SEPARATING ACTION AND KNOWLEDGE*

Mikayla Kelley
University of Chicago

Penultimate Draft. Forthcoming in *The Philosophical Quarterly*

April 4th, 2024

Abstract

Intentional action is often accompanied by knowledge of what one is doing—knowledge which appears non-observational and non-inferential. G.E.M. Anscombe defends the stronger claim that intentional action always comes with such knowledge. Among those who follow Anscombe, some have altered the features, content, or species of the knowledge claimed to necessarily accompany intentional action. In this paper, I argue that there is no necessary connection between intentional action and knowledge, no matter the assumed features, content, or species of the knowledge. Further, rather than follow the usual methodology in this debate of arguing by counterexample, I present an argument that explains why we continue to find counterexamples: intentional action and knowledge are regulated by thresholds under distinct pressures; in particular, the threshold of control regulating intentional action is disparately influenced by the role of intentional action in practical normativity.

1. INTRODUCTION

When we act intentionally, we often know what we are doing. Moreover, we seem to know what we are doing in a special, non-observational way. Say I am walking to the store. While walking to the store, I know that I am walking to the store. I do not know this by looking down at my flip-flops landing one in front of the other moving in the direction of the store. Nor does it seem like I infer that I am walking to the store. I seem to just know what I am up to. Is intentional action *always* accompanied by this special, non-observational knowledge of what one is doing?

In her groundbreaking monograph *Intention*, Anscombe (1957) argues 'yes'. She defends the *Knowledge Thesis*: while intentionally φing, one knows that one is φing without observation. There is, to quote Setiya (2016b), "a profound conflict in action theory between those who follow Anscombe in seeing a close connection between intentional action and knowledge of what one is doing and those who do not" (p. 65). The reason for the conflict stems at least in part from what looks like a wealth of counterexamples to the Knowledge Thesis and related claims. For example, it seems to some that I need not know

*For helpful comments and discussion, I thank Michael Bratman, Dan Friedman, Wesley Holliday, Nadeem Hussain, Taylor Madigan, Sarah Paul, Juan Piñeros Glasscock, Calum McNamara, Antonia Peacocke, Sam Rogers, multiple anonymous referees, and audiences at Stanford University and the 2023 Pacific APA.

or even believe that I am scoring a penalty shot in order to be intentionally doing so; and if I do know or believe that I am scoring a penalty shot while doing so, surely that could be in part on grounds of my observing, e.g., the ball moving past the goalie into the net. Some take these examples to show that the Knowledge Thesis is too strong, and others—whom I will call *Anscombeans*—try to retain the core insight of the Knowledge Thesis in the face of purported counterexamples.[1]

There have been attempts to modify the Knowledge Thesis in at least four ways.[2] First, some have challenged the claim that knowledge of our actions is non-observational,[3] some suggesting that we know what we intend without observation, but we know—or at least might know—what is actually happening observationally.[4] Others have suggested alternative contents for the knowledge necessary for action; the content of the knowledge that is necessary for intentionally φing is not that one is φing but rather, e.g., that one is ψ-ing for some description 'ψ' of φ.[5] Some have changed the species of knowledge in question. Anscombe is clear that she means for her thesis to be about propositional knowledge—one knows *that* one is φing while intentionally φing—but some have suggested that the knowledge which necessarily accompanies intentional action is instead, e.g., knowledge how.[6] Finally, some have weakened the generality of Anscombe's claim, allowing that the Knowledge Thesis admits of exceptions.[7]

In this paper, I argue that there is good reason to think that there is no necessary connection between intentional action and knowledge, no matter the assumed features, content, or species of the knowledge. That is, we have good reason to think that the following schematic thesis has no true precisifications:

> **Schematic Knowledge Thesis**: while intentionally φing, one knows$_T$ p_φ in way F, where knows$_T$ is some species of knowledge (e.g., propositional or knowledge how), p_φ is some content about one's φing of the appropriate form for the species of knowledge know$_T$, and F specifies certain features of the knowledge (e.g., non-observational or non-inferential).

Therefore, I argue that any plausible weakening of the Knowledge Thesis must give up on its full generality. For the sake of space and because I am most confident of the argument in the case of propositional knowledge, I will limit my discussion to versions of the Schematic Knowledge Thesis that posit necessary propositional knowledge—referring to this restricted thesis as the *Propositional Schematic Knowledge Thesis* (PSKT). I will also sometimes focus on the Knowledge Thesis, offering only a proof of concept. I try to note throughout how to generalize the discussion.

Many have recognized the need to weaken the generality of the Knowledge Thesis and related theses, but this recognition

[1] Those who have defended the Knowledge Thesis in its original (or stronger) form include: Hampshire (1959), Newstead (2006), Rödl (2011), Thompson (2011), Small (2012), Stathopoulos (2016), Pavese (2021b), and Beddor and Pavese (2021).

[2] Some modify the Knowledge Thesis in multiple of the following ways.

[3] Throughout, I use 'action' to refer to intentional action.

[4] See, e.g., Jones 1960, Donnellan 1963. See also Grice (1971) and Paul (2009a) for inferentialist accounts of knowledge of intentional action.

[5] See, e.g., Davidson 1980 (see also Hunter 2015), O'Brien 2003, Pavese 2020.

[6] See, e.g., Setiya 2016a, Setiya 2016b, Kearl forthcoming.

[7] See, e.g., Velleman 1989, Falvey 2000, Schwenkler 2019.

is often in response to purported counterexamples, e.g., Davidson's much discussed ten carbon copies example (Davidson 1980, p. 50).[8] Here I try to offer more. I try to explain why we continue to find counterexamples to the Knowledge Thesis and related theses. The explanation is in the spirit of Harman's (1976) separation of intentional action and belief and Bratman's (1987) separation of intentional action and intention in that the explanation relies on the influence of practical normativity on the classification of behaviors as intentional action. More precisely, using 'ethical' permissively to refer to all of practical normativity, I will argue that ethical facts fix the threshold of control required for acting intentionally, and insofar as ethical considerations affect the threshold of warrant required for knowledge, they do so differently. I argue that this differential influence from ethics on the thresholds that regulate the classification of behaviors as intentional action and beliefs as knowledge, respectively, makes any precisification of PSKT unlikely to be true.

The plan for the rest of the paper is as follows. In Section 2, I present the Argument from Distinct Thresholds, the central argument of this paper. I present a related explanatory challenge for the Anscombean in Section 3. In Section 4, I defend a key premise of the Argument from Distinct Thresholds, namely that ethical facts fix—or at least influence—the threshold of control required for acting intentionally. In Section 5, I respond to an objection from encroachment.

2. The Argument from Distinct Thresholds

Anscombeans have been concerned mostly with answering the following two questions:

- Why think that one believes that one is φing whenever one is intentionally φing?
- Even if one believes that one is φing whenever one is intentionally φing, why think there is warrant for the belief?

