ACTING AS A REASON:
AN ACCOUNT OF PRACTICAL KNOWLEDGE

Practical knowledge is knowledge of what one is doing. Without it, what one does is not intentional. For instance, I might, unbeknownst to me, be dialing your number by sitting on my phone. I am, indeed, dialing your number, but I am not doing so intentionally (or on purpose). And the reason seems to be that the dialing that is happening is *unbeknownst* to me. Further, if I heard the phone ringing and came to realize that I was calling someone, the possession of this knowledge would not make my calling intentional. The knowledge of what one is doing that makes one's action intentional is not derived from observation or evidence of any kind. That is to say, one does not *discover* what one is doing, when that knowledge is practical. Also, intentional action is not just action accompanied by practical knowledge. If I am calling you because I sat on my phone and, independently, I decide to call you and so reach for my phone, I will have practical knowledge of calling you, but this will not make the calling that is already happening into an intentional calling. Nor would it do so, if my decision happened to coincide with the moment I sat on my phone. An action is intentional only if the action is happening *because of* the practical knowledge. The unintentional calling would have happened whether I had the knowledge or not. The intentional calling would not be happening were it not for the knowledge. Practical knowledge is productive of what's known. In sum, practical knowledge is necessary for intentional action, non-evidential, and productive.

We want to understand these three central thoughts about practical knowledge because we want to understand human agency. Such agency is the capacity to not just act but to act *intentionally*. In order to understand this capacity, therefore, we must understand intentional action and the role of knowledge in it. Cognitivists about intention (according to which an intention is a belief) argue that knowledge or at least belief is central to intentional action (Velleman, 1989; Setiya, 2008; Marusic & Schwenkler, 2018). While non-cognitivists argue that knowledge and belief are, at best, peripheral (Paul, 2009a; Sinhababu, 2013). In this paper, I will argue that knowledge is indeed central, but I will do so by giving a novel explanation of our three central thoughts about practical knowledge. This explanation will avoid a number of problems with the dominant cognitivist and non-cognitivist explanations.

Roughly, the dominant cognitivist explanations make forming an intention look irrational, and the dominant non-cognitivist explanations overly weaken our three central thoughts. Their shared mistakes, I will suggest, are assuming that intentional action is acting *for* a reason and that practical knowledge is about the agent's *token* intentional action. I will argue that we should flip our understanding of intentional action: an intentional action is not an action explained by the agent's reason for action; rather, it is an action that explains why the agent is acting (i.e., it is an explanatory reason for action).[1] For instance, an intentional arm-raising is not intentional because it is raised (e.g.) out of a desire to hail a taxi or out of an intention to turn on a light. Instead, the arm-raising is intentional because some other action is done out of it—e.g., it explains why the agent has moved her body into a certain position.

We should flip our understanding in this way because this can help us explain our three central thoughts about practical knowledge better than the dominant cognitivist and non-cognitivist explanations and, in turn, can give us a more attractive picture of intentional action and agency.

In the end, it will emerge that practical knowledge is knowledge about *how to perform* an action (not about the agent's token intentional action).

1. Cognitivism

The cognitivist explanation of practical knowledge goes like this. According to cognitivism, the intention to do A is the belief that one is doing A. If one is doing A intentionally, then one intends to do A and thus believes that one is doing A. Practical knowledge is necessary for intentional action, therefore, because practical knowledge is embodied by this belief.

Further, an intentional action is not just an action accompanied by an intention, but an action done because of an intention. Since practical knowledge is embodied by this intention, an intentional action is an action done because of practical knowledge. That is to say, practical knowledge is productive of action because intentions are productive of action.

[1] A similar theory is suggested by Michael Thompson (2008). In this paper, I will use this theory to develop a novel explanation of practical knowledge and thereby a novel argument for the theory.

And finally, intentions are not formed out of consideration of the evidence. Rather, intentions are formed out of consideration of the (normative) reasons for action. For cognitivism, an intention is a special kind of belief that aims at being true not by getting itself in line with the facts but by getting the facts in line with it. This is why practical knowledge is a special kind of non-evidential knowledge: because intentions are a special kind of non-evidential belief, and practical knowledge is embodied by them.[2]

The problem with this explanation is that it makes forming an intention irrational. Prior to forming an intention to do A, the agent has insufficient evidence that she is doing A because she is not doing A yet. On the one hand, we know that, despite this, it is (normally) perfectly rational to form the intention to do A. It is (normally) perfectly rational to form the intention to do A out of the rational desire to do A more than anything else. On the other hand, it is (normally) irrational to form a belief in the face of insufficient evidence for it but out of a desire for its truth. That is wishful thinking. However, according to cognitivism, the intention to do A is the belief that one is doing A, and thus characteristically, when an agent forms an intention to do A, she is forming a belief in the face of insufficient evidence for it but out of a desire for its truth. Cognitivism entails that something normally rational is characteristically irrational (Grice, 1972).

