Introduction

Knowledge-first epistemology places knowledge at the normative core of epistemological affairs: on this approach, central epistemic phenomena are to be analysed in terms of knowledge. First put forth in the groundbreaking *Knowledge and Its Limits* (Williamson 2000), the knowledge-first approach to epistemological theorizing is one of the most successful and prolific research programmes in epistemology and philosophy in general, having given rise to novel theories of the epistemic justification ((Bird 2007), (Ichikawa 2014), (Lasonen-Aarnio MS), (Kelp 2016), (Millar 2010), (Miracchi 2015), (Schellenberg 2018), (Silva 2017), (Simion 2019), (Sutton 2005), (Williamson 2000)) the nature and normativity of evidence and defeat ((Lasonen-Aarnio 2014), (Kelp 2023), (Dutant & Littlejohn 2021), (Simion 2024c,d), (Williamson 2000)), understanding ((Kelp 2015), (Sliwa 2017)), the basing relation (Carter & Miracchi 2024), permissible suspension ((Miracchi 2017), (Simion 2024d)) norms of inquiry ((Kelp 2023), (Friedman 2020), (Willard-Kyle 2023)), rationality ((Dutant and Littlejohn Forthcoming), (Miracchi Forthcoming)), KK principles (Goodman and Salow 2018), epistemic functions ((Kelp 2018), (Simion 2019)), know-how and intentional action ((Pavese 2015), (Stanley & Williamson 2001)), epistemic responsibility and blame (Lasonen-Aarnio 2014, Williamson Forthcoming), norms of speech acts (Benton (2011), Kelp (2018), Kelp & Simion (2021), Simion (2021), Turri (2016), Williamson (1996)), action (Hawthorne & Stanley 2008), practical and theoretical reasoning ((Fantl & McGrath 2012), (Simion 2021), (Williamson 2000)), perceptual entitlement (Millar 2010, Schellenberg 2018), testimonial entitlement (Simion 2021), disagreement ((Hawthorne & Srinivasan 2013), (Simion & Broncano-Berrocal 2024)), group belief and group justification (Bird (2010), Simion, Carter, and Kelp (2022)), the epistemology of science (Bird 2022) and the epistemology of law (Bloome-Tillmann 2017, Littlejohn 2017, Moss 2019).

Claims paradigmatically associated with the knowledge-first programme include that knowledge is a mental state, that it is unanalysable, that it is distinctively valuable, that knowledge is the aim of inquiry and belief formation, that other epistemic normative phenomena (e.g. justification, understanding, responsibility, evidence, defeat) are to be unpacked in terms of knowledge, that knowledge is the norm of assertion and action, that it is essential to intentional action, that social epistemic phenomena (e.g. testimony, disagreement, group belief) afford unpacking in terms of knowledge. Crucially, many of these claims are theoretically independent from each-other; unsurprisingly, while some knowledge-first theorists subscribe to the full research programme, others only champion some of its claims.

This study focuses on the normative core of the programme: it offers a defence of an integrated, naturalistic knowledge-first account of justified belief, reasons, evidence, and defeat, permissible assertion and action, and the epistemic normativity of practical and theoretical reasoning. On this account, *the epistemic* is an independent normative domain organized around one central function – knowledge – which generates epistemic norms.

Here is the game plan: in Section 1, I argue that the main epistemic etiological function of our epistemic practice of inquiry is generating knowledge. In turn, this epistemic function generates epistemic norms governing moves in our epistemic practice, such as forming and maintaining beliefs, asserting, and reasoning. Section 2 discusses the epistemic normativity of belief and the nature of justification. Section 3 integrates the account of justification with a theory of the nature and normativity of reasons, evidence, and defeat. Section 4 discusses the normativity of action and assertion, and Section 5 defends a unified account of epistemically good reasoning.

1 The Knowledge Function

Introduction

Tim Williamson (2000) famously argues for the knowledge-first programme on abductive grounds: the best explanation for the fact that an analysis of knowledge has proven elusive for several decades is that knowledge does not afford a non-circular dismantling analysis. If so, we should use knowledge as a primitive in analysing other central epistemic phenomena.

This section's ambition is to offer further support to the knowledge-first programme. I argue that generating knowledge is the function of our epistemic practices, and that this function generates epistemic norms of proper functioning that constitute the epistemic domain.

1.1 Knowledge Is the Function of Inquiry

What is the function of our epistemic practice of inquiry? A clarification: In asking this question, I follow tradition[1] and speak of inquiry broadly conceived: that is, I take all of our epistemic endeavours to fall under the general epistemic practice of inquiry: automatically formed attitudes (beliefs that I form as I walk down the street) as well as careful reasonings and judgments, suspensions,

[1] E.g. Kelp (2014a, 2021), Sosa (2023).

credence formations, as well as epistemic exchanges – assertions, tellings, conjectures, and so on.

Inquiry is an epistemic practice. Practices have main intrinsic goals, or functions: the practice of driving mainly aims at safely getting one to one's destination in a reasonable amount of time; that's its main function. The practice of cooking mainly aims at producing tasty and nutritious food. The practice of medicine aims at generating health, and so on. The restriction to intrinsic here is crucial; often, but less importantly for my purposes here, practices have extrinsic aims too. Plausibly, the practice of inquiry and the practice of medicine share their extrinsic aim: the aim of survival. Compatibly with that, these two practices will have different intrinsic aims. Ideally, when the intrinsic aim of the practice is achieved, it reliably serves the achievement of the corresponding extrinsic aim. This, however, may not be the case. Some of our practices are bad for us.

Now, note that, when it comes to their goals, practices tend to strike a good balance between value and achievability; they are often aimed at the most valuable achievable intrinsic goal. The value at stake is not good *simpliciter*, to be sure; many of our practices have horrible goals! What is at stake is domain-specific goodness. Driving does not aim at teleporting you to your destination. Nor does it aim at safely getting you one mile away from your destination. It aims at getting you to your destination – which is the most valuable target in the domain – in a reasonable amount of time – which is an achievable goal. Similarly, cooking aims at producing tasty food, not at producing, say, self-generating nutrients (not achievable), nor half-baked goods (not the most valuable). Medicine aims at generating health, not at making us immortal (not achievable) nor at keeping us barely alive (not the most valuable goal). Practices' intrinsic goals tend to strike a good balance between value and achievability.

It is widely accepted that knowledge is, epistemically, more valuable than any lesser epistemic standing. The way in which knowledge is more valuable has been a topic of debate in the past years (e.g. Pritchard (2010)). The challenge, however, is to explain how, not to argue that it is: little doubt[2] has there been expressed – since Plato's *Meno* – with regard to the superiority in value thesis itself. Furthermore, several people think that knowledge is not only more valuable than lesser epistemic standings, but also distinctively so. That is, the difference is one in kind rather than in degree (Pritchard 2010, Simion and Kelp 2019). I have argued in previous work that this is so because knowledge marks a jump on the epistemic value continuum, in that it stands in a non-Archimedean

[2] But see e.g. Pritchard (2009) and Kvanvig (2003) on scepticism about the value of knowledge.

value relation to states that fall short of knowledge: for the aim of leading a cognitively flourishing life, for creatures like us, some amount of knowledge is better than any amount of mere true beliefs. A life of cognitive flourishing is a life rich (enough) in knowledge.

Second, note that, in most epistemic walks of life, knowledge is readily available (Kelp & Simion 2017): we are fairly well-equipped epistemic agents, living in an extremely friendly epistemic environment. Furthermore, knowledge is much more readily available, and thus much more easily achievable, than both lesser and stronger epistemic standings. When it comes to stronger states – such as certainty or understanding –, that should be easy to see. It takes us quite a bit of time and effort – and, arguably, a fair amount of knowledge – to achieve understanding. Furthermore, (epistemic) certainty is hard.

To see it for states that fall short of knowledge too, consider first perceptual beliefs about middle-sized dry goods. My belief that there is a computer on the desk before me qualifies as knowledge: it is produced by a highly reliable ability to recognize computers in an epistemically hospitable environment. The crucial point is that belief formation by suitable processes in hospitable environments is the norm; formation of beliefs by unsuitable process, or in inhospitable environments is the exception. If this isn't immediately clear, consider again my belief that there is a computer on the desk before me and ask yourself what would have to be the case for my belief to remain true but fall short of knowledge: I mistake a hologram for a computer, whilst unbeknownst to me there is a computer somewhere else on the desk, I acquire my belief by a highly unreliable process such as a coin-toss. While any of this might come to pass, it is undeniable that, as a matter of fact, it only rarely does.

Consider, also, testimonial belief about propositions of practical importance: propositions about bills that need to be paid, the nature of your sickness and the medication that will cure it, what's available at the local restaurant, and so on. Or consider inferentially supported beliefs that exploit a variety of natural and social regularities: that my couch is still in my living room, that Paris is still the capital of France, and so on. Here too, when beliefs in these ranges are formed by suitable processes in sufficiently hospitable epistemic environments, they will qualify as knowledge.

These considerations suggest that, in a wide range of cases, knowledge is easily achievable. All we have to do to acquire knowledge is open our eyes, listen to what other people tell us, attend to our feelings, and so on. In contrast, true belief that falls short of knowledge is a rare commodity that exists only in very special environments. What is readily available is true belief *that is* knowledge (Kelp & Simion 2017).

Why not say that the aim of belief is truth (not true belief that falls short of knowledge, but truth *simpliciter*)? If we aim at true beliefs, every instance of knowledge will accomplish this aim, and so will some true beliefs that aren't knowledge. So true beliefs will be an easier aim to achieve than knowledge. Furthermore, plausibly, in the vast majority of cases, true belief just is knowledge.[3]

If I am right and the goal of a practice is often defined by the best score on both achievability and value, however, in virtue of the distinctive value of knowledge, knowledge and not true belief is plausibly the goal of the practice, even if the latter is present whenever the former is present and, sometimes (but not often) when the former is absent. To see this, note that, similarly, getting one to one's destination is more plausibly the goal of driving than getting one to a point that's situated precisely one inch before one's destination, in virtue of the fact that the former is more valuable than the latter (in spite of the fact that the latter is always achieved when the former is, and, likely, in virtue of being more easily achievable, on some cases when the former is not).

To take stock: We have seen that goals of practices tend to strike a good balance between value and achievability: they are the most valuable achievable goals. I have argued that knowledge is more valuable than lesser epistemic standings and much more easily achievable than both weaker and stronger states.

I submit, in the light of all this, that we have good reason to believe that knowledge is the most valuable achievable epistemic state, and thus *the* function of the epistemic practice of inquiry.

Importantly, that is not to say that some particular species of inquiry can't aim at stronger epistemic states: to the contrary, it is plausible that, for instance, scientific and moral inquiry aim at understanding the relevant phenomena. Note, though, that this is perfectly compatible with our general practice of inquiry aiming at knowledge, since species of a genus are bound to have extra-properties, on top of those of the genus itself.[4]

Perhaps more surprisingly, my view is also compatible with a particular species of inquiry aiming at less than knowledge. Here is why: a well-researched category in normative ethics is that of 'contrary-to-duty imperatives'. Very roughly, these are norms that step in when one is in breach of a norm. You break your neighbour's window; that is in violation of quite a few types of norms – moral, social, prudential, and so on. Now, the contrary-to-duty

[3] See also Hetherington (2000) for a view that uses a similar claim to put forth an analysis of knowledge as true belief.

[4] I am assuming a view on which understanding implies knowledge (see e.g. Kelp (2015)). For non-factive views of understanding, see e.g. Elgin (2017).

imperative asks you to apologize: it's what one ought to do given that not breaking the norm is not an option that's still on the table. It's the next best thing to do.

Similarly, there are domains where and situations when knowledge is not attainable. Some think philosophy is a domain like that;[5] I disagree. However, if this is true, we should expect this sub-domain of inquiry to be aimed at less than knowledge – roughly, to be aimed at the next best epistemic standing. Alternatively, one might be in a situation where urgency makes it unlikely that knowledge is achievable in the available amount of time: reporting is often like that. In these situations, as I have argued at length elsewhere,[6] one's reporting should be based on the best achievable epistemic standing. Again, all this is perfectly compatible with the framework I am developing here.[7]

1.2 Etiological Functions

I borrow the general normative picture I rely on from philosophy of biology – that is, from the normativity of etiological functions: in traits, artefacts and practices alike, functions generate norms. There is such a thing as a properly functioning heart, a properly functioning can opener and a proper way to make coffee. If that is so, what we need in order to identify a particular type of norm governing a particular type of practice, is to first identify its function.

[5] See e.g. Goldberg (2017).

[6] Simion 2016.

[7] Crucially, note that this observation (in conjunction with the value-theoretic picture defended so far) leaves it open whether in these fields, belief that falls short of knowledge will be permissible: after all, one need not hold that one should always end inquiry with belief: justified credence might do. Furthermore, what one wants to say here will also depend on one's commitments concerning justified belief and suspension. For people liking the claim that justified belief is knowledge, suspension will be permissible only when one does not have enough epistemic support for knowledge (e.g. Hawthorne and Srinivasan 2013). On my preferred view (Simion 2024d), suspension is justified insofar as one does not have the support needed for forming a justified belief, and justified belief is belief formed via a properly functioning knowledge-generating cognitive capacity (see Section 2). Views like this might raise some eyebrows (see e.g. McGlynn 2014): say that I know that I cannot know that p, but I also know that I have pretty strong support for p (just not enough for knowledge). Lottery cases are paradigmatic cases in this respect. Why should I not believe? After all, a true belief is surely more valuable than suspension. Note, however, that (1) just because belief is also epistemically valuable, it does not yet follow that it is more valuable than suspension, and (2) views like these do not predict that full neutrality is required in these cases – but mere belief suspension: one can always form the corresponding credence, without seemingly missing out on much. A high credence that one's preferred philosophical theory is true seems to be sufficient to get by, without full belief on the topic. For arguments against the Lockean thesis (i.e. the thesis that justified belief maps on to sufficient justification for high credence short of 1), see e.g. Kelp (2014b) and Douven and Williamson (2006).

On the etiological theory of functions,[8] functions turn on histories that explain why the item exists or operates the way it does. Take my heart: tokens of the type pumped blood in my ancestors. This was beneficial for my ancestors' survival, which explains why tokens of the type continue to exist. As a result, my heart acquired the etiological function (henceforth also e-function) of pumping blood. Acquiring an etiological function is a success story: traits, artefacts and actions get etiological functions of a particular type by producing the relevant type of benefit. My heart acquired a biological etiological function by generating biological benefit. Through a positive feedback mechanism – the heart pumped blood, which kept the organism alive, which, in turn, ensured the continuous existence of the heart – our hearts acquired the etiological function of pumping blood.

Not all functional items follow the model of the heart: there will be cases where a requirement of selection over generations for function acquisition will be implausibly strong.[9] The paradigmatic case is that of beneficial macro-mutations, also known as hopeful monsters[10] (Graham 2014, 30). Most mutations are harmful (think of extreme birth defects); once in a while, though, a happy accident happens: someone is born with an almost entirely new trait or organ, very different in kind from its ancestral trait, which actually benefits the recipient. Since they are mutations, they don't have an evolutionary history; they are 'first generation' traits. Still, they can acquire functions. What matters is that the existence/continuous existence of a trait is explained via a history of positive feedback.

While etiology does require some history, then, it does not require an awful lot of it; there are several ways to cash out the etiological requirement that do not presuppose directional selection, that is, selection over generations.[11] A trait can also acquire a particular function by ongoing, maintenance selection, or through a learning process, or even by the metabolic activity of the organism itself. What it all amounts to is explaining the existence/continuous existence of a trait through a longer or shorter history of positive feedback (Graham 2014, 35).

Here is, then, on a first approximation, what it takes for a trait to have an etiological function of a particular type:

E-Function: A token of type T has the e-function of type B of producing effect E in system S iff (1) tokens of T produced E in the past, (2) producing E resulted

[8] Defended by e.g. David J. Buller (1998), Ruth Millikan (1984), Karen Neander (1991). The etiological theory of functions is, by far, the most widely endorsed view in the literature. Its main competitor is the systemic theory of functions defended in Cummins (1975).
[9] See e.g. Sosa (1993) on Swampman.
[10] Davidson's Swampman is the epistemic incarnation thereof (Sosa 1993).
[11] See Buller (1998) for an excellent overview of etiological theories of function.