

Knowledge as a Social Kind¹

Tammo Lossau

ABSTRACT: I argue that knowledge can be seen as a quality standard that governs our sharing and storing of information. This standard satisfies certain functional needs, namely it allows us to share and store trustworthy information more easily. I argue that this makes knowledge a social kind, similar in important ways to other social kinds like money. This provides us with a way of talking about knowledge without limiting ourselves to the concept of knowledge. I also outline three ways in which this view of knowledge can shed light on familiar epistemological problems: it can explain why knowledge is the norm of assertion, it can help us carve out the harm associated with testimonial injustice, and it can provide us with a clear analysis of the dangers associated with spreading misinformation.

KEYWORDS: conceptual analysis, social kind, knowledge norm of assertion, testimonial injustice, misinformation

We agree on a variety of standards. Some are defined by institutions, such as standards of building safety or standards for what produce counts as organic. Some are only carried by an implicit understanding within a group or society, such as standards of politeness or standards for what counts as infidelity in a relationship. While these standards can be changed and sometimes opted out of, they at least by default create norms governing our actions. In this paper, I want to argue that knowledge is such a standard: it is an implicitly agreed on standard for the quality of information that governs which

¹This version of the article has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <https://link.springer.com/article/10.1007/s12136-023-00561-4>. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use <https://www.springernature.com/gp/open-research/policies/acceptedmanuscript-terms>

information we disseminate and store. Approaching knowledge in this way allows us to find a plausible and fruitful methodology for the epistemological analysis of social phenomena.

I will begin by outlining a methodological challenge: epistemology can either confine itself to a conceptual analysis, or it needs to find a way of productively talking about knowledge “itself”. The social kind view can offer a way of doing the latter while also providing a story that justifies our statements about knowledge. Part two will provide a motivation for the idea that knowledge is a social kind. I argue that knowledge, understood as a standard for the dissemination of knowledge, fulfills typical conditions for kindhood and occupies a functional role that is similar to that of other social kinds. Finally, I will sketch three applications of the social kind view: its explanation of the knowledge norm of assertion, its take on testimonial injustice, and the perspective on the risks of spreading misinformation it can provide. My aim here is to argue that the social kind view not only provides an answer to a methodological question, but also gives us useful new explanatory resources.

1. The methodological challenge

Philosophers often appeal to intuitions. This is particularly true for epistemologists in the Gettier debate, where proposed definitions of knowledge were rejected because they did not align with proposed counterexamples. But “post-Gettier” epistemologists also tend to support their claims with intuitions. For example, intuitions about the deductive closure of knowledge or whether higher-order knowledge goes along with regular knowledge are discussed in relation to skepticism; case studies are debated with respect to epistemic norms; and even discussions about epistemic injustice are often motivated by expository cases. The use of intuitions is also particularly important to philosophers drawing from experimental methods: John Turri (2016, 4) even suggests that we proceed “in the manner of an ethologist” in our epistemological studies. But if intuitions are supposed to provide

justification for our philosophical claims, we need a story that connects those intuitions to the truth of the relevant claims (cf. De Cruz 2015).

Such a story is readily available when we concern ourselves with conceptual analysis. Arguably, our concepts are collective cognitive representations. One of the ways in which these representations reveal themselves is through our verdicts about whether a given application of a word is felicitous in a well-described scenario. Of course, problems can arise: we may not share the same intuitive response in some cases; or we may find that our intuitions are lead by pragmatic features of a given case. But in general, intuitions are plausibly indicative of the nature of our concepts. And, arguably, conceptual analysis is what most participants in the Gettier debate were concerned with. While the project of finding a definition of knowledge ultimately failed, the debate about the semantics and pragmatics of knowledge ascriptions is still continuing and has yielded some plausible suggestions.

But for most of us, there is more we want than a conceptual analysis. Discussions about skepticism are typically framed in terms of the question whether we have *knowledge*, not in terms of whether the concept of knowledge applies to our beliefs about the external world. Similarly, we don't ask for when the concept of justification applies to the relation between our perceptual states and certain proposition – we ask for when our perception *justifies* us. And we ask about the conditions for us *having epistemic rights*, not for the conditions under which the concept of an epistemic right to assert, believe, or act on P applies to an agent. Can we support these claims by appealing to intuitions as well? Timothy Williamson (2004) has argued that we should approach intuitions as simply judgments, and that we are *prima facie* entitled to trust our judgments. But for most of our judgments, we are able to offer some kind of a story that allows us to find them likely to be correct.² So the absence of such a story still looks to be a warning sign.³

2 It is also worth noting that our intuitions are very unreliable in some domains, such as statistics.

3 One may add: even if Williamson is right and our intuitive judgments are reliable, this will at least raise the question *how* they can be reliable. That is, we will then wonder what sort of facts in the world bring it about that our intuitions are reliable guides to the questions mentioned above.

I think it is instructive to look at a model for what an answer to this methodological challenge can look like: Hilary Kornblith's (2004) program of naturalist epistemology. Kornblith argues that knowledge is a *natural kind* which is uncovered by cognitive ethology. According to this view, philosophers trying to understand what knowledge is should do so by paying attention to the circumstances in which ethologists ascribe knowledge to animals. The mental states of these animals form a "well-behaved category, a category that features prominently in causal explanations, and thus in successful inductive predictions. If we wish to explain why it is that members of a species have survived, we need to appeal to the causal role of the animals' knowledge of their environment in producing behavior which allows them to succeed in fulfilling their biological needs" (Kornblith 2004, 60).⁴

If we assume that knowledge is a natural kind and the surrounding naturalist picture, it addresses the methodological issue described above. First, we gain a way of determining what knowledge is: it is a particularly adaptive type of belief that can be understood by engaging with cognitive ethology. Second, we also are able to relate this idea to other epistemological issues. For example, skepticism will appear like a chimera, because it seems clear that we have a fairly wide range of belief that fall under the natural kind knowledge – beliefs that are adaptive, even if we cannot rule out certain remote scenarios of error. Epistemic norms can be traced back to the adaptive benefits of retaining and disseminating knowledge among one's peers. Perception may be approached as a knowledge-producing mechanic, defining its proper functioning. Intuitions take a back seat on this view in favor of a more science-oriented methodology.

⁴ In the background of this view is the idea that natural kinds do not always depend on immediate physical similarity, but rather on functional organization (Millikan 1996). Thus, for two mental states to count as knowledge of P, they do not need to be highly physically similar, but rather must have similar causal origins and/or effects.

