Episteme (2025), 1–17
doi:10.1017/epi.2025.5

episteme

ARTICLE

Knowing Your Commitments in Action

Merve Rumeysa Tapınç

Philosophy Department, Boston University, Boston, MA, USA
Email: mervertapinc@gmail.com

(Received 9 April 2024; revised 13 October 2024; accepted 3 January 2025)

Abstract

An interesting class of intentions is *commitments*: diachronic intentions that are especially representative of human agency. I argue that the justification conditions for knowing our commitments differ from those for knowing ordinary intentions, and I propose an externalist view according to which knowing one's own commitments is much like knowing those of others. I discuss Sarah Paul's transparency view, according to which, we know our intentions by making a conscious decision, even when we do not follow through on them. This is because, as she argues, the function of decisions and intentions is similar, and changing one's mind in the face of weakness of will does not defeat or undermine that knowledge of intention acquired through a conscious decision. In contrast, I show how the same evidence from weakness of will undermines or outright proves a lack of commitment, as commitments require more than decisions; they demand consistent patterns of action to resist temptations and follow-through over time. Therefore, a conscious decision is not a reliable way to know our commitments in the same way it is for ordinary intentions. I conclude that we cannot justify our commitments merely by referring to a mental item, without also referring to our regular patterns of action.

Keywords: Epistemology of commitments; intentions; Sarah Paul; the transparency view; weakness of will; externalism

1. Introduction

Many of us make commitments that we know we might someday break: we commit to writing a book, we commit to a relationship, or we commit to a job. How do we know we are *truly* committed to those pursuits? How can we tell the difference between a genuine commitment which we are capable of sustaining, and a mere intention from which we might deviate at a whim? In this paper, I argue that we cannot do so merely by knowing our immediate mental states: we must also evaluate our regular patterns of action. As the fallible agents we are, "intending" and "acting in accordance with one's intentions" are not same thing, and the gap between the two grows when it comes to commitments, especially as the realization of one's intended goal becomes more complex and temporally spread. As this gap grows, so does the epistemic gap between what we *believe* we commit to do and what we *know* we are committed to do.

Let us consider an example. Suppose Zoe consciously decides to be vegan but fails to put in adequate long-term effort. She knows that if she regularly prepared a vegan lunch,

she wouldn't go to the closest restaurant to eat her favorite steak burger. Yet she frequently procrastinates preparing her lunch and changes her mind whenever she wants to hang out with her nonvegan friends, something that happens very often. After enough lapses, she comes to question whether she made a genuine commitment at all. Did she simply want to signal virtue to her friends and colleagues? When there is a breakdown between what one decides to be committed to, and what one does, not only do all relevant goal-directed actions remain an open possibility for self-interpretation of our desires, but such breakdowns can weaken and threaten the existence of a true commitment. That is the epistemic gap I wish to explore.

This paper addresses the following question: what does it take to know one's commitments successfully? This question matters greatly for self-knowledge, because the most important kinds of self-knowledge do not focus merely on one's momentary synchronic intentions (e.g. to drink water, to stretch one's limbs, to go for a run, etc.), but on the strength of one's *diachronic* aims: to write a book, to run a marathon, or to do well at one's new job. That is, some of the most important forms of self-knowledge center on knowing one's *commitments*.

Accordingly, I will use the term "ordinary intention" to refer to rather momentary (and not necessarily explicit) mental states directed at basic actions (e.g. to stretch one's limbs).[1] I will use the term "decision" to refer to discrete conscious mental events, following Paul (2012), which I will explain further in the next section. On my view, a commitment to Φ is a decision to Φ which is accompanied by the understanding that we are liable to be tempted not to Φ, and, thus, accompanied by the goal of doing what it takes (developing skills, disciplining ourselves, etc.) to resist those temptations.

Current theories of self-knowledge are ill equipped to explain how we know our commitments, in part because the literature on self-knowledge has not made a clear distinction between the epistemology of commitments and that of intentions aimed at ordinary actions. We may mistakenly assume that knowledge of the former closely resembles that of the latter. In this paper, I clarify why the prevailing view on the epistemology of intention does not apply to commitments. To this end, I will consider two mutually exclusive ways of understanding how we know our commitments:

(1) knowledge of one's commitments does not, *like ordinary intentions*, require an understanding of the actions one takes to fulfill them

or,

(2) knowledge of one's commitments does, *unlike ordinary intentions*, require an understanding of the actions one takes to fulfill them.

My strategy is to defend the (2) view by critically discussing the (1) view. To illustrate the view (1), I will focus on Sarah Paul's (2012) account of the epistemology of intentions which she developed by revising the transparency theory of intention. Transparency theorists argue that questions regarding self-knowledge of certain mental states are transparent to questions about non-mental topics.[2] We acquire knowledge of our mental states by "looking outward" at the facts of the world, rather than by observing our own

[1]I do not take a specific stance on the nature of basic actions, which are famously characterized by Hornsby (1980) as actions we simply do without doing anything else to bring them about. I will follow the common practice of referring to both ordinary short-term actions and complex long-term actions, which require commitment to accomplish.

[2]The main mover of this view is Evans (1982).

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

minds or actions. Although transparency theorists differ in how they explain the transparency of intentions and in addressing the question of how deliberating on an action can generate self-knowledge, the key feature of this view is that the first person is epistemically authoritative in knowing her own present intentions.[3] This is because each of us is in a unique position to self-ascribe our intentions through deliberation, without needing to observe our own mind or actions.

Paul (2012, 2015) denies that this world-to-mind model can be generalized to intentions. She argues that we do not know our intentions by reflecting outward or evaluating reasons for action. Instead, Paul argues that we acquire knowledge of our intentions through a distinct capacity: the ability to make decisions, which provides reliable knowledge of what we intend to do. As she puts: "Our intentions are transparent, but they are transparent to something mental: conscious events of deciding what to do" (Paul 2015: 1532). The capacity to decide enables us to acquire knowledge of our intentions without needing to observe or collect evidence about our actions. As a result, on Paul's view, the subject is in a better position to know her intentions than anyone else, and this knowledge is not based on observed data as it is with third-person ascriptions of mental states (Paul 2012: 328). Thus, although Paul decidedly rejects the idea that intentions are transparent to what we have good reason to do or some worldly, non-mental fact in the same way beliefs are, her account retains the fundamental feature of the transparency view. On her view, self-knowledge of our intentions is privileged because we each are in an epistemically authoritative position to know our intentions by exercising our capacity to decide.

