




ARTICLE

A Practice-based Account of The Truth Norm of Belief

Xintong Wei 

Philosophy Department, University of Stirling, Stirling, UK
Email: l.weixintong@gmail.com

(Received 2 December 2021; revised 30 May 2022; accepted 30 June 2022)

Abstract

It is a platitude that belief is subject to a standard of correctness: a belief is correct if and only if it is true. But not all standards of correctness are authoritative or binding. Some standards of correctness may be arbitrary, unjustified or outrightly wrong. Given this, one challenge to proponents of the truth norm of belief, is to answer what Korsgaard (1996) calls ‘the normative question’. Is the truth norm of belief authoritative or binding regarding what one ought to or may believe? If so, what grounds its authority? The aim of this paper is to offer a novel answer to the grounding challenge on a reason-based normative framework. I develop and defend a practice-based account of the truth norm, according to which, the authority of the truth norm of belief is grounded in what I call the T-practice, a justified social practice that functions to facilitate knowledge production, maintenance, and social cooperation.

Keywords: The Truth Norm of Belief; Belief; Truth; The Nature of Belief; Epistemic Normativity; Social Practices; Practice-based Norms; Evaluative Norms; Constitutive Norms

1. Introduction

It is a platitude that belief is subject to a standard of correctness: a belief is correct if and only if it is true, and many hold the view that truth is a norm of belief. A norm, intuitively, is something authoritative. But not all standards of correctness are authoritative. Judith Thomson (2008) warns us against the ‘tendency to take it that the appropriateness of the word “correct” or “incorrect” in a context is, by itself, a conclusive sign that there is normativity at work in that context’ (2008: 108). Some standards of correctness may be arbitrary or outrightly wrong. Consider a patriarchal society in which the way women speak is subject to some standard of correctness. But that standard has no normative authority over the way women may or ought to speak. Correctness and normativity can come apart.

Given this, one challenge to proponents of the truth norm of belief is to answer what Korsgaard (1996) calls ‘the normative question’. The issue is *not* whether truth is the standard according to which some beliefs count as correct, but rather, whether that standard of correctness is normative such that it can make a claim on us about what we may or ought to believe. Is the truth norm of belief authoritative regarding what

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

we ought to or may believe? If so, what grounds its authority? Call this the *grounding challenge*.

On a reason-based normative framework, normativity is a matter of normative reasons. If a standard of correctness is genuinely normative, then the fact that a response is correct must be a normative reason that counts in favour of that response, and the fact that a response is incorrect must be a normative reason that counts against that response. The normative question with respect to the truth norm is best understood in terms of whether there is any normative reason to believe truly (correctly) and not to believe falsely (incorrectly).

Two views are popular in the current literature. According to the *evaluative* account of the truth norm, the truth norm of belief is authoritative because there are value-based or desire-based reasons to believe truly and not to believe falsely (e.g., Alston 1985; Goldman 1999; Sosa 2003; Lynch 2004, 2009; Fassio 2011; McHugh 2012). According to the *constitutive* account of the truth norm, the truth norm of belief is authoritative because it is constitutive of the attitude of belief that truth is a reason for belief (e.g., Korsgaard 1996; Velleman 2000; Shah 2003; Gibbard 2005; Shah and Velleman 2005; Wedgwood 2007; Boghossian 2008).

I argue elsewhere that the existing attempts to ground the truth norm by appealing to value-based, desire-based, and constitutive reasons fail (Wei 2021). The aim of this paper is to develop and defend a novel response to the grounding challenge. I will argue that the authority of the truth norm is grounded in what I will call the T-practice, a justified social practice that functions to facilitate knowledge production, maintenance, and social cooperation.

The plan is as follows. I will first set up the debate on a reason-based normative framework and put forward a practice-based response to the grounding challenge (section 2). To develop and defend the practice-based account of the truth norm, I will first say more about social practices and what makes them justified (section 3). I will then argue that the T-practice is a justified social practice. The reason to believe truly and not to believe falsely is practice-based (section 4). Finally, I will respond to two potential worries for my view (section 5).

2. Authoritative norms and normative reasons

Normative reasons are facts that count in favour of various responses. By responses I refer to things that are responsive to reasons, such as actions, intentions and attitudes. I will assume that for a fact R, a response φ , R constitutes a normative reason in favour of φ -ing if one of the following reason grounding facts obtains:¹

(Desire-based)	If R were to obtain, φ -ing would promote desired states of affairs.
(Value-based)	If R were to obtain, φ -ing would promote valuable states of affairs.
(Constitutive)	If R were to obtain, it is constitutive of a certain act/attitude that R would count in favour of φ -ing.
(Practice-based)	If R were to obtain, φ -ing would constitute a justified social practice.

¹For our present purposes, I will leave aside the ongoing debate about what it takes for a fact to constitute a normative reason. Very briefly, reason fundamentalists hold that at least some facts about reasons are normatively fundamental, in the sense that they are not fully grounded in other normative facts (e.g., Scanlon 1998, 2014; Raz 2010; Parfit 2011). Others think facts about reasons are ultimately grounded in facts about desires (e.g., Smith 1994; Schroeder 2007); in facts about values (e.g., Finlay 2006; Maguire 2016); or in facts about rationality (e.g., Korsgaard 1996).

Here are some examples. Life with dignity constitutes desired and valuable states of affairs. The fact that refugees need food and shelter constitutes a desire-based and value-based reason to help them. It is constitutive of the act of promising that making a promise counts in favour of keeping that promise. The fact that one made a promise is a constitutive reason to keep that promise. Punishing criminals in accordance with the law constitutes a justified social practice. The fact that someone commits a crime is a practice-based reason to punish that person in accordance with the law.

On the reason-based normative framework, whether the truth norm is authoritative is a matter of whether for any p , the truth of p is a normative reason to believe p and not to believe not- p , which in turn depends on whether some of the reason grounding facts obtains. This paper will develop and defend the idea that the reason to believe truly and not to believe falsely is practice-based – for any p , if p were true, believing p and not believing not- p would constitute a justified social practice, the T-practice.

The practice-based account is in contrast with the evaluative and the constitutive accounts of the truth norm. The evaluativist thinks that the reason to conform to the truth norm is grounded in value-based (or desire-based) facts.² For any p , if p were true, believing p and not believing not- p would promote valuable (or desired) states of affairs. To begin with, true beliefs are instrumentally valuable. From a practical standpoint, true beliefs are useful for getting what we want, after all, true beliefs map the world in such a way that they enable us to navigate through it successfully. Second, true beliefs are arguably valuable for their own sake. Part of the epistemic ideal (no matter what else it might involve), is believing all truths and avoiding all falsehoods (see, e.g., Kvanvig 2008: 209–10; Lynch 2009: 227). Third, we seem to have a basic desire for true belief. As Lynch (2004: 504) puts it, ‘I want my beliefs to track reality, to “accord with how the world actually is” – which is to say I want them to be true.’

However, note first that true beliefs do not always pay off – although true belief is, *in general* useful for getting what we want, it is unlikely that *every* true belief would be instrumentally valuable. There is a myriad of true propositions that if one were to believe them, would not produce any desired/valuable outcomes. Further, we (ordinary epistemic agents) do not seem to be disposed to react to uninteresting or trivial true beliefs as if they are valuable. If a true belief is intrinsically valuable, then at least some agent will value that belief for its own sake, which, plausibly, involves some dispositions to act as if that belief is valuable, for example, by expressing positive attitudes towards that true belief, among other things. But we do not seem to have such dispositions towards uninteresting or trivial true beliefs. Moreover, it is also doubtful that we have a basic desire for true beliefs. While we often prefer true belief to false belief, there is little evidence that we have a basic desire for true belief *simpliciter*. In fact, some true beliefs are not desirable – they might be harmful to friendship (e.g., Keller 2004; Stroud 2006), detrimental to self-esteem (Kelly 2003), or reduce the chances of surviving challenging disease (Reisner 2009; Rinard 2015).³ So, we have good reasons to doubt that for any p , if p were true, believing p and not believing not- p would promote valuable (or desired) states of affairs. The evaluativist has not met the grounding challenge.