The motivation for trying to answer the first question stems from the uncontroversial view that one can only know what one believes. And without an answer to the second question, we should be—and many have been—skeptical that said belief neecssarily has the epistemic stature of knowledge.[9] Note that I use the idea of warrant as a theoretical placeholder for whatever makes the difference between true belief and knowledge. The only thing I assume about warrant is that it is gradable.

There is an important further question that can go unanswered even if these two questions are successfully answered. Because warrant comes in degrees, it is not enough to establish that *some* warranting conditions for a belief about one's behavior necessarily hold if the behavior is intentional. For example, it is not enough to establish that there is necessarily some evidence in favor of the belief, necessarily a somewhat reliable connection to the truth, necessarily some safety, necessarily a defeasible right to the belief, or necessarily some degree of intellectual virtue that produces the belief. To establish that the

[8] For an exception, see Piñeros Glasscock 2019.

[9] See, e.g., Grice 1971, Langton 2004, and Paul 2009b for objections to Anscombeans along these lines.

belief is necessarily knowledge, one needs to establish that there is necessarily *enough* evidence, necessarily a reliable *enough* connection, necessarily *enough* safety, necessarily not *too* strong countervailing evidence, or necessarily *enough* intellectual virtue. So there is a third question that Anscombeans should be concerned with:

- Even if one necessarily believes that one is φing while intentionally φing and even if the belief is necessarily warranted, why think that the belief is necessarily sufficiently warranted to count as knowledge?

I am assuming then that a crucial feature of (the concept of) knowledge is that it has what I will call a *regulating threshold*: for some belief to count as knowledge, it must meet a threshold of sufficient warrant.[10] More generally, call a concept with a regulating threshold—that is, a minimal degree of some gradable feature that something must have in a given circumstance to fall under that concept—a *threshold concept*.

The concept of intentional action is a threshold concept as well: in order for some behavior to be an intentional action, the agent must exert sufficient control of the right kind over the behavior. Making a basket that is really far away as a novice basketball player such that one's making the basket is largely a matter of luck does not count as intentionally making a basket. One might attempt and succeed at baking a delicious cake but not do so as an intentional action because one fails to exert enough control over the process, making the resulting success too accidental. So the concepts of knowledge and intentional action are threshold concepts, the former having a *warrant threshold* and the latter having a *control threshold*.

Moreover, if we use 'control' schematically to pick out whatever property marks the difference between a mere φing and an intentional φing, then control is also sufficient for intentional action. And indeed, I will assume a highly schematic characterization of intentional action on which to intentionally φ *just is* to successfully φ with a sufficient degree of the right kind of control over one's φing, where "the right kind of control" and what counts as a sufficient amount of that kind of control are theoretical placeholders.[11] All that I assume about control is that it is gradable. Examples of the right kind of control include causal guidance by a practical representation or irreducible agent, the utilization of skill to advance one's goals, irreducible goal-directedness, manifestation of epistemic goods like knowledge (though, of course, my goal is to ultimately argue against such a theory of control), or a disjunctive criterion requiring at least one of these kinds of control. The right kind of control could even be a (gradable) unanalyzable notion, thus leaving significant room for anti-reductionist approaches (e.g., Levy 2013). Given how schematic this characterization of intentional action is, I take it to be common ground among many in the debate over the Knowledge Thesis.

From these minimal claims about knowledge and intentional action, we can derive an argument against PSKT. Recall

[10] That there is such a threshold is made explicit in discussion of "the threshold problem for knowledge" (Hetherington 2006).

[11] I directly defend this schematic characterization of intentional action in Kelley msa. But many theories of intentional action are consistent with it, for many theories of intentional action appeal to something gradable and control-like as marking the difference between what you do intentionally and what merely happens to you. And this is for good reason: the idea of control is one of the first ideas one might reach for in trying to capture the difference between a φing that you perform intentionally and a φing that merely happens to you. Theories of intentional action then represent attempts to more determinately describe the relevant notion of control.

again PSKT: while intentionally φing, one knows that p_φ in way F. Any claim of this form implies a necessary connection between the control threshold and the warrant threshold: whenever an agent's φing meets the control threshold, the agent has an accompanying belief with some content about one's φing that meets the warrant threshold. But, as I will argue, we have good reason to think that there are no necessary connections between the control threshold and the warrant threshold. And so there is good reason to think that there is no true precisification of PSKT. I call this the *Argument from Distinct Thresholds*. In schematic form:

P1. Knowledge requires meeting a certain threshold, namely a threshold of sufficient warrant.

P2. Intentional action requires meeting a certain threshold, namely a threshold of sufficient control.

P3. Any true precisification of PSKT implies some necessary connection between these two thresholds. More precisely, any true precisification of PSKT implies the following thesis: whenever an agent's φing meets the control threshold, some belief with some proposition about one's φing as content meets the warrant threshold.

P4. (*Separation Thesis*) We have good reason to think that there are no necessary connections between these two thresholds.

C. Therefore, we have good reason to think that there is no true precisification of PSKT.

In this essay, I assume P1, P2,[12] and P3 and focus attention on the Separation Thesis.[13]

3. An Explanatory Challenge for the Anscombean

A first thing to note with regards to the Separation Thesis is that if there *is* a necessary connection between the warrant and control thresholds, then this calls out for explanation. Indeed, the two thresholds govern the application of two distinct concepts which have their respective homes in what at least appear to be two distinct domains of our lives—the practical and the theoretical. This makes it at least prima facie surprising that they would be connected as a matter of necessity. As noted, there have been plenty of explanations for the existence of a warranted belief associated with intentional action, but this does not get us to a precisification of PSKT; we need, in addition, an explanation for why this warranted belief always meets the threshold of sufficient warrant required for knowledge. Thus, we have isolated an *explanatory challenge* for the Anscombean.

[12] I therefore leave open the (in my view, theoretically costly) option of rejecting one of the highly schematic and orthodox characterizations of knowledge and intentional action captured in P1 and P2.

[13] The Argument from Distinct Thresholds extends to the Schematic Knowledge Thesis insofar as any species of knowledge has a regulating threshold which is under distinct pressures from those shaping the control threshold, such that a thesis analogous to the Separation Thesis is true with respect to it. So if, for example, possessing knowledge of how to do something has a regulating threshold under distinct pressures from those shaping the control threshold, then the Argument from Distinct Thresholds applies to a precisification of the Schematic Knowledge Thesis in terms of knowledge how. Thus, if one is an intellectualist about knowledge how, then the extension to knowledge how is relatively clear; if one is an anti-intellectualist, the extension is less obvious since it is not clear that knowledge how is attributed via some regulating threshold.