One cognitivist response is to claim that when forming an intention, it is rational to form a belief in the face of insufficient evidence because such a belief is self-fulfilling and self-justifying. An intention-belief does not merely represent a fact; it also causes that fact to obtain. For instance, if I believe that I am raising my arm and this belief is an intention, then that belief will be the cause of my arm-raising and thus I will be raising my arm and thus my belief will be true. The belief is self-fulfilling. Further, the belief is the cause of my arm-raising and thus its existence provides evidence for my arm-raising. The belief is self-justifying. Therefore, forming such a belief is rational because even though there is insufficient evidence for the belief prior to forming it, one is assured that once the belief is formed it will be true and justified. Forming an intention, on this response, is a leap of faith (Velleman, 1989).

The problem with this response is that it trades one form of irrationality for another. Initially cognitivism looked unattractive because it made forming an intention in the

[2] See Velleman (1989) and Setiya (2007, 2008) for a more detailed version of this explanation.

face of insufficient evidence irrational. The response is to claim that forming an intention is a leap of faith. But leaps of faith are also irrational. Even if one is assured that one's belief will be true and justified once one forms it, prior to forming it one's belief is not true or justified and thus epistemic norms dictate that one should not form the belief. The leap of faith response trades one unattractive feature of cognitivism for another (Langton, 2004).

A different cognitivist response is to claim that when forming an intention, it is rational to form a belief in the face of insufficient evidence because the formation of such a belief is underwritten by one's know-how. Normally, as we noted, forming an intention in the face of insufficient evidence is rational. According to this response, this is rational only if one knows how to perform the relevant action. For example, were one to form the intention to bake bread without knowing how to bake bread, then one's formation of this intention would be irrational. This is because one needs to know how to bake bread intentionally in order to bake bread intentionally; thus, if one forms an intention to bake bread without knowing how, one will not be able to execute this intention. One will be forming an intention to do something that one cannot do. The rational thing to do in such cases is to form the intention to try to do the action or to form the intention to learn how to do the action. According to this response, then, the formation of an intention is made rational by one's know-how. This fact, the response continues, is an *epistemic* fact. Know-how is a kind of knowledge and thus it makes forming a belief in the face of insufficient evidence *epistemically* rational. It gives one permission, despite the lack of sufficient evidence, to form a belief. This is why forming an intention, even though (according to cognitivism) it involves forming a belief in the face of insufficient evidence, is rational (Setiya, 2008).

There are two problem with this response. First, it seems false that were one to form the intention to do something without knowing how to do it, then one's formation of this intention would be irrational. It seems perfectly rational, for instance, to form the intention to dance the tango at one's wedding, even if one does not know, right now, how to dance the tango. In fact, it seems that such an intention would, typically, motivate one to learn how to do it. Second, this response can't explain how it is supposedly rational to form a belief *out of a desire for its truth*. In a case where one knows how to perform two incompatible options, one must decide to perform just one. The option an agent chooses will be the one she rationally desires more. However, although it seems rational to form the intention to perform an action out of a desire to perform it over another, it seems irrational to form the belief that she is performing

an action out of a desire to perform it over another. And since she knows how to perform both actions and thus could be performing the other if she so chose, it is not clear how her know-how could give her epistemic permission to form the belief that she is performing the one action over the other. In sum, it is doubtful that know-how underwrites the formation of an intention or that it could if intention were a belief (Paul, 2009b).

So, the cognitivist explanation of practical knowledge faces problems. None of them are decisive against cognitivism. Perhaps, cognitivism can overcome these difficulties. However, cognitivism itself faces problems. According to cognitivism, an intention to do A is a belief that one is doing A, but the problem is that the latter does not seem required for the former. For instance, I can intend to *be properly putting on makeup with my eyes closed* but not believe that I am (because such an action is very difficult) (Davidson, 1980b [1978]). I can even intend to be doing something while not believing that I'm taking any means toward doing it. I can intend to be repairing my television by beating on it while quite confident that I'm making things worse (i.e., that my beating is a means to breaking the television).[3]

Thus, not only is there problems with the cognitivist's explanation of practical knowledge, but also it is doubtful that cognitivism is the right place to start. This recommends, at least, that we assess the advantages of an alternative, *non-cognitive* explanation of practical knowledge.

2. Inferentialism

The non-cognitive inferentialist explanation of practical knowledge goes like this. For the reasons outlined in the previous section, we should abandon the thought that practical knowledge is non-evidential. This also avoids the rather mysterious claim that practical knowledge or belief is formed by considering the reasons to act and not the reasons to believe. Instead, we should accept the more modest proposal that practical knowledge is *non-observational*. That is to say, practical knowledge is not derived from observation, but it is derived from evidence of some kind. An agent does discover (or infer) what she is doing, but she discovers it in a way that is unique to agents and unlike how an outside observer would discover what she is doing.

[3] See McCann (1991) and Paul (2009b) for even more problems.

For example, when I accidentally call you by sitting on my phone but hear the phone ringing, I come to know that I am calling someone in the same way that the person next to me might come to know that. I know it by observation. But when I decide to call you and reach for my phone, I come to know that I am calling you in a way others around me could not. I come to know it in a way that only the agent of the action could. I know it agentially.

And this is because, according to inferentialism, an agent can infer what she is doing from knowledge of her intention and ordinary background knowledge. This is something that only the agent of the action can do because only the agent has special access to knowledge of her intention. Ordinarily, if an agent decides to do A, she forms the intention to do A and comes to have self-knowledge of this intention. It is an easy inference from this knowledge and ordinary background knowledge (e.g., knowledge of her abilities and her present circumstances) to at least the justified belief that she is doing A. So, practical knowledge is inferred in this way, a way unique to agents, and not by observing what one is doing.