The constitutivist, on the other hand, thinks that the reason to conform to the truth norm is grounded in constitutive facts. It is constitutive of belief that, for any p , if p were true, p would count in favour of believing p and refraining from believing not- p . Here is

²For the ease of presentation, I will not distinguish between evaluativists who think the truth norm is value-based and those who think the truth norm is desire-based.

³For a recent overview of such cases, see, Maguire and Woods (2020: 212–13).

one way to motivate the idea. At the level of attitude, different mental attitudes can share the same content. If you can believe p , then you can also imagine p , fear p , desire p , etc. What individuates belief from imagination, fear, desire, etc., is the relation the subject stands to the mental representation of p . To use the well-known metaphor, we may imagine that there is an ‘Imagination Box’, ‘Fear Box’, ‘Hope Box’ alongside the ‘Belief Box’, where we can place tokens of the same type of mental representation. Take an attitude you have towards p . What it takes for that attitude to be the ‘Belief Box’, is for the following normative condition to obtain: if p were true, p would count in favour of having that attitude towards p and refraining from having that attitude towards not- p . The normative condition individuates belief from other types of mental attitudes. Given the plausible assumption that a condition C is constitutive of some type of mental attitude M if C individuates M from non- M , the normative condition is constitutive of the attitude type of belief.

However, the normative condition may not be *sufficient* to distinguish belief from other attitudes. For example, Owens (2003) argues that although the normative condition can distinguish belief from attitudes such as desire and hope, it fails to distinguish belief from guessing. For any p , if p were true, p would also count in favour of guessing p and refraining from guessing not- p .⁴ Furthermore, the normative condition does not seem to be *necessary* to distinguish belief from other mental attitudes. The most common way to individuate the ‘Belief Box’ is *not* by appealing to a normative condition, but by engaging in functional analysis, i.e., by identifying the relevant causal structure of the type of attitude in a system. For example, a typical analysis of the ‘Belief Box’ may include the following causal structure, which has perceiving and inferring p as inputs (which tends to cause the belief that p), and has actions and behavioural dispositions as outputs (which are apt to be caused by possessing the belief that p in combination with some desires). The normative condition appears unnecessary – it is the causal/functional roles belief plays in a system that individuate belief from other propositional attitudes. So, we have good reasons to doubt that for any p , it is constitutive of the attitude of belief that the truth of p would count in favour of believing p and not believing not- p . The constitutivist has not met the grounding challenge.

I discussed these arguments in much more detail in Wei (2021). For our present purposes, it suffices to note that both the evaluator and the constitutivist have more work to do to meet the grounding challenge. In the remainder of this paper, I will argue that the practice-based account can meet the grounding challenge by showing that the T-practice is a *justified social practice* that can ground the reason to believe truly and not to believe falsely. Let us now turn to the notion of social practice.

3. Three earmarks of social practices

A practice is a regular pattern of behaviour and attitude exhibited by the participants of that practice. Playing tennis, drinking coffee after each meal, keeping a diary, celebrating the Spring Festival are some everyday examples of practices. But not all practices count as *social practices*. In this section, I will identify three earmarks of social practices. Our discussion will set the groundwork for my defence of the practice-based account of the truth norm in the next section.

⁴This thought motivates Owens (2003) to make a case for a norm of knowledge as the correct constitutive norm of belief, because he thinks a norm of knowledge, but not a norm of truth can distinguish belief from mere guessing.

3.1. *The first earmark: social practices are rule-governed*

Consider the differences between the examples we intuitively take to be individual practices and those we take to be social practices. The first thing to notice is that to engage in a social practice is *not* simply for participants of that practice to exhibit shared, regular patterns of behaviour. The fact that you, me, and some other likeminded coffee enthusiasts regularly drink coffee after each meal does not make it a social practice. Drinking coffee after each meal, without further qualification, is something an individual does for reasons of her own and lacks shared social meaning.⁵ So, what makes some shared patterns of behaviour a social practice?

One popular answer among philosophers is to appeal to a rule-governed conception of social practices (see, e.g., Rawls 1955; Brandom 1994; Haslanger 2018).⁶ Playing tennis, and celebrating the Spring Festival, are things participants engage in together while adhering to some mutually acknowledged rules governing those practices.

Importantly, rules governing a practice need not always be explicit or strictly defined. Some rules may be implicit in ways participants of that practice engage one another. For instance, practices of festivities tend to be governed by rules that are more loosely defined and more open to interpretation. Implicit rules can evolve over time, gradually and spontaneously to reflect various changes in the patterns of behaviour exhibited by participants of that practice. Some rules are explicit and strictly defined. For instance, there are 31 rules governing the game of tennis according to the International Federation of Tennis, explicitly defined and documented. Explicit rules do not evolve spontaneously but can be revised. Change in explicit rules usually occurs through a formal procedure where participants or their representatives agree on the change made to the rules governing the relevant social practice.

By contrast, individual practices are usually not rule-governed. Drinking coffee after each meal is something likeminded coffee enthusiasts do. But there is no rule, explicit or implicit, about when and how often one should drink one's coffee that likeminded coffee enthusiasts mutually acknowledge and adhere to.

The first earmark of social practices, I suggest, is that social practices are rule-governed. Participants of a social practice exhibit shared and regular patterns of behaviour in accordance with rules mutually acknowledged by participants of that practice, which may be implicit or explicit.

3.2. *The second earmark: social practices are interactive*

The rule-governed conception of a social practice naturally leads us to another earmark of social practices. If participants engage in a rule-governed social practice, then we would expect them to hold one another mutually responsible to act in ways conforming to the relevant rules.

Take the practice of festivities. Participants not only are expected to exhibit patterns of behaviour conforming to the implicit rules governing that practice, but also to engage, interact and respond to one another's behaviour. They are expected to approve

⁵I am not ruling out that in some cultures or communities the practice of drinking coffee after each meal may be counted as a social practice.

⁶Let me be clear here: rules need not be norms for not all social practices are justified. The practice-based account of the truth norm does not beg the question against the anti-normativist by appealing to a rule-governed conception of social practices.

and reward behaviours that are in accordance with the implicit rules and to disprove and sanction behaviours that are in violation. By contrast, coffee enthusiasts who engage in the practice of drinking coffee after each meal, without further qualification, are not expected to exhibit such patterns of interaction.

One distinctive type of interactive behaviour exhibited by participants of a social practice is the act of giving and asking for (normative) reasons when participants perceive irregularities, deviations or violations in that practice.⁷ Participants of a social practice attribute obligations and permissions to themselves and other participants of the practice in accordance with their understanding of the relevant rules governing that practice. But participants do not always meet one another's expectations about what they are required or permitted to do in a given situation. It may be due to a disagreement about what is the correct application of the relevant rules in a given situation. Or perhaps they agree about how to apply the relevant rules in a given situation, but some nevertheless fail to do the right things. In such cases, participants of a social practice are expected to engage in the act of giving and asking for reasons to negotiate their differences and bring one another into mutual conformance. The participant who fails to meet other participants' expectations would either be committed to alter their behaviour to meet their demand or defend their behaviour with reasons.

