

What is the role of affective forecasting in knowing what we value?

Diana Craciun

To cite this article: Diana Craciun (19 Sep 2024): What is the role of affective forecasting in knowing what we value?, Philosophical Psychology, DOI: [10.1080/09515089.2024.2403509](https://doi.org/10.1080/09515089.2024.2403509)

To link to this article: <https://doi.org/10.1080/09515089.2024.2403509>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 19 Sep 2024.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

What is the role of affective forecasting in knowing what we value?

Diana Craciun

Philosophy Department, University College London (UCL), London, UK

ABSTRACT

Generally, we confidently ascribe valuing states to ourselves. We make statements such as “I value democracy” or “I value my best friend” - our sense of who we are depends on doing so. Yet what justifies that confidence? If you were asked “Do you value philosophy, or are you just doing it for the money?”, how might you go about generating such knowledge? I will operate with the notion that valuing involves, at a minimum, a set of distinctive emotional dispositions toward the valued object. Given this view of valuing, a plausible explanation stemming from Tooming and Miyazono is the following: to know whether we value, we need to predict our potential or future emotional reactions. That is, they suggest that affective forecasting is *necessary* to generate knowledge of valuing states. I argue that, despite its intuitive appeal, affective forecasting is not necessary. I then consider a more modest claim, namely that affective forecasting is still explanatorily powerful when it comes to self-knowledge of valuing states, as it explains the difficulty of generating such knowledge. I reject this further claim, arguing that there is a stronger available explanation: the use of theories about valuing.

ARTICLE HISTORY

Received 17 July 2024

Accepted 6 September 2024

KEYWORDS

Valuing; self-knowledge;
affective forecasting;
emotions; theory of mind;
theory theory

Generally, we confidently ascribe valuing states to ourselves – and it matters to us that we can. We make statements such as “I value kindness in a partner”, “I value literature”, “I value democracy,” “I value my best friend” and so on. Our sense of who we are depends on doing so. What justifies that confidence, especially given instances in which we simply do not know whether we value something? We often face questions such as: “Are you sure you value our friendship? It doesn’t seem like it” or “Do you value philosophy or are you just doing it for the money?”. How do we go about coming to know whether we value something?

CONTACT Diana Craciun  cristina.craciun.15@ucl.ac.uk  Philosophy Department, University College London (UCL), 36 Quadrant Close, The Burroughs, London NW4 3BY, UK

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

To understand how a subject has such knowledge, we need first to understand what valuing is. Following a prevalent trend in the literature (Scheffler, 2010; Seidman, 2009), I will operate with the notion that valuing involves, at a minimum, a set of distinctive emotional dispositions toward the valued object. For example, valuing an object will involve being disposed to feeling joy upon interaction with it, or to feeling excitement at the prospect of interaction with it, or to feeling worried if the object is in danger. To be clear, I will here focus on valuing something in itself, rather than valuing it instrumentally. So, I will not be concerned with whether we value a friend because they can provide e.g., social benefits, but rather whether we value them in themselves.

Understanding how we have knowledge of what we value is not easy; generating knowledge of such states seems, *prima facie*, to face several potential difficulties (liability to biases, self-deception, the desire to maintain a certain self-image and so on), to be potentially corrigible (others might be better placed to know whether we value something), and to therefore be difficult to come by. Unlike knowing whether your leg is itchy, knowing whether you value something poses several epistemological challenges.

How, then, do we gain knowledge of valuing states? In a recent paper, Tooming and Miyazono (2023) suggest that such self-knowledge will require what they call “affective forecasting”. Affective forecasting is, roughly, predicting our future emotional states. Their claim is that affective forecasting is *necessary* to know that one values something.¹ They use the fact that we are generally not good at affective forecasting, relying on recent empirical studies,² to explain why such knowledge is generally harder to come by than more trivial forms of self-knowledge. Knowing our values, for them, poses the evidential challenges of generating evidence about the future.

In this paper, I reject the idea that knowing what we value *requires* affective forecasting. Instead, I suggest there is something else that is required: having correct theories about valuing. To that aim, I will proceed by clarifying the notion of valuing in [section 1](#). In [section 2](#), I will contextualize the discussion by explaining Cassam’s distinction between substantive and trivial self-knowledge, which Tooming and Miyazono make use of. In [section 3](#), I outline Tooming and Miyazono’s argument for why affective forecasting is necessary for some types of substantive self-knowledge and valuing in particular. I reject their argument in [section 4](#).

Having rejected the idea that affective forecasting is necessary for knowing what we value, I discuss whether affective forecasting might still explain, as Tooming and Miyazono claim, why typically, knowledge of what we value is harder to come by. I reject this more modest claim as well, arguing that there is a better alternative to explaining the obstacle to knowing what we value that they overlook. I will conclude, in [section 5](#), that the use of theory

by the subject figuring out her values provides an equally common, typical account of the obstacle, and moreover, is an element that is necessarily involved in all cases of coming to know what we value.

1. What is valuing?

First, what is valuing? Recent trends in the literature seem to suggest that, among other things, valuing involves having a set of emotional disposition toward its object. I begin by outlining one account that I find most promising, namely Seidman's (2009).

Seidman's starting point is Jaworska's account of caring. On Jaworska's account, caring is:

A structured compound of various less complex emotions, emotional predispositions, and also desires, unfolding reliably over time in response to relevant circumstances. Typical components of caring include: joy and satisfaction when the object of one's care is flourishing and frustration over its misfortunes; pride in the successes of the object of care and disappointment over its failures; the desire to help ensure those successes and to help avoid the failures; fear when the object of care is in danger and relief when it escapes unharmed; grief at the loss of the object, and the subsequent nostalgia. (Jaworska, 2007, p. 560)

Thus, on this account caring about something is to be understood as having certain emotional dispositions toward an object. This, Jaworska notes, is not enough. The reason is that there might be patterns of emotional dispositions toward something that are wholly separate from one's caring about it (e.g., one's irritation at one's father's humming). She therefore introduces the idea that what links the relevant emotional networks is the fact that the subject sees the object as important. Seidman (2009) wholeheartedly agrees with this suggestion; however, he argues that we need to arrive at a precise sense for "seeing as important". This seems right, after all "seeing an object as important must not be just valuing it" on pain of circularity.