Further, for the reasons outlined in the previous section, we should abandon the thought that practical knowledge is necessary for intentional action. Instead, we should accept the more modest proposal that practical knowledge *often* (but not always) accompanies intentional action. When an agent is doing A intentionally, this is not because (in part) she has practical knowledge that she is doing A. Rather, it is because she is acting out of an intention to do A, which is a non-cognitive state. And yet it is true that knowledge of what she is doing often accompanies her intentional actions. So this deserves an explanation. For the inferentialist, it is because an agent often makes the easy inference from knowledge of her intention (and ordinary background knowledge) to a belief about what she is doing.

So far, the inferentialist has explained why practical knowledge is non-observational and often accompanies intentional action, but what about why practical knowledge is productive of what's known? Recall, the thought is that an action is intentional only if the action is happening because of the practical knowledge. For the cognitivist, this is because practical knowledge is embodied by the intention out of which the action is done. However, for the *non-cognitive* inferentialist, intention is not a belief, and thus this explanation is unavailable: the intention out of which the action is done cannot be knowledge because it is a non-cognitive state. How, then, is the productivity of practical knowledge to be explained?

In fact, it cannot. For the non-cognitivist, the intentional action is happening because of the intention (not any special knowledge of the intention), and for the inferentialist, as we have just seen, practical knowledge is not necessary for intentional action. Therefore, there is no reason to think that an intentional action is intentional only if it is happening because of the practical knowledge. The non-cognitive inferentialist must deny that practical knowledge is productive of what's known. She can, however, explain how it is a special sort of agential knowledge that often accompanies intentional action.[4]

The problem with this explanation is that it abandons all three central thoughts about practical knowledge. It does so in a valiant effort to avoid the problems of cognitivism and give us an epistemically familiar explanation of the most compelling components of these thoughts. It avoids the mystery of forming a belief by considering reasons for and against action and explains (at least) the apparent necessity of and distinctive basis for practical knowledge in terms of the ordinary evidential process of inference. This is laudable, but disappointing.

In Sections 4 and 5, I will argue that my theory of intention and intentional action can explain all three central thoughts (without abandoning them), while at the same time avoiding the problems with cognitivism and remaining epistemically familiar. Thus, we should prefer my explanation over both the cognitivist's and the inferentialist's.

3. Conservatism

I want to propose an explanation of practical knowledge that is epistemically familiar but does not face the problems of cognitivism or inferentialism. It does, however, come with a problem of its own. One advantage of an explanation is to be conservative, especially when what's being conserved has proven explanatory power. My proposal, unfortunately, will not be conservative. In particular, my proposal will deny that an intention is a mental state. Worse, the claim being denied has proven explanatory power. In what follows, therefore, my argument will have to work as follows. I will argue that the avoidance of all the above problems and the ensuing explanation of practical knowledge outweighs the fact that this explanation is non-conservative. This will not be decisive, but my hope is that the resultant picture will reveal

[4] See Paul (2009a) for a more detailed version of this explanation.

its promise. A thorough defence of its explanatory power will have to wait for another time.

Before this, it should be noted that the claim that intention is a mental state (statism) is not without its own problem. In particular, if an intention is a state, then the most natural account of acting with an intention is a causal one. When one acts *with* an intention, one's intention does not merely accompany the action; one's intention explains the action. For instance, if I am pouring flour with the intention to bake cookies, then I don't just happen to have the intention to bake cookies; rather, my intention to bake cookies explains why I am pouring flour. If this intention is a state, then the most natural explanation of why the intention explains the action is because the intention causes the action.

However, this explanation faces the problem of causal deviance. Suppose a burglar has the intention to distract a security guard by knocking a glass and breaking it, but this intention makes the burglar so nervous that his arm shakes, knocking a glass and breaking it. Thus, the burglar's intention causes his action, yet he doesn't shake his arm *with* the intention to distract the security guard. This is the problem of causal deviance. Although it afflicts the causal account of acting with an intention and not statism directly, the causal account seems irresistible once we accept statism. Either the statist has to deal with the problem of causal deviance or she must explain how an intention is state yet acting with an intention does not involve causation (Wilson, 1989; Sehon, 1997).

Again, this is not a decisive objection against statism, but it does suggest that being non-conservative about a theory of intention has a further advantage. Not only will my non-conservative explanation of practical knowledge avoid the problems from Sections 1 and 2, it can also avoid the problem of causal deviance—a problem we must face if we remain conservative.

4. Acting as a reason

On my view—the acting-as-a-reason view—an intention to do E is (among other things) an intentional action of type E. An intentional action is an action done *as a* reason (not for a reason). Thus, an (explanatory) reason for action is an intentional action. When an agent does M out of an intention to do E or (i.e.) with the intention to do E, this intention is her reason for doing M and an intentional action of type E. The explanatory relation between a reason and an action is not causal but mereologi-

cal. The reason is a wider action of which the given action is a narrower part. When an agent does M out of an intention to do E, her action of type M is a narrower part of her action of type E. And this makes her action of type E an intentional action.

There are three major attractions of this view. First, it can give us a uniform explanation of acting for a reason. A want, an intention, an attempt, and a doing can all be reasons for action. On this view, this is because all are intentional actions in progress (Thompson, 2008). Whether (e.g.) Jane is raising her arm because she wants to hail a taxi, or intends to hail a taxi, or is trying to hail a taxi, or is hailing a taxi, Jane's reason for action is her intentional taxi-hailing in progress.

Second, it can explain acting with an intention without positing a sui generis explanatory relation distinct from causation. Teleologists solve the causal deviance problem by claiming that the explanatory relation between a reason and an action is not causal but irreducibly "teleological" (Wilson, 1989; Sehon, 2016). The acting-as-a-reason theory solves the problem by claiming that this relation is mereological—a kind of relation as ordinary and commonplace as causation.

Third, it can explain intentional action without positing a sui generis mental state of intention, intrinsically distinct from ordinary beliefs or ordinary desires.[5] And even better, it can explain intentional action without positing sui generis inner mental causes of outward behaviour. This theory does not need to posit that the desires and intentions that characteristically explain intentional action are special states of an agent.[6] Instead, desires and intentions are kinds of material processes, individuated not by their causal history but by the way they explain their parts.

Despite these attractions, the major objection to this view is that an agent can intend to do E without actually doing E. I have defended this view against this objection in Russell (2018). Roughly, intention is to action as tadpole is to frog. An intention to do E is an intentional action of type E at a certain developmental stage, just as a tadpole

---

[5] Davidson (1980b [1978]) aimed to avoid positing a sui generis mental state of intention but, by his own admission, failed. Mele (1992), Bratman (1999), and Setiya (2007), in different ways, have simply embraced the position.

[6] The theory that desires and intentions are sui generis inner mental causes of outward behaviour is so entrenched in contemporary analytic philosophy that it may not be obvious that its rejection is an explanatory advantage. It is more parsimonious to explain intentional action in terms of a material process with parts that play certain roles than in terms of a material process that is caused by special states of an agent specially suited to produce sophisticated behaviour.

is a frog at a certain developmental stage. Consequently, an agent can intend to do E without actually doing E, even though an intention to do E is an intentional action of type E: in the same way that a thing can be a tadpole without being a frog, even though a tadpole is a frog. A tadpole is not a mature frog, but it is a frog (at an immature developmental stage). And an intention to do E is not a doing of E or (i.e.) a mature intentional action of type E, but it is an intentional action of type E (at an immature developmental stage). There is a parallel objection for wants and attempts but also a parallel response.

So the theory is this: an intentional action of type E is an action of type E done as a reason. An intention to do E is an intentional action of type E (at a certain developmental stage). This explains why an intention is a reason for action. It also helps explain what it is to be an explanatory reason for action: it is for the reason (be it a want, an intention, an attempt, or a doing) to be the wider intentional action of which the action explained is a narrower part. The explanatory relation between explanatory reason and action explained, when the reason is the agent's reason for wanting, intending, attempting, or doing the action, is parthood.

But what makes it the case that one action is a part of another? Suppose that I am engaged in two projects: I am baking cupcakes, and I am baking bread. Right now, I am pouring flour into a bowl. In fact, this flour-pouring is being done with the intention to bake bread. That is, I am pouring flour in the bowl out of my intention to bake bread, not out of my intention to bake cupcakes. On the above theory, this is because my flour-pouring is a part of my bread-baking and not a part of my cupcake-baking. But what makes it the case that my flour-pouring is a part of the former and not the latter?

The answer to this question is the key to understanding practical knowledge. Practical knowledge is the knowledge that, characteristically, plays a role in making it the case that one action is a part of another (when the one action is the agent's reason for wanting, intending, attempting, or doing the other action).

But before we can understand this, we must understand practical reasoning.

Practical reasoning is the exercise of practical reason, and practical reason is the cognitive capacity that agent's characteristically have to advance action using knowledge. Practical reasoning has the following structure: it returns an action of type M, given an action of type E and given knowledge that action-type M is a way of instantiating

action-type E. For example, an agent with the cognitive capacity of practical reason might exercise this capacity and, as a result, pour flour, given that she is baking bread and that she knows that flour-pouring (in general) is a way of instantiating bread-baking (in general).

Some might find it useful to think of this structure as a syllogism where the first premise is an action of type E, the second premise is the knowledge that action-type M is a way of instantiating action-type E, and the conclusion is an action of type M. We could write it as follows:

1. an action of type E

2. action-type M is a way of instantiating action-type E

C. an action of type M

The agent is performing premise 1, knows premise 2, and concludes C by performing it. Practical reason is the cognitive capacity to conclude C as a result of performing 1 and knowing 2.[7]

For instance, consider the following:

1. baking bread

2. flour-pouring is a way of instantiating bread-baking

C. pouring flour

In an exercise of practical reason, the agent is baking bread (or wants to, intends to, or is attempting to) and knows that flour-pouring is a way of instantiating bread-baking. She exercises her practical reason and, as a result (if all goes well), pours flour (or wants to, intends to, or is attempting to).

Practical knowledge is, characteristically, the knowledge of premise 2. That is to say, practical knowledge is the knowledge characteristically involved in practical reason-

[7] Following Anscombe (2000 [1953], 2005 [1989]), this view posits a rational capacity distinct from a capacity to form propositional attitudes from other propositional attitudes, thus the premises are not all propositions and the conclusion is not a proposition. It is plausible that human rational capacities are not exhausted by those of forming propositional attitudes from other propositional attitudes or those that are reducible to these.

ing. It is the knowledge that action-type M is a way of instantiating action-type E *given in an exercise of practical reason*.[8] The knowledge that action-type M is a way of instantiating action-type E is not, alone, practical knowledge. Practical knowledge is this knowledge *used in practical reasoning*.

This is my theory of what it is for knowledge to be practical knowledge. I will now argue for this theory by showing how it can explain all three central thoughts about practical knowledge (without abandoning them), while at the same time avoiding the problems with cognitivism and remaining epistemically familiar.[9]

5. Practical knowledge

Recall that practical knowledge is necessary for intentional action, non-evidential, and productive. Without it, what one does is not intentional. One does not discover what one is doing, when that knowledge is practical. And, an action is intentional only if the action is happening because of the practical knowledge.

Following the schema from Section 4, practical knowledge is necessary for intentional action because if the belief that action-type M is a way of instantiating action-type E does not amount to full knowledge, then the action given to practical reasoning will not advance intentionally. If the belief is true, then the resultant action of type M will instantiate an action of type E. If the belief is full knowledge, then the resultant action of type M will not only instantiate an action type E but also this token action of type E will count as an addition to the given action of type E. For instance, if the agent reasons with knowledge that dough-kneading is a way of bread-baking, then the resultant token dough-kneading will instantiate a bread-baking and will count as an addition to the token bread-baking given to the practical reasoning.

---

[8] For simplicity, I will continue to claim in what follows that practical knowledge is knowledge that action-type M is a way of instantiating action-type E, but if this is distinct from know-how knowledge, the theory should allow that knowing how to instantiate action-type E by instantiating action-type M can also play the role of practical knowledge. Thus, practical knowledge is either knowledge that action-type M is a way of instantiating action-type E or know-how knowledge of how to instantiate action-type E by instantiating action-type M. This would allow for a kind of intentional action without propositional knowledge.

[9] This theory is further supported by the Davidsonian thought that intentional action requires an instrumental understanding of the action explained (Davidson, 1980a [1963]). To have a reason for action, it is not enough to desire some property of the action, or intend to bring about that property, or be attempting to bring about that property, or even be bringing about that property. To have a reason for action, the agent must understand that the action explained can instantiate that property. On my view, this understanding is practical knowledge.

If the belief is false, then the action of type M will not instantiate an action of type E and thus cannot count as an addition to the given action. If, for instance, flipping the switch is not in these circumstances a way of turning on the light, then a token act of flipping the switch will not count as turning on the light and thus cannot count as an addition to some ongoing process of turning on the light. Further, if the belief is true but only as a matter of luck, then although the resultant action of type M will instantiate an action of type E, this will be only as a matter of luck. Thus, the action of type M will not count as an *intentional* advancement of the action given to practical reasoning. If one were simply lucky in thinking that a certain switch-flipping is a way of turning on the light, then one's token switch-flipping would count as turning on the light but this advancement of the ongoing process would not be intentional: one would have turned on the light by accident. Therefore, without full knowledge, the resultant action of practical reasoning will not count as an intentional advancement of the action given to practical reasoning—i.e., it will not count as something the agent is doing (or wanting, or intending, or attempting to do) intentionally.[10]

Next, practical knowledge is productive because were it not for the knowledge, the action given to practical reasoning would not be advancing. When an intentional action of type E is advancing, if the knowledge that action-type M is a way of instantiating action-type E were completely absent, the practical reasoning would have no value to return and thus would not be returning an action of type M. The reasoning would have nothing but the given action to work with. If the knowledge were, instead, mere belief, the practical reasoning would be returning some action, but that action would not count as an advancement of the given action. Therefore, an intentional action, when it is advancing, is advancing (in part) because of the knowledge given to practical reasoning.

Finally, practical knowledge is non-evidential because its role is neither to be evidence nor to be gained from evidence. Practical reason is the capacity to advance action using knowledge, while theoretical reason is the capacity to derive knowledge from other knowledge. The knowledge characteristically used in practical reasoning is practical knowledge. The knowledge characteristically used in theoretical reasoning is evidence. The result of practical reasoning is a narrower action derived from a wider action and practical knowledge. The result of theoretical reasoning is knowledge derived from evidence. Practical knowledge is neither characteristically the

[10] For further defence in the same vein, see Gibbons (2001) and Pavese (2018).

knowledge used in theoretical reasoning nor characteristically the result of theoretical reasoning. Thus, as practical, it is not evidence nor derived from evidence. Syllogistically, we would say that practical knowledge is neither knowledge of a premise nor the knowledge concluded in theoretical reasoning.

