Being Wrong in an Extra-Moral Sense

Establishing Degrees of Accountability for Epistemic Failures

by
Pavlo Kostin

Submitted to the Department of Philosophy, Central European University
In partial fulfillment of the requirements for the degree Master of Arts

Supervisor: Katalin Farkas

Vienna, Austria 2023

Abstract

The aim of this work is to explain the conditions that warrant different degrees of epistemic accountability for inappropriate doxastic stats with respect to agents with unequal access to epistemic resources. Contemporary theories of epistemic blame and accountability largely ignore the question of degrees of accountability, as well as the role social factors play in shaping our epistemic agency. In including these factors, this work supplies an account of epistemic normativity that takes stock of disparate distribution of epistemic resources. Such distributive inequality adversely affects individuals' epistemic agency and, by extension, negatively influences their doxastic states. Bearing in mind the corrective purpose of practice of holding individuals to account, I argue that the degree of accountability is relative to the influence available resources exert on individual's epistemic agency.

Acknowledgments

First and foremost, I want to thank my father for the immense support he gave me, and for protecting a place I am calling home.

For whatever glimpses of clarity a reader will encounter in this work, I have to thank Katalin Farkas, who helped to bring order to the calamity of my ideas. I would want to thakn her for many other things as well.

I am glad that Ferenc Huoranszki, Tim Crane and Asya Passinsky intellectually inspired and challenged me throughout all these years.

I am grateful to S, J, and A, who were near whenever I needed them the most.

Lastly, I cannot fail to mention N, who is the sole person the most responsible for my philosophical journey becoming possible.

Table of contents

Introduction

Consider the following remark. A remark belongs to K1, my distant acquaintance, and a manual worker. "Historically speaking, people cannot handle themselves alright without a Big Brother watching, I mean, look at those cannibals and such". The remark sounds extremely familiar. Many of us encountered something similar when debating with older relatives at a family dinner, or when being forced to listen to an annoying classmate who discovered their passion towards political philosophy. This time, however, a person saying something similar was Yuval Noah Harari, a holder of the D.Phil. degree in history from the university of Oxford. He writes that "*in time, human violence was brought under control through the development of larger social frameworks – cities, kingdoms and states*" (Harari, 2014). Yet another passage states that "*the decline of violence is due largely to the rise of the state*" (Ibid.). The reason for this assertion is a study of death rates in tribal societies, featuring eleven tribes where violent deaths are, or were, a prominent cause of loss of life at the time the studies were conducted (Walker, Bailey, 2013). From the data collected about a fraction of contemporary tribal societies, he extends the conclusion that such societies are inherently violent to similar societies in the past. What strikes me about Harari's assertion is how poorly substantiated it is. The evidence for such an assertion is extremely scant, and its content contradicts the findings of the very researchers Harari cites – they claim that they cannot extrapolate their findings to other contemporary tribal communities, let alone to those of ancient past (p. 33). I would expect him to present more statistical data on other societies or examine testimonies of ethnographers and travelers of the past who encountered peaceful tribal communities, and so on. Instead, his claim is another iteration of "look at those cannibals" remark – stateless societies are extremely violent. The fact that Harari presents some evidence for it does not make the assertion justified, or true. And, on top of that, he makes a lot of similar omissions.

Among the range of "sins" committed by our characters, this situation exemplifies a failure of *epistemic* nature. Both agents make assertions based on insufficient evidence that does not support either of

those claims. Both agents seem to *believe* that what they assert is true. And, in both cases, the beliefs they ended up with do not live up to the scrutiny. One can argue that these claims seem to be so disquieting because of ethical or political implications. However, the violations both agents commit are of primarily epistemic nature. They make assertions without sufficient *epistemic* reason for holding the beliefs they thus express. Intuitively, one can say that they ought not believe those things without there being a sufficient reason to do so, and if what they believe is not quite true.

We are often facing situations where agents commit mistakes of epistemic nature. These mistakes are often considered to be significant independently of practical matters they bear upon. When such mistakes are made, we often find ourselves compelled to sanction agents for making them – or realize that we deserved a sanction ourselves. However, people who fail in these respects come from backgrounds with extremely disparate epistemic resources. Yuval Harari (let me anonymize him as K2) supposedly spent many years learning how to conduct historical research, how to work with the sources, which inferences from the available data are justified and which are not. He, one can assume, should know which evidence counts as sufficient and which does not. K1, on the contrary, was never exposed to sophisticated methods of research, to the standards of research in humanities and, overall, does not have access to knowledge that K2 has. These epistemic agents possess different *epistemic resources*. I will provide a more detailed account of epistemic resources in chapter 3. For now, I will loosely identify epistemic resources as the means necessary for engaging in the practice of diligent acquisition of knowledge and/or true beliefs. The availability of such resources, for the large part, is a matter of social circumstances agents find themselves in.

In cases like the one above, the disparity in epistemic resources available to a knower is a result of underprivileged social standing of one of them. According to a commonly shared intuition, the degrees of accountability of either of those agents would differ. As a result, I would hold a less privileged

knower accountable to a lesser degree. For example, a mild rebuke pointing out the deficiencies of the evidence he uses will suffice for K1. With an epistemic agent like K2, one would be justified in a harsher reaction for a mistake he makes – he is in possession of all the resources necessary to avoid it. One can express indignation when pointing out the deficiencies of K2's epistemic stance, to his face or, say, in the critical piece. Absence of justification in his case is not a result of the lack of access to proper evidence or of inability to carry out an inquiry, but of a failure to attend to the resources already within his reach.

I think that the above reasoning has some intuitive appeal. Yet, intuitions go only so far. *My aim in this work is to explain the conditions that warrant different degrees of epistemic accountability with respect to agents with unequal access to epistemic resource*s. I will primarily focus on agents and how disparity in availability of epistemic resources affects accountability for mistaken beliefs. The existing theories of epistemic blame and accountability do not comprehensibly address this issue. Some of these theories originate in the debates concerning theoretical questions, like the externalist charge against the New Evil Demon problem (Williamson, forthcoming; Littlejohn, 2016; Madison, 2017) and are lacking in detail when it comes to supplying an account of blame or accountability for real-life scenarios. Others largely focus on whether there are purely epistemic norms and, hence, grounds for sanctions for purely epistemic shortcomings (Kauppinen, 2018; Schmidt, 2021). Another significant chunk of the literature attempts to define what epistemic blame is, and what tasks theories attempting to provide such an account are facing (Boult, 2020; Rettler, 2017; Brown, 2018; Piovarchy, 2020). However, they all are lacking in a viable agent-centric explanation of what exactly warrants varying degrees of sanctions. Cameron Boult (2023) attempted to address this issue, arguing that the harshness of criticism depends on two components – degree of justification and degree of culpability. A degree of culpability is determined by how far a knower is from a possible world where their belief is justified, as well as their care towards epistemic goods (Boult, 2023, p. 14). Though interesting in its own right, I want to

concentrate on a somewhat different phenomenon – *privilege with respect to access to epistemic resources*. At the same time, I want to avoid the counterfactual analysis of culpability with respect to epistemic faults.

In what follows, I will provide an account of how epistemic resources available to an agent affect the degrees of accountability and severity of our responses to their epistemic shortcomings.

This is how I will proceed. In Chapter 1 I will introduce the notion of epistemic accountability and epistemic norms. In section 1.1, I will explain what we want from a theory of epistemic accountability and what is the purpose of the practice of holding people to account for their mistakes. In section 1.2, I will present an account of norms violations of which warrant accountability. I will argue that we can be held to account for violation of norms that prescribe us to care about fundamental intellectual goods, as well as derivative norms, such as the norms of epistemic practice. In section 1.3, I will address the issue of control we exercise over our doxastic states. My contention will be that we exercise direct control over epistemic actions which, in turn, influence our doxastic states.

In Chapter 2, I will briefly focus on conditions that either excuse or exempt agents from accountability. Sections 2.1 and 2.2 will deal with exemptions and excuses respectively.

Chapter 3 will deal with the question of how availability of epistemic resources affects the degrees of agents' accountability. In section 3.1, I will establish a connection between epistemic resources and dispositions to act in compliance with the norms of epistemic practice. Lack of relevant resources would negatively influence an agent's ability to perform epistemic actions in compliance with the norms governing such practice. In section 3.2, the connection will be made between availability of epistemic resources and degree of accountability. To evaluate the degree of agent's accountability, one needs to factor in 1) the dispositions an agent should have relative to availability of epistemic resources and 2) epistemic difficulty, I.e., the influence such dispositions would have on a belief in question.

Chapter 1. A theory of epistemic accountability and epistemic norms

As noted above, our intuitions, however reliable they seem, need explanation to see whether they are viable or not. The task of a theory of epistemic accountability is to provide such an explanation. Namely, what grounds our practice of sanctioning agents for epistemic wrongs. The answer to this question is, doubtless, a complicated one. It would require much more space than the scope of the present work could possibly allow. In this chapter, I will outline the main issues any theory of epistemic accountability has to address. After that, I will proceed with the account of the epistemic norms violations of which are of interest for this work.

1.a. Why do we need a theory of epistemic accountability?

What do we want from a theory of epistemic accountability in the first place? To address this question, one should look at what singles out epistemic accountability from other related notions, as well as at desiderata to be satisfied by any such theory. Due to the lack of space, I will not delve in depth into this question. I will limit my exposition to the most concise explanation of relevant criteria for such theories to date, as well as argue why this set of criteria is sufficient for my purposes.