In Sections 4 and 5, I will show in detail why I am skeptical that the challenge can be met. In the rest of this section though, I show that two strategies for defending PSKT fail because they fail to meet this explanatory challenge.

First, some have tried to explain why intentional action requires knowledge by developing an account of control in terms of knowledge. For example, Beddor and Pavese (2021) argue for the Knowledge Thesis by giving an epistemic theory of control: someone is in control of their φing at t if and only if they know that they are φing at t, and they know this in virtue of exercising their knowledge of how to φ (p. 5).[14] Since intentionally φing requires control over one's φing, intentionally φing requires knowledge that one is φing. Beddor and Pavese are not forthcoming about the fact that control is gradable, but it is *sufficient* control that is required for intentional action, not just any amount of control whatsoever. Thus, we should assume their epistemic theory of control is a theory of *sufficient* control. Now we should ask: why think that one controls one's φing to a degree sufficient for intentional action only if one's belief that one is φing is sufficiently warranted to be knowledge?[15]

Beddor and Pavese put forth an inference to the best explanation to answer this question. They offer a number of examples where intentional action and sufficient control, on the one hand, and knowledge, on the other, come together and fail together. While these examples might give us good grounds to infer that intentional action and knowledge *often* come together and fail together—or that sufficient control and knowledge often come together and fail together—why should we infer from a limited range of cases the far stronger claim that they *always* come together and fail together? Put in terms of thresholds, these examples might give us good grounds to infer that the control threshold and the warrant threshold often line up in the way needed for PSKT, but why should we infer that they *always* so line up? These questions are particularly pressing given that critics of PSKT offer examples that look like cases where the control threshold and the warrant threshold *do not* line up (e.g., controlling making ten copies enough to do so intentionally while lacking sufficient warrant to know that you are). It is thus not clear that an epistemic theory of control actually explains the data once the data is enlarged, let alone is the best explanation of it.[16]

Taking a different strategy, Pavese (2021b) points out that control requires something like modal robustness—a φing is not controlled if one might easily have failed to φ; but then the belief that one is φing could not easily have been wrong either, which secures a level of safety for the belief. But this fails to answer the crucial question of why we should think that the degree of modal robustness needed for intentional action is precisely the degree of modal robustness needed for the relevant belief to be knowledge.[17] Why couldn't it be that at least sometimes one's φing is sufficiently modally robust to be intentional

[14] See also Pavese and Beddor 2023.

[15] For additional criticisms of epistemic theories of control, see Carter and Shepherd 2023.

[16] Moreover, there are alternative theories of control on offer that explain why control often comes with knowledge but need not (see, e.g., Shepherd 2021 and Kelley msa). Or if Carter and Shepherd (2023) are right that knowledge does not explain control but rather control explains knowledge, we'd also have an explanation of Beddor and Pavese's examples.

[17] Shepherd and Carter (2023) and Carter and Shepherd (2023) suggest that intentional action and knowledge have different standards for safety. See also Schwenkler 2019. I see the Argument from Distinct Thresholds as generalizing and further bolstering this basic insight about safety thresholds.

while it is not sufficiently modally robust for one's belief that one is φing to be knowledge?[18]

Valaris (2022) also argues for the Knowledge Thesis by arguing for an epistemic theory of control. He, like Pavese, suggests that both control and knowledge require one's behavior to be modally robust; and he is sensitive to the worry that these modal robustness conditions may not be the same in the required way. He asks: "still, why suppose that the possibilities of failure excluded by the control condition must include all the possibilities that need to be excluded to ensure safety?" (p. 742). He then offers an argument based on the practical nature of the knowledge involved in intentional action to the conclusion that the possibilities for failure which threaten safety also threaten control. But even if the possibilities for failure which threaten saftey also threaten control, intentional action does not require *complete* control (and presumably the same goes for knowledge and warrant/safety).[19] Thus, there are crucial (albeit perhaps vague and contextual) points at which too many threatening possibilities of failure rule out sufficient control for intentional action and analogously for warrant/safety and knowledge. And again we are left with the same unanswered question: why think that these points are always the same?

I consider one final way to develop an epistemic theory of control in terms of knowledge which might look to have more promise.[20] Consider the following precisification of PSKT: if one intentionally φs, then one knows for some means ψ to φ, that they are ψing in order to φ. Now, consider a theory of control according to which one controls one's φing to the degree that they have knowledge of the means to their φing; so the more means they have knowledge of, the more control they have over φing. Setting aside whether this is a plausible theory of control in intentional action, note that it runs into similar issues already raised when it is used to defend PSKT. Consider a case where for each means ψ to φing, the agent has (at best) a true, somewhat but not sufficiently warranted belief that one is ψing in order to φ. Given there is only a difference in degree (of warrant) between this case and another case where the agent has knowledge about means and so some degree of control, this case presumably involves some degree of control as well. The same explanatory challenge arises: why think this minimal degree of control could not possibly be enough for intentional action?[21] So the basic explanatory challenge I am pressing is robust against changes in the details of the epistemic theory of control.

There is a similar explanatory gap in another influential kind of argument for PSKT. This approach—which I will call

[18] There is a similar gap in Pavese 2021a. Pavese here argues that an intentional φing requires a true belief about how to φ that is modally robust because the belief explains the intentional success and satisfactory explanations must be modally proportional to their explanandum (Greco 2016). Still, we are not given reason to think that the level of modal robustness that the φing needs to be intentional matches the level of modal robustness that the explanatory belief needs to be knowledge. A similar gap remains in Pavese 2021c, Sec. 8.3 and Pavese and Beddor 2023 too.

[19] Valaris does not seem strongly committed to this, but he does assume that *any* threatening possibility of failure rules out both (sufficient) control and safety (see his footnotes 8 and 14). As he notes, this assumption is a matter of widespread debate in the case of safety. This assumption has not, as far as I know, been discussed in the case of intentional action, and I find it implausible (in part, for reasons discussed in the next section).

[20] This strategy was suggested by Carlotta Pavese in conversation.

[21] Relatedly, one might try to defend PSKT by claiming that (i) intentional action entails some non-zero degree of knowledge how and (ii) defeat of knowledge how and propositional knowledge are the same (see, e.g., Pavese 2021c; Kearl 2023). These premises work as a defense of PSKT only in the presence of the following implausible principle: if defeat conditions are such that all of one's beliefs about φing are somewhat warranted but insufficiently warranted for propositional knowledge, then defeat conditions are such that one has no degree of knowledge how. Now, an intellectualist about knowledge how might accept this principle. But in that case, the intellectualist needs to explain why some degree of knowledge how *per se* is always necessary for intentional action. Why couldn't these insufficiently warranted beliefs about φing—whatever you want to call them, if not a minimal degree of knowledge how—sometimes suffice?

the *Imperfective Approach*—emphasizes that the content of the knowledge that one has while intentionally φing involves an imperfective act-description: while intentionally φing, one knows that one is *φing* (not necessarily that one φed or will φ). And in light of the openness and broadness of the progressive (Galton 1984), one's representing oneself as φing is sufficient for it to be the case that one is (intentionally) φing; thus, one's representation constitutes sufficient warrant for itself.

