



Received: 4 October, 2024

Accepted: 1 July, 2025

Published: 27 October, 2025

On Knowing What One Really Wants: Emotions, Valuing, and Substantial Self-Knowledge

Alba Montes Sánchez, PhD – Universidad Carlos III de Madrid (Spain)

Abstract

The contribution of emotions to our knowledge of ourselves remains undertheorised in the philosophical literature on self-knowledge, even among authors who focus on “substantial” self-knowledge. This is a shortcoming because emotions are crucial to obtaining the sort of self-knowledge that can play a transformative role in our lives (or engaged self-knowledge). To show this, here I rework Krista Lawlor’s (2009) example of Katherine, a woman who is wondering whether she wants a second child. Lawlor offers an empiricist account of substantial self-knowledge, where purely epistemic self-knowledge, gained through introspection and inference, is enough to satisfy Katherine and “stick” (i.e., halt her investigation). While Lawlor conceives of sticking as a mere absence of the motivation to investigate further, in my analysis sticking is an affective phenomenon that involves the epistemic emotion of certainty based on emotional insight, i.e., the right evaluative responses to what one knows. But achieving emotional insight in cases of opaque substantial self-knowledge is not easy. Here, I go beyond Lawlor by proposing that investigating what we value through an imaginative exploration of our emotions and other evaluative attitudes is a particularly promising route.

Keywords: emotions, substantial self-knowledge, engaged self-knowledge, emotional insight, desires, valuing, imagination

1. Introduction

The contribution that emotions make to our knowledge of ourselves remains undertheorised in the philosophical literature on self-knowledge. This is a shortcoming because, as evaluative phenomena that spring from what we care about, emotions are crucial for investigating and obtaining substantial self-knowledge, i.e., self-knowledge of the effortful kind that can be considered a valuable epistemic and existential achievement (Cassam 2014). To start articulating the importance of emotions in this domain, in this paper I re-examine Krista Lawlor’s (2009) question: how does one know what one wants? In formulating her account of how we (often) get to know our opaque desires, Lawlor deliberately presents us with a familiar but difficult example: that of Katherine, the mother of a young toddler, who is wondering whether she wants another child. What sort of question is this, and what sort of answer can satisfy it? According to Lawlor, this is a purely epistemic question about an opaque desire, and a satisfactory answer amounts to a correct self-ascription of desire. In my view, however, something more is crucially at play in this sort of case.

In fact, Lawlor herself seems to point beyond epistemic considerations in insisting that the investigation process ends when one finds a self-ascription that “sticks,” i.e., when one ceases to be inclined to interrogate it. But, as I argue

in Section 3, it is rather unclear what sticking has to do with the epistemic status of the self-ascription: sticking rather seems to amount to an affective change of stance on Katherine's part towards her mental life, something that resembles the sort of emotional insight that psychotherapists talk about (see e.g. Bell and Leite 2016). While *epistemic knowledge* is purely a matter of belief, of arriving at a true, grounded belief that one has a particular desire or emotion, *emotional insight* entails not only epistemic knowledge of one's mental state, but also a grasp of its significance and implications, and an appropriate emotional response to what one knows. The latter therefore is a much better candidate for the kind of phenomenon that has enough motivational force to halt an investigation such as Katherine's and make a self-ascription "stick." This element of Lawlor's own solution to the problem thus suggests that another sort of question, beyond the purely epistemic, is in the background of her investigation, namely: how can we obtain self-knowledge that is efficacious, or, in other words, that plays the kinds of roles that many of us want it to play in our lives?

In this paper, I aim to highlight the crucial contribution that emotions can make to the acquisition of efficacious substantial self-knowledge by reworking Lawlor's example and spelling out the roles that affective states play in such a case. The paper starts by explicating the notion of substantial self-knowledge, as introduced by Cassam (2014), and arguing that the reason why many people seek it and find it valuable is its transformative potential. Nonetheless, for a piece of self-knowledge to have this potential, the subject must relate to it in an engaged rather than a detached way (Corbí 2023). Both empiricists and agentialists about self-knowledge (see e.g. Gertler 2018; Sorgiovanni 2019; more on this in Section 2) offer detached methods of investigation, and these cannot guarantee engaged, efficacious self-knowledge. To address this problem, in Section 3 I summarise Lawlor's account of her Katherine example and her proposal that one finds out what one wants by causal inference to a self-ascription that sticks. Sticking, for Lawlor, is a mark of self-knowledge and consists in the absence of a motivation to keep investigating. In my view, however, sticking is a complex affective phenomenon encompassing an epistemic feeling of certainty as well as emotional insight (i.e., the appropriate affective and motivational responses to what one knows). But how does one achieve this? The last two sections of the paper are devoted to sketching a particularly promising route. I start Section 4 with the idea that Katherine doesn't want to know about just any desire, but a *deep* desire. To obtain sticky knowledge of such a desire, one needs to grasp other aspects of oneself—namely, what one cares about and values—by engaging one's emotions. As I show in Section 5, emotions, together with desires, are one of the key elements of the patterns of evaluative responses that cluster around what we care about and value. Thus, exploring the patterns of emotional responses to one's future possibilities is a particularly conducive way to achieve engaged (i.e., efficacious) self-knowledge of one's opaque deep desires.

2. The Value and Efficacy of Substantial Self-Knowledge

To begin with, let me stress, for the sake of clarity, that this paper does not aim to contribute to debates surrounding the epistemic privilege of the first-person, on whether a kind of self-knowledge exists that is immediate, authoritative, or self-intimating. The kind of self-knowledge at stake here is what has been called "substantial" (Cassam 2014) or "Socratic" (Renz 2017). In *Self-Knowledge for Humans*, Quassim Cassam (2014) distinguishes between two sorts of self-knowledge: trivial and substantial. This distinction is not categorical, but a matter of degree: different bits of self-knowledge can be more or less substantial. In a nutshell, trivial self-knowledge (e.g., knowing that I am wearing socks) is easy to obtain and usually reliable, but it is typically perceived to be of no great value by its subject. By contrast, substantial self-knowledge (e.g., knowing that I'm fastidious) is difficult to obtain and easy to get wrong, i.e., epistemically opaque, but usually conceived of as much more valuable by its subject. Since Cassam proposed this distinction, several philosophers have challenged the idea that trivial self-knowledge is of little value, and pointed out its importance for grounding our rationality and enabling all sorts of crucial

everyday endeavours, so trivial self-knowledge might be more valuable than Cassam assumes (see e.g. Boyle 2015; Jongepier 2021). Nonetheless, he is right that in the search for meaningful self-understanding, wisdom, or the like, most of us tend to think of his “trivial” instances of self-knowledge as negligible. This is because they lack a certain personal and existential relevance, they don’t contribute much to one’s sense of self or one’s self-concept (Cassam 2014, 30), and they are thus unable to play some important roles in our lives, such as giving us a sense of purpose, guiding us in decisions or allowing us to reflectively correct our vices.