Paul does not distinguish between commitments and intentions, but her view on the epistemology of intention could also apply to commitments, which are essentially a type of intention. Whereas Paul insightfully argues that the transparency view fails to explain how we know our beliefs because breakdowns between what we judge and what we believe are quite common, she argues that those breakdowns do not lead to an epistemic gap in the case of intention as they do in the case of beliefs. This is because, as she claims, by making a conscious decision, we can reliably self-ascribe the relevant intention, since, as I will clarify in her argument, decisions and intentions – unlike judgments and beliefs – are functionally similar (Paul 2012: 332–35). Further, crucially, intentions do not require a "follow-up" and are compatible with changes of mind in the face of temptation, boredom, or in the presence of better options (342). In contrast, I argue that there is a class of intentions where breakdowns are not only common, but where changes of mind or weaknesses of will act as evidence against their existence. I do not deny that Paul's view might work well in ordinary intentions, but I deny that commitments are the kinds of intentions that one comes to know transparently in the way Paul proposes – namely, by referring only to a mental item as sufficient evidence, without needing to see

[3]Some transparency theorists argue that just as we self-attribute beliefs by focusing on what is true, we self-attribute intentions by deliberating on our actions, not by observing our own minds. In both cases, our minds are transparent to the object of these states. For Moran, the deliberative question of what to do is answered by reflecting on what is "choiceworthy," or what is "worth pursuing" (Moran 2001: 126). Boyle proposes "Moran-inspired reflectivist view," according to which self-knowledge is considered privileged because the agent constitutes her mental states through rational assessment of the matter in question (Boyle 2011: 237). Byrne (2011) rejects Moran's normative proposal and argues that intention to Φ is transparent to and inferred from non-observational knowledge that one will Φ. Setiya (2011), while also rejecting the normative approach, argues that our knowledge of intentions is non-inferential. Transparency theorists' views on intention are influenced by Anscombe (1963). See also Paul (2014) for a helpful discussion of how transparency theorists differ in their understanding of the transparency of intentions.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

what one does to accomplish the intended goal. Moreover, the cases where Paul's transparency view proves inadequate are more crucial for the purposes of self-knowledge.[4] We are interested in knowing whether one has a commitment to writing a book, pursuing a PhD, maintaining a certain diet, mastering a musical instrument, and so forth. Therefore, I suggest an alternative externalist picture of commitments in favor of this view:

(2) knowledge of one's commitments does, *unlike ordinary intentions*, require an understanding of the actions one takes to fulfill them.

For related reasons, I also reject "inner-sense" models, according to which we directly observe (Chalmers 2003; Locke 1689/1975), or utilize an introspective faculty to know our mental states (Armstrong 1981; Lycan 1996). Although I do not present a complete and finished account, I propose that some form of non-introspective, externalist view will best explain how we know our commitments, according to which we must use evidence about how we tend to act.

It is important to clarify from the outset that the externalist view need not be "behaviorist" or even "dispositionalist." According to dispositionalists (following Gilbert Ryle 1949), we know ourselves through observing what we do, how we reason, and how we act, much like how we infer the attitudes of others (Ryle 1949; Schwitzgebel 2002, 2012; Cassam 2015). On an extreme reading of this view, any case of failing to act in accordance with one's expected patterns would be a result of one's dispositions. However, it should be possible for one's actions to vindicate the existence of a sincere commitment. Although I argue that knowledge of commitments requires external evidence from how we act, it is the agent who makes a series of decisions and who resists to temptations to maintain the commitment. Knowledge of commitments is not simply something we find out by gathering data about ourselves, it is something we do.

My view shares features with Richard Holton's (2011) account of "resolutions," though my starting point is epistemological.[5] In Holton's view, resolutions involve resistance to temptation and weakness of will. My goal is rather to show that our actions are a significant source of evidence for the existence (and strength) of our commitments, and that this surprising epistemic fact reveals significant insights into the nature of commitment.

I will proceed as follows. In Section 2, I will critically discuss Paul's argument that intentions are transparent to our minds and that we know them by making conscious decisions. In Section 3.1, I will identify the key characteristics that distinguish commitments from ordinary intentions. In Section 3.2, I will consider an objection from weakness of will and argue why yielding to weakness can undermine or even invalidate a commitment. In Section 4.1, I will propose an externalist view, arguing that our knowledge of commitments goes beyond mere introspection and requires external evidence. In Section 4.2, I will show why we do not have epistemically authoritative access to our commitments and offer insights into justifying our commitments to ourselves and others.

[4]Cassam (2015) raises this criticism of the transparency view in the context of beliefs.

[5]Yet, I should note that my view is not merely an epistemic view. Compare this to Alex Byrne's (2018) view, which is an epistemological account of how we know our mental states. He does not address the nature of these mental states or what the epistemology reveals about them.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

2. Sarah Paul on conscious decisions and knowledge of intentions

Let me begin by discussing why Paul claims that while breakdowns between our judgments and beliefs are sufficient to challenge the transparency view of beliefs, such breakdowns do not pose an epistemic problem for knowledge of our intentions.

Paul argues that judgments "aim at the truth," but reflection on one's judgments (and on the reasons for one's belief) does not substantially influence what one actually believes, since the nature of judgments and beliefs differs significantly. Although we merely consider what is true in making judgments, believing one's judgment requires dispositions and tendencies. Therefore, breakdowns between what we judge and what we actually believe are quite common. As she compellingly points out:

> We are all familiar with the ways that conscious judgments can be radically disconnected with broader dispositions of thought and action: people judge that "he's never going to leave her," but wait for decades for it to happen; they judge that women and men have innately equal musical talent, but hire a much higher proportion of male candidates for their orchestra than they would if auditions were blind; they judge that tipping well is the right thing to do, but never tip well; they judge that heaven is far superior to this terrestrial existence, and that it is easier for a camel to pass through the eye of a needle than for a rich man to get into heaven, but do everything they can to avoid departing for the afterlife and to amass a fortune in this one (Paul 2012: 341).