There are serious objections to the idea that knowledge is a natural kind and I am not advocating this position here.⁵ But I do think it provides a useful model of how to advance epistemology beyond conceptual analysis. It comes with a very clear ontology of what knowledge is (a natural kind) and how it may relate to other epistemic phenomena (namely in virtue of its adaptive capacity). It also provides resources for justifying one's epistemological claims, namely by appealing to natural science, specifically ethology. I think we should aim to provide a methodological framework for epistemology that can compete with this model.

In a discussion of Kornblith's work, Martin Kusch (2005, 414) suggests that Kornblith should have considered the idea that knowledge can be understood as a *social kind*. This alternative would allow for knowledge to be studied not just by cognitive ethology, but by a range of social sciences. Unfortunately, Kusch's suggestion did not receive much attention at the time. The only other (apparently independent) discussion of this possibility that I am aware of comes from Michael Hannon (2019, 31-3). Hannon expresses openness to the idea of viewing knowledge as a social kind, but he argues that the social kind view ultimately blurs the distinction between concept and kind and appears to consider this claim unnecessary for his philosophical program. I will aim to show below that the idea that knowledge is a social kind is not only plausible, but also offers unique explanatory resources that make this claim attractive.

2. Knowledge as a social kind

Before making the case that knowledge can be seen as a social kind, we will first need to look at some criteria for social kind. I will discuss the case of money, which is both a paradigmatic example of

⁵ One such objection is that naturalist epistemology makes reference to truth, but it remains unclear whether true beliefs are generally conducive to evolutionary fitness (e.g. Pernu 2009). There may also be questions about Kornblith's claim that knowledge as a category has a significant predictive efficacy, which may suggest that other categories are better candidates for being a natural kind (Roth 2003).

a social kind and shares interesting similarities with knowledge. I will then discuss how we develop a standard for the exchange of information and to what extent this standard is subverted by contextual variation through pragmatic encroachment. Finally, I will argue that this standard meets the criteria for a social kind.

2.1 Social kinds

Social kinds belong to the same category as natural kinds, but they are constituted by human activity. While natural kinds include water, ostriches, electrons or planets, social kinds include laws, universities, marriages and money. Ásta (2018, 291) explicates the idea of a kind (whether natural or social) using two features, plus an optional third one:

- (a) Kinds are stable across contexts, as is the membership in a kind.
- (b) Kinds are useful in explanations.
- (c) Kinds play a substantive causal role in the world (and correspondingly in explanations).

For example, the natural kind of an ostrich is stable across contexts: what counts as an ostrich to me will count as an ostrich to anyone else, including in the past and the future. It also helps to categorize animals as ostriches when we want to explain ostrich-typical behavior, or the fact that they cannot be cross-bred with non-ostriches. Feature (c) is optional in the sense that a deflationary account of kinds does not need to embrace (c), but can be satisfied with (b) instead – we can stay neutral on this issue here. Given such an understanding, it follows by definition that if knowledge is a kind – natural or social – it will be useful in explanations. This makes it reasonable to explore the question whether we can see knowledge as a social kind.

Social kinds are an attractive way to categorize knowledge because they provide a stable middle ground between two extremes. On the one hand, entities like natural kinds are completely independent

of what we think of them: water is water, and no amount of change in our thinking about it can change what belongs to this category. As a result, categories at this extreme cannot be explored by relying on evidence like intuitions, at least not in any straightforward way.⁶ For epistemology, this would mean that we would have to radically reorient our methodology. At the other extreme, there are categories that behave relativistically. For example, Searle (1995, 34) suggests that whether a given event counts as a cocktail party depends on our attitudes towards that particular event. But if the criteria for knowledge were this soft, it would be difficult to defend many general claims referring to knowledge. Social kinds are located in between these two extremes because their nature depends on our attitudes, but normally not our attitudes in a specific given case.⁷ A core example of social kinds (also due to Searle 1995) is money. Broadly speaking, it depends on our attitudes whether such a thing as money exists: if we stopped believing in its existence, then there would not be such a thing.⁸ But whether a given piece of paper counts as money does not depend on our attitudes towards that piece of paper. Rather, it will depend on whether it has the right causal history: it must have been produced in a way that is sanctioned by the institutions that we have tasked with producing and managing our money.

Is there a social kind we can call knowledge? I want to argue that there is, namely in the form of a social standard for the quality of information. The idea is that we have created ourselves a standard for when information is reliable enough to pass on and store. Such a standard is set up to create a norm that

6 This is because what constitutes water, for example, does not depend on what we think about it. However, Kornblith (2007) points out that our epistemic intuitions may still be fairly reliable, but only insofar as they have benefited from either expert education or an implicit sensitivity to the natural state of affairs. But, as he points out, it would nevertheless be a mistake to rely on these intuitions *as evidence*. Rather, if these intuitions are indeed reliable, they are so *because of independent evidence*.

7 Muhammad Ali Khalidi (2015) distinguishes three types of social kinds: (a) those that depend on some of our attitudes, but not on our attitudes towards that kind (such as recession); (b) those that depend on our attitudes towards that kind in general, but not on our attitudes in a specific case; and (c) those that depend on our attitudes in specific cases. On his view, cocktail parties are social kinds. I will here work with a more robust notion of a social kind that does not allow cases of type (c). I will be discussing money and knowledge, which both belong to type (b): if we thought that knowledge or money did not exist, they would indeed disappear.

8 This kind of connection is why Ian Hacking (1995) talks of a “looping effect” with respect to “human kinds”. The idea is that our classificatory practices lead to changes in our behavior, which in turn can lead to revisions in the classificatory practices. In the case of money, this effect can describe how we have sometimes narrowed our classification of money (e.g. by no longer accepting tender money) and sometimes widened it (e.g. by accepting cryptocurrencies).

stops us from sharing false and poorly justified information. This standard is contextually stable and useful in explanation, fulfilling the basic conditions for a social kind noted above.

But these criteria are fairly minimal. I think it will make our case much stronger if we take a look at the question *why* social kinds are contextually stable and important in explanations. To do so, let us return to the case of money. Why do we have money? In economics, money is often associated with three functions:⁹

Medium of exchange: money allows us to exchange goods and services in a way that would not otherwise be possible. For example, without money a baker would only be able to exchange her bread for goods that are more or less commensurate with bread in terms of their cost in labor and resources. For example, she may be able to exchange bread for vegetables, but a shoemaker will not accept a hundred loafs of bread for a pair of shoes.