But first, I would like to briefly explain the use of the term "accountability" in this work. Colloquially, accountability is not necessarily negatively valanced. Accountability for a politician, for example, might presuppose transparency, integrity and answerability before their voters. Yet, whenever we are speaking of accountability for a wrong, either moral or epistemic, what we typically have in mind is an agent's liability for a violation of a certain norm. That is, in a bad case where a certain norm - moral, prudential or epistemic - has been violated, we hold an agent to account by applying sanctions (Kauppinen, 2018, p. 3). By sanctions, I understand reactive responses in the likes of criticisms, rebukes, expressions of indignation or mistrust, etc. In this work, I will treat accountability as a ground for such sanctioning responses. Subsequently, the degree of accountability will determine the

harshness of sanctions – how harsh should one be in holding the culprit to account for a wrong committed. Holding someone to account may manifest itself in a cool disapproving judgement as well as in an expression of indignation.

But why not use the notion of blame instead if one of the concerns is the degree of harshness of sanctions for epistemic wrongs? The main problem with relying on the notion of blame has to do with the definition of blame itself. In moral philosophy, blame is associated with a wide range of often conflicting responses. Emotional theories of blame describe it as engaged emotional responses (Strawson, 1962 (2008); Wallace, 2011); conative theories associate blame with attitudinal responses, like modifications of certain attitudes in the face of a wrongdoing (Sher, 2006; Scanlon, 2008); conversational theories of blame associate the notion with functions blame performs, as derived from paradigmatic cases of face-to-face communication (Fricker, 2014). Thus, I would prefer to use a notion that is not subject to extensive debate.

Second, we have to consider what exactly a theory of epistemic accountability is supposed to explain. We hold agents accountable for their failures on different grounds (for moral or epistemic failures); we (should) hold them accountable only when they are an appropriate target of accountability; and holding them accountable for their failures fulfills a certain purpose. I will outline the desiderata for such a theory borrowing from Adam Piovarchy's (2020) account of the demands for a theory of epistemic blame. On his account, a theory of epistemic blame should answer (at least) three questions if it is to live up to its explanatory pretensions. Even though I am concerned with the broader notion of epistemic accountability, Piovarchy's account can be used for the purposes of this work without much modification.

According to the account thus modified, a theory of epistemic accountability has to answer three questions:

(1) What is epistemic accountability?

(2) When is someone an appropriate target of epistemic accountability?

(3) What is a justification for a practice of holding knowers accountable for their epistemic failures? (Piovarchy, 2020, p. 793)

Answers to these questions are interconnected. Answer to question (2) will not be satisfying without providing answers to (1) and (3). For example, to determine (2) when someone deserves to be held accountable for an epistemic failure, we should know (1) why accountability is epistemic and (3) what makes the practice of holding someone accountable justified. Similar considerations apply to questions (1) and (3).

Question (1) presents a requirement to distinguish epistemic accountability from other forms of accountability. What warrants purely epistemic accountability, rather than any other form of accountability? This question requires us to distinguish between epistemic norms and other normative requirements, such as the norms of morality, prudence or political rationality. One should explain how the violation of epistemic norms gives rise to a distinctively epistemic kind of accountability and sanctions.

Question (2) presents a requirement to explain the cases where an agent can justifiably be held accountable. The answer to the question is threefold. First, one should limit the pool of beliefs to those where epistemic failures are in some sense significant. We would not sanction an agent for making a trivial mistake there are no reasons to care too much about. Second, we should limit the pool of agents who in principle can be held accountable. Some agents are completely exempt from any form of accountability – small children, people with mental disabilities or people in states that significantly hinder their rational capacities. Third, one should supply an account of conditions that

determine when rational agents are merely excused from sanctions. A derivative question is what determines the degrees of accountability for different agents. When it comes to holding knowers accountable, it is not enough to determine whether they are excusable or not. Our responses come in degrees of severity, relative to the circumstances knowers find themselves in. Availability of epistemic resources is a key for determining circumstances under which the agent is either excused, or deserves a lesser or harsher sanction.

Question (3) presents a requirement for a justification of a practice of epistemic accountability. In other words, it does not make sense to hold a person to account if the practice itself does not have a justification for its existence. Actual or potential epistemic sanctions may be appropriate or misplaced. One of the criteria for deciding whether a sanction is appropriate or not is whether such a sanction fulfills a certain purpose. In this work, I will argue that the purpose of the practice of holding knowers to account is to cultivate the right sort of epistemic agency required for believing correctly (Piovarchy, 2020, p. 800; similar argument for the purpose of moral blame was developed by Vargas, 2013; see also McGeer, 2013). In other words, the purpose of accountability is primarily corrective - to improve agent's epistemic conduct in a way that would positively influence their epistemic standing. If an instance of holding a knower accountable for the epistemic failure does not conform to this requirement, an actual or potential sanction to be applied to a knower is misplaced. We should choose a sanction that best fits the purpose of such a practice. Harsher sanctions might be conducive towards cultivation of the right sort of agency in some cases, but not in others. Thus put, the agency-cultivating purpose of a practice of sanctioning knowers must be sensitive towards circumstances of a knower – whether certain epistemic resources are available or not. Without the availability of certain epistemic resources for some knowers, harsher sanctions do not satisfy the agency-cultivating purpose of holding knowers accountable.

Importantly, holding someone to account does not always require a second-personal response. The corrective purpose of holding someone accountable still acts as criterion for responses to epistemic shortcomings even when there is no prospect to sanction a culprit in person. No doubt, second-personal responses to someone's failure are paradigmatic cases of a sanctioning response. When we directly hold a wrongdoer accountable, we would rather do so in order to fulfil the corrective purpose of sanctioning them. From the paradigmatic example of face-to-face responses, other cases of holding someone to account are derived. These examples inherit residual traits of a second-personal response. What is at stake, even when we consider accountability of a "wrong-believer" who is very distant from us, as is K2, is the residue of this corrective purpose – this person would rather not believe what they do (for a similar account for moral blame see Fricker, 2014). That is, a corrective purpose of holding someone to account is inherited in other forms of reactive responses as well. For example, there is no point in trading insults at K2 when discussing his epistemic shortcomings with a friend. Both in second- and third-personal responses to epistemic shortcomings, what is at stake is a corrective purpose of the *practice* of holding individuals to account as a whole. Inconsistency between the second- and third-personal responses not only has an air of hypocrisy to it, but also undermines the consistency of the entire practice.

Answer to the main question of this work will be the focus of Chapter 3 – how the availability of epistemic resources affects the degrees of accountability for epistemic failures. The answer will hinge on the assumption that the purpose of holding knowers accountable is corrective – sanctions are to be conducive towards cultivation of the right sort of epistemic agency. In the next three sections, I will attempt at answering the first question – what is epistemic accountability? In section 1.b, I will provide an account of epistemic norms themselves. In section 1.c, I will address the question of control that we exercise over our doxastic states that enables the practice of holding agents to account for their epistemic faults.

1.b Which norms?

The notion of epistemic normativity is, undoubtedly, a murky one. There exists an extensive debate whether epistemic norms exist at all. Due to the lack of space, I will have to bypass this undoubtedly extremely important discussion. At this point, it will suffice to say that there exist constraints on how we should or should not go on with our epistemic lives. After all it is a truism, though a powerful one, to say that you *should not* believe falsehoods and that you *should* believe truly.

What I will concentrate on is the notion of a norm of epistemic *practice*, the violation of which may warrant accountability. I will explain what I mean by the norm of epistemic practice in greater detail below. For now, it will suffice to say that not all false or unjustified beliefs warrant accountability for their violation. First, I will briefly address the problem of why treating epistemic norms that are not concerned with the practice of acquisition of beliefs as warranting accountability is problematic. Second, I will provide a positive account of norms the violation of which can warrant accountability for purely epistemic shortcomings.

In the literature, the most common conceptions of norms for belief are Truth norm and the Knowledge norm. The core of the Truth norm, bearing in mind its various iterations, can be summarized as follows: For any proposition p, you should believe p only if p is true. The Knowledge norm for belief can be summarized as follows: for any proposition p, you should believe p iff you know p. On top of that, another requirement featuring in almost every epistemological work is a requirement for sufficient justification. This norm can be laid down as follows: for every proposition p, you should believe p only if you have a sufficient justification for p. For the sake of simplicity, let's call these norms theoretical norms. A lot of ink has been spent on debating whether these are genuine norms and whichever is more explanatorily powerful. For my purposes, however, it is enough to state that a violation of either of these norms is not sufficient to hold an agent accountable for their mistake.

Let's take an example and contrast it with the predicament K1 and K2 got themselves into. Imagine your professor sees a dog on the street and exclaims "Look how huge this Husky is!" Yet, she points not at a husky, but at malamute – the former's bigger relative. In any case, she formed a belief that is neither true, nor constitutes knowledge. One can also say that her belief is lacking in justification. All in all, a theoretical norm has been violated. Yet, one can hardly hold her to account for such a mistake – this belief is not worth a criticism, let alone an indignant response. There is no indication that one should care about forming correct beliefs about everything in the world. After encountering this mistake, I do not suspend trust in her credibility as a knower, and there is no need to sanction her – I still perceive her as a reliable and conscientious knower.