Understood in this way, the opacity and existential value of substantial self-knowledge are independent from each other, they are both a matter of degree, and they do not always match; in other words, more opacity is not automatically equivalent to more existential value. For example, it might be very easy for a particular person to know they love their child, and nonetheless this piece of self-knowledge might be one that is central to structuring their self-understanding and guiding their decision-making in cases of explicit value conflicts. Thus, for them, this piece of self-knowledge would be more existentially than epistemically substantial. This doesn’t rule out the possibility that all self-knowledge has some intrinsic epistemic value, or that some bits of self-knowledge might be prudentially disvaluable for some individuals (e.g. because they might cause counterproductive pain). Self-knowledge can have different kinds of value, not all of them aligned. But Cassam’s distinction between trivial and substantial self-knowledge highlights the existentially transformative potential that many of us seek in some instances of self-knowledge; and this is what interests me in this paper.

The literature that deals with substantial self-knowledge, Cassam and Lawlor included, is mostly preoccupied with epistemic opacity and how to overcome it. This is a vital endeavour, since epistemic soundness is what allows us to avoid self-deception and ensure that what we obtain in the process is, well, *knowledge* rather than something else. But as Corbí (2023) eloquently defends, if our relation to the evidence that grounds our self-knowledge is detached rather than engaged, the knowledge we obtain from it won’t be able to play the transformative roles that many of us want it to play in our lives. This is a problem that, according to Corbí, plagues both empiricist and agentalist views of self-knowledge alike.

There are many different versions of both empiricism and agentalism about self-knowledge, but let me try to offer some broad characterisations. Empiricists all share the idea that our self-knowledge involves detecting what’s on our minds in one way or another (Gertler 2018; Boyle 2015). Cassam’s and Lawlor’s inferentialism, where self-knowledge comes from drawing inferences based on evidence, is a version of this type of view. Agentalism, on the other hand, rests on the intuition that we are not merely passive observers of our minds: many of our mental states are of our own creation or, when they aren’t, we can deliberate about them and endorse them as ours, and this is what gives us special epistemic authority over them.¹

1 The division of the self-knowledge debate into two camps is a simplification that misses, at least, a third big player: expressivism, as defended by Finkelstein (2008) or Bar-On (2004). Roughly, the central tenets of expressivism are that statements like “I am so happy!” or “I want a glass of water” are avowals that flow directly from our mental states and express them in roughly the same way that shouting “ouch!” expresses pain (they are not based on evidence), and this grounds some forms of epistemic privilege or at least a presumption of truth. Whether these avowals count as self-knowledge or just as privileged ground for it is a matter of debate. But importantly for my purposes, the cases of opaque substantial self-knowledge that I examine here are cases where avowal seems impossible or is plagued by obstacles and ambiguities, i.e., where expressions seem not to “flow” as they are supposed to. Finkelstein might retort that uncertainty is still a mental state that can be expressed and known about, but this doesn’t capture the process of probing and searching for an unambiguous desire. Bar-On might in turn suggest that the search is carried out through a process of self-interrogation, consisting of a series of expressive acts of tentative desire, aiming at a clear answer. This would make one’s mental states more concrete and easier to deliberate about. At this point, however, the process of finding out what one wants powerfully resembles what Moran proposed. Therefore, and given that the main authors I dialogue with (Lawlor and Corbí) do not address expressivism separately, I do not to give it a separate treatment in this paper either.

Despite their differences, Corbí (2023, 5–9) argues that both types of views share the “Detachment Assumption” that self-knowledge is obtained from the perspective of a detached observer. Interestingly, something akin to Corbí’s concern inspired Richard Moran’s (2002) original criticism of empiricism. According to Moran, the mere detection of mental states can only render third-personal, alienated, and motivationally inert bits of information about oneself. In line with this, Corbí (2023, 8–9) shows that Cassam’s and Lawlor’s inferentialism rests on the assumption that the self-knowledge seeker bears a detached relationship to the evidence. In fact, Cassam (2014, 149–50) insists that the only significant asymmetry between the subject of self-knowledge and an external observer is the kinds of evidence they have available.

Moran worried that examining oneself in this way only produces alienated self-knowledge, and proposed that the problem could be solved through deliberation on one’s reasons to believe, want, and so on. For Moran, such deliberation renders knowledge of what one actually believes or wants because both questions are transparent to each other: we find out that we believe something by looking at the world, not at ourselves, and the same, with some modification, holds for other attitudes like desire. For example, we find out whether we want to eat dinner at the Indian restaurant down the road by examining the reasons to desire, or the desirability of, such an option (cf. Byrne 2011, who articulates this idea within the empiricist framework). What deliberation adds, according to Moran, is the capacity to endorse one’s attitudes by finding the reasons to have them, thus making one’s attitudes properly one’s own. Generally, empiricists believe that an account that equates making up one’s mind with self-knowledge is wrong, because acting upon our minds seems like a way of changing them, not like a way of knowing what’s on them: a detective should respect the evidence, not tamper with it (Doyle 2022, 134). Furthermore, as Lawlor (2009, 71) remarks, transparency is dubious: wanting something is not the same as formulating one’s reasons to want it through deliberation. In fact, in many cases it is easier to know one wants something than to find any reasons to want it (Cassam 2014, 103–4). Finally, and most importantly for my purposes, deliberating doesn’t rule out alienation, since one can be as alienated from the conclusions of one’s deliberation as from the knowledge obtained by detection and inference (Corbí 2023, 9). In Corbí’s diagnosis, this is because deliberation is also carried out from the perspective of a detached observer.

In sum, both approaches involve a detached relation to the evidence, and thus risk delivering a result that, even if epistemically correct, might still fail to produce the desired effects. One might infer one has a particular vice and remain unwilling to correct it. One might conclude one has excellent reasons to accept a job offer abroad and still feel sure one doesn’t want the job. My point here is not that there is something epistemically wrong with the methods proposed by different self-knowledge theorists, but rather to point out that many fail to address the value and efficacy issue: how do we overcome not only the opacity challenges, but also the efficacy challenges, in order to get to self-knowledge that can be transformative?

In this paper, I focus on the efficacy issue and defend the idea that emotions are indispensable here, not simply as mental states we can get to know or as bits of evidence for inference, but as engaged responses to changes in our environments that evaluate such changes on the basis of our cares and concerns and prepare us to act in specific ways. Affective responses are always engaged, since they spring from what we care about, and they are—I will argue—a key ingredient of efficacious self-knowledge. Lawlor, for example, offers a fine account of a route to overcome the opacity challenges to substantial self-knowledge, but her example points beyond the epistemic. So let’s delve into it.