Store of value: Money allows us to store value for later use. For example, our baker must trade her bread today or it will rapidly depreciate. Money does not perish and can (normally) be expected to have about the same exchange value over time.

Unit of account: Money allows us to measure the value of things. It would be much more difficult to name the value of a pair of shoes on terms of loafs of bread for many reasons, whereas putting a monetary value on it is much easier. The ability to do so is useful for things like calculating taxes or insuring goods.

These three functions can give us a lot of insight into the nature of money. For one thing, they can explain why we have money in the first place: because it offers us a range of practical benefits. They can also explain the contextual stability of money. When the value of money floats, as in cases of hyperinflation, this takes away our ability to enjoy any of the benefits of these three functions. And finally, these functions lend money as a category its explanatory power: money can explain how certain exchanges happen, how value is taken out of a system or injected into it, and how the value of different things compare to each other.

2.2 Knowledge as a Standard

9 This characterization goes back to Wiiliam Jevons (1875).

I think we can say quite similar things about the function of a standard of information. To see this, we need to consider the functional role of such a standard as well. Fortunately, important resources for such an analysis can be drawn from the work of Edward Craig (1990) and the research tradition following him. While this research tradition is concerned with explicating the *concept* of knowledge, I think that most of the core ideas discussed within it will translate nicely to the idea of a social standard. On the Craigian view, the ur-function of a concept of knowledge is to evaluate information¹⁰ under the guise of a decision we need to make: we are selecting information that is good enough for us to assume it to be accurate in making that decision. As Craig points out, the information in question must not only be likely to be true, but must be *recognizable* as being likely to be true – otherwise, we cannot evaluate it positively. This gives rise to a notion of “protoknowledge” (using the terminology of Kusch 2009): information that is recognizably of sufficient quality to base our current decision on.

The notion of protoknowledge is entirely local: it is used to evaluate information only with the current decision problem in mind. But different decisions will impose different standards to hold information against. One well-studied aspect of this are the “stakes” of a decision: in Keith DeRose’s (1992) famous bank cases, the question whether we count ourselves as knowing that the bank is open on Saturday turns on how urgently we need to visit it before the end of the week. DeRose is concerned with the concept of knowledge, but his examples also demonstrate that “stakes” can affect which information we are happy to rely on and which not. In addition to “stakes”, it also matters what range of information is available to me and whether I have time and resources to seek out more or better

10 Craig’s original position is that this ur-function is to evaluate potential *informants*. I frame the discussion here in terms of evaluating *information*, roughly following Klemens Kappel (2010) and Christoph Kelp (2011), who advocate what Patrick Rysiew (2012, 275) calls the “certification view”. On this view, the point of the concept of knowledge is to “certify” information. The general idea is that focusing on information allows for a broader view, which accommodates acts like evaluating one’s own perception. Informants can be evaluated too, of course, but on this conception they are evaluated as potential providers of information. Not much turns on the question whether protoknowledge is primarily concerned with informants or information for our purposes here – if anything, Craig’s informant-based conception supports a more social view of knowledge.

information. To take a version of Craig's (1993, 98) examples: if I know that my train is departing in two minutes, but I am unsure about its platform, I will take a stranger's best guess and rush up the stairs to that platform; but if I have ten minutes, I will take a moment to consult the posted schedule. In essence, protoknowledge is fine-tuned to the needs and means of the current context.

Such a local evaluation may be feasible and useful in some contexts, but it also has serious drawbacks: life is a series of decisions, many of which will differ in their needs and the means available in the current context. But this means that we cannot rely on our previous evaluation of information once we are in a new context. This alone gives us reason to evaluate information in a way that will allow us to make use of that evaluation in the future. We may imagine a system that classifies some information as "certain", other as "likely true", yet others as "viable hypotheses", and so forth. Once we have built up a kind of database of information classified in these ways, we could easily access the relevantly "certain" information when we encounter a new decision problem. So in the long run, we are better off developing classifications that evaluate information independently of our local situation of needs and means.

But there is an even more significant factor that drives us away from the local evaluation of information: our informational state can be massively improved if we establish a practice of sharing useful information with others. To use another of Craig's (1990, 11) examples, Fred sitting up in a tree is better positioned to tell whether a tiger is approaching than Mabel sitting in a cave – and she would benefit from receiving such information from him. But sharing information is only useful if we can count on it to be reliable. More precisely, we need to be able to classify the reliability of the information we received: can we count on it to be "safe", or at least to be likely correct? If we receive information that we are unable to classify in any way, we cannot use it in our own decision-making. Some system of norms or practices is required for the exchange of information to be useful.

What can such a system look like? We could imagine a normative system that requires us to provide each other only with information that meets the standard of protoknowledge. But such a system would have several drawbacks: first, it would require us to familiarize ourselves with the decision problem the potential recipient of information is facing – only then would we be able to assess whether the information we are about to provide is sufficiently reliable. And second, it would make it difficult to provide information outside of the context of a particular decision, because there would not be a local standard to hold that information against. All we could do is to provide information alongside a fairly extensive amount of “metainformation”, information about the origin or reliability of the original piece of information. Both of these problems make a protoknowledge-based norm of assertion a practical impossibility.

What we need here is a simpler standard: a widely recognized stable threshold for the quality of information that we can apply to all shared and stored information.¹¹ Such a standard can allow us to simply state information, without including metainformation or familiarizing ourselves with the recipient’s epistemic needs, and will allow the recipient to assume that the information is safe to use – if the speaker is abiding by epistemic norms. This standard will need to be fairly high: the information it permits must be such that it is indeed useful in most, if not all, decision contexts. Because the system described here gets rid of the need to include metainformation, it makes it very hard to re-evaluate the information in situations that require more than the common standard. For example, suppose I am a safety inspector at a nuclear power plant. In doing my job, I will only want to rely on information that is highly reliable. But if our social standard of good information was lower than what I require, I would not be able to freely draw from the body of information shared with me under that standard. Instead, I would need to re-evaluate each piece of information individually, which would require me to memorize the relevant metainformation – did I learn it in college, read it in a textbook, etc.?. To avoid situations

11 A very nice account of how such a threshold can be developed that also draws from Craig and runs along similar lines has been given by Hannon (2017).

like this, establishing a very high standard of information is desirable. Craig describes the process of broadening the usefulness of such a standard as the “objectivization” of knowledge.