On the other hand, K1's and K2's mistaken beliefs discussed earlier do not seem to be that trivial. These mistakes compel us to care about them, and subsequently warrant a response - "How come you came to believe this so carelessly, contrary to the abundance of facts pointing to the contrary?" In this case, one can trust another knower's credibility less, and be entirely justified in imposing sanctions.

In both cases, neither of the theoretical norms is satisfied. Yet, intuitively, only in the second case can the agents be justifiably held to account. How to explain the difference in these two cases?

Some tend to argue that what makes a difference is the change in the attitude of trust towards an agent who espouses false beliefs (Kauppinnen, 2018; Schmidt, 2021; see also Boult, forthcoming, p. 43). They argue that the violation of genuine epistemic norms would lead to a reactive response of blame or criticism. Suspension of trust is taken to be a basic reactive attitude that amounts to *holding agents accountable or blameworthy*. The norms violation of which warrants such reactive responses can be treated as sufficiently authoritative norms. Such norms ground distinctively epistemic "oughts". In other words, these are genuine epistemic norms.

This response, however, is not very satisfactory. It is hard to shake off the worry about circularity. I need to explain in which cases violation of theoretical norms warrants accountability. To this end, I need a normative criterion that will prescribe in which cases we should strive towards truth or knowledge, and in which cases no such obligation exists. I cannot simply state that the violation of norms that warrant accountability warrants accountability. Therefore, something else is needed to explain why we *should* care about truth or knowledge, and why we *should* do so selectively.

What makes a professor's mistake trivial is that it does not violate the principles that we hold dear in our intellectual pursuits. The mistake in her case does not disprove her claim to being a conscientious knower, or that she cares about whether she believes truly, or whether she is willing to proceed carefully and diligently in her inquiries. On the other hand, false and poorly supported beliefs of K1 and K2 provide evidence to the fact that they do not exhibit these laudable qualities. What we see from their assertions is that they treat the complex topic of human history and forms of societal organization with a handy-wavy attitude. Undoubtedly, the topic itself, as well as the careless approach to inquiry, bear more significance for our intellectual pursuits than the taxonomy of dog breeds, or the way we form simple perceptual beliefs.

One can say that the significance of a mistaken belief that prompts a sanctioning response is determined by how a mistake affects the fundamental intellectual goods. Furthermore, positing a fundamental intellectual good might explain why we *should* care about some truths while being justified in not caring too much about others. I will follow the lead of others and will call this fundamental good *intellectual flourishing* (Brogaard, 2014). On Brogaard's account, the norm for belief (and for believers) will take the following form: *you should believe p only if believing p does not hinder intellectual flourishing* (2014, p. 15). This formulation, however, seems to entail unwanted consequences. Arguably, our professor's belief does not hinder the goal of intellectual flourishing. Yet, in this iteration the

fundamental norm would posit an obligation to believe a proposition only if it does not hinder the goal of intellectual flourishing. And it is hardly conceivable that she should mistakenly believe that she is seeing a husky before her. In the light of the above considerations, I want to propose an alternative. The alternative formulation will look as follows: *you should believe in compliance with theoretical norm(s) only if thus believing serves the goal of intellectual flourishing*. In such a case, a positive obligation to believe truly, or to know a proposition, or to possess justification for a belief exists only if it serves the attainment of most important intellectual goods. In the end, our professor did not have an obligation to believe that it was a malamute before her since it was hardly one of the goods that is constitutive of the ideal of intellectual flourishing. In the end, the fundamental norm has not been violated.

Intellectual flourishing is not simply an individual good. Sure, one can flourish intellectually on their own, but the subsequent standards for what it is to thus flourish depend on what it means to flourish intellectually for both individuals and the community - "to live the best possible intellectual life and to pursue the most important intellectual goals" (p. 19). The generality of the principle requires positing derivative norms that would allow us to satisfy this fundamental norm. For example, one should not believe things based solely on hearsay, should not ignore the available evidence, or should always consult reliable sources prior to forming beliefs *because* doing otherwise would not serve the goal of intellectual flourishing. Intellectual flourishing comprises many things.

Sure, one might find this proposal somewhat vague. What I want to say is that we care about purely intellectual goods, and the contingent goods we find important here and now exemplify our adherence to some such general principle. No doubt the norms derivative from a norm of intellectual flourishing might change over time, domain-relative norms might appear and disappear (for example, scientific "code of epistemic conduct" is prone to change). However, the general notion that we cherish success in our intellectual pursuits, often regardless of practical consequences, is necessary for making sense

of why we *should* care about truth or knowledge. Similarly, it is important to have a general principle to explain why, sometimes, it is better not to care about some truths, or not to know certain things.

This argument extends to accountability as well. We do not always care about false beliefs people have – we believe falsehoods too often. By the same token, it does not make much sense to hold people accountable for something they were under no obligation to do or believe. Trivial mistakes, akin to the one made by our professor, may serve as a case in point. There is no reason to care about her expertise in the breeds of dogs precisely because in her case this knowledge does not correspond to the ideal of intellectual flourishing.

With this in mind, one can explain why K1 and K2's beliefs warrant accountability, and our professor's incorrect belief does not. First, K1 and K2 encroach on an intellectual good the possession of which is necessary for the intellectual flourishing of these individuals, as well as their communities – history. Getting facts about our past wrong does not only affect the political and moral aspects of our existence but is detrimental on purely epistemic grounds. After all, some false beliefs are detrimental not because they bear on the present, but because they adversely affect the communal intellectual goods. Once again, the norm of intellectual flourishing requires that we set priorities regarding our intellectual pursuits. Certain things are more important for the communal well-being than others. Expertise in dog breeds, beyond the professional confines of zoologists and veterinarians, rarely, if ever, has any impact on communal intellectual well-being.

The second aspect of K1's and K2s failure is manifested in the way they proceeded with their respective inquiries. Our epistemic conduct is governed by principles that in one way or another derive from a general norm of intellectual flourishing. Ofttimes, attaining knowledge or believing truly and with sufficient justification serves this purpose. Whenever true belief/knowledge and justification are necessary for satisfaction of a fundamental norm, one has to exhibit due care in satisfaction of

theoretical norms. And to this end, it is important to know *how* exactly one is to comply with a Truth/Knowledge norm, or a requirement for a sufficient justification. In other words, what one *should do* to believe correctly.

For instance, one ought to distinguish between high-quality and low-quality evidence, consider alternative sources of information that might disqualify the available evidence, assess whether an inference one makes is valid, or abstain from believing things based on hearsay etc. These requirements are absent when a false belief does not encroach on our intellectual well-being – we do not have to waste our time and energies on diligently adhering to these norms in pursuit of outcomes that are not worth it. On the contrary, adherence to the norms of epistemic practice greatly influences one's ability to attain truth or knowledge in cases where it matters.

In pursuit of "the most important intellectual goals", diligence in inquiry is necessary to satisfy the norm of intellectual flourishing. And this diligence is manifested in attendance to the rules governing our epistemic practice. Failure to adhere to these norms threatens the intellectual well-being of an individual and the community. And in the case of K1 and K2, their mistake is indicative of a flawed approach to epistemic practice. Thus, one can say that they failed to comply with the general norm of intellectual flourishing by carelessly approaching the topic of importance – they had little evidence, they failed to consider alternative sources and the inference from the evidence they had to their respective conclusions lacks in warrant. Our professor, on the other hand, was not under a clear obligation to conform to the derivative norms of epistemic practice as specified above. Taking steps to comply with a Truth/Knowledge norm, or with a requirement for sufficient justification does not serve the goal of intellectual flourishing. Sometimes one indeed should follow the rules of practice in pursuit of an intellectual good, while in most cases even if we take things wrong, this does not affect our intellectual prosperity.

1.c Norms of practice and indirect doxastic influence

So far, I have tried to show that there exist independent normative considerations to adhere to theoretical norms – one should adhere to these norms only if such adherence will serve the goal of intellectual flourishing. However, apart from theoretical norms, I posited practical norms that serve this fundamental goal. These norms prescribe what an agent should do to satisfy theoretical norms whenever it is required by the fundamental norm. When one violates the norms of epistemic practice, which in turn serve the goal of intellectual flourishing, one can be justifiably held to account.

The stress on the norms of practice serves another purpose as well. This purpose derives from the extensive debate on epistemic responsibility. We cannot be held to account for things we are not responsible for. In his influential work William P. Alston (1989) argued that the deontological conception of epistemic justification faces a challenge. On deontological conception, an agent being justified in believing that p means that there are no norms that prohibit the said agent from forming this belief. Yet, if epistemic norms of justification exist, one should be able to follow them in line with the principle that "ought" implies "can" (OIC principle). For an agent to be able to follow a norm of justification for belief they should "have an effective choice" whether to believe something or not (Alston, 1989, p. 259). That is, one has to exercise direct voluntary control over their beliefs. However, belief is not analogous to action. We do not change our beliefs according to the previously formed intentions, and we are not facing alternative possibilities about what to believe. Thus, the ability to exercise direct voluntary control is lacking, and justification cannot be conceived in normative terms – one *cannot* satisfy the "ought" of justificatory norms.