So, practical knowledge is made practical by being knowledge of a premise in practical reasoning. It is not made practical by being the knowledge *concluded* in practical reasoning. That is, it is not made practical by being derived from some special, practical source. Also, practical knowledge is not made *knowledge* by being derived from some special, practical source. Rather, it is made knowledge in the familiar ways that propositions are made knowledge. For instance, it may be made knowledge by being derived from observation or evidence. But this knowledge is of a special kind because it takes a characteristic place in practical reasoning, not because it takes a characteristic place (premise or conclusion) in theoretical reasoning—that's why it is notably non-evidential.

On this picture, practical knowledge has a direction of derivation (we might say) that runs reverse of theoretical knowledge. Roughly, with theoretical knowledge, the knowledge (both content and correctness) are derived from the fact of the matter, while with practical knowledge, the fact of the matter (both content and correctness) are derived from the knowledge. More precisely, with theoretical knowledge, the result of theoretical reasoning is determined by the content of the evidence, which depends on the fact of the matter, and a belief in disagreement with this determination is a mistaken judgment. For instance, if the fact of the matter is that one is raising the temperature of the Earth, then sufficient evidence used in theoretical reasoning will determine that one is raising the temperature of the Earth and to believe otherwise would be a mistake. With practical knowledge, the result of practical reasoning is determined by the practical knowledge and an action in disagreement with this determination is a mistaken performance. For instance, if the agent has practical knowledge that stroking is a way of swimming, then this will determine a stroking and for the agent to be performing otherwise would be a mistake.

In sum, practical knowledge is the knowledge used in practical reasoning, which is distinct from theoretical reasoning. It is not, *characteristically*, the result of practical reasoning or theoretical reasoning, although as knowledge, it may be the result of the latter. Practical knowledge, when it is genuine knowledge and not mere belief, advances the action given to practical reasoning by determining the action returned by

practical reasoning. The action that results from the reasoning becomes a part of the given action and thereby the given action advances intentionally. When practical knowledge is absent or not genuine knowledge (i.e., false or accidentally true belief), the given action does not advance intentionally, and thus if it does advance, this is not intentional.

When I call you intentionally by reaching for my pocket, the reaching is the result of wanting to or intending to or attempting to call you and from this, reasoning practically using the knowledge that reaching for my pocket is a way of instantiating a call to you. Absent this knowledge, I would not be reaching for my pocket and thus not be calling you intentionally in this way. And if the knowledge were accidentally true belief, I would be reaching for my pocket and thus calling you but not intentionally.

When I call you unintentionally by sitting on my phone, the sitting is not the result of reasoning practically from a given action, and thus the sitting, although it advances an action of calling you, is not a part of the advancement in the same way that the reaching is. Knowledge that sitting on my phone is a way of calling you is not producing the advancement by making it a part of some action I am performing.

This explains our three central thoughts about practical knowledge. Unlike the cognitivist explanation, practical knowledge is not knowledge gained by forming an intention. And unlike the inferentialist explanation, practical knowledge is not knowledge gained by reasoning from knowledge of intention. Practical knowledge is not, characteristically, knowledge gained but knowledge used. Practical knowledge is part of a special way of reasoning with knowledge, not a result of a special way of reasoning. Relatedly, the content of practical knowledge is not that the agent is performing her intentional action. Rather, the content of practical knowledge creates the fact that the agent is performing her intentional action.

This explanation of practical knowledge is epistemically familiar because it does not claim that practical knowledge is some form of licensed wishful thinking or licensed leap of faith or otherwise epistemically unordinary belief. Rather, practical knowledge, as knowledge, conforms to the very same evidential standards as ordinary knowledge. As knowledge, it is ordinary knowledge. It is knowledge of some instantiation relation between two types of action. Its status as practical does not mean that it is ungrounded or grounded in some special way; rather, it simply means that it plays a special role in action.

And this explanation does not abandon our three central thoughts about practical knowledge.

6. Diagnosis of the error

What led the cognitivist and inferentialist into their mistaken explanations? In the first instance, it is the ambiguity in the expression 'what one is doing'. Practical knowledge is knowledge of *what one is doing*. Read one way, this says that practical knowledge is knowledge of the *token* action one is performing intentionally. Call this 'the token interpretation'. Read another way, it says that practical knowledge is knowledge of the *type* of action one is performing intentionally. Call this 'the type interpretation'. On the token interpretation, the content of practical knowledge includes the particular action one is performing specified by the type under which it is intentional. Thus, if an agent is doing A intentionally, then the content of her practical knowledge is that she is doing A. On the type interpretation, the content of practical knowledge need not include the particular action one is performing. To say that practical knowledge is knowledge of what one is doing is to say that such knowledge is of a certain general entity—namely, the type itself under which the agent's action is being performed intentionally. Thus, if an agent is doing A intentionally, then the content of her practical knowledge includes action-type A.

The cognitivist's and inferentialist's initial mistake is to understand practical knowledge under the token interpretation. Therefore, they read 'knowledge of what one is doing' as 'knowledge that one is doing A' (where A is the type under which the action is intentional). If an agent is baking bread intentionally, then she knows that she is baking bread. Thus, the challenge becomes to explain how such knowledge is necessary for intentional action, non-evidential, and productive. The problems with the cognitivist's explanation arise because the agent must be epistemically permitted to believe this content prior to having sufficient evidence for it. The problems with the inferentialist's explanation arise from correcting this by finding a source of sufficient evidence for the very same content. My view avoids these problems by correcting the content of practical knowledge. I read 'knowledge of what one is doing' as 'knowledge of action-type A' (where action-type A is the type under which the action is intentional).