2.3 Practical Factors

The line of reasoning above points us to the standard of knowledge being stable across a broad range of circumstances. Only this kind of stability will allow us to pass on information without first familiarizing ourselves with the decisions in which it may be used. But as debates about pragmatic encroachment have shown, certain practical factors of a situation *are* relevant to what counts as knowledge. In particular, the “stakes” of a context seem to have an impact on whether we are happy to credit ourselves and others with knowledge.¹² In Keith DeRose’s (1992) classical example my judgment whether I know that the bank will be open tomorrow will in part depend on how important it is to me to be able to deposit a check tomorrow. The more important the deposit, the better my evidence needs to be in order to grant me knowledge.

Practical factors push us back to what I have called a “local evaluation” above: we go back to judging information in light of our needs and means in the current circumstances, rather than in light of what constitutes an acceptable threshold of reliability in general. But the existence of a standard is not inconsistent with the existence of specific circumstances in which the threshold specified by that standard is temporarily raised or lowered. In a sense, the standard of knowledge is what practical factors encroach upon.

12 This statement – that “stakes” can affect our judgments about the felicity of knowledge ascriptions – connects pragmatic encroachment with a form of contextualism. In standard usage, contextualism about knowledge ascriptions is concerned with the context in which the knowledge ascription is uttered. Examples of sensitivity to “stakes” are associated with so-called subject-sensitive invariantism (SSI). At the heart of SSI, though, is the claim that the truth-conditions of knowledge ascriptions are sensitive to the *subject’s* context. Nevertheless, it is important to separate SSI and the idea of pragmatic encroachment from contextualism (in the sense described above), because the context of utterance and the subject’s context need not be identical.

Consider, by way of analogy, our standard of safe drinking water. This specifies to what extent water may (and must) contain other elements and germs in order to be deemed safe for unfiltered consumption. It is specified in light of what is best for human health, but also in light of what can be realistically expected of water treatment plants, and therefore allows for a small degree of contamination. However, there are circumstances in which we require water to be uncontaminated beyond this threshold – for example, we may want to supply certain vulnerable patients with water that is uncontaminated to an even higher standard. There may also be contexts in which we are willing to compromise on the quality of drinking water, for example because exceptional circumstances have impacted our ability to purify a sufficient amount of water. Doing so does not require us to reject the general standard or to endorse a context-sensitive version of it: the benefits of a general standard remain in place as long as the exceptions to it are relatively isolated.

Proponents of pragmatic encroachment may argue that this type of encroachment happens regularly, not just in exceptional circumstances. If this were indeed the case, it would be a problem for the idea of a fixed standard. We would not gain the benefits of a stable standard if it were regularly disregarded or altered – this goes for an epistemic standard as much as it goes for a standard of money or safe drinking water. However, I do not think it is plausible that we constantly shift around our standard of information. It is important to distinguish three things that may be subject to change here: (a) our threshold for knowledge ascriptions, (b) the threshold for passing on information to others and (c) the threshold for using information in our practical reasoning. Pragmatic encroachment can affect all three of these¹³, but the standard I am interested in here is concerned with (b) and can affect (a) and (c) only

13 Contributions in the debate about pragmatic encroachment have focused on different aspects. A good number of papers address the semantics or pragmatics of knowledge ascriptions (e.g. De Rose 1992; Hawthorne 2004; Stanley 2005; Fantl & McGrath 2009). In the more recent debate, practical reasoning has played a bigger role (e.g. Hawthorne & Stanley 2008; Brown 2013). The requirements for passing on information seem to feature less prominently, although they curiously are addressed in Fantl & McGrath's (2002) foundational treatment of pragmatic encroachment. Hannon's (2017) recent treatment of the "threshold problem" addresses this aspect, but is criticized for lacking an account of the conditions for practical reasoning (Langford 2023).

indirectly. And with respect to the way we pass on information, it seems that we only differ from the quality standard after some conscious deliberation. We would also typically hedge our assertions when passing on sub-standard information (“I’m not sure, but I seem to remember that...”) or when we feel that information of higher quality is required (“I would normally tell you that P, but in this case we should probably check to make sure.”). So, it seems that cases in which we adjust our standard are relatively rare. However, I will discuss the danger of increasing amounts of misinformation and the possible dissolution of the standard in section 3.3.

2.4 Knowledge as a Social Kind

Craig’s story about the emergence of protoknowledge in a “state of nature” and its “objectivization” from there on does not need to be understood as an origin story. There are rival interpretations of his account¹⁴, and I would like to stay as neutral as possible on how exactly to read it. At a minimum, I think Craig’s account can provide a rationalization for establishing a standard for the quality of information and can illuminate the pressures that drive the negotiation of this standard. The extent to which this story can also be seen as a viable theory of the historical development of the standard can be left open here.

Of course, Craig thought that what was being developed was a concept, not a standard.¹⁵ I have been tracing back an account derived from Craig and the literature in his tradition, but I have framed it in terms of the development of a social standard rather than a concept. This translation is rather straightforward, because Craig’s genealogy is implicitly committed to the development of a standard as

14 A good overview is provided by Queloz (2021, 12-13).

15 At one point, Craig (1990, 3) asks: “Couldn’t it be that knowledge, like water, is common and important stuff, and that the purpose of the concept is simply to enable us to think and talk about it?” However, he then rejects this hypothesis, because knowledge, unlike water, “is not a given phenomenon, but something that we delineate by operating with a concept[.]” It may be arguable that Craig may have been more sympathetic to the idea that knowledge is a social kind, but that he lacked this idea as a conceptual resource.

well, namely the standard for information that the concept of knowledge names. On the Craigian view, the point of the concept of knowledge is to “flag” good information. But this presupposes that we are interested in collectively separating good from bad information. Such a separation must happen according to some standard or other. Giving a Craigian account of the concept of knowledge therefore assumes that there is an underlying standard which is developed alongside it.¹⁶

With this, we have developed a functional understanding of knowledge as a standard of information. Is it a social kind? It seems that our standard of knowledge, on our understanding, not only satisfies the basic criteria of contextual stability and being a category that is useful in explanations; it even shows striking similarities to money. Just like money, this standard facilitates a certain type of exchange, namely the exchange of information. As we saw, assuming that our conversational partner is adhering to a quality standard in her assertions allows us to take that information on board very easily. And just like in the case of money, it also helps us store information. That is, once a piece of information is understood to be a piece of knowledge, it can be safely remembered for future use. That is, our standard of knowledge has features that correspond to two of the three functions of money.

