlds is sensitive to what conditions are as *normal as the conditions at the world under consideration*, then many (if not all) counterexamples to the sufficiency of safety simply disappear. Moreover, a sufficient modal condition on knowledge can be shown to fall out from a parallel between skillful action and knowledge. Hence, there are independent strong theoretical reasons to think that Premise (d) is true.

Finally, Premise (e) is the uncontroversial claim that if an action were based on a Gettiered belief, it could have failed in some of those close worlds where the belief is false. With Premise (a) to Premise (e) in play, the conclusion deductively follows by *modus ponens* and *modus tollens*.

### 3. The Defeasibility of Know-How

#### 3.1 From a Knowledge Centered Psychology to Intellectualism about Know-How

---

<sup>9</sup> This is in fact a widespread assumption in the debate on moral responsibility and moral luck (Nagel 1979, Williams 1981, 1993). I cannot discuss a parallel argument for moral responsibility requiring knowledge but I do so elsewhere (Pavese 2019).

Hence, a knowledge-centered psychology can be motivated on the basis of intuitions about cases. But there are, also, more theoretical arguments for a knowledge-centered psychology, ones that rely on rather minimal assumptions about the nature of Gettiered beliefs, control, and intentional actions.<sup>10</sup>

A knowledge-centered psychology motivates intellectualist views of know-how, skills, and action (cf. Pavese 2018, 2020, Beddor & Pavese 2019). First note that the kind of knowledge that, on a knowledge-centered psychology, explains intentional action is exactly the same kind of knowledge that, on a broadly intellectualist picture, is required by know-how. To see this, consider the kind of knowledge that would be needed for explaining intentional action. Start with Goldman's (1970) action theory, according to which one intentionally  $\phi$ s when one has a plan to  $\phi$ , where a plan to  $\phi$  is a belief that specifies the means to  $\phi$  (cf. also e.g., Audi 1986; Bratman 1987; Ginet 1990; Harman 1976; Velleman 1989/2007; Mele & Moser 1994):

*(Intentionality/Belief)* If  $s$  intentionally  $\phi$ s, then there are some means  $m_1, \dots, m_n$  to  $\phi$  such that  $s$  truly believes that  $m_1, \dots, m_n$  are means for oneself to  $\phi$ .

*(Intentionality/Knowledge)* can be formulated along the same lines:

*(Intentionality/Knowledge)* If  $s$  intentionally  $\phi$ s, then there are some means  $m_1, \dots, m_n$  to  $\phi$  such that  $s$  knows that  $m_1, \dots, m_n$  are means for oneself to  $\phi$ .

Now, according to standard formulations of intellectualism (Stanley and Williamson 2001; Stanley 2011; Pavese 2015, 2017), one knows how to  $\phi$  only if, for some means  $\psi$  to  $\phi$ , one knows that  $\psi$  is a means for one to  $\phi$ :

*(Intellectualism about Know-How)*  $s$  knows how to  $\phi$  is at least in part a matter of knowing, for some means  $\psi$  to  $\phi$ ,  $s$  knows that  $\psi$  is a means for oneself to  $\phi$ .<sup>11</sup>

Hence, the knowledge that *(Intentionality/Knowledge)* requires for intentional action is the same that intellectualists require for know-how.

In fact, a knowledge-centered psychology and intellectualism about know-how are strictly connected views, supporting each other. Start from *(Know-How/Intentionality)*, endorsed by many scholars, intellectualists and anti-intellectualists alike (Ryle 1949;

---

<sup>10</sup>For other sorts of arguments for a knowledge-centered psychology that highlight the modal profile of explanations of action, see Greco (2016) and Pavese (2018). For more theoretical considerations in favor of a knowledge-centered psychology, see Nagel (2013).

<sup>11</sup> I am stating intellectualism as the view that know-how *requires* knowledge of the means for action. Stating intellectualism as a fully reductive claim would require talking about practical modes of presentation, which I cannot discuss here.



Stanley and Williamson 2001; Stanley 2011; Hawley 2003; Hornsby 2011; Setiya 2013; Pavese 2018):

*(Know-how/Intentionality)* If  $s$  intentionally  $\varphi$ s,  $s$  knows how to  $\varphi$ .

Among the motivations behind *(Know-how/Intentionality)*, there is that operations which cannot be performed intentionally, such as digesting, are ones that one cannot know how to perform (Williamson & Stanley 2001) and the consideration that manifestations of know-how seem to be characteristically intentional: as Ryle (1949) put it, what distinguishes the clumsy person, who by accident falls and tumbles, and the skillful clown is that the latter, but not the former, falls and tumbles on purpose.

