AFFECTIVE EXPERIENCE AS A SOURCE OF KNOWLEDGE

Lara Apolline Jost

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

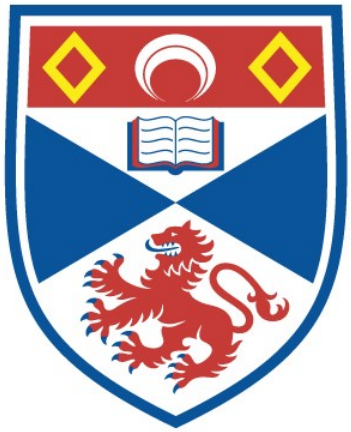

2023

Full metadata for this thesis is available in
St Andrews Research Repository
at:
http://research-repository.st-andrews.ac.uk/

Identifiers to use to cite or link to this thesis:

DOI: https://doi.org/10.17630/sta/387
http://hdl.handle.net/10023/27364

This item is protected by original copyright

This item is licensed under a
Creative Commons License

https://creativecommons.org/licenses/by-nc-nd/4.0

Affective Experience as a Source of Knowledge

Lara Apolline Jost

University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

March 2023

Abstract

In analytic philosophy, appropriate sources of knowledge do not include affective experience. In this thesis, I oppose this commonly accepted view and argue in favour of valuing the contribution of affective experience to knowledge. My argument unfolds in two parts. In the first part of the thesis, I propose a positive argument favouring the inclusion of affective experience in our sources of knowledge. My main contribution is the Reflective Model, a novel account of affective experience as a source which responds to the common objections towards the epistemic value of affective experience. This new model is built on the idea that affective experience is a reason-tracking ability, which incorporates a reflective element through internal and social calibration. I illustrate this phenomenon by discussing how members from marginalised groups identify microaggressions through their emotions. In the second part of the thesis, I offer a negative argument: given the negative ethical and practical consequences of excluding affective experience, we should include it in our sources of knowledge. I conceptualise a new form of epistemic injustice, *source based epistemic injustice*, which highlights unfairness towards knowers who utilise unapproved sources of knowledge like affective experience. In addition, I discuss two applied cases where testimony based on affective experience is unfairly downgraded: testimony about lesser-known illnesses and testimony about everyday racism. I argue that the first case leads to the epistemic exploitation of patients, whilst the second gives rise to dialectical white scepticism, a pernicious form of white scepticism which has not been previously theorised. Altogether, I argue that legitimising affective experience as a source of knowledge would enable us to incorporate a greater number of epistemic contributions in our epistemological system, thus reducing instances of injustice and silencing experienced by members of marginalised groups.

Candidate's declaration

I, Lara Apolline Jost, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 58,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2019.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 22/3/23 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 22/03/2023 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the

University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Lara Apolline Jost, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 22/3/23 Signature of candidate

Date 22/3/2023 Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Lara Apolline Jost, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 22/3/2023 Signature of candidate

Cette thèse est dédiée à mes parents,
Bernard et Nathalie, qui m'ont appris à embrasser
mes émotions dès mon plus jeune âge

This thesis is dedicated to my parents
Bernard and Nathalie, who taught me to embrace
my emotions from a young age

Contents

Acknowledgements

I want to thank first and foremost Justin Snedegar for being the most supportive and incredible primary supervisor one could have hoped for. I don't think my doctoral experience would have been this enriching without him, and I'm glad he encouraged me to take my project in various unexpected directions. I want to also thank Catarina Dutilh Novaes who has been an amazing secondary supervisor. I deeply appreciate the way she challenged my ideas during our meetings and helped me improve my work. Thanks to Alison Duncan Kerr who served as my secondary supervisor early in my PhD. Her invaluable comments and feedback helped shape what I wanted to say about emotions and microaggressions. The feedback from these three individuals has been essential to shape most of these ideas in this thesis, and I am deeply grateful for their help, advice, and encouragement.

Thanks to my close friends and partner who offered comments, helped me proofread this thesis, cheered me on and always supported me in the best ways possible: Tessa Carver, Leah Liora Cohen, Fiona Ramage, Petronella Randell, and Luca Stroppa. I love you all so much and I am so lucky to have you in my life.

Thanks to everybody who has discussed ideas from this thesis with me, whether it was at conferences, seminars, the pub, at lunch in Arché or over coffee breaks. I would like to thank everybody at the Arché Epistemology Current Themes seminar, the Arché Language and Mind seminar, the Arché Feminist Philosophy and Social Theory seminar, the CEPPA work-in-progress sessions, and the SEECRs Network for their helpful questions and comments on the work I've presented there throughout the years. A special thanks to those who have taken the time to give me comments on the material that made its way in this thesis: Christina Friedlaender, Sandy Goldberg, Alex Lloyd, Joe Millum, Lilith Newton, Philipp Schönegger, Ravi Thakral, and Emilia Wilson. Thanks to my annual reviewers Lisa Jones, Adrian Haddock, and Simon Hope for their advice and reassurance on my progress. A special thanks to Regina Rini who first introduced me to the topic of microaggressions and without whom this project would never have existed, and to Kristin Andrews who encouraged me to do a PhD. Thanks to various staff

members at SASP who have mentored or helped me throughout my PhD: Jessica Brown, Derek Ball, Alexander Douglas, Philip Ebert, Claire Field, Jade Fletcher, Katherine Hawley, Theron Pummer, Greg Restall, Gillian Russell, and Benjamin Sachs-Cobbe.

Thanks to all my philosophy friends whom I have not mentioned yet and who made this PhD so much more enjoyable: Nick Allen, Lisa Bastian, Katharina Bernhard, Robert Brown, Gwendoline Choi, Matt Clark, Jimena Clavel Vazquez, Rasa Davidavičiūtė, Jann Paul Engler, Camille Fouché, Enrico Galvagni, Eline Gerristen, Jean Gové, Nick Kuespert, Ethan Landes, David MacDonald, Christopher Masterman, Sophie Nagler, Quân Nguyen, Johannes Nickl, Thomas Randriamahazaka, Francisca Silva, Deryn Thomas, Colin Troesken, and Xintong Wei. Thanks to my friends outside of philosophy who made my time in St Andrews unforgettable: Avery Hawkins, Anushrut Ramakrishnan Agrwaal, and Natalie Smith.

Thanks to the St Andrews and Stirling Philosophy Graduate Programme, the Stirling Philosophy Department, and to the St Andrews Philosophy Department. I have benefitted tremendously of being in such stimulating environment, and this has helped shape my research for the best. Thanks to the St Andrews and Stirling Minorities and Philosophy Chapter, for providing a supporting community and improving my experience as a student. Thanks to the administrative staff from the School of Philosophy, Anthropology and Film Studies who have been extremely helpful during my studies. Thanks to Amy, Jan, and Liz from Estates for the lovely chats we had in the Arché Building. Thanks to Kate Leavy and Michelle Petrie from the Disability Team for their support throughout my studies, especially during the pandemic.

Finally, and most importantly, thanks to my parents Bernard and Nathalie, and my sister Clothilde for encouraging me to pursue my dreams and supporting me every step of the way. Without you, none of this would have been possible. Merci pour tout.

Funding

This work was supported by an AD Links Foundation Postgraduate Research Scholarship from the University of St Andrews and a Fondation Zdenek et Michaela Bakala Scholarship.

Introduction

Affective experience is a core part of how we experience the world, whether it is by being afraid when chased by a big dog, being sad after losing one's parent, or being in pain after breaking one's leg. These mundane affective experiences appear to highlight important evaluative features: our fear highlights the dangerousness of the dog, the sadness draws attention to the loss incurred, and the pain indicates that there is something wrong with the body. In recent years, attention has been paid to the interesting array of emotional display towards injustice and oppression. Pain and rage were expressed by many participants in the viral #MeToo movement in 2017 with regards to the ubiquity of sexual violence (Page and Arcy 2020) and protestors of the Black Lives Matter movement displayed rage at racial injustice, especially after the murder of George Floyd in May 2020 (Cherry 2021).

After working on the topic of microaggressions in 2017, I started to get interested in the role of affective experience in knowledge acquisition. Microaggressions refer to "brief, everyday exchanges that send denigrating messages to certain individuals because of their group membership" (Sue 2010, xvi). Microaggressions are particularly pernicious because they are usually highlighted through the lens of an affective experience, like anger or a gut feeling that something is wrong. However, such affective responses are often dismissed due to being too subjective and thus incapable of proving that the phenomenon is occurring. For instance, one might be told that they should learn to take a compliment when they try to confront the person who made a subtle sexist remark. Some scholars even deny that microaggressions exist in the first place, and rather argue that people are becoming too sensitive (Lilienfeld 2017). The debate surrounding microaggressions and other subtle forms of oppression magnified issues around affective states and their contribution to knowledge acquisition. Indeed, the lack of consideration for the role affective experiences could play impacts real people negatively and can even perpetuate injustice. If affective experiences were recognised as epistemically valuable, and people using them were flagged as good informants, more instances of injustice could be recognised.

However, before being able to accept affective experience and its contribution to knowledge, one needs to explain why affective experience should count as an appropriate source of knowledge. Knowledge, throughout this thesis, will be defined according to its functional role in our societies. According to Craig (1990), the concept of knowledge is used to flag approved sources of information, whether those sources are other people, the senses or other means, like reflection. Whether knowledge is justified true belief, the result of a reliable cognitive process, or defined in other ways, my aim is to question what should count as a good source of knowledge or justification, and if the affective experience of human agents could enter this category.

Whether affective experience should be included in our sources of knowledge is an old debate in Western philosophy. In the Antiquity, Plato saw emotions as something that gave willpower but needed to be guided by reason. Emotions were thus important to knowledge, even though they needed to be reined in to be useful (Jaggar 1989). In the Middle Ages, the status quo was that women were more susceptible to affective experiences which compromised their mental capacities. This made them less likely to be reliable testifiers (Van Dyke 2018). In the seventeenth century, the empirical tradition saw emotion as in opposition with reason, and affective experience was excluded as a source (Jaggar 1989). Following the empirical turn contemporary analytic epistemology continues to be wary of affective experience, which does not currently count as a source of knowledge. Indeed, there are four standard basic sources of knowledge or justification: perception, memory, consciousness and reason (Audi 2002). Contrastingly, the literature produced by late modern and contemporary feminist thinkers, like Myisha Cherry (2021), Patricia Hill Collins (1986, 2000), bell hooks (1996), Alison Jaggar (1989), Audre Lorde (1984), or Virginia Woolf (1929) explains the importance of emotions, in particular anger, to recognise injustice and oppression. They highlight how the white colonial patriarchy has pushed women and people of colour to repress their emotions and has used their purported emotional nature as a tool to discredit them. The current landscape about the contributions of affective experience to knowledge is thus divided:

on the one hand, contemporary analytic philosophy[1], constituted mostly of philosophers from dominant backgrounds, tends to ignore it, whilst the scholarship produced by members of marginalised groups, in particular Black[2] women, points out its crucial role.

As explained above, affective experience is not part of the approved sources of knowledge in contemporary analytic epistemology. Indeed, it is intuitively too subjective, unreliable, and often hinders knowledge acquisition, hence why it shouldn't be included. This leads to a tension: on the one hand there are many cases in the real world where agents appear to obtain information through their affective experience, which favours including affective experience in our sources of knowledge; on the other hand, epistemologists have good theoretical reasons to exclude affective experience from the appropriate ways of obtaining knowledge.

In this thesis, I argue that we should resist excluding affective experience as a source of knowledge, by exploring the problem from the perspective of contemporary analytic epistemology and applied ethics, whilst keeping in mind the work of feminist scholars. However, I want to preface this thesis by pointing out that the literature covering these aspects is wide-ranging. To start, the literature in epistemology on what counts as knowledge, what justification is, or how people obtain knowledge is extremely vast, having expanded tremendously post Gettier. When it comes to emotions and affective experience, a similar observation can be made. The literature on emotions contains many theories, often opposed, and so does the literature on pain and pleasure, and gut

[1] Contemporary analytic philosophy is the main tradition present in the Anglo-Saxon academic world and is the dominant tradition within what are considered leading journals in philosophy. This influences greatly what is considered to be relevant scholarship, which new works are published in places that reach out a lot of members in the profession, and which authors gain traction with their ideas and influence other scholars.

[2] I decided to capitalise the "B" in Black women and Black people through the whole thesis, as it represents a group of people linked by a common experience within our society, rather than being just a descriptive adjective.

feelings. Finally, the intersection of these two topics has also gained traction in recent years. Whether it is the literature on emotions like curiosity and how they help guide inquiry (Brun, Doguoglu, Kuenzle 2008; Candiotto 2019) or what emotional cognition is (Colombetti 2014; Thagard 2006), or the role of intuitions in moral knowledge (see for instance Audi 2013; Bergqvist and Cowan 2018; Huemer 2005; Lyons 2018; McGrath 2004; Setiya 2012; Sinnott-Armstrong 2008). In addition, the topic of the rationality of emotions has led to many interesting discussions, on whether emotions can be assessed along rational dimensions or on whether emotions represent the world accurately (see for example Ben-Ze'ev 2000; Brady 2013; D'Arms and Jacobson 2000; Deonna and Teroni 2012; Döring 2007; Furtak 2018; Pelser 2014; Prinz 2006; Roberts 2003; Solomon 2004; de Sousa 1987)[3]. There are also many articles or books written about specific emotions like anger, indignation, sadness, love or blame. Unfortunately, my engagement with these literatures will necessarily be limited, given how wide ranging each of them is. However, I do engage with the most relevant pieces throughout this thesis. My project is primarily in epistemology and applied epistemology, hence my focus on these literatures. Moreover, I aim to provide a theory and a discussion which would work for most accounts of affective experience and justification. This means that readers should feel free to incorporate their more fine-grained views of justification and affective experience in my argument, and this wouldn't work if I fleshed out the details from the start. Some readers will, of course, want more details, but I hope that approaching the topic in this manner will enable us to see the big picture and make relevant links rather than getting lost in a myopic focus and peripheral debates.

The thesis is divided in two parts. In the first part, made of chapters 1, 2 and 3, I focus on the epistemic aspects of the problem and explain why affective experience should be included as an appropriate source of knowledge. In the second part, which includes

[3] I explore in more details the current literature about the nature of emotions and how they provide direct justification for knowledge in chapters 1 and 2.

chapter 4, 5 and 6, I focus on the ethical aspects of the current exclusion of affective experience and the practical problems it creates for real agents.

I begin in **Chapter 1** by introducing what I take affective experience to be and the main arguments against the epistemic value of affective experience. Affective experience faces three major issues: first, it seems to provide us with false information and have a lower reliability than other sources; second, it doesn't appear to be sensitive to defeaters; finally, affective experience leaves open the possibility of conflict without contradiction, a problem not faced by the other sources. These problems motivate the exclusion of affective experience from our sources of knowledge. Nonetheless, I argue that we shouldn't be too quick to exclude it, as there are applied cases which turn the tables in favour of affective experience. I finish the chapter by introducing two applied cases, microaggressions and the diagnosis of endometriosis, where the exclusion of affective experience from our ways of knowing has negative consequences. This leaves us with a problem: either we follow our theoretical intuitions and exclude affective experience, but risk losing a valuable way of knowing, or we follow practical concerns and include it, but then risk getting away from the truth.

In **Chapter 2**, I propose to solve this problem, by providing a conception of affective experience as a source of knowledge which avoids the three issues presented in Chapter 1. I first discuss two currently available theories, Sabine Döring's perceptual model (2007, 2009) and Michael Brady's view of emotions as focusing attention (2013). I show that both fail to give a full picture of the direct justificatory role of affective experience. Given these shortcomings, I introduce my own model, the Reflective Model of affective experience. On my account, certain types of affective experience can be sources of knowledge, because rational agents would stop using affective experience to obtain information were it to fail to track reasons. Within this conception, affective experience is regulated through four calibration mechanisms, which help rational agents know when to stop trusting it as a source of knowledge. I discuss how my model fits within broader discussions in epistemology like the internalism/externalism debate and address objections.

After explaining this main theoretical insight, I propose in **Chapter 3** to dive back into the applied case of microaggressions, which was introduced in the first chapter. This chapter illustrates some of the mechanisms of the Reflective Model of affective experience and how agents use their emotions to obtain knowledge. I present one core unresolved issue in the debate around microaggressions: how does one determine if a microaggression has occurred? This question, central to the topic of microaggressions, is difficult to answer, as it requires us to explain both the nature of microaggressions and the epistemic access that agents have to them. In the literature, two types of definitions stand out: an experiential one, according to which a microaggression has occurred when the target of the exchange identifies it as such and a structural one, according to which a microaggression has occurred when a link between the act and an objectively existing form of structural oppression can be established. Due to the shortcomings of those definitions in giving both a satisfying ontology and epistemology of microaggressions, I offer a hybrid definition, the Emotional Account, which combines the most promising features of those two existing approaches. According to the Emotional Account, microaggressions occur if the target of the exchange could have had a warranted emotion towards the structural oppression present in that exchange. Moreover, the Emotional Account explains the link between the perspective of individual targets of microaggressions and group knowledge claims about structural oppression. Ultimately, the goal of the discussion is to show that the Emotional Account addresses the concerns of people who think that microaggression is just a concept used by easily offended people and should therefore be abandoned.