And just like in the case of money, these functional features are the grounds for having the features of a social kind. That is to say, the functional usefulness of knowledge *is why we have it*. We create a standard of information not as an arbitrary rule, but because it allows us to do certain things and makes our lives easier. The functional features are also responsible for the contextual stability: only a stable standard can allow us to assert information without having to evaluate the current parameters; and only a stable standard can allow us to access information stored earlier and be confident that it is still in compliance with our standard. And thirdly, the functional features are at the heart of the explanatory power of knowledge. For example, our standard of information can explain why we assert some things

¹⁶ Note that the converse is not true. It would be possible to give a functionalist account of a standard for information that remains silent about the concept of knowledge. Of course, it is *plausible* that our concept of knowledge tracks our standard of information. But the social kind view is at least in principle open to some differences.

and not others, or why we sometimes hedge our assertions by adding things like “I suspect that...”. It can also explain why we may get upset about someone violating this standard by spreading unreliable information. And it can explain our practices of storing information in containers like books; why we value libraries that retain these books indefinitely, but also why we set up procedures to control the quality of the information contained in these books.

What all these observations come down to is this: knowledge can be seen as a social standard for the quality of information that has all the typical features of a social kind. This provides us with an alternative to viewing knowledge as a concept: we can view it as a social kind. This claim can give us grounds for talking about “knowledge itself”, as contrasts with the concept of knowledge or individual pieces of knowledge. It can also give us a story about the way in which we have access to what constitutes knowledge. Knowledge is the outcome of an implicit social negotiation for a quality standard of information. This means that we should have some sense of what the outcome of this process should be. But it also allows us to analyze the process of negotiation itself, which opens up a critical perspective: we can explore what sort of features a standard of information would ideally have and compare this with the evaluation of information in our actual behavior. I will explore the usefulness of this perspective more in the next section.

3. Some Applications

I have argued that we can plausibly view knowledge as a social kind, namely as a socially constructed standard of information. One may take this claim to be simply a methodological note: as I outlined in section 1, those philosophers concerned with more than the conceptual analysis of “knowledge” are under pressure to provide an ontological clarification that allows them to justify either their use of intuitions or provides them with another methodology. The claim that knowledge is a social

kind can provide such a clarification. But at the same time, it keeps knowledge tied in with our ordinary thinking about it: it seems arguable that the standard of information I have outlined is tracked by our concept of knowledge, perhaps sparing for some of the contextual variation in that concept. If that is so, one may take the idea that knowledge is a social kind to be a small methodological note, perhaps even one that we should only take on as “baggage” to our theories if we are directly concerned with epistemological methodology. This is close to Hannon’s (2019) view, and it is the reason why he leaves the idea that knowledge may be a social kind aside after a few pages.

In this section, I want to show that the social kind view is more than a methodological note. I will discuss three examples of areas in which this view can provide new perspectives that other approaches, such as those concerned with concepts, cannot offer.

3.1 The Knowledge Norm of Assertion

In its standard form, the knowledge norm of assertion (KNA) states that one has the epistemic right to assert that P only if one knows that P. The restriction to an *epistemic* right is important here: asserting that P may be wrong for all sorts of other reasons such as being impolite or irrelevant. Having the epistemic right to assert that P merely means that one’s *epistemic standing* is good enough to make that assertion. It is also worth noting that the KNA only covers what Keith DeRose (2002, 180) calls the “primary propriety” of an assertion: it specifies which assertions one must *attempt* to avoid; but one may faultlessly violate that norm when, for example, one asserts a false proposition which one had very good reasons to think to that one knew.¹⁷ However, in cases like these one is at least obligated to retract

¹⁷ By contrast, most of the competing accounts of a norm are concerned with what DeRose calls the secondary propriety of assertions, i.e. the conditions under which we can blamelessly assert that P (at least as far as epistemic norms are concerned). These competing suggestions do not require the truth of P, but instead require things like it being reasonable to believe that P (Lackey 2007), having a certain kind of justification for P (Kvanvig 2009), having appropriate context-relative warrant for P (Gerken 2012), or being justified to believe that one knows that P (Neta 2009). DeRose’s distinction seems to dissolve at least part of the disagreement between advocates of these accounts and advocates of the

one's assertion once one learns that it was untrue or poorly justified. Finally, the KNA is typically only stated as giving a necessary condition for having the epistemic right to assert. This is due to counterexamples stemming from what Jennifer Lackey (2011, see also Brown 2010) calls "isolated second-hand knowledge". To take an example from Adam Carter and Emma Gordon (2011): a professor writing a letter of recommendation for a student may not have the epistemic right to praise the student's writing skills without having actually sampled her writing – even if the professor has it on very credible testimony from a colleague that the student's writing skills are indeed excellent.

If we accept the idea that knowledge is a standard for the quality of information that determines which information should be shared or stored, the knowledge norm of assertion drops out of this idea directly. On the view that I am suggesting, knowledge simply is whatever standard for sharing information we implicitly agree on. But given that we agree on it to be the standard, we are thereby bound to adhere to it, so as to not violate our agreement. So the social construction of a norm of assertion and the social construction of knowledge are one and the same. This account of the KNA may not seem like much of an explanation, but it does shed light on two of its features mentioned above.

First, we can compare the focus on primary propriety to the case of money: if I pay you, I need to attempt to do so with genuine money, rather than counterfeit. However, it is possible that I nevertheless pay you with counterfeit money which I take to be genuine. In this scenario, my mistake is blameless – but it arguably does give rise to an obligation on my part to accommodate your losses, for example by paying you again with genuine money. So it makes sense to distinguish between primary propriety and secondary propriety (i.e. the question whether an action carries fault in one's own) in the case of paying money. Why, then, should we not do the same in the case of assertion?

KNA. In addition, James Beebe (2012) argues that it is beneficial for us to distinguish blameful and blamelessly harmful epistemic behavior, and that the concept of knowledge allows us to draw this distinction. This would support focusing on the "primary propriety" of assertions and using the KNA to capture it.

Second, we can revisit the case of “isolated second-hand knowledge”. In the case of the professor writing a letter of recommendation, it intuitively seemed like she did know that the student has excellent writing skills, but was not allowed to state this due to her lack of first-hand experience. One possible explanation is this: the context of a letter of recommendation places expectations on assertions that supersede the general standard. That is to say, in these cases we expect letter writers to be in an epistemic position that is, at least in some respect, stronger than the standard we have commonly negotiated. The situation would then be somewhat like this: there is a global norm of assertion which requires knowledge (and nothing else), and there is a more localized norm which applies specifically to the context of letters of recommendation and which requires us to provide only first-hand information. Similar norms may be in place, for example, in situations like giving testimony in court. But this would still leave the KNA in a special position: it would be the only global norm of assertion, because it (and only it) applies to all assertions, regardless of context.¹⁸

Overall, the social kind view can give a much more direct account of why knowledge is the norm of assertion. On our view, knowledge is a standard that is specifically established for the purpose of governing assertion. On other views, the KNA is a self-standing norm: we agree, by an implicit understanding, that knowledge (rather than some other benchmark) should be a requirement for assertion. This is by no means absurd: such an agreement would have obvious benefits some of which we have discussed in the previous section. But viewing the KNA as a self-standing norm requires an additional explanatory move in comparison to the social kind view I am advocating here. The idea that

18 A problem for the KNA would arise if there were norms of assertion that overrode the requirement that assertions must at least be knowledge. That is, if there were norms that allowed unhedged genuine assertions without knowledge in certain contexts, this would raise questions about whether the KNA is in place as a general norm, rather than just as a somewhat more extended norm. Note, however, that not every utterance of a declarative sentence counts as a genuine assertion. For example, the following are not assertions: utterances made while acting; utterances made in *propria persona*; utterances recognizably made as a conjecture or hypothesis; etc.

the expectations on assertion *are part of what knowledge consists in* is, in my view, a more elegant and direct explanation of the KNA.¹⁹

3.2 Testimonial Injustice

Let me turn to an application of the social kind view that brings out a different side of its benefits which concerns Miranda Fricker's (2007) discussion of *testimonial injustice*. This is a natural place to look for a fruitful area of applying this idea because Fricker relies on Craig's functionalist account in her discussion. Her general goal is to discuss testimonial injustice as a distinctively epistemic type of injustice that cannot be reduced to distributive injustice (Fricker 2007, 1). She achieves this by presenting testimonial injustice as a kind of injustice that harms someone *in their capacity as a knower* (Fricker 2007, 20). The idea is that the harmed person is assigned an unjustly low level of credibility, paradigmatically as a result of what Fricker (2007, 27-8) calls an *identity prejudice*. The idea is that identities are assigned stereotypes which can negatively affect the credibility assigned to carriers of that identity, leading to what she calls "identity-prejudicial stereotypes" (Fricker 2007, 35). Testimonial injustice occurs when these stereotypes lead to an unjustified systematic under-rating of credibility on members of that group. She presents the case of Tom Robinson in *To Kill a Mockingbird* as a central example of this: a black man is, despite good evidence in his favor, not believed and falsely convicted

19 The account presented here shares some features with Hannon's account of the apparent coincidence between the concept knowledge and assertibility. Hannon (2019, 115) argues:

The point of the concept of knowledge, I claim, is to certify reliable informants to members of our epistemic community. [...] Our practical concerns generate a standard that is fitting to certify information that is good enough for the practical reasoning situations of many people with diverse interests. [...] The level of warrant needed for knowledge is that which puts the agent in a strong enough epistemic position for her to serve as a reliable source of actionable information for members of her epistemic community.

The main differences between Hannon's account and mine are that (a) Hannon begins with the concept of knowledge, taking a detour to the idea of a normative standard and (b) that Hannon thinks that the relevant standard is a simple threshold that needs to be "overruled" in cases like the examples of "isolated second-hand knowledge" referred to above. On my view, standards are more than simple thresholds: they can have built-in exceptions that demand that extra epistemic features can be demanded in situations like these.

of murder. However, Fricker (2007, 44) argues that things like being falsely convicted is only a practical secondary harm caused by an epistemic injustice; the primary harm consists in wronging someone in their capacity as a knower, thereby treating them as less than fully human.²⁰ That is not to say that the primary harm is more significant – in the case of Tom Robinson, it is arguably less significant. But the primary harm is a more direct result of the action, one that is independent of the broader circumstances.

Fricker uses Craig's functionalist genealogy to bring out why testimonial injustice causes a fundamental intrinsic harm. She writes (Fricker 2007, 145):

If the core of our concept of knowledge is captured in the concept of the good informant, because (as the State of Nature story shows) essentially what it is to be a knower is to participate in the sharing of information, then another dimension to the harm of testimonial injustice now comes into view. When someone is excluded from the relations of epistemic trust that are at work in a co-operative practice of pooling information, they are wrongfully excluded from participation in the practice that defines the core of the very concept of knowledge.

This is what Fricker means by the phrase “being wronged in one’s capacity as a knower”: it is an exclusion from a practice that lies at the core of the concept of knowledge. So, whatever the practical offshoots of that may be, there is a type of harm associated with the exclusion itself which comes out based on Craig’s understanding of knowledge.

While Fricker relies on the concept of knowledge, rather than a social kind knowledge, I think that her account can be enriched by viewing knowledge as a social kind. The core issue for Fricker is to bring out the nature of the “epistemic harm” that occurs when one’s testimony is wrongfully not taken seriously. To make the case that epistemic injustice is more than a form of distributive injustice, she needs to make the case that this harm does not ultimately boil down to some form of unfair distribution

20 Aside from this perhaps abstract-sounding dimension of the primary harm, Fricker (2007, 51-5) argues that being harmed in this way will mean that the subject can no longer engage in what Bernard Williams calls “trustful conversation” with the person who wronged them. The idea of trustful conversation is that it is a deeply personal type of conversational reflection that “steadies the mind” and allows one to gain an understanding of oneself. If the wronged subject lacks a community to engage in trustful conversation with, they are inhibited in forming their identity.

of goods. She pins the epistemic harm of testimonial injustice to the exclusion from a practice of sharing information. But it is not immediately clear that this kind of exclusion would constitute a harm that extends beyond potential material disadvantages, which would mean that it could be seen as a case of distributive injustice.

However, when we recognize knowledge as a social kind, we gain a resource for articulating the harm in question more clearly: it may be that the epistemic injustice we are looking at consists in the wrongful denial of recognition of the membership of one's beliefs in the kind of knowledge. That is to say, even though one's beliefs meet all the commonly negotiated criteria for knowledge, they are not treated as such – and this constitutes a metaphysical miscategorization. Such a miscategorization can constitute a harm even beyond its material consequences.

To bring this out, consider the following example: Adil was trained as a doctor in Syria at a highly reputable university, passing all his exams and practicing at the hospital for many years. With the outbreak of the civil war, he is forced to flee to Europe. While he is able to secure refuge and a work permit in Germany, his qualification as a doctor, despite being well-documented, is not recognized. Instead, he is granted a certification that will allow him to work as a nurse. This is not in view of the quality of his training – which holds up to German standards – but simply in view of the fact that the qualification was gained in the wrong country. To be sure, this means that Adil suffers a material harm: he will not be able to work as a doctor, which means that he misses out on a higher salary and some other potential benefits, like reputation. But even beyond this, it seems that not granting him the status as a doctor, even though he meets the relevant standards, constitutes an inherent injustice; one that would exist even if Adil had no intention of working as a doctor. There appears to be an injustice that consists purely in the denial of a status despite meeting the eligibility criteria that should apply to all people equally. If we view knowledge as a social kind similar to that of a doctor, we can see how the same kind of injustice could be in play in cases of wrongful denial of knowledge.

We can also make some sense of Fricker's remark that testimonial injustice amounts to treating someone as "less than fully human": knowledge is socially constructed by humans as a category primarily applied to humans – animals may possess "knowledge" in some sense, but they are not part of the original practice of sharing information that lead to the establishment of the standards for knowledge, and thus those standards may not straightforwardly apply to them. But to deny the application of those standards to a human is to place them outside of the boundaries of the construction of those standards – which means: outside of humanity. By giving a metaphysical dimension to this type of injustice, the social kind view can provide us with a better explanatory resource for understanding testimonial injustice than approaches grounded in the conceptual analysis of knowledge alone.

3.3 Dissolving Standards

In the preceding discussion I have argued that knowledge is a socially constructed standard for the quality of information that is implicitly accepted in our society. I believe this to be true; but it need not be true forever. Standards can vanish or change in nature, and there is at least a risk that some such change might happen with respect to knowledge. There seems to be an increase in the (sometimes deliberate) sharing of misinformation and there are at least some circles in which asserting claims with little justification is accepted, as long as those claims feed into a larger narrative. While I am not attempting to give anything like an empirical account of these changes here, I want to provide some cursory remarks on how the social kind view of knowledge can help us better understand the mechanics of dissolving knowledge as a standard of information.

One reason to worry about the spread of misinformation is its instrumental damage. If we hear that the government is run by a clandestine organization with evil goals, we may end up believing this to be

true and may conclude that these circumstances justify political violence. Similarly, telling others that drinking bleach can cure Covid-19 may potentially lead to life-threatening decisions. These are serious worries, but they do not apply to all types of misinformation equally. For example, the direct effects of spreading the claim that Loch Ness is inhabited by an ancient sea monster appear to be less serious. At worst, it seems, a credulous listener may spend a relatively uneventful vacation in Scotland searching for said monster.

But there is nevertheless something problematic about the spreading of even such seemingly benign misinformation, and the social kind view brings this out. The threat lies in the dissolution of an epistemic standard. Standards in general only continue to exist as long as they are by and large being adhered to. It may be true that certain safety standards exist only “on paper”, but are not actually implemented. But even if such “standards” can be said to exist as actual standards, it is only because they are maintained (or at least “kept on the books”) by an institution. Epistemic standards are not explicated or enforced by institutions. As such, they depend on whether they are perceived to be the norm by our society at large. Therefore, the standard of knowledge depends on whether the information we actually exchange *normally* meets certain quality criteria.

The word “normally” here does a lot of work here. What matters is not just the quantity of misinformation (or information that is of poor quality) being spread. It also matters to what extent those having spread misinformation are willing to take it back when it has been exposed; and to what extent we are willing to push back when misinformation is disseminated. We may be able to maintain our standard even in the face of relatively frequent transgressions if those transgressions are at least reprimanded and corrected. While the overall picture is complex and merits future research, the general trend is, I hope, clear: if the dissemination of information that does not meet our standard becomes normalized, the quality standard will either change or completely disappear. Importantly, this

normalization can happen even through the spreading of misinformation that is otherwise inconsequential, such as claims about “Nessie”.²¹

In this way, the social kind view can explain the danger associated with spreading misinformation without having to appeal to the immediate consequences of belief in that information. It also does not need to appeal to an apparatus like Quassim Cassam’s (2016) epistemic vices. It is more closely aligned with the idea that epistemic trust is important to a society (Fuerstein 2013) and that the spreading of misinformation can lead to an erosion of such trust (e.g. Reglitz 2022). Once I do no longer assume that others are adhering to quality standards in what they assert, I will also stop trusting the information they deliver to be accurate. And it should come as no surprise that the failure of a standard of information means that the goals that it was created to achieve will also no longer be reached.

A useful aspect of the social kind view here is that it allows us to compare knowledge with other standards. Let us consider dress codes in academic philosophy: while these have always been evolving and differ, it seems that they have shifted towards being quite informal over the past decades. What drove this development? Very roughly, I think some philosophers, predominantly in the analytic tradition, have made a choice to wear T-shirts to conferences. This, I think, made a statement about being unwilling to “dress up” not only oneself, but also one’s philosophical content. Such a statement may have expressed a desire for simplicity or a contempt for rhetorical devices. Following this, more have adapted to this dress code, until eventually wearing everyday clothing was no longer sub-standard, but entirely accepted. I think it is plausible that a development like this is possible for knowledge: while initially violations of our standards of information may “stick out”, this will no longer be the case once they are widespread enough. We will attend to a few individual cases of standard violation, but

21 In two recent papers, Christoph Blake-Turner (2020) and Jeroen De Ridder (2021) also try to explicate the damage done by the spread of misinformation, beyond the direct effects on our decisions. They both argue that it leads to a degradation of our epistemic environment to the effect that we cannot rely on the testimony of others as much as we would like. I essentially agree with their core argument. However, I think the idea that an epistemic standard is dissolving can explain how exactly our epistemic environment is degraded.

once these former violations become commonplace, we may at best lament how things used to be better in the old days.

Are there other models for changes in standards that can make us more hopeful? Some standards are institutionally enshrined and can be enforced by those institutions, such as building safety standards. As long as those standards are maintained and enforced by the institutions, they usually stay in place. There is a parallel for this in the case of knowledge, namely the way we penalize lying under oath. But even the enforcement of this quite limited standard often turns out to be difficult. And what is more, this does not seem to be a feasible model for a democracy, because implementing something like this would provide institutions with tools to silence speech that is politically inconvenient. A perhaps more hopeful comparison case are standards that are maintained through education. Many forms of etiquette, such as table manners, are maintained through explicitly teaching them to children or novices. It is at least conceivable that educational campaigns could have an effect in maintaining the standard of knowledge.