Further, suppose that *(Intentionality/Knowledge)* is true and so that the intentionality of an action is to be explained at least in part in terms of propositional knowledge. Then by *(Know-how/Intentionality)* and *(Intentionality/Knowledge)*, we get that if one intentionally  $\varphi$ s, one both knows how to  $\varphi$  and one has propositional knowledge of some means to  $\varphi$ :

*(Know-how, Intentionality, Knowledge)* If  $s$  intentionally  $\varphi$ s,  $s$  both knows how to  $\varphi$  and for some means  $m_1, \dots, m_n$ ,  $s$  knows that means  $m_1, \dots, m_n$  are means for oneself to  $\varphi$ .

The intellectualist picture provides the best explanation for why *(Know-How, Intentionality, Knowledge)* should hold. According to this explanation, *(Know-How, Intentionality, Knowledge)* is true not just out of a coincidental aligning of propositional knowledge and know-how in intentional action. Rather, its truth is grounded on the very nature of know-how.

For these reasons, although intellectualists have mostly appealed to a linguistic argument for motivating their views (Stanley & Williamson 2001, Stanley 2011), I am inclined to think that the main motivation for the view does not come from linguistics but, rather, from the sort of action theory that a knowledge-centered psychology recommends.<sup>1213</sup>

There is also an argument that goes in the other direction — i.e., from intellectualism to a knowledge-centered psychology. Suppose that both *(Intellectualism about Know-How)* and *(Know-How/Intentionality)* are true. Then, the intentionality of an

---

<sup>12</sup> For worries concerning the linguistic argument on behalf of intellectualism, see Brown (2014).

<sup>13</sup> Much more could be said about how best to understand the sort of knowledge that one must possess, in accordance with *(Intentionality/Knowledge)*. In Pavese (2020), I propose that in order to overcome several worries about overintellectualizing intentional action (cf. Setiya 2011), we should understand knowing that certain means are means for one to perform an action in *probabilistic terms*, as knowing that one is sufficiently likely to perform that action by those means, and I propose that this knowledge ought to be understood along the lines of Moss' (2018) probabilistic knowledge.

action by an agent requires knowledge of the means to perform it. Hence, intentional action is guided by knowledge — i.e., (*Intentionality/Knowledge*) follows. A knowledge-centered psychology not only invites an intellectualist view of know-how and action but is also motivated by it.

### 3.2. *Knowledge, Defeasibility, and Know-How*

So far I have argued on behalf of a knowledge centered psychology by looking at the role of knowledge in explaining intentional action. Knowledge-centered psychology naturally goes together with an intellectualist picture of know-how and vindicates the relation between know-how and intentional action.

Now, the intellectualist picture motivated by a knowledge-centered psychology and outlined in the last section makes a very clear prediction: that both knowledge and know-how are defeasible to the extent to which the corresponding ability to intentionally act is defeated. Against this prediction, Carter & Navarro (2018) have argued that the defeasibility of know-how does not go together with the defeasibility of knowledge. They use this claim to argue against the intellectualist claim that know-how consists in a state of propositional knowledge. We are now in position to address the challenge, for it turns on a failure to appreciate the relation between knowledge and intentional action.

Carter & Navarro (2018, 666) propose the following example:

**Ana and the Grenade Factory.** Ana and Maria work in a grenade factory during the Spanish Civil War. They are thoroughly instructed when hired, with examples and practical explanations. By controlled trial and error, they learn their job, and both continue working at the factory for years, believing they are making working grenades. However, one day each comes to realize that the other is making grenades in an importantly different way, and they identify the origin of the problem: as it turns out, the instructions were ambiguous and allowed for two different interpretations. The instructors were not aware of this, and there is nobody above them now who may say who is right. Given that the grenades may only be used in battle, which is very far away, neither Ana nor Maria knows whose grenades actually work, and so there is no way to find out who is making them the right way. As a matter of fact, Ana got the instructions right (she produces grenades in way *w*, which is the correct way); she is very successful in producing grenades that later work perfectly. It is Maria who got something wrong (she makes them in *w'*, the possible interpretation of the instructions that the instructors did not foresee), and her grenades are always duds. Unaware of this, both have reasonable doubts they did not have before but they have to keep on working.

