

Practical Knowledge without Luminosity

Forthcoming in *Mind*

BOB BEDDOR

National University of Singapore, Singapore

phibrs@nus.edu.sg

CARLOTTA PAVESE

Cornell University, USA

cp645@cornell.edu¹

Abstract: According to a rich tradition in philosophy of action, intentional action requires practical knowledge: someone who acts intentionally knows what they are doing while they are doing it. Piñeros Glasscock (2020) argues that an anti-luminosity argument, of the sort developed in Williamson (2000), can be readily adapted to provide a *reductio* of an epistemic condition on intentional action. This paper undertakes a rescue mission on behalf of an epistemic condition on intentional action. We formulate and defend a version of an epistemic condition that is free from any luminosity commitments. While this version of an epistemic condition escapes *reductio*, it comes with substantive commitments of its own. In particular, we will see that it forces us to deny the existence of any essentially intentional actions. We go on to argue that this consequence should be embraced. On the resulting picture, intentional action is not luminous. But it still entails practical knowledge.

¹ The order of the authors is alphabetical.

1. Introduction

What is the difference between acting intentionally and acting unintentionally? A number of philosophers have sought to answer this question in terms of knowledge. According to an epistemic condition on intentional action, an important part of acting intentionally is to know what one is doing while one is doing it.

An epistemic condition on intentional action has a rich history: it has roots in Aristotle, and it plays a starring role in the work of Anscombe (1958), Hampshire (1959), Gibbons (2001), Newstead (2006), Thompson (2011), Rödl (2011), Small (2012), Wolfson (2012), among others. But recently it has fallen under attack. Here we focus on one of the most important challenges in the recent literature, due to Piñeros Glasscock (2020, this journal). In effect, Piñeros Glasscock argues that defenders of an epistemic condition are committed to the idea that intentional action is *luminous* — that is, whenever one acts intentionally, one can know that one is acting intentionally. But a well-known style of argument seems to show there are no (non-trivial) luminous conditions (Williamson 2000; Srinivasan 2015). Piñeros Glasscock argues that these anti-luminosity arguments can be readily adapted to provide a *reductio* of an epistemic condition on intentional action.

This paper undertakes a rescue mission on behalf of an epistemic condition. We start by motivating an epistemic condition, highlighting its ability to capture the conceptual connections between intentional action and control (§2). Our preferred formulation of an epistemic condition is importantly weaker than the formulation targeted by Piñeros Glasscock, allowing it to escape his version of the *reductio* (§§3-4). But not all is smooth sailing. Another tradition in action theory — which also features prominently in Anscombe (1958) — holds that some actions are *essentially intentional*. Given this assumption, we show that Piñeros Glasscock's *reductio* can be revived against even our weaker epistemic condition (§5). Consequently, the real lesson of the

anti-luminosity argument is that philosophers of action must choose between an epistemic condition and the doctrine of essentially intentional actions.

Having laid out this choice, we explore how to resolve it (§6). We argue that the main motivations for positing essentially intentional actions can be captured using the weaker claim that some actions are intention-entailing. Moreover, the conceptual connections between intentional action and control provide reason to doubt the existence of essentially intentional actions, since it is doubtful whether any actions are essentially under our control. By rejecting the doctrine of essentially intentional actions, we can preserve the idea that intentional action entails knowledge, while denying that it is luminous.

2. Motivating an epistemic condition on intentional action

One platitude about intentional action is that it is, in some sense, under the agent's control. As an illustrative example, Ryle contrasts the clown, who tumbles intentionally for an audience, with a klutz, who tumbles inadvertently (Ryle 1949, p. 33). It seems that their actions are distinguished, in part, by the fact that the clown is in control of their tumbling, whereas the klutz is not. This motivates:

Control Constraint: Whenever an agent ϕ s intentionally, they are in control of their ϕ -ing.

Further motivation for the Control Constraint comes from its ability to explain a variety of cases that have loomed large in the philosophy of action literature. First, consider deviant causal chains. It is a familiar observation that an agent can intend to ϕ , and this intention can lead them to ϕ , but they nonetheless do not ϕ *intentionally*, because the intention and its execution are not connected in the 'right way' (Davidson 1973). A standard example:

Inadvertent Crash. A philosopher intends to knock over his glass to distract his commentator. However, his intention so upsets him that his hand shakes uncontrollably, striking the glass and knocking it to the floor. (Adams and Mele 1989, p. 519)

Even though the philosopher intends to knock over the glass, and even though this intention causes him to knock over the glass, he does not knock over his glass intentionally. The Control Constraint explains this intuition. When the philosopher's shaking hand makes contact with the glass, his action is not under his control.