Alston's argument sparked a sprawling debate about the nature of responsibility for belief. One of the corollaries of his claim is that if one cannot follow the norms of justification in a familiar sense of

"following", one cannot be held accountable for their beliefs – we are not responsible for what we do not control. This conclusion seems counterintuitive to many, including myself. From the moral standpoint, people are responsible for espousing racist beliefs even though they might never harm another human being. From the standpoint of intellectual goods, it seems that we are responsible for not believing falsehoods, even though these might not have any bearing on moral on practical matters. As a result, it seems that one can be held to account for both morally and intellectually significant falsehoods.

I will go along with Alston and others and assume that we do not have direct voluntary control over our beliefs. Nevertheless, I want to claim that we exercise control over what we *do* to increase our chances of believing correctly. This leads me to assume that we exercise indirect doxastic influence over our beliefs. We form intentions as to how to proceed with the inquiry, and the result of the inquiry is to a large extent determined by the steps we both should and can perform in such a pursuit. And the norms governing the steps one should perform on the quest of attaining intellectual goods are the norms of epistemic practice.

It is a commonplace among the theorists of epistemic responsibility to glean approaches from those who write on responsibility in moral domain. A significant chunk of literature on epistemic responsibility (Steup, 2008; McCormick, 2011; McHugh, 2013, 2014) draws from Fischer and Ravizza's (1998) conceptions of guidance control and reasons-responsiveness. In a nutshell, guidance control does not presuppose the availability of alternatives to do or believe otherwise – it requires an agent to freely perform an action (or form a belief in the epistemic case) without the possibility to freely choose between the alternatives. This form of control is different from what Fischer and Ravizza call regulative control - "the dual power" to freely choose between alternative courses of action (or belief) (Fischer and Ravizza, 1998, p. 31).

All in all, guidance control is a species of direct doxastic control that does not presuppose prior formulation of intention for acting or believing in a certain way (Kruse, 2014, p. 2814). In the epistemic case, a belief is under guidance control if and only if it meets certain conditions: "*belief is formed or sustained by a mechanism that is the believer's own, and that is receptive and reactive to epistemic reasons*" (McHugh, 2013, p. 143). That is, perception, reasoning or memory are the relevant mechanisms that an agent treats as their own. On top of that, an agent should be *in principle* responsive to reasons for believing otherwise regardless of the actual availability of such reasons, and to be able to *accept* the reasons to believe otherwise in a counterfactual scenario.

For the purpose of determining the target of epistemic accountability this approach seems to be unsatisfactory. The reasons-responsiveness approach is too strong. Given that what justifies holding agents to account for their mistakes is the violation of the norm of intellectual flourishing, reasons responsiveness account presupposes that agents who do not violate this norm will be justifiably held accountable for their mistakes as well. Let's take our professor who confused husky and malamute. She made a mistake that derives from the workings of a cognitive mechanism which is her own (and she accepts ownership herself), and should the appropriate reasons been available to her, she would have recognized those reasons and accepted them. The same holds for K1 and K2 in the case of their mistaken beliefs. However, the professor did not violate a fundamental norm of intellectual flourishing, which means that there is no independent reason to hold her accountable for this mistake. Yet, if we accept reasons-responsiveness view both our professor and two other culprits will be accountable for their mistakes. If reasons-responsiveness cannot explain the discrepancy between the professor on the one hand and K1 and K2 on the other, it does not suffice as a ground for epistemic accountability.

What are the alternatives? I want to claim that a form of control necessary for accountability to arise is indirect doxastic influence we exercise over our beliefs (for similar account of doxastic responsibility see Meylan, 2013). Put crudely, we can *do* things in order to come to a correct belief, and what we thus *do* is, rather uncontroversially, under our control. For example, one can look for additional pieces of evidence for a target of inquiry, or one can question the premises that lead to a conclusion affirming one's biases, or one can abstain from believing something without evidence, or one can consult a peer reviewed article over a tabloid piece on the same subject. We perform these epistemic actions in pursuit of truth, knowledge or sufficient justification for our doxastic states. The outcome of these actions is not certain, even though the diligent performance of those drastically increases one's chances of believing correctly. In other words, our doxastic states are influenced by epistemic actions we perform while conducting an inquiry. And, without sowing too much controversy, these actions can be said to be under our direct control.

For example, sometimes, we are faced with the choice of whether to perform a certain action or not. While browsing an article on a topic of interest, one is faced with alternatives – should I read an article written by the author's contender who disagrees with them or not? I can either succumb to my natural laziness or go check the source that presents evidence to the contrary. As a consequence, these actions cause the formation of beliefs, albeit indirectly. The case is analogous to manslaughter caused by drunk driving: one is indirectly responsible for the deed, because the actions that played a significant causal role in the outcome were under direct control of the agent. Similarly, one is accountable for a mistake if an epistemic action causes one to adopt a mistaken or unjustified belief.

Stress on epistemic action allows to circumvent another complication encountered by many theories of epistemic normativity – norms of epistemic practice satisfy the "ought" implies "can" principle. These epistemically significant actions are regulated by derivative norms of epistemic practice. The

latter prescribe how one is supposed to proceed in complying with theoretical norms whenever it serves the goal of intellectual flourishing. The norms of epistemic practice provide general guidelines for action: consider alternative sources of evidence; do not reason by way of associations; do not believe things based on a hearsay; consult high-quality evidence rather than low-quality evidence. These norms are many, some of them domain specific - norms for sociologists and physicists would differ. And, it goes without saying, the roster I presented is in no way complete. The stress on norms of epistemic practice is important, since we want to exclude cases where non-epistemic actions and events, like excessive drinking or oversleeping, preclude an agent from attaining intellectual goods.

The intended upshot is that we form beliefs by performing a set of epistemic actions, whether mental or not. And we are supposed do so in compliance with the norms of practice whenever the fundamental norm of intellectual flourishing thus demands.

Chapter 2. Like adults: responsibility, excuses and exemptions

We have already started answering the second question – who qualifies as an appropriate target of epistemic accountability – by positing the fundamental norm of intellectual flourishing and norms of epistemic practice that agents must follow in order to satisfy the fundamental norm. Those who violate derivative norms of epistemic practice are appropriate targets of accountability. Two unresolved issues remain. First, who, in principle, cannot be a target of any kind of accountability for an epistemic failure – who is to be *exempted* from any kind of sanction for epistemic wrong. Second, one should explain what are the *excusatory* conditions for epistemic agents who committed an epistemic wrong (Strawson, 1962 (2008); Littlejohn, 2016). Let's group these two under the heading of *exculpatory conditions*. This distinction is important because different kinds of exculpatory conditions for rational and not (fully) rational agents obtain on different bases. Furthermore, the agents who do not have access to relevant epistemic resources are not to be treated as children or mentally ill. This is so not only because this would be profoundly unethical, but also because the agency cultivating purpose of holding agents accountable would not be fulfilled once one does not treat a 'culprit' as (fully) rational.

2.a. Exemptions.

I will have to address the issue of exculpatory conditions very briefly – the space I am allotted is limited. But it is important to exclude certain categories of agents from the pool of those who can be held accountable for their epistemic mischiefs. In addition to that, it is important to treat those who are appropriate targets of accountability as rational – the purpose of holding people to account is to cultivate a right sort of agency which also happens to be beneficial for attaining intellectual flourishing. We do not expect this from mentally ill, or in the moments of failure of rational capacities in rational

agents, and the expectations we have from children are slightly different from what we expect from adults.

It is commonplace to argue that exemptions from any kind of accountability obtain for not (fully) rational agents (Strawson, 1962 (2008); Littlejohn, 2016; Boult, 2020). One way to put it, if an individual does not possess rational capacities required to adhere to epistemic norms, they cannot be criticized or blamed for their failures. After all, "ought" implies "can", and those who cannot adhere to the requirements of epistemic nature are exempted from the relevant obligations and prohibitions. However, these are not the likes of the Truth norm, Knowledge norm or requirement for sufficient justification that are of interest here. First, it is hard to see in what sense one "can" adhere to theoretical norms simpliciter – *how exactly* one is to believe truly, or justifiably, or to know things? To cash out the sense of "can", as I argued before, one should posit epistemic actions that are, uncontroversially, under our direct control, and these actions can be performed successfully by following the norms of epistemic practice.

One *can* perform them in order to increase the chances of believing correctly. Therefore, the presence or absence of capacities to adhere to the norms of epistemic practice (and to perform epistemic actions) determines whether exemptions obtain. Second, stress on the norms of practice rather than theoretical norms is supported by the fact that mentally ill, intoxicated or distressed individuals can, occasionally, satisfy theoretical norms. What they cannot do is to adhere to the norms of epistemic practice. An individual with schizophrenia, or impairment of cognitive development, or a rational agent with severe depression would hardly be able to successfully scrutinize the inferences one makes, or consistently choose reliable sources over unreliable one's, or abstain from believing something in the light of insufficient evidence. For the lack of a better term, I will call this condition a condition of

practical epistemic rationality – if a person is incapable of doing things that positively influence their doxastic states, they are to be exempt from accountability.