My goal here is not to solve our misinformation crisis. But I hope to have provided a hint of how viewing knowledge as a quality standard for information can provide us with a useful perspective on changes in the way we share information and carve out the significance of these changes. In general, I hope to have made the case that it is not only plausible to view knowledge as a social kind, but also that this view addresses both an important methodological question and, perhaps more importantly, can provide us with impulses and new perspectives on some other epistemological problems.

Acknowledgments

The core ideas of this paper stem from my dissertation at Johns Hopkins University. I would like to thank anyone who has provided input on it, particularly Michael Williams. For valuable comments on

this specific paper I would like to thank Elanor Taylor, Ben Winokur and the audience at the Social Ontology 2022 conference at the University of Vienna.

References

- Ásta (2018). Social Kinds. In M. Jankovic and K. Ludwig (Eds.), *The Routledge Handbook of Collective Intentionality* (pp. 290-299). New York: Routledge.
- Beebe, J. (2012). Social Functions of Knowledge Attributions. In J. Brown and M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 220-242). Oxford: Oxford University Press.
- Blake-Turner, C. (2020). Fake news, relevant alternatives, and the degradation of our epistemic environment. Forthcoming in *Inquiry*. DOI: 10.1080/0020174X.2020.1725623
- Brown, J. (2010). Knowledge and Assertion. *Philosophy and Phenomenological Research* 81(3), 549-566.
- Brown, J. (2013). Impurism, Practical Reasoning, and the Threshold Problem. *Noûs* 47(1), 179-192.
- Carter, A. and E. Gordon (2011). Norms of Assertion: The Quantity and Quality of Epistemic Support. *Philosophia* 39(4): 615-635.
- Cassam, Q. (2016). Vice Epistemology. *The Monist* 99(2), 159–180.
- Craig, E. (1990). *Knowledge and the State of Nature*. New York: Oxford University Press.
- Craig, E. (1993). *Was wir wissen können: Pragmatische Untersuchungen zum Wissensbegriff*. Frankfurt a. M.: Suhrkamp.
- De Cruz, H. (2015). Where Philosophical Intuitions Come From. *Australasian Journal of Philosophy* 93(2), 233-249.
- De Ridder, J. (2021). What's so bad about misinformation? Forthcoming in *Inquiry*. DOI: 10.1080/0020174X.2021.2002187.

- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52(4), 913-929.
- DeRose, K. (2002). Assertion, Knowledge, and Context. *The Philosophical Review* 111(2), 167-203.
- Fantl, J. and M. McGrath (2002). Evidence, Pragmatics, and Justification. *The Philosophical Review* 111(1), 67-94.
- Fantl, J. and M. McGrath (2009). *Knowledge in an Uncertain World*. New York: Oxford University Press.
- Fuerstein, M. (2013). Epistemic Trust and Liberal Justification. *Journal of Political Philosophy* 21(2), 179-199.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. New York: Oxford University Press.
- Gerken, M. (2012). Discursive justification and skepticism. *Synthese* 189(2), 373-394.
- Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 351–394). New York: Oxford University Press.
- Hannon, M. (2017). A solution to knowledge's threshold problem. *Philosophical Studies* 174(3), 607-629.
- Hannon, M. (2019). *What's the Point of Knowledge?: A Function-First Epistemology*. New York: Oxford University Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. New York: Oxford University Press.
- Jevons, W. (1875). *Money and the Mechanism of Exchange*. New York: D. Appleton and Company.
- Kappel, K. (2010). On Saying that Someone Knows: Themes from Craig. In: A. Haddock, & A. Millar (Eds.), *Social Epistemology* (pp. 69-88). Oxford: Oxford University Press.

- Khalidi, M.A. (2015). Three Kinds of Social Kinds. *Philosophy and Phenomenological Research* 90(1), 96–112.
- Kornblith, H. (2002). *Knowledge and its Place in Nature*. New York: Oxford University Press.
- Kornblith, H. (2007). Naturalism and Intuitions. *Grazer Philosophische Studien* 74, 27-49.
- Kusch, M. (2005). Beliefs, Kinds and Rules: A Comment on Kornblith's *Knowledge and Its Place in Nature*. *Philosophy and Phenomenological Research* 71(2), 411-419.
- Kvanvig, J. (2009). Assertion, Knowledge, and Lotteries. In P. Greenough and D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 140-160). Oxford: Oxford University Press.
- Lackey, J. (2007). Norms of Assertion. *Noûs* 41(4), 594–626.
- Lackey, J. (2011). Assertion and Isolated Second-Hand Knowledge. In J. Brown, & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 251-276). Oxford: Oxford University Press.
- Langford, S. (2023). Pragmatic Encroachment and the Threshold Problem. *Erkenntnis* 88(1), 173–188.
- Millikan, R. (1996). On Swampkinds. *Mind & Language* 11(1), 103-117.
- Neta, R. (2009). Treating Something as a Reason for Action. *Noûs* 43(4), 684-699.
- Pernu, T. (2009). Is knowledge a natural kind? *Philosophical Studies* 142(3): 371-386.
- Queloz, M. (2021). *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*. Oxford: Oxford University Press.
- Reglitz, M. (2022). Fake News and Democracy. *Journal of Ethics and Social Philosophy* 22(2), 162-187.
- Roth, P. (2003). Review of “Knowledge and its Place in Nature”. *Notre Dame Philosophical Reviews*. Available online. <https://ndpr.nd.edu/reviews/knowledge-and-its-place-in-nature/>. Cited 23 May 2023.
- Searle, J. (1995). *The Social Construction of Reality*. New York: The Free Press.

Stanley, J. (2005). *Knowledge and Practical Interests*. New York: Oxford University Press.

Turri, J. (2016). *Knowledge and the Norm of Assertion*. Cambridge: OpenBook Publishers.

Williamson, T. (1996). Knowing and Asserting. *Philosophical Review* 105(4), 489-523.

Williamson, T. (2000). *Knowledge and Its Limits*. New York: Oxford University Press.

Williamson, T. (2004). Philosophical 'Intuitions' and Scepticism about Judgement. *Dialectica* 58(1), 109-153.

Conflict of Interest Statement

There are no conflicts of interest with respect to this paper.