According to Carter & Navarro (2018), before receiving information of how her knowledge has been acquired (call this piece of information ‘MISLEADING’), Ana

might know, for some means to make grenades, that it is a means to make grenades, but her knowledge is defeated as soon as the misleading evidence is acquired. On the other hand, they think that Ana still knows how to make grenades after receiving MISLEADING. If they were right, this would be a case where know-how stands undefeated when the corresponding knowledge is instead defeated. They conclude (2018, 669):

If know-how really were a case of know-that, we should expect it to be defeasible by the same kinds of mechanisms by which propositional knowledge is defeated. But it is not. In other words: garden variety defeaters of knowledge-that do defeat the knowledge agents have about the ways in which they do what they do.<sup>14</sup>

Carter & Navarro (2018)'s argument hinges on the claim that Ana still knows how to make grenades, upon receiving MISLEADING. Why think that? They argue that Ana retains her knowledge how to make grenades, for evidently she still maintains *the ability to make grenades* (2018, 666):

The claim that Ana preserves her know-how along all the variations of the case is supported by the fact that she is still able to make grenades proficiently, and the doubts she acquires do not seem to imperil this ability in any relevant sense.

Now, Carter & Navarro (2018) are going *far* too quick here. Granted, Ana still preserves *some ability* that is relevant to grenades-making. What is much less obvious, and as I argue, ultimately incorrect, is to assume that Ana preserves the sort of ability that goes together with know-how. Let me explain.

As the discussion in the previous section already suggests, know-how does not just go together with *any ability*. It goes together with *the ability to intentionally perform a certain task*. For example, knowing how to make risotto does not merely go together with the ability to make risotto but with the ability to *intentionally* make risotto. For if one had the ability to make risotto but lacked the ability to intentionally make it, one would not count as knowing how to make risotto. This point is well-known in the literature since Hawley (2003) and is accepted by both intellectualists and anti-intellectualists (Ryle 1949, Setiya 2011, Pavese 2017). The clumsy person has the ability to fall and tumble, as they reliably do so. But only the clown has the ability to intentionally do that. As another illustrative example, Susie may have the ability to irritate Ben, for she would succeed at irritating him if she tried. But suppose she falsely believes that it is the smell of the smoke, rather than the noise she makes whenever she smokes, that irritates Ben. In this

---

<sup>14</sup>Before we start assessing the argument, note that the argument assumes that knowledge can be easily defeated by high-order evidence. This assumption is controversial and is not granted by prominent epistemologists (Aarnio 2010, 2014). Let us play along, however, and see that the challenge rests on other false assumptions about the nature of know-how.

case, she does not intentionally irritate Ben: her success is too coincidental to count as intentional. Nor does she have the ability to irritate Ben. Because of this, it seems that she does not know how to irritate Ben.

On the bases of similar examples, Intellectualists and anti-intellectualists alike endorse the claim that know-how goes together with the ability to *intentionally* perform the task. The crux of the recent debate has been precisely whether one can characterize the ability to intentionally perform the task independently of knowledge (cf Setiya 2011; Pavese 2018, 2020).

Suppose it is true that know-how goes together with the ability to intentionally perform a task. If so, it is independently plausible that there is an important sense in which upon receiving MISLEADING Ana does not know any longer how to make grenades. For upon receiving MISLEADING, when asked to make grenades, she will be at loss. She will even stop performing at the workplace, until she is told that she has been making grenades correctly all along. Suppose she were forced to reproduce whatever process she would initiate before MISLEADING. She would unknowingly succeed at making grenades. But the success would be too out of her control to count as intentional. She is still able *in some sense* to make grenades but in an important sense she now lacks the ability to intentionally make grenades. If so, then she also lacks her know-how.

If the reader is not yet ready to grant this conclusion, it is because, actually, things are more complex and some additional distinctions are called for. Ascriptions of ability to intentionally perform a task are *opaque*, for as it is well known in action theory, “intentionally” is an intensional operator (Davidson 1968, Goldman 1970, Pavese 2015a). Lois might intentionally kiss Superman but not Clark Kent. Because of the opacity of intentionality reports, it is paramount that we distinguish between (*de re ability*) and (*de dicto ability*):

(*de re/ability*) There is some task *t* that is in fact the task of making grenades such that Ann has the ability to intentionally perform *t*.

(*de dicto/ability*) Ana has the ability to intentionally make grenades.

Now, with this distinction in play, consider again Ana’s situation upon receiving MISLEADING. Ana does not have (*de dicto/ability*) any longer. For one to have the (*de dicto/ability*), one needs to be able to make grenades *on demand* (if asked to make grenades, she would). Ana does not have that ability, for if, after MISLEADING, she were asked to make grenades, she would now be at loss. If know-how goes with the ability to intentionally perform a task, then, to this distinction between (*de dicto/ability*) and (*de re/ability*), there corresponds the following distinction:

*(de re/know-how)* There is a task  $t$  that happens, unknown to Ann, to be the task of making grenades, such that Ana still knows how to perform  $t$ .