The act of giving and asking for reasons is often accompanied by what Strawson (1962) famously called 'reactive attitudes'. Reactive attitudes, as I understand them here, are attitudes of approval and disapproval participants hold towards themselves and other participants of that practice. Participants tend to exhibit attitudes of approval towards someone (including themselves) when they perceive that agent as conforming to the relevant rules governing the practice, which may include praise, satisfaction, gratitude, and other affirmative attitudes. Participants tend to exhibit attitudes of disapproval towards someone (including themselves) when they perceive that person as violating the relevant rules governing the practice, which may include blame, guilt, remorse, resentment, disappointment, anger, and other critical attitudes.

The second earmark of social practices, I suggest, is that social practices are interactive. Participants of a social practice exhibit a distinctive cluster of behaviour and attitude designed to hold one another mutually responsible for conforming to the rules governing that practice.

3.3. *The third earmark: social practices have social functions*

What underpins the interactive conception is that a social practice typically has some social function which affects its participants. The examples we intuitively take to be individual practices, e.g., the practice of drinking coffee after each meal and keeping a diary, without further qualification, do not seem to have any social function, although they may serve some individual needs. By contrast, examples we intuitively take to be social practice such as playing tennis and celebrating the Spring Festival, function to serve some collective social ends. Crudely speaking, for instance, we might say that

⁷My characterization of the second earmark of social practice draws on what Brandom (1994) famously labels as the *game of deontic scorekeeping*. According to Brandom, to engage in a social practice is not simply for participants of that practice to exhibit shared, regular patterns of behaviour. Crucially, they must engage in a game of deontic scorekeeping. The implicit scorekeeping can be made explicit by giving and asking for reasons.

the function of the practice of playing tennis is to organize the way we spend our leisure, to produce and distribute pleasure and good health.⁸

But of course, it is not always possible to read off the function of a social practice directly from a foray into the effects of a paradigmatic instance of that practice. In section 5, I will introduce a genealogical method to uncover the function of a social practice. But for now, let me just say that a third earmark of social practices is that a social practice typically produces, distributes, or organizes some social goods which have a practical impact on the participants of that practice.

Importantly, not all social practices can be justified by their social functions. This raises the difficult question of what it is for a social practice to count as a justified practice. I do not have a full answer here. But reflecting upon some examples might help. History is full of examples of unjustified social practices. Some social practices discriminate and oppress particular groups of people, e.g., the practice of slavery, patriarchal gender roles, voter suppression, etc. Some social practices exploit workers and natural resources, e.g., the practice of child labour, unregulated farming, deforestation, etc. Such practices perform its social function at the cost of causing substantial harm to participants involved in those practices.⁹ For our purposes, we can take on board the following plausible characterization of a justified social practice. A justified social practice is a practice that functions to produce, distribute, or organize social goods that overall benefit the participants of that practice without causing substantial harm to some participants involved in that practice.

The third earmark of social practices, I suggest, is that social practices have social functions which have a practical impact on the participants of those practices. But not all social practices are justified by their social functions.

We can summarize the points I made as follows: social practices are rule-governed patterns of interactive behaviour and attitude exhibited by a group of agents for producing, distributing, or organizing some social goods. With this characterization of social practices in place, in the next section I will flesh out the T-practice in more detail and show that the T-practice is a justified social practice.

4. Why the T-practice is a justified social practice

4.1. Characterizing the T-practice

The T-practice, at its core, refers to the shared pattern of belief management where one believes *p* and refrains from believing not-*p* when *p* is true. An immediate objection to characterize the T-practice as a social practice is the following observation. Isn't it true that we have many false beliefs and lack beliefs about many truths? If, in many cases, when *p* is true, one does not believe *p* or refrain from believing not-*p*, then in what sense is the T-practice a shared *regular pattern* of behaviour exhibited by ordinary epistemic agents?

There are two things to say in response. First, we should distinguish the question of whether the T-practice can be reasonably characterized as a social practice from the question of whether the T-practice is an actual social practice most communities engage in. The question concerning us here is the former not the latter. To justify the claim that

⁸This, of course, is speculative and my main point here is that social practices typically function to satisfy some collective ends of participants who engage in those practices.

⁹Causing substantial harm to the environment and animals may also make a social practice unjustified. But for our purposes, I will just focus on the practical impact a social practice has on people.

the T-practice is a social practice, it suffices to show that, the T-practice can be reasonably characterized as a practice involving rule-governed patterns of interactive behaviour and attitude exhibited by epistemic agents for some collective ends.

This, of course, is not to say that the latter question is not important. I do think the T-practice is a common epistemic practice most epistemic agents engage in.¹⁰ It is important to bear in mind that the exhibition of irregularities, deviation or violation of rules are to be expected in a social practice. Irregularities, etc., are expected to happen and trigger interactive patterns of behaviour (e.g., giving and asking for reasons) and attitude (e.g., reactive attitudes such as blame), which I identified as an earmark of social practices. As such, irregularities, etc., are in fact evidence of an underlying social practice. The T-practice is not an exception. The fact that we have many false beliefs and lack beliefs about many truths does not imply that the T-practice is not an actual practice ordinary epistemic agents participate in.

Having said that, if the irregularities become widespread and persistent, and if a large number of members of a community continuously fail to exhibit rule-governed patterns of interactive behaviour and attitude, then we should perhaps refrain from attributing the T-practice to the community in question. Indeed, I think a partial breakdown of the T-practice can occur in some communities and cause substantial harm, a point I will discuss in more detail shortly.

The T-practice, i.e., the shared pattern of belief management where one believes *p* and refrains from believing not-*p* when *p* is true, satisfies all three earmarks of social practices.

First, the T-practice is rule-governed. The practice of believing *p* and refraining from believing not-*p* when *p* is true is a shared pattern of mental behaviour in accordance with the truth norm of belief, which is mutually acknowledged by ordinary epistemic agents like you and me. As producers and transmitters of belief, we understand that we rely on each other to adhere to the standard to manage our beliefs correctly so that they can be shared with other epistemic agents in my epistemic community. The fact that the truth norm of belief is widely accepted as a self-evident platitude is a good indicator that it is a rule governing our everyday practice of belief management. As I have said, a rule needs not be explicitly endorsed by a community of participants of the relevant practice for that rule to be operating in that community.

Second, the T-practice is interactive. It is common for epistemic agents to exhibit interactive behaviour and attitude toward one another to hold one another mutually responsible for conforming to the truth norm of belief. We attribute epistemic obligations and permissions to ourselves and other epistemic agents in accordance with our understanding of the correct application of the truth norm of belief in a given situation. But, of course, we do not always meet one another's expectations. We often fail to believe the correct things. Sometimes we do not meet one another's expectations because we may disagree about whether a proposition is true, or whether we have an obligation or merely a permission to believe that proposition. In such cases, we engage in the act of giving and asking for reasons to negotiate our differences with the aim of bringing one another into conformity with the truth norm of belief. The epistemic agent who fails to meet others' expectations would either be committed to alter her belief to meet their demands or defend her belief with reasons. It is also common for epistemic agents to respond to one another with various reactive attitudes depending

¹⁰In fact, as I shall argue, the survival and flourishing of human societies hinge on the T-practice. The T-practice is not something we can be even partially dispensed with without paying a hefty price.

on their judgments about how well the others' belief management conforms with the truth norm of belief.

Third, the T-practice has a social function. In our previous discussion of the social dimension of belief, we have encountered one idea which provides initial support the claim. We rely on a social network to acquire and maintain knowledge, which calls for the T-practice. But more work is required to uncover the exact function of the T-practice. This is the task I shall turn to now.

4.2. *Uncovering the function of the T-practice*

I will begin by introducing the method of genealogy: a method to uncover the function of a social practice. To uncover the function of a social practice is to ask what it does for the participants of that practice. What needs does it answer to? What problems does it promise to solve? And how exactly does the practice function to meet the needs of the participants of that practice? Many philosophers have become attracted to the idea that we can uncover the function of a social practice by studying its structural origin.¹¹ The thought is that the general facts about ourselves and our environment which lead to the emergence of a social practice will illuminate the needs the practice is supposed to address and explain how that practice might help to satisfy those needs –since the best explanation for the emergence of a social practice is that the practice in question is disposed to respond to those needs of the participants of that practice. Thus, we can uncover the (current) function of a social practice by investigating its aetiology. This philosophical method of analysing a social phenomenon by investigating its aetiology is known as genealogy.

What does a genealogy involve? According to Williams (2002: 22), a genealogy of a social practice is a narrative describing a way in which 'it came about, or could have come about, or might be imagined to have come about.' An investigation into the historical origins of a social practice can reveal the effect for which it was or could have been selected. Following Williams, we can distinguish between historical genealogies and imaginary genealogies. Let me illustrate the distinction with some examples.

Historical genealogy is a narrative of how a social practice came about based on historical facts. It describes the historical conditions under which a social practice developed over time. An accurate historical genealogy depicts the historical process and the evolution of the social practice as it happened, supported by historical facts.

However, using historical genealogy as a method to uncover the function of a social practice faces a few difficulties. First, in view of the complexity of a social practice and its causal connection to a diversity of events and forces, it is difficult to pin down the exact causal chain leading to the emergence of the social practice in question. Second, what justifies a particular selection of historical events as opposed to others? How do we determine whether the historical narrative in question is accurate? A tentative historical genealogy can in general enrich our understanding of a social phenomenon. However, a historically inaccurate genealogy would significantly undermine the justificatory force of the genealogical method in illuminating the function of the social practice in question.¹²

¹¹See Queloz (2020b) for a recent defence of genealogy as a method to reveal the social function of a social practice.

¹²My point here is not to reject the value of historical genealogies or deny the possibility of a responsible and accurate historical narrative of how a social practice emerged in history. Rather, I merely want to stress its difficulty and why the fictional narrative is to be preferred.

Imaginary genealogy, by contrast, avoids those difficulties. Imaginary genealogy is a fictional narrative which describes ways in which a practice could have come about, abstracting away from historical events and forces. Imaginary genealogies often involve state-of-nature stories, which describe a simplified, imaginary situation in which our ancestors lived. It reconstructs how a social practice could have developed under those initial conditions. One influential example of imaginary genealogy in epistemology is Craig's (1990) genealogical analysis of the concept of knowledge. In Craig's state of nature, a small community of language-using human beings with unequal skills and talents were dependent on one another for their survival. Resources were scarce and threats from predators were constant. They had certain basic needs, including a need for cooperation which relies on obtaining true information effectively. Craig offers an account of how these basic needs could have led individuals to develop a prototypical concept of knowledge, *protoknowledge*, as a tool to flag good informants. He then describes how this prototypical concept is gradually objectivized – it loses its connection to the satisfaction of subjective needs and is developed into a concept with the properties associated with our concept of knowledge.¹³

Imaginary genealogy is abstract and idealised. Its epistemic force does not rely on the accuracy of its narrative, but on the plausibility of the counterfactual reasoning involved. Its main advantage is that, by constructing narratives in idealised conditions, it makes the alleged functional relation between the practice in question and the relevant needs stand out. Williams (2002: 32–4) points out that imaginary genealogies are particularly apt at revealing the point or function of a phenomenon that is not obviously functional. Take Hume's (2000: 3.2.2) genealogy of the artificial virtue of 'justice' (understood as respect for property). Intuitively, we do not regard the virtue of justice or our disposition to respect for property as functional. But if Hume's analysis of the origin of justice is correct, it reveals that the amelioration of social conflict is the primary function of the virtue of justice. Craig's genealogy also illustrates this point. If his analysis is correct, then it reveals that the concept of knowledge has a certain social function, when we do not expect it to be so. Furthermore, Williams argues that, despite being fictional, the explanations imaginary genealogies offer are unmysterious. Imaginary genealogies explain what is functional in terms of human motivations and needs that are widely taken for granted. Hume's genealogy offers a game-theoretical reconstruction of the emergence of the virtue of justice based on a conception of humans as creatures that are self-interested and have limited sympathy. Craig's analysis can also be read as a project of explaining the concept of knowledge in terms of the basic need for cooperation which requires obtaining true information effectively.

For these reasons, the genealogical analysis I offer in this section will rely mostly on fictional narratives and abstract away from specific historical conditions.¹⁴ Using the method of genealogy, I will now turn to the task of uncovering the function of the T-practice.

To see how the T-practice could have emerged, consider first the epistemic needs of early humans living in the state of nature (cf. Craig 1990: Ch. 1; Williams 2002: Ch. 3).¹⁵ Early humans lived in small communities with face-to-face interactions; they faced scarce resources and constant threats from predators. They relied on one another for

¹³For more detailed reconstruction and further discussion of Craig's genealogy, see, e.g., Queloz (2019).

¹⁴For further discussion on the methodology of genealogy, see Queloz (2020a).

¹⁵My description of the epistemic needs of early humans living in the state of nature is broadly in line with Craig (1990) and Williams (2002).

survival. To meet the needs for food and shelter, they had to cooperate with each other. Cooperative social endeavours require information such as where the food is located, whether a predator is approaching, etc. In the state of nature, the need for survival gives rise to the need for social cooperation, which in turn gives rise to the epistemic need for acquiring information about food, predators, shelter, etc.

Each member of the primitive society, regardless of their skills and talents, had what Williams (2002: 42–3) calls *positional advantage*. Positional advantage is an epistemic advantage enjoyed by any epistemic agent simply in virtue of occupying a place at one time. Imagine an individual, Ada, living in the state of nature. Ada enjoys an epistemic advantage by occupying a spatial-temporal point which enables her to acquire information via observation available from that specific spatial-temporal point, e.g., whatever she perceives in her immediate environment.

It is worth noting, however, the positional advantage one enjoys is quite limited. When Ada's epistemic network is a single node, i.e., herself, the positional advantage she enjoys can only yield limited knowledge about the world (Figure 1). She can only acquire beliefs from what she manages to perceive in her immediate environment and from what she manages to infer on her own.

We can imagine that in the state of nature there were many epistemic singletons just like Ada, each enjoying some limited positional advantage. Positional advantage enables a primitive form of division of epistemic labour – by occupying a place at one time, each individual is in a position to acquire information by observation that may not be available to other members in their community. By sharing such information, the community as a whole can make the most out of each other's positional advantage to meet their collective epistemic needs for information about food, shelter, etc.

Imagine Ada, Ben, Chen, and Dario, a small community of four individuals living in the state of nature. Suppose that each one of them, in virtue of their positional advantage, is in a position to acquire some important information about shelter, predators, food and water source. To acquire the vital set of information [shelter, food, water, predators] required for successful cooperation and their survival, they must engage in a division of epistemic labour and form an epistemic network such as that shown in Figure 2.

Early epistemic communities such as Ada's face a problem. To have a pool of information maximizing positional advantage, each individual must do their bit. They must acquire and maintain true beliefs from their environment and from their social network so that the vital set of information [shelter, food, water, predators] can become available to all members of the community.

Very likely, the T-practice could have emerged in response to this kind of challenge in primitive communities in the state of nature. Ada, Ben, Chen, and Dario would have learned that to cooperate and survive, they must hold one another accountable with their belief-production and maintenance. It is in the interest of the whole community that everyone should do their best to acquire and retain true and only true beliefs. As the division of epistemic labour became more sophisticated and as more nodes were added to the epistemic network, the emergence of rule-governed patterns of interactive behaviour and attitude would have enabled communities to meet their collective epistemic needs and to ensure their survival. The genealogy of the T-practice in the state of nature suggests that it functions to pool information in a community to facilitate knowledge production, maintenance, and social cooperation.