Yet another consideration against restricting exemptions to theoretical norms, is that in some cases those who, intuitively, are exempt from accountability can "adhere" to the norms of epistemic practice all right. Imagine a person who counts grass stems in their garden every day. Furthermore, this person invents a highly accurate probabilistic model that allows them to predict the number of grass stems for any given day of the year. They possess a set of properly justified, true beliefs about the subject matter of interest. Furthermore, in pursuit of their bizarre goal, this individual successfully and consistently adheres to the norms of epistemic practice. But something is out of place with this person – this is the only thing they are interested in. What is wrong is that this person fails to adhere to the norm of intellectual flourishing (Brogaard, 2014, p. 16). If they are in principle incapable of pursuit of intellectual goods and demonstrate the incapacity to assign intellectual priorities, one should exempt them from accountability. After all, this goes beyond the interest in obscure topics and activities that sometimes allow individuals to flourish – we can confidently say that this person would rather invest their energies in something else! Let's call this a condition of intellectual rationality – even if a person demonstrates excellency in following the norms of epistemic practice, a pathological inability to pursue worthwhile intellectual goals in adherence to the norm of intellectual flourishing exempts from accountability.

In the end, we have two conditions that have to be jointly satisfied for an agent to be accountable for their mistakes. First, one has to demonstrate the capacity to engage in the epistemic practice – to be, in principle, capable of acting in a way that would positively influence one's doxastic states. Second, one has to be, in principle, capable of assigning priorities to one's intellectual pursuits – to invest their limited energies into worthwhile endeavors. If an individual demonstrates lack of rational capacities

required to fulfill either of those conditions, it simply does not make sense to hold them to account. Bearing in mind the corrective purpose of holding agents to account for their mistakes, it would not be feasible to do so with respect to those who are simply incapable of improving in these respects.

2.b Excuses

What about excuses? Excuses obtain whenever an agent violates a certain norm, yet at the same time the violation in itself does not present grounds for holding this person to account. That is, for an agent to be excused, there should exist some positive factor that makes sanctioning them for a mistake unnecessary and unjustified. There exist different positions about excusatory conditions for epistemic failures. Some argue that excuses obtain whenever an agent forms a false belief, but has a general disposition to believe truly, and has the same belief as someone with a relevant disposition would have (Williamson, forthcoming). Others stress that excuses obtain whenever a person demonstrates rational excellency despite having a false belief, or forming a belief without sufficient reason (Littlejohn, 2016). ,I am not going to decide on which approach is better. First of all, the discussion mentioned above pertains to debates on theoretical matters, and I have no pretensions to present any contribution to those. Second, I want to be consistent. Thus, I once again call the notion of intellectual flourishing to the rescue.

Recall our professor. As we have seen before, she does not quite qualify as someone who is to be held to account for her mistake. Unlike in the case of exemptions, our professor did not exhibit failures of practical epistemic or intellectual rationality. However, she still violated a theoretical norm – her belief is false and is apparently lacking in justification. What excuses her is that theoretical norms, in the absence of an independent obligation to follow them, do not have authoritative normative force. A requirement to care about intellectual goods – the norm of intellectual flourishing – does possess such

an authoritative force. And in the absence of the obligation to believe truly for the sake of fundamental intellectual good, a rational agent is excused. After all, it is hard to think of justification for sanctioning someone for mischief that is in no way harmful. Thus, given that there is no obligation for the professor to educate herself on the breeds of dogs, she has a good excuse for not believing correctly – it does not undermine intellectual flourishing.

The same cannot be said of K1 and K2. Their beliefs, in turn, hinder the ideal of intellectual flourishing. First, they formed false and unjustified beliefs about the topic that we deem important to get right for the sake of our intellectual well-being. Thus, believing falsely or without justification would not simply violate a theoretical norm, but also a fundamental intellectual norm. Second, since the subject matter of inquiry is itself an important intellectual good, one should adhere to the norms of epistemic practice in order to form an appropriate doxastic state. Both K1 and K2 violated those norms. All in all, they cannot be excused on the same grounds as our professor – their beliefs and the way they came to those beliefs are harmful for their, as well as communal, intellectual well-being.

To conclude, not all cases of unwarranted doxastic states require sanctions. First, some agents should not even be included in the pool of those who deserve sanctions for their epistemic missteps – they are entirely exempt from accountability. That is, if one demonstrates a significant failure of rationality – either practical epistemic, or intellectual – one may conclude that they are *incapable* of navigating a diverse set of norms governing our epistemic conduct. If they, in principle, cannot adhere to such norms, there is no obligation for them to do so. In the absence of such an obligation, they cannot be held to account for whatever mistakes they make. Second, not all unwarranted doxastic states held by rational agents deserve sanctions. Some mistakes are insignificant enough for us to still regard a person in a positive light. In cases where an unwarranted doxastic state does not adversely affect intellectual well-being one is to be excused from accountability – sometimes, no harm is being done by believing

falsely. In the end, for sanctioning agents for their epistemic mischiefs they have to fulfil two criteria: 1) to demonstrate intact rational capacities and 2) to violate the norm of intellectual flourishing.

Chapter 3. Degrees of accountability

In the previous chapter I tried to lay down the groundwork for addressing the main question of this work: how to explain our intuition that we should hold people accountable to different degrees relative to the epistemic resources available to them. I isolated the pool of epistemic norms that warrant accountability for their violation – these are the norms of epistemic practice in the service of the fundamental norm of intellectual flourishing. Next, I tried to show who and in which cases is excused or exempted from accountability for epistemic faults. Excuses obtain in cases where the fundamental norm of intellectual flourishing is not violated. Exemptions obtain in cases where a person is incapable of exhibiting intellectual and/or practical epistemic rationality. In all other cases, the violations of epistemic norms that are characteristic of a failure to adhere to the fundamental norm of intellectual flourishing warrant accountability. That is, we are justified in holding agents to account for the unwarranted doxastic states they form – when they believe things they should not. Now, let's turn to what determines the degrees of accountability, and how the availability of epistemic resources affects the degree to which we can hold agents accountable.

I will proceed, first, by unpacking the notion of epistemic resources. Second, I will address the issue of how availability of epistemic resources influences *dispositions* to follow the norms of epistemic practice. Lastly, I will show how availability of epistemic resources influences the degree to which we hold agents accountable.

Before I go any further, I would like to underline that the account I present is but a rough schema. Real life is much more complicated than this philosophical abstraction. What I hope to achieve is to lay down the general principles which, more or less reliably, would allow us to see how privileges (or lack thereof) of epistemic nature influence our accountability for epistemic mischiefs.

3.a Epistemic resources - the question of disparity

I invite the reader to recall their academic path. One can start from the school years, moving gradually to the bachelors, and after that to graduate school. One can also recall what their intellectual environment looked like, and with whom did they spent their pastime, and the intellectual standards applied to even the casual conversations. Then, let's compare what you managed to gain from the opportunities you had to the less fortunate individuals. Here, I will lay down such an *approximate* comparison.

In rather vague terms, one's intellectual success – the breadth of erudition, ability to perform complex epistemic tasks, ability to assign degrees of importance to intellectual deliberations etc. - is largely influenced by the availability of epistemic resources. *What* you have been taught is equally important to the *know-how* of epistemic practice you have been taught. In other words, epistemic resources can be classified as factual and practical.

When it comes to factual epistemic resources, a person from a more privileged background, by way of availability of better education and having well-educated people around simply has an opportunity to know more facts. The institutional environment available to me, for example, provides me with much more reliable information than someone who was orphaned by the educational system earlier in life. On the practical side, I had an opportunity to see *how* people proceed with their inquiries, or what they *do* to influence their doxastic states in order to comply with theoretical norms; or how they decide whether such compliance is important. I can observe them working, they can tell me how one should operate epistemically, and I can practice myself, subjecting my deliberations to evaluation by my peers and superiors.

Both comparative wealth of facts available and the opportunity to practice in one's intellectual deliberations make an agent more prone to believe correctly. It is hard to draw a clear distinction

between factual and practical epistemic resources. After all, acquaintance with high-quality work can not only provide us with relevant facts, but also influence epistemic practice – in the light of the facts I learn, I can guide my actions accordingly. For instance, after acquiring knowledge or true beliefs about certain facts, in the future I would want to strive towards similar standards of evidential support and argumentation in the information I consume. Consequently, I will be more prone to *do* things to increase the probability of me getting to the correct belief – looking up different sources, comparing those sources, or just saying to myself "ok, this piece seems too sloppy, not sure I should take what is written here for granted".

On top of that, when it comes to epistemic "know-how", one also learns what one *should* do to believe correctly. The norms that govern successful epistemic actions are usually learned, so that one would be able to practice in their epistemic agency according to a normative standard. Whenever one practices operating in the complex epistemic environment with the possibility of external evaluation, one has a chance to grasp what one *should* do to achieve success. That is, the range of norms that govern epistemic actions is internalized through learning and practicing according to the standards employed in intellectually promiscuous environments. And as a result of continuous learning and practice, individuals get a chance to achieve consistency in successful manifestations of their epistemic agency.

On the other hand, those who are less fortunate cannot enjoy the same privileges. First, the unavailability of high-quality education presupposes that a person would not be exposed to the same fact-rich information. On top of that, the lack, or the lesser quality of education would entail the lack of practical epistemic skills exhibited by superiors and, by extension, by less privileged learners themselves. This, in turn, would mean that the external evaluation of epistemic success is going to be based on much shakier foundations. And in an environment like this, a person would be highly unlikely

to possess high-quality factual information or to develop skills necessary for conducting inquiry in a diligent manner. In other words, the lack of resources makes it less likely that a less privileged person will know enough, and, on top of that, it decreases the chances that they will know what they *should* do to attain intellectual goods. In the end, in most cases consistency and success of one's epistemic agency would be seriously undermined.