*(de dicto/know-how)* Ann knows how to perform grenades.

*(de re/know-how)* and *(de dicto/know-how)* are different know-hows, as they go along with different dispositions to behavior. One might have *(de re/know-how)* even if one has no idea that what one is doing when doing  $t$  is making grenades. Suppose, for example, one is simply instructed to follow a certain procedure but has no idea of its outcome. In this case, one might have *(de re/know-how)* without *(de dicto/know-how)*. This is plausibly Ana's quandary: Because Ana still knows how to execute whatever task she was executing before MISLEADING, she plausibly still have *de re* know-how. After all, if she were told at the workplace to do whatever she was doing before she received MISLEADING, and she obeyed the order, she would intentionally perform a task, which, unknown to her, is the task of making grenades. Hence, Ana preserves *(de re/ability)* upon receiving MISLEADING. So Ana plausibly retains her *(de re/know how)* upon receiving MISLEADING.

However, Anna does lose *(de dicto/know-how)*. For one to possess *(de dicto/know-how)*, it is not sufficient to possess *(de re-ability)*; one would need in addition to have the corresponding *de dicto* ability, which as we have seen she lacks. While Ana loses *(de dicto/know-how)* and *(de dicto/ability)* upon receiving MISLEADING, Ana preserves *(de re/know-how)* and *(de re/ability)*, for she still knows how to do whatever it was that she was doing before (which as far as she knows, it is not grenades!), and she still preserves the *de re* ability to make grenades upon receiving the misleading information.

Crucially, intellectualists can accept all of this. According to intellectualism, *(de re know-how)* requires *(de re knowledge)*, while *(de dicto/know-how)* requires *(de dicto knowledge)*:

*(de re/knowledge)* There is a task  $t$  that is in fact, but unknown to Ana, the task of making grenades such that Ana knows for some way  $w$  that  $w$  is a way to execute  $t$ .

*(de dicto/knowledge)* Ana knows for some way  $w$  that  $w$  is a way to make grenades.

Although Ana loses *(de dicto/knowledge)*, Ana arguably still also preserves *(de re/knowledge)*. MISLEADING only defeats (if anything) *(de dicto/knowledge)* — i.e., the knowledge that the procedure she was implementing was for making grenades. *(De re/knowledge)*, instead, is not at all defeated by MISLEADING: Ana continues to have it, as she might continue to know what procedure she was implementing, before receiving

MISLEADING, when she was intending to make granades. And so, by intellectualism's lights, she might continue to know how to make whatever she was making when she thought (correctly, it turns out!) that she was making granades: by intellectualism's lights, because Ana preserves (*de re/knowledge*), Ana can preserve (*de re/know-how*) as well as (*de re/ability*).

In conclusion, Carter & Navarro (2018) fail to raise a challenge against intellectualism. Their presumed challenge turns on the failure to appreciate the relation between knowledge, know-how, and intentional action. Because of that, they fail to distinguish between different sorts of abilities that go together with know-how. Once one appreciates that know-how goes with the ability to act intentionally, then, because ascriptions of this sort are opaque, it becomes paramount to distinguish between *de re* abilities and *de dicto* abilities. With this distinction comes the corresponding distinction between different sorts of know-hows and between the different sorts of knowledge that Ana does preserve upon receiving MISLEADING. As we have seen, against Carter & Navarro (2018), the sort of *de re* abilities Ana does preserve can be fully accounted for on a picture on which know-how is knowledge. And those *de dicto* abilities that she does lose are also correctly predicted to go lost on the same intellectualist picture.

Far from coming apart in their pattern of defeasibility, know-how and knowledge go hand in hand, just as one would expect on the sort of intellectualist picture that is motivated by a knowledge-centered psychology.

## References

- Aarnio, M. L. (2010). "Unreasonable knowledge". *Philosophical Perspectives*, 24, 1–21.
- , (2014). "Higher-order evidence and the limits of defeat," *Philosophy and Phenomenological Research*, 88(2), 314–345.
- Carter, A. J., & Navarro, J. (2018). "The Defeasibility of Knowledge-How," *Philosophy and Phenomenological Research*, 95(3), 662–685.
- Audi, R. (1986). "Intending, Intentional Action, and Desire," in Joel Marks, ed., *The Ways of Desire*, Chicago: Precedent: 17–38.
- Beddor, B., & Pavese, C. (2018). "Modal Virtue Epistemology," *Philosophy and Phenomenological Research*, Early View.
- , (2019), "Skill as knowledge", manuscript.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information.
- Brown, J. A. (2013). "Knowing-how: Linguistics and cognitive science," *Analysis*, 73(2), 220–227.
- Carter, A. & D. Pritchard (2015a). "Knowledge-How and Epistemic Luck," *Noûs* 49 (2015), 440–53.

- . “Knowledge-How and Cognitive Achievement,” *Philosophy and Phenomenological Research* 91, 181–99.
- . “Knowledge-How and Epistemic Value,” *Australasian Journal of Philosophy* 93, 799–816.
- Cath, Y. (2015). “Revisionary Intellectualism and Gettier,” *Philosophical Studies* 172.1: 7–27.
- Dretske, F. (1988). *Explaining Behavior*. MIT Press.
- Ginet, C. (1990). *On Action*. Cambridge: Cambridge University Press.
- Gibbons, J. (2001). “Knowledge in Action,” *Philosophy and Phenomenological Research* 62(3):579–600.
- Goldman, A. (1970). *A Theory of Human Action*. Englewood Cliff, NJ. Prentice Hall Inc.
- (1976). “Discrimination and perceptual knowledge,” *The Journal of Philosophy*, 73(20), 771-791.
- Greco, D. (2016). “Safety, Explanation, Iteration,” *Philosophical Issues* 26.1: 187–208.
- Hampshire, S. (1959). *Thought and Action*, London: Chatto and Windus.
- Harman, B. (1976). “Practical Reasoning,” *The Review of Metaphysics* 79: 431–63.
- Hawley, K. (2003). “Success and knowledge-how,” *American philosophical quarterly*, 40(1), 19-31.
- Hawthorne, J. & Stanley, J. (2008). “Knowledge and Action,” *Journal of Philosophy* 105.10: 571–590.
- Levy, Y. (2013). “Intentional Action First,” *Australasian Journal of Philosophy* 91(4): 705–718.
- Lewis, D. (1974). “Radical Interpretation,” *Synthese* 27:331–344.
- , (1979). “Attitudes *de dicto* and *de se*.” *The Philosophical Review*, 88(4), 513-543.
- Mele, A. & Moser, P. (1994). “Intentional Action,” *Noûs* 28:39–48.
- Moss, S. (2018). *Probabilistic Knowledge*. Oxford/NY: Oxford University Press.
- Nagel, T. (1979). *Mortal Questions*, New York: Cambridge University Press;
- Nagel, J. (2013). “Knowledge as a Mental State,” *Oxford Studies in Epistemology* 4:275–310.
- Pavese, C. (2015a). “Knowing a rule”, *Philosophical Issues, A Supplement to Nous*, 25(1): 165–188
- , (2015b). “Practical Senses,” *Philosophers’ Imprint* 15.29:1–25.
- , (2017). “Know-How and Gradability,” *Philosophical Review* 126.3:345–83.
- , (2018). “Know-How, Action, and Luck,” *Synthese*, 1–23.
- , (2020). “Probabilistic Knowledge in Action,” *Analysis*.
- , (2019). “Knowledge and mentality”, manuscript.
- Pritchard, D. (2005). *Epistemic luck*. Clarendon Press.
- (2012). “Anti-luck Virtue Epistemology.” *The Journal of Philosophy*, 109(3), 247–279.

- Setiya, K. (2008). "Practical Knowledge," *Ethics* 118(3): 388–409.
- , (2012). "Know-How," *Proceedings of the Aristotelian Society* 112 (3): 285–307.
- Sosa, E. (1999). "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13: 141–153.
- , (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). OUP Oxford.
- Stalnaker, R. (1984). *Inquiry*. The MIT Press.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- & Williamson, T. (2001). "Knowing How," *The Journal of Philosophy* 98.8: 411–444.
- Stich, S. (1978). "Autonomous Psychology and the Belief-Desire Thesis," *The Monist* 61.
- Velleman, D. (1989/2007). *Practical Reflection*. CSLI Publications.
- Weisberg, J. (2013). "Knowledge in Action," *Philosophers' Imprint*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library.
- Williams, B. (1981). *Moral Luck*, Cambridge: Cambridge University Press; page reference is to the reprint of chapter 2 in Statman 1993b.
- (1993). "Postscript", in *Moral Luck*, D. Statman (ed.), Albany: State University of New York Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.