What I said so far, I believe, gives a plausible account of how the T-practice could have emerged in early epistemic communities. But I can imagine someone raising the following two objections.




Figure 1. Epistemic Singleton.

First, granted that the brief genealogy offers some insight about the function of the T-practice in primitive societies. But why should the fictional narrative have any bearing on the *current* function of the T-practice in real, non-idealized conditions?

In a recent paper, Smyth (2017) makes a similar point in his criticism of the genealogical analysis of the function of morality developed in Joyce (2006) and Kitcher (2011). Smyth points out that functions are sensitive to environment. There is no such thing as ‘intrinsic function’. The white fur of a polar bear has the function of camouflage only if its environment is white. But if that condition is altered, say, due to climate change, then the white fur may not have the same function. So, an entity E can perform its function in a system S only relative to a set of conditions. Call S* the current system in which E is said to have a certain function, Smyth formulates the requirement for a warranted inference from the genealogy of E to its current function as follows:

Continuity: ‘Given the set of enabling conditions that obtained in S when E developed, the same set of conditions obtains in S* and in the recent history of S*’ (Smyth 2017: 1132).

Smyth goes on to argue that the set of conditions which enables morality to perform its function in the state of nature (e.g., people living in a small community with face-to-face interactions, facing scarce resources and constant threats from predators and so on) no longer obtains in modern societies. The continuity condition is not

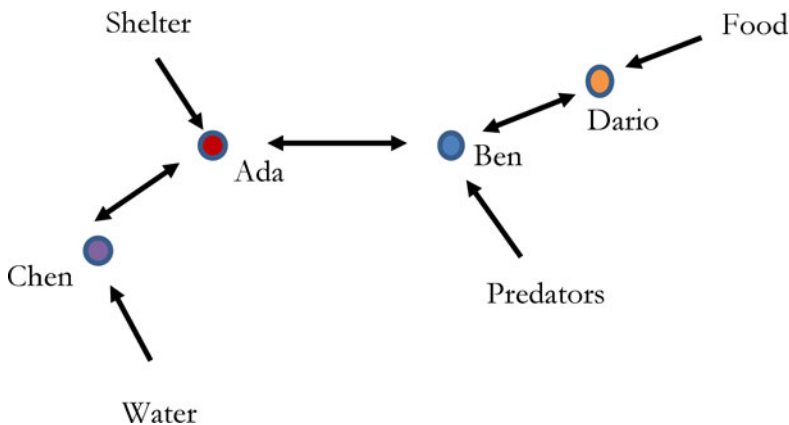


Figure 2. Epistemic Quad.

satisfied. Hence, Smyth argues, the inference from the genealogy of morality to the claim about the current function of morality is significantly weakened.

Could one raise a similar objection regarding the T-practice? To begin with, some of the conditions which enable the T-practice to perform its function in the state of nature, e.g., people living in a small community with face-to-face interactions, facing scarce resources and constant threats from predators, enjoying positional advantage and so on, no longer obtain in modern societies. In the state of nature, our hunter-gatherer ancestors rely on the maximization of their positional advantage to generate a pool of information about food, shelter, etc. But in modern societies the positional advantage one enjoys in virtue of being at a place at one time plays limited role in guiding cooperative endeavours. Social cooperation in modern societies requires highly specialized knowledge. Obtaining such knowledge requires specific training and skills. The division of epistemic labour takes a different form. What is needed is expert input and an effective way to disseminate the relevant information to the relevant parties. Hence, it may be argued that the inference from the genealogy of the T-practice to the claim about the current function of the T-practice is significantly weakened due to the change of circumstances.

Smyth makes an important point that a function is relative to a set of conditions. But I do not think the exact same set of conditions must obtain for a warranted inference from the genealogy of E to its current function. For instance, it is not the case that the same climate and geographical conditions must obtain for the white fur of a polar bear to function as a camouflage. Suppose that polar bears are relocated to the Himalayas and survive. Even if the enabling conditions that obtained in the Arctic do not obtain in the Himalayas, the inference from the genealogy of the white fur of polar bears to its function is not undermined. What is required for a warranted inference, it seems to me, is a *structural continuity*: polar bears living in both environmental systems have a need to camouflage and the conditions in both systems enable their white fur to function as a camouflage. We should revise the Continuity condition as follows:

Continuity*: Given the needs in response to which E emerged in S, the conditions in S* and in the recent history of S* must enable E to satisfy the same needs.

The move from the genealogy of the T-practice to the claim that it functions to facilitate knowledge production, maintenance and social cooperation does not violate Continuity*. Conditions in modern societies continue to enable the T-practice to carry out the same function as it first emerged in the state of nature. Let me explain.

As societies develop, of course, our specific epistemic needs change dramatically over time and the division of epistemic labour grows ever more sophisticated. In the state of nature, the kind of knowledge our hunter-gatherer ancestors needed for social cooperation was limited to mostly observational knowledge about the location of food, predators, etc. While in modern societies, the kind of knowledge we need for social cooperation is often systematic and highly specialized. But the expansion of our specific epistemic needs and the more sophisticated forms of division of epistemic labour do not violate Continuity*. There is a structural continuity between the state of nature and modern societies in that we continue to form epistemic networks to pool our epistemic resources and to engage in division of epistemic labour to produce and maintain knowledge and to facilitate social cooperation. Each individual epistemic agent in modern societies must also do their bit and hold one another accountable with their belief-

production and maintenance. It is in the interest of the whole community that everyone should do their best to acquire and retain true and only true beliefs. The T-practice could be reasonably expected to persist in modern societies to meet our needs for knowledge production, maintenance, and social cooperation.

Another concern one might have for the genealogy of the T-practice is uniqueness. Suppose we define *E-practice* as a social practice of belief management governed by a norm of evidence and define *K-practice* as a social practice of belief management governed by a norm of knowledge. It seems that one could provide a similar genealogical analysis of the E-practice and the K-practice along the lines I have put forward for the T-practice. If so, it seems plausible to say that there may be a few alternative social practices that also arise in response to meet our needs for knowledge production, maintenance, and social cooperation.

But I do not think the observation that there might be other social practices sharing the same social function as the T-practice raises any issue for my argument that the T-practice has the function of facilitating knowledge production, maintenance, and social cooperation. In fact, it would be surprising if the T-practice is the *only* social practice emerged over time to meet those needs, given how central they are to our survival and flourishing. Furthermore, I take the fact that similar analysis may be provided to account for the E-practice and the K-practice a booster to my view, for it shows the practice-based approach I developed in this chapter has the potential to not only account for the truth norm of belief, but also for the evidence and knowledge norm of belief. Indeed, a practice-based account of epistemic norms seems a promising project for further research.

In sum, the genealogy of the T-practice, i.e., how the T-practice could have emerged in early epistemic communities, reveals the current function of the T-practice, which is to facilitate knowledge production, maintenance, and social cooperation.

4.3. Is the T-practice justified?

Earlier I emphasized that not all social practices can be justified. A justified social practice is a practice that functions to produce, distribute, or organize social goods that overall benefit the participants of that practice without causing substantial harm to some participants involved in that practice. Now, is the T-practice justified? Does it facilitate knowledge production, maintenance, and social cooperation at the cost of causing substantial harm to some epistemic agents involved in that practice?

Recent discussions on epistemic injustices (e.g., Fricker 2007) and epistemic wrongdoings (e.g., Basu 2018, 2019) provide some grounds for concern here. For instance, there is some evidence that facts about group membership affect how knowledge is produced and distributed in an epistemic community. If so, one might worry that the T-practice in its current form (and given the existing structural issues in our societies), may systematically favour members who belong to the dominant groups and disfavour members who belong to the disadvantaged groups.¹⁶ Worse still, there is some evidence that the T-practice may even cause harm to some epistemic agents involved in that practice. For example, believing certain things about a marginalized community, even when they are true, can often amplify an unjustified narrative about the community and put members of that community in harm's way.

¹⁶It should be noted, however, it remains highly contentious whether it is the T-practice itself to be blamed here for epistemic injustices.

These cases, no doubt, raise critical issues. Indeed, I am quite sympathetic to the idea that the T-practice as it is currently being practiced is subject to reform and improvement. However, it is beyond the scope of my current project to address these issues directly here. Having said that, I believe there is a compelling case to be made for the T-practice. Even if the T-practice as it is currently being practiced is not flawless, there is overwhelming reason to think that it is justified, for even a partial breakdown of the T-practice would bring substantial harm to every epistemic agent involved in that practice. In the remainder of this section, I will use George Orwell's *1984* as a thought experiment and suggest that the story brilliantly illustrates what a community would look like when the T-practice is in a partial breakdown. It would be marked by an impoverished epistemic environment and cause substantial harm to epistemic agents living in that community.

The story describes an imagined dystopia in the year 1984, when the world is divided into three totalitarian superpowers and is under perpetual war. The superpower Oceania is ruled by the Party under the dictatorship of Big Brother. The protagonist, Winston Smith, is a Party official and a secrete dissident. Winston works in the Ministry of Truth, which ironically controls the production and dissemination of misinformation and disinformation. In Oceania, the Party's machinations corrupt all standard epistemic resources. The Party controls all channels of news media – newspapers and radios reported 'news' according to the party lines: from economic data such as the quarterly output of pairs of boots (which always overfills quota) to the situations on the battlefield (which always glorifies the Oceanian army). The Party machinery also constantly destroys and revises historical records. The Party frequently changes its alliance with one of the other two superpowers. Although Oceania was at war with Eastasia for years, once the Party declares that it is Eurasia who they had been fighting, all records concerning Oceania's relations with Eastasia and Eurasia are 'corrected'. And all traces of tampering are erased. There is no way to find out what happened by consulting historical records or documents. Furthermore, the entire society is under mass surveillance and ruled by terror. Thought Police are omnipresent through a surveillance device called telescreen and everyone is a potential victim of arbitrary arrest and torture. Most people are either brainwashed by propaganda or dare not deviate from the Party line overtly. It is impossible to tell what others genuinely believe and no one is reliable. One's memory becomes unreliable due to long-term immersion in such a corrupt epistemic environment. Winston's lover Julia, for example, does not remember that Oceania was at war with Eastasia and sincerely believes along the Party line that Oceania has been on friendly terms with them. Even sense perception becomes an unreliable epistemic source of belief regarding sensitive issues, because there is no other way to confirm what one takes oneself to see and hear.

When it concerns issues that are sensitive to the Party, such as politics, economy, warfare, etc., the epistemic environment in Oceania is so corrupt that none of the standard epistemic sources – testimony, memory, and sense perception has a stable connection to truth. As such it is difficult for an Oceanian to correctly manage her beliefs even if she is committed to adhere to the truth norm of belief. Consequently, it is not surprising that Oceanians would persistently fail to exhibit patterns of interactive behaviour conforming to the truth norm of belief with respect to a large set of beliefs. A partial breakdown of the T-practice is marked by the impoverished epistemic environment caused by the Party's systematic manipulation.

Turning to another insight I want to draw from Orwell's *1984*. The daily running of the Party's machinations relies on Party members such as Winston, who oversee the

tampering of records, the production of propaganda and so on. So how could officials who have knowledge of the Party's manipulative tricks and manifold lies still fall victim to the Party's manipulation? Orwell offers a brilliant explanation by appealing to what he calls 'doublethink', a capacity to believe two contradictory claims simultaneously. It is through doublethink, that a Party member can alter their memories and believe in things that they know to be false as long as the Party demands it. A new language – Newspeak – is invented to reinforce people's ability to doublethink. In addition to the Party's active indoctrination of doublethink and the tools made available to facilitate doublethink, there is also a strong incentive for a Party member to doublethink since this seems to be the only way to retain some sanity while running the machinations of the Party.

What Orwell describes as doublethink can be understood as a form of intentional self-deception: one deliberately brings about a belief that one knows or believes to be false (e.g., Davidson 1985). Orwell's story illustrates how it is possible for the Party to partially break down the T-practice even for those 'well-off' participants who do have some epistemic access to reality with respect to sensitive issues and are able to distinguish truth from falsity in a corrupt epistemic environment. It is the persistent and socially embedded self-deception which makes it difficult for them to adhere to the truth norm.¹⁷ Consequently, it is not surprising that even Party officials, the epistemic elites so to speak, also persistently fail to exhibit patterns of interactive behaviour conforming to the truth norm of belief with respect to a large set of beliefs. A partial breakdown of the T-practice is marked by pervasive self-deception induced by the Party's systematic manipulation.

Using Orwell's *1984* as an example, we see how the T-practice may be broken down in a community by a coercive, manipulative political force. When the T-practice is under attack, we pay a hefty price at both a personal and societal level. A person living in Oceania may have thoughts (e.g., thoughts along the Party lines), emotions (e.g., hatred towards the enemy and love of Big Brother) and relationships (e.g., they have families and colleagues). But none of them seems to mean anything anymore. They think, feel, and socialize only as dictated by the Party. They are neither curious nor creative. The living standards in Oceania are poor. Most people live in poverty. The productivity is low. Art and literature vanish because for those living in Oceania, there is nothing in themselves that seeks expression. People are entertained by pornography, cheap alcohol and state lotteries that never pay out. A life in Oceania is deprived of everything a human life is worth living for. At the personal level, one who loses contact with a large part of reality as a result of manipulation and self-deception is deprived of meaningful thoughts, emotions and relationships; at the societal level, we see a termination of knowledge production, a regress in technology, science, and arts. Without the T-practice, a human life is hardly worth living.

Of course, Orwell's *1984* depicts a worst-case scenario unfolding in a totalitarian society. But a partial breakdown of the T-practice using the same kind of mechanisms described in the novel (e.g., propaganda, misinformation campaigns, etc.) is common in real life. Think about epistemic agents living in a religious cult, citizens in an authoritarian regime, or an online community of conspiracy theorists. Their epistemic situation is similar to someone living in the Oceania of *1984*, in that none of them gain (hardly any) genuine knowledge from their social network. The substantial harm caused by the partial breakdown of the T-practice is self-evident: cult members, oppressed citizens and victims of conspiracy theorists could not flourish in their impoverished epistemic

¹⁷See Wei (2020) for further discussion on the social dimension of self-deception.

environment. Orwell's dark tale of what a society would look like when the T-practice is under attack for political control is hardly an exaggeration.

So, given the crucial social function of the T-practice, and given that even a partial breakdown of the T-practice would cause substantial harm to epistemic agents at both individual and societal level, we have strong reasons to think that the T-practice is a justified social practice capable of grounding the reason to believe *p* and to refrain from believing not-*p* when *p* is true. The reason to believe truly and not to believe falsely is practice-based. In the next section, I will respond to two potential worries for the practice-based account of the truth norm.

5. Objections

5.1. *The wrong kind of reason problem*

One concern for the practice-based account is that it seems to share the wrong kind of reason problem facing the evaluative account. Some beliefs are valuable independent from their truth values. Believing that the pandemic has ended promotes valuable states of affairs because it makes me feel happy. Believing that my friend Lena will have a successful writing career promotes valuable states of affairs because it strengthens our friendship. If believing *p* promotes goods such as happiness and friendship, it seems that I have value-based reasons to believe *p* regardless of its truth value. If one is committed to the view that the reason to believe truly and not to believe falsely is value-based, one seems to be forced to accept other value-based reasons for belief. But facts such as that believing *p* makes me happy or strengthens a friendship are the wrong kind of reasons for believing *p*. The right kind of reasons for belief must bear upon the truth of the proposition in question.

On the face of it, the practice-based account seems to avoid the wrong kind of reason problem. According to the practice-based account, the reason to believe truly and not to believe falsely is grounded in the fact that believing truly and not believing falsely constitutes a justified social practice, i.e., the T-practice. The T-practice-based reason to believe *p* and to refrain from believing not-*p* is distinctive of our epistemic practice of belief management. It does not force one to accept other desire-based or value-based reasons for belief and thus avoids the wrong kind of reason problem.

However, one might argue that if the reason to believe truly and not to believe falsely is practice-based, then one must also accept other practice-based reasons for belief. It is possible that there are practice-based reasons for belief that are independent from its truth values.¹⁸ For example, suppose that the epistemic practice of believing things that would promote happiness and friendship is a justified social practice. On the practice-based account, facts such as that believing *p* makes me happy or strengthens a friendship would constitute practice-based reasons for believing *p*. But they are the wrong kind of reasons for belief.

One response to the problem is to reject the idea that non-evidential reasons are the wrong kind of reasons for belief. Some philosophers endorse what is known as the pragmatist view about epistemic reasons, according to which, there are reasons for belief which do not bear upon the truth of that belief (e.g., Reisner 2009; Papineau 2013; Rinard 2015, 2017, 2018, 2019; Steglich-Petersen & Skipper 2019; Maguire & Woods 2020). For example, one putative non-evidential reason against believing *p* is the fact that believing *p* will take up cognitive resources such as storage space and processing

¹⁸I thank an anonymous referee of this journal for raising this objection.

power (e.g., Harman 1986). The fact that I have no interest in *p* seems to be another non-evidential reason against believing *p* (e.g., Papineau 2013). Furthermore, it has been argued that it is a distinctive virtue of the pragmatist view that it unifies epistemic normativity and practical normativity (see, e.g., Rinard 2017, 2018, 2019). Epistemic and practical reasons are of the same species. There is nothing distinctive about epistemic reasons. Finally, the pragmatists can also explain why in most cases evidential considerations are reasons for belief. This is because true beliefs are pragmatically beneficial. If pragmatism is correct, the wrong kind of reason problem does not bite, for there is no deep distinction between the right and wrong kind of reasons for belief.

But pragmatism about epistemic reasons is not an insignificant commitment and it is beyond the scope of this paper to settle the debate between pragmatism and evidentialism. But fortunately, unlike the evaluator, I think proponents of the practice-based account need not rely on pragmatism to respond to the wrong kind of reason problem. For the wrong kind of reason problem to arise for the practice-based account, one must offer an account of a justified social practice capable of grounding non-evidential reasons for belief. It is one thing to say that such practice is conceivable (e.g., the practice of believing things that would promote happiness and friendship), it is quite another thing to show that such practice is indeed a justified social practice. In the absence of such an account, I do not think the wrong kind of reason problem bites.

It is also worth pointing out that the practice-based account can preserve our intuition that value-based and desire-based reasons are the wrong kind of reasons for belief. That you will receive a large reward if you believe the moon is made of cheese is the wrong kind of reason to believe that the moon is made of cheese. That being said, I stay open to the idea that there may be non-evidential reasons for belief, if it can be shown that there is some justified social practice which grounds non-evidential reasons for belief. If such practice-based non-evidential reasons do exist, they are the right kind of reasons for belief for they are distinctive of our epistemic practice of belief management which grounds those reasons.

5.2. The problem of limited authority

Another concern for the practice-based account is that it seems to share the problem of limited authority facing the constitutive account. Consider the constitutive rules of chess. As far as you are playing the game of chess, there is a reason against moving your knight horizontally. But you have no reason to refrain from moving your knight horizontally if you are not playing a game of chess. Whether the rules of chess are authoritative depends on whether one is playing the game of chess. Likewise, as far as you are playing the game of belief, the truth of *p* is a reason for you to believe *p* and to refrain from believing not-*p*. But you have no reason to do so, if you are not playing the game of belief with respect to *p*. Whether the truth norm of belief is authoritative depends on whether one is playing the game of belief.¹⁹ So, the constitutive grounding fact could explain why the truth of *p* is a reason to believe *p* and to refrain from believing not-*p*, only if one were to form a *belief* about *p*. But this fails to account for the full scope of the authority of the truth norm, according to which, for any *p*, the truth of *p* is a reason to believe *p* and not to believe not *p*, regardless of whether one were to form a *belief* about *p* or not.

Perhaps it is all right to just bite the bullet. After all, unlike the game of chess, the game of belief is not something we can easily opt out of (e.g., Feldman 2000). The

¹⁹For a recent and rich discussion of the game of belief, see Maguire and Woods (2020).

constitutive grounding fact could explain why the truth of p is a reason to believe p and to refrain from believing not- p since playing the game of belief is just what we do as epistemic agents. But even if the game of belief is something I play to navigate through my daily life, it seems unclear why, on the constitutive account, the truth of p is a reason for me to believe p , if I had no intention to form a belief about p .

On the face of it, the practice-based account seems to avoid the problem of limited authority. What grounds the reason to believe p and to refrain from believing not- p when p is true is the fact that the T-practice is a justified social practice, regardless of whether one is playing the game of belief. Even for individuals who are not playing the game of the T-practice, the practice-based reason is nonetheless normatively binding. This is because as far as the T-practice is a justified social practice operating in their community, the truth of p is a practice-based reason for the individual to believe p and to refrain from believing not- p . No individual epistemic agent lives outside the authority of the T-practice. So, the practice-based account avoids the problem of limited authority. The truth of p is a reason to believe p and to refrain from believing not- p , even if one is neither playing the game of belief nor playing the game of the T-practice – the truth norm of belief remains authoritative.

However, one could argue that it is not inconceivable that the T-practice is not operative in some community. On the practice-based account, if the T-practice is not operative in a community, then there is no normative reason to believe p just in case p is true, since the reason is grounded in the T-practice.

I have argued in section 4.3 that even a partial breakdown of the T-practice would be disastrous for a community, since it involves systematic external manipulation and self-deception. But one might argue that we can conceive of some practice T^* such that (1) for most p , believing p and not believing not- p when p is true constitutes the T^* -practice; (2) for some q , it is *not* the case that believing q and not believing not- q when q is true constitutes the T^* -practice; (3) a community could survive and flourish with T^* -practice in place. For a community that opts for the T^* -practice (which is a minor deviation from the T-practice), at least for some q , there isn't a normative reason to believe q and not to believe not- q when q is true. So, the practice-based account fails to explain the full scope of the authority of the truth norm, according to which, for *any* p , the truth of p is a reason to believe p and not to believe not- p .

However, it is one thing to say that T^* -practice is conceivable, it is quite another thing to show that such practice is indeed a justified social practice. As I have argued in section 3.3, a justified social practice is a practice that functions to produce, distribute, or organize social goods that overall benefit the participants of that practice without causing substantial harm to participants involved in that practice. While it is possible that T^* -practice might bring benefits to some individuals, it is doubtful whether the T^* -practice has a social function that benefits its participants without causing some substantial harm to some individuals.