So, his aim is to flesh out the cognitive aspect of caring by clarifying the notion of "seeing as important". According to him, caring necessarily involves:

seeing certain X-related considerations as reasons (with the kind of considerations that are relevant varying with the kind of object X is). These "seeings-as" dispose an agent to believe relevant considerations to be reasons; and where no countervailing considerations defeat these dispositions, the subject will do so. (p. 290)

For Seidman, then, caring involves both emotional and cognitive dispositions and seeing the object as important involves cognitive ones.

With Jaworska's view of caring in mind, he argues in his paper that valuing turns out to be just caring. I am not going to consider those arguments; I remain agnostic regarding whether valuing just is caring, having introduced

this account for the purposes of illustrating a contemporary view of valuing. Further, I am not committed to his account being correct across the board. Rather, I am going to adopt the minimal claim that valuing involves a set of emotional dispositions.³ This is an aspect that his view shares with several others in the literature (e.g., Anderson, 1993; Helm, 2001; Scheffler, 2010; Tooming & Miyazono, 2023). This will be the minimal assumption I make about the nature of valuing for the purposes of my discussion.

Yet before moving away from this section, I will briefly discuss a dominant alternative account of valuing. Doing this will make it clear that even if we reject my minimal assumption about valuing, there is scope to be concerned with whether affective forecasting is involved in self-knowledge with respect to what we value. Contrary to accounts that make valuing an essentially or partially affective matter, there are theories that construe valuing as a solely intellectual attitude toward some object. An illustrative example of such an account would be Scanlon's, who defines valuing as follows: "To value something is to take oneself to have reasons for holding certain positive attitudes toward it and for acting in certain ways in regard to it" (Scanlon, 1998, p. 95). Crucially, emotional dispositions are not constitutive of valuing for him – rather, we should focus on what the subject takes themselves to have reasons for.

Even if one prefers a Scanlonian approach, one might still agree that the epistemological question of generating knowledge of what we value might be answered via the notion of affective forecasting, even though the metaphysical question of whether emotional dispositions (at least partly) constitute valuing may have a negative answer on Scanlonian accounts. That is, the two questions may come apart. One might push back at this stage: how, then, do the metaphysical and epistemological questions interact? Would it not be odd if affective forecasting was *required* for knowledge even though emotional dispositions are not constitutive of valuing? Yet note that, despite being counterintuitive, it is not an entirely implausible suggestion. Consider a parallel: perhaps the only way to know that somebody else is in pain is via their behavior, but behavior is not constitutive of pain. In that case, the epistemological and metaphysical questions are clearly distinct.

The following dichotomy seems to emerge: either we are intellectualist⁴ about valuing, or we are not. On both counts, it makes sense to ask about the role of affective forecasting for generating knowledge of that state. However, you might think that, if one were an intellectualist about valuing, then whatever role affective forecasting may play, it would be harder to argue that it is necessary. Or, at the very least, that it would make sense for Tooming and Miyazono to argue for their position on an affective, rather than intellectualist view of valuing, as it would be the stronger position. So, I will grant Tooming and Miyazono an affective view of valuing (which I also in fact endorse). I will then show that, even on such an account, affective forecasting cannot be necessary for self-knowledge of our valuing states. Since,

then, affective forecasting cannot be necessary on their approach to valuing (which is stronger than the alternative), it is not necessary *tout court*. Before diving into that argument, however, I will clarify the relevant notions of self-knowledge.

2. Substantive self-knowledge and trivial self-knowledge

Prior to arguing that affective forecasting is not necessary for knowledge of our valuing states, we should first integrate this type of knowledge within a broader understanding of self-knowledge, which is the framework of Tooming and Miyazono's discussion. This detour will also provide the necessary context for the weaker claim they make (namely that affective forecasting explains the overall difficulty of knowing what we value, even if it is not necessary), which I reject in [section 5](#).

The framework Tooming and Miyazono operate in is based on Cassam's (2014) distinction between "trivial" self-knowledge (TSK) and "substantial" self-knowledge (SSK). Standard examples of the former include knowing that one believes it is raining, that one intends to go to the kitchen, that there is an itchy sensation on one's leg, that one sees trees and so on. The latter, substantive self-knowledge, consists of knowledge of what we value, of our character traits, aptitudes, emotions, and so on. The two are not strictly separate, rather being conceived on a spectrum with cases in the middle.

Cassam suggests that the distinction between trivial and substantive self-knowledge is a matter of degree, not of kind (2014, p. 29). So, he offers a set of criteria that are meant to help differentiate between the two categories. These include: "The Fallibility Condition" (there is always the possibility of error with SSK), The Obstacle Condition (the possibility of error is linked with the presence of several potential obstacles such as self-deception, repression, bias, embarrassment), The Corrigibility Condition (SSK is corrigible, which is connected to the fact that we may not be the most authoritative source regarding our own substantial states, e.g., a spouse may be better placed), The Non-Transparency Condition (SSK cannot be arrived by via the Transparency Method, i.e., you cannot find out whether you are kind by considering whether you ought rationally to be kind), The Evidence Condition (SSK is based on evidence, and what counts as evidence will depend on the type of substantive state), The Cognitive Effort Condition (obtaining SSK requires cognitive effort, i.e., you don't "just know" whether you are kind), The Indirectness Condition (SSK is both psychologically and epistemically mediate), etc. (pp. 30–32).

Self-knowledge of valuing, then, seems to be a particularly good example of substantive self-knowledge. It is indirect, it requires effort, it can be reasonably challenged etc. One might therefore argue that, since knowledge of our valuing states is an instance of substantive self-

knowledge, the answer to how we might acquire the former is in some sense to be derived from our understanding of knowledge of the latter. It is in this context that the notion of affective forecasting comes in: to explain why substantive self-knowledge is systematically more difficult to come by than trivial self-knowledge, Tooming and Miyazono argue that the difference lies in the evidential state of the subject. They note that substantive self-knowledge tends to require affective forecasting, though they are not claiming it is necessary in every case. Instead, they argue that, within substantive self-knowledge, we can trace a division between “strong” and “weak” substantive self-knowledge. They talk of “strong” substantive self-knowledge when affective forecasting is required, such as knowledge of what we value or of character traits such as kindness. On the other hand, “weak” substantive self-knowledge does not, for them, require affective forecasting, for instance knowledge of aptitudes such as being able to speak Spanish.

So, we might be able to explain why strong substantive self-knowledge is overall more difficult to come by when there is a need for affective forecasting. The fact that strong substantive self-knowledge satisfies most of the criteria outlined by Cassam could be explained by affective forecasting. Gaining knowledge of our valuing states, one might argue, relies on an indirect process, it requires effort, it faces significant obstacles, and so on, because it requires use of affective forecasting, which we are not particularly adept at (Wilson & Gilbert, 2005). Thus, affective forecasting emerges as a source of significant obstacles for obtaining most types of substantive self-knowledge, particularly valuing. Contrary to this view, I will argue that not only is affective forecasting not necessary for knowing what we value, but it also fails as an explanation for the difficulty of strong substantive self-knowledge as a whole. This is because, as I argue in [section 5](#), there is a stronger alternative they overlook (that is explanatorily prior to affective forecasting).

For now, however, I will focus on the claim that, at least in the case of valuing (which is their primary example), affective forecasting is necessary for knowledge. I will now outline their view, which I reject in [section 4](#).

3. The case for affective forecasting

Tooming and Miyazono (2023) claim that the substantiality of a case (its location on the trivial to substantial spectrum) depends on the kind of evidence needed. As they put it, “As a first approximation, case X is more substantial than case Y when an agent in X needs to possess the kind of evidence that an agent in Y does not need to make a knowledgeable self-attribution” (p. 20). A further crucial point is that the relevant evidence, i.e., the evidence that contributes to the substantiality of an attitude, is about one’s *future and counterfactual*

affective reactions. They explain their view by reference to Lawlor's (2009) case about the self-knowledge of desire, as Cassam himself also engages with it in his discussion about substantive and trivial self-knowledge.

So, let us proceed with Lawlor's example, which focuses on Katherine, a woman who is trying to figure out whether she wants another baby. Lawlor outlines Katherine's total evidence, and hints at how it might justify the inference to the presence of a desire. Tooming and Miyazono use this example as a starting point and then move on to knowledge of valuing states. I will follow their dialectic in what follows.

So how will Katherine find out what she wants? Now that the question has been called, Katherine starts noticing her experiences and thoughts. She catches herself imagining, remembering, and feeling a range of things. Putting away her son's now-too-small clothes, she finds herself lingering over the memory of how a newborn feels in one's arms. She notes an emotion that could be envy when an acquaintance reveals her pregnancy. Such experiences may be enough to prompt Katherine to make a self-attribution that sticks. Saying "I want another child", she may feel a sense of ease or settledness. (Lawlor, 2009, pp. 12–13)

Based on this example, Lawlor concludes the following: "In sum, it seems that causal self-interpretation, specifically, inference from internal promptings, is a routine means by which we know what we want. In some cases, the internal promptings are simple sensations; in other cases, where desires are more complex, one's internal promptings may include imaged natural language sentences and visual images (which in turn may figure in specific kinds of imaginative rehearsal)" (p. 17). As she puts it, these internal promptings serve as evidence for the relevant inference. Of course, the existence of internal promptings is not enough, as the self-attribution needs "to stick". We need not go into more depth at this stage.

Tooming and Miyazono introduce an elaboration of the Katherine case. Instead of asking herself whether she desires having another baby, what she now wants to know is whether she *values* having another child. In this variation, Katherine concludes *based on the same evidence* that she does value it.⁵ However, upon having another baby, she distressingly experiences none of the joy anticipated. Instead, she feels disappointed and regrets her choice. "That's not what I wanted", she concludes. At least insofar as her affective responses are concerned, her imaginings (and further evidence about the present such as jealousy) seem to have clearly missed the mark, point out Tooming and Miyazono.

This is the crux of their argument: "When Katherine turns out to be disappointed at having another child, this indicates that her evaluative attitudes toward that prospect are fickle: before giving birth, she seemed to value having another child, while Katherine's subsequent affective responses indicate that her valuing is at most only surface-level or that

she did not value it in the first place. *It makes sense to say that the disappointment defeats her earlier belief that she valued having another child*“ (p. 22, my emphasis).

First, this relies on a view of valuing like the one discussed in [section 1](#), i.e., one that considers dispositions for emotional reactions toward the valued object as constitutive of valuing. Based on this starting point, if her affective dispositions do not match the ones required of valuing (she does not exhibit valuing dispositions and exhibits non-valuing ones such as disappointment), then, in their view, this gives her reason to doubt her previous belief that she values having another baby. In other words, by “defeat” they mean she is no longer justified in holding that belief, not necessarily that it is false. I discuss defeat more in [section 4.1](#). This, they argue, shows that evidence about present mental states is not sufficient for knowing what one values. To know that one values, they conclude, some other type of evidence must be necessary: evidence about one’s future affective reactions.

To see why they think evidence about the future is the best alternative, consider what they later state: “In order to know one’s affective dispositions one has to know what affective responses one would have across a large variety of possible circumstances. One therefore needs to predict those responses, i.e., one needs to resort to affective forecasting in which, as the data indicate, people tend to fail“ (p. 25). In other words, knowing whether one values is, to them, knowing whether certain dispositions will manifest in future circumstances. So, Katherine needs to be able to predict how she would react in relevant circumstances, such as having a baby. Her evidence about present mental occurrences is simply not enough. I will now outline the steps of their argument in more depth, with an eye to revealing what I take to be its significant flaws.

If Katherine values, then, she should have a set of affective dispositions (say, dispositions d_1, d_2, \dots, d_n). At T_1 , based on a set of evidence E_1 (evidence about her present states generated from internal promptings), she concludes that she values. Nevertheless, her later affective reactions (at T_2 , after having a baby) give her a reason to think that she was mistaken at T_1 . This is because her current reactions indicate that she does not have the affective dispositions that are required for valuing. Instead of experiencing d_1, d_2 , etc., she finds herself experiencing the opposite of these dispositions. This, then, shows that she did not have sufficient evidence at T_1 , and so that evidence about present mental occurrences sometimes does not suffice.

Since evidence about present mental occurrences does not suffice, then some other evidence must be necessary. They do not think evidence about past mental events can be sufficient, as I discuss in [section 4.3](#). The only remaining option, if we eliminate evidence about the present and the past, is evidence about the future. Knowledge of one’s future affective states is, in their view, sometimes necessary for substantial self-knowledge. Consider the following version of their argument:

P1: If one knows that one values, then one has sufficient evidence.

P2: Evidence about the present (E1) is not sufficient.

P3: If evidence about the present (E1) is not sufficient, and one does know that one values, then one must have further evidence E2 (about the future).

C1: Evidence about the future (E2) is necessary for knowing you value.

P4: To obtain evidence about the future, affective forecasting is necessary.

C2: Affective forecasting is necessary for knowledge of values.

I will now evaluate their position.

4. The case against (the necessity of) affective forecasting

4.1. Unjustified (general) conclusion

In this section I make three related claims. First, I argue that Tooming and Miyazono unjustifiably draw conclusions about *types* of evidence from considerations about *particular sets* of evidence. They argue that, since Katherine's particular set of evidence is defeated at some later stage, the type of evidence she relies on cannot be sufficient in principle. Secondly, I argue that this leads to unreasonably demanding requirements for self-knowledge: to know what we value, we need to predict future defeating scenarios. Third, I question whether the transition from the putative insufficiency of evidence about the present to the necessity of evidence about the future is justified. Let us start with a reconstruction of their main position:

P1: If one's evidence is defeated at a later stage, then it turns out it was not sufficient for knowing whether one values.

P2: Katherine's initial evidence (E1)* was defeated at T2.

C: Katherine's evidence (E1) was not sufficient for knowing whether she values.

*Recall E1 is a set of evidence about present mental states.

First, I should clarify how I understand the notion of "defeat" being deployed here. By "defeat", I take Tooming and Miyazono to mean that, if there is some piece of evidence in the world that shows one's belief is false, then one is not entitled to holding that belief. In the given example, Katherine's emotional reactions at T2 constitute evidence against Katherine valuing. Since her belief at T1 was that she values, at T2 she is

no longer entitled to holding that belief because she now has defeating evidence. Note that this is a strong understanding of defeaters: whilst a defeater can more modestly be considered to merely undermine or threaten claims to knowledge, here they are taken to be more damaging (which, one might add, is not a particularly plausible notion).

Tooming and Miyazono may, and I suspect do, further imply that her belief is already defeated at T1. Even if the defeaters are not yet present at T1, the fact that they will, as a matter of fact, manifest, means that Katherine's evidence at T1 is, unbeknownst to her, defeated. Therefore, at no stage is her belief justified, and she does not have knowledge, both at T2 and T1. That is, Katherine is not *in a position to know* her valuing state – even if she does in fact value, epistemically she is not entitled to that belief.

The above reconstruction of their argument brings out, I suggest, a subtle equivocation on the word “evidence” in their overall reasoning. On the one hand, we have *a set* of evidence, i.e., the specific pieces of evidence in one's possession at a point in time. In this case, this is Katherine's set E1, consisting of independent evidence about her present mental states. On the other hand, we have *types* of evidence at a more general level. Here, we are considering two specific types: evidence about the future and evidence about the present. I want to suggest that Tooming and Miyazono are moving from a claim about the insufficiency of a particular set to the insufficiency of a particular type: that defeat of a particular set of evidence in a particular instance is sufficient reason for us to conclude that that type of evidence is never sufficient for the type of belief it is used as evidence for.

Katherine's set is almost exclusively made up of evidence about the present (there may be evidence about the past, too, however it does not make a difference for our current purposes). Based on this set, she forms the belief that she values. This belief is later defeated (or there is a future defeater for her belief), and so we may infer that her set is not sufficient (barring issues about her having changed her valuing state in the meantime). Whilst it makes sense to conclude that her set E1 is not sufficient for her conclusion to amount to knowledge, we cannot thereby also conclude that the kind of evidence that makes up her set is never sufficient for that aim. For one, there may be an issue with her particular set E1. Perhaps it has very little evidence, or perhaps what she takes to be evidence was arrived at hastily and so may actually not contain evidence at all. Whatever the reason, set E1 seems insufficient for a knowledge claim. That does not mean that an amended set E1* could not amount to sufficient evidence. Further, and crucially, it does not mean that, in general, evidence about the present cannot be sufficient for knowing what we value.

Consider a parallel: you look at a clock at a train station and conclude it is 10am. A few moments later (at T2), a passerby kindly informs you that most clocks at that station are stuck and do not work. As it happens, you were

consulting the only clock that works (it had been fixed that morning), and so your belief (that it is 10am) was true. However, you are no longer entitled to that belief, because your evidence (the clock indicating it) has been defeated.

That being said, I think it is unreasonable to conclude that clocks indicating the time are not sufficient evidence for what time it is just because they are defeated in some scenarios. Tooming and Miyazono's position seems to imply that, at the time of reasoning (T1), you should have predicted that no one would tell you that most clocks do not work. That does not seem like something we should require in these cases – it is too demanding, and there does not seem to be a principled way to know in which cases we need to predict clock-defeating scenarios and in which ones we do not.

Not only is the conclusion that evidence about the present is not sufficient unjustified, but the transition to evidence about the future being necessary is also questionable. What justifies that transition? For Tooming and Miyazono, it is the idea that valuing is a state that requires a stable disposition to react a certain way in certain situations. As a result, they would argue that knowing whether you have a disposition now involves knowing whether that disposition will manifest in the future. Hence, it is necessary to know our future emotional states for them.

Nevertheless, I take it that knowledge about the present is primary, and knowledge about the future is only secondary (or derived from the former). That is, it is based on knowledge of a current valuing state that we might derive knowledge about our future affective reactions. Whilst knowledge of present dispositions may lend credence to beliefs about the future (e.g.: it is likely that I will enjoy holding my second baby in my arms), there is no good reason to believe that we can have knowledge about the future that is independent of knowledge about the present. Indeed, it is questionable whether we can have knowledge about the future without first having knowledge about the present, which I explore in more detail in [section 4.3](#)).