Yet, Paul argues that the analogous account of the transparency of intention does not face this problem because, while we use judgment to evaluate reasons and might not have the relevant dispositions to actually believe what we judge, we can still know our intentions, even when making bad or akratic decisions. This is because whereas judging that p and believing that p require two separate capacities, both the intention and the decision require the same functional capacity – not two separate ones. The difference between an intention and a decision is that, in her view, an intention might be formed without an explicit decision. Decisions, in contrast, are conscious mental events (rather than a mental state) that we bring about by using our capacity to decide (Paul 2012: 336). Still, both decisions and intentions serve to direct our agency and activate standing dispositions, leading to action without further deliberation. We can self-ascribe an intention by deciding to act among the various options available, rather than by endlessly evaluating what is worth doing, and this is evident in the absence of a single norm to which decisions can be transparent (334). Whereas judgments are typically guided by the singular norm of "truth," which we use to close deliberation after assessing reasons for whether P, it is the capacity to decide that allows us to choose between equally good yet incommensurable actions, opt for the worse option, or decide based on our strongest desires or inclinations. In all cases, we bring potentially endless deliberation to a close through the act of deciding, rather than by continually weighing practical reasons for what to do. This does not imply that we simply disregard normative reasons when making decisions; rather, her argument is meant to conclude that the capacity to decide is a distinct capacity from one's capacity to assess reasons. So, whereas we cannot reliably self-ascribe the belief when we judge, we can reliably self-ascribe the intention when we decide to Φ.

Paul's account significantly contributes to the transparency view, as she demonstrates how to retain its strengths while rejecting the claim that intentions are transparent to some worldly, non-mental fact in the way beliefs are. Her conception of a *decision-based*

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

self-ascription of intention[6] shares the core features of transparency views, as she shows that we have epistemic authority to self-ascribe propositional attitudes by making conscious decisions and deny that we know our minds by observing external evidence about ourselves (Paul 2012: 336). Our conscious decisions are "epistemically authoritative" because they provide a significant source of evidence that is only available to the agent (Paul 2015: 1546). She concedes that a conscious decision is not the *only* way to know one's intentions, and that it is not a "failsafe way" (Paul 2012: 337). We can have intentions formed without an explicit decision or self-ascribe an intention after considering that we engage in complex means-ends reasoning and bring about goal-directed actions. Yet, she claims that in most cases, decision-based self-ascriptions are sufficiently reliable to count as knowledge.[7] As she puts: "... the mental act of deciding to Φ is a paradigmatic way of entering into the state of intending to Φ" (336). The subject has a direct, authoritative access to her intention via conscious event of deciding what to do.[8]

Paul acknowledges that there can be instances where a conscious decision to Φ does not align with a diachronic intention to Φ, yet she maintains that such instances are relatively uncommon. She gives a breakdown example from *Anna Karenina*, who decides to break up with Vronsky after considering all the bad consequences of seeing him again. However, when Anna has a second chance to see him, she reconsiders and decides that she does want to see him. Because Anna's conscious decision to break up with Vronsky has no bearing on her diachronic intention, Paul interprets this situation as a breakdown between the two. Nonetheless, she claims that these breakdowns are quite rare (Paul 2012: 338).

Crucially, Paul further argues that failing to carry out a decision does not negate its existence, because decisions are compatible with changes of mind and weakness of will. Although we may criticize the agent for being impulsive, unreliable, or overly feeble, these faults are insufficient to disprove the presence of an intention in the first place. She adds: "Even with Anna, it is difficult to find clear grounds on which to insist that she truly never intended to break off the affair, rather than concluding that the rush of emotion brought on by seeing V causes her to decide at the last minute not to break it off with V" (Paul 2012: 342).

I think Paul's version of the transparency view provides a compelling account of how we know our ordinary intentions. However, given that in the literature on self-knowledge of intention, the epistemic status of commitments is not distinguished from intentions, it might be tempting to conclude that both intentions and commitments could be seen as transparent to our own minds in similar ways. Therefore, I take her view on the epistemology of decisions (and by extension, intentions) to also apply to commitments, which are essentially a form of intention. We might be inclined to think that both can be known simply by deciding without needing to know our actions. So, I want to challenge the implicit assumption that all types of intentions can be known by simply knowing our own minds, the following view I mentioned in the introduction:

[6]Paul credits her insight to Peacocke's theory of consciously based self-ascription Peacocke and Gertler 2003). See Paul (2012: 339, footnote 22). However, she disagrees with Peacocke that we can self-ascribe the belief using a similar method.

[7]Paul (2015) proposes "conscious decisions" as a solution to an overly skeptical account that Carruthers (2011) defends. She criticizes Carruthers' account for collapsing all agency into unconscious goal-directed activity.

[8]Paul's view of the authority we have over decisions differs from Moran's (2001) view. While Moran primarily discusses beliefs, he also addresses intentions, linking first-person authority to our rational capacity to form intentions based on perceived good reasons.See Gertler (2010) for a helpful discussion on why Moran provides rationalist account of self-knowledge. In contrast, Paul ties this authority to the ability to decide independently of reasons.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

(1) knowledge of one's commitments does not, *like ordinary intentions*, require an understanding of the actions we take to fulfill them.

Before I propose an alternative epistemic account of commitments, let me clarify why I am skeptical that the breakdowns Paul discusses can be dismissed as easily as she suggests. I am particularly hesitant to accept that Anna's change of heart does not negate the existence of her intention. This is because Anna's task, given that she can foresee her desire to meet Vronsky again, is not merely to make an ordinary decision but to adopt the attitude of commitment to resist the temptation to see him. The more considered conclusion, I shall argue, is that Anna may not be in a good epistemic position to know her intention. This is partly because she fails to recognize the difficulty of the action – one that requires a firm commitment and follow-through despite the temptation. In this case, where it was foreseeable that she would succumb to weakness of will, she could be more thoughtful about how to resist that temptation. If Anna did not seriously reflect on the likelihood of changing her mind and how to stay committed, she could not effectively decide to stop seeing Vronsky. Analogously, someone who does not know how to even attempt to go to Mars cannot intend or decide to go to Mars. Someone who does not consider how to resist the temptation to realize an action cannot make a sincere decision to be committed.

To be clear, it is important to note that Paul need not deny that intentions are often stable and motivating. As Bratman has shown us, intentions are plans which "... help guide our later conduct and coordinate our activities over time, in ways in which our ordinary desires and beliefs do not" (Bratman 1984: 376).[9] Yet, we can consistently accept both Bratman's claim that intention requires planning and stability, and Paul's claim that intentions are compatible with changes of mind. For example, I can intend to go to the market to buy some chips, and while I might need to plan for that, my intention to go to the market to buy some chips would not be negated even if I change my mind and buy yogurt instead. Paul does not dispute that a deliberate decision or intention often leads to an action or deny Bratman's functionalist account, when she argues that intentions are compatible with changes of mind in the face of better options or boredom. She concedes that intention ordinarily activates standing dispositions (Paul 2015: 1543), and that we often do what we intend. Her point is rather an epistemic one: when there is no follow-through or a change of mind, that fact does not necessarily give us evidence that there was no intention in the first place. My disagreement lies in this very epistemic point. There are certain intentions where a change of mind or a preventable lapse in willpower serves as evidence that those intentions do not truly exist, and I take Anna's weakness of will as a case that falls into this category. Anna herself cannot simply know her intention by only considering what is internally available to her own mind, without factoring in the effort she puts into resisting temptation. Her inability to resist temptation should be taken as evidence that her intention may not truly exist.