3. Lawlor on Knowing What One Wants: Stickiness, Certainty and Emotional Insight

In her suggestive paper, Lawlor (2009, 56) approaches her question by trying to be as faithful as possible to the phenomenology of the search for self-knowledge. She describes in detail Katherine, a woman who is trying to find out whether she wants another child. To do this, she spends many days occasionally reminiscing about her son as a baby, imagining that she holds a new baby in her arms, feeling envious when an acquaintance reveals her pregnancy, Sometimes her memories and imaginings appear spontaneously, sometimes she engages in them on purpose. This seems like a reasonably realistic description of the way in which many of us go about exploring our minds when we are in doubt about what we want, but as Lawlor remarks, it doesn't look like a process of Moranian deliberation to make up one's mind. Katherine is not making efforts to *decide* whether to try to have another child. She is trying to *find out* whether she desires one or not.

To do this, Lawlor argues, Katherine employs a process of causal inference from "internal promptings," especially her imaginings. She uses these experiences as evidence for her desire: these promptings must have a cause. Over time, by inference, Katherine reaches the conclusion that the cause is her desire for a second child. This self-ascription might or might not stick and be stable. If unstable, the investigation will restart. But if the self-ascription sticks, then Katherine has reached self-knowledge by causal inference: she knows her desire. This, for Lawlor, amounts to a purely cognitive achievement. Furthermore, purely epistemic self-knowledge (a correct self-ascription) is all Katherine seeks in examining the question and all that she needs to satisfy her pursuit.

While Lawlor's account has many merits, I disagree with her contention that in these cases, when we look for self-knowledge of an existentially relevant and potentially transformative nature, all that is at stake are correct self-ascriptions and purely epistemic achievements. To see this, we need to pay close attention to the notion of sticking. For Lawlor, a self-ascription sticks when it settles one's question and puts an end to the investigation. But what sort of attitude is sticking such that it can have these effects? And why is sticking necessary to mark the achievement of epistemic self-knowledge? What is the relationship between sticking and the proper acquisition of a correct belief about one's mental state?

Starting with the second question, note that what makes something an instance of knowledge has to do with the sort of relation that holds between one's belief and the world and the aptness of the procedure to form such a belief. But when Lawlor describes the process of arriving at a sticky self-ascription, she doesn't portray Katherine as checking the range and quality of her evidence and the soundness of her inferences to arrive at a judgment that her ascription is well founded. What Lawlor describes is a process of trial and error that relies on Katherine's affective responses to the self-ascription. Lawlor (2009, 59) presents us with three possibilities: "trying on" a self-ascription to see how one feels about it, believing one has landed on a correct self-ascription just to find oneself reopening the investigation a little later, and, finally, full-blown sticking, where the motivation to reopen the question disappears for good.

This sounds to me like a sensible proposal: I find it hard to imagine most people systematically checking their self-ascriptions for epistemic soundness. So most of us tend to rely on trying them on to see how they feel. This is sensible, first, because the opacity challenges of substantial self-knowledge make it extremely difficult to ascertain that one has consulted all the relevant evidence, but secondly, also because judgments can be detached and motivationally inert. So we need an attitude that halts the investigation through some form of heuristics before it becomes cognitively too costly, and which has motivational power. These are precisely the

roles that epistemic emotions play in our mental life (de Sousa 2008; Morton 2009; Candiottio 2019). Sticking, therefore, is best conceived as an epistemic emotion of certainty (de Sousa 2008, 147, 149–50) or a feeling of clarity (Nguyen 2021): a change of attitude towards one's mental life, which includes an evaluation of it ("This is clear," "I'm sure of it," or the like) and an action tendency (to stop questioning and open up the space for action).

This is increasingly recognised and literature on the epistemic emotions has exploded in recent years, but how are these states characterised? First, they are emotions: intentional evaluative states that typically imply action tendencies (Scarantino and de Sousa 2021). In other words, emotions are directed at specific objects or states of affairs, which they place under a characteristic evaluation (fear presents its intentional objects as threatening, anger, as offensive, and so on), and prepare us to act in specific ways (in fear one tends to flee, fight, or freeze, in anger one tends to attack, and so on). Both the evaluations and the action tendencies are aspects of a subject's response to a situation where something they care about (to a higher or lesser degree) is at stake; if we don't care at all, we don't emote. Emotions express the manner and intensity of our cares and concerns: the same event, let's say a ground-breaking scientific discovery, will trigger very different emotions in different people (the main scientist in charge of the project, her husband, the head of a rival lab, her next-door neighbour, the president of the university, a science journalist, and so on), depending on whether, how, and to what extent they care or don't care about this event. Second, these emotions are epistemic: they spring from our caring about knowledge and thus they motivate, fuel, or end our search for it. Emotions such as doubt, interest or certainty play essential roles in initiating, sustaining, and halting our epistemic endeavours, because they have a motivational force that purely cognitive states lack. In Lawlor's example, Katherine gradually moves from doubt to certainty, and this is what changes her motivation from one of questioning and examining her mental life to one of stopping the investigation process and opening the space to act on the basis of what she knows.

Now, how does Katherine move from one to the other? If all that Katherine was after was purely epistemic self-knowledge, reflecting on her investigation process and concluding that she has considered all the relevant evidence and interpreted it through a correct procedure (e.g. inference) would be all that was needed for sticking. Something like this can be what triggers certainty in cases of purely epistemic knowledge (although, of course, we can be manipulated into wrongly believing that all of that is in place, or that it is unnecessary because we trust a purported authority; see Nguyen 2021), and sometimes it might be all we seek. In Katherine's case, however, this is not enough: recall that Lawlor readily admits the possibility that certainty (or sticking) can misfire and be provisional. Katherine might wrongly think she has reached self-knowledge, only to find herself returning to doubt and unexpectedly reopening the question (Lawlor 2009, 59). Her emotions might go against her conclusion; in other words, new evidence might surface that questions the deliverance of her emotional heuristics (i.e., certainty) and reveals it as unwarranted.

This, in my view, is particularly likely to happen in a case such as this one because Katherine is looking for *substantial* self-knowledge, i.e., knowledge that is not only difficult to obtain but that can contribute to her self-understanding and play a transformative role in her life, guiding important decisions or allowing her to change, for example. For this to happen, the self-ascription needs to trigger not only temporary certainty, but a wider range of responses that support and confirm it: a practical-affective change of stance that subtly alters her space of possibilities for experience and action (where "action" here includes deliberating about what to do next, formulating intentions, and making decisions). Thus, in my view, *when potentially transformative substantial self-knowledge is at stake*, sticking amounts to an affective attitude that reflects not only certainty but also emotional insight (Bell and Leite 2016).