There are a number of versions of the Imperfective Approach,[22] but any one worth considering will deny that representing oneself as φing is in all cases sufficient for it to be the case one is (intentionally) φing. If I represent myself as currently flying to the moon, that does not in itself imply that I am currently (intentionally) flying to the moon.[23] Such a sufficiency claim reasonably only holds when certain background conditions are in place, and so it is only *defeasibly* true that a representation that one is φing constitutes sufficient warrant for itself.

Falvey (2000) defends an approximate version of PSKT using a plausible version of the Imperfective Approach. He argues that in light of the openness and broadness of the progressive, one has an epistemic entitlement to assume that one is doing what one intends to be doing *unless one has reason to think one is not*. But as soon as we admit that one's representation constitutes sufficient warrant for itself only defeasibly, the explanatory challenge under discussion becomes applicable. There will be cases where the defeaters of the relevant entitlement—e.g., reasons to think that one is not doing what one intends—reach the strength at which they undermine one's entitlement, and so one fails to know that one is φing; but why think that in every one of these cases, one also fails to be intentionally φing?

My claim at this point is only that this question needs answering by the Anscombean, not that it has no answer. But to get the sense that there is something difficult to explain here, consider a kind of example one might raise against the Knowledge Thesis. You are attempting to solve a Rubik's cube for the first time within a specified time limit. Solving a Rubik's cube is difficult, and you have no reason to think you'd succeed if you attempted. It turns out, unbeknownst to you, that you're a natural—you solve the Rubik's cube in the allotted time. Moreover, while the solving of the Rubik's cube is a bit sloppy—your skills, while natural, are not developed—your attempt is on track to be successful throughout. Plausibly, all while solving the Rubik's cube, you are intentionally doing so; and you are not just intentionally *trying* to solve the Rubik's cube. Indeed, your success—not just your attempt—is explained by your (minimal) natural skill rather than luck, skill which you justifiably do not know you have; moreover, you'd be credited with the success—not just an attempt—and be praised for its impressiveness. But because you recognize the difficulty of solving a Rubik's cube, you are justifiably ignorant of your natural Rubik's cube skills, and you have no further reason to think you'll successfully solve the Rubik's cube in the allotted time, your entitlement to believe that you're in fact solving the Rubik's cube will, at least toward the beginning, be defeated.

More generally, I have not offered truth conditions for it to be the case that one is intentionally φing at t, but here is a

[22] For a notable example, see Thompson 2011.
[23] See, e.g., Paul 2020, Sec. 6.3b for more examples of this kind.

natural sufficient condition given the truth conditions I have given for it to be the case that one intentionally φs: at t, the agent represents themselves as φing and is on track to φ with a sufficient degree of (the right kind of) control over their φing. Then even if the defeaters of the entitlement are related to (potential) failures in control, the agent may still be on track at t to successfully φ with a sufficient degree of control if the level of control needed to intentionally φ is relatively low. The Imperfective Approach, to be a full defense of PSKT, needs to include an explanation of why this could not possibly be the case. In fact, it is telling that Falvey defends only the approximate necessity of knowledge to action, suggesting that he does not think the relevant explanatory challenge can be met.

In the next two sections, I give a positive argument in favor of the Separation Thesis. I argue that ethical facts shape the control threshold (Section 4) and that there is no similar influence from ethics on the warrant threshold (Section 5). In light of the disparate pressure from the ethical domain, we should be skeptical of any necessary connection between the two thresholds. I thereby show that it is no mistake that existing defenses of the Knowledge Thesis fail to meet the explanatory challenge under discussion—there is good reason to think that it cannot be met.

4. The Ethical Dimension of Intentional Action

Intentional action has a central place in our ethical thinking and practices. It is a central kind of behavior to which we ascribe ethical properties such as rightness and wrongness or rationality and irrationality, that may be the subject of reactive attitudes (Strawson 1962), and based on which we assess a person's character. As Bratman (1987) puts it, our concern with intentional action is "not limited to description and explanation of actions, but extends to the assessment of agents and to the appropriateness of reactive attitudes like indignation and resentment" (p. 124).

Others have isolated intentional action from necessary connections to certain states of mind in light of the place of intentional action in ethical thought and practice. Harman (1976) recognizes a separation between intentional action and belief stemming from the fact that when one has reason *not* to φ, one can intentionally φ without believing that one is φing.[24] In a similar spirit, Bratman (1987) recognizes a separation between intentional action and intention grounded in the disparate pressures affecting the extensions of both, the former influenced by pressure to classify behaviors as intentional actions insofar as the agent may be held morally responsible for them and the latter influenced by pressure to classify states as intentions insofar as they satisfy the characteristic roles in motivation and practical reasoning highlighted by Bratman's planning theory of intention.[25] Thus, both Harman and Bratman argue against a necessary connection between intentional action and some state of mind in part by appealing to the connection between intentional action and ethics.

[24] Velleman (2007) agrees with Harman here and thinks this is why what one does intentionally is not perfectly co-extensive with what is known without observation (see his Footnote 7); though he does still think that when one intentionally φs, one's intention to φ "probably" amounts to knowledge (p. 203).

[25] More precisely, Bratman argues that in intentionally φing, one need not intend to φ. He does tentatively accept what he calls the Single Phenomenon View on which intentionally φing necessarily involves an intention to do something. (To see why I say "tentatively", see pp. 126-127.)

I continue this tradition of thought in defending the Separation Thesis. In this section, I sketch an account of the ethical dimension of intentional action on which ethical facts fix the level of control required for intentional action.[26] I therefore claim that intentional action is not merely a phenomenon that plays an important role in our ethical lives—intentional action is, in a sense, intrinsically ethical. I give two arguments for this proposal, one which appeals to existing experimental literature on the folk concept of intentional action and another which appeals to the function(s) of the concept of intentional action. I conclude the section by pointing out that, for those unconvinced by the two arguments, an even weaker connection between intentional action and ethics would support the Separation Thesis as well.

The key idea underlying my proposal is that part of what makes a behavior an intentional action is that some sort of assessment of the agent on account of the behavior is or would be (in certain circumstances) appropriate. It is no surprise that it *follows* from the fact that a behavior is intentional that the agent is in some sense assessable for it, but I am suggesting that the connection to assessability is stronger: part of *what it is* for a behavior to be an intentional action is that the agent is in some sense assessable for it. The extension of the kind *intentional action* is fixed, in part, by ethical facts about assessability in the relevant sense.

Recall that in Section 2, I gave a schematic theory of intentional action: to φ intentionally just is to φ with a sufficient degree of the right kind of control over one's φing, where the right kind of control and sufficient control are theoretical placeholders. So I have now proposed two application conditions for the concept of intentional action. Both conditions can be met if the right kind of control is a kind that is *necessary and sufficient* for assessability in the relevant sense when exerted to a sufficient degree and a sufficient degree of that right kind of control for intentional action is *fixed by* the degree required for assessability in the relevant sense (call this the *assessability threshold*). This is precisely the view that I propose. The ethical dimension of intentional action on this proposal is therefore embodied in the control threshold: we cannot determine the degree of (the right kind of) control that is enough for intentional action—and thus what counts as an intentional action—without turning to the ethical facts and, in particular, to facts about what an agent is in the relevant sense assessable for.

I have been referring vaguely to *the relevant sense of* assessability. This is because determining the exact sense of assessability that undergirds the application of the concept of intentional action is a difficult task that I cannot undertake here. To determine it is, on my view, to fully solve what might be called "the threshold problem for intentional action" which is analogous to the threshold problem for knowledge (see Footnote 10). To give a flavor of the space of options though, consider the following proposals:

- (moral responsibility) intentional action is behavior for which we may be held morally responsible;
- (deontic assessability) intentional action is the kind of behavior which could be assessed with deontic terms, i.e., it is

[26] To repeat, throughout I am *not* using 'ethical' and 'moral' interchangably; by 'ethical facts', I mean normative facts about practical life broadly construed.

not a category mistake to deem the behavior something you ought or ought not have done.

- (character evaluation) intentional action is behavior which is most relevant for character evaluation;
- (creditability) intentional action is behavior which the agent gets credit for, where getting credit for a behavior is understood normatively;
- (priority evaluation) intentional action is behavior which is a prioritized candidate for ethical evaluation broadly construed.

The first option faces two immediate potential counterexamples: cases of moral luck (Nelkin 2021) and unintentional but negligent actions for which we hold one another responsible (Rosen 2003). However, the last four options are at least prima facie plausible. Rather than attend to the important task of figuring out exactly what sense of assessability is constitutively connected to intentional action—a task which I am pursuing in other work—I leave assessability as a theoretical placeholder. The core of my suggestion is that any complete theory of intentional action must provide an account of this relevant sense of assessability because facts about it shape the control threshold which shapes the boundaries of intentional action as a kind.

In fact, experimental work on the folk concept of intentional action goes some way toward supporting this proposal. There is a large literature about how the folk conception of intentional action is morally loaded. One important study for the argument here is summarized as follows. Nadelhoffer (2005) gave 40 undergraduate students three vignettes inspired by an example of Mele and Moser (1994). In Case #1 (C_1), a nuclear reactor is in danger of exploding and thereby killing thousands of people. Fred must type in a certain ten-digit number to shut it down. He has no idea what the code is but manages to guess the correct code, stop the reactor from exploding, and thereby save thousands of lives. Case #2 (C_2) is like C_1 except that Fred is angry about being fired from his job at the nuclear power plant and is inputting the ten-digit code in order to cause the reactor to explode. Again, he has no idea what the code is but manages to guess the correct code, cause a nuclear meltdown, and thereby kill thousands of people. In Case #3 (C_3), Fred is playing a lottery machine where he must type in a certain ten-digit code to win $1,000,000. Again, Fred has no idea what the code is but manages to guess the correct code and win $1,000,000. Among these three cases, far more participants said that Fred intentionally punched in the correct code in C_2 (67%) than in C_1 (38%) or C_3 (20%). Further, far more people said that Fred intentionally caused the explosion (83%) and intentionally prevented the explosion (73%) than intentionally won the lottery (33%). These results are suggestive that the control threshold is lowered where moral stakes are high, as in cases C_1 and C_2, compared to where moral stakes are low, as in case C_3. Nadelhoffer (2004) presents results from a similar experiment with 160 undergraduate students that confirm these findings.[27]

[27] See also Knobe 2003a, Knobe 2003b, and Knobe 2004.

Nadelhoffer claims that these results show that skill and control are not necessary for an action's being intentional (Nadelhoffer 2004, p. 283). However, I think something like Mele's (2001) conjecture is more likely to be true: while skill and control are still necessary for an action's being intentional, the degree of skill and control necessary for intentional action is the degree necessary for the agent to be assessable in the relevant sense for the behavior, where being assessable in the relevant sense is sensitive to moral stakes.[28] Note that while the kind of assessability intrinsic to intentional action may not be distinctively moral, whether the agent is assessable in the relevant (perhaps non-moral) sense may still be affected by moral stakes. For example, there may be more permissive conditions on a token behavior being creditable to the agent when it is embedded in a morally loaded context.

There are a number of concerns one might have about these and related experimental results, concerns which have generated a large literature. For example, one might reasonably wonder whether Nadelhoffer's findings are merely evidence of ethically sensitive norms on the *ascription* of the concept intentional action.[29] Or one might ask whether the respondents are distinguishing clearly enough between attributing intentionality and attributing an intention.[30] Additionally, there is a large body of continually growing work on how the folk understand the relationship between intentional action, knowledge, and the ethical (e.g., Pavese and Henne 2023). Rather than seriously engage with these important concerns or with the more general debate about the philosophical import of experimental work on ordinary language concepts—something that I do not think I could do sufficiently well in this context—I offer a second independent argument for my proposal that intentional action is intrinsically ethical.

In order to understand the features of intentional action and its connection to the ethical, we might look to the *function* of the concept of intentional action, that is, the kinds of roles that the concept plays in our lives. The concept of intentional action, I claim, has two functions as a representational tool, one (plausibly) non-ethical and one ethical. It has the function of tracking distinctive kinds of controlled behavior and the function of tracking behavior that one is assessable for in a certain sense. This dual function explains why the control threshold and the assessability threshold coincide: if not, both tracking functions could not be fulfilled by the single representational device.

Why think that the concept of intentional action has this ethical function? Well, it is no surprise that *some* concepts of ours would function so as to flag forms of agency which are requiring of our ethical attention in various ways. We would benefit greatly in our ethical practices—from our practice of moral responsibility to our practice of rationality assessment—if we had communal methods for delineating and referring to the phenomena (most) relevant to those practices. In particular, we are

[28] For an opinionated overview of other responses to the presented and related empirical work, see Cova 2016.

[29] As an initial response, the experimental results confirming the relevance of moral considerations to folk ascriptions of the concept of intentional action are robust across many different settings (Cova 2016).

[30] But Knobe (2004) shows that even when respondents are able to ascribe both intention and intentionality, the latter is far more influenced by the moral stakes (see also Knobe and Burra 2006).

acquainted with a vast number of behaviors at any given moment, and so it would be cognitively efficient to have concepts for filtering ethical attention toward those behaviors which are requiring (or perhaps *most* requiring) of our attention. Given the clear centrality of intentional action to our ethical thinking and practices, it would also not be surprising if the concept of intentional action is one such concept. While only a sketch, this seems a promising form of argument in support of the proposed ethical dimension of intentional action.[31]

Finally and importantly, while I see good reason to think that the connection between intentional action and the ethical is an intrinsic one, I recognize that this is an ambitious claim. Thus, it is crucial to note that the Separation Thesis could be supported by a weaker connection. Indeed, the control threshold need not be *fixed* by considerations from the ethical domain; the control threshold need only be *influenced* by ethical considerations, where this influence is not identical to the influence exerted on the warrant threshold. So even if there is only a defeasible connection between intentional action and assessability (in the relevant sense), insofar as the connection is not identical to that between knowledge and assessability, there is support for the Separation Thesis. And again, given the role that intentional action plays in ethical thought and practice—a role which is in many ways deeper and clearer than that played by knowledge—we should not be surprised if this were so.

To put the last point slightly differently, moral and pragmatic encroachment on epistemic concepts are taken very seriously, even if controversial. Intentional action, on the other hand, is a quintessential *practical* concept, and so for there to be influence from ethics (understood to encompass all of practical normativity) on the control threshold should be taken even more seriously. In fact, such influence should not be seen as a form of encroachment at all—ethical facts seem to be the natural and expected kind of facts to influence the level of control necessary for intentional action.

5. The Encroachment Objection

In the previous section I suggested that the control threshold is fixed by the assessability threshold. I suggested that this grounding relation is entailed by a plausible picture of the function of the concept of intentional action in our ethical conceptual scheme and could partly explain why actions that are largely a matter of luck yet morally significant are classified by the folk as intentional. But now I must face an important objection: might the warrant threshold also track the assessability threshold? More precisely, might the level of warrant required for knowledge about what one is doing just be the level of warrant required for one's behavior to be assessable in the relevant sense? If so, then I have no argument for the Separation Thesis, as I have failed to show that ethical considerations *differentially* influence the two thresholds. In fact, a natural first thought is that pragmatic and/or moral encroachment might establish precisely this analogous influence. Call this the *Encroachment Objection*. I begin by arguing that directly applying existing views on pragmatic and moral encroachment to defend the

[31] I am therefore suggesting that we will come to see the intrinsically ethical nature of intentional action if we study the concept of intentional action in a manner broadly similar to Craig's (1990) functional study of the concept of knowledge. I am currently pursuing this suggestion in other work (Kelley msb).

Encroachment Objection likely fails.

As I will understand it, pragmatic encroachment is the view that a difference in pragmatic circumstances can constitute a difference in knowledge (Ichikawa and Steup 2018, Sec. 12). For our purposes here, I assume a more specific version:

> **Pragmatic Encroachment**: a difference in pragmatic circumstances can constitute a difference in knowledge through a change in the warrant threshold.

A common motivation for Pragmatic Encroachment is the following link, which Bolinger (2020) calls the *Knowledge-Action Link*: if an agent knows that p, then she is rationally permitted to deploy p as a premise in practical deliberation.

A first point against the Encroachment Objection is that insofar as pragmatic considerations affect the warrant threshold in line with the Knowledge-Action Link, it looks like they will sometimes affect the warrant threshold in the opposite direction than they affect the control threshold. Indeed, consider Fred in C_1 who is guessing a ten-digit code in an attempt to save thousands of people from a nuclear explosion. We saw that compared to the morally neutral case in C_3, more people thought that Fred intentionally entered the correct code and many people thought Fred intentionally prevented the explosion. Assuming there is an actual lowering in the assessability threshold in light of the heightened moral stakes, the folk intuitions correctly notice a lowering in the control threshold. But did Fred *know* that he was preventing the explosion or that he was correctly entering the code while doing so? The Pragmatic Encroachment defender asks: was Fred rationally permitted to assume that he was preventing the explosion or entering the correct code as a premise in practical deliberation while typing in the code? It looks like certainly not. In fact, the practical cost of being wrong is so high in C_1 (as compared to C_3) that Fred better have especially strong warrant for his belief to rationally practically reason assuming he is right. (I am not thereby suggesting that Fred ought to practically reason assuming he is *wrong*, however, for presumably this would lead him to do nothing which is also costly.) For example, if someone offers him an alternative method of stopping the nuclear reactor from exploding while he is typing in the ten-digit code, surely he shouldn't reason about this offer on the assumption that he's entering in the correct code. Thus, Pragmatic Encroachment implies that the warrant threshold *increases* in C_1 as compared to the low stakes version C_3, while the control threshold (if anything) *decreases* in C_1 compared to the low stakes version C_3.

Moral encroachment as it is usually understood is also insufficient to defend the Encroachment Objection.

> **Moral Encroachment:** Holding fixed the degree of probabilistic support the evidence gives for p, whether an epistemic attitude about p has some positive epistemic status can depend importantly on its moral features.

A standard case thought to support moral encroachment involves forming a belief about a person on the basis of their race in a way that reinforces racial stereotypes, even if the person's race makes the belief very likely to be true.[32] We hesitate

[32] See Gendler 2011 for a specific example.

to call such a stereotypical belief knowledge, suggesting that the standard for knowledge depends on the moral credentials of the belief. However, Moral Encroachment does not establish the Encroachment Objection, for many cases where moral considerations affect the control threshold are not cases where *believing* that one is φing is of any moral significance—it is the φing itself which is of moral significance. For example, there is nothing morally significant about Fred believing or failing to believe that he is successfully causing or preventing the explosion in typing the ten-digit number; it is rather the actual causing or preventing the explosion that is of moral significance.[33]

So given how pragmatic and moral encroachment are commonly understood and defended, it is unlikely that either could be straightforwardly wielded into an objection to the Separation Thesis. Still, there may be a plausible novel account of encroachment that is tailored to establish the Encroachment Objection. It is to this suggestion that I now turn.

First, we might first look to why the control threshold is identified with the assessability threshold. As discussed in Section 4, this came from the fact that intentional action as a concept, along with having the function of tracking certain kinds of control, has the function of tracking behavior which agents are assessable for in a certain sense. Notice that an analogous story cannot be told about knowledge: we know about many of our behaviors that we are not in *any* substantive sense assessable for, e.g., our stomach rumbles and leg spasms. Taking inspiration from Anscombe, one might object that the knowledge I have of my stomach rumbles and leg spasms is of a different kind than the knowledge I have of my intentional actions. The former is *speculative* (propositional) knowledge and the latter, which functions to track knowledge about behaviors for which we are assessable in the relevant sense, is *practical* (propositional) knowledge. However, it is hard to see any independent motivation for thinking that we operate with two concepts of (propositional) knowledge and that one of them functions to track knowledge with contents about assessable behavior. Importantly, a connection to assessability was not part of Anscombe's own account of practical knowledge.