To spell out the difference between the two categories of agents, I want to introduce the notion of disposition – agents are in possession of dispositions to act in accordance with the norms of epistemic practice. Reference to dispositions allows to highlight the two aforementioned factors which influence our epistemic lives – consistency and success in manifestations of epistemic agency. Of such dispositions there is a variety, for instance: the disposition to choose high- over low-quality sources; disposition to consider alternative evidence; disposition to abstain from believing things based on hearsay; or a disposition to assess validity of one's inferences. Possession of such dispositions is largely influenced by the availability of relevant resources. Subsequently, I will provide an account of how the presence or absence of relevant dispositions affects the *degree* of epistemic accountability.

Among the philosophers venturing into the domain of epistemic normativity, the notion of disposition to follow epistemic norms is by no means new (Williamson, forthcoming; Boult, 2021). Dispositions to follow epistemic norms are typically understood as a propensity to believe truly, or with sufficient justification (Williamson, forthcoming, p. 13). That is, one has a disposition to follow a truth-related epistemic norm if "one is the sort of person who complies with" such a norm in most cases (Ibid.). And, the story goes, those who have a disposition to follow truth-related norms for belief are excused in cases where they are mistaken – they made a mistake despite being epistemically virtuous individuals.

However, in line with what I have been saying above, I want to distance myself from the approaches that see dispositions simply as propensities to follow theoretical norms. Recall, that theoretical norms require that one's belief should aim at truth, constitute knowledge or be sufficiently justified. However, adherence to theoretical norms is not always necessary, and sometimes may hinder attainment of worthwhile intellectual goods. When adherence to theoretical norms is warranted by the norm of intellectual flourishing, it largely concerns inquiry into the subject matter we deem important for our intellectual well-being. And to exhibit warranted care towards the subject matter of importance, one should follow the practical epistemic norms guiding that inquiry. The disposition to follow theoretical norms as an excusatory condition does not track the difference between intellectually significant and insignificant cases where theoretical norms are applicable.

On top of that, presence of a disposition to follow theoretical norms is not suitable for being an excusatory condition on my account – I tried to argue that a person is excused on different grounds, namely when they do not violate the fundamental norm of intellectual flourishing. That is, the function dispositions to follow epistemic norms perform in the literature is different from the function I want to assign to them. In the end, I just borrow the term, but not in the way that bears connection to how it is currently treated by other philosophers.

But how are dispositions to follow norms of epistemic practice related to availability of epistemic resources? To acquire a disposition to follow norms of epistemic practice is to get hold of an intellectual habit, ability or any other kind of propensity to successfully and consistently do things that positively influence one's doxastic states. And, as trivial as it sounds, lack of the relevant epistemic resources hinders the prospects of acquiring the relevant dispositions. If one does not know what one should do to increase the chances of forming proper doxastic states, then one cannot acquire a general propensity to follow the norms that prescribe what one should do to this end. Furthermore, if one's

environment does not allow to practice one's epistemic agency, the likelihood of acquiring relevant dispositions diminishes. For example, it is a truism with which hardly anyone would disagree that one *should* consult alternative sources prior to forming a belief on a complex subject matter. Yet, if one lacks knowledge as to where and how one is to conduct this search, or which sources would be reasonable to consult, the likelihood that this individual will be able to develop a disposition to act in adherence to this norm decreases. Similarly, if an agent did not have an opportunity to hone their epistemic "skills" by practicing under "supervision", it also makes attainment of a disposition to follow the aforementioned norms less probable.

To sum up, abundance of epistemic resources within one's reach has a positive influence on one's ability to follow norms of epistemic practice successfully and consistently. The lack thereof, on the other hand, makes the likelihood of an agent lacking relevant dispositions much higher.

3.b Dispositions and degrees of epistemic accountability

As I tried to argue above, the lack of epistemic resources manifests itself in the lack of dispositions to follow the norms of epistemic practice and, thus, to perform epistemic actions that positively influence attainment of appropriate doxastic states. Provided that these dispositions depend on the access to relevant knowledge and the opportunity to practice, the intellectual environment is extremely important for an agent to be able to acquire and cultivate those dispositions. Now, let us take a look at how the lack of opportunities to acquire and cultivate relevant dispositions translates into a lesser degree of accountability.

Before we start, let's recall how the intellectual portraits of our characters look like. K1, at the time of a conversation, worked in a tram depot. After the ninth grade in a public school, he started his studies

in a technical school. Despite modest educational opportunities, an absolute lack of systematic education in humanities and austerity of his intellectual environment, he retained commendable level of curiosity and consumes a lot of information that he deems true and reliable. The problem is, he cannot always tell good sources from bad ones, and forms a bunch of false beliefs based on sources that just seem plausible. In terms of epistemic actions needed for attainment of appropriate doxastic states, he could not perform those successfully enough. That is, he could not consistently follow the norms regulating epistemic practice. When he expresses his belief that all stateless societies were extremely violent, he did so based on plausibly looking sources he encountered on Facebook. He hardly knew any better – some citations from historians (or "historians") would do as a sign of a decent work. That is, the epistemic resources in his possession were too scarce to influence consistency and success in conducting epistemic actions.

Given that both success and consistency condition in his epistemic conduct are lacking, one can say that he lacks relevant dispositions. For instance, even though he seems to be choosy as to sources he uses, he cannot successfully and consistently follow the practical norm that requires to choose high-quality evidence for beliefs – the notion of high-quality evidence he has is too scant to ensure success in adherence to the said norm. His ability to follow the requirement to consult alternative evidence is similarly hindered – the resources in his possession do not provide knowledge about where to look for the alternatives, and how to reliably tell good sources from the bad ones. For the sake of the argument, let's assume that the same holds true of his ability to perform other epistemic actions in compliance with the practical norms.

In his early life, K2 received his education in high-end educational institutions. He also received his *D.Phil.* from the University of Oxford, where he was taught to conduct research in history. On top of that, he did a post-doc a couple of years after his *D.Phil.* This suggests he reaped the fruits of being in

the intellectually rich environment – he had an opportunity to meet a lot of intelligent people, acquaint himself with a lot of filtered and reliable information and subject his doxastic states and epistemic practices to external evaluation by peers and superiors. In the end, however, he still expresses his belief that stateless societies were extremely violent, based on only one study of a limited pool of contemporary tribal communities.

It is hardly feasible he did not know any better – for some reason he did not act on what he knew, or at least was supposed to know. It is reasonable to assume that the epistemic resources in his possession – both factual and practical – are sufficient to ensure possession of dispositions to successfully follow the norms of epistemic practice. For example, he could easily comply with the requirement to consult high-quality sources before forming beliefs – his notion of high-quality sources is much more elaborate than that of K1. To K2's credit, the source he consults is a peer-reviewed article, with extensive evidence, reliable research methods and carefully drawn conclusions. Where K2 fails, it is in drawing an appropriate inference from the evidence he has, as well as in consulting other sources that could have influenced his belief.

Recall, that the article at hand states that in some contemporary tribal societies in South America violent death toll is extremely high. The authors overtly state they cannot extrapolate their findings to other tribal communities. K2, in turn, claims that stateless societies were extremely violent throughout human history. Assuming he acted in good faith, he failed to draw a correct inference from the data supplied by his source, and he failed to consult further evidence that could support or undermine his belief. Furthermore, throughout his book we can see a lot of instances of sweeping statements with little evidential support, bold conjectures about the future of humanity and overall lack of diligence in conducting his inquiry. That is, K2 demonstrates the lack of consistency and success in attending to the norms of epistemic practice – he does not manifest relevant dispositions. And provided the

abundance of epistemic resources at his disposal, these are very serious failures, seemingly uncharacteristic of richness of his epistemic environment.

What we can say of K2 – and cannot say of K1 – is that his performance is uncharacteristic of how someone with the same resources would approach inquiry. That is, a person with similar access to epistemic resources – his ideal counterpart - would do better than K1. By "ideal" I do not mean a perfect epistemic agent, but someone who possesses the set of dispositions that one could reasonably expect from them, taking into consideration the availability of epistemic resources.

To unpack which factors, contribute to evaluation of degree of one's accountability, let's start with K2 and his ideal counterpart. In purely quantitative terms, we can confidently assume that the epistemic resources at his disposal would allow to develop dispositions to follow a wide array of practical norms. The sheer abundance of facts that one has an opportunity to learn in academic environment, as well as multiple opportunities to hone epistemic "know-how" are supposed to positively influence one's ability to successfully and consistently act in ways that would positively influence his doxastic states. With such a background, an ideal counterpart of K2 would possess a disposition to consult alternative evidence, evaluate the validity of one's inferences, to use sources that rely on a robust research methodology rather than intuition, etc.

Possession of such dispositions would not guarantee the impossibility of failure, but it would drastically increase the chances for an ideal agent to believe correctly. On top of the ideal set of dispositions, another factor in evaluation of the degree of accountability is the influence such dispositions have on the formation of a given belief in each particular case. In other words, how difficult it would be for an agent to believe correctly with a set of dispositions they ideally should have. After all, even the most diligent knowers make mistakes.

In the case of K2, the dispositions an ideal agent has would almost certainly guarantee absence of the mistake. After all, it is not that hard to consult other anthropological studies in order to avoid a false belief, or to perform other actions in adherence to the norms of epistemic practice. Thus, we can evaluate the degree of accountability by factoring in (1) the dispositions an agent should have relative to availability of epistemic resources and (2) epistemic difficulty[1], I.e., the influence such dispositions would have on a belief in question.

The verdict for K2 is not very flattering. First, when it comes to the dispositions he should have had, it seems hardly likely that they correspond to those of his ideal counterpart. Recall that this is not the only instance of the violation of the norms of epistemic practice by K2 – there are a number of indications of a sloppy adherence to such norms. Arguably, a person with epistemic resources equaling K2's should demonstrate a much more consistent pattern of successful epistemic action and would form beliefs much more carefully. Second, the ideal set of dispositions that K2 should have would not allow the formation of a false belief K2 ended up with.

Performance of epistemic actions in accordance with the norms of practice would almost guarantee a correct belief: K2 could have reviewed a bigger chunk of existing literature on tribal communities; directly consulted specialists in the field; or at least questioned whether the very source he uses supports his belief. In the end, K2 fails on both counts – he demonstrates the lack of relevant dispositions that correspond to the wealth of epistemic resources at his disposal and, on top of that, should he be in possession of relevant dispositions, it would almost be impossible to form a false belief in question.

[1] I am borrowing the term "epistemic difficulty" from others (Boult, forthcoming) without adherence to the existing analysis of the said notion.

I add a simplified illustration of how these factors are playing out in our evaluation of the degree of K2's accountability. I add a schematic graph which presents a very rough approximation of how availability of epistemic resources affects the degree of accountability. Numerical values serve a purely illustrative purpose – they roughly correspond to our colloquial notions of "more" and "less". Let's start with K2.

The numerical value of the vertical axis – ID - represents the ideal set of dispositions an agent should have relative to the epistemic resources available. The more epistemic resources are available to the agent, the higher numerical value on the vertical axis. The horizontal axis – ED – represents the epistemic difficulty of acting in accordance with the dispositions one ideally should have. Here, the maximum numerical value remains constant – 1. Provided that K2 demonstrates the overall lack of dispositions to follow the norms of epistemic practice, and it would almost be impossible to fail to comply with those norms in the case at hand, both on ID and ED axis he would score rather low.

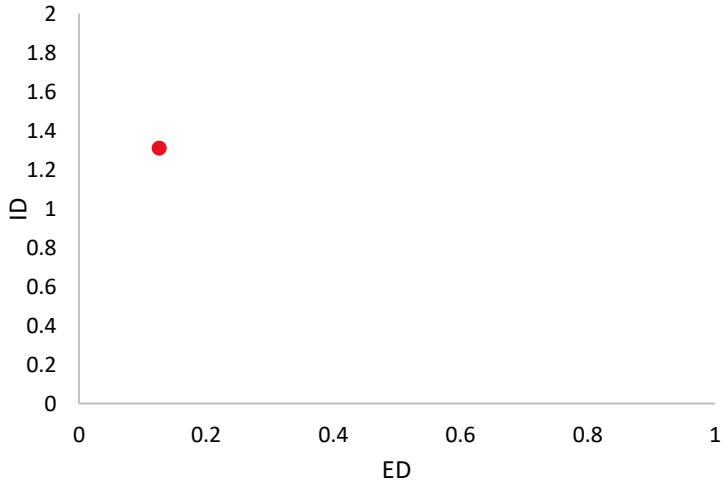

The lower one scores on both counts, the higher degree of accountability for the epistemic mischief. If an agent who had all the opportunities to acquire a certain disposition fails to manifest it on multiple occasions and, subsequently, to act accordingly in the absence of factors that exempt them from accountability, they deserve a harsher response. Recall that the purpose of holding knowers to account

is to enhance their epistemic agency. And in the cases of those who are in possession of the relevant epistemic resources, mild criticism might not be sufficient to make them change their ways. Simply telling K2 that he should have consulted more sources or evaluated whether the source he uses supports his belief is not enough. It is hard to imagine he did not know that one should have consulted more evidence or considered whether the evidence he uses supports his belief. As we established above, K2 is guilty on both counts – he failed to demonstrate possession of relevant dispositions despite having the means to acquire them, and failed to act as he should have – and could have - in the case at hand. In the end, this kind of failure would hinder intellectual well-being to a larger extent than if his mistake was just a misstep. And given the serious intellectual consequences of his mischief, he deserves harsher treatment – harsh enough to alter his epistemic agency in a manner that would be conducive to fulfilling the goal of intellectual flourishing.

Let's take a look at K1. When it comes to the first factor, his ideal counterpart would not look too different from him. With secondary school education, technical school and errand manual jobs his ideal counterpart would have roughly the same dispositions. Scarce factual landscape, lack of sufficiently educated superiors to evaluate his epistemic success, and other factors would adversely influence both K1's and his ideal counterpart's ability to success and consistency of their epistemic performance. In the end, the range of epistemic actions K1 could successfully perform would be roughly the same as what he should ideally be able to do – the dispositions in their possession would not differ much, if at all. For instance, he would not be able to consistently and successfully choose the high-quality over low-quality evidence, or consistently find alternative sources of specialized information, or evaluate the feasibility of the research methods employed by the sources he uses. Thus, K1 scores higher on this count – his dispositions are in line with what they should be, considering the epistemic resources at his disposal. As per the second factor, the ideal set of dispositions would highly likely land him in the same belief. With the resources at hand, reviewing up to date research by

anthropologists and historians might not be feasible, due to the language barrier and lack of exposure to the texts in social sciences – how to tell good research from bad? Provided the scarcity of means, the set of dispositions in K1's disposal would be very unlikely to positively influence his belief. Here as well, K1 scores higher than K2 – the correspondence to his ideal counterpart is much more proximate. Consequently, the failure on behalf of K1 is less severe – he ended up believing what a person with access to the same epistemic resources normally would. In our rough visualization, the numerical value for the ID axis would be less than in the case of K2, since the range of epistemic resources available to him is much more modest.

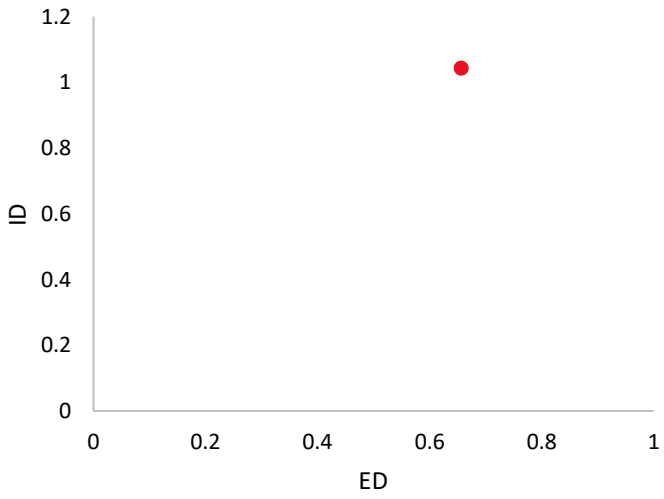

As we can see, he scores higher on both counts. In the end, K1 seems to be accountable to a lesser degree than K2. The set of dispositions at his disposal corresponds to that of his ideal counterpart, and, on top of that, the dispositions at hand would be unlikely to positively influence his doxastic state. Bearing in mind the corrective purpose of holding agents to account, indignant response with an emotional sting attached to it would be an overkill. The corrective purpose of holding agents to account requires to use the means that would enhance one's epistemic agency, not undermine it. And emotionally charged responses targeting those who do not have sufficient means to excel in epistemic

practice will not do the job – the more pressing task is to acquaint a culprit with what they had to do epistemically. In other words, if one is in the task of holding someone to account for an epistemic mischief and the target does not possess relevant resources, one should supply them with some. To this end, a mere criticism would do: K1 should be let to know that he is in the wrong, but also how to avoid those situations in the future. Unlike K2, he needs additional means to improve in his epistemic agency, such as factual inputs, prompts as to which sources to use and how to tell good ones from the bad ones, and so on.

Undoubtedly, there would be many cases in between. For example, agent's dispositions would correspond to those of their ideal counterparts yet fail to positively influence their doxastic states despite the high likelihood of doing so. For example, one can imagine that K2 is not as prone to be mistaken in his judgements, and the belief that we want to hold him accountable for is influenced by a solitary misstep. Let's call this more diligent twin-brother of our character K3. In such a case, he would be accountable to a lesser degree, and deserve a less harsh response than K2, since he fails on only one account of the two – despite him being a diligent epistemic agent whose dispositions correspond to the epistemic resources in his possession, he failed to act in a way that would positively influence his belief. Our schema for K3 will look something like this:

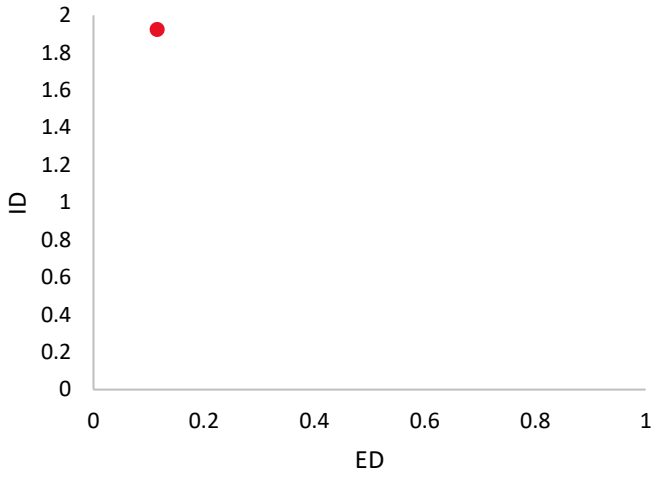

Such failures, no doubt, deserve criticism. If an agent's failure is a result of neglect, laziness or any other sin we as knowers often commit, one's response should be such as to forestall failures of epistemic agency in the future. Here, as in the case of K2, a mere criticism might not be sufficient – the agent has all the resources to act accordingly. On the other hand, most of the time K3 puts these resources to good use and is typically disposed to follow the norms of epistemic practice. We might add some force to the criticism - "K3, this is so uncharacteristic of you, what a rookie mistake you've made" - that would go beyond a mild criticism towards K1 and fall short of an indignant response reserved for K2. Such a rebuke would suffice for correcting K3's epistemic agency and, as is characteristic of rebukes like this, it would have some emotional charge attached to it. On the other hand, if the set of dispositions has a scanter chance to influence the doxastic state of an agent – it scores higher in terms of epistemic difficulty - we might mute our criticisms to an extent since in such a case the agent is accountable to a lesser degree.

The schema I propose is, no doubt, extremely austere. One can say that it does not account for other reasons as to why human beings demonstrate neglect towards epistemic practice and intellectual goods. What about natural stubbornness, or arrogance, or any other vice that makes some of us approach inquiry with neglect? I want to say that this worry does not present a threat to my account.

My task was to try to explain how social condition affects agent's accountability for epistemic faults. This task is rather modest, and it does not include enumeration of every vice out there, whether intellectual or moral. Furthermore, it seems, intellectual or moral vices already affect the degree to which an agent is accountable, even though there is no need to stress the role of each and every of them. When a person with abundant epistemic resources in their possession fails to develop relevant dispositions in the absence of exculpatory conditions, we can suspect that something is wrong with their intellectual or, perhaps, even moral character. Similarly, if a person demonstrates neglect towards the norms of epistemic practice despite possessing relevant dispositions, we can suspect some kind "vicious" trait occasionally manifesting itself – something of which we are all guilty. In any event, those manifesting those traits in their epistemic conduct would score low both on account of correspondence to their ideal counterpart, as well as on account of epistemic difficulty. And the more vices influencing one's epistemic standing can be detected in an individual, the worse they would score on both counts.

Conclusion

I hope to have provided an account of epistemic accountability which, however schematically, manages to address the issue of inequality of access to epistemic resources. That is, people often fail not because they possess "bad" intellectual character, even though some of us might. Rather, the less privileged knowers do suffer from the lack of opportunities to develop relevant abilities, habits or other features that greatly help individuals to successfully engage in epistemic activities.

To cash this out, I had to start from the basics – what do we need a theory of epistemic accountability for, and what is the purpose of holding agents to account for their epistemic mischiefs. As I tried to argue, the most plausible purpose one can find for sanctioning people for their mistakes is to make them improve as knowers.

The talk about this overarching purpose is supposed to point in the direction of why we should care about other's performance as epistemic agents. There is a reason why we care about truth, knowledge or justification for doxastic states we – and individuals around us – form. And, as I tried to show, there is a sense in which we *should* care about those goods. And this "should" must be authoritative enough to warrant accountability for mistakes we make. To cash out the sense of authoritative "should" I had in mind, I introduced the norm of intellectual flourishing – *you should believe in compliance with theoretical norm(s) only if thus believing serves the goal of intellectual flourishing*. That is, we should care about correctness of our doxastic states whenever it affects our intellectual well-being. To this end, we should *act* in ways that are conducive towards forming correct beliefs in cases where it matters.

The possibilities to successfully and consistently act in the ways that positively influence our doxastic states are not always there. Sometimes, people cannot consistently act in a way they should do to believe correctly – the standards of compliance with the norms of epistemic practice are not being made clear to all, and not all get a chance to hone their epistemic skills. It is a truism, but the lack of

factual information, as well as the lack of practicing in epistemic pursuits significantly affects one's chances to come to correct beliefs when it matters. And, as I hope to have shown, this affects the degree to which we can justifiably hold less privileged agents to account.

Where does that leave us? It seems that accounts of epistemic accountability, blame or excuses that want to rely on dispositions (Boult 2019; Williamson n.d.), rational excellency (Littlejohn 2016) or, perhaps, intellectual virtues as excusatory conditions miss one important thing. People are not simply intellectually virtuous, or disposed to believe truly, or demonstrate rational excellency. These are not default settings. We learn those things. And the circumstances in which we learn are disparate and unequal. And when discussing social matters, one has to factor in the unpleasantries of social reality for one's theory to be explanatory for actually existing phenomena.

References

Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–299. https://doi.org/10.2307/2214077

Boult, C. (2019). Excuses, exemptions, and derivative norms. *Ratio*, 32(2), 150–158. https://doi.org/10.1111/rati.12228

Boult, C. (2020). There is a distinctively epistemic kind of blame. *Philosophy and Phenomenological Research*, *103*(3), 518–534. https://doi.org/10.1111/phpr.12726\

Boult, C. (2023). Degrees of Epistemic Criticizability. *The Philosophical Quarterly*. https://doi.org/10.1093/pq/pqad053

Boult, C. (forthcoming) *The Puzzle of Epistemic Blame: On the Nature and Norms of Epistemic Relationships*

Brogaard, B. (2014). Intellectual Flourishing as the Fundamental epistemic Norm. In C. Littlejohn & J. Turri (Eds.), Epistemic Norms New Essays on Action, Belief, and Assertion (pp. 11–33). Oxford University Press.

Brown, J. (2018). What is Epistemic Blame? *Noûs*, *54*(2), 389–407. https://doi.org/10.1111/nous.12270

Fischer, J. M., & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.

Fricker, M. (2014). What's the Point of Blame? A Paradigm Based Explanation. *Noûs*, *50*(1), 165–183. https://doi.org/10.1111/nous.12067

Kauppinen, A. (2018). Epistemic Norms and Epistemic Accountability. *Philosophers' Imprint*, 18, 1-16.

Kruse, A. (2017). Why doxastic responsibility is not based on direct doxastic control. *Synthese*, 194(8), 2811–2842. https://doi.org/10.1007/s11229-015-0951-6

Littlejohn, C.M. (2016). A Plea for Epistemic Excuses.

Madison, B. J. C. (2017). On justifications and excuses. *Synthese*, *195*(10), 4551–4562. https://doi.org/10.1007/s11229-017-1418-8

McCormick, M. (2011). Taking control of belief. *Philosophical Explorations*, 14(2), 169–183. https://doi.org/10.1080/13869795.2011.569745

McGeer, V. (2013) Civilizing Blame. In D. Coates, J. & Tognazzini, N. A. (Eds.), *Blame: Its Nature and Norms*. Oxford University Press. pp. 162—188

McHugh, C. (2013). Epistemic Responsibility and Doxastic Agency. *Philosophical Issues*, 23, 132–157. http://www.jstor.org/stable/26382646

Meylan, A. B. (2013). Foundations of an Ethics of Belief. In De Gruyter eBooks. https://doi.org/10.1515/9783110327816

Piovarchy, A. (2020). What do We Want from a Theory of Epistemic Blame? *Australasian Journal of Philosophy*, *99*(4), 791–805. https://doi.org/10.1080/00048402.2020.1833353

Rettler, L. (2017). In defense of doxastic blame. *Synthese*, *195*(5), 2205–2226. https://doi.org/10.1007/s11229-017-1332-0

Robert S. Walker, Drew H. Bailey, Body counts in lowland South American violence, *Evolution and Human Behavior*, Volume 34, Issue 1, 2013, Pages 29-34, ISSN 1090-5138, https://doi.org/10.1016/j.evolhumbehav.2012.08.003.

Scanlon, T. M. (2010). *Moral Dimensions: Permissibility, Meaning, Blame*. Amsterdam University Press.

Schmidt, S. (2021). Epistemic Blame and the Normativity of Evidence. *Erkenntnis*. https://doi.org/10.1007/s10670-021-00430-9

Sher, G. (2006). *In Praise of Blame*. Oxford University Press.

Steup, M. (2008). Doxastic Freedom. *Synthese*, 161(3), 375–392. http://www.jstor.org/stable/27653701

Strawson, P. F. (2008). *Freedom and Resentment and Other Essays*. Routledge.

Vargas, M. (2011). *Building Better Beings: A Theory of Moral Responsibility*. Oxford University Press.

Wallace, R. J. (2011). Dispassionate Opprobrium: On Blame and the Reactive Sentiments. *Reasons and Recognition Essays on the Philosophy of T.M. Scanlon*, 348–370. https://doi.org/10.1093/acprof:oso/9780199753673.003.0015

Williamson, T. (forthcoming). Justifications, excuses, and sceptical scenarios. In J. Dutant (Ed.), *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*. Oxford University Press (forthcoming).