As I use it here, the term “emotional insight” comes from psychotherapy, where experts and practitioners distinguish between clients being able to know and recognise “in theory” *that* they have a certain character trait, or that they reacted in a certain way because they were experiencing a specific emotion, for example, and relating to that knowledge in the right way, i.e., truly understanding what it means for them (Britton 1998; Levine 2011). According to these authors, one cannot count as *fully* possessing self-knowledge in the transformative sense unless one is affected by it in the right way. Purely epistemic self-knowledge, or knowledge “in theory,” can allow us to explain our past actions and reactions, to describe our character perhaps, but it doesn’t have an impact on one’s motivations and actions, and therefore it affords no growth. From a therapeutic perspective, what allows the individual to develop and deal with their conflicts constructively is emotional insight (Bell and Leite 2016). Genuine, transformative, self-understanding requires a change of stance that includes the right emotional responses to what one has learnt about oneself and the motivations to act in consistent ways. This is why I argue that, in order to satisfy Katherine, the self-knowledge she acquires at the end of the process must include emotional insight. Without it, her self-ascription will have a reduced impact on her life and it will not stick.

In sum, when we search for substantial self-knowledge with a transformative potential, the stickiness of one’s answer includes two emotional components: an epistemic feeling of certainty and emotional insight. These are two attitudes that allow us to relate to our self-knowledge in an engaged and motivationally efficacious way. Certainty is a part of many purely epistemic endeavours and can do the important work of halting them before they consume excessive cognitive resources, but emotional insight is crucial to achieve stickiness when we are looking for opaque substantial self-knowledge, since it provides the self-ascription with a host of affective and motivational support beyond certainty.

In our search for substantial self-knowledge, there are any number of methods that can yield epistemically correct self-ascriptions (e.g. Lawlor’s causal self-interpretation based on inference, talking to friends, going to therapy, reading books about psychology, and so on), but not all of them are equally likely to lead to stickiness. Because I take this phenomenon to require emotional insight, I argue that a self-investigation method that centrally involves emotion is better positioned than many other methods to fit the bill. Sketching such a method will be my task for the rest of the paper.

4. Deep Desires

How can one achieve emotional insight in a case like Katherine’s? To answer this question, it is first necessary to say a bit more about the sort of desire Katherine wants to know about. First, she wants to know about a desire for something that comes very close to what L. A. Paul (2014) calls a “transformative experience”: an experience that radically changes its very subject at the epistemic and personal levels, rendering its consequences for one’s life unforeseeable. At the epistemic level, transformative experiences teach us something we cannot get to know before having them, like tasting a durian fruit for the first time (no amount of propositional knowledge about durians can teach me what I learn by tasting one). At the personal level, transformative experiences change who one is and one’s preferences in unpredictable ways. This means that, before undergoing such an experience, there is no way to know how it will change one and one’s preferences. I think that this lack of familiarity with the character of the experience and its potential consequences poses some specific challenges to getting to know one’s desire, making it more opaque in ways that might require more than inference to achieve self-knowledge (see Tooming 2020). But since here I focus on the existential rather than the epistemic aspect of the substantiality of self-knowledge, I’ll leave this issue aside. The point I want to highlight is that

Lawlor deliberately chooses to investigate the desire for an experience that has the potential to alter the subject's identity and sense of self by impacting on what she values. While having a second child might not be as radically transformative as having the first, it is still an event that deeply alters the family's life and what its members value, since now there is a new person to love and care for, and a new constellation of personal relations. In contrast with mundane desires for non-transformative experiences, like my desire for a cup of coffee when I wake up in the morning, desires like Katherine's involve experiences that can alter one's sense of self, one's life and one's future possibilities; in other words, they have an existential impact. If a desire for such an experience is to stick and satisfy one's inquiry, it must have a stability and depth that not all desires possess: shallow, volatile desires will not allow for the sticky self-ascription that Lawlor insists on.

No matter which account of desire one subscribes to (and there are several; see Schroeder 2020), these mental states can be volatile, superficial, and short-lived; they can come and go and quickly disappear, sometimes even despite not being satisfied. You might go to the supermarket thinking you want one thing for dinner, and then discard that desire and decide to shop for something else that seems more appealing. Katherine, by contrast, is looking for a self-ascription that prevents her self-investigation from being reopened over and over again. But a fickle desire that falters and dissipates, that doesn't consistently exert its motivational power in the appropriate circumstances, would seem to give one good cause to reopen the investigation on questions such as wanting to have a second child, or to change careers. In these cases, we seem to be talking about a motivation that is deeper and more persistent than the desire to eat a specific thing for dinner, a motivation that can sustain a long-term project. After all, the road from finding out one wants a child to deciding to try and actually trying to have one, to achieving a pregnancy, giving birth, and raising that child is a temporally extended process with many turning points at which one might change one's mind (Callard 2018, 34–35). A motivation that is able to sustain such a process cannot be vulnerable to whimsical changes; it requires a certain solidity, a certain depth, that everyday desires lack. Therefore, in my view, Katherine doesn't just want to know what she wants: she wants to know what she *really* wants, and this is crucial.

Now, you might retort, are you saying that it is impossible to experience a volatile desire for a child? Can't one experience such a desire for a short while and then change one's mind? This is indeed possible, but it isn't what Lawlor is after. Everything she says is compatible with Katherine experiencing an array of contradictory volatile desires about having a second child and knowing about them. Nonetheless, given their volatility, none of these desires would settle the question, because they won't allow any self-ascription to stick for too long. The amount of time and effort that Katherine devotes to investigating her question shows that she is not merely after what she finds appealing at a specific moment, but after what she *really* wants: her desire must have a dispositional character, it must be able to make its motivational power felt whenever the appropriate circumstances arise.

At this point, however, one might be tempted to object that Lawlor is wrong in talking about desires. Perhaps what is at stake in Katherine's search of self-knowledge is not a mere desire, but something else. The fact that actually carrying out the project of having a baby involves a stable source of motivation doesn't mean that this motivation must come from a desire. In practice, such projects require more than desire to fuel them: they require normative states like intentions and decisions, which imply personal and often also joint commitments (Gilbert 2013; Bratman 2014). So, the need for stickiness might indicate that Katherine is in fact trying to formulate her intentions through deliberation.

My reply here is that yes, in real life, intentions and commitments are surely necessary at some point in these processes. But if the desire to carry on with the project vanishes completely, this might constitute a good reason to revise one's commitment, so desires *per se* are not irrelevant. Furthermore, Lawlor explicitly states that finding out what one wants and deliberating about what to do are different things, and I think this is intuitively right (although the former might be an important ingredient of the latter). One might know full well that one wants another child, and finally decide to give up the project because of considerations that have nothing to do with one's desires (perhaps one has financial difficulties, or health problems, or one cares more about one's career). Still, the desire at stake here is not a volatile desire: it's a dispositional desire that one gives up at a cost. It is a deep desire. Just as emotions can be deep or shallow (see Pugmire 2005, ch. 2 on emotional profundity), I suggest that the same is true of desires. So, what does it mean for a desire to be deep?