After discussing this application of the theory and wrapping up the discussion around the epistemic reasons to include affective experience in our sources of knowledge, I transition to the second part of the thesis. In this part, my concerns are mostly moral and practical, as I look at the consequences of dismissing the knowledge obtained through affective experience. To start us in this direction, I introduce in **Chapter 4** a new expansion to the concept of epistemic injustice, source based epistemic injustice (SBEI). SBEI refers to the unfairness towards knowers who use epistemic methods based on sources of knowledge currently unapproved by the dominant epistemological

framework. I focus my analysis on the downgrading of knowledge obtained through affective experience, like emotions or pain, to illustrate the phenomena. I explain why the insistence on intersubjectively verifiable justification and objectivity as an epistemological methodology leads to the downgrading of affective experience. I argue that this downgrading is unfair and disproportionately affects members of marginalised groups like women, people of colour and people with lesser-known illnesses. I show how this novel theoretical tool offers a new insight into the origins of some forms of hermeneutical injustice and enables us to explain some cases of epistemic injustice which could not be fully explained with the existing theories. I finish by sketching avenues that our current epistemological practices could take if we wish to reduce instances of SBEI.

Having taken stock of how SBEI operates and how we should conceptualise it, I put forward another practical issue caused by the exclusion of affective experience in **Chapter 5**. In this chapter, I investigate the problems arising when patients testify based on their phenomenal bodily experience, in particular their pain, and cannot provide objective evidence for their illness. I cast doubt over the benefits of a shift towards greater patient activation, the measure of a person's skills, confidence and knowledge to manage their own health. I explain two types of labour - hermeneutic labour and epistemic labour- that patients have to engage in to get appropriate care. I argue that the healthcare system requires epistemic and hermeneutic labour from various patients, not just to improve some already good health outcomes, but rather to survive and avoid medical neglect and dismissal, especially in the case of patients with lesser-known illnesses. Using Berenstain's three core features of epistemic exploitation, I show how demanding greater patient activation might lead to the epistemic exploitation of patients. Given that these exploitative epistemic mechanisms are not being discussed within healthcare, I argue that there is a risk that championing patient activation might lead to patient exploitation.

Having explored dynamics about what counts as good testimony and evidence in the context of healthcare, I propose in **Chapter 6** to shift contexts and look at the testimony

about everyday instances of racism and why the burden of proof for beliefs about racism falls on racialised people. I identify the phenomenon of dialectical white scepticism, a subtype of white scepticism where the resistance towards testimony about racial oppression is camouflaged as being motivated by purely epistemic concerns, rather than white privilege. I draw a parallel between the case of scepticism about common knowledge and dialectical white scepticism based on their use of an initially plausible epistemic closure principle. I detail how it puts unsatisfiable demands for defence of assertions about racism on racially oppressed people, by building a structural parallel with scepticism about common knowledge. I conclude by highlighting how this leads to dialectical white scepticism not being recognised as a form of white scepticism, which contributes to the epistemic burden carried by racially oppressed people.

What emerges from this multifaceted discussion is an argument in favour of including affective experience in our sources of knowledge. Without it, the contribution of agents from marginalised groups are particularly at risk of being excluded from our common epistemic enterprise. It is only by including affective experience that we can move towards bringing justice, given the important role that it plays in recognising the oppressive features of our society. By building the argument on the assumptions of contemporary analytic philosophy, I hope to show that the exclusion is not warranted, and that affective experience should be accepted within mainstream and dominant epistemology. Following this, maintaining that affective experience always subverts knowledge would only perpetuate unfair dominant standards that benefit privileged groups. Finally, including affective experience in our appropriate ways of obtaining knowledge enables us to build a bridge between the contemporary analytic tradition and the insight obtained from feminist epistemology, in particular standpoint theory. Through this reconciliation, the voices from the margin can finally take a centre place in our mainstream epistemological discussions.

Chapter One:
Against Affective Experience

So why should I care about a bad reputation anyway?

Joan Jett

Introduction

Affective experience, which refers to a variety of embodied mental states like emotions, pains/pleasures or gut feelings, occupies a peculiar place within epistemological discussions. Indeed, there has been a long debate about whether affective experience should count as a source of knowledge, and whether it hinders or facilitates the acquisition and maintenance of knowledge.

The value of emotions for knowledge has been widely discussed with regards to the role emotions play in our knowledge acquisition, by motivating inquiry or on the contrary, by hindering it due to emotional biases or implicit prejudice. These discussions focused on defining a category of emotions, namely epistemic emotions, such as curiosity, intellectual courage or the love of truth, characterised by having epistemic values as their objects, like acquiring truth and justification (Candiotto 2019; Morton 2009; Slote 2014). However, little has been said about how emotions themselves can provide us with knowledge or have epistemic value in themselves. Various authors mention that emotions can justify some evaluative beliefs (Döring 2007; Johnston 2001; Pelser 2014), but few explain in detail how emotions can provide immediate justification rather than just help direct our attention (Harrison 2021). Moreover, the contemporary analytic tradition tends to exclude affective experience as a source of knowledge. Whilst perception, introspection, memory and reasoning are mentioned as ways of obtaining knowledge and justifying knowledge, no introductive textbooks in epistemology mention affective experience and the role it plays in knowledge acquisition.

In contrast, the feminist and afro-feminist literatures are rich with examples illustrating the importance of affective experience. As Audre Lorde explains, “our real power comes from the personal; our real insights about living come from that deep knowledge within

us that arises from our feelings" (Lorde and Tate 1983, 106). Alison Jaggar highlighted that our emotional responses "stimulate us to new insights" (Jaggar 1989 170). Recently, Myisha Cherry dedicated a book to the role of rage and its importance in our lives, in particular to obtain racial justice (Cherry 2021). This chapter and the following aims to reconcile these important insights with classical theories of knowledge from contemporary analytic philosophy.

This chapter focuses on the problems that affective experience faces. It is divided into three sections. In section 1, I start by giving an overview of what affective experience is. In section 2, I present the reasons why affective experience has not been considered to have much epistemic value in the literature. In section 3, I challenge these reasons by briefly introducing two applied cases, pelvic pain in endometriosis and microaggressions. I use these to intuitively motivate the appeal of affective experience as a source of justification and the need for a theory on the role of affective experience in knowledge.

1. Affective Experience

I use the term 'affective experience' to refer to a cluster of embodied mental states, which include emotions, gut feelings, or pains/pleasures. Those states are individualistic, subjective, embodied and exhibit a particular phenomenology. Indeed, emotions, gut feelings and pains/pleasures are more than simply mental states, they are accompanied by an internal sensation, where the agent feels a change in their body. In this chapter, and throughout the rest of the dissertation, I will remain agnostic as to how to define affective experiences more precisely than this. The analysis to come is meant to appeal to a broad range of theorists and uses an intuitive conception of affective experience which most people, philosophers and laypeople alike, would accept. For this reason, I take the intuitive defining features of affective experience to be the following:

1) *Subjectivity*- Affective experiences are subjective mental states.

2) *Phenomenology*- Affective experiences are embodied and have a felt character.

3) *Valence*- Affective experiences are valenced, i.e., they can be positive or negative.

Moreover, the repetition of those states is an important aspect of such an experience, as agents learn to refine their analysis of what those states refer to. Some of these states are reactive states inasmuch as they react to changes in something other than themselves. Some affective experiences are states which are triggered by an external input and are 'about' something. In this sense, they exhibit intentionality. For instance, being afraid of a tiger is an affective state which exhibits intentionality, as the fear is directed towards an object, here the tiger (Scarantino and de Sousa 2018). On the contrary, an affective state like feeling generally depressed doesn't exhibit intentionality. In some cases like depression, it is hard to point to the object of the depression and what causes it. For simplicity, I will focus on affective experiences that arise from stimulus from the environment, the environment being understood broadly as anything that isn't the affective experience itself; this includes thoughts, memories or the physical world, to give a few examples. The reason for this choice is that my focus is on how affective experience can give us information about the world, and how we can learn about the world through our affective experience. Given this focus, it simplifies the argument to focus only on affective experiences with intentionality.[4]

In what follows, I characterise the three main types of affective experience, emotions, pain/pleasures, and gut feelings and discuss their relation to standards of rationality.

1.1 Emotions

There are many theories in Philosophy explaining what emotions are (Deonna and Teroni 2012; Scarantino and de Sousa 2018). These theories can be divided in three main traditions: emotions as feelings, emotions as evaluations and emotions as motivations. Within each of these traditions, the focus lies on one component of emotions. I now briefly introduce the main theories in each of these traditions.

[4] I do not consider that affective experiences with intentionality are the only affective states which are epistemically valuable. Indeed, some affective moods might be useful for knowledge, but given their lack of intentionality, it might be hard to explain how they provide direct and immediate justification.

According to the feeling tradition, emotions are conscious subjective experiences with a phenomenal character (Scarantino 2016). Within the feeling tradition, the James-Lange theory places emphasis on the bodily dimension of emotions. Here, an emotion is defined as the agent's feeling of bodily modifications, reacting to changes (James 1890; Lange 1922). The motivational tradition, on the other hand, defines emotions as motivational states or behavioural patterns of a certain type (Scarantino 2016). One theory of this type, the attitudinal theory of emotions, identifies emotions as feelings of action readiness (Deonna and Teroni 2012). Scarantino argues that emotions should be understood as central motives states or behavioural programs (2014). This echoes older theories by Fridja (1986) on emotions as action tendencies. Finally, the evaluative tradition holds that emotions are individuated by the evaluations that they involve, whether they are themselves evaluations, or because they are caused by evaluations of various types (Scarantino 2016). The evaluative tradition includes two main categories of theories: emotions as judgements and emotions as perceptions. In the first category, emotions are defined as judgement that the formal object of an emotion is instantiated. For instance, to be fearful of the tiger is to judge that the tiger is dangerous. Nussbaum (2004) and Solomon (1993) offer theories along these lines, inspired by the Stoics. In the second category, the perceptual model takes emotions to be perceptions of values (Deonna and Teroni 2012). These can either be direct perceptions of values (Döring 2009; de Sousa 1987; Tappolet 2000), where emotions are a form of affective perception which, like visual perception, presents the properties of the world to the agent. This is also known as the perceptual model of emotions. I will get back to the perceptual model of emotions in the next chapter, as these theories are often associated with attempts to explain how emotions provide us with evaluative knowledge. Otherwise, emotions can be considered as indirect perceptions of value (Prinz 2004), where the awareness of relevant changes in the body is given an evaluative judgement. This type of theory is sometimes considered to be building on the feeling tradition, whilst incorporating evaluative components.

Due to their intentional nature, emotions can be assessed on two rational dimensions that resemble two important ones in epistemology: truth and justification (within the

emotion literature, those are called *fittingness* and *warrant*). For instance, emotions can be assessed on the basis of whether the emotion targets the right object. This is called the fittingness of the emotion (Deonna and Teroni 2012; Scarantino and de Sousa, 2018; de Sousa 1987). Take the example of someone who is afraid of a very realistic lion animatronic. In this case, the emotion is not fitting, because the lion animatronic is not dangerous. However, that does not mean that the person has no reason to be afraid of it. Indeed, if the animatronic is realistic, it could have easily been mistaken for a real tiger, and the person has reasons to be afraid of it. In this sense, their emotion is 'justified', or warranted if one is using the proper terminology from emotion theorists.

1.2 Pains and Pleasures

Pains and pleasures refer to the category of sensory affective experiences. Following Bain and Brady, we can define pleasure as positive sensory pleasures and pain as physical pains, which are experiences that the agent has in parts of their body (2014).

Can pain and pleasure be assessed along standards of rationality? Insofar as pain and pleasure are reactions to positive or negative inputs, one can say that they could be fitting or not. Indeed, experiencing pain after a negative input (like putting one's hand on a hot stove or breaking one's leg) seems fitting and warranted, whereas experiencing pleasure after such negative input appears to be unfitting and unwarranted. However, there are cases, like the ones of some sexual masochists or people who like very spicy food, where this simple analysis is more complicated (Bain and Brady 2014). Indeed, some people do experience pleasure after receiving a negative input, and this pleasure is pleasant in virtue of the bodily pain experienced. While a thorough discussion of this issue would be outside the scope of this chapter, we can nonetheless analyse those cases by explaining that the agent, in those specific contexts of sexual cooperation or culinary adventure, experiences both pain and pleasure, by separating the bodily sensation from the affective experience, and thus separating the two objects provoking the reaction (Feldman 2004; Mason 2007). On the one hand, the bodily input of the harm done to one's body is provoking pain, while the painfulness of the sensation

produces the pleasure. Thus, the rational standards to assess those cases can still be applied to pleasures and pains.

1.3 Gut Feelings

Gut feelings can be understood as bodily states which offer a valenced representation (Prinz 2004) or as mental processes which are unconscious and embodied (Gigerenzer 2008). Gigerenzer explains that gut feelings are primitive feelings which help us make snap decisions (2008).

Can we say that gut feelings are fitting or warranted? One could worry that due to their spontaneous and instinctive nature, it might be difficult to say that gut feelings are fitting or warranted. When it comes to fittingness, it does seem that we can evaluate them. Indeed, gut feelings are directed towards an object. Thus, if the gut feeling does not correspond to the object, we can say that it is unfitting. For instance, someone having a gut feeling that they are about to die, while being in a perfectly safe situation does not have a fitting gut feeling. However, making sense of the warrant dimension of gut feelings might prove to be more complicated. Indeed, it seems plausible that someone could have a gut feeling about an object, without possessing reasons to justify or explain such a reaction. Nonetheless, it is possible that agents might be able to reflect on their gut feelings and try to modify their future spontaneous reactions. For instance, stereotypes and implicit biases are a good example of gut reactions which can be reflected on by the agent exhibiting them, or at the very least which can be shifted or reprogrammed through new associations (Huebner 2016). Various studies showed that attitudes can be shifted when changing the pattern of stimuli one is exposed too. For instance, Dasgupta and Greenwald (2001) showed that participants who are presented with biographies and images of admired African Americans and despised white Americans had a decreased implicit racial bias, an effect which persisted 24h after the exposure to the stimuli. Moreover, there exists a variety of strategies to help regulate implicit biases. Those strategies are often divided into two categories: change-based strategies, which purport to modify the associations underlying an agent's biases; control-based strategies, which aim at suppressing the effect that the associations have

on the agent's behaviour (Brownstein 2019). Thus, it seems that is possible to influence one's gut feelings to make them respond to situations in a different, more warranted way.

2. Against the Justificatory Power of Affective Experience

In the last section, I provided a brief general overview of the nature of affective experience and discussed how its main types are assessable along various standards of rationality. In this section, I propose to synthesise the main reasons against accepting affective experience as a source of justification or as contributing to knowledge directly. As explained in the introduction, the epistemic value of affective experience has been mostly restricted in contemporary analytic epistemology to the benefits of epistemic emotions. Indeed, there seems to be a consensus against its use as a source of knowledge or justification more generally. I will now highlight the main worries motivating this consensus in favour of the exclusion of affective experience from our ways of knowing. In general, sources of knowledge or justification are taken to generate beliefs constituting knowledge for an agent (Audi 2002). Given that sources are fallible, the beliefs delivered by the source are subject to commonly accepted standards of rationality. For instance, it should be sensitive to defeaters and generate true beliefs in most cases. This means that an agent cannot maintain contradicting beliefs. For instance, the beliefs generated by the source should not be maintained if they conflict with other beliefs that the agent holds, unless the agent would rather give up on some of their old beliefs to avoid the conflict with the new ones.

In what follows, I highlight how affective experience appears to fail when assessed based on these standards of rationality. I start with the problem of false information and low reliability, then discuss whether affective experience is sensitive to defeaters and how its lack of intersubjectivity undermines it as a source of knowledge. Finally, I introduce the problem of conflict without contradiction.

2.1 False Information and Low Reliability

Affective experience can provide us with false information. It is a fact that I will not try to contest, and which is not unique to affective experience, see for instance perceptual illusions. Indeed, many times, people form false beliefs following an affective experience. I now present three paradigmatic cases where affective experience provides us with false information and why it is concerning.

The first common case is the one of phobias:

> **Spider phobia**: Dorian sees a tiny spider, of the type *oonops domesticus*, in the upper corner of his attic. The spider is about 2 mm in size, but nonetheless, Dorian reacts with fear and runs to grab a broom to kill it because he believes that the spider is dangerous.

In the **spider phobia** scenario, Dorian's fear led him to believe something false. Indeed, the spider is not dangerous at all. Due to his phobia, Dorian misinterprets the situation and forms a false belief. If Dorian had not been irrationally afraid of spiders, he would not have formed the belief that the spider is dangerous. Thus, in the **spider phobia** case, it seems that Dorian's fear concretely plays against his ability to competently assess the parameters of the situation such as the species, size and location of the spider, thus providing him with the false information that the spider is dangerous. Therefore, affective experience can be detrimental to knowledge acquisition, as shown by the **spider phobia** case.