The cognitivist's and inferentialist's concomitant mistake is to conflate practical knowledge with a different phenomenon. In particular, often intentional action is

accompanied by knowledge that one is doing A (where A is the type under which the action is intentional), and often this knowledge is known in a special way. If an agent is baking bread intentionally and an observer (hearing some clanging in the kitchen) asked, "what are you doing?", the agent can often answer and will believe that she is baking bread. By understanding practical knowledge under the token interpretation, the cognitivist and inferentialist conflate practical knowledge with this phenomenon. The cognitivist, thus, gives a problematic explanation of a phenomenon that is not practical knowledge. The inferentialist gives a much better explanation of this phenomenon, but this explanation is not an explanation of practical knowledge. This is why the inferentialist's description of practical knowledge seems so far removed from our initial description of practical knowledge.

So, the cognitivist and the inferentialist explain the wrong kind of knowledge. Now I will argue that my explanation of practical knowledge gives us the resources for a better explanation of even this other kind of knowledge.

I have already outlined the problems with the cognitivist's explanation of agential knowledge. Let me now demonstrate a problem for the inferentialist's explanation. For the inferentialist, agential knowledge (that one is doing A) often accompanies intentional action because an agent often makes the easy inference from knowledge of her intention to do A and ordinary background knowledge to a belief that she is doing A. This inference is easy because what she comes to believe is often made true simply by being the right sort of agent (ready, willing, and able) in the right sort of circumstances (nothing stands in her way). When an agent is reliable, intends to do A, has the ability to do A, and nothing stands in the way of completing the action, there is a minimal sense in which she is in the process of doing A. This is the content of agential knowledge for the inferentialist: that she is doing A (in this very minimal sense that is made true simply by being the right sort of agent in the right sort of circumstances) (Paul 2009a).

The problem with this explanation is that agential knowledge is not so minimal. When the agent answers that she is baking bread, she believes and asserts more than that a bread-baking is minimally underway. She believes and asserts that *she* is advancing a bread-baking (i.e. intentionally). The agent would rethink what she was asserting if it were pointed out that the steps she was taking were ineffective or setting the process back, and the agent would not do so if what she believed and asserted depended merely on being the right sort of agent in the right sort of circumstances.

Advancing an action intentionally involves more than being a certain sort of agent in certain circumstances. It also involves knowing (not merely believing) some way to advance the action and using this knowledge in an exercise of practical reason. My theory gives the inferentialist the resources for a better explanation of agential knowledge: such knowledge often accompanies intentional action because an agent often makes the easy inference from knowledge of her intention to do A, ordinary background knowledge, and knowledge of her practical reasoning to the belief that she is advancing an action of type A intentionally. She knows this in a way unique to the agent of the action.[11]

Thus, my account not only gives us a better explanation of practical knowledge than what we have seen. It also helps us give a better (inferentialist) explanation of agential knowledge than what we have seen.

7. An objection

Some might object as follows. An agent can want to do E (even, intend to do E) without exercising practical reason. For instance, I can suddenly want to eat ice cream without thinking at all about it. But, on my view, a want is an intentional action, and an intentional action is an action done as a reason or (i.e.) as the given action in an exercise of practical reason. Therefore, if an action is intentional, it is given in an exercise of practical reason. Thus, if an agent wants to do E, her action of type E must be given in an exercise of practical reason. But this conflicts with the fact that an agent can want to do E without exercising practical reason. (A parallel objection can be given about intention.)

My response is to deny that an agent can want to do E without exercising practical reason. An agent can want to do E, *in a certain sense*, without exercising practical reasoning. And an agent can want to do E without exercising practical reason, *in a certain sense*. But an agent cannot want to do E without exercising practical reason, in the sense I mean 'want' and 'practical reason' throughout my account.

An agent can want to do E without exercising practical reason, when this want is not the kind of thing that can be an agent's (explanatory) reason for action. For instance,

[11] This is an additional benefit, but my account is not committed to it. My account could also provide helpful resources for a non-inferential explanation of agential knowledge, but I'll provide no argument for this here.

I can crave some ice cream, and thus, in some sense, I want to eat ice cream. But such a thing cannot be an agent's reason for action. If an agent does something out of a craving, this is not the agent's reason for action (as agent). Instead, it explains why the agent acts (as animal, or as biological being). An explanatory reason for action, in the sense that is relevant to intentional action and agency, is one that explains why the agent acts (as agent), and such an explanation must be in terms of an agent's characteristic capacity to advance action using knowledge. There is a kind of want that can provide such an explanation. It is this kind of want that I mean to refer to throughout my account and that an agent cannot have without exercising practical reason.

Further, an agent can want to do E without exercising practical reason, in the sense of an explicit presentation of the reasoning in one's mind's eye (or ear) or in written or audible speech acts. In other words, an agent can want to do E while the practical reasoning is implicit—the mere exercise of one's capacity to advance action using knowledge without any explicit presentation of the reasoning to oneself or others. But, on my view, an agent cannot want to do E, in the relevant sense, without exercising this capacity at all.