Beyond issues with Tooming and Miyazono's argument, this section brings to light something crucial for our discussion: whether the future confirms or defeats a present claim to self-knowledge, this does not mean that knowledge of the future is necessary. Affective forecasting does not constitute evidence in itself; at most, it might challenge the present evidence available, however there does not seem to be any reason to think it is ever necessary for knowing what we value. As such, affective forecasting is an unlikely necessary contender for generating knowledge of present valuing states.

4.2. *Assuming closure*

Let us assume Tooming and Miyazono are right and affective forecasting is necessary for knowledge of our valuing states. In this section I show that, if they are right, then they are committed to (unrestricted) epistemic closure, which has several unsavory implications. Note that this does not mean they explicitly endorse this principle. Rather, it can be derived from their position. I show this as follows: I will offer a straightforward implication of their argument, then I point out that it would be unsound without closure, and so I conclude that their view assumes closure. First, here is an argument they are bound to endorse based on their view:

P1: If you value x , you (will) react a certain way to it.

P2: If you know P1 and you know you value x , then you know you will react a certain way to it.⁶

P3: Katherine knows P1 and that she values x .

C: Katherine knows she will react a certain way to x .

I bring up this argument structure for the following reason: formulated this way, the argument relies on closure. Closure has been formulated in several ways, and it is beyond the scope of this paper to adjudicate which one is more plausible. However, as a starting point, consider Dretske: “if S knows that P is true and knows that P implies Q , then, evidentially speaking, this is enough for S to know that Q is true” (2005, p. 27). Without accepting closure, the above argument would not work: P2 would lack justification. Only if we introduce a premise 0 (presumably a tacit or hidden premise) stating the principle of closure will P2 make sense. I now show that accepting this principle leads to unsavory implications in this case.

Accepting closure seems to imply that, since Katherine knows she values having children, and she knows that valuing children implies feeling happy when having children, then she is in a position to know that she will be happy when having children. The same could be said for a slew of other emotions. She knows that if someone values having children, they will not feel hatred toward them (or generally have unusual negative emotional reactions toward them). For the purposes of this discussion, assume that Katherine does know these things. She does value and she knows it. And she

knows that if one values having children, then one will feel happy when having children. Or, that if they value them, they will not have unusually negative emotional reactions to them.

This seems too strong. It seems that, based on what Katherine knows about valuing and what it implies, she is in a position to know significant amounts of her future emotional behavior. She knows she will not regret having them, ever feel hatred toward them, she will feel extremely joyful to hold them, she will not feel angry with them except in reasonable situations and so on. Crucially, she can know all of this *based on her knowledge about her valuing state alone*. But shouldn't more evidence be required for such varied instances of knowledge?

Even if she feels exactly the way she expects based on her beliefs about what valuing entails (mostly having a plurality of intense positive emotions and few unjustified negative ones), this does not mean that she *knows* she will feel these ways. Knowing that she values *now* may put her in a position to reasonably come to think certain reactions likely, but we need not agree that she *knows* she will feel these ways. Particularly, we need not agree she knows all of this based on her present valuing state alone. More evidence should be required for these conclusions. She should, presumably, know whether she will be financially stable when having the baby, whether her mental health will be in a good place, whether she will be a single parent, whether she will have to give up something important to her (e.g., a career), and other facts which might significantly bear on her reactions to her baby. I am not suggesting knowledge of the future is not possible, but that if it is at all possible, it should require more evidence than the presence of a single present mental state.

So, I contend, licensing closure here would be unnecessarily strong, as we would grant too much future knowledge too readily. Consider: if Katherine will be happy when she has a baby, then she will not be hit by a bus before giving birth. She knows this. Does she therefore know she will not be hit by a bus? It seems that we are potentially legitimizing too much knowledge about the future, knowledge based on very limited pools of evidence that do not consider relevant future factors.

This section argued as follows: if Tooming and Miyazono are right, then the picture of self-knowledge of valuing they paint requires closure. However, accepting closure would lead to rather unsavory implications about what Katherine knows. She would know, for instance, that she will not be hit by a bus or that she will not have a miscarriage – all based on her knowledge about her valuing state. This, I think, is an extremely unpalatable implication. We should, then, reject Tooming and Miyazono's position.

4.3. Evidence about the past

Tooming and Miyazono anticipate a potential objection to the idea that affective forecasting is necessary for many types of substantive self-knowledge such as courage or valuing. The objection appeals to evidence from past experiences (p. 34). Specifically, one might argue that evidence about one's future reactions is not necessary, because evidence about the past can be independently sufficient. That is, to know whether I am kind I do not need to anticipate how I would react when facing someone in need. Instead, I need only look at how I felt and behaved in those situations in the past.

They disagree with this claim, stating that the evidence from the past is not sufficient because, at least in the case of some mental states like valuing, they are “supposed to be stable and constant across time and situation” (p. 34). So, in most cases our “limited past experiences in limited situations” are not sufficient for knowing certain mental states such as valuing. That is, they are not denying that past evidence might play some evidential role, but they believe that knowing whether we have a disposition now *requires* knowing whether it will activate in the near future (we may have stopped having the disposition between the past and the present). Furthermore, they argue, past experiences are open to different interpretations “you might argue that you will be different the next time, in the right kind of situation” (p. 34.). But to test these interpretations, we need to consider different possible future scenarios – and as such affective forecasting appears to be necessary even if past evidence is present.

That being said, they do think that past experiences can be sufficient at least for some people in some cases.⁷ The example they provide is that of elderly people, pointing out that a person with extensive life experience can easily know whether they are courageous (the army general has sufficient evidence about their performance in wars etc. to confidently draw that inference). On the other hand, they note, those with far less life experience are not in a position in which they have the same amount of evidence and therefore require affective forecasting.

There are several issues one could discuss at this stage. However, I will cast aside the questionable claim that only elderly people can sufficiently rely on past evidence. Instead, I will focus on the idea that knowing whether we value amounts to knowing whether a disposition will manifest in the future, such that affective forecasting is necessary. Could this be right?

First, it is not clear how they construe affective forecasting – what kind of mechanism is it exactly? To what extent would such a mechanism rely on evidence about the past? One might reasonably ask, in other words, how one could predict a future emotional state without relying on past evidence. To

know whether it might rain in two weeks, you need to know how the weather tends to evolve. Likewise, to know whether I will be happy or scared when I hop on a rollercoaster depends on how I have reacted in the past. Alternatively, if I have never had any relevantly similar experience, my forecasting will be based on ideas I might have (e.g. my friends have fun on rollercoasters, so I will too), but it would not amount to knowledge – they would just be mildly educated guesses.

Perhaps, they might claim, affective forecasting is more simulationist⁸ in manner: you imagine yourself on a rollercoaster, and however your imagination fills in the details (e.g., screaming in terror), that is the emotional reaction you predict. Since they do not commit to any particular mechanism of affective forecasting, this is a route they could plausibly take. Even on this kind of picture, however, I think the point about reliance on past experience still stands. The imagination, especially if we are allowing it to run free (i.e., we are not imagining what we want to see, e.g., being unperturbed on the rollercoaster because we want to think of ourselves as courageous), tends to be constrained by what we already know. To showcase how that might work, consider Van Leeuwen's (2016) theory of the pathways from imagination to emotions, which compellingly explains how our imagination is constrained.

Specifically, Van Leeuwen paints a compelling picture of two parallel pathways, one from perception to emotions (P-C-E-C) and one from imagination to emotions (I-C-E-C). The two pathways exhibit certain similarities, particularly what van Leeuwen calls “reality congruence” (p.93). By this he means that what we imaginistically imagine⁹ are generally things that could happen in our environment, given our beliefs about that environment. He thinks that our imagination is by default constrained by our environment, which he defends with three distinct claims.

One point in favor of this is that the content of our imaginings is already largely composed of representational constituents already instantiated by perception. Two, factual beliefs “typically constrain inferences from one imagining to the next” (pp. 93–94). For instance, if one imagines a person being drunk, then one will imagine that person stumbling and slurring their words. Three, there is empirical evidence that children strongly prefer imagining realistic story continuations for realistic stories (and fantastical continuations for fantastical stories). By “realistic” is meant “*in conformity with familiar patterns of causation*” (p. 94).

Thus, it seems plausible that what I simulate is still constrained by past experiences, and if these past experiences are considered insufficient when used as evidence in reasoning, then surely, they would render our affective forecasting likewise insufficient. Instead, we can simply use our evidence about the present and the past to obtain knowledge about our present mental states. We do not need to go through an additional move of using that data to predict future emotional states to only then infer our present

states. If we follow Tooming and Miyazono in rejecting evidence about the past as being sufficient for substantive self-knowledge, then we should think that our evidence about the future, since it is plausibly reducible to our evidence about the past, is also insufficient.

One might here object that, whilst it is plausible that affective forecasting heavily relies on evidence about the past, it might also require additional inferential steps. In that case, if the additional elements are not reducible to past evidence, then affective forecasting remains unchallenged. First, to the extent that other considerations are necessary, they need not be about the future. For instance, the majority of additional elements might be beliefs either about ourselves (e.g., “I have a reliable support network”) or more general beliefs (e.g., “People grieve for about a year”). If this is so, then since we are not relying on evidence about the future, affective forecasting can still be argued to be unnecessary.

Second, assume for the sake of the argument that we in fact sometimes do rely on evidence about the future. What kind of future evidence could that be, given what has already been discussed about the reliance of future evidence on past evidence? If we exclude evidence about future emotional states, presumably the other relevant type of evidence will concern our future context. For instance, Katherine might know that she will be financially stable, allowing her to enjoy spending time with her baby. If we set aside the fact that, plausibly, even this kind of evidence is reducible to either past events or beliefs (e.g., “I have a rich relative with a short life expectancy”), this kind of future evidence still does not do much to help the Tooming and Miyazono position. This is because evidence about future contextual factors solely determines whether we are likely to be in the right circumstances for the disposition to manifest, but it does not show that we have that disposition. In other words, evidence about the future does not contribute to the question of whether we have the relevant affective dispositions, and so there seems to be nothing that affective forecasting brings that cannot already be inferred from beliefs about the past and present.

Thus it seems to me that, their position that evidence about the future is necessary for substantial self-knowledge, and thus that use of affective forecasting is necessary, does not hold up to scrutiny. Evidence about the future can be reduced to evidence about the past and the present (and to beliefs about general states of fact), such that we can generate the relevant knowledge without the additional step of affective forecasting. Given the limitations I outline here and in the other parts of [section 4](#), I have argued that they fail in showing that affective forecasting is necessary for substantive self-knowledge in general, and for valuing states in particular.

5. Positive proposal: use of theories

Even if I am right and affective forecasting is not necessary for obtaining self-knowledge of what we value, Tooming and Miyazono might still be right in suggesting that it explains the overall asymmetry between substantive and trivial self-knowledge. In other words, they might still be able to explain the difficulty in generating substantive self-knowledge by appealing to the fact that, although not necessary, affective forecasting is still evidence we in fact make use of. Since we are not adept at it, then, we may be making it more difficult for us to generate substantive self-knowledge in general and knowledge of valuing states in particular.

In effect, Tooming and Miyazono are offering an argument to the best explanation: affective forecasting best explains why substantive self-knowledge is systematically more difficult to come by, so we should take it as the main criterion for the epistemic asymmetry between the two types of self-knowledge (although they agree it is not the only factor). This would then mean that affective forecasting does play a significant role in knowing what we value, since it accounts for the difficulty in generating such knowledge. I think even this claim fails, and I think it does so because it overlooks an important alternative: the use of theories.¹⁰

Recall that, for them, if a type of substantive self-knowledge requires affective forecasting, then it is categorized as strong, rather than weak substantive self-knowledge. Yet this tells the story from the middle: it skips the introduction. That is, difficulties concerning obtaining evidence x or y *come only after* establishing whether x or y would even count as evidence. Exactly why that is so will become clearer in what follows.

First, when I use the notion of “theory” I loosely borrow from the literature on other minds, where Theory Theory is roughly conceived of as “a number of ‘folk’ laws or principles connecting mental states with sensory stimuli, behavioral responses, and other mental states” (Barlassina & Gordon, 2017). In other words, by “theory” I mean some association between mental states on the one hand and certain outputs (behavioral, emotional, etc.) on the other hand. For instance, I may connect the mental state of love with certain feelings (warmth, butterflies in my stomach, desire to spend time with a person), such that when I experience any of these feelings, I associate them with the mental state of love. Likewise, I may associate kindness with certain behaviors (sharing, donating, paying attention to others) and certain emotions (pity toward beggars, compassion for those in need etc).

These associations tend to function in an epistemically rigorous way: they inform what counts as evidence for a mental state, and they allow us to draw certain inferences from a body of evidence. However, these theories are not employed consciously. Building on this notion of “theory”, then, I think that inferring anything about a valuing state (and other types of substantive self-

knowledge in general) requires i) theories about what constitutes that mental state and ii) theories about how to generate evidence and what counts as evidence.

Let us return to Katherine. Whether she wants to find out if she desires or values having another baby, she cannot blindly use every reaction as a piece of evidence. She cannot use her excitement to spend time with friends or her sadness upon reading *Anna Karenina* as evidence that she values having another baby. Indeed, when she focuses on her jealousy about other mothers or her joy upon imagining holding a new-born, she can use these reactions as evidence because she has certain theories. She has theories linking, say, jealousy about people having x with desire to have x oneself, or joy upon imagining x with desire to have x . Without these theories, she would not be able to infer the presence of some mental state like desire or valuing.

Furthermore, when actively looking for evidence, she looks for certain reactions. She specifically watches movies about parenting, or vlogs of new moms, and imagines herself with a new baby. That is, the way she goes about generating evidence is guided by what she thinks would count as evidence. Throughout these actions and imaginations, she carefully monitors her reactions. If, whilst genuinely asking herself whether she wants another child, she starts imagining dinosaurs, that would be odd and irrelevant. Were she to imagine dinosaurs, she would not be carefully monitoring her reactions (perhaps awe if she is a fan) – she would interrupt the reasoning process (or change to a different one).

Kloosterboer (2022) expresses similar considerations about what may or may not count as evidence, in the process highlighting a crucial point worth considering: there are significant differences amongst individuals, such that what counts as evidence for your valuing state may not count as evidence in my case. As she puts it:

How does Katherine know that, e.g., her envy reveals a deeper truth about herself and isn't due to, for example, just having a grumpy day? The envy – the internal prompting – does not itself tell us whether it is a symptom of a deeper desire. After all, emotional episodes have different kinds of significance for a person. (p. 11)

The point is an important one: what may count as evidence for me may not count as evidence for you. This suggests that one cannot blindly seek evidence but needs instead to be guided by a theory. Further, however, this suggests that there might be two different types of theories we employ in different contexts: whilst we all presumably have theories of happiness in general, we might also have theories about happiness in relation to oneself (theories about what makes *me* happy, rather than what makes *you* happy). These personal theories explain why evidence seems to sometimes behave differently from subject to subject. Tooming and Miyazono do not have a plausible explanation for why different sets

of evidence may not allow different individuals to reach the same conclusion.

I thus think that, without the use of theories (be they more personal or general), we would be unable to purposefully generate evidence, and we would not be able to separate relevant from irrelevant evidence. How does Katherine know that awe about dinosaurs does not mean she wants another baby, but jealousy over other mothers does? The best explanation, I contend, is that she has a theory to that effect. In other words, the use of theories is prior to the use of affective forecasting: indeed, how we use affective forecasting (and to a certain extent how helpful the results are) will depend on the accuracy of our theories.

Furthermore, having false theories or misapplying them is more explanatorily powerful for the difficulty of generating substantial self-knowledge than the need for affective forecasting. Indeed, we could explain many of our difficulties with affective forecasting in terms of failures at the level of theory. For instance, flawed theories about how to use the imagination to generate evidence are a noteworthy culprit in cases such as Katherine's. She may not know what kind of imaginative scenario she should imagine (how much detail, what details), and she may not know the intentional object of her resulting emotion (what the emotion is actually about). She may also be employing the wrong theory: a general rather than personal one. Thus, I argue that Tooming and Miyazono do not offer a satisfactory explanation for the difficulty of substantial self-knowledge, particularly of valuing states. Instead, what explains the difficulty of generating this kind of knowledge is more likely connected to the subject's use of theories.

One might here worry that, whilst it is true that theories are necessary, they are only necessary *conditions*, rather than *grounds* for one's knowledge.¹¹ The distinction can be clarified through further examples: one might say that having a brain, or not being profoundly impaired are pre-requisites for knowledge. However, it is not particularly informative about the nature of self-knowledge to be told that one needs to have a brain or to be alive. Likewise, one might hold, having certain theories is necessary in the way that having a brain is. On the other hand, since affective forecasting plays a role at the level of content of the target knowledge claim (i.e., is a part of the inference), then it can be a ground for, rather than merely a pre-condition for self-knowledge.

However, I disagree with this position, since theories also play a role at the level of content. If we are to represent the target phenomenon via an inference from premises to conclusions, then theories are present among the premises. Katherine's evidence is not only of individual pieces of evidence, for evidence does not speak for itself. Instead, she must use premises about how certain pieces of evidence connect to a given conclusion. One cannot infer Q just from knowing that P, one

must also know $P \rightarrow Q$. As such, I think arguing that theories are only necessary in the way in which having a brain underestimates the impact of theories. Theories are an intrinsic part of the reasoning process that results in knowledge claims, and as such not mere pre-conditions.

This section has argued that, whilst we may occasionally employ affective forecasting when deliberating on whether we value, we cannot say that affective forecasting plays a crucial explanatory role in such knowledge.

6. Conclusion

In this paper, I followed Tooming and Miyazono in understanding valuing as (at a minimum) a set of emotional dispositions toward the valued object. I then argued against their position that affective forecasting is necessary for knowledge of our own valuing states. I have showed that their position is untenable by arguing that it relies on implausible epistemic principles and by suggesting that affective forecasting itself can be replaced with inferences based on present and past evidence.

Yet Tooming and Miyazono may still argue that we do generally employ affective forecasting to obtain substantive self-knowledge such as valuing. They could therefore claim that our unreliability with affective forecasting then explains the general asymmetry between substantive and trivial self-knowledge: the fact that the substantive type is systematically harder to come by and faces more obstacles. The upshot of that position would be the idea that affective forecasting plays an important explanatory role when it comes to self-knowledge of valuing, as it would explain why it is generally harder to obtain this kind of knowledge. I argued that even this explanation fails, as it overlooks a powerful alternative: the use of theories – having false theories or misapplying them is more explanatory powerful than the mere absence of sufficient evidence, because it can explain why we may be lacking the evidence. As such, I conclude that knowing what we value does not require affective forecasting, nor is affective forecasting the reason why this type of knowledge tends to be harder to come by.