As Paul convincingly showed us, beliefs are generally "impervious" to judgments, since reflecting on reasons does not substantially influence what one actually believes (Paul 2012: 333). Yet, similar breakdowns are also common in the case of intentions that hold more significance for self-knowledge. Just as our diachronic beliefs are sometimes impervious to conscious judgments, the kinds of intentions which are more interesting objects of self-knowledge are similarly impervious to conscious decisions. People decide to write books, but they do not develop the writing habits necessary to finish them; much as others decide to adopt a certain diet, without putting in enough work to keep it up. Half of marriages end in divorce. If people have that kind of low success rate with what is

[9]One of the positions that Bratman decidedly rejects is the desire-belief model, which was famously defended by Davidson (1963).

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

probably the single most important commitment they make in their entire lives, how well can we know whether we are genuinely making other sorts of commitment by only considering our conscious decisions? Crucially, such breakdowns are quite common in the cases of commitments which are more significant kinds of intentions for the purposes of self-knowledge.

Further, it is a common phenomenon that when one does not follow through on the significant commitments one makes, one questions whether one made a genuine commitment after all. People invest a great deal of money in therapy to understand why they cannot stick to their commitments. We think about the requirements for making a sincere commitment or often question whether we have already made one. If we could simply know our commitments in the same way as our ordinary intentions – transparently, by making a conscious decision – why would we feel the need to question whether we made a genuine decision when there is a persistent gap between what we decide and what we do?

Anna Karenina's struggle with weakness of will sheds light on the complexities of commitment. Paul notes that it is difficult to assert that Anna lacked the intention in the first place. Here is how I characterize the problem: Anna consciously decides not to see Vronsky. Then, there she goes again, failing to resist seeing him. Therefore, she did not really know that she intended not to see him, since that intention was merely illusory: in which case, making a conscious decision does not always guarantee adequate self-knowledge. This is not explained by "unconscious" intentions either: indeed, the decision to resist seeing him again is conscious. Does Anna feel out of control or coerced? If so, then why her subsequent guilt? I think the correct interpretation must be that she feels guilty because a commitment not to see Vronsky again requires resisting her weakness, and she lacks that commitment. She could easily foresee the possibility of changing her mind upon encountering Vronsky again, but she fails to take steps to avoid this. By changing her mind based on her strongest desires, Anna reveals an inadequate understanding of herself, her motivations, and her dispositions, which ultimately prevents her from making a genuine commitment in the first place.

These common breakdowns and the fact that we learn about our commitments by considering whether we follow it through indicate that knowledge of our commitments cannot be determined transparently through internal observation alone and that the knowledge of commitments is not as cheap as it is for ordinary intentions. While it may be true that, in the case of ordinary intentions, we do not question their existence if we fail to follow through, the uncertainty we feel about our commitments – especially when we notice regular patterns of breakdown – presents an epistemic problem for the existence of our commitments.

I will now return to explain why this epistemic difficulty arises primarily from the significant differences between commitments and ordinary intentions and argue for the following claim:

(2) Knowledge of one's commitments does, *unlike knowledge of ordinary intentions*, require an understanding of the actions one takes to fulfill them.

3. Commitments, values, and weak-willed action

3.1. The difference between commitments and intentions

While commitments and intentions are often closely related, they differ in important ways. A key difference is that commitments are made with consideration of what is

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

genuinely worth pursuing. When we intend to go for a walk or pour another cup of coffee, we are not expressing deeply held values or judgments about the worth of such actions. However, commitments are closely linked to our sense of who we are and what we value. After all, we would not invest our valuable time or forgo other pleasurable pursuits unless we genuinely believed that our commitments align with what matters most to us. In the case of commitments, we not only simply use our capacity to decide, but we also use our capacity to assess reasons in considering whether to make a commitment. But we are also good at tricking ourselves into believing that we will genuinely put the effort required to keep our promises. This reveals a crucial point: the disparities between judgments and beliefs are also observable in the case of commitments because commitments are closely linked to value judgments. Our self-bias and optimism about our abilities create a gap between identifying a good reason for a commitment and following through with the necessary effort to honor our goals.[10] We make decisions based on what we consider valuable to achieve, but we cannot reliably predict if we will develop the habits needed to stay committed. While all commitments may involve decisions, not all decisions qualify as commitments. So, the breakdowns in the cases of judgments and beliefs are also common in the case of decisions and commitments. Just as judgment and belief can diverge, so too can our decisions and our ability to sustain our commitments because commitments (unlike ordinary intentions) are not immediately permeable by conscious decisions.

This brings me to the second related point. Unlike our ordinary intentions, which we might form out of compulsion or whim, commitments require us to resist psychological challenges and maintain consistent self-discipline. While the function of a decision and intention, as Paul proposes, is to activate standing dispositions, the function of a commitment is to resist them. Thus, while a decision and an intention to Φ might overlap, the decision to commit and the state of having a commitment cannot overlap, because the function of a commitment is to counteract our standing dispositions. We make a commitment when we know that the intended goal cannot be brought about without some diachronic self-discipline and regular patterns of action. We choose our commitments freely, but in doing so, we bind ourselves to a series of decisions and actions needed to sustain the commitment and achieve the desired outcome. Because we are weak-willed agents, making a commitment is a way of holding ourselves to a standard, disciplining ourselves, organizing our desires across time, and putting ourselves to the test – to see if we will show ourselves the resolution to continue, despite wavering motivations. When we commit to being in a relationship, we promise not to give up after a trivial disagreement; when we take responsibility for writing a book, we make the commitment not to give up on it after a minor setback, or whenever there are more interesting things to do. If someone claims to be committed to writing a book but keeps changing his mind when a more interesting topic arises, we might judge that he lacks real commitment and has mistaken desire for a sincere decision. Insisting he is committed in this situation would either show a misapplication of the concept or an inability to understand one's commitment.