Deep desires are more than dispositional desires: they are closely linked to one's sense of who one is, to one's self-relevant values and projects, in a way that not all dispositional desires are. After all, if chocolate ice cream is your favourite, this might mean that you have a reliable disposition to desire it in various circumstances. This doesn't make your desire for chocolate ice cream deep, though. Intensity alone doesn't cut it either: food cravings or sexual desires can be very intense without this indicating any psychological depth. For a desire to be deep, it has to be well integrated with other mental states and responses of the subject, especially their beliefs, their emotions, their life projects and what they value, all of which importantly contribute to making one who one is (Pugmire 2005, ch. 2).

Now, before I proceed, a few words on the notion of self at stake here are in order. Debates on the nature of selfhood and personal identity have occupied philosophers for centuries. Many dimensions of selfhood have been highlighted, among others, memory, embodiment, self-consciousness, the first-personal givenness of experience, sociality, recognition, or agency and freedom, i.e., the capacity to project oneself towards the future. This is no place to go into the details; here I simply want to highlight the dimension that I see as most relevant in cases like Katherine's, namely: freedom, agency, and what one values, or what some philosophers call one's "moral self" (see e.g. Prinz and Nichols 2016). What is more prominently at stake here is an effort to adjust or update one's sense of self (or self-concept) in a way that can inspire motivations and guide decisions about life projects. Thus, in talking about the "self" or "who one is," I mainly refer to the values that one endorses and that effectively steer one's motivations, actions, decisions, and projects. Now, this doesn't mean that what we value is a matter of free choice; on the contrary, external influences such as upbringing, culture and socialisation are essential. Neither do I wish to deny that self and identity are crucially shaped by factors beyond one's control, such as social norms, but these are not foregrounded in our search for engaged sticky self-knowledge. Our agency and what we value are not the full story of who we are, but they are crucial for the sort of stickiness Lawlor describes. While this might seem to bring me closer to the agentalist camp of the self-knowledge debate, I must emphasise that efficacious substantial self-knowledge is not just a matter of finding reasons to want, but of experiencing the right responses and motivations.

A deep desire, therefore, is one that is integrated with and expressive of the effective values and projects that are central to who one is, but it doesn't have to be phenomenologically very salient or intense—although of course it can and often will be. To see what this means, imagine a well-adjusted immigrant, who is satisfied with their situation in their host country, and who nonetheless harbours a desire to eventually return to their country of origin, as is often the case. Even if this desire is only felt saliently once in a while and in low intensity, it will still be in line with their beliefs about the place where they belong, the things that make them feel at ease, and so on, and will be constantly present in the background, shaping this person's entire horizon

of future possibilities. For example, this person might be unable to imagine themselves growing old in their host country, and they might be likely to give serious thought to circumstances that could facilitate their return. Given this link to other central elements of one's psychological life, deep desires are connected to one's sense of who one is and manifest something important about oneself. This, I contend, is the kind of desire that Katherine is looking for in her self-investigation.

Sticky self-knowledge of deep desires can be difficult to obtain because one is not clear about the way in which their object fits within one's projects, one's beliefs, one's emotions, and one's sense of who one is. Oftentimes, in cases like Katherine's, the affective or motivational states—call them desires, inclinations, or what you like—that are present at the beginning of the inquiry will lack the degree of depth that can allow for stickiness. Thus, one will need to engage in a process of exploring one's emotional responses and one's mental life in general in search of a sufficient degree of integration. Emotions are quintessentially evaluative states that can provide us with key insights into what we value, and this is why I see them playing a crucial role in cases like Katherine's.

5. Emotion, Imagination, and Knowing What One Really Wants

To see the role that emotions can play in achieving stickiness, it is important first to say a bit more about affective intentionality and the link between emotions and our cares and concerns. For the sake of simplicity, here I rely on Bennett Helm's formulations, since from the start he has aimed at offering a framework that explains the relation between emotions and our other evaluative attitudes, namely desires and evaluative judgments. According to Helm's holistic theory of affective intentionality, emotions are at the same time responses to and constitutive of personal value or, as he calls it, of import (see Helm 2001). In a nutshell, import is what matters to a particular person, the nexus that links together their rational evaluative responses. For Helm, emotions individually register import (like perceptions): they alert us to features of a situation in which something we care about is at stake. At the same time, they are collectively constitutive of import because, for Helm, caring about something just is evaluatively responding to it in the appropriate ways in the relevant situations, and emotions are quintessentially evaluative responses. Furthermore, import entails a preparedness to act, when relevant or necessary, in favour of that which one finds important, so the motivational component of emotions is explained.

Helm (2001) analyses the intentionality of emotional episodes by distinguishing between their targets and focuses. The target is the ostensible intentional object of the emotion, what it is thematically about. The focus is the background object one cares about that makes one's emotional response intelligible. For example, in hearing the news that my friend suffered an accident and is in hospital, I will get sad that they suffered injuries (this is the target of my sadness), because I care about my friend's wellbeing (the focus of my emotions). In hearing the same news about a stranger, my emotions are likely to be much less intense, or I might feel no emotion at all, revealing my indifference. In other words: without a focus of concern, no emotional response.

But there is more: the fact that I care about my friend rationally calls for my experiencing a range of emotions towards them in various situations: joy when I see them after a long absence, sadness when I hear about the accident, relief when I later learn that the injuries are not serious, and so on. This is rationally required because these responses are all entailed by my caring about my friend. Thus, emotions tend to form rational patterns that cluster around focuses of concern. But emotions are not the only responses required by these focuses: other evaluative responses, namely desires and evaluative judgments, are also a part of these patterns (for similar

views, see Callard 2018; Scheffler 2011). And, for Helm, predictably responding to something through such patterns is what constitutes its import for us. Caring about something just is predictably experiencing a range of evaluative responses that form rational patterns with that something as their focus. Note, however, that even though every focus of concern calls for a range of evaluative responses, such as desires and judgments, only emotions rely on focuses to be intelligible. Having a desire for a cup of coffee, or holding a judgment that democracy is a good form of government, doesn't rationally require that one cares about these things. But my fear about my friend's accident is unintelligible without my caring for them, so the link between emotions and caring is much tighter. Furthermore, barring cases of delusion, recalcitrance, or other forms of irrationality or gross error, every emotional response can be taken as a commitment to the import of its focus, as exerting a certain normative pressure to emote consistently with it in the future.

Now, emotional episodes are also situational—they arise as a response to changes in our immediate environments—which means that on occasion they will deviate from these patterns. For instance, I might fail to respond with the usual anticipation to the prospect of seeing my friend if I just got a call alerting me that my house is on fire. The fact that my emotional responses can on occasion be overridden by a competing concern doesn't mean that they are not required by my caring about their focus: the rational demand is still there, it is just that in this occasion I have a good reason to overlook it (just like there are valid excuses for other sorts of normative non-compliance). But if my emotions consistently diverge from this pattern for no reason (if whatever the circumstances I rarely ever feel joy when I see my friend or sadness when things go badly for them), I can justifiably be said not to care about my friend at all.