Notes

1. Note that this is a qualified claim which will become clearer in [section 2](#), after introducing the distinction between substantive and trivial self-knowledge.
2. See Wilson and Gilbert (2005).

3. These dispositions might be seen as connected by what the subject sees as important, or by something else – for the purposes of my argument, I do not need any such commitment. Furthermore, Tooming and Miyazono do not make such commitments, and I follow their dialectic.
4. A non-Scanlonian version of intellectualism could emerge from the literature on commissive attitudes. One might argue that valuing something just is being committed to it. However, this notion is distinct from valuing as I understand it. For one, Coliva (2009) notes that commissive beliefs are the result of a judgment: we judge that something is true, and in so judging commit ourselves to it. There are more criteria, but this one will suffice to highlight a disparity between commitments and values: values are seldom developed in such a rational manner. Indeed, values are often the result of upbringing and circumstances: I have never *judged* that my family or friends are good, but rather developed certain attitudes toward them over time. Indeed, it is precisely because of this way of coming to value that we do not always know our values. I will therefore not address commitments here, yet note that, if one thinks commitments are values, and that they are easy to come by, then that is consistent with my position, as I argue that affective forecasting is not necessary for knowing we value. Thanks to a referee for raising this worry.
5. For the sake of the argument, I will grant Tooming and Miyazono that the same kind of evidence should *prima facie* suffice for knowledge of desire and valuing respectively.
6. If affective forecasting is necessary for knowing we value, then one who knows one values will have already used affective forecasting. Therefore, it is plausible that they know they will react a certain way to x.
7. Note that this contradicts their claim that, at least in the case of valuing, affective forecasting is always necessary. To be charitable to this portion of their argument, we might interpret them as saying that *in most* cases, knowing what we value requires affective forecasting. This would not be a problem for my arguments against them: if their claim is that in most (but not all) cases affective forecasting is necessary, my position is that affective forecasting is *never* necessary. However, I do think that their paper is inconsistent or at the very least vague when it comes to these two claims. Concerning valuing, they state: “there are cases of SSK [substantive self-knowledge], such as knowing one’s values, where knowing one’s affective reactions to future events is necessary” (p. 23). This suggests that, for them, affective forecasting is always necessary for knowledge of valuing. It could, however, be interpreted as “in most cases” to make sense of their latter claim about past evidence sometimes being sufficient.
8. For an overview of simulationism in the literature on other minds see Barlassina and Gordon (2017).
9. His term for a type of sensory imagination.
10. They do consider whether the use of theories is the correct explanation, but they reject it claiming that it is not informative. I think they discard this option too quickly.
11. Thanks to a referee for raising this worry.

Acknowledgements

I am very grateful to Rory Madden and Lucy O'Brien for invaluable feedback on earlier versions of this article. I would also like to thank the audience of the Apophansis Journal Symposium at King's College London (June 2024) for helpful comments on this work, as well as the anonymous reviewers from Philosophical Psychology, whose careful comments helped make my position clearer and stronger.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributor

Diana Craciun Currently a PhD student at University College London, Diana Craciun's research focuses on self-knowledge of valuing states, particularly on how our emotional reactions to imagined scenarios can result in such self-knowledge. Prior to this current project, Diana's MPhil thesis (also at University College London) has focused on emotional recalcitrance and whether such emotional experiences are irrational. During her BA, Diana studied Philosophy and German, successfully completing a year abroad at Humboldt University in Berlin, where she studied Philosophy and German Linguistics.

References

- Anderson, E. C. (1993). *Value in ethics and economics*. Harvard University Press.
- Barlassina, L., & Gordon, R. M. (2017). Folk psychology as mental simulation. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition). Retrieved January 5, 2024, from <https://plato.stanford.edu/archives/sum2017/entries/folkpsych-simulation/>
- Cassam, Q. (2014). *Self-knowledge for humans*. Oxford University Press.
- Coliva, A. (2009). Self-knowledge and commitments. *Synthese*, 171(3), 365–375. <https://doi.org/10.1007/s11229-008-9322-x>
- Dretske, F. (2005). Is knowledge closed under known entailment? The case against closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 13–26). Blackwell.
- Helm, B. W. (2001). *Emotional reason: Deliberation, motivation, and the nature of value*. Cambridge University Press.
- Jaworska, A. (2007). Caring and internality. *Philosophy and Phenomenological Research*, 74(3), 529–568. [10.1111/j.1933-1592.2007.00039.x](https://doi.org/10.1111/j.1933-1592.2007.00039.x).
- Kloosterboer, N. (2022). Substantial self-knowledge and the necessity of avowal. *The Journal of Value Inquiry*, 1–21. <https://doi.org/10.1007/s10790-022-09916-3>
- Lawlor, K. (2009). Knowing what one wants. *Philosophy and Phenomenological Research*, 79(1), 47–75. <https://doi.org/10.1111/j.1933-1592.2009.00266.x>
- Scanlon, T. M. (1998). *What we owe to each other*. Belknap Press of Harvard University Press.
- Scheffler, S. (2010). Valuing. In *Equality and tradition: Questions of value in moral and political theory*. Oxford University Press.
- Seidman, J. (2009). Valuing and caring. *Theoria*, 75(4), 272–303. <https://doi.org/10.1111/j.1755-2567.2009.01049.x>

- Tooming, U., & Miyazono, K. (2023). Affective forecasting and substantial self-knowledge. In A. M. Sánchez & A. Salice (Eds.), *Emotional self-knowledge* (pp. 17–38). Routledge.
- Van Leeuwen, N. (2016). The imaginative agent. In A. Kind & P. Kung (Eds.), *Knowledge through imagination* (pp. 85–109). Oxford University Press.
- Wilson, T. D., & Gilbert, D. T. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science*, 14(3), 131–134. <https://doi.org/10.1111/j.0963-7214.2005.00355.x>