In contrast, ordinary intentions are compatible with being tempted by alternatives. Indeed, for example, if I intend to go to the shop to buy a skirt I saw online, but if I change my mind after seeing another one I like better, this does not negate the existence

[10]As Scheffler 2010) argues that valuing goes beyond recognizing something as deserving of value; it also requires characteristic actions and feelings toward it. While commitments cannot be reduced to valuing – one can value past memories without being committed to them – commitments do share similarities, because we commit to the actions we value. Both involve not only cognitive approval but also affective inclinations toward the goals we aim to promote.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

of my initial intention. However, unlike these (ordinary) intentions, knowing one's commitments requires whether one typically has the determination to *resist* constant changes of mind.[11] After all, commitments like writing a book or completing a PhD involve taking on challenging tasks that demand the development of self-discipline to resist self-sabotaging tendencies.

Third, commitments often require *regular patterns* of action to make progress or complete a task. While we can fulfill an ordinary intention with a single action, commitments involve taking on challenging tasks that demand ongoing effort through repeated actions to reach the goal. Thus, as I will further argue, knowing a commitment requires stronger evidence than knowing an ordinary intention.

A potential objection to this view is that failing to follow through on an action might be attributed to a mere lapse in willpower rather than an actual lack of commitment. I will now discuss why this is not necessarily the case and show how weakness of will can, in fact, surprisingly reveal a fundamental failure to uphold the commitment.

3.2. Weakness of will as a defeater of the existence of commitments

As many have pointed out in the case of intentions, when there is a breakdown between a decision or an intention about what one ought to do and one's *action*, that breakdown does not impugn the existence of the former (Broome 2002; Raz 1978; Shah 2008).[12] According to Shah, what we achieve through deliberation is forming an intention to act, while the ability to carry out that intention lies beyond the scope of deliberation itself. He rejects the idea that deliberation and intention are limited to matters within one's "voluntary control" (Shah 2008: 4). That is, we can intend what we ought to do, and whether we believe we will succeed in it is a separate matter.[13] Sometimes, the breakdown merely reflects fear, which might hinder the initiation of an action, but it does not mean that one misjudged or changed one's mind about what one ought to do. This view may be accurate regarding our understanding of ordinary intentions, but it does not apply to commitments. When we fail to act on something within our control, it can cast doubt on whether we genuinely made that commitment in the first place. After all, the whole point of a commitment is often that we undertake concrete efforts to resist weakness of will.[14]

Here, I want to flag two features of weakness of will that challenge the existence of commitments: first, *foreseeable and preventable weaknesses*, and second, *regular patterns of failure* to manage such weaknesses.[15] Specifically, foreseeable, preventable forms of

[11]This is not to say that one should never change their mind about their commitments or simply value the trait of being resolute. I am not claiming that being committed is always a good thing, regardless of how much our values change. See Calhoun (2009) for an insightful discussion as to whether the integrity following through the commitment is always a good thing.

[12]These views are in contrast with the traditional Socratic view, according to which clear-eyed akrasia is impossible (though not all the mentioned philosophers draw the contrast in this way). As Plato's Socrates claims "... no one who knows or believes there is something else better than what he is doing, something possible, will go on doing what he had been doing when he could be doing what is better" (Protagoras 358b7-c1, trans. Cooper 1997).

[13]As Shah discusses, his view contrasts with that of some philosophers who argues that intending to Φ involves believing that one will Φ (Shah 2008: 6).

[14]Compare, for instance, the Kantian maxim that "whoever wills the end also wills (in so far as reason has decisive influence on his actions) the indispensably necessary means to it that is in his control" (Kant 1785/1998 4:417), which Kant claims is an analytic truth.

[15]I am grateful to an anonymous referee for encouraging me to further elaborate on my views in this section.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

weakness of will that are in tension with our commitments serve as crucial evidence against the commitment, whereas an accidental uncontrollable impulsion or overwhelming desire does not. For example, if I intend to give a talk but am unable to attend due to a debilitating panic attack, my deliberation and initial decision remain intact, despite my physical inability to follow through. In such cases, the issue is not a failure of agency but rather external factors such as compulsion, fear, or biological constraints.

However, it'd be misleading to treat weakness of will as too close to physical inability or biological factors that are completely external to our will. Not all kinds of weakness of will are beyond our agential control. To be sure, we can encounter an unpredictable physical or psychological impediment that we could not foresee. However, weakness of will cases are not those in which one is simply overwhelmed by emotion or impulse, in which there is no internal failure of agency. While some forms of compulsion and intense, sudden emotions that are wholly uncontrollable may in fact be caused by physical or biological factors outside of one's will, we can also form commitments precisely to combat such weaknesses and to make our dispositions consistent with our commitments. As Odysseus shows us, it can be necessary to deprive oneself of the opportunity to choose wrongly such as by tying himself to his ship so that he cannot join the sirens. We tend to judge that one only made a shallow commitment if he allows his dispositions or other desires to easily get in the way of realizing his goals. Weakness of will cannot be explained by "unconscious" intentions or dispositions either, because when we experience it, we feel guilty for not exerting adequate self-control.

Even when weakness of will results from a physical or psychological issue, we can still consider whether we can manage these weaknesses to prevent them from hindering our ability to make progress on a task. If I am indeed committed to making public speeches, I could look into the solutions about what I could do to prevent my panic attacks, whether I could use therapy or more practice beforehand. We might have psychological or physical impediments that will foreseeably sabotage our goals, and we can reasonably take precautions to prevent them. This is not to say that we are always responsible for weakness of will, regardless of our efforts. To be sure, even with therapies and interventions, one might not fully overcome a panic attack. The key point is that to evaluate the degree of one's commitment, we consider the efforts we put to resist or prevent foreseeable and preventable weaknesses of will, rather than merely considering whether one has made a conscious decision. Therefore, it is necessary to distinguish between wholly uncontrollable emotions and impulses, and foreseeable and preventable weakness of will, which is within one's control and can be prevented by appropriate countermeasures. Only the latter impugns the existence of a commitment.

Considering a different example of weakness of will that is less physical in nature might be more helpful. Take, for instance, a writer who decides to finish her book but finds that spending excessive time in coffee shops impacts her writing productivity and recognizes that this habit interferes with her goal. If she believes that she is committed to working on her book project regularly but continues to spend too much time in coffee shops, knowing full well that this habit distracts from her writing time and focus, her consistent failure to alter her routine reveals that her commitment to the project is not particularly strong.

That said, we should concede that a single instance of accidental weakness of will may not be enough to undermine the commitment, as occasionally yielding to weakness of will might actually be beneficial for regaining energy and sustaining motivation.[16]

[16]See Baumeister et al. (1998), who argue that willpower is a finite resource that can be depleted with use, and allowing oneself to succumb to weakness of will from time to time could help conserve energy and maintain willpower over the long term.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

However, as I pointed out in the previous section, unlike ordinary intentions, commitments involve undertaking challenging actions that require *regular patterns of actions* to make progress or complete a task. Therefore, the claim is that it is the *regular patterns of foreseeable, preventable failure to resist one's weakness of will* that ultimately calls into question the existence of a genuine commitment. And this is supported by the widespread experience that when agents fail to follow through on significant commitments, they often find themselves questioning the existence of those commitments. When weakness of will becomes a recurring theme, it compels us to critically assess the sincerity of our commitments, urging us to confront whether our initial resolutions were genuinely made or merely aspirational.

Of course, one might do everything to control one's temperament, take precautions, yet one still might fail. In such cases, it would be unfair to think that these entirely accidental weaknesses of will immediately defeats the existence of a commitment. However, even if we do not fully succeed in overcoming these weaknesses, when assessing our commitments, what matters is the effort put into resisting to our weaknesses or dispositions. The existence of a commitment would not, and should not, be judged solely on a mental state or a single moment of decision.

4. How do we know our commitments?

4.1. An externalist proposal

Given that regular patterns of weakness of will undermine the existence of commitments, I can now sketch why we should adopt an externalist view to understand knowledge of commitments, drawing a surprising contrast with the recent view of intentions introduced in (1) at the outset of the paper. My argument for this externalist view appeals to the following claim: if S systematically fails to Φ in relevant circumstances, then S should infer that S is not committed to Φ-ing in those circumstances. We may *believe* that we are committed to Φ in particular circumstances, but we cannot know this until we have sufficient evidence that we actually *will* reliably Φ.

Knowledge of commitment cannot depend only on one's internal mental events, a traditional assumption that is associated with Descartes. I argue for an externalist view because a commitment to Φ has its own external correctness conditions, which are objective insofar as they cannot be altered purely through acts of will or thought. For instance, there are external conditions of satisfaction for whether a person is committed to writing a book, which cannot be altered through pure acts of will. While we may discover and explore new ways to write a book, the introduction of new correctness conditions is appropriate only if these adjustments enhance efficiency in making progress toward our goals. Self-knowledge of commitments requires a "reality check": one must look to external evidence to determine whether one typically reason and acts accordingly.

Contrast this with our self-knowledge of ordinary intentions. Ordinary intentions are linked to short-term actions, and in many cases, my intention to Φ and my knowledge of that intention are inseparable; for example, my intention to stretch my limbs is typically intertwined with my knowledge of that intention. In contrast, commitments are not merely one-time actions; they involve a broader range of actions and attitudes over appreciable periods of time. And there are correct and incorrect, feasible and infeasible, ways to go about realizing one's goals. They are the kinds of actions which, if one is disastrously wrong or ignorant of them, or fails to act in accordance with them, undermine one's pretention of being truly committed to Φ. So, knowledge of commitments can be confirmed or undermined by agents' regular patterns of action.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

Whether you truly are committed to the relationship depends not just on what you intend at the time of undertaking it: it depends also on how you act and react over time (in which case, there might not yet be a fact of the matter about whether you are committed or not).

One might object that even if S's failure to reliably Φ implies that S is not committed to Φ-ing, it does not follow that, in order to know in advance that S is committed to Φ-ing, S must be certain that she will reliably Φ in relevant circumstances.[17] In other words, one might contend that I have demonstrated only that if an agent fails reliably to Φ, it undermines her knowledge of her commitment to Φ – not necessarily the commitment itself. If I believe I am committed to Φ-ing but then observe that I am not reliably doing so, I should abandon that belief, not the commitment.

My response to this objection is that commitments function precisely to make us think about how to overcome temptations, weakness of will, and other challenges that we are capable of preventing. So, if an agent does not start to resist her inclinations or trying to improve her relevant skills, then it is hard to say that she still has or did make a genuine commitment. That said, I acknowledge that consistently acting in ways that are incompatible with Φ-ing is a more conclusive defeater than merely failing to act in ways that would bring about Φ-ing. Acting in ways that directly contradict our main goal is a stronger indication of its defeat than simply not trying enough to achieve the goal.

Further, we might *desire to have a commitment,* and we often make irrational commitments, knowing full well that we might fail. In those cases where we do not try enough, the commitment remains merely aspirational rather than a genuine one.[18] To be sure, this does not imply that we cannot (or should not) make commitments if we are unsure of our ability to follow through. As Marušić (2015) argues, we can resolve to do something even if the evidence from our own past makes it likely that we will not follow through. We can commit to actions on the basis of practical reasons, and despite the evidence, even when we are unsure of our chances of success. Marušić does not aim to explain how one can know that one has a commitment. His point is that one can resolve to do things, despite contrary evidence from one's past conduct. This is compatible with my claim that knowing whether one is committed requires evidential support related to the actions necessary to achieve one's goals. My claim is that independent of whether we are rational or irrational in making commitments, the epistemic story about how we discover them will be the same. We learn about the extent of our commitments by observing our regular patterns of action, how much we struggle with our weakness of will, how well we control our behavior, and what steps we take to maintain them. Whereas ordinary intentions are not defeated by a lack of follow-through, commitments are undermined or confirmed by what we do.

One might also object as follows: suppose someone is committed to helping the poor and starts an organization that sends college students to build houses in impoverished areas.[19] It may turn out that this effort further impoverishes the area by taking jobs away

[17]I thank an anonymous referee for a careful presentation of this objection.

[18]Velleman (1989), Setiya (2008), and Paul (2009) each offer distinct answers to the question of how much belief an agent needs regarding the likelihood of success for her intention to be classified as truly intending a goal, rather than simply desiring or hoping for it. Velleman (1989) argues that we intend to act on our strong desires because we are motivated to predict our actions. Setiya (2008) argues that self-knowledge regarding the likelihood of success is grounded in know-how knowledge. Paul (2009) critiques these views, and introduces different criteria. Whether these views can adequately address the nature of commitments and our belief in the likelihood of success requires a separate discussion.

[19]I'd like to thank Walter Hopp and Eric Schwitzgebel for encouraging me to clarify this point.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

from local construction workers. Do we then want to say that this person was not truly committed to helping the poor?

In response, I identify two forms of extra-mental evidence about propositions of the form "S is committed to Φ." The first facts are about what we intentionally do: that is, whether S acts in such a way as to intentionally Φ in a consistent fashion. The second facts are regarding which actions *actually* promote achieving Φ effectively. In the example above, our good doer is ignorant of the second set of facts. But the first body of facts strongly supports the view that they are committed: their attempts are serious and consistent with their stated aim, and with what they think is necessary or conducive to it. Yet, I hold that these discrepancies are less likely to occur in the case of commitments than in ordinary intentions directed toward single, one-off actions because commitments require learning from trial and error through regular patterns of action. That is, commitments demand staying attuned to reality and the effects of our actions in the world, as this is the only way to succeed or make progress toward our goals. Further, ideally, discrepancies between these two types of facts will be rare, as committed agents should be sensitive to external feedback about the efficacy of their efforts. If one truly cares about and commits to Φ, they will consider which actions are most conducive to achieving it. While we may sometimes fail to choose wisely, what matters is the body of evidence showing that one sincerely, consistently, and justifiably believes that actions x, y, and z will bring about Φ.

It is true that some features of our actions are strongly external. For example, whether you end up in Boston depends on whether the plane actually takes you there as promised. You might intend to go to Boston but end up in Moscow due to completely external, unforeseen circumstances. Other conditions are only weakly external. In such cases, what typically justifies one's commitment is whether the agent made sufficiently reasonable, determined, and consistent efforts to accomplish her goal and whether any mistakes made about how to achieve it were non-culpable (and, if mistaken about how to get there, whether those mistaken beliefs were formed non-culpably).

My claim is that self-knowledge of commitments is external in a weak sense. The relevant external conditions are the foreseeable norms for fulfilling an action. For example, if we are committed to running a marathon, we need to follow a training regimen, maintain proper nutrition, and build endurance. There is shared, intersubjective knowledge about how to train for a marathon effectively. While we cannot control every external factor – such as weather or injuries – and might need to adjust or even abandon our intentions when those conditions are unmet, we are still bound by the norms for preparation if we intend to complete the race. Evidence about how well those norms or conditions are met justifies our beliefs about our commitment.

4.2. Denying first-personal privileged epistemic access

Now that an externalist proposal is on the table, I can elaborate how this view helps us evaluate and justify our commitments both to ourselves and others. If knowledge of commitments can be confirmed or undermined by agents' regular patterns of action, then these actions are, in principle, observable by others. This is why, while we might have first-personal authoritative epistemic access to our ordinary intentions, we lack such access to our commitments.

This is evident in our tendency to challenge the commitments of others when they repeatedly fail to take adequate steps toward achieving their stated goals or meet the external conditions necessary for success – conditions that, in principle, can be observed by others. Consider the case of a CEO who publicly claims to prioritize inclusivity but consistently neglects critical issues like pay gaps between men and women, reports of

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

discrimination, and promotes the same homogeneous leadership team. While their public statements may suggest a commitment to inclusivity, the absence of concrete actions raises serious doubts about the existence of that commitment.

To be clear, the focus of this paper is on *intrapersonal commitments* rather than *interpersonal* ones, as the latter is not a form of intention.[20] For example, if A tells B, "I promise to attend your talk," A has made a promise, regardless of whether A shows up or consistently breaks promises. A's patterns of behavior, past or future, are simply irrelevant to whether or not he has promised B that he will come to her talk. In such cases, the key fact is that the promise-maker took responsibility to keep his promise, not for his actions. Interpersonal commitments are performative statements that might have no bearing on achievement, achieving any kind of accomplishment, or intending to achieve any such success. My analysis focuses only on *intrapersonal commitments*, which require a different analysis. However, it is still worth noting that A might still be criticized, not just for breaking the promise, but also for failing to make a genuine effort to keep it. This distinction between lacking commitment and breaking a promise highlights the relevance of intrapersonal commitment, even in cases of interpersonal unfulfilled promises.

In some cases, an agent might insist that he is committed but consistently takes the wrong means to achieve his goal. This situation is less likely to occur with commitments than with ordinary intentions because commitments involve regular patterns of action that can be evaluated and adjusted over time; they thus allow more room for a reality check on whether we are on the right track. In fact, we would and should also assess whether an agent is committed, considering to what extent they are sensitive to reality checks and the relevant outside feedback in making progress toward their goals. If we care about our commitments, we engage in ongoing evaluation which makes it harder to overlook the constraints that affect the effectiveness of our actions. Commitments require us to be in touch with reality and assess the effects of our actions in the world to make progress toward our goals. Just as we risk being irrational if we maintain a belief despite contrary evidence, we risk being irrational if we insist that we are committed while consistently using ineffective or contradictory means to our ends. Authoritative first-personal privileged epistemic access to our momentary decisions does not eliminate this risk.

I noted at the outset that my view is neither behaviorist nor dispositionalist. Dispositionalists hold that we know ourselves through observation, just as we observe and infer what others believe and intend. As Schwitzgebel insightfully argues, believing primarily involves having *stereotypical dispositions* to reason and behave in ways that we would typically associate with having that attitude (Schwitzgebel 2002: 250–53). And, indeed, others might be better judges of what we believe. Although I propose an externalist view according to which knowing one's own commitments is much like knowing those of others, I do not endorse a dispositionalist account of commitments. On an extreme version of a dispositionalist analysis, any time I act against an expectation that I will do A, it is because I have a disposition to do B instead. Yet, this would be problematic for a theory of commitment. It would fail to explain why we regret those deviations or pride in our successes, or "reactive attitudes"[21] toward others' failures to uphold commitments, when their commitments matter to us.

By contrast, commitments are quite different from other goal-directed actions to which we are not committed. So, the evidence for them is not purely observational. As a type of intention, our commitments are not simply facts we uncover by observing

[20]I thank an anonymous referee for making this distinction salient.

[21]As Strawson (1962) uses the term.

https://doi.org/10.1017/epi.2025.5 Published online by Cambridge University Press

ourselves; they are actions we choose to perform. While such observational data can undermine a claim of being committed, our reactive attitudes, such as regret when we fail and pride when we succeed, indicate that we take responsibility for our commitments, recognizing them as more than just dispositions or habits.

Denying the privileged epistemic authority does not require choosing between a first-personal and a dispositionalist approach. An extreme dispositionalist view would imply that any failure to act in accordance with our commitments implies we never truly had the relevant dispositions. However, our actions might vindicate the presence of a genuine commitment as we can work on overcoming the dispositions that hinder our goals. By adjusting our dispositions or acquiring new ones, we prove our commitment. While I agree with the transparency view that commitments are first-personal because it is up to the agent what to commit to and initiate the action, I reject the idea that the first person is in an epistemically privileged or authoritative position to know her commitments. Therefore, we can still preserve some of the virtues of the transparency view without losing the epistemic merits of an externalist, evidentialist account.

5. Conclusion

In this paper, I focused on an especially important class of persisting intentions, which I called "commitments." I began by discussing Sarah Paul's transparency account, according to which conscious, decision-based self-ascriptions of our intentions are sufficiently reliable to count as knowledge. Paul argues that such knowledge is not defeated by changing one's mind in the face of temptation, boredom, or better options. I argued that this model cannot extend to commitments, because the nature of commitments differs significantly from that of ordinary intentions, and, crucially the function of commitments is to resist weakness of will. Therefore, the breakdowns between our avowed commitments and our actual actions plausibly undermines knowledge of commitments. In contrast to ordinary intentions, the knowledge of which is complete at the time of decision, I argued that we cannot know our commitments merely by making a conscious decision. Rather, we must look to external evidence about our regular patterns of action and reasoning if we are to gain the requisite sort of self-knowledge.[22]

References

Anscombe G.E.M. (1963). *Intention*. Cambridge: Harvard University Press.

Armstrong D.M. (1981). *The Nature of Mind and Other Essays*. Ithaca, NY: Cornell University Press.

Baumeister R.E., Bratslavsky E., Muraven M. and Tice D.M. (1998) 'Ego Depletion: Is the Active Self a Limited Resource?' *Journal of Personality and Social Psychology* **74**(5), 1252–1265.

Boyle M. (2011). 'Transparent Self-Knowledge.' *Aristotelian Society Supplementary Volume* **85**(1), 223–241.

Bratman M. (1984). 'Two Faces of Intention.' *The Philosophical Review* **93**(3), 375–405.

Broome J. (2002). 'Practical Reasoning.' In J.L. Bermúdez and A. Millar (eds), *Reason and Nature: Essays in the Theory of Rationality*, pp. 85–111. Oxford: Oxford University Press.

[22]I am deeply grateful to Sarah Paul for her invaluable guidance, comments at various stages of this paper, and many conversations. I am also indebted to Walter Hopp for his extremely helpful comments throughout the development of this paper. I have benefited from the insights and feedback of Richard Moran, Eric Schwitzgebel, Erin Seeba, and Allen Speight on earlier stages of this work. I thank Zara Amdur and Aja Watkins for their editorial contributions to different sections of this paper and Matthew Macdonald for editing the entire manuscript. I also thank anonymous referees for their helpful feedback. I extend my gratitude to an anonymous referee at *Episteme* for the insightful suggestions on restructuring and presenting the work.

Byrne A. (2011). 'Transparency, Belief, Intention.' *Aristotelian Society Supplementary Volume* **85**(1), 201–221.
Byrne A. (2018). *Transparency and Self-Knowledge* Oxford: Oxford University Press.
Calhoun C. (2009). 'What Good is Commitment?' *Ethics* **119**(4), 613–641.
Carruthers P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford: Oxford University Press.
Cassam Q. (2015). *Self-Knowledge for Humans*. Oxford: Oxford University Press.
Chalmers D. (2003). 'The Content and Epistemology of Phenomenal Belief.' In A. Jokic and Q. Smith (eds), *Consciousness: New Philosophical Perspectives*, pp. 220–272. Oxford: Oxford University Press.
Davidson D. (1963). 'Actions, Reasons, and Causes.' *The Journal of Philosophy* **60**(23), 685–700.
Evans G. (1982). *The Varieties of Reference*, Oxford: Oxford University Press
Gertler B. (2010). *Self-Knowledge* New York: Routledge.
Holton R. (2011). *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
Hornsby J. (1980). *Actions*. London: Routledge and Kegan Paul.
Kant I. (1785/1998). *Groundwork for the Metaphysics of Morals*. Cambridge: Cambridge University Press.
Locke J. (1689/1975). *An Essay Concerning Human Understanding*. Oxford: Oxford University Press.
Lycan W.G. (1996). *Consciousness and Experience*. Cambridge, MA: MIT Press.
Marušić B. (2015). *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford: Oxford University Press.
Moran R. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University press.
Paul S.K. (2009). 'How We Know What We're Doing.' *Philosophers' Imprint* **9**(11), 1–24.
Paul S.K. (2012). 'How We Know What We Intend.' *Philosophical Studies* **161**(2), 327–346.
Paul S.K. (2014). 'The Transparency of Mind.' *Philosophy Compass* **9**(5), 295–303.
Paul S.K. (2015). 'The Transparency of Intention.' *Philosophical Studies* **172**(6), 1529–1548.
Peacocke C. (2003). 'Conscious Attitudes, Attention, and Self-Knowledge.' In B. Gertler (ed), *Privileged Access*, pp.63–98. Aldershot: Ashgate Publishing Company.
Plato. (1997). 'Protagoras.' In J. Cooper (ed), *Plato: Complete Works*. Indianapolis, IN: Hackett Publishing Company
Raz J. (1978). *Practical Reasoning*. Oxford: Oxford University Press.
Ryle G. (1949). *The Concept of Mind*. London: Hutchinson's University Library.
Scheffler S. (2010). 'Valuing.' In *Equality and Tradition: Questions of Value in Moral and Political Philosophy*, pp. 15–40. Oxford: Oxford University Press.
Schwitzgebel E. (2002). 'A Phenomenal, Dispositional Account of Belief.' *Noûs* **36**(2), 249–275.
Schwitzgebel E. (2012). 'Self-Ignorance.' In J.L. Liu and J. Perry (eds), *Consciousness and the Self*, pp. 163–174. Cambridge: Cambridge University Press.
Setiya K. (2008). 'Practical Knowledge.' *Ethics* **118**(3), 388–409.
Setiya K. (2011). 'Knowledge of Intention.' In A. Ford, J. Hornsby and F. Stoutland (eds), *Essays on Anscombe's Intention*, pp. 170–197. Cambridge: Harvard University Press.
Shah N. (2008). 'How Action Governs Intention.' *Philosophers' Imprint* **8**, 1–19.
Strawson P.F. (1962). 'Freedom and Resentment.' *Proceedings of the British Academy* **48**, 187–211.
Velleman J.D. (1989). *Practical Reflection*. Princeton: Princeton University Press.

Merve Rumeysa Tapınç earned her doctorate from the Department of Philosophy at Boston University, after completing an MA in Philosophy at Boğaziçi University. Her research interests include epistemology, with a particular focus on self-knowledge, philosophy of action, and moral psychology. She is currently a lecturer at Tufts University.

Cite this article: Tapınç M.R. (2025). "Knowing Your Commitments in Action." *Episteme* 1–17. https://doi.org/10.1017/epi.2025.5