Note also that the action given to practical reasoning, the practical reasoning itself, and the action resulting from practical reasoning can all occur simultaneously, just as a tree can produce a blossom, exercise its capacity to reproduce, and produce a petal all at the same time. The production of the petal is an exercise of the tree's capacity to reproduce and part of the production of the blossom all at once and at the same time. Similarly, an action of type M can, sometimes, be an exercise of an agent's capacity to advance action using knowledge and part of an action of type E all at the same time.

A typical case of wanting might look like this: an agent wants to eat ice cream, exercises her practical reason, and deliberates about whether to eat ice cream. That is, the action given to practical reasoning is a token ice-cream-eating. The knowledge given is that deliberating about whether to eat ice cream (in general) is a way of instantiating an ice-cream-eating (in general). And the resultant action is a token deliberating about whether to eat ice cream. Like the tree, the deliberating is an exercise

of the agent's capacity to advance action using knowledge and is part of the ice-cream-eating, all at the same time.[12]

Once we recognize that (a) a want, in the relevant sense, must be the kind of want that can be a reason for action and that (b) an exercise of practical reason can be implicit and that (c) a want can be composed of the very implicit exercise of practical reason it is given to, it is much more plausible to claim (as my theory implies) that an agent cannot want to do E without exercising practical reason. (A parallel response can be given to the parallel objection about intention.)

Conclusion

An agent, characteristically, has practical reason: the capacity to advance action using knowledge. The action that advances by the exercise of this capacity just is an intentional action. And the knowledge used in the exercise of this capacity just is practical knowledge.

Consequently, practical knowledge is necessary for intentional action. With anything less than full knowledge, the resultant action of practical reasoning will not count as an intentional advancement of the action given to practical reasoning. Further, practical knowledge is productive of what's known. An intentional action, when it is advancing, is advancing (in part) because of the knowledge given to practical reasoning. And practical knowledge is non-evidential. Practical knowledge, as part of practical reasoning, is not evidence for an agent's reason for action but creates an agent's reason for action by making the action given to practical reasoning intentional.

Because this explanation of practical knowledge avoids the problems with the dominant cognitivist explanation and the dominant inferentialist explanation and it is made possible by the theory that acting intentionally is acting as a reason, we should accept this theory. We should flip our understanding of intentional action from action performed for a reason to action performed as a reason. Not only does this theory give us an attractive picture of practical knowledge, it also gives us an attractive picture of an explanatory reason for action: one that avoids the problems for statism

[12] As explained earlier, the ice-cream-eating is not a doing but a want. So the agent is not eating ice cream but merely wants to eat ice cream. Yet the want is still an action—in this case, realized by a deliberating.

and causalism and that explains why wants, intentions, attempts, and doings are all explanatory reasons for action.

REFERENCES

Anscombe, G. E. M. (2000). *Intention*. Harvard University Press.

Anscombe, G. E. M. (2006). Practical inference. In Geach, M. & Gormally, L. (Eds.). *Human life, action and ethics: essays by G. E. M. Anscombe*. Luton: Andrews UK.

Bratman, M. E. (1999). *Intention, plans, and practical reason*. Harvard University Press.

Davidson, D. (1980a). Actions, reasons, and causes. In *Essays on actions and events*. Oxford: Clarendon Press.

Davidson, D. (1980b). Intending. In *Essays on actions and events*. Oxford: Clarendon Press.

Gibbons, J. (2001). Knowledge in action. *Philosophy and phenomenological research*, 62(3), 579–600.

Grice, H. P. (1972). *Intention and uncertainty*. Oxford University Press.

Marusic, B. & Schwenkler, J. (2018). Intending is believing: a defense of strong cognitivism. *Analytic philosophy*, 59(3), 309-340.

McCann, H. J. (1991). Settled Objectives and Rational Constraints. *American Philosophical Quarterly*, 28(1), 25–36.

Mele, A. R. (1992). *Springs of action: understanding intentional behavior*. Oxford University Press.

Paul, S. K. (2009a). How we know what we're doing. *Philosophers' Imprint*, 9(11), 1–24.

Paul, S. K. (2009b). Intention, Belief, and Wishful Thinking: Setiya on "Practical Knowledge." *Ethics*, 119(3), 546–557.

Pavese, C. (2018). Know-how, action, and luck. *Synthese*, doi: 10.1007/s11229-018-1823-7.

Russell, D. (2018). Intention as action under development: why intention is not a mental state. *Canadian journal of philosophy*, 48(5), 742–761.

Sehon, S. R. (1997). "Deviant Causal Chains and the Irreducibility of Teleological Explanation." *Pacific philosophical quarterly*, 78(2), 195–213.

Sehon, S. R. (2016). *Free will and action explanation: A non-causal, compatibilist account*. Oxford University Press.

Setiya, K. (2007). *Reasons without rationalism*. Princeton University Press.

Setiya, K. (2008). Practical Knowledge. *Ethics*, 118(3), 388–409.

Sinhababu, N. (2013). The desire-belief account of intention explains everything. *Noûs*, 47(4), 680–696.

Thompson, M. (2008). *Life and action: elementary structures of practice and practical thought*. Harvard University Press.

Velleman, D. (1989). *Practical reflection*. Princeton University Press.