Moreover, even if the T^* -practice can be shown to be a justified social practice, it is doubtful that *in the absence of the T-practice*, the T^* -practice can fulfil the function to produce, maintain knowledge production and social cooperation, which is crucial to the survival and flourishing of a community. It is one thing to say that the T^* -practice might be operative in some community, it is another to say that the community can *replace* the T-practice with the T^* -practice and continue to flourish. So, if I am right that a community cannot survive and flourish in the absence of the T-practice, the T-practice is not something a community can opt out of. So, the practice-based account does not face the problem of limited authority and can account for the full scope of authority of the truth norm.

6. Conclusion

To conclude, in this paper I argued for a novel practice-based account of the truth norm on a reason-based normative framework. On this view, the authority of the truth norm is grounded in the T-practice, a justified social practice that functions to facilitate knowledge production, maintenance, and social cooperation. The reason to believe *p* and to refrain from believing not-*p* when *p* is true is practice-based.²⁰

References

- Alston W.P. (1985). 'Concepts of Epistemic Justification.' *The Monist* 68(2), 57–89.
- Basu R. (2018). 'Can Beliefs Wrong?' *Philosophical Topics* 46(1), 1–17.
- Basu R. (2019). 'What We Epistemically Owe to Each Other.' *Philosophical Studies* 176(4), 915–31.
- Boghossian P.A. (2008). 'Epistemic Rules.' *Journal of Philosophy* 105(9), 472–500.
- Brandom R.B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Craig E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.
- Davidson D. (1985). 'Incoherence and Irrationality.' *Dialectica* 39(4), 345–54.
- Fassio D. (2011). 'Belief, Correctness and Normativity.' *Logique et Analyse* 54(216), 471–86.
- Feldman R. (2000). 'The Ethics of Belief.' *Philosophy and Phenomenological Research* 60(3), 667–95.
- Finlay S. (2006). 'The Reasons that Matter.' *Australasian Journal of Philosophy* 84(1), 1–20.
- Fricke M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Gibbard A. (2005). 'Truth and Correct Belief.' *Philosophical Issues* 15(1), 338–50.
- Goldman A.I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Harman G. (1986). *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.
- Haslanger S. (2018). 'What is a Social Practice?' *Royal Institute of Philosophy Supplement* 82, 231–47.
- Hume D. (2000). *A Treatise of Human Nature* (D.F. Norton & M.J. Norton, eds). Oxford: Oxford University Press.
- Joyce R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Keller S. (2004). 'Friendship and Belief.' *Philosophical Papers* 33(3), 329–51.
- Kelly T. (2003). 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research* 66(3), 612–40.
- Kitcher P. (2011). *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Korsgaard C.M. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Kvanvig J.L. (2008). 'Pointless Truth.' *Midwest Studies in Philosophy* 32(1), 199–212.
- Lynch M.P. (2004). 'Minimalism and the Value of Truth.' *Philosophical Quarterly* 54(217), 497–517.
- Lynch M.P. (2009). 'The Value of Truth and the Truth of Values.' In A. Haddock, A. Millar and D. Pritchard (eds), *Epistemic Value*, pp. 225–42. Oxford: Oxford University Press.
- Maguire B. (2016). 'The Value-Based Theory of Reasons.' *Ergo: An Open Access Journal of Philosophy* 3.
- Maguire B. and Woods J. (2020). 'The Game of Belief.' *Philosophical Review* 129(2), 211–49.
- McHugh C. (2012). 'The Truth Norm of Belief.' *Pacific Philosophical Quarterly* 93(1), 8–30.
- Owens D.J. (2003). 'Does Belief Have an Aim?' *Philosophical Studies* 115(3), 283–305.
- Papineau D. (2013). 'There Are No Norms of Belief.' In T. Chan (ed.), *The Aim of Belief*, pp. 64–79. Oxford: Oxford University Press.
- Parfit D. (2011). *On What Matters, 1*. Oxford: Oxford University Press.
- Queloz M. (2019). 'Genealogy and Knowledge-First Epistemology: A Mismatch?' *Philosophical Quarterly* 69(274), 100–20.
- Queloz M. (2020a). 'From Paradigm-Based Explanation to Pragmatic Genealogy.' *Mind* 129(515), 683–714.

²⁰This work was supported by the 2021–2022 Analysis Trust Studentship, Royal Institute of Philosophy 2020–2021 Jacobsen Studentship, the Aristotelian Society 2020–2021 Student Bursary, and the John Templeton Foundation PhD Studentship under Grant 58450.

- Queloz M.** (2020b). 'Revealing Social Functions through Pragmatic Genealogies.' In R. Hufendiek, D. James and R. Van Riel (eds), *Social Functions in Philosophy: Metaphysical, Normative, and Methodological Perspectives*, pp. 200–18. London: Routledge.
- Rawls J.** (1955). 'Two Concepts of Rules.' *Philosophical Review* **64**(1), 3–32.
- Raz J.** (2010). 'Reason, Reasons and Normativity.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 5*. Oxford: Oxford University Press.
- Reisner A.** (2009). 'The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem.' *Philosophical Studies* **145**(2), 257–72.
- Rinard S.** (2015). 'Against the New Evidentialists.' *Philosophical Issues* **25**(1), 208–23.
- Rinard S.** (2017). 'No Exception for Belief.' *Philosophy and Phenomenological Research* **94**(1), 121–43.
- Rinard S.** (2018). 'Believing for Practical Reasons.' *Noûs* **4**, 763–84.
- Rinard S.** (2019). 'Equal Treatment for Belief.' *Philosophical Studies* **176**(7), 1923–50.
- Scanlon T.** (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scanlon T.M.** (2014). *Being Realistic About Reasons*. Oxford: Oxford University Press.
- Schroeder M.** (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Shah N.** (2003). 'How Truth Governs Belief.' *Philosophical Review* **112**(4), 447–82.
- Shah N. and Velleman D.J.** (2005). 'Doxastic Deliberation.' *Philosophical Review* **114**(4), 497–534.
- Smith M.** (1994). *The Moral Problem*. London: Blackwell.
- Smyth N.** (2017). 'The Function of Morality.' *Philosophical Studies* **174**(5), 1127–44.
- Sosa E.** (2003). 'The Place of Truth in Epistemology.' In L. Zagzebski and M. DePaul (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 155–80. Oxford: Oxford University Press.
- Steglich-Petersen A. and Skipper M.** (2019). 'An Instrumentalist Account of How to Weigh Epistemic and Practical Reasons for Belief.' *Mind*, fzz062.
- Strawson P.** (1962). 'Freedom and Resentment.' In *Proceedings of the British Academy* **48**, 1–25.
- Stroud S.** (2006). 'Epistemic Partiality in Friendship.' *Ethics* **116**(3), 498–524.
- Thomson J.J.** (2008). *Normativity*. Chicago, IL: Open Court.
- Velleman D.** (2000). 'On the Aim of Belief.' In D. Velleman (ed.), *The Possibility of Practical Reason*, pp. 244–81. Oxford: Oxford University Press.
- Wedgwood R.** (2007). *The Nature of Normativity*. Oxford: Oxford University Press.
- Wei X.** (2020). 'The Role of Pretense in the Process of Self-deception.' *Philosophical Explorations* **23**(1), 1–14.
- Wei X.** (2021). *In Defence of the Truth Norm of Belief*. PhD thesis, University of St Andrews, St Andrews.
- Williams B.** (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton, NJ: Princeton University Press.

Xintong Wei is the recipient of the 2021–2022 Analysis Trust Studentship, hosted by the University of Stirling. She completed her PhD at the St Andrews and Stirling Graduate Programme in Philosophy (SASP). Her research situates in the intersection of epistemology, philosophy of mind and metaethics.