The second case is one of affective experience being elicited without the relevant target object which leads the agent to form mistaken beliefs (inspired by Correll et al. 2002):

> **Shooter bias**: Tomer participates in an experimental study using a first-person shooter task. In the task, he is asked to shoot or not shoot various human targets who appear holding either a neutral object (wallet, phone, soda can) or a gun. The goal is for him to only shoot at the individuals holding a gun. During the task, Tomer mistakenly shot unarmed Black targets more often than unarmed white targets, and erroneously did not shoot white targets with a gun more often than Black targets with a gun. Due to his unconscious racial bias, Tomer is more likely to be afraid of a Black target and form the mistaken belief that the target is dangerous, even when the target does not hold a gun.

In the **shooter bias** case, Tomer's fear of the unarmed Black targets is elicited even though the targets are not holding a gun. In this case, it seems that an emotion, here fear, can be elicited due to the presence of implicit bias and without the object (here the unarmed Black targets) exhibiting the relevant property (being dangerous), leading to the formation of mistaken beliefs (I believe that this target is dangerous). Thus, emotions can be externally elicited, without being directly linked to the relevant object and lead to the formation of mistaken beliefs.

The final case is one of confirmation bias:

> **Jealous partner:** Kathryn and Annette are a couple. Annette is a faithful partner, but Kathryn has a gut feeling that Annette is cheating on her. She thus develops jealousy towards Annette's friends. Due to her jealous emotional state, Kathryn looks for any clues indicating that Annette might be cheating on her with one of her friends. Moreover, she also overlooks any clues that might go against her belief that Annette is unfaithful. While she has no evidence that Annette is unfaithful, she still maintains this belief due to her emotional state.

In the **jealous partner** case, Kathryn obtains false information through a confirmation bias that emerges and is maintained by her affective experience. Due to her emotional state, she will only look for clues that prove her hypothesis and will disregard clues that go against it. Moreover, her emotional state distorts her perception of various clues and causes her to maintain the false belief that Annette is cheating on her.

Through these cases, we can see that affective experience can be elicited without the relevant object present, motivates and sustains confirmation bias, and in general can give us false information. *Prima facie*, this contrasts with other commonly accepted sources like perception or reasoning. Whilst perception is not infallible, it appears to have a greater success rate compared to affective experience. Indeed, affective experience seems particularly prone to distortions. I will get back to this point and explain why in the next two sections. Overall, it seems that we should not accept affective experience as being epistemologically valuable. Indeed, doing so would be epistemologically too risky, and we are better off using other sources to obtain information and justify our beliefs.

2.2 Sensitivity to Defeaters

The picture painted in the last section gives clear reasons to exclude affective experience from our ways of knowing. If we often acquire false information due to our affective experience, we could argue that we should not use affective experience as a way of acquiring information. However, this reasoning, if applied to other commonly accepted sources of information, sounds quite pessimistic. For instance, there are plenty of cases where our perception can be mistaken, like optical illusions. However, we don't abandon perception as a source of information for this reason. Thus, what is the reason for accepting something like perception as a source of knowledge or justification, but rejecting affective experience?

The difference between the two might reside in how well they respond to defeaters and evidence. According to Pollock and Cruz, reasons in epistemology are defeasible, "in the sense that, while they can justify us in believing their conclusions, that justification can be "defeated" by acquiring further relevant information" (1999, 10) For instance, learning that the lion is an animatronic instead of a real one is an information that constitutes a defeater for the belief that *the lion is dangerous*. Indeed, if one possesses the information that the lion is an animatronic, thus is innocuous, and forms a relevant belief on the basis of this newly acquired information, they cannot rationally maintain their belief that *the lion is dangerous*. While I do not wish to delve into the debates about rationality and what makes an agent epistemically rational, I think that most people will accept a minimal sense in which an agent is epistemically irrational if they hold a belief even when presented with an undercutting or rebutting defeater. A rebutting defeater is evidence which contradicts their belief, like learning that the lion is an animatronic. An undercutting defeater is new information which means that the evidence previously held no longer supports the belief. As Pollock and Cruz illustrate for an undercutting defeater, "I may justifiably believe that something is red because it looks red to me. I have a justified belief here. But if I then acquire a further belief, viz., that the lighting is peculiar in certain ways, that may justify me in believing a defeater for the original defeasible reason and may have the result that I am no longer justified in believing that the object in question is red. If I am rational, I will then reject that belief" (1999, 44). In

most cases, when one is presented with a defeater to the justification of their belief, they are no longer justified in believing, and to continue to believe would be irrational.

Being sensitive to defeaters seems to be an important part of how one can reliably acquire knowledge. The problem is that beliefs produced by affective experience seem to be more prone to being 'bad beliefs', beliefs which are not aiming for truth, and which are not representing the world as it is, because they are not sensitive to defeaters. Looking back at our cases, we can see that is the explanation as to why those cases give the agent false information, and most importantly, why agents are not able to change their belief even when presented with defeating evidence. For instance, in the **spider phobia** case, Dorian believes that the spider is dangerous due to his fear. Given that it is a phobia, it is highly likely that Dorian would not abandon the belief that the spider is dangerous, even when told by an arachnologist that the spider is innocuous. In the **shooter bias** case, Tomer probably already knows that targets holding a phone are not dangerous. However, due to his unconscious bias, he is not able to amend his belief that the target is dangerous when presented with visual evidence that the target is not dangerous. Finally, in the **jealous partner** case, Kathryn is insensitive to defeaters and truth because of her jealousy. Indeed, she is not responding to evidence that Annette is faithful. Thus, the problem is not as much that the agents obtain false information because of their affective state. The problem arises because the agents are more likely to maintain their problematic beliefs formed based on affective experience, even when presented with defeating evidence.

Which features make the process of acquiring information through affective experience less likely to be responsive to defeaters? To answer this question, I discuss a characteristic of affective experience which usually differentiates it from commonly accepted sources of information: its personal nature.

Affective experience is a subjective state. Indeed, emotions, gut feelings, pain and pleasures are all subjective mental and bodily states. For this reason, it is difficult to compare one's affective experiences with somebody else's. Usually, many people can react to the same stimulus in similar ways, but there is no way to assess whether

somebody else is going through the same exact affective experience as we are. Moreover, there are plenty of cases where two people presented with the same stimulus can have contradictory affective responses. For instance, imagine that Abigail is telling a joke to Basim and Constanza. Basim finds the joke extremely funny and is amused at it, while Constanza is bored by it and finds the joke unfunny. In this case, Basim and Constanza were exposed to the same stimulus, but formed two opposed beliefs about the joke: *the joke is funny* or *the joke is not funny*. If Basim and Constanza decide to compare their beliefs to see who is right about the joke, they will probably run into an issue. Indeed, it seems difficult to decide that either Basim or Constanza is wrong when their feelings are so subjective.

This leads us to the core tension that exists between affective experience and knowledge. On the one hand, we have a form of experience which is very personal. Gut feelings, pains, pleasures, or emotions are all subjective and private, and it is impossible to explicitly compare whether two affective experiences have the same intensity or phenomenal attributes. While it is possible to compare whether two agents have the same reasons or the same object for their experience, there is an important part of such an experience which will remain opaque to an external observer. On the other hand, there is a phenomenon, knowledge, where epistemic reasons should be transparent to multiple agents, and knowledge is objective. Indeed, within the contemporary analytic tradition, knowledge doesn't rely on any phenomenal features, and the reasons of various agents can be shared, transmitted and compared as they are perfectly impersonal and transferable.[5]

When examining the three cases presented in Section 2.1 and why the agents discussed in them are insensitive to defeaters, there is a noticeable influence from the phenomenal part of the affective experience on the unresponsiveness to defeaters. In the **spider phobia** case, Dorian continues to be afraid of the spider when told that the

[5] One should note that this conception of knowledge was greatly criticized by feminist epistemologists, especially standpoint theorists (Harding 2004).

spider is innocuous, because he has a negative phenomenal experience of the spider: sweating at the sight of it and his body being tensed. Knowing that the spider is innocuous does not reduce this physical component of the affective experience, thus leading him to maintain his belief that the spider is dangerous, because his body tells him so. Tomer, in the **shooter bias** case, has an unconscious racial bias against Black targets. Implicit biases have often been described as attitudes, which have a valence and an intensity, and are composed of cognitive, emotional and behavioural components (Frankish 2016). If Tomer's beliefs are influenced by some non-cognitive features like emotional or behavioural features, it could explain why it is difficult for him to stop the formation of false beliefs during an exercise which requires him to act fast. Finally, in the **jealous partner** case, Kathryn has an emotion, jealousy, which does not go away even when being presented with counterevidence. Kathryn, in this case, is experiencing a recalcitrant emotion, which is characterised by Brady as an emotion involving the "mobilisation of cognitive resources in the service of a question that has, by the subject's own lights, already been answered" (Brady 2009, 428). Kathryn's jealousy makes her construe Annette as unfaithful, because the phenomenal part of her emotion (the feeling) inclines her to search for reasons which justify her jealousy (Brady 2009). This process mobilises her attention and makes her less prone to respond to reasons that go against such a construal.

The phenomenal component of affective experience thus complicates how an agent responds to defeaters. This component is what renders affective experiences deeply personal and subjective. It is also the aspect of affective experiences which can interfere with an agent's response to the world. The personal nature of affective experiences is their defining feature. It is also the feature that can be used to justify its exclusion from reliable sources of knowledge or justification. Indeed, if an affective experience cannot be assessed from an external perspective nor is sensitive to defeaters then it becomes difficult to see how it can track truth and represent the world as it is. Thus, we should accept that affective experiences cannot be used to gain knowledge.

However, one might wonder: why is the phenomenal component of affective experience hindering knowledge acquisition, rather than something else? Couldn't we associate this failure to the agent being irrational instead of criticising affective experience? To answer this objection, let us now discuss a feature of affective experience which differentiates it from other sources of justification and knowledge: the possibility of conflict without contradiction. This will help us understand why an agent can be rational whilst still not responding to defeaters when some of their beliefs are formed through affective experience.

2.3 Conflict Without Contradiction

Let us consider a story. Anushrut believes that his cat Fabio is sleeping on the sofa. At the same time, Anushrut holds the belief that Fabio is in the kitchen eating his kibbles. It seems unintuitive that Anushrut can hold both the beliefs that *Fabio is on the sofa* and *Fabio is not on the sofa* at the same time if he is a rational agent. In this case, we would say that Anushrut has two contradictory beliefs, which leads to a rational conflict for the agent. In other words, he is faced with a conflict with contradiction. If Anushrut is rational, he should aim to resolve this conflict by abandoning one of these beliefs. If he doesn't, he will contradict himself.

This story holds when an agent has two conflicting judgements or beliefs. However, when there is a conflict between an affective experience and a judgement, the story is different. Indeed, there is a possibility that the agent could form a judgement that has a content contradictory to the content represented by an affective experience. For instance, when looking at a gorilla at the zoo, it does not seem contradictory for Tara to judge that she is both safe, because she is protected by the enclosure, while feeling fear at the gorilla and experiencing the situation as dangerous. If Tara were to hold both the judgement that she is safe and the judgement that she is unsafe, at the same time, she would be experiencing a rational conflict. However, for her to have one belief state that she is safe, while having an emotional experience that she is unsafe doesn't seem to lead to a contradiction, even though the two mental states are about the gorilla and conflict about how they represent the gorilla. As Döring explains, emotions leave open

the possibility of conflict without contradiction (Döring 2009), and this can be generalised to other affective states. This problem of affective experience contradicting beliefs in rational agents arises similarly in the literature about phantom limb pains. For instance, in the seventeen and eighteenth centuries, natural philosophers like Descartes, Locke and Porterfield were fascinated by the problem of people experiencing pain in amputated limbs (Skuse 2021, chapter 6). In these deceptive cases, the agent's power to correct deception seemed absent, which caused a problem for conceptions of rationality. Porterfield himself had a missing limb and explained that despite all his reasoning, he could not resist scratching his missing leg, even though he considered himself to be a rational being (Skuse 2021, 155-156). The problem of phantom limb pains echoes the one of emotional conflict without contradiction. Even though the agent believes that he does not have a leg, he still has an affective experience of having a leg. Here the representations conflict with each other, even though there is no contradiction. The agents in these cases consider themselves to be rational and would generally be considered rational, and we can describe these situations as cases of conflict without contradiction.

The possibility of conflict without contradiction explains why in some cases affective experience persists even when the agent is presented with defeating counterevidence. If rational agents don't experience conflict with contradiction for beliefs formed through affective experience, then affective experiences don't seem to fit well within a model of knowledge acquisition where agents need to be rational and sensitive to defeaters to acquire knowledge.

We have now reached a roadblock. If affective experiences cannot fit within a traditional conception of rationality and knowledge acquisition, then we should exclude it as a source of knowledge and justification. From a theoretical point of view, this seems to be the natural progression from what I have discussed so far. However, is that the right move when it comes to the way agents obtain knowledge in the real world? In the next section, I address this question by discussing two applied cases which turn the tables in favour of affective experience.

3. Cases in Favour of Affective Experience

Can we exclude affective experience as a source of knowledge or justification while avoiding undesirable consequences? In this last section, I argue that we shouldn't be so prompt to do so. I propose to examine two applied cases where affective experience is epistemically valuable. The first one is the case of pain for the diagnosis of endometriosis. The second one is the role of emotions in recognising microaggressions[6].

3.1 Pain and Endometriosis

Pain plays an important role in learning about our body and its states. If there is a negative sensation, this sensation is registered by the agent as a painful experience, and it often leads to the agent forming a related belief about a problem within the body. In many cases, this belief leads to a knowledge state, and the object of this knowledge is the state of the body. For instance, when one breaks a leg, the pain radiating through one's leg is a crucial information that something is wrong. More formally, S knows that something is wrong with their leg in cases where: 1) their leg hurts; 2) the belief has either been produced through a reliable process or is justified; and 3) that belief is true. In a sense, the agent knows that something is wrong with their leg through multiple ways. First, usually, when they feel pain, it correlates reliably with the fact that something is wrong with their body. For instance, a painful ankle correlates with a twisted ankle, a painful thigh to an enormous bruise after a fall, etc. Moreover, the pain gives the agent a justification for their belief that something is wrong with their leg. If asked about why they think something is wrong with their leg, they could reply, 'I know something is wrong with my leg, because it hurts'. There are counter cases to this general rule, like phantom limb pains, but in general pain is a good indicator that something is wrong with one's body.

[6] These two cases are discussed in more details in other chapters of this dissertation. For pain and endometriosis, see chapters 4 and 5; microaggressions and emotions are discussed in chapter 3.

In general, pain therefore seems to be an interesting way to know or to acquire important information. Indeed, it is one of the best ways for an agent to learn about their bodies, in particular about what might be wrong with them. However, one might object that pain is not enough to provide knowledge. It might be useful to direct one's attention to a problem and to push the agent to look for additional evidence, as pain could be insufficient to justify the belief. For instance, one might prefer to seek professional advice, like the one of a physician, or use technologies, like X-rays or thermometers, to corroborate one's belief. In short, to use sources of justification that are external to the agent and might be subject to intersubjective assessment.

Nonetheless, there are cases where gaining justification through intersubjective means seems difficult or not a straightforward way to justify one's belief. One of these cases is the one of endometriosis. Endometriosis is a painful gynaecological condition affecting 1 in 10 people with a uterus in which tissue similar to the lining of the uterus starts to grow in other places of the body. Every month, this tissue bleeds in the same way as the tissue in the uterus, but unlike the latter, it has no way to escape, causing a painful and/or heavy period. Moreover, it causes inflammation, pain and the formation of scar tissue, impacting the well-being of people who suffer from this condition. It is also one of the main causes of female infertility (Endometriosis UK 2022). Unfortunately, there are no easily available diagnostic tools for it, as the only way to pose a diagnosis is through an invasive surgery.

There is an important aspect about endometriosis that is worth discussing: it takes patients an average of 7.5 years to obtain a diagnosis, after seeing a doctor for the first time to discuss symptoms (Endometriosis UK 2022). This could be due to the absence of non-invasive diagnostic tools, the state of the medical system, systematic injustice, sexism, or the interests of researchers. Putting aside these concerns, there is one particularly interesting characteristic about endometriosis diagnoses: on average, patients consulting for pelvic pain experience a delay in the diagnosis of up to 3.4 years compared to patients who consult for infertility (Arruda et al. 2003). Why do patients with pelvic pain wait longer than infertile patients when the same diagnostic methods

are used for both groups? Endometriosis with pelvic pain symptoms offers an interesting problem: periods are considered to be painful (or at least inconvenient) and thus painful periods are not considered problematic by healthcare practitioners. Why do patients then seek a diagnosis if they are told that their pain is normal?