A second class of cases involves agents who accomplish their intention non-deviantly, but only due to a stroke of luck. For example:

Nuclear Reactor. A nuclear reactor is in danger of exploding. Fred knows that its exploding can be prevented only by shutting it down, and that it can be shut down only by punching a certain ten digit code into a certain computer. Fred is alone in the control room. Although he knows which computer to use, he has no idea what the code is. Fred needs to think fast. He decides that it would be better to type in ten digits than to do nothing. Vividly aware that the odds against typing in the correct code are astronomical, Fred decides to give it a try. He punches in the first ten digits that come into his head, in that order, believing of his so doing that he 'might thereby' shut down the reactor and prevent the explosion. What luck! He punched in the correct code, thereby preventing a nuclear explosion. (Mele & Moser 1994, p. 40)

Intuitively, Fred does not punch in the correct code *intentionally*. A Control Constraint captures this intuition: punching in the correct sequence was not under Fred's control.

Given these examples, it is no wonder that the Control Constraint has been endorsed by so many action theorists (e.g., Mele & Moser 1994; Gibbons 2001; O'Brien 2007; Wu 2016), including Piñeros Glasscock (2020, p. 1245-1246). But what does it mean to have control over one's action?

Here is a big picture hypothesis: control can be cashed out in epistemic terms. While there are different ways of fleshing out this hypothesis, we think the following proposal has considerable promise:

Epistemic Theory of Control: Someone is in control of their φ -ing at t if and only if they know they are φ -ing at t , and they know this in virtue of exercising their knowledge of how to φ .

This theory combines two ideas that have played a prominent role in the literature. The first is that control requires practical knowledge of what one is doing. The second is that control requires know-how or skill (Mele & Moser 1994; Small 2012; Pavese 2018; Piñeros Glasscock 2020). The Epistemic Theory of Control integrates these ideas, requiring that one knows what one is doing in virtue of exercising know-how.²

To unpack this theory, let us see how it handles our cases. Start with the contrast between the clown and the klutz. While tumbling, the clown knows that they are tumbling. Moreover, they arrive at this knowledge in virtue of exercising their knowledge

² A couple of remarks may help forestall some immediate objections. First, we should not require that the knowledge in question is consciously entertained by the agent: a skilled athlete might be implicitly aware of their movements (cf. Paul 2009). In such a case, the Epistemic Theory of Control will count them as having control over their movements. Second, some may object that the Epistemic Theory of Control delivers the wrong results in Davidson's (1978) carbon copier case, in which someone writes heavily on a page, intending to produce 10 legible carbon copies. According to Davidson, if they succeed in making 10 copies, they did so intentionally, even though they did not know what they were doing. Two responses to this sort of case are available. The first is to dispute Davidson's claim that the copier intentionally made 10 carbon copies. According to this response, if the copier is not checking what they are doing, then they do not exhibit the right sort of control over their action for it to be intentional. (See Thompson 2011, p. 210; Small 2012, p. 199; Piñeros Glasscock (2020, p. 1245-1246) expresses sympathy for this response.) A more concessive response is to weaken the Epistemic Theory. Taking a cue from Davidson's own remarks, we might propose that in order for someone to be in control of their φ -ing, they only need to know that they are performing some action ψ that is a possible means of φ -ing, and that they know this in virtue of exercising their knowledge of how to φ .

of how to tumble. Not so with the klutz. Mid-tumble, they might realize that they are tumbling, but they do not have this knowledge in virtue of exercising their knowledge of how to tumble.

Next, take Inadvertent Crash. As our rattled philosopher's shaking hand bumps into the glass, they may or may not recognize that they are knocking over their glass. But even if they do, they do not have this knowledge in virtue of exercising their knowledge of how to knock over the glass. Finally, the Epistemic Theory of Control nicely handles lucky successes. In Nuclear Reactor, Fred does not know that he is punching in the correct code. He also does not know how to punch in the correct code.

Some might wonder whether we could explain these cases equally without invoking knowledge. Why not instead hold that someone is in control of their action just in case they are able to adjust their course of action to ensure it goes according to plan (cf. Piñeros Glasscock 2020, p. 1245)? But, on closer examination, it is not clear this proposal is an *alternative* to the Epistemic Theory of Control. In order for someone to be able to adjust their course of action to ensure it goes according to plan, they will need to successfully monitor what they are doing. But what does it mean to successfully monitor one's actions? A natural thought is that to successfully monitor one's actions is just to *know* what one is doing. If so, this approach will entail an epistemic condition on intentional action after all.

Perhaps, some might suggest, successful monitoring only requires *true belief* about what one is doing, or perhaps *justified true belief* (cf. Piñeros Glasscock 2020, p. 1255). In order to address this, consider cases where someone has a justified true belief about what they are doing, but this belief is Gettiered. Consider, for example, the following variant of Chisholm's (1966) 'sheep in the field' case:

Deer Hunting. Artemis, the goddess of the hunt, spies what appears to be a deer. She expertly notches her bow, aims, and releases her arrow. It turns out that she

was aiming at a lifesize papier mâché statue of a deer — one that had been constructed so artfully that it would fool even the most discerning eye. But, as luck would have it, a deer was standing directly behind the cervine statue. Artemis' arrow passes directly through the papier mâché sculpture and into the hapless deer.

Artemis successfully shot a deer. But she did not do so intentionally. Of course, there are various actions she performed intentionally in the course of shooting a deer: she intentionally released her bowstring; she intentionally shot at what appeared to be a deer, etc. While these further actions are indeed under her control, it seems purely accidental that Artemis succeeded in *shooting a deer*. But if we analyze control in terms of successfully monitoring what one is doing, and we take successful monitoring to only require justified true belief, then we will wrongly predict that this action is under Artemis' control. (After all, Artemis has a justified true belief that she is shooting a deer.) By contrast, the Epistemic Theory of Control delivers the correct prediction. Artemis' action is not under her control since she does not *know* she is shooting a deer.³

So the Epistemic Theory of Control has considerable explanatory power. Given the Control Constraint, this entails an epistemic condition on intentional action:

Epistemic Condition (EC): Whenever an agent ϕ s intentionally, they know that they are ϕ -ing, and they have this knowledge in virtue of their knowledge of how to ϕ .

³ Some might protest that even if some Gettier cases favor the Epistemic Theory of Control, others count against it. Consider a variant of a case discussed by Setiya (2008, 2012): a patient has been given a drug that paralyzes their hand; they are told by a doctor it will wear off at noon. The patient forms the intention to clench their fist at noon. It turns out that the doctor was looking at the wrong charts. But, by coincidence, everything the doctor said was correct. Still, one might think, when the patient successfully clenches their fist at noon, their action is under their control. However, this intuition is compatible with the Epistemic Theory of Control. Normally, someone who is clenching their fist gains strong proprioceptive evidence that they are clenching their fist. If our patient has this proprioceptive evidence, it can enable them to know that they are clenching their fist *while* they are doing so. Thus before noon, the patient didn't know that they could clench their fist at noon; but at noon, mid-clench, they can know that they are clenching. Of course, we can imagine a variant case where the patient lacks this proprioceptive evidence, perhaps due to the lingering effects of the drug. But once we add this stipulation, it becomes less clear that the agent's clenching is intentional.

Thus far we have motivated an epistemic condition on intentional action by considering the connections between intentional action and control. A further motivation comes from its intuitive plausibility. Take a well-known example from Anscombe (1958, §6): if someone is (intentionally) sawing a plank, it would be odd for them to claim that they did not know that they were sawing a plank. EC naturally explains this oddity. Moreover, some version of EC is arguably reflected in our legal and judicial practices. According to the Model Penal Code (2.02), one of the requirements for determining whether an agent acted culpably is whether they acted *knowingly*. A plausible explanation for this requirement is that acting culpably entails acting intentionally, and one acts intentionally only if one knows what one is doing.

Having outlined some considerations in favor of an epistemic condition, let us now consider the trouble that lies in store.

3. Piñeros Glasscock's anti-luminosity argument

Piñeros Glasscock argues that an epistemic condition on intentional action leads to disastrous consequences. In developing his argument, Piñeros Glasscock focuses on the following epistemic condition:

Practical Knowledge Principle (PKP): Necessarily, if an agent is ϕ ing (intentionally and under that description), they know that they are ϕ ing (intentionally and under that description).

Before delving into Piñeros Glasscock's argument, we should flag an important difference between PKP and the epistemic condition defended above (EC). PKP requires that whenever one intentionally acts, one knows that one is acting *intentionally*. By contrast, EC only imposes the weaker requirement that, whenever one intentionally acts,

one knows that one is acting; one need not know that one's action is intentional. As we will see shortly, this difference is crucial.

On to the argument itself. Piñeros Glasscock offers the following case:

Cleaning Sisyphus. The floors of the Underworld are getting filthy with blood and bile. Hades notices, and decides to give Sisyphus a more useful task than pushing a rock, repeatedly, up a hill. Handing him a mop, he orders him to clean the floors using the waters from the river Acheron, the cleanest in the Underworld. There is one problem: while at noon the waters of Acheron are clean as a spring, they slowly and gradually get dirtier and dirtier as the day goes by — by midnight it is just filth, much dirtier than the floors of the Underworld. Hades thus tells Sisyphus that at midnight he will be punished proportionally to his efficiency: his punishment will be worse if he fails to keep mopping when the water is still clean enough or if he keeps mopping when the water is dirtier than the floors. Each day, therefore, Sisyphus grabs his mop in the morning trying to mop as much of the Underworld as he can... At noon, he is fully confident that he is intentionally acting under the description *cleaning the floors*. As the day goes by, however, he loses more and more confidence. But by midnight on those days when he keeps mopping to defy Hades, he is certain that he is *not* cleaning the floors of the Underworld (but instead making them dirtier). (Piñeros Glasscock 2020, p. 1248-9)

From the description of the case, we have:

- 1) Sisyphus is intentionally cleaning the floors at t_0 (noon).

From 1) and PKP we get:

2) Sisyphus knows he is intentionally cleaning the floors at t_0 .

Next, Piñeros Glasscock appeals to a general modal constraint on knowledge, due to Williamson (2000):

Margin for Error: If S knows that S is in condition C in a case σ , then in all nearby cases σ' , S is in C .

A special instance of Margin for Error is:

Margin for Error for Action (Mara): If S knows that S is φ ing intentionally at t , then S is φ ing intentionally at $t+1$.

2) and Mara yield:

3) Sisyphus is intentionally cleaning the floors at t_1 .

By another application of PKP, we get:

4) Sisyphus knows that he is intentionally cleaning the floors at t_1 .

Another application of Mara lets us conclude that Sisyphus is intentionally cleaning the floors at t_2 . By repeated applications of this mode of reasoning, we arrive at the conclusion that at midnight, Sisyphus is intentionally cleaning the floors of the

Underworld. But this contradicts the stipulation that he is not cleaning the floor at midnight. So PKP combined with Margin for Error leads to a contradiction.⁴

4. Escaping the *reductio*

Piñeros Glasscock concludes that we must jettison PKP. We agree.⁵ In a sense this conclusion is unsurprising, given the way PKP is formulated. In effect, PKP claims that intentional action is luminous — i.e., whenever one is intentionally acting, one is in a position to know that one is intentionally acting. If Williamson (2000) is right that there are no (non-trivial) luminous conditions, we should expect that PKP fails for precisely the sort of reasons Piñeros Glasscock identifies.

Does this mean that there is no connection between intentional action and knowledge? Not necessarily; one might opt for a weaker connection between the two. There are at least two ways this might go.

One possibility is briefly suggested by Piñeros Glasscock, who floats the hypothesis that the function of the will is to produce practical knowledge (Piñeros Glasscock 2020, 1262). On this view, PKP holds in all normal circumstances — that is, all circumstances where the will fulfills its function. But it leaves open the possibility that there are abnormal circumstances in which one acts intentionally while lacking practical knowledge.⁶

⁴ Piñeros Glasscock actually uses a slightly different formulation of Mara, which omits the qualification that the cleaning is intentional:

Piñeros Glasscock's Mara: If Sisyphus knows that he is cleaning the floors at a case α_i , he is cleaning the floors at α_{i+1} , for all times in the series $t_0...t_n$. (Piñeros Glasscock 2020, p. 1248)

However, notice that while this formulation allows us to infer that Sisyphus is cleaning the floors at t_2 , it does not yet license the inference to 3) (Sisyphus is cleaning the floors *intentionally* at t_2), unless we make the further assumption that the cleaning is intentional. And 3) is crucial to the argument, since without it we cannot get the *reductio* by further applications of PKP. We have opted for the formulation of Mara in the main text in order to make the structure of the argument perspicuous.

⁵ Another option would be to reject Margin for Error. For example, Berker (2008) worries that Margin for Error encodes an implausibly strong requirement on knowledge. For relevant discussion, see Srinivasan (2015) and Goldstein and Waxman (forthcoming). For the purposes of this paper, we will not pursue this line of objection. Instead, we will show that even granting Margin for Error, an epistemic condition on intentional action is defensible.

⁶ See O'Brien (2007, p. 159-160), Setiya (2009, p. 131), and Schwenkler (2019, chp. 6) for similar suggestions.

While this is an intriguing option, it forfeits some of the explanatory advantages that motivated an epistemic condition. One goal of a theory of intentional action is to predict our judgments *even when practical knowledge is absent*. For example, in §2 we saw that an epistemic condition explains why the protagonists of Inadvertent Crash, Nuclear Reactor, and Deer Hunting do not act intentionally. The explanation was that in all of these cases, the agents are not in control of their action, because they do not satisfy the epistemic requirements on control. By contrast, if we restrict PKP to only hold in normal circumstances, we will not be able to make any predictions about these cases. After all, these are cases where practical knowledge is lacking. Hence, on the view under consideration, these do not qualify as normal circumstances. To put the point another way: we would like a theory that tells us whether an action is intentional or unintentional even in circumstances where the will is not fulfilling its function.

For this reason, we think a more promising option is to retain a fully general epistemic condition on intentional action, but one that avoids the paradoxical consequences of PKP. Luckily, we already have such a principle: EC. Notice that EC is not equivalent to a luminosity thesis: it does not say that if one intentionally φ s then one knows that one is *intentionally* φ ing. According to EC, if an agent intentionally φ s, they only need to know that they are φ ing.

This difference enables EC to avoid the *reductio*. To see this, let us try reconstructing Piñeros Glasscock's argument using EC. As before, we have:

- 1) Sisyphus is intentionally cleaning the floors at t_0 .

From 1) and EC we can no longer derive 2), but only:

- 2') Sisyphus knows he is cleaning the floors at t_0 .

Next, we might appeal to the relevant instance of Margin for Error:

Mara': If S knows that S is ϕ ing at t , then S is ϕ ing at $t+1$.

From 2') and Mara', we can derive:

3') Sisyphus is cleaning the floors at t_1 .

But, crucially, we cannot draw the stronger conclusion that Sisyphus is *intentionally* cleaning the floors at t_1 . So we cannot use EC to derive 4) from 3'):

4) Sisyphus knows that he is intentionally cleaning the floors at t_1 .

Without 4), the *reductio* does not get off the ground.⁷

5. Reviving the *reductio*?

Piñeros Glasscock briefly acknowledges that one might weaken PKP, along the lines suggested by EC. But, in a footnote he argues that this is not sufficient for resolving the problem, since either principle will range over actions that are 'by their very nature intentional' (fn 8, p. 1240). While Piñeros Glasscock does not spell this out in detail, we think he is onto something important. Unpacking this will take some work.

⁷ As we have seen, EC is motivated by the Epistemic Theory of Control. However, some might wonder whether the motivations for EC support the stronger principle, PKP. After all, it would seem that when someone is in control of their action, they are typically not just aware of their action; they are also aware that they are acting *intentionally*. While we agree this is typically the case, we take Piñeros Glasscock's argument to provide a powerful reason to doubt that it is always the case. For example, in Cleaning Sisyphus, there will be some last moment t when Sisyphus is intentionally cleaning the floors. A moment later ($t+1$), Sisyphus is still cleaning the floors, but he is not doing so intentionally. This is because, at $t+1$, he does not *know* he is cleaning the floors, since at $t+1$ he does not satisfy the Margin for Error Principle. This diagnosis is compatible with Piñeros Glasscock's suggestion that PKP holds in normal circumstances. But it differs in insisting that some minimal epistemic condition on intentional action (i.e., EC) holds in all circumstances, normal or otherwise.

First, some background: Anscombe (1958, §47) famously argued that there is an important distinction between actions such as *offending*, *dropping*, and *kicking*, and actions such as *greeting* and *marrying*. The first class of actions can be done either intentionally or unintentionally. By contrast, Anscombe contends that the second class of actions are *essentially intentional*: whenever they are performed, they are performed intentionally.

Suppose for now that this is correct. Then we can run a version of the anti-luminosity argument using one of these essentially intentional actions, without relying on PKP. To illustrate, assume *greeting* is an essentially intentional action. Now consider the following scenario:

Hades' Greetings. Upon descending to the Underworld, Diogenes is given a job to make up for his life of indolence. He is in charge of greeting new souls as Charon ferries them to their final resting place. At 9am he starts the day eager and fresh, and gives the first unfortunate soul who steps off the ferry a hearty greeting. Throughout the day, there is a steady stream of souls, conveniently spaced a minute apart. As the day goes on, he grows tired, and his greetings become increasingly lackluster. By 3pm, he barely smiles at each new soul. At midnight, when the last arrival of the day descends from the ferry, he does not greet them.

From the description of the case, we have:

- 1) At 9am, Diogenes greets someone.

By the assumption that greeting is essentially intentional, we get:

- 2) At 9am, Diogenes intentionally greets someone.

Hence by EC:

3) At 9am, Diogenes knows that he is greeting someone.

Next we apply Mara' to derive:

4) At 9:01, Diogenes is greeting someone.

By another application of the assumption that greeting is essentially intentional, we get:

5) At 9:01, Diogenes is intentionally greeting someone.

5), together with another application of EC, yields:

6) At 9:01, Diogenes knows that he is greeting someone.

Repeat enough times, and we can derive the conclusion that at midnight, Diogenes is greeting someone, which contradicts the description of the case.

This shows that we can get the *reductio* up and running without relying on PKP. However, this version of the *reductio* depended on the assumption that there are essentially intentional actions. So the real upshot of the argument is that, given Margin for Error, we must choose between two claims that have been widely endorsed by action theorists: an epistemic condition on intentional action and the doctrine of essentially intentional actions.

Which of these two claims should we retain? We have already seen (§2) that EC comes with considerable explanatory advantages. To recap: EC follows from the Control

Constraint, together with the Epistemic Theory of Control. And the Epistemic Theory of Control explains our judgments about a wide variety of cases (e.g., Inadvertent Crash, Nuclear Reactor, Deer Hunting). These advantages should not be abandoned lightly. They thus provide grounds for taking a closer look at the doctrine of essentially intentional actions. Does this doctrine enjoy any comparable explanatory advantages?

6. The Doctrine of essentially intentional actions

The doctrine of essentially intentional actions traces back to Anscombe's claim that certain action descriptions require that the action is performed intentionally. Anscombe's list includes *telephoning*, *greeting*, and *marrying* (§47). Other philosophers have echoed Anscombe's claim. For example, Bennett (1988) discusses what he calls 'intention-drenched' verbs such as *hunting* and *fishing*. More recently, Moore (2010) takes up Anscombe's claim in the context of philosophy of law, adding *impersonating* and *concealing* to Anscombe's list.

At first glance, Anscombe's claim seems compelling. However, given the Control Constraint, the doctrine of essentially intentional actions has a striking consequence — namely, that certain actions are *essentially under our control*. On closer scrutiny, this consequence seems doubtful. For any action, it seems we can concoct cases where the action is successfully performed but the action is not under the agent's control, due to some deviance in the causal path, or some luck in the circumstances surrounding its performance.

To illustrate the general strategy, let us start with *greeting*. For a case where a greeting is deviantly caused, consider:

Cultural Confusion. Glen is visiting a new country and asks his taxi driver about the local etiquette for greeting people. As a matter of fact, the local etiquette is to greet by waving one's right hand; waving with the left is considered offensive.

However, Glen's driver has a mischievous streak: he tells Glen that the local custom is to wave with one's left hand, never the right. Glen thanks the driver for the advice, and sets out to implement it. The next day, Glen spots a new acquaintance. Glen intends to greet them, and intends to do so by waving with his left hand. However, Glen is prone to confusing his left with his right, and he does so on this occasion. He notices with consternation that he is waving with his right hand. Thankfully, the acquaintance smiles and returns the greeting in kind.

Does Glen greet his acquaintance? It seems to us the answer is 'Yes'. This intuition can be supported on theoretical grounds. *Greeting* is standardly classified as an illocutionary speech act that is used to express welcome or recognition. In the literature on speech acts, there is some debate about whether uptake suffices for the performance of an illocutionary speech act, or whether an intention on the speaker's part is also required. But on virtually all views, a sufficient condition for an illocutionary speech act to be successfully performed is that *both* i) the speaker intends to perform the act, ii) their action is received as such by their audience (e.g., Austin 1975; Searle 1985; Hornsby and Langton 1998). In *Cultural Confusion*, both of these conditions are satisfied.

But does Glen greet his acquaintance *intentionally*? The Control Constraint provides a principled reason for answering 'No'. After all, it was purely accidental that Glen confused his left and his right hand on this occasion. If he had not confused the two, he would not have successfully greeted his acquaintance; rather, he would have insulted them. To reinforce this judgment, note that Glen does not even *know how* to greet his acquaintance. As we saw in §2, it is widely agreed that intentional action requires know-how (a point granted by Piñeros Glasscock (2020, p. 1242-1243). This gives us a further argument for denying that Glen's greeting was intentional.⁸

⁸A referee raises the possibility that we should distinguish between different senses of greeting. Perhaps there is an external sense of greeting, according to which any action that conforms to the social norms for greeting qualifies as a greeting. And perhaps there is also an internal sense, according to which one greets one's acquaintance if and only if one performs the specific action that one believes constitutes a greeting. Glen's action is a greeting in the external

Of course, *greeting* partly depends on social conventions. But our argument readily generalizes to less conventional actions. To illustrate this general point, consider Bennett's example of *hunting*. For a case where someone hunts successfully but this action was not under their control, we need look no further than our Deer Hunting scenario (§2). Artemis succeeds in hunting a deer. But she does not do so intentionally, since her success is due to luck.⁹

Some might balk at a blanket rejection of essentially intentional actions. Surely, one might protest, Anscombe and Bennett were onto something when they distinguished between actions like *dropping* and *offending* and actions like *greeting* and *hunting*. However, there is a way of accommodating this distinction without subscribing to the doctrine of essentially intentional actions. Distinguish between:

Essentially intentional actions: Actions that are intentional whenever they are performed.

Intention-entailing actions: Actions that are accompanied by an intention whenever they are performed.¹⁰

sense, but not the internal sense. However, even if we are willing to make this distinction, our main point still goes through. Let us amend our case. As before, Glen intends to greet someone by performing some particular action φ . This time, suppose that Glen does end up φ -ing successfully, but only by way of a deviant causal chain. In this variant, Glen greets his acquaintance in the internal sense, but still does not do so intentionally. So the point remains: greeting — on either the external or the internal sense — is not essentially intentional.

⁹ These considerations extend to more passive and immediate actions. Consider *watching*. Suppose you are playing hide and seek. You're the seeker, and you have your eyes closed, so as not to see where your friends are hiding. A noise startles you; you inadvertently open your eyes; your gaze fixates on your friend, who is crouching behind the sofa. Or consider the 'eyeball-torturing' scene in Kubrick's *A Clockwork Orange* where Alex is forced to watch violent films. In neither case is the watching intentional.

¹⁰ There are a couple ways of refining the notion of being 'accompanied by an intention'. One option is to gloss this in causal terms: we might maintain that the intention causally contributes to the performance of the action. Another option is to adopt a doxastic conception of intention and propose that an action is accompanied by intention provided it is performed *in the belief* that one is or may be so acting. For our purposes, we can remain neutral on this issue. (Thanks to a referee for raising this point.)

If there were any essentially intentional actions, they would be intention-entailing, given the plausible assumption that intentional actions are accompanied by an intention.¹¹ But the converse does not hold: an action might be accompanied by an intention, but still fail to be intentional, as illustrated by many of the cases we have considered (e.g., Inadvertent Crash, Deer Hunting, Cultural Confusion).

This suggests a way of preserving a version of Anscombe's distinction while holding onto EC: one can deny that *greeting*, *hunting*, *fishing* and the like are essentially intentional, but grant that they are intention-entailing. Often when philosophers invoke essentially intentional actions, intention-entailing actions will serve their purposes just as well. For example, Bennett points out that if someone dangles a hook in the water and happens to snare a fish, it does not follow that the person fished. Bennett concludes that 'S fished' entails that S performed some action that S intended to result in catching a fish (Bennett 1988, p. 206). But this only shows that *fishing* entails a corresponding intention, not that it is essentially intentional. Similarly, Moore (2010) argues that Anscombe's distinction helps us to single out a class of actions whose performance entails *mens rea* (i.e., criminal intent). This use of Anscombe's distinction is perfectly consistent with the hypothesis that these actions are merely intention-entailing.

Now, one might well question whether all of Anscombe's putative examples of essentially intentional actions are even intention-entailing. Take Anscombe's example of *telephoning* (Anscombe 1958, §47). It seems possible to telephone someone without intending to do so, as revealed by the common experience of inadvertently placing a call by sitting on your phone. For a more controversial example, consider Anscombe's example of *marrying*. A common trope in fiction revolves around 'accidental marriages.' For example, in Wilkie Collins' novel *Man and Wife*, a plot point hinges on whether a woman has inadvertently become married to an acquaintance by spending a night under

¹¹ Bratman (1984) rejects the idea that intentionally φ -ing requires an intention to φ . But Bratman still accepts that intentionally φ -ing requires having an intention to perform some action that is appropriately related to φ -ing.

the same roof as him in a pub.¹² Similar qualms can even be raised about *greeting*. In particular, speech act theorists who regard uptake as sufficient for the performance of a speech act will deny that greeting requires intention. For our purposes, we need not take a stand on how to classify such cases. The important point is that *telephoning*, *marrying*, and *greeting* are not essentially intentional; we can remain neutral on whether they are intention-entailing.

Summarizing: intentional action requires control. But, for any action, we can devise cases where the action is successfully performed, but it was not properly under the agent's control. This suggests that, at least for agents like us, there cannot be essentially intentional actions. There can only be (at most) intention-entailing actions.

On reflection, this conclusion comports nicely with the main lesson of Williamson's anti-luminosity argument. The upshot of the anti-luminosity argument is that subjects epistemically like us lack any 'cognitive home': there is no condition C such that being in C guarantees that you can know you are in C . This is because, for every condition C , it could be a matter of luck that you are in C : you could easily have failed to be in C , had circumstances been slightly different. Similarly, to deny that there are essentially intentional actions is to deny that we have any 'practical home': there is no action φ such that φ -ing (even with the intention to do so) guarantees that you φ intentionally. This is because, for every action, it could be a matter of luck that you fulfilled your intention to perform it: you could easily have failed to perform that action, had circumstances been slightly different. So, the anti-luminosity argument fits very naturally with the idea that no action is essentially intentional.

Conclusion

This paper explored how to salvage an epistemic condition on intentional action in the face of Piñeros Glasscock's anti-luminosity argument. Our advice has been to abandon

¹² Thanks to the editors for this example.

PKP (according to which intentional action requires knowing that one is acting intentionally) and retreat to EC (according to which intentional action merely requires knowing what one is doing). As we have seen, the threat of *reductio* resurfaces for even this weaker condition, given the assumption that some actions are essentially intentional. However, we went on to argue that the very sort of considerations that render an epistemic condition attractive provide reason to doubt the existence of essentially intentional actions. Moreover, the important philosophical applications of essentially intentional actions can be equally well achieved by intention-entailing actions. For action theorists willing to go this route, an epistemic condition on intentional action remains a viable contender.¹³

References

- Adams, Frederik and Mele, Alfred 1989, 'The Role of Intention in Intentional Action' *Canadian Journal of Philosophy* 19 (4): 511–531.
- Anscombe, Gertrude Elizabeth Margaret 1958, *Intention* (London: Basil Blackwell)
- Austin, John Langshaw 1975, *How to do things with words*. Vol. 88 (Oxford: Oxford University Press)
- Bennett, Jonathan 1988, *Events and their Names* (Oxford: Oxford University Press)
- Berker, Selim 2008, 'Luminosity Regained' *Philosopher's Imprint* 8 (2): 1–22
- Bratman, Michael 1984, 'Two faces of intention' *The Philosophical Review* 93(3): 375–405
- Chisholm, Roderick 1966, 'The status of appearances' *Theory of knowledge* 91–98
- Davidson, Donald 1973, 'Freedom to act', in Ted Honderich (ed). *Essays on Freedom of action* (London: Routledge)
- 1978, 'Intending' In Yovel (ed.), *Philosophy of History and Action* (Dordrecht: D. Reidel)

¹³ Many thanks to the editors of *Mind* and two anonymous referees for very helpful comments on a previous draft.

- Gibbons, John 2001, 'Knowledge in action' *Philosophy and Phenomenological Research* 62(3): 579–600.
- Goldstein, Simon & Waxman, Daniel 2020, 'Losing confidence in Luminosity' *Noûs*.
- Hampshire, Stuart 1959, *Thought and Action* (University of Notre Dame Press)
- Hornsby, Jennifer, and Rae Langton 1998, 'Free speech and illocution' *Legal Theory* 4: 21–37.
- Mele, Alfred and Paul Moser 1994, 'Intentional action' *Nous* 28(1): 39–68.
- Moore, Michael 2010, *Act and crime: The philosophy of action and its implications for criminal law* (New York: Oxford University Press)
- Newstead, Anne 2006, 'Knowledge by Intention? On the Possibility of Agent's Knowledge' in Stephen Hetherington (ed.), *Aspects of Knowing* (Amsterdam: Elsevier Science)
- O'Brien, Lucy 2007, *Self-knowing agents* (Oxford: Oxford University Press)
- Paul, Sara 2009, 'How We Know What We're Doing', *Philosophers' Imprint* 9: 1–24.
- Pavese, Carlotta 2018, 'Know-how, action, and luck' *Synthese* 1–23.
- Piñeros Glasscock, Juan 2020, 'Practical Knowledge and Luminosity' *Mind* 129(516): 1237–67.
- Rödl, Sebastian 2011, 'Two forms of practical knowledge and their unity' In Ford, Hornsby, & Stoutland (eds.), *Essays on Anscombe's Intention* (Cambridge: Harvard University Press), pp. 1–23.
- Ryle, Gilbert 1949, *The Concept of Mind* (London: Hutchinson & Co.)
- Schwenkler, John 2019, *Anscombe's Intention: A Guide* (New York: Oxford University Press)
- Searle, John 1985, *Expression and meaning: Studies in the theory of speech acts* (Cambridge: Cambridge University Press)
- Setiya, Kieran 2008, 'Practical knowledge' *Ethics*, 118(3): 388–409.
- 2009, 'Practical knowledge revisited' *Ethics*, 120(1): 128–137.

- 2012, ‘Knowing How’ *Proceedings of the Aristotelian Society* 112(3): 285–307.
- Small, Will 2012, ‘Practical Knowledge and the Structure of Action’ *Rethinking Epistemology* 2: 133–228.
- Srinivasan, Amia 2015, ‘Are we luminous?’ *Philosophy and Phenomenological Research*, 90(2): 294–319.
- Thompson, Michael 2011, ‘Anscombe’s Intention and Practical Knowledge’ In Fort, Hornsby, and Stoutland (eds.), *Essays on Anscombe’s Intention* (Cambridge: Harvard University Press)
- Williamson, Timothy 2000, *Knowledge and its Limits* (Oxford: Oxford University Press)
- Wolfson, Ben 2012, ‘Agential Knowledge, Action and Process,’ *Theoria* 78: 326–57.
- Wu, Wayne 2016, ‘Experts and deviants: The story of agentic control.’ *Philosophy and Phenomenological Research* 93(1): 101–126.