While it seems unlikely that any concept has the function of tracking knowledge about behavior that agents are assessable for in the relevant sense, it is nonetheless clear that knowledge is intimately linked to assessable behavior. Indeed, one might think that while we know about many behaviors beyond what we are assessable for, we are (at least defeasibly) assessable only for the behaviors that we know about.[34] Call this the *Knowledge-Assessability Link*.[35] If the Knowledge-Assessability Link is true, then (for beliefs about behaviors) the warrant threshold is met whenever the assessability threshold is met which

[33] Finally, much of the discussion of moral and pragmatic encroachment has focused on how moral and pragmatic factors make it more *difficult* for a belief to be knowledge.

[34] This defeasibility clause is both plausible and potentially problematic. It is plausible in light of behaviors we do not know we are performing when we should not have been performing those behaviors without such knowledge; for example, shooting a prisoner (who is in fact trying to escape) on a baseless hunch that they are trying to escape (Hawthorne and Stanley 2008). In this case, you are (in many senses) assessable for shooting an escaping prisoner but do not have knowledge that you are. The defeasibility clause is potentially problematic, because it introduces a gap between knowledge and intentional action.

[35] There is some empirical support for that ascriptions of knowledge about what one is doing are sensitive to the moral features of the behavior (see, e.g., Beebe and Buckwalter 2010). It is perfectly consistent with the Separation Thesis that the warrant threshold is dependent on moral features of the behavior in question; what is inconsistent with the Separation Thesis is that the warrant threshold is dependent on moral considerations *in the precise way* that the Knowledge-Assessability Link purports.

implies, with the fact that the control threshold is identical to the assessability threshold, that the warrant threshold is met whenever the control threshold is met. This necessary connection between the warrant and control thresholds suffices to support the Knowledge Thesis.

There is considerable plausibility to the Knowledge-Assessability Link. We do not usually normatively assess or hold agents responsible for behaviors that are cognitively inaccessible to them. If you thought that your friend did not know that they were hurting your feelings (and perhaps some other *ceteris parabus* conditions held, e.g., there was no reason that they should have known that they were), you would not take them to be assessable for hurting your feelings in the way that is constitutive of intentional action. One is also similarly not assessable in this sense for the behaviors that one performs while unconscious or asleep (again, assuming *ceteris parabus* conditions hold). However, I claim that the Knowledge-Assessability Link is false, though its apparent plausibility stems from a similar and true principle. Indeed, it is a result of confusing the necessity of knowledge for assessability with the necessity of some cognitive representation or other, e.g., (insufficiently warranted) belief or intention.

To see why, imagine that your friend knows that they are hurting your feelings, intends to be hurting your feelings, and is aware that they are hurting your feelings and these representations are explanatorily connected to their actually doing so. Now compare this to a case where the piece of knowledge is replaced by a true, almost sufficiently warranted belief. Perhaps in this second case, your friend has lost some reason to think they've picked an effective insult, though their belief that they have remains unchanged. It seems plausible that your friend is assessable in the sense that undergirds intentional action in both cases, despite lacking knowledge in the second case. I conjecture that this is because while knowledge is sometimes sensitive to small differences in the quality of epistemic warrant, assessability is not correspondingly sensitive in many of these instances; and the latter is because often, what is of primary concern for assessability is what one *takes oneself* to be doing, rather than whether one had epistemically good reasons for doing so. It seems that one does not need to know per se *anything* about one's φing to be asssessable for it for similar reasons—beliefs will often suffice. For example, it does not seem that one even needs to know per se that one is *trying* to φ to be assessable for φing—it is often enough that one believes that one is trying to φ (I assume that it is possible to believe without knowing that you are trying to φ). So one could not use the Knowledge-Assessability Link restricted to trying behaviors to argue that knowing that one is trying to φ is necessary to be intentionally φing. In this way, the argument here extends to precisifications of PSKT beyond Anscombe's original formulation.

Without a more substantive account of assessability on the table, it is difficult to argue conclusively that assessability sometimes depends primarily on one's belief-like attitudes, rather than knowledge *per se*. But I will note that assessability in the five senses mentioned in Section 4 are often insensitive to small differences in epistemic warrant. For example, the appro-

priateness of holding your friend morally responsible for their behavior, the appropriateness of making deontic or evaluative judgments about their behavior, the appropriateness of evaluating your friend's character on the basis of the behavior, the crediting of the behavior to your friend, and whether the behavior is a prioritized candidate for ethical evaluation do not seem to differ in the two cases. Moreover, it is clear that the reason why we are not assessable in the relevant sense for what we do while asleep or unconscious is not that we do not know about these behaviors *per se* but that we do not cognitively represent these behaviors at all. I conjecture, therefore, that the motivation for the Knowledge-Assessability Link really only establishes the weaker thesis that substantive cognitive access explanatorily connected to one's behavior is necessary for assessability in the relevant sense.

6. Conclusion

If the Argument from Distinct Thresholds holds up, then we have good reason to think that no precisification of the Propositional Schematic Knowledge Thesis is true. This is because we have good reason to think the thresholds governing the application of the concepts of knowledge and intentional action do not always line up as required. I further conjectured that an argument of a similar form can be employed to undermine any precisification of the Schematic Knowledge Thesis more generally. Nonetheless, even if acting intentionally does not *necessarily* involve knowledge, the special way we know what we do intentionally when we do know is a crucial domain of inquiry for understanding human agency. The discussion here does not at all undermine the importance of the epistemology of action. It only implies that understanding the nature of intentional action requires going beyond epistemology into ethics in order to grapple with the ethical function of the concept of intentional action.

References

G. E. M. Anscombe. *Intention*. Harvard University Press, 1957.

Bob Beddor and Carlotta Pavese. Practical Knowledge without Luminosity. *Mind*, 10 2021.

James R. Beebe and Wesley Buckwalter. The epistemic side-effect effect. *Mind and Language*, 25(4):474–498, 2010.

Renée Jorgensen Bolinger. Varieties of moral encroachment. *Philosophical Perspectives*, 34(1):5–26, 2020.

Michael Bratman. *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information, 1987.

J. Adam Carter and Joshua Shepherd. Intentional action and knowledge-centered theories of control. *Philosophical Studies*, 180(3):957–977, 2023.

Florian Cova. The folk concept of intentional action: Empirical approaches. In Wesley Buckwalter and Justin Sytsma, editors, *Blackwell Companion to Experimental Philosophy*. 2016.

Edward Craig. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press, 1990.

Donald Davidson. Agency. In *Essays on Actions and Events*. Clarendon Press, 1980.

Keith S. Donnellan. Knowing what I am doing. *Journal of Philosophy*, 60(14):401–409, 1963.

Kevin Falvey. Knowledge in intention. *Philosophical Studies*, 99(1):21–44, 2000.

Antony Galton. *The Logic of Aspect: An Axiomatic Approach*. Clarendon Press, 1984.

Tamar Szabo Gendler. On the epistemic costs of implicit bias. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 156(1):33–63, 2011.

Daniel Greco. Safety, explanation, iteration. *Philosophical Issues*, 26(1):187–208, 2016.

H. P. Grice. Intention and uncertainty. *Proceedings of the British Academy*, 57:263–279, 1971.

Stuart Hampshire. Thought and action. *Philosophy*, 36(137):231–233, 1959.

Gilbert Harman. Practical reasoning. *The Review of Metaphysics*, 29(3):431–463, 1976.

John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008.

Stephen Hetherington. Knowledge's boundary problem. *Synthese*, 150(1):41–56, 2006.

David Hunter. Davidson on practical knowledge. *Journal for the History of Analytical Philosophy*, 3:1–19, 2015.

Jonathan Jenkins Ichikawa and Matthias Steup. The Analysis of Knowledge. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2018 edition, 2018.

O. R. Jones. Things known without observation. *Proceedings of the Aristotelian Society*, 61:129–150, 1960.

Timothy Kearl. What we know when we act. *Philosophical Studies*, pages 1–19, forthcoming.

Timothy R. Kearl. Knowledge-how and the limits of defeat. *Synthese*, 202(2):1–22, 2023.

Mikayla Kelley. A control theory of action. msa.

Mikayla Kelley. The normative function of intentional action. msb.

Joshua Knobe. Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16(2): 309–325, 2003a.

Joshua Knobe. Intentional action and side effects in ordinary language. *Analysis*, 63(279):190–194, 2003b.

Joshua Knobe. Intention, intentional action and moral considerations. *Analysis*, 64(2):181–187, 2004.

Joshua Knobe and Arudra Burra. The folk concepts of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture*, 6(1-2):113–132, 2006.

Rae Langton. Intention as faith. In H. Steward and J. Hyman, editors, *Agency and Action*, pages 243–258. Cambridge University Press Press, 2004.

Yair Levy. Intentional action first. *Australasian Journal of Philosophy*, 91(4):705–718, 2013.

Alfred R. Mele. Acting intentionally: Probing folk notions. In Bertram Malle, L. J. Moses, and Dare Baldwin, editors, *Intentions and Intentionality: Foundations of Social Cognition*, pages 27–43. MIT Press, 2001.

Alfred R. Mele and Paul K. Moser. Intentional action. *Noûs*, 28(1):39–68, 1994.

Thomas Nadelhoffer. The butler problem revisited. *Analysis*, 64(3):277–284, 2004.

Thomas Nadelhoffer. Skill, luck, control, and intentional action. *Philosophical Psychology*, 18(3):341–352, 2005.

Dana K. Nelkin. Moral Luck. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition, 2021.

Anne Newstead. Knowledge by intention? on the possibility of agent's knowledge. In Stephen Hetherington, editor, *Aspects of Knowing*, page 183. Elsevier Science, 2006.

Lucy F. O'Brien. On knowing one's own actions. In Johannes Roessler and Naomi M. Eilan, editors, *Agency and Self-Awareness*. Clarendon Press, 2003.

Sarah K. Paul. How we know what we're doing. *Philosophers' Imprint*, 9:1–24, 2009a.

Sarah K. Paul. Intention, belief, and wishful thinking: Setiya on "practical knowledge". *Ethics*, 119(3):546–557, 2009b.

Sarah K. Paul. *The Philosophy of Action: A Contemporary Introduction*. Routledge, 2020.

Carlotta Pavese. Probabilistic Knowledge in Action. *Analysis*, 80(2):342–356, 05 2020.

Carlotta Pavese. Know-how, action, and luck. *Synthese*, 198(Suppl 7):1595–1617, 2021a.

Carlotta Pavese. Knowledge and mentality. *Wiley: Philosophical Perspectives*, 35(1):359–382, 2021b.

Carlotta Pavese. Knowledge, action, defeasibility. In Jessica Brown and Mona Simion, editors, *Reasons, Justification, and Defeaters*. Oxford: Oxford University Press, 2021c.

Carlotta Pavese and Bob Beddor. Skills as knowledge. *Australasian Journal of Philosophy*, 101(3):609–624, 2023. doi: 10.1080/00048402.2022.2056753.

Carlotta Pavese and Paul Henne. The know-how solution to kraemer's puzzle, 2023.

Juan Piñeros Glasscock. Practical knowledge and luminosity. *Mind*, 129(516):1237–1267, 2019.

Sebastian Rödl. Forms of practical knowledge and their unity. In Anton Ford, Jennifer Hornsby, and Frederick Stoutland, editors, *Essays on Anscombe's Intention*. Harvard University Press, 2011.

Gideon Rosen. Culpability and ignorance. *Proceedings of the Aristotelian Society*, 103:61–84, 2003.

John Schwenkler. *Anscombe's Intention: A Guide*. Oxford University Press, 2019.

Kieran Setiya. Practical knowledge. In *Practical Knowledge: Selected Essays*. Oxford University Press, 2016a.

Kieran Setiya. Practical knowledge revisted. In *Practical Knowledge: Selected Essays*. Oxford University Press, 2016b.

Joshua Shepherd. *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford: Oxford University Press, 2021.

Joshua Shepherd and J. Adam Carter. Knowledge, practical knowledge, and intentional action. *Ergo: An Open Access Journal of Philosophy*, 9:556–583, 2023.

Will Small. Practical knowledge and the structure of action. In Günter Abel and James Conant, editors, *Rethinking Epistemology Volume 2*, pages 133–227. Berlin, Germany: De Gruyter, 2012.

Alexander Stathopoulos. Knowing achievements. *Philosophy*, 91(3):361–374, 2016.

Peter Strawson. Freedom and resentment. *Proceedings of the British Academy*, 48:187–211, 1962.

Michael Thompson. Anscombe's intention and practical knowledge. In Anton Ford, Jennifer Hornsby, and Frederick Stoutland, editors, *Essays on Anscombe's Intention*. Harvard University Press, 2011.

Markos Valaris. Knowledge out of control. *Philosophical Quarterly*, 72(3):733–753, 2022.

David Velleman. *Practical Reflection*. Princeton University Press, 1989.

David Velleman. What good is a will? In Anton Leist and Holger Baumann, editors, *Action in Context*. de Gruyter/Mouton, 2007.