In these cases of pelvic pain, it seems that the first-person experience of the agent might prove to be epistemically valuable in a way that cannot be replaced by other ways of belief formation. Indeed, their pain assessment is the best way to justify their belief that something is wrong with their body, and the best way for them to obtain this information. It is only through their embodied experience that they can compare past experiences and new ones, which enables them to gain information. In cases of endometriosis with pelvic pain, the epistemic gain brought by using one's affective experience is immense. Indeed, it seems that endometriosis sufferers come to know that something is wrong with their body on average 7.5 years before healthcare professionals, which suggests that their way of knowing in this case might be the most efficient. This also explains why these patients continue to seek a diagnosis, as their pain is giving them reasons to search for an explanation about the changes in their bodies.

Pain has in general a good epistemological value, as it enables agents to gain information about their bodies. Furthermore, in some cases, pain is the best way to acquire justified beliefs about one's body, being even better than intersubjective justificatory methods, like technical medical examinations, as shown by the case of endometriosis. It is worth noting that I am not arguing that the intersubjective justificatory methods are not good, or that if healthcare practitioners could look inside the person's body in a simple way to assess the nature of the problem that it would not constitute a good way to obtain knowledge. The argument here is that endometriosis sufferers currently rely on their pain to know that something is wrong with them. This makes pain a good way to obtain knowledge, and at least in the real world, one that is very important. Indeed, if patients didn't trust the information gained through their pain, they wouldn't be diagnosed and could suffer disastrous health outcomes by being left untreated.

3.2 Microaggressions

The second case highlighting the epistemological value of affective experience is the one of microaggressions. Derald Wing Sue, one of the main contributors to the topic, defined microaggressions as "brief, everyday exchanges that send denigrating messages to certain individuals because of their group membership" (Sue 2010, xvi). These exchanges are often done without the intention from the aggressor to send a denigrating message nor to harm the other person. Even worse, some microaggressions are done with good intentions, just like when students of colour are complimented for being "so articulate" by their teachers (Sue et al. 2007, 276). In the current philosophical literature on the topic, microaggressions are often characterised through their functional role within a structure of oppression: microaggressions are recognised as such when one is able to link a particular act or situation to a pattern of oppression (Friedlaender 2018; McTernan 2018).

To recognise individual instances of microaggression, it appears that the affective experiences of an agent play an important role[7]. As Saba Fatima explains, their affective experiences enable the agent to link the particular act or situation to the pattern of oppression, as agents know what happened with their "brains and [their] hearts connecting seemingly disparate life experiences in a split second" (Fatima 2020, 166). The idea that emotions, especially anger at injustice, can pick out patterns of oppression was recently discussed by Myisha Cherry. As she explains, "emotions are often responses to a pattern of events over time" (Cherry 2021, 39). Moreover, she highlights that the frequency of experiences of anger or discomfort about injustice can help the agent understand the source of their discomfort (Cherry 2021, 42). This mechanism helps targets of injustice recognise it through their emotions, as their emotions react to the pattern of oppression, even if subtle. Contrastingly, the lack of emotions and personal experience can make one doubt the existence of injustice or microaggressions.

[7] See Chapter 3 of this dissertation, *The Emotional Account of Microaggressions*, for a new characterisation of microaggressions as a phenomenon identified by the emotions of the target.

These agents will need more convincing to believe in the existence of the pattern because they lack the adequate affective experiences.

Let us examine an example of this. It is taken from the personal experience of Saba Fatima, who is a professor of philosophy and an immigrant woman of colour. When she tells her colleagues about one of her students who is disrespectful towards her, they say to her that they "all get disgruntled students" (Fatima 2017, 153). Fatima thinks that this disrespect is due to the general disrespect people have towards immigrants. However, all her colleagues think that this student is just another annoying student, and that this has nothing to do with Fatima's status as an immigrant (Fatima 2017). In this case, the colleagues are lacking the affective experience that would enable them to link this individual incident to a wider pattern of oppression. Because of this, they deny that this particular situation is a microaggression, even if they agree with Fatima on the general definition and existence of microaggressions. Here, having an emotional response to structural injustice is one that becomes automatic in agents who are regularly microaggressed, like Fatima. by contrast, her colleagues lack the emotional access to the oppressive features of environment because they are not experiencing it as victims. This creates a gap in the knowledge they can acquire about microaggressions and leads them to doubt the experience of the victim[8].

The case study of microaggressions shows us that affective experiences, here emotions, play a central epistemic role in the recognition of individual instances of microaggressions, because they highlight patterns[9]. It shows that emotions can be directly epistemologically valuable in themselves and lacking them can hinder the acquisition of knowledge in some agents. Without the valuable pattern recognition brought by emotions, agents are worse off epistemically.

[8] I will get back to this case in Chapter 3.

[9] This idea will be defended further in Chapter 3.

4. Conclusion

Affective experience has a 'bad reputation' for sensible reasons: it gives us false information in plenty of cases, it isn't sensitive to defeaters, and there is the possibility of conflict without contradiction. Therefore, it may seem right to exclude it from our sources of knowledge. However, as we have seen in the last section, excluding affective experience from our ways of knowing can have unwanted consequences. Without it, patients with endometriosis go undiagnosed longer, and targets of microaggressions might fail to recognise oppression. We are now pulled in two different directions: either we follow our theoretical arguments that we should exclude affective experience, but risk excluding valuable ways of knowing, or we follow our practical concerns, include affective experience, but risk being misled. Choosing either option currently appears suboptimal.

However, perhaps there is a way to explain the epistemic contribution of affective experience whilst safeguarding it from the problems discussed in this chapter. We could solve the problem if we can find a way to explain when affective experience should count as a source of knowledge and when it should not. But to this end, we need a conception of how affective experience provides justification or knowledge.

Chapter Two:
The Reflective Model of Affective Experience

Our feelings are our most genuine paths to knowledge.[10]

Audre Lorde

Introduction

We left the previous chapter with a problem: we need a conception of how affective experience may provide justification or knowledge, otherwise, we cannot motivate why it should be included as a source of knowledge. Furthermore, we cannot explain why affective experience makes valuable epistemic contributions in the cases of endometriosis diagnosis or microaggressions without it. This chapter aims to discuss several options available, including a novel account, the Reflective Model of affective experience.

Various emotion theorists have highlighted the importance of emotions to gain evaluative knowledge (Cowan 2018; Döring 2007; Pelser 2014; Tappolet 2000). However, there are currently few discussions about whether—and if so, how—affective experience can provide direct justification or knowledge. Most of these theorists have perceptual models of emotions, given that the perceptual model can easily explain that evaluative knowledge is gained through a form of affective perception akin to visual perception.

This chapter is divided into four sections. In section 1, I introduce two current theories, Michael Brady's view of emotions as helping focus our attention, and Sabine Döring's perceptual model. I argue that both fail to give a full picture of the epistemic value of affective experience and downplay the direct justificatory role that affective experience plays. Given these shortcomings, in section 2, I introduce my own model, the Reflective Model of affective experience. I argue that the resistance against accepting emotions as

[10] (Lorde and Tate 1983, 106)

sources of justification disappears with a suitable reconceptualisation of the notion of rationality and how rational agents interact with reasons. If rationality includes reflective self-monitoring, then emotions that are reason-tracking can provide agents with justification for their evaluative beliefs. This model builds on the work by Brady and Döring by appealing to Karen Jones's framework of reflective rational agents. In section 3, I discuss the model and address objections. Finally, in section 4, I highlight how the model fits within various theories of justification from contemporary analytic epistemology.

1. Current Theories

I now discuss two of the most influential accounts of emotions which address how emotions can provide justification and knowledge. The first is Sabine Döring's perceptualist model; the second is Michael Brady's theory of emotions as promoting evaluative understanding[11]. After presenting both positions, I explain why both remain unsatisfactory or unconvincing.

1.1 Döring and the Perceptual Model

Let us start with Sabine Döring's perceptualist model. Döring's theory belongs to a wider attempt to analyse emotional experience in an analogous manner to perceptual experience, so as to make sense of the epistemological appeal of emotions and how they provide us with evaluative knowledge. According to the perceptual model of emotions, emotions involve a perceptual experience of an object as representing an evaluative property. In other words, the representational content of the emotion is an evaluative content (Brady 2013; Döring 2007). Given the analogy with perception, perceptual theorists of emotions such as Döring argue that this representational content

[11] I am purposefully not discussing Rick Furtak's view of emotions (2018). Indeed, he adopts a view which cannot be reconciliated with the idea that knowledge is transparent. I aim to propose a theory which would fit within contemporary theories in analytic epistemology, and Furtak's work isn't compatible with this aim.

is subjected to correctness conditions, just like perception. Indeed, an emotion can be correct or incorrect in the way it represents the world (Döring 2007). For instance, one's anger is correct when one is being lied to by a friend, as it accurately represents the situation as being unfair.

One appealing aspect of the perceptual model is how easy it makes the link between the emotion and the justification of the evaluative belief. Just like perceptual experience justifies perceptual beliefs, emotional experience justify evaluative beliefs, given the representational content of emotions. As Döring explains, "a judgement made rational by the representational content of the subject's emotion is not a matter of an inference from the occurrence of the emotion" (2009, 89). She explains that the belief does not require any inferential reasoning to justify the belief arising from the emotional experience. For instance, a tiger jumps in Julieta's direction, and she is terribly afraid of it. Here, her belief and judgement that the tiger is dangerous is directly justified by her fear, as Julieta takes her fear's content- *this tiger is dangerous*- at face value. To summarise this position, Julieta can rely on her fear to justify her belief that the tiger is dangerous, and she doesn't need to rely on any other inferences, like *this tiger has big claws*, *tigers are carnivorous and deadly animals*, etc., as her fear is a sufficient source of justification (Döring 2009).

According to the perceptual model, and in particular the version defended by Döring, the representational content of emotions, which is what makes them similar to perception, is what enables emotions to play a direct justificatory role. By building on the similarities with perceptual experience, Döring adds that the content of emotion and perception is often a correct representation of the world, even though it can be false, due to emotional experience and perception being fallible. Thus, agents often consider their emotions to generally provide an accurate representation of the world, even though agents might not consider a particular individual instance as being necessarily accurate (Döring 2009).

While the perceptual model is appealing due to the ease with which it explains how emotions provide justification to evaluative beliefs, it encounters various issues. Michael

Brady, who agrees with the perceptual model's core ideas that emotions tell us about evaluative features and are epistemically valuable, provides some cogent criticism of the perceptual model. His first main issue with the perceptual model is that it fails to explain how emotions can provide reasons for evaluative judgements, without the emotional experience appearing to justify itself. For instance, the fact that the tiger is running towards me and has sharp claws is a good reason for me to be afraid of it and a good reason for me to form the belief that the tiger is dangerous (Brady 2013, 113). Here the non-evaluative features (the claws and the run of a predator towards me) are important to evaluate whether the emotion is accurately representing the situation and is thus rational. Moreover, these features are also reasons for the evaluative judgement that the tiger is dangerous. However, when this is the case, then this emotional experience cannot provide reason for the evaluative judgement. Returning to the tiger example, the fear of the tiger cannot be a reason to believe that the tiger is dangerous, because it would amount to saying that I have a good reason to be afraid of the tiger because I am afraid of the tiger (Brady 2013, 114).

The second main criticism from Brady is that emotions cannot stop the demand of justification, even though the perceptual model affirms it can do so. Indeed, emotions motivate the search for reasons rather than put a halt to the search. This, Brady argues, is because the goal of evaluative knowledge and ethical thought is not simply knowing, but rather understanding. For instance, when someone is angry, their anger leads them to search for additional reasons about why they are currently angry (Brady 2013; Milona 2015). Thus, emotions do not provide strong reasons that suffice to provide justification for one's evaluative belief.

1.2 Brady and Attention Capture

From this second criticism, Brady builds his own theory, according to which "emotional experience captures and consumes our attention, motivating the search for reasons that bear on the accuracy of our initial emotional responses, and helping to satisfy an epistemic and a practical need for reappraisal" (Brady 2013, 116). When one is afraid, they look for clues about the situation: is the tiger a real one, is it coming in Julieta's

direction, ready to kill? Here, Julieta's fear motivates her to search for reasons that her belief that *the tiger is dangerous* is justified. In this situation, even though Julieta's fear does not justify her belief that the tiger is dangerous, it is nonetheless useful. As Brady explains, emotions can work as proxies for genuine reasons. They are "useful stand-ins or surrogates for genuine reasons for evaluative beliefs, but lack that status themselves" (Brady 2013, 129). Emotions, according to this view, have two dimensions: first, they are an initial, reflective and automatic evaluative response, which helps the agent capture their attention on the situation; second, the emotions enhance the evaluative response, as the attention is consumed, which forces the agent to search for additional reasons for the initial evaluation. This leads to a reappraisal of the initial evaluation, to either approve it or reassess the situation. For this reason, Brady argues that emotions "raise rather than silence the justificatory demand with respect to our initial emotional appraisal" (Brady 2013, 110).

The attention model offered by Brady contrasts with perceptual models, such as Döring's, as emotions cannot provide justification for evaluative beliefs themselves. Thus, while these two approaches agree that emotions play an important role in the acquisition of evaluative beliefs, they disagree fundamentally on what this role is. Let us now discuss two problems with these theories which motivate the search for an alternative explanation about the role of emotions in knowledge acquisition.

1.3 The Dilemma

Given the current theories available to explain the epistemic import of emotions, one is faced with a dilemma: if one wants to say that emotions offer direct justification for evaluative knowledge, they have to choose the perceptual horn, and have to subscribe to a perceptual model of emotions. But this comes with two problems: first, the perceptual account of emotions is only one of the various theories of emotions available. Thus, if one wants to offer an explanation about the epistemic role of emotions, they cannot remain neutral about which theory of emotion to ascribe to. This means that the force of the argument will be limited to people who already accept a perceptual model of emotions. More worryingly, the perceptual model fails to explain how emotions can

justify the rationality of emotions without falling into a circular argument, as Brady explained. The alternative is to go for the Brady horn of the dilemma.

But if one chooses this side of the dilemma, then they have to give up the idea that emotions can provide justification themselves, and they have to accept that emotions are just proxies for genuine reasons, rather than providing reasons themselves. However, this horn of the dilemma is not satisfactory given that the perceptual model seemed to have something right: we do trust our emotions at face value in certain situations where they seem to provide justification directly. Additionally, we saw with the applied cases of endometriosis and microaggressions that searching for additional reasons is not always possible, or that emotions seem to play a more direct justificatory role. Affective experience more generally would suffer from the same problems as emotions, as reducing it to a form of perceptual model would be too restricting and fall into the same circularity problem. And when it comes to the Brady horn, we want more from affective experience, as explained before.

This dilemma is a serious one. However, there is still the possibility of finding an alternative theory, where one can remain agnostic about the underlying theory of affective experience, while arguing that affective experience provides immediate justification. The rest of this chapter proposes to offer such an alternative. To build this alternative, I look at theories present in other domains of philosophy, specifically practical rationality. In chapter 1, I explained that we want affective experience to fulfil the commonly accepted standards of epistemic rationality. In what is to come, I propose that this conception of rationality is not suitable to make sense of the epistemic value of affective experience.

2. The Reflective Model of Affective Experience

To exit the dilemma, we need to take a step back from epistemology. In this section, I propose to reframe the issue at stake by borrowing from discussions in action theory and practical rationality. Questions discussed in the domain of action theory and practical rationality revolve around how agents select action options, given the reasons

they have; in short, about what the norms of practical rationality are (Jones 2003). While the reasons at stake in these discussions are practical reasons, and the concept of rationality discussed there is not purely epistemic, I argue that there are some interesting parallels that can be drawn and applied to discussions of rationality in contexts that are more clearly epistemic, and to the case of affective experience.

I focus on Karen Jones's conception of normative agency and her discussion of what reflective rational agents are. I use Jones' model to define affective experience as a reason-tracking ability rather than a reason-responding one, thus explaining why affective experience seems inadequate to fit the requirements of rationality in analytic epistemology. By using this new conception of rationality, I build a new model, the Reflective Model, which incorporates the possibility of diachronic calibration in rational agents.

2.1 Jones' Rational Agency

According to Jones (2003), agents think that they can guide their actions as they have a conscious deliberative self. Indeed, if agents are engaging in conscious deliberation, then they think that they can rationally guide their actions via reasons. Moreover, these reasons are "not merely registered, but also understood as reasons, that is, understood as justifying the performance of an action" (Jones 2003, 188-189). This means that agents are not simply viewing themselves as a system that passively registers and responds to the environment, nor do they view their reasoning abilities as operating as a self-contained epistemic mechanism that adds to perceptual or emotional capacities. On the contrary, a reflective rational agent is able to engage in conscious deliberation and have a conception of reasons as reasons. If the agent does not meet these requirements, then the agent is not conceiving of themselves as a rational agent (Jones 2003).