This is in line with what moral philosophers working in value theory have recently been proposing. According to Callard (2018, 64–65; see also Scheffler 2011), valuing has four essential and interconnected dimensions: affect, motivation, judgement, and self-monitoring. When I value something, I make myself emotionally vulnerable to it, as I just explained, following Helm. But valuing something also means being motivated to act in ways consistent with one's value (which includes experiencing certain desires and otherwise seeing relevant reasons for action), so my valuing my friend's wellbeing will lead me to desire to help them feel better after the accident, and as a result decide to come visit them with flowers at the hospital. Thirdly, valuing something means having the relevant beliefs and making the relevant judgements that something (in this case, our friendship) is valuable. And lastly, it means self-monitoring and endorsing one's emotional responses, motivations and judgements: evaluating them as appropriate. In self-monitoring, one might endorse or detach oneself from one or several of these aspects of valuing, revealing different sorts of internal conflict, but full-blown valuing entails an alignment of the four dimensions (Callard 2018, 65). How does this help us to get to know what we want in the hard cases?

5.1. Finding Out What One Really Wants

Recall that I've argued that Katherine is not investigating a desire *tout court*: she is investigating an epistemically opaque and deep desire, i.e., a desire that is well integrated with central aspects of who she is. Of course, deep desires are not always opaque: sometimes we find it easy to know that we want a second child or want to move to a different city. But when deep desires are opaque, one needs to investigate other aspects of oneself, especially what one values, in order to get to know them. My proposal—which takes Lawlor's as a starting point, but amends and expands it—is that this is done through counterfactual narrative imagination, i.e., imagining how one's life will unfold if one chooses one path rather than another.

For Lawlor, Katherine's imaginings are merely a source of evidence to ground inference. Nonetheless, according to Lawlor's description, Katherine not only examines her spontaneous internal promptings, but she also *deliberately prompts* further imaginings in order to explore her desire (2009, 57–58). This active and creative element tells against Lawlor's idea that what is going on is solely a detached cognitive process of inferring one has a mental state from its symptoms (i.e., internal promptings), or at least it suggests that more is going on than Lawlor realises. If Katherine herself is intentionally producing some of these imaginings, then she cannot *only* be using them as a basis for inferring their cause, since she already knows this: the cause is her intention to imagine. More importantly, for the sole purpose of inferring causes, what would be the advantage of prompting again something that she already knows appeared in her mind? What would be the advantage of re-imagining things—with all the phenomenal richness that imagination offers—rather than simply recalling *that* one imagined something in the past? Yet Lawlor is right in this: in such situations, many of us deliberately imagine how things might turn out if we take one course of action rather than another. In my view, this makes sense because we do more than infer causes: we are exercising our capacities for counterfactual imagination to explore our motivational and valuational psychology. Thus, I propose that a large part of what Katherine does when she investigates her mental life is to monitor her emotional responses to the various scenarios she imagines; in other words: she is testing her appraisals of those situations and thereby exploring the evaluative structure around any focuses of concern that might be connected to her question. This is what can tell her which one of her potential competing desires can go deep enough to stick.

So, what role does imagination play in helping Katherine get to know her deep desire? To answer this question, let me first offer a brief account of imagination that draws on the phenomenological tradition (Sartre 1986; Casey 2000; Szanto 2017). Imagination, like perception, memory, belief, and so on, is one of the modes in which objects can be present in experience. In contrast to perception, imagination is characterised by making present to us objects that are absent. Furthermore, acts of imagination do not represent (or “posit”) their intentional objects as real, but the reality status of objects is suspended. This doesn't mean that when we imagine, we always experience things as unreal: this is only one possibility. Rather, in imagination, objects can be posited in four characteristic ways: (1) as fictional or non-existent, as when you imagine a character from a novel; (2) as absent or non-actual, for example when, on your way to hospital, you imagine how your friend might look after the accident; (3) as existentially neutral, neither existent nor non-existent; and (4) as pure possibilities, as when Katherine imagines herself with a new baby in her arms. Furthermore, imaginings are highly sensorial or quasi-sensorial phenomena, which means that not only can they present emotions as objects of imagination (e.g., I can imagine myself feeling joy), but they can also easily trigger affective responses (e.g., my imagined happiness can make me feel joy right now). In other words, in imagining, we “perceive” in as-if mode and respond emotionally as if imagined objects and events obtained (Casey 2000). Importantly, whichever view one takes on the status of emotional responses to fiction or imagination, these responses are not illusory. They are actual emotional responses, actual engagements of our valuational psychology with the imagined stimuli, which means that they carry crucial information about what we care about. This is why imagination allows for an exploration of one's cares and concerns.

Thus, imagination is particularly suited to arriving at sticky self-ascriptions of epistemically opaque deep desires *precisely* because it is so evocative of emotion, not (only) because it might offer a basis for inferring causes. Indeed, research in empirical psychology suggests that exploring our current emotional responses to imagined scenarios is a better route to self-knowledge than rationalising methods like lists of pros and cons, which are third-personal and detached (Wilson 2004, 174, 177–78). I propose that an important reason for this is that imagination allows us to gauge the depth of our responses to temporally extended events: how and to what extent our desires, emotions, and beliefs around a specific object or situation are integrated with each other.

The strength of this method doesn't rely on its power to accurately predict the impact of a certain event or decision on our lives, nor our emotional responses to it: any endeavour in this regard will reveal itself importantly limited (see Paul 2014; Tooming and Miyazono 2023). In other words, the point is not to get our imagined emotions—the emotions that are a part of the represented content—right. Rather, the point is engaging our evaluative responses, putting our current emotions to work to inform us of what we care about now, of whether our current structure of cares and concerns impels us to project ourselves forward in one direction or another. Exploring our ongoing emotional responses to imagined scenarios can indicate whether the specific path we're envisaging is in consonance with our deep motivations and values (Goldie 2012, ch. 4). One single instance of emotion is not likely to be enough to settle on a sticky self-ascription of desire, but if one can discern, or even more inchoately sense, a pattern of emotions and other evaluative responses around a particular focus, one will gain crucial insight into what matters to one, in which way and to what extent. Eventually, the process will allow one to see whether the desire under investigation fits or not within the pattern, and to articulate a desire that fits. A desire that fits will be deep: it will have a stronger motivational power and be more stable than a desire that doesn't, simply because one's entire valuational structure will act to support it. But what if the answer is still not clear, or if no desire fits well enough to count as what one *really* wants?

Recall that emotions are not just responses, they are simultaneously commitments to the import of their focuses (they rationally call for consistent evaluations of that focus). Thus, activating our emotional responses can help us reduce the number of options that present themselves to us as serious, meaningful possibilities. When we get clear on what we value, the courses of action that go against it are less likely to appear to us as worth considering and, in the limiting cases, they won't even show up as possibilities for us (Goldie 2012, 91–92). As I suggested above, Katherine's problem is not that she cannot know her shallow, volatile desires. The problem is that, in order to get to the deep, sticky one, she needs to demote or discard some of them, to reduce the number of conflicting desires she experiences, or at least their pull. Imagination helps us do this not by willingly repressing what does not concur with the result of our logical deliberation, but by activating our valuational system to reinforce what we care about. In replaying different counterfactual scenarios, one will experience emotions, desires, and so on connected to them, thus re-enlivening one's commitment to the import of their focuses. Over time, one's responses will exert their normative pressure and lead one to discard certain scenarios. Thus, over time one of them is likely to take precedence; in other words, the favourable pattern of emotions and other evaluative responses attached to it will get stronger and clearer than others, and with it the desires that fit within it (Goldie 2012, 92–93).

Achieving emotional insight on what one really wants entails finding a sufficiently clear pattern where some desire(s) fit, and some have no place. If the pattern is clear and solid, desires, emotions, and judgements that contradict it will be rare (or significantly less frequent than those that fit the pattern), weak, or even non-existent. But the complete absence of contradicting evaluative responses is not necessary for one to know what one really wants and for that knowledge to play its practical role in our lives: the pattern might be strong, and one might still experience contradictory emotions and desires on occasion. The point is: over time one pattern will come to dominate, with discordant responses standing out as dissonances. Accordingly, in this process our options will be narrowed down to the one or few that better express our personal values, because our evaluative responses will gradually steer us away from considering certain paths that go against them as actual possibilities. This combination of fewer options with the support of other evaluative attitudes is what will result in emotional insight, which in this view is not mere epistemic certainty that a self-ascription is correct, but the experience that one's evaluative attitudes concerning the matter at hand are sufficiently well aligned and predominantly point in the same direction. This is what will prevent the question from being reopened and confer stickiness to the answer.

5.2. Worries and Objections

Let me now address a few worries that might have arisen for readers. Firstly, my claim is that emotions are crucial to achieving sticky knowledge of what one really wants in the hard cases. But emotions are not always easy to identify and understand. Sometimes we don't know how we feel. We might have difficulties labelling our affective responses or feel several emotions at the same time. Sometimes certain emotions are so unpleasant or difficult to acknowledge, perhaps because they have been stigmatised in our upbringing or in our culture, perhaps because of what they might reveal about ourselves, that we repress them or transform them. And as empirical research has shown, thinking about the reasons why we experience a particular emotion can push us to confabulate and rationalise (Wilson 2004, chs. 7, 8).

In response to this, I readily acknowledge that affective self-knowledge is fallible, but that does not mean it is useless or irrelevant. First, emotions are not the only mental processes that can occasionally be unreliable, suspect, or hard to know. The fact that we cannot always easily identify which emotion we are experiencing does not universally disqualify emotions as sources of (self-)knowledge (see Brady 2013). Second, and most importantly, my claim is not that one single emotional episode holds the key to knowing what one really wants and to relating to that knowledge in the right way. Since evaluative responses cluster around focuses of concern, it is the *patterns* that will be truly informative. When emotional episodes reliably form a pattern around a focus of concern, this means that one does care about it in the way that is consistent with the pattern. Furthermore, when seen in light of the pattern, single episodes of emotion (or desire, for that matter) can be informative too: do they fit within the pattern or deviate from it? Why? What does that mean in terms of their depth? (Thomason 2023). The patterns that our emotions form therefore constitute a direct route to knowing what matters to us.

To this you might retort: emotions are often messy and contradictory. When investigating her desire for another child, Katherine might find herself simultaneously longing for the feeling of a baby in her arms and scared about what maternity leave might do to her career. Or, perhaps more contradictorily, she might find joy in the prospect of playing with the baby and seeing it develop, and simultaneously cringing at the prospect of diaper changes and sleepless nights. Where is the pattern there? My answer is that talk of rational patterns might make things seem more rigid than they are. The pattern, the caring, is not prior to the evaluative responses or invariable. It evolves as subjects evolve and it can change and adapt over time. Besides, people typically don't have one single thing we care about; we have many, but we don't care about all of our concerns with the same intensity and depth, and some of these concerns can be in conflict. Sometimes we might be genuinely ambivalent about something. Our emotions will reflect this. Think about two people who for different reasons decide not to have children. There is a difference between the person who always knew they didn't want to procreate and goes through life without second-guessing this intention, and the person who at some point felt a desire to have children but decided against it because it conflicted with other important cares and concerns of theirs.

As for Katherine's contradictory feelings about babies, one can see this as two contradictory patterns of evaluation that co-exist, with none of them clearly dominating over the other. Katherine would be breaking up the situation "caring for a baby" into different aspects and evaluating each of them separately and consistently: she loves playing but hates changing diapers. However, if we see this as a single pattern, what we have is ambivalence. And nothing in this account requires patterns of emotions to be clear-cut and pure: a pattern of ambivalence is also a pattern, and what one learns by paying attention to it is precisely that one is ambivalent. This might not be a completely satisfactory or easy to handle piece of self-knowledge, it might

not be knowledge of a deep, sticky desire, but it is self-knowledge nonetheless. Some people might be content with the conclusion that they are ambivalent; others, perhaps most, might be uneasy with this situation and keep re-starting the inquiry until they manage to shape a sticky desire (see Lawlor 2009, 57–58). Whether the end result is satisfactory or not, exploring the whole range of the evaluative attitudes that are relevant to the question, especially our emotions, remains a crucial path to getting to sticky substantial self-knowledge.

6. Conclusion: Emotional Insight and the Transformative Potential of Self-Knowledge

In this paper, I have argued that emotions should have a larger role in the philosophy of substantial self-knowledge than has been assigned to them so far. On the one hand, according to philosophers of emotion and theorists of value alike, paying attention to our emotions is crucial for gaining insight into what we care about and value. This knowledge is essential for us to shape and get to know our epistemically opaque deep desires. On the other hand, in order to find answers to difficult questions about substantial self-knowledge, such as the one depicted in Lawlor's Katherine example, it is not enough to acquire epistemic self-knowledge. For a desire self-ascription to stick, for us to acquire a belief that can play the desired roles in our lives, an affective change of stance, reflecting emotional insight, has to take place. Emotional insight goes over and above epistemic self-knowledge of one's desire or emotion. While the latter merely consists in a grounded belief about one's mental state, the former also encompasses an understanding of that state's significance and implications and the appropriate emotional responses to what one knows. Emotional insight thus involves a self-reconciliation that allows self-knowledge to stick and imbues it with its motivational, transformative potential. Halting the philosophical investigation at the point where one achieves detached self-knowledge, as Lawlor does in her paper, without further analysing what "sticky" self-knowledge requires, thus represents a missed opportunity to illuminate the profound, reciprocal interaction between self-knowledge and emotional insight. When we invest so much time and effort in the pursuit of substantial self-knowledge as Katherine does, we typically do so in pursuit of something more than a correct belief: we seek meaningfulness and transformational power. A route that engages our emotions in the investigation process is particularly conducive to this end.

Acknowledgements

I thank Alessandro Salice, Simon Høffding, Uku Tooming, Kengo Miyazono, Rui Kubota and the anonymous reviewers for their extremely helpful feedback on previous drafts of this paper. This paper has been presented at academic events in Graz, Leipzig, Cork, Copenhagen and Madrid. I thank the audiences for their insightful comments. The research leading to this piece has been funded by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 890316, the Austrian Science Fund (FWF) project P32392-G, and the Autonomous Government of Madrid (Spain), Talent Attraction grant number 2022-T1/HUM-23931.

References

- Bar-On, D. 2004. *Speaking My Mind: Expression and Self-Knowledge*. Oxford University Press.
- Bell, D. and A. Leite. 2016. "Experiential Self-Understanding." *The International Journal of Psychoanalysis* 97 (2): 305–32.
- Boyle, M. 2015. "Critical Study: Cassam on Self-Knowledge for Humans." *European Journal of Philosophy* 23 (2): 337–48.
- Brady, M. S. 2013. *Emotional Insight: The Epistemic Role of Emotional Experience*. Oxford University Press.
- Bratman, M. 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press.
- Britton, R. 1998. "The Suspension of Belief and the 'As-If' Syndrome." In *Belief and Imagination*, edited by R. Britton, 59–68. Routledge.
- Byrne, A. 2011. "Knowing What I Want." In *Consciousness and the Self: New Essays*, edited by J.L. Liu and J. Perry, 165–83. Cambridge University Press.
- Callard, A. 2018. *Aspiration: The Agency of Becoming*. Oxford University Press.
- Candiotto, L. 2019. *The Value of Emotions for Knowledge*. Palgrave Macmillan.
- Casey, E. S. 2000. *Imagining: A Phenomenological Study*. 2nd ed. Indiana University Press.
- Cassam, Q. 2014. *Self-Knowledge for Humans*. Oxford University Press.
- Corbí, J. E. 2023. "Distinctive Substantial Self-Knowledge and the Possibility of Self-Improvement." *Synthese* 201 (6): 201.
- De Sousa, R. 2008. "Epistemic Feelings." In *Epistemology and Emotions*, edited by G. Brun and U. Doguoglu, 185–204. Routledge.
- Doyle, C. 2022. "Knowing Your Mind by Making Up Your Mind Without Changing Your Mind, Too Much." *Journal of Philosophical Research* 47 (November): 133–46.
- Finkelstein, D. H. 2008. *Expression and the Inner*. Harvard University Press.
- Gertler, B. 2018. "Self-Knowledge and Rational Agency: A Defense of Empiricism." *Philosophy and Phenomenological Research* 96 (1): 91–109.
- Gilbert, M. 2013. *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Goldie, P. 2012. *The Mess Inside: Narrative, Emotion, and the Mind*. Oxford University Press.
- Helm, B. W. 2001. *Emotional Reason. Deliberation, Motivation, and the Nature of Value*. Cambridge University Press.
- Jongepier, F. 2021. "The Value of Transparent Self-Knowledge." *Ethical Theory and Moral Practice* 24 (1): 65–86.
- Lawlor, K. 2009. "Knowing What One Wants." *Philosophy and Phenomenological Research* 79 (1): 47–75.
- Levine, H. 2011. "'The Consolation Which Is Drawn from Truth': The Analysis of a Patient Unable to Suffer Experience." In *Bion Today*, edited by C. Mawson, 188–211. Routledge.
- Moran, R. 2002. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.
- Morton, A. 2009. "Epistemic Emotions." In *The Oxford Handbook of Philosophy of Emotion*, edited by Peter Goldie, 385–400. Oxford University Press.
- Nguyen, C. T. 2021. "The Seductions of Clarity." *Royal Institute of Philosophy Supplements* 89 (May): 227–55.
- Paul, L. A. 2014. *Transformative Experience*. Oxford University Press.
- Prinz, J. J. and S. B. Nichols. 2016. "Diachronic Identity and the Moral Self." In *The Routledge Handbook of Philosophy of the Social Mind*, edited by J. Kiverstein, 449–64. Routledge.
- Pugmire, D. 2005. *Sound Sentiments*. Oxford University Press.
- Renz, U. 2017. "Socratic Self-Knowledge in Early Modern Philosophy." In *Self-Knowledge: A History*, edited by U. Renz, 146–63. Oxford University Press.
- Sartre, J.-P. 1986. *L'imaginaire: Psychologie Phénoménologique de L'imagination*. Gallimard.
- Scarantino, A. and R. de Sousa. 2021. "Emotion." In *The Stanford Encyclopedia of Philosophy*. Edited by E. N. Zalta. Summer 2021 edition. <https://plato.stanford.edu/archives/sum2021/entries/emotion/>.
- Scheffler, S. 2011. "Valuing." In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*, edited by R. J. Wallace, R. Kumar, and S. Freeman, 23–42. Oxford University Press.
- Schroeder, T. 2020. "Desire." In *The Stanford Encyclopedia of Philosophy*. Edited by E. N. Zalta. Summer 2020 edition. <https://plato.stanford.edu/archives/sum2020/entries/desire/>.

plato.stanford.edu/archives/sum2020/entries/desire/.

Sorgiovanni, B. 2019. “The Agential Point of View.” *Pacific Philosophical Quarterly* 100 (2): 549–72.

Szanto, T. 2017. “Collective Imagination: A Normative Account”. In *Imagination and Social Perspectives*, edited by M. Summa, T. Fuchs, and L. Vanzago, 223–46. Routledge.

Thomason, K. K. 2023. “Alienated Emotions and Self-Knowledge.” In *Emotional Self-Knowledge*, edited by A. Montes Sánchez and A. Salice, 39–55. Routledge.

Tooming, U. 2020. “Being Familiar with What One Wants.” *Pacific Philosophical Quarterly* 101 (4): 690–710.

Tooming, U. and K. Miyazono. 2023. “Affective Forecasting and Substantial Self-Knowledge.” In *Emotional Self-Knowledge*, edited by A. Montes Sánchez and A. Salice, 17–38. Routledge.

Wilson, T. D. 2004. *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Belknap.



This work is licensed under a Creative Commons Attribution 4.0 International License.