This particular type of agency, where agents understand reasons as reasons in order to guide their actions is what Jones calls *agents as reason-responders* (Jones 2003, 189). Under this definition, reason-responders need to be actively engaging in complex

reasoning and reflection. As Jones explains, “to be able to respond to reasons as reasons, an agent requires critical reflective ability, dispositions to bring that ability to bear when needed, and dispositions to have the results of such reflection control their behaviour” (Jones 2003, 190). When an agent considers themselves to be rational, it is insofar as they take themselves to be able to understand reasons as reasons and to be able to act on them to guide their actions; in short, they take themselves to be a reason-responder (Jones 2003). For instance, when I decide whether to eat the last cookie in the jar, I think about my sister who won’t have any snack if I eat it. If I reflect on this and decide that this reason justifies me in not eating the last cookie, I have acted as a reason-responder.

However, being a reason-responder is not the only conception of agency available when it comes to interacting with reasons. Indeed, agents can interact with reasons without responding to them as reasons. In this second case, the “agent is capable of registering reasons and behaving in accordance with them, but it need possess neither the concept of a reason nor have a self-conception” (Jones 2003, 190). These agents are what Jones calls *reason-trackers.* Reason-trackers do not need higher-order reflective capacities, nor do they need to display robustness in their capacity to track reasons. Reason tracking is possible through innate and learned behaviours, which enables animals to be reason-trackers. Overall, reason-trackers will be worse at modifying their course of action when changes in the environment occur compared to reason-responders, given that they engage less in critical reflection (Jones 2003). The idea of reason-tracking appears more plausible when we consider reasons outside of action theory. For instance, perception is a reason-tracking ability. When using their perceptual experience, agents simply register reasons, without reflecting on them, unless they are asked to. When moving towards the domain of beliefs, many of them appear to be the result of reason-tracking rather than reason-responding.

While there are two ways for agents to interact with reasons, namely by simply registering them, which amount to tracking, or by responding to them in virtue of them being reasons, only the latter is what agents themselves consider to be rational agency. From a first-person perspective, the agent has to be a reason-responder, otherwise, they

are not an agent (Jones 2003). This first-person perspective creates a hierarchy between these two conceptions of agency, where being a reason-responder is at the forefront, while reason-tracking is relegated to the background. This happens even for belief formation: even though most belief formation is reason-tracking, agents think of themselves as responding to reasons when they are asked to justify their beliefs.

However, this doesn't mean that reasons-tracking agents cannot be considered rational. Indeed, for Jones, agents are rational as reason-trackers insofar as they are committed to the ongoing cultivation and exercise of reflective self-monitoring (Jones 2003, 194). Reflective self-monitoring is defined by Jones as 'such that [the agent] would not rely on that first order subsystem were it reasonable for her to believe that it failed to reason-track" (Jones 2003, 195). This reframing of rationality helps us include fast and habitual action as something that can be rational, i.e., if the agent wouldn't use it if they were doubtful of their ability to track reasons when using that cognitive system.[12]

Highlighting that rationality is possible when reason-tracking opens the door to new possibilities. Moreover, as I have shown, forming beliefs is more of a reason-tracking ability than a reason-responding one. This gives us an additional reason to move towards reason-tracking to explain the rationality of affective experience. If what is important is the cultivation and exercise of self-monitoring, then it is possible that some fast and habitual systems could fit this bill, and so could affective experiences. What we now need to establish is whether affective experience is reason-tracking and whether reflective self-monitoring is possible.

2.2 Affective Experience as Reason-Tracking

Before I dive into how affective experiences interact with reasons, it is important to discuss the differences between judgement/belief and affective experience when it comes to their content and how they relate to truth. Our previous discussion of rational

[12] I will come back to the idea of fast and habitual action in section 2.3, as I will discuss the work of Kahneman and Klein (2009).

agency was linked to choosing the best course of action. However, in what will come, it is reasons linked to choosing what represents the world best, or truth, that interests us. Thus, we should now look at how different mental states relate to truth.

Beliefs and the propositions contained in judgements can be assessed as true or false because their formal object is truth, or how close they do represent the world. For affective experience, the formal object is not truth; it is the value represented by objects in the world which elicit the affective experience. Therefore, while it is possible to discuss whether an affective experience does represent its object correctly, it is not the primary function of affective experience to give an accurate representation. However, as Döring (2009) explains for emotions, "this need not, and does not, prevent the subject from regarding these contents as true *by default*" (245). What Döring means by this is that emotions are seen as generally reliable cognitive systems which can give accurate representations of the world. When an agent has a belief or a judgement, they often ascribe a truth value to it in order to make it intelligible. On the contrary, with affective experiences, a truth property does not need to be ascribed to the object to make sense of it. For instance, when facing a gorilla, I can make sense of my fear that the gorilla is dangerous without having to ask myself whether I should be afraid because the gorilla is truly dangerous. The same goes with pains; I can make sense of the pain in my leg without asking myself whether I should be in pain because my leg is broken. This difference explains why affective experience is not the type of cognitive system that pushes the agent to always engage in critical reflection and reason-responding. Indeed, when having an affective experience, one can simply register the reasons for the affective experience to arise. There is no need to understand the reasons as reasons to experience the affective experience and form an evaluative judgement. Rather affective experience is reason-tracking, where agents will use the reasons provided by their affective experience to inform their beliefs or actions.

Affective experience understood as a reason-tracking ability explains the possibility of conflict without contradiction[13]. Indeed, because agents don't consider that their affective experience provides them with reasons that justify their beliefs, they might still hold on to their beliefs acquired through affective experience. And these beliefs might conflict with reasons acquired through other means. As the agent is less critical when they track reasons, the conflict is not immediately registered and analysed at a higher-order level. Moreover, this explains why agents struggle to recalibrate immediately and respond to defeaters when using their affective experience to form beliefs. The reason-tracking agent trusts their first order system to generally deliver. In occasions where the system fails, the agent doesn't recalibrate immediately. The calibration of beliefs obtained via reason-responding can occur immediately for a single output because reasons are understood as reasons, and this occurs in the same cognitive system as the acquisition of the reason. For tracked reasons, calibration occurs at a different level of the cognitive system, namely the meta-cognitive level, as the agent needs to reassess whether they trust the whole system rather than the single output made by the system. This explains why affective experience doesn't respond well to immediate defeaters if the calibration is to occur synchronously with the acquisition of the defeater.

In the next sub-section, I discuss the consequences of framing affective experience as a reason-tracking ability, and whether rational agents could use it to acquire justified beliefs. For that, we need to see whether affective experience can be calibrated through reflective self-monitoring.

2.3 Reflective Self-Monitoring and Affective Experience

As explained in section 2.1, an agent is rational as a reason-tracker if they cultivate and exercise reflective self-monitoring. We've established that using affective experience for belief formation should be understood as a reason-tracking ability rather than a reason-responding one, as it is fast, habitual and agents don't need to register the reasons

[13] See chapter 1 section 2.3 for a discussion of the problem of conflict without contradiction.

provided by their affective experience as justifying their beliefs. However, for it to be an ability that rational agents would use, we still need to prove that agents wouldn't use it if they were doubtful of tracking reasons with their affective experience. In what follows I discuss whether agents using affective experience can engage in reflective self-monitoring, as this will enable us to decide whether affective experience can be considered an ability that rational agents would use.

I explained in the previous chapter that agents are unable to calibrate their affective responses in various scenarios, even when presented with defeaters. Moreover, given the possibility of conflict without contradiction with affective experience, agents are unable to respond to conflicting reasons and adjust their affective responses accordingly. Given that affective experience fails as a reason-responding ability, it tends to be excluded from the range of ways of knowing. However, I have now introduced an alternative: affective experience should be understood as a reason-tracking ability. This means that the standards of rationality are different, and calibration based on defeaters occurs at a different cognitive level. As Döring explains, when "an emotion is revised in the light of better judgement and knowledge, this is not a matter of contradiction but of calibration" (2009, 246). This means, we are better off switching to a diachronic model, where agents can recalibrate the whole system immediately or at a later point based on their experiences and feedback. Based on this, we can then decide whether to use affective experience or not as a way of obtaining knowledge or justification.

As calibration is at the core of the rationality of affective experience, let me explain the various ways in which affective experience can be calibrated through reflective reasoning. I argue that there are four main ways to calibrate affective experience: immediate internal regulation, metacognitive internal regulation, extrinsic regulation and collective resistant regulation. I will use the term regulation based on the work of James Gross, where emotional regulation "refers to the processes by which individuals influence which emotions they have, when they have them, and how they experience and express these emotions" (Gross 1998, 275). I take affective regulation to mean the processes by which individuals influence which affective experiences they have when

they have them and how they experience and express these affective states. Two of the types of affective regulation are personal and performed by the agent themselves, whilst the other two are social and initiated by other agents.

Immediate intrinsic regulation occurs when the agent regulates their affective experience and the belief it produces *during* the experience, as they realise that it stops tracking reasons. For instance, someone stops fearing an extremely realistic animatronic of the tiger when they realise that it is innocuous. Additionally, imagine a daughter who has a gut feeling that something terrible happened to her father. When the father calls her and reassure her that he is fine, she feels relieved. Her gut feeling and belief that something happened to her father have disappeared, as she immediately calibrated based on the new information. This form of calibration is the one we would expect from a rational agent who is also a reason-responder. It is immediate and can take into account any new information that appears. However, as I discussed, it doesn't work in the case of recalcitrant emotions and in cases where there is conflict without contradiction. Thus, we need to look at other types of calibration to explain how less-responsive affective experiences can still be used rationally as ways of knowing.

The second type of calibration is *metacognitive intrinsic regulation*. In these cases, the agent realises after the experience that it didn't track reasons properly. They subsequently try to regulate the system before having new experiences. Meta-cognition can be understood as the "self-evaluation of one's own predicted or acquired mental properties" (Proust 2013, 5). As Proust explains, our metacognitive processes include retrodictive evaluation, which compares whether the present outcome fits the simulation of how the world should be like. This is what she calls "post-evaluation", and its aim is to determine whether a mental action is successful (Proust 2013, 189). This echoes Kahneman and Klein's proposal about how intuitive judgements are explained on the two-system model: system 1 operations are automatic, involuntary and effortless; whereas system 2 includes controlled operations and self-monitoring. They highlight that system 2 ensures the "continuous monitoring of the quality of reasoning" (Kahneman and Klein 2009, 519). Checking the intuitive judgements produced by system

1 therefore requires the effortful use of system 2. Given that this dual process can be difficult, it happens that people do not always perform this higher-level monitoring and don't always reflect on whether the intuitive outputs from system 1 are correct (Kahneman and Klein 2009). It is likely that the intuitive outputs produced by affective experience follow this model: after having an affective experience which leads to the production of a belief, the agent can reflect on whether their affective experience is fitting or warranted.

Let us examine the cases from the previous chapter in this new light. In the case of Tomer in the **shooter bias** case, he can ask himself why he was afraid of targets holding phones. He might be able to reflect on this and determine that his fear was misleading. This might lead him to decide to stop trusting his fear if he is put in a similar situation again. In the **spider phobia** case, Dorian might realise that his fear is irrational. If he does, he might seek some therapy for his phobia to retrain his intuitive fear reactions. Finally, in the **jealous partner** case, Kathryn could reassess her jealousy: she could try to discard future beliefs about Annette's faithfulness given that she knows that her intuitive reaction is misleading. For all three cases, there is a possibility that all agents could engage in a form of internal metacognitive regulation. By controlling the quality of the beliefs produced by their affective experience, they could try to influence and regulate future affective responses. This could help them make their affective experience a system which tracks reasons better, thus making it more reliable. Internal metacognitive regulation is central for rational agents: it is through meta-cognitive reasoning that they can assess whether their affective experience fails to track reasons, and thus whether they should rely on it to form beliefs or not. However, as we have seen, internal metacognitive regulation can be foregone by agents who consider it too be too difficult or effortful. This means that some agents might fail to engage in calibration on their own and might fail to use their affective experience in a rational manner. However, as I will now explain, there are two types of social calibration which play an important role in regulating affective experience, as not all calibration is initiated by the agent.

The first type of social calibration is *extrinsic regulation.* As Archer and Mills explain for emotions, extrinsic regulation occurs "when people try to regulate other people's emotions. For example, someone may try to make someone happy by buying them a gift or to make someone feel ashamed by highlighting their deficiencies in front of others" (Archer and Mills 2019, 4). Extrinsic regulation can also occur when the agent already has the affective experience: in these cases, the external person tries to regulate the agent's affective experience by either pushing them to experience another affective state, or by trying to stop them from experiencing their current affective state. For instance, Abby could tell Bernard that he shouldn't be angry at his sibling who just ate the last cookie. She could discuss with him why he is angry and encourage him to stop being angry as his emotion is not fitting. In these cases, the regulation seems positive. Indeed, the agent shouldn't be experiencing their current affective state as it fails to track genuine reasons. With the help of the external person, they might be able to track better. For instance, in the **jealous partner** case, Kathryn could be told by her friend Paula that her jealousy is leading her astray and that she shouldn't trust it. After her discussion with Paula, Kathryn knows that her affective state is misleading and that she shouldn't be trusting any information that can be influenced by her jealousy. By reflecting on her jealousy, she becomes rational because she knows it fails to track and thus stops using it.

However, these positive cases of extrinsic regulation come alongside negative cases. Archer and Mills highlight that the extrinsic regulation of affective experience can be used to further oppression (2019). People experiencing anger at injustice might be forced to hide or repress their fitting emotions due to an unfair environment: they might be told that they shouldn't be angry at the oppressive situation. For instance, after being repeatedly told that everyday sexism is normal, a woman might stop feeling angry when she witnesses sexist microaggressions. In these situations, extrinsic regulation can lead to affective injustice when the affective experience is appropriate to the situation, as the agent is told to suppress it for the wrong kinds of reasons (Archer and Mills 2019). Overall, extrinsic regulation can both benefit and hinder the agent who is asked to engage in it: if their affective experience fails to track reasons, it is beneficial for an

external person to point this mistake; on the contrary, if the external person is mistaken and the affective experience was tracking the right reason, then asking the agent to regulate can be detrimental.

The final form of calibration is one that is useful to counteract the negative types of extrinsic regulation. *Collective resistant regulation* refers to the regulation of affective experience by marginalised agents who critically reflect on their experience within their marginalised communities. This last type of calibration has been widely discussed within the literature on epistemic resistance and feminist epistemology (Cherry 2021; Collins 2000; Lorde 1984). As Saba Fatima highlights, sharing experiences within marginalised communities is crucial, "for it is the sharing of our stories within marginalised communities that allows us to critically reflect on our cumulative experiences" (Fatima 2020, 166). Indeed, marginalised agents can discuss events that affected them, which might appear innocuous when taken individually, but form into a pattern of oppression when tied together. For instance, the experience of sexual harassment might have stayed underplayed had women not come together to discuss their experiences and discomfort. Through discussing it with their peers, they were able to see that their anger was appropriate.

The importance of resistant spaces for knowledge and reflection has been highlighted in the work of multiple afro-feminists scholars like Audre Lorde (2019) and Patricia Hill Collins (2000). When reflecting in these spaces, agents can analyse their affective experience. As Audre Lorde highlights, to analyse one's anger "implies peers meeting upon a common basis to examine difference, and to alter those distortions which history has created around our difference" (Lorde 2019, 122). She also highlights that Black women need to share their mistakes with each other, as it is "the only way learning occurs." (Lorde and Tate 1983, 105) Patricia Hill Collins echoes this sentiment. For her, "Black women's ability to forge these individual, often unarticulated, yet potentially powerful expressions of everyday consciousness into an articulated, self-defined, collective standpoint is key to Black women's survival" (Collins 2000, 36). Similarly, Black psychiatrist Chester Pierce makes use of collective resistant regulation as his response

to microaggressions. Indeed, he explains that we need street therapists who "become a consultant to the ever-growing efforts by the poor to organise themselves, to articulate their wants, to negotiate their demands" (Pierce 1970, 279-280). More importantly, the street therapist has the tools to "point out and clarify feelings" (Pierce 1970, 280). Here the street therapist plays a crucial role, because they help highlight features of the environment that are not easily accessible because racism has become subtle (Pierce 1970). Not only does exchanging with one's community help develop and confirm the appropriateness of one's feelings towards injustice; it also helps correct affective experience which misfires. As Myisha Cherry argues, "the solidarity technique provides a trusted community that people can rely on to challenge them and hold them accountable in their anger. If at any point anger becomes unproductive, a moral critic can point it out" (Cherry 2021, 152-153).

In short, resistant communities create spaces where affective experience can be productive, but also criticised and regulated in ways that won't feed into the harms experienced by agents who undergo inappropriate extrinsic regulation. This is due to the particular features of these communities: in these spaces, solidarity is presupposed, and the people challenging the affective experience are likely to possess the same information and share a close standpoint to the challenged. This helps circumvent issues around information being missed by the challenger (Cherry 2021, 153). *Collective resistant regulation* is more likely to be used for affective experience that responds to the social features of the environment, in particular its unfairness. This explains why most of the literature highlighting this form of calibration focuses on anger in response to racism. However, *collective resistant regulation* could happen in any resistant communities where solidarity is a core feature, and it doesn't preclude that it could be applied to affective experiences other than anger.

These four types of affective regulation help explain how affective experience can be used by rational agents. Indeed, if agents engage in reflective calibration, it means that they wouldn't continue to use their affective experience if it fails to track reasons. The calibration helps them do two things: first, they learn through reflection that they failed

to track, which means that they should stop using this type of affective experience; second, once they know that, they can either work to improve how their affective experience tracks or discard any belief coming from a certain type of affective experience if it cannot be improved to track. Moreover, through social forms of calibration, agents are made to reflect on their affective experience in a way which fits within a wider group of people. This helps bring a form of intersubjectivity to affective experience. Indeed, due to social calibration, affective experience is not a purely individual experience and is rather interconnected with other agents.

This leads me to the introduction of the *Reflective Model of Affective Experience* (shortened to the "Reflective Model" in the rest of this thesis), which explains why affective experience should be considered a source of knowledge. Affective experience is a source of knowledge if the agent has rational control over using it, and a rational agent would stop using it if they didn't consider it to be reliable in the relevant conditions. Given the numerous ways affective experience can be regulated, if an agent is going through the forms of calibration explained above, then they can be said to be making a rational use of affective experience as a way of knowing[14]. An agent who successfully tracks every time they use a certain type of affective experience in a particular case doesn't have to be aware that they are doing so, until they fail. And when it fails, the agent could reflect on why it failed and whether the failure could be avoided in future cases. This echoes the use of other accepted sources of knowledge: for perception, agents trust it until they discover conditions in which it fails. For instance, people know to stop trusting their perception if they have taken hallucinogenic drugs like LSD, given that it fails to track in most cases. The same applies for affective experiences: I know I shouldn't trust my fear to provide me with reliable information about a spider when I have a phobia of spiders.

[14] I will return to this in the next section when I discuss the problem of people who don't realise that the tracking fails, even when all four types of calibration are present.

The Reflective Model avoids the first problem raised by Brady to the perceptual model: on the perceptual model, it was impossible to explain how emotions can provide reasons for the evaluative judgement without the affective experience justifying itself. Remember the tiger example: the fear of the tiger cannot be a good reason to believe that the tiger is dangerous, because the non-evaluative features of the situation are both justifying the rationality of the emotion and the associated evaluative judgement (Brady 2013, 114). On the Reflective Model, this problem vanishes: someone's fear of dangerous animals is usually considered to be a reliable cognitive system, hence why the fear provides the justification. If a rational agent were to fail to track reasons with their fear, they would stop using it to justify their belief that the tiger is dangerous. The fact that the fear is rational and justifies the belief that the tiger is dangerous is not dependent on the non-evaluative features, i.e., the claws and run of a predator towards the agent, of the situation. With the Reflective Model, the non-evaluative features only trigger the emotion, without justifying the associated belief, hence avoiding the circularity issue pointed out by Brady. Given this feature, the Reflective Model has another advantage over the perceptual model.

3. Discussion

Let me now address various objections and problems that might arise for the Reflective Model of affective experience. These can be grouped in two strands: first, affective experience doesn't track reasons with a high enough degree of success to count as a valid source of knowledge; second, affective experience doesn't bear immediate justificatory power, and thus cannot count as a source of knowledge.

Let us start with the first strand. One might argue that affective experience should not count as a source of knowledge, even given these considerations because the cases where it appears to track reasons appropriately are scarce; while a source should yield a high propensity of true beliefs (Audi 2002). Whilst it is true that affective experience fails in a variety of cases, this should not preclude it from being included. Indeed, it still provides some agents with a good or the best way of obtaining certain kinds of knowledge or knowledge of certain kinds of things, thus being risky but worthy of

epistemological consideration. As we have seen in the cases of pain or microaggressions, affective experience plays a crucial role, and excluding it entirely can lead to negative consequences. This response will be fleshed out in more detail for a variety of applied cases in later chapters of this thesis. Moreover, it is not clear that affective experience is the only source facing the problem of low success rate. Indeed, memory retrieval is fallible and fails in a good number of cases too. Perception similarly fails: for instance, there are many optical illusions in which the agent is tricked. Affective experience is imperfect, but so are other commonly accepted sources. Moreover, affective experience differs from other methods which sometimes yield true beliefs but which we wouldn't count as sources, like tarot card readings, clairvoyant visions or bone readings. In all these cases, there aren't any subtypes which seem directly linked to the world and evaluate it truthfully. Moreover, the agent would be unable to distinguish good cases from bad cases in these, as there is no reflective element, nor is it regulated internally and externally.

In the problem cases introduced in the previous chapter, agents would stop relying on their affective experience to obtain knowledge, because they know or learn that it fails to reliably track. Dorian knows in the **spider phobia** case that his fear fails to track, and he shouldn't use it as a source of knowledge if he is rational. Tomer in the **shooter bias** case discovers that he has a racial bias. He learns through reflection that he shouldn't trust his fear in these cases or future ones due to his bias. Finally, in the **jealous partner** case, Kathryn can reflect on her jealousy with someone external, a friend or a therapist, and realise that she shouldn't trust her jealousy. This leads her to stop trusting her current and future beliefs that are influenced by her jealous state. When looking at the general objections against affective experience, we can now see how they can be circumvented. For the low reliability and the false information problem, a rational agent would not use their affective experience in these cases. Indeed, they would realise that it stops tracking reasons. For sensitivity to defeaters, a rational agent would track reasons and would use their reflective abilities. Given the many opportunities to regulate affective experience, an agent would need to act in bad faith to refuse to regulate their affective experience if they discover that it stops tracking. Finally, when it

comes to conflict without contradiction, in these cases, the reflective rational agent would weigh the beliefs against each other, and they would reflect on whether their affective experience accurately tracks reasons. All problems now boil down to problems about rational agency rather than problems about affective experiences being a source of knowledge or not. Indeed, a reflective rational agent would be able to use affective experience productively as a source of knowledge, because of their ability to reflect on their use of it, leading to recalibration.

This reflective aspect is crucial to explain why certain categories of affective experience, like the anger felt by targets of microaggressions or patients with pelvic pain, are good sources of knowledge. In these cases, a high degree of reflection has occurred, which makes the agents more confident in their ability to accurately track reasons through their affective experience. For instance, with microaggressions, epistemic certainty is gained because all types of regulation occur: the agent reflects on their own experience when it is occurring, asking themselves whether their emotion or gut feeling is right, searching for additional reasons that confirm their affect. Moreover, other people might question their affective experience: are they sure that they were microaggressed? Are they sure that they are not overreacting? Was it a microaggression or just a poorly formulated compliment? External regulation happens through this questioning. Finally, by coming together with people from their community, they can reflect on their experience in an environment which helps cumulate epistemic resources as they can compare their experience with the one of others (Fatima 2020).

However, someone might object that there are cases where agents seem to use all four types of calibration and yet fail to recognise that they are irrationally or incorrectly using their affective experience to track reasons. For instance, the anti-vaccine movement during COVID-19, saw millions of people use their fear to justify their decision not to get vaccinated, given that this new vaccine could include some unknown side effects (Burki 2020; Palmer 2022). Many anti-vaccine activists appeared to engage with the arguments of scientists yet found objections, thus engaging in external regulation. For instance, they used the Vaccine Adverse Event Reporting System (VAERS), a governmental

database administered by the Centre for Disease Control and Prevention (CDC) in the United States, to extract data about the reactions to the COVID-19 vaccine and sustain claims that the vaccine could be dangerous (Jarry 2021). Moreover, they used a form of social resistant calibration, by engaging with other agents from the anti-vaccine community online (Burki 2020). From the perspective of an anti-vaccine activist, their fear of the vaccine tracks the danger of vaccination and seems rational given the data obtained on the VAERS database. How could we distinguish this type of case from ones like racialised people using their anger to track racial microaggressions? In both contexts, an agent engages with other agents from the majority group (people who are pro-vaccination/people who are not the target of the microaggression) yet they find reasons to maintain their belief gained through affective experience through collective resistant regulation with a smaller group, even when confronted by the wider group. I have argued that microaggressed agents are justified to use their anger to identify microaggressions, yet I would ideally want to resist that people against vaccinations are justified to use their fear to justify not getting vaccinated. Nonetheless, both cases appear extremely similar from the perspective of the agent which appears to create a problem for my argument.

How could we explain the difference between the case of anti-vaccine activists and people who are microaggressed? There are two ways to answer this. First, we could try to show that the anti-vaccine population is not fully engaging with the evidence presented to them by the wider community. Given their refusal to engage in extrinsic regulation properly and their unresponsiveness to positive counterevidence, we can say that they fail to be rational. On the contrary, targets of microaggressed are often told that they are imagining the microaggression without being presented with positive counterevidence. In that case, they are not failing to engage with counterevidence, given that counterevidence is not presented. However, this explanation might remain unconvincing to some people. This leads me to the second answer: not all instances of affective experience which should be conducive to reliable belief-formation leads to a true belief. Affective experience remains a fallible way of obtaining knowledge, and there will be bad cases. Specifically, it will be hard for affective experience to be a good

process for belief formation in cases of collective delusions, even if the agent goes through all types of calibration. Nonetheless, we should not let this problem lead to the exclusion of affective experience. Indeed, no source is infallible. For all the sources, there are types of cases where they are often reliable but fail on some occasions. To say that we should not accept affective experience as a source of knowledge because it has some bad cases like the anti-vaccine movement would be equivalent to saying that we shouldn't accept memory or perception because they fail on some occasions. We don't do that with the currently approved sources as we accept that they can be fallible, and affective experience should be afforded the same treatment.

The model I have proposed will not uniquely determine whether affective experience is producing a true belief or not, but it makes it likelier to be reliable if the agent engages in reflection and calibration. By large, if agents do this, then it captures a large proportion of cases where affective experience tracks reasons accurately. In some cases, they might still engage in these reflective mechanisms and get it wrong. If agents engage in reflective self-monitoring when using their affective experience to form and justify their beliefs, then the result is epistemically valid. The belief formed through affective experience might not be true, but that doesn't make it an inappropriate method of knowledge acquisition. There is room to explore what characterises the cases where affective experience used by self-reflective agents fails to track beliefs. For instance, what are the specific characteristics of the regulation around the anti-vaccine movement that leads to the failure of the reason-tracking? What distinguishes the anti-vaccine group from other epistemic bubbles like anti-racism activists? These interesting questions are unfortunately ones that I will not answer in this thesis. However, I intend to address them in future research.

One further phenomenon is worth mentioning: when people only have access to a certain type of regulation, rather than all, there is the possibility that regulation becomes unproductive. For instance, targets of microaggression who don't have access to collective resistant regulation, and are not aware of the patterns of microaggressions can be gaslighted by extrinsic regulation associated with epistemic injustice (Fatima

2020). This leads to their own sense of reality being warped (Fatima 2020, 171). In this case, the agent would stop trusting their affective experience, even if it accurately tracks reasons, because they are made to think that it is unreliable by other agents. This shows the importance of balance between the different types of regulation, as they can epistemically influence agents for better or worse. In the case of pelvic pain and endometriosis, a similar tension is visible: on the one hand, patients are often told that their pain is normal and that they shouldn't worry. This form of extrinsic regulation is telling them to discard their belief based on their pain. On the other hand, their internal regulation (both immediate and metacognitive) is confirming that their pain appears to be justified. When patients meet other patients with the same condition who have been able to get a diagnosis, they enter the stage of collective resistant regulation. This helps them counteract harmful extrinsic regulation and helps them gain epistemic certainty about the accuracy of their pain.

One might object that my view accepts too many cases where affective experience fails, especially now that I have conceded that affective experience can still fail to track reasons even with the four types of regulation at play. This is a price I am willing to pay to support the acceptance of affective experience as a source of knowledge, given the benefits it yields for certain agents, in particular ones from marginalised groups. This is something that will become apparent in chapter 4, where I introduce the idea of source-based epistemic injustice. For now, I am hopeful that the many ways we regulate affective experience, both internally and socially help us safeguard it as a good way of knowing. Moreover, I am hopeful that reflective agents who engage in these various forms of calibration can use their affective experience in a productive manner. The question then is, what if most people are unreflective and stubborn? This problem is a real one, but not for my account. Indeed, I have argued that if people engage in reflective self-monitoring, then they will be rationally using affective experience as a source of knowledge. It might happen that there are few cases where this happens, but it shouldn't preclude affective experience from counting as a source. Moreover, I hope

that the way we are embedded as agents would prevent this feature[15]. Indeed, there are chances that our wider epistemic community could help the agents who use their affective experience in a bad way, as it would be pointed out, and they would learn not to trust themselves. Excluding a way of knowing because some agents don't know how to use it would come at a great cost for our collective epistemic enterprise, as this would mean giving up a useful way of knowing because we are afraid of people who refuse to engage in reflection and the exercise of rationality.

The more a type of affective experience is regulated by all four types of regulation, the more likely it is that it is rational to use it as a source of knowledge. Indeed, it has withstood many tests, and with each of these, the agent gains epistemic certainty. Just like an expert birdwatcher knows that they can trust their perception, an agent knows that they can trust their affective experience because they have put it to the test. One should note, however, that it might be hard to put it to the test in some domains. Failure and success help us know when we should trust a source, and the same applies to affective experience. The focus of current contemporary analytic epistemology on the cases where it fails gave us a myopic focus. Instead, looking at the wider picture, where we try to see what a rational agent would do, when they would trust it or not turns out to be much more productive. Indeed, we can now see why affective experience counts as a source in some cases and not in others. Affective experience is not just a wild card: it is a cognitive ability which is tested through many types of regulation, and evaluations are improved with new information. Reflective rational agents have epistemic confidence to use it in some cases to gain knowledge. Conversely, agents who fail to ask themselves whether their affective experience tracks reasons can be led astray by it.

[15] This personal view might appear overly optimistic given the current trends in populistic politics and the rise of fascism. Fortunately, the argument does not rest on this optimism given that the Reflective Model is formulated as a conditional. If it happens that there are only a few cases which fit my model, it does not preclude affective experience from counting as a source of knowledge even if it is rarely used as such.

Now for the second strand: if we have this reflective element, does this mean that affective experience cannot provide immediate justification? This objection can be formulated drawing on Harrison's work on emotional dogmatism (Harrison 2021). Harrison defines emotional dogmatism as the idea that "an emotional experience that makes it seem to you that *e* immediately and defeasibly justifies you in believing that *e*" (2021, 2536). For Harrison, emotional dogmatism is problematic because it over-generalises and starts including bad cases. Compare two cases: in the first one an interviewer thinks that an interviewee is acting suspiciously, based on their emotional experience which picks up on subtle mannerisms and micro-behaviours; in the second one, the interviewer has sexist biases and is unknowingly biased, and has an emotional experience that the interviewee is acting suspiciously. Harrison argues that both emotional experiences have the same justificatory power, even though one is epistemically legitimate whilst the other isn't (Harrison 2021). Harrison explains that answers which rely on the use of knowledge-how could explain this difference: one emotional seeming arises from the exercise of a knowledge-how capacity, whilst the other is the result of illicit biases do not provide immediate justification. However, these responses solve the problem at the cost of the immediacy of justification. Indeed, because the knowledge-how capacity is determined by the possession of background information, this means that the seemings provided by the emotional experience are not doing the core epistemic work and need to be mediated by the background information (Harrison 2021). Does the Reflective Model face a similar issue? In a way, the Reflective Model also relies on background information due to its reflective aspect, which might make its justificatory power less immediate. However, it is not the case: Harrison's worry mostly applies to the phenomenal component of affective experience bearing the justificatory power. If we take affective experience as a cognitive capacity to provide justification, we can avoid this problem. Indeed, by considering all aspects of affective experience, we can see that it incorporates this background information within the capacity itself. Given that the background information is incorporated within the cognitive capacity, then affective experience can provide immediate justification, because it is not just its phenomenal component at play. Going back to the problematic

case of the interviewer: the Reflective Model can make the difference between the two cases: the biased interviewer could be called out by their colleagues. If they were a reflective rational agent, this would defeat the justification they acquired through their affective experience. By contrast, the good interviewer would not have defeaters and could continue to use the justification acquired through their affective experience.

4. The Reflective Model Within Epistemology

After introducing the Reflective Model, let me situate it within wider conceptions of epistemic justification. One might wonder how my conception of affective experience as a source of knowledge relates to some theories available in epistemology. For instance, is the form of justification provided by affective experience internalist or externalist? Moreover, could we consider my account to assume a form of reliabilism, phenomenal conservatism or virtue epistemology?

Let us start with the dichotomy internalism versus externalism. Epistemic internalism is generally understood as the idea that agents have internal access to the relations between their beliefs (Pollock and Cruz 1999, 24) and that the epistemic status of an agent "depends wholly on matters which are 'internal' to that thinker, rather than at least partially on matters which are 'external' to her, such as her relations to her environment" (Brown 2007, 14). This general direction includes theories such as foundationalist theories: dogmatism, coherence theories or direct realism (Pollock and Cruz 1999). By contrast, externalism is defined as the denial of this thesis, as externalists take that justification depends partly on matters which are external to the agent (Brown 2007).

Justification acquired through affective experience could be understood both under an internalist and an externalist conception, depending on which theories within the dichotomy one ascribes to. For instance, it could be a form of phenomenal

conservatism.[16] Phenomenal conservatism is the idea that justification depends on seemings (Moretti 2015). As Moretti explains, "the central intuition of the phenomenal conservative is that one should grant that things are the way they appear to be unless one has reasons for doubting it" (Moretti 2015, 296). This is an internalist theory, because the justification based on the seeming is dependent only on the mental states of the agent. The Reflective Model of affective experience could be spelled out in these terms. Indeed, we could say that affective experience is a form of seeming. The currently most popular view in phenomenal conservatism is to take seemings to be experiences, where seemings have a distinctive phenomenology which presents the truth of a proposition (Moretti 2015). Given that affective experience presents agents with evaluation of the environment, it is easy to see how they could count as seemings. If one is already a phenomenal conservatist, they could easily incorporate affective experience as a way of obtaining justification. Brogaard proposes a version of phenomenal conservatism which incorporates memory seemings in addition to other types of seemings to explain how inferential knowledge is possible. In these cases, memory seemings help connect two experiences which are stored separately in memory (Brogaard 2018). When someone uses affective experience, they can have two types of seemings which help them form their belief: the seeming provided by their affective experience, and the memory seeming which links this experience with previous ones. This explains how affective experience can track reasons, as it is helped by memory.

Going the internalist route is not the only option available to include affective experience in our theories of justification. Another sensible option would be to ascribe to reliabilism. Reliabilism is the idea that "a belief is justified if and only if it is produced by a reliable psychological process, meaning a process that produces a high proportion of true beliefs" (Goldman 1998). One core aspect of reliabilism is that the agent doesn't have to know whether the belief was reliably produced for it to be justified (Goldman

[16] One should note that phenomenal conservatism is also sometimes called phenomenal dogmatism. These two terms refer to the same theory.

1998). This feature is what makes reliabilism an externalist theory of justification. The Reflective Model of affective experience can easily be fleshed out in reliabilist terms: a type *t* of affective experience is reliable to produce true beliefs. Here the reliability of various types of affective experience can be evaluated through the reflective features. However, it isn't necessary for the agent to know whether it is reliable or not. The reflective element is not necessary to make affective experience reliable. Including the reflective element doesn't add epistemic value in beliefs that are reflective compared to unreflective beliefs (Kornblith 2012).

This may seem odd, given that the reflective element of the Reflective Model is an important part of it. However, we should understand it in this way: the reflective element doesn't participate in the justification of the belief; rather it helps the agent know when their ability stops tracking reasons and thus when it should be abandoned as a belief forming mechanism. It is because the ability (here a certain type of affective experience) is reliable that it provides justification. The reflective element helps integrate new information into the system, rather than gain second-order beliefs about the first-order ones. We can make an analogy to Kornblith's example of the plover protecting its nest, who stops engaging in protective behaviour after repeated interactions with a particular individual (Kornblith 2012, 55). With the repeated interaction, the plover registers that no threat is present, and new information is added to its understanding of the environment. A human agent using their affective experience acts analogously: through repeated interactions and various forms of regulation, the new information about their environment is registered. This helps them calibrate their affective response in future similar situations if the environment is cooperative and allows for calibration.

However, one could wonder to what extent affective experience would reliably produce a high number of true beliefs. It is true that various types of affective experiences, like some types of fear or jealousy, tend to produce a high number of false beliefs. In these cases, these affective states would not count as reliable processes. However, there are other affective states which, if used in specific contexts, tend to have high reliability.

Lordean rage, a type of anger with targets racial injustice, is one of these (Cherry 2021). Thus, not all affective experience can provide knowledge, but some types of it can. One could nonetheless wonder: how can we distinguish which types of anger provide knowledge from the types of anger which don't? How can an agent know that they should trust their Lordean rage, but not their other forms of anger?

There are two ways to answer this question. First, many affective experiences are sorted into broad categories, like pain, anger, sadness, fear, etc. due to the broader formal object they respond to, like bodily harm, offensiveness, loss or danger. However, if we can specify them based on a more fine-grained formal object and their action tendency, we can differentiate them. For instance, both Lordean rage and rogue rage respond to injustice, but they are very different when it comes to the target of the rage and the action tendency. With Lordean rage, the target of the rage is perpetrators of racial injustice, and the action tendency is to channel the power of the anger to fight injustice (Cherry 2021, 24). Rogue rage on the contrary is directed towards the world, as the world appears to be responsible for the injustice. Its action tendency is to retaliate against the world, often through violence, to find relief even though it is not seeking resolution (Cherry 2021, 17). In both cases, we can see how both types of rage are very different, due to their different targets and action tendencies. The same form of reasoning applies to other forms of affective experience. I know my pelvic pain signals a problem in my uterus when all of my periods have been painful, especially if the pain has lasted for years. However, the chest pain I feel when breathing might be caused by COVID-19 or stress, but I cannot trust it to tell that I am sick. In this latter type of chest pain, it cannot provide me with the knowledge that something is truly wrong with my body, as it might not be reliable. On the contrary, my pelvic pain can be trusted to provide me with justification that something is wrong with my body. Both experiences are pain, but they can be individuated in types, which have varying degrees of reliability, which means some should be counted as sources of justification, and others would not.

Finally, one could put aside the internalist/externalist debate and go down the route of virtue epistemology. By shifting away from belief-based epistemology, the virtue

epistemologist characterises belief evaluations in terms of intellectual virtues and vices (Battaly 2008). For Zagzebski, a justified belief is "what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, might believe in like circumstances" (1996, 241). Intellectual virtues can be understood as reliable faculties by virtue-reliabilists or as states of character by virtue-responsibilists (Battaly 2008). For virtue reliabilism, "virtuous thinkers are reliable truth-producers" (Battaly 2008, 645). For virtue-responsibilism, "what makes an agent an excellent thinker are active features of her agency" (Battaly 2008, 648). Both accounts can be seen as complementary, rather than competing, as they highlight different qualities of good thinkers: they can reliably get to the truth and act appropriately, and care about getting the truth (Battaly 2008). The Reflective Model could be understood under this light: agents who engage in reflective behaviour care about tracking reasons properly; moreover, they are reliable truth-producers, as they would stop trusting their affective experience was it to stop tracking reasons accurately.

The Reflective Model of affective experience can accommodate many different types of intuitions about what justification and knowledge are. I do not want to subscribe specifically to any one of these theories, as my goal is primarily to answer a wider question about whether affective experience has direct epistemic value and should count as a source of knowledge. Moreover, I do not want to choose which one of the internalist, externalist or virtue epistemology routes is the best, because this means that these different types of epistemologists have the potential to be convinced of the significance of affective experience for justification and knowledge. As we know, all these theories are highly contested, and each has their supporters for different reasons (Pollock and Cruz 1999). To show that my theory is compatible with various of these options appears is the best way to show that affective experience is a source of knowledge, full stop. Its status as a source is not dependent on a particular theory of justification, in a similar fashion to the other sources like perception, memory or reasoning.

5. Conclusion

By shifting to a conception of agent as reason-trackers, where rational guidance only requires that the agent monitors the system so that it continues to track reasons reliably, affective experience can provide knowledge. Previous discussions about affective experience assumed a model where agents had to respond to reasons as reasons with their affective experience. However, this led to the various problems of insensitivity to defeaters or conflict without contradiction. By conceptualising responsiveness to affective experience as a reason-tracking ability which incorporates new information through regulation, we can explain how affective experience may provide us with knowledge or justification. Affective experience gives us justification if reflective agents would stop using it were it to fail to track reasons appropriately. Given the many ways affective experience is regulated, agents are likely to evaluate whether their affective experiences fail to track reasons adequately. The Reflective Model of affective experience incorporates an aspect of affective experience which had not been incorporated previously: we regulate our affective experience both internally and socially. Affective experiences happen in a deeply social world, which is both a feature that helps us regulate it for good and for bad. Importantly, affective experience is a fallible source of knowledge. Not all types of affective experiences should be trusted to give justification, like phobias, biases, or some negative emotions like jealousy. Moreover, belonging to certain groups can calibrate affective experience in a misleading way, as seen with the anti-vaccine movement. However, in some cases, it provides us with the best ways of knowing, and excluding it can have dire consequences. This last idea is what the rest of this thesis will focus on, through the theorisation of a new type of epistemic injustice, and the discussion of applied cases where affective experience plays a crucial role in the acquisition of knowledge.

Chapter Three:
The Emotional Account of Microaggressions

I have suckled the wolf's lip of anger
and I have used it for illumination, laughter, protection, fire in places where
there was no light, no food, no sisters, no quarter.[17]

Audre Lorde

Introduction

Imagine being a same-sex couple who gets stared at for holding hands in public[18]. While this situation might seem unimportant at first and not worthy of any discussion, it is not. Many couples hold hands in public, but only a small portion of them gets stares from strangers. Unsurprisingly, the couples who often get stared at for public display of affection are those who do not fit the traditional idea of the heterosexual couple. Being stared at in this situation makes the individual feel that others see their relationship as an odd one, or worse, that strangers subtly signal that the gay person should not be holding the hand of their partner in public. Moreover, the individual might start to feel uneasy at the situation and it might discourage them from holding hands with their partner in future similar situations.

Staring at a same-sex couple is a form of oppressive action. As Pierce explains, "most offensive actions are not gross and crippling. They are subtle and stunning" (1970, 265-266). This idea developed by Pierce, a Black psychiatrist studying the offensive tactics used by white people against Black people, constitutes the original account of microaggressions. Microaggressions are opposed to "gross, dramatic, obvious macro-aggressions such as lynching" (1970, 266) and "the complications they cause can be

[17] (Lorde 1984, 126)

[18] Some of the text from the introduction and section 1 appears in a paper I submitted for the course *PHIL 6470 Topics in Applied Ethics: Microaggressions and Moral Responsibility* at York University (Toronto, Canada) in 2017, when I first began thinking about these issues.

appreciated only when one considers that these subtle blows are delivered incessantly. Even though any single negotiation of offence can in justice be considered of itself to be relatively innocuous" (1970, 266). The sum of collective microaggressions therefore paves the way to oppression and constitutes an effective strategic manoeuvre to maintain power dynamics. Going back to the introductory case, a singular instance of being stared at while holding hands looks innocuous. However, if one is being stared at every time they hold hands with their partner it acts as a form of policing, preventing same-sex couples from being visible in the public space. This microaggression, when added to many others, participates in the oppression of same-sex couples, which ultimately opens the door to macro-aggression like insults, physical aggression or murder.

More recently, Sue, one of the main contributors to the topic, defined microaggressions as "brief, everyday exchanges that send denigrating messages to certain individuals because of their group membership" (Sue 2010, xvi). Those exchanges are often done without the intention from the aggressor to send a denigrating message nor to harm the other person. Even worse, some microaggressions are done with good intentions, just like when students of colour are complimented for being "so articulate" by their teachers (Sue et al. 2007, 276). The complicated nature of microaggressions has prompted a lot of discussion recently (Campbell and Manning 2016; Dover 2016; Fatima 2017; Freeman and Stewart 2018; Freeman and Weekes Schroer 2020; Friedlaender 2018; Lilienfeld 2017; McClure 2019; McClure and Rini 2020; McTernan 2018; Rini 2018, 2022). However, despite increased interest in this topic, one important question remains unsatisfactorily answered: how do we determine that a microaggression has occurred? This question can be read in two ways: first, as an ontological inquiry, about what microaggressions are; second, as an epistemological query, about how agents know that microaggressions happen. When two people argue about whether a microaggression has occurred, they usually do so on one of two different grounds: they disagree about the nature of microaggressions and/or one of the agents lacks the knowledge that the microaggression happened, while the other person has such knowledge. Moreover, two agents might agree on the ontological definition, but

disagree on the occurrence of a microaggression. Thus, a successful theory of microaggression should answer both questions, in order to provide a framework that aims at settling disputes between agents in the real world.

In a recent paper surveying the existing literature, McClure and Rini identified three main types of definitions: psychological accounts, where microaggressions are defined "based upon an (unconscious) mental state in the microaggression perpetrator, as a causal antecedent of the act" (2020, 3); experiential accounts defining "microaggressions based upon the phenomenological mental state in the microaggression victim, as a causal consequence of the act"(2020, 4); and finally, structural accounts, where microaggressions are defined "based upon their functional role within an oppressive social system, with multiple possible causal realisers" (2020, 5). In the same paper, shortcomings of the different options are presented, and they conclude that future research should go forward by aiming to hybridise an experiential account with either a psychological account or a structural account, and keep their most promising features (2020, 7).

This chapter aims to offer such a hybrid alternative, by drawing from the main experiential account available, defended by Sue (2010), and building upon existing structural accounts by Friedlaender (2018) and McTernan (2018). Sue's account relies only on the subjective experience of the microaggressed agent, whereas Friedlaender's and McTernan's accounts tie the exchange to a form of structural oppression. Section 1 demonstrates the main shortcomings and good aspects of those various accounts. In this discussion, I present two main challenges to conceptualising microaggressions based on Lilienfeld (2017): first a good definition has to address the possibility of victims overreacting and misidentifying oppression; second a good definition has to be applicable at the level of singular events and from the perspective of individual agents. My main point will be to highlight that while Sue's account might respect the importance of the subjective experience, it invites worries about overreaction and the growth of victimhood culture, thus failing the first challenge; structural accounts avoid these problems, as they include objective components. However, their focus is not explicitly

including the role that subjective experience plays in the way in which victims know about microaggressions, making it difficult to apply those accounts at the level of an individual event, thus failing the second challenge. Building on this discussion, I argue in section 2 for a new hybrid definition that avoids the same difficulties: the Emotional Account of microaggressions, where the warranted emotions of the target of the action should be taken into account to determine whether a microaggression has occurred. The Emotional Account is the missing link between an analysis of microaggressions from an individual agent's perspective and the group level knowledge of such a phenomenon. Indeed, it constitutes the first account of microaggressions which explicitly frames the dynamics between those two levels of analysis and shows how they influence each other. Finally, the Emotional Account is put to the test, to show how it ties in all the aspects that a convincing account of microaggressions should exhibit.

1. Existing Accounts

1.1 Sue's Experiential Account

According to Sue et al. (2007), it is the agent towards whom the action is targeted who determines whether this action was a microaggression or not. They defend this experiential account because social psychological research shows that many people belonging to dominant groups have biases, often implicit (Jones 1997; Keltner and Robinson 1996). Those biases, which target members of socially stigmatised groups, like people of colour, women or LGBT people are outside of someone's conscious awareness and control, and act as a mechanism to sustain inequalities between different social groups. Unfortunately, many people in position of privilege are unaware that they have such biases in the first place and that those biases impact their social behaviour (Brownstein and Saul 2016). Therefore the most accurate assessment of whether a microaggression has occurred in a particular situation is "most likely to be made by those most disempowered rather than by those who enjoy the privileges of power" (Sue et al. 2007, 278). For racial-based microaggressions, Sue et al. explain that "people of colour rely heavily on experiential reality that is contextual in nature and involves life experiences from a variety of situations" (2007, 278). This experiential account

emphasises the role of the experience of the real world person who suffers from microaggressions, as they seem to be the best placed to evaluate the situation. People from the majority group, being those who generally gain power from social structures, tend not to see the repetitive pattern of those negative brief exchanges and why those brief exchanges can be harmful because those mechanisms participate in maintaining the superiority their social group over another, thus being in their interest (Mills 2007, 35). Another advantage of this experiential account is that when the publicly available facts are insufficient to tell exactly what happened, then the experiential account can be applied as it relies solely on the experience of the agent (Sue et al. 2007).

As Lilienfeld, who critiques this experiential account, puts it, we can describe the experiential approach as the view that "microaggressions necessarily lie in the eye of the beholder" (2017, 141). However, it is this exact feature of Sue's experiential account that makes it prone to worries about overreaction on the part of microaggressed people and the possible emergence of a victimhood culture. As Lilienfeld explains, seeing oneself as a victim, might lead to a malevolent interpretation of an exchange, instead of treating it as a benign exchange (2017, 157). This general worry has been widely discussed, especially in the context of microaggressions and freedom of speech on university campuses (Campbell and Manning 2014). If people who claim to be microaggressed are prone to see themselves as victims, even when that is not justified, perhaps the concept of microaggressions should be abandoned. I will call this worry the misidentification challenge.

The misidentification challenge highlights the first problem of Sue's theory: the theory sets itself up for failure regarding this challenge because it does not answer the ontological question in a satisfactory way. If anything that is interpreted as sending a negative message due to one's group membership can count as a microaggression, then the concept does not characterise a useful category of phenomenon. For instance, it can even include cases that ought not to be included as microaggressions, mostly for the reason that the vast majority of people from the group targeted would not see it as one.

Consider the case of a small minority of Christians feeling microaggressed when people say “Happy Holidays” instead of “Merry Christmas” during the holiday season. They see it as sending a denigrating message. If the experiential account is right, then this counts as a microaggression. However, it does not seem that those Christians are microaggressed, as it does not seem that people saying “Happy Holidays” are sending a denigrating message. But dismissing the case of the Christians because of the lack of a denigrating message would be to lose some advantages of the experiential account. In fact, no one sends a denigrating message intentionally in most microaggressions, and the experiential account is meant to get around this by saying that it does not matter what message the microaggressor meant to communicate (Lilienfeld 2017, 147). But this generates a problem, as we don't want to count those Christians as being the target of a microaggression, as it seems that microaggressions are not simply any harm targeting someone because of their group membership but are somehow related to structural oppression. Therefore, the experiential account as Sue presents it cannot address the misidentification challenge and is prone to putting the label of microaggressions on cases that might not be microaggressions, in addition to failing to properly frame what makes an exchange a microaggression and not simply an instance of harm, harassment or bullying.

The second problem for Sue’s account is a reverse of the misidentification challenge: the definition fails to identify some genuine instances of microaggression, because the victim fails to interpret the situation as a microaggression. For example, there is a pattern of actions of people wanting to touch Black women’s hair (Cut 2016). Most Black women experience this as a microaggression. However, in a video focusing on this issue (Cut 2016), one woman reports being happy to have her hair touched[19]. In the same video, the women who feel microaggressed explain their position by saying that it is

[19] The video features 21 Black Women. In the video one woman looks happy and she says that "most people say no, but you can touch mine, hey." She then describes having many hairstyles and that her hair is for her "like an accessory". See the video between 1.46 and 2.03 minutes for her testimony (Cut 2016).

disrespectful of their body integrity and that it exoticizes their body. This pattern is actually a historical one and the small act of the person attempting to touch their hair does not respect their body autonomy. When a Black woman does not see this pattern and judges that she is not microaggressed, do we want to say that she was not microaggressed? It would seem clear to other members of her group who experience the same thing that she was microaggressed, even though she disagrees with them. Considering cases like this, Lilienfeld points out that the experiential account can lead to "potential logical contradictions" (2017, 143). He gives a similar example: "if Minority Group Member A interprets an ambiguous statement directed toward her – such as "I realise that you didn't have the same educational opportunities as most whites, so I can understand why the first year of college has been challenging for you" – as patronising or indirectly hostile, whereas Minority Group Member B interprets it as supportive or helpful, should it be classified as a microaggression?" (Lilienfeld 2017, 143). To avoid this logical contradiction, we must advocate for a form of relativism regarding microaggressions. In that example, if we relativise it to the subject, then Member A was microaggressed, whereas Member B was not. However, by doing so, the concept becomes of little use to identify microaggressions in a systematic manner. A non-relative concept is needed, if one wants to be able to talk about the moral response to microaggressions, responsibility attribution or blame. There is something unattractive about being lucky whether one acts wrongly or not[20], for instance by asking to touch the hair of the only Black woman in the video does not consider it to be a microaggression. Moreover, having such a relative concept cannot serve the goals of theorists like Friedlaender (2018) or McTernan (2018). Both argue for a systematic definition of the concept to serve normative goals in order to propose global strategies to reduce the occurrences of such a harmful phenomenon. Finally, if one wants to look into empirical research of microaggressions and develop methodologies to measures the

[20] I recognize that moral luck is a complicated topic, which would require much more discussion than this chapter can provide. Nevertheless, the kind of luck that this sort of relativism builds into the account should be avoided for the other reasons I will mention in the rest of this paragraph.

phenomenon, a concept with clear boundaries over the set of relevant cases is needed (Lilienfeld 2017).

As this discussion shows, the experiential account of microaggressions is insufficient to address the misidentification challenge and its reverse, because it invites relativism and overreaction, due to a wrong ontology of microaggressions. However, instead of suggesting to stop using the concept in real-world applications due to those problems, like Lilienfeld does (2017, 163), one could take microaggressions seriously by describing them in terms of objective features, thus shedding light on the difference between microaggressions and other instances of everyday harmful exchanges, like bullying or harassment.

1.2 Structural Accounts

In this section, I propose to discuss the two main structural accounts currently available, proposed by Friedlaender (2018) and McTernan (2018) respectively. These accounts aim to conceptualise microaggressions through some objective characteristics.

Friedlaender explains that "an act is a microaggression if and only if we can establish a link between the act and an objectively existing form of structural oppression" (2018, 8). To explain the notion of structural oppression, Friedlaender uses Cudd's account of oppression. According to Cudd "Oppression names a circumstance in which four conditions are satisfied:(1) *The harm condition*: there is a harm that comes out of an institutional practice; (2) *The social group condition*: the harm is perpetrated through a social institution or practice on a social group whose identity exists apart from the oppressive harm in (1); (3) *The privilege condition*: there is another social group that benefits from the institutional practice in (1); (4) *The coercion condition*: there is unjustified coercion or force that brings about the harm. These conditions, I claim, are jointly necessary and sufficient for oppression" (2006, 25).

Similarly, McTernan explains that "microaggressions […] are those minor or subtle details of interactions between people that fulfil a certain role: namely, functioning as a particular kind of degradation or putdown while appearing innocuous and plausibly

unintentional" (2018, 266). The particular kind of putdown present in microaggressions is characterised by two features: first it is committed by those belonging to the dominant group with respect to the attribute to which the microaggression calls attention; second, it can only occur towards members of subordinated groups: those that are oppressed or marginalised (McTernan 2018, 265). McTernan defines a 'subordinated group' as one whose members face systemic injustice, where either structures or exercises of power by others reduce their chances or exclude them, following the definition found in Iris Marion Young's work (McTernan 2018; Young 1988).

Both accounts put an emphasis on structural oppression and are driven by normative concerns. In both cases, the definition of microaggressions does not rely on the experience of the target, but rather on linking the particular act to a pattern of oppression. In both cases, oppression is understood as institutional discrimination, where one dominant group benefits from another group being oppressed (Friedlaender 2018; McTernan 2018). However, while Friedlaender accepts that members of oppressed groups can commit microaggressions (Friedlaender 2018, 19), McTernan (2018) allows only members of dominant groups to be perpetrators. By focusing on the pattern, singular microaggressions are defined through their collective functional role within a structure to maintain inequality (McClure and Rini 2020). With structural accounts, the focus shifts from individual acts to a collective appraisal of the phenomenon.

Structural accounts succeed in addressing the misidentification challenge, because they provide a clear ontological definition of microaggressions, in which the nature of microaggression is clearly linked to structural oppression, and not just any form of denigrating everyday exchange. This highlights the difference between this phenomenon and other similar harmful everyday exchanges like bullying or harassment. One should note that both Friedlaender and McTernan provided their definition in order to be able to talk about the harm of microaggressions and did not have as a primary goal to answer both the ontological and epistemic queries. Nonetheless, they offer a significant alternative to Sue's account, which is why I examine their positions in detail.

For instance, in the case of the Black woman who doesn't think that she is microaggressed, we can nonetheless say that the act of someone asking her to touch her hair is microaggressive as it is linked to the structural oppression of people of colour and acts as a putdown. In addition, in the case of the Christians who feel microaggressed during the holiday period, the act of saying 'Happy Holidays' can be ruled out of the microaggression category, as it is not linked to a pattern of structural oppression and does not consist in a putdown. Thus, structural accounts accurately capture what microaggressions are.

Structural accounts, by preventing the misidentification of microaggressions, as well as relativism in determining whether an act was microaggressive, are better than the experiential account defended by Sue (2010). Nonetheless, even though structural accounts provide answers to both the ontological and epistemological queries of microaggressions, they do not provide a full story about how agents come to know about oppression when it comes to particular instances of microaggressions.

In order to make this criticism clear, here are two examples. The first one is the example of Saba Fatima, a professor of philosophy and an immigrant woman of colour, which I already introduced in chapter 1 of this thesis. When Fatima tells her colleagues about one of her students who is disrespectful towards her, they say to her that they "all get disgruntled students" (Fatima 2017, 153). Fatima thinks that this disrespect is due to the general disrespect people have towards immigrants. However, all her colleagues think that this student is just another annoying student like they all get, and that this has nothing to do with the fact that Fatima is an immigrant (Fatima 2017). The colleagues deny that the behaviour of the student is a microaggression because they deny that this particular situation is linked to the structural oppression of immigrants, even if they agree with Fatima on the definition of microaggressions.

The second example is a common one for individuals assumed to be women who date other people read as women (abbreviated WDW).[21] Imagine two lesbians out on a date and they meet someone who mistakes them as being best friends. This is an instance of the 'Gal Pal'[22] phenomenon, which describes the tendency to think that two women who are dating are actually just close friends (McBean 2016, 2018). In this case, the lesbians were not showing obvious romantic affection towards each other when the mistake was made, therefore making it difficult to link this singular case to the systematic oppression of WDW, and homophobia in general. Here, only a WDW who has had that experience might recognise that the Gal Pal phenomenon exists and is a form of homophobia. Indeed, she knows that if she had been with a man, most observers would have assumed that they were dating, even without an obvious display of affection. As Sue et al. say, people belonging to majority groups tend to evaluate their behaviours through a singular event, therefore failing to see a pattern of bias that cause harm when those small acts are reiterated numerous times (2007, 278). And the Gal Pal phenomenon is a clear example of that, as many heterosexual people don't see why it's a problem when they automatically assume that two women are friends. In many cases of subtle oppression, lived experience gives some epistemic insight about the existence of the structural oppression, thus making lived experience important.

Fatima and WDW have their experience and expertise doubted, because it is hard to prove that the exchange that has just happened is not just a one-off disgruntled student or another benign mischaracterisation of one's relationship. Moreover, they are not the

[21] This category usually includes lesbians, bisexual women or feminine presenting individuals who date other people assumed to be women by an external observer. The acronym WDW stands for 'women dating women', where 'woman' is understood as 'assumed to be a woman by an external observer' independently of the gender identity of the individual.

[22] See McBean (2016) for a discussion of the phenomenon with regards to the media coverage of celebrities. For an example of the Gal Pal phenomenon, see the article by Courtney Enlow (2015) on the media's reaction to Kristen Stewart and her girlfriend.

only ones that have such things happen to them. How many women have been told that the joke they thought was sexist was just a harmless joke and has nothing to do with sexism? How many Black people have been told that it is a compliment if people want to touch their hair, and that it is not related at all to racist oppression? Unfortunately, Fatima's case and the WDW's case are not exceptional in their nature. In general, structural accounts can invite such minimising and denying responses, if the different agents involved in the discussion disagree about the particular instance being a microaggression, because those accounts do not focus enough on the critical role that the experience of the victim plays in linking the particular act to a larger pattern of oppression. On both structural accounts, the agent would need to be able to link what is currently happening to them to other similar past instances, in order to explain how they know that this particular case is a microaggression. However, structural accounts, by focusing on structural explanations fail to explain the epistemological mechanisms behind the recognition of individual instances by individual agents and how the level of collective appraisal of microaggressions hinges on the individual level. For this reason, they fail to respond to the second challenge, because they do not clearly explain how a singular agent identifies a singular instance of microaggression.

This is where the experiential account had some good insight into the phenomenon: the lived experience of the victims, as agents who have to go through oppression on a daily basis is epistemically valuable in itself, because it gives priority to such lived experience in the assessment of microaggressions. The experience itself is important because it can shed light on the existence of oppressive structures and the link that those structures have to particular instances. As Frye explains, "it is now possible to grasp one of the reasons why oppression can be hard to see and recognise: one can study the elements of an oppressive structure with great care and some good will without seeing the structure as a whole, and hence without seeing or being able to understand that one is looking at a cage and that there are people who are caged, whose motion and mobility are restricted, whose lives are shaped and reduced" (1983, 5). Microaggressions, due to their nature as subtle instances of structural oppression, are like those individual wires: by looking at them, one might not see that they are all part of a bigger structure of

oppression that resembles a cage. Nonetheless, when the cage traps a victim, then it becomes clear to them that all those individual wires are part of the bigger structure, and this knowledge is acquired through their lived experience as a person being trapped by said structure.

As we can see, structural accounts prevent labelling cases that are not oppressive harms as microaggressions, and does not lead to relativism, resulting in an ontological definition of microaggressions which people use in normative discussions. However, they do not give a full picture about how agents come to know about the link between the exchange and the structure of oppression, thus failing to explain how the collective level of appraisal is built from the individual experience of microaggressions. Lived experience can help fill this explanatory gap. This leads me to my own account, which aims to hybridise structural definitions with a central idea in Sue's theory: whether an action is a microaggression is dependent on the experience of the victim. To do this, I propose to use warranted emotions.

2. The Emotional Account

2.1 Warranted Emotions and Microaggressions

As shown in the previous section, lived experience plays a central part in the epistemic access to microaggressions. One component of this experiential reality is the affective response. In the case of microaggressions, people feel anger or indignation at the structural oppression they are facing (Cherry 2021, 50; Sue et al. 2007, 279). In many microaggressive scenarios, the victim has an affective response, which is the reason I want to explore it as a new way of determining if a microaggression has occurred.

Emotions have been described in many ways (Deonna and Teroni 2012). One should note that the theory of microaggressions to come is consistent with any theory of emotion, as long as it views emotions as having some intentionality (emotions are about something in the world) and that the formal object of emotions are values instantiated by particular objects. Emotions can give us some information about the evaluative characteristics of a situation. For instance, fear can inform us about the dangerousness

of a given situation, like the fear that I experience when standing in front of a tiger. The same goes for anger, which can highlight the injustice of a situation. How might we use that to give an account of the occurrence of microaggressions? The simple answer to this question is that a microaggression has occurred when one has a warranted emotional response that highlights the structurally oppressive character of the situation. The important point here is the notion of warrant, which makes all the difference with Sue's account.

Emotions have what is called cognitive rationality, which is their ability to give an accurate representation of the world (Scarantino and de Sousa, 2018; de Sousa 1987). Cognitive rationality has three parts: fittingness, warrant and coherence. The fittingness corresponds to the emotion accurately representing a value that is present in the particular object that the emotion is directed towards and taking a corresponding shape and size (D'Arms and Jacobson 2000). For instance, Deepika's great sadness towards the death of her father is fitting, because the death of her father is a great loss, and loss is the value that is represented in sadness. Had Deepika been as sad in a situation where she had suffered no loss, her sadness would not have been fitting. One should note that the fittingness of an emotion is only concerned with the presence of certain relevant evaluative features, and not at all with whether it is morally right to feel this emotion. For instance, Deepika's sadness is fitting, even in cases where it would be morally objectionable to delve into those sad feelings (D'Arms and Jacobson 2000). An emotion is warranted if there is evidence for the presence of the value in the particular object. In this chapter, I will assume an externalist account of warrant, as it relies on evidence that comes from things in the world, rather than from the mental states themselves. For example, my disgust at my glass of milk is warranted when I see that the liquid that I just poured is green instead of white, which gives me relevant cues to think that the milk is rotten and therefore disgusting. Finally, the coherence of the emotion deals with the consistency of the emotion with other beliefs and representations of the world. For instance, one might believe that snakes are dangerous, and so it would be coherent to fear snakes. An emotion can have none, some or all of those dimensions. For instance, one might see a coral snake and be afraid of it. The fear is fitting if the snake is a real

coral snake and not a king snake, which looks eerily similar but is not venomous. Whether the snake is a coral snake or a king snake, the emotion is warranted, because one has some evidence by looking at the colour of the snake that it might be highly venomous. Finally, the emotion is coherent with the belief that coral snakes are a certain pattern of colour, and that those snakes are venomous. As I will explain, not all aspects of cognitive rationality are relevant to give an account of microaggressions. Indeed, only warrant will feature in this account. Thus, the initial thesis that I will now explore is that warranted emotions are what indicates that a microaggression has occurred:

Emotional Account: a microaggression has occurred when the agent who is the target of the brief everyday exchange has a warranted emotion towards the structural oppression present in that exchange.*

When a microaggression occurs, structural oppression is present, as Friedlaender and McTernan show. According to a lot of emotion theorists, the emotional response to injustice, of which structural oppression is a subtype, is anger or indignation[23] (Cherry 2021; De Sousa 1990; Neblett 1979; Nussbaum 2016). This fits perfectly with the psychological picture that Sue et al. described, when they explained that people get angry when they are microaggressed (2007). However, if we simply say that what counts for a microaggression is the presence of anger or indignation, the suggestion will fall into the same traps as Sue's experiential definition. Without any constraints, an emotional theory of microaggressions cannot give a better description of the phenomenon of microaggression than Sue's account can. This is why the Emotional Account* needs to appeal to the warrant dimension present in the cognitive rationality of emotions, in

[23] As one will notice by reading the rest of this paragraph, I slighly disagree with this claim, insofar as some other types of emotions, like contempt or amusement could also be emotions that respond to injustice and are warranted by the circumstances. Indeed, microaggressions provide some support to accept that the range of emotions responding to injustice can be broadened, as responding to continued and frequent microaggressions with contempt or resignation can be appropriate. The only important thing is that the emotions are warranted by the circumstances and are responding to the injustice present in the situation.

order to distinguish between emotions that respond to structural oppression, and those that don't. Emotions responding to structural oppression do not have to be anger or indignation in all cases, because it is possible to imagine that one might instead get amused at microaggressions, because they got tired of being angered at them, after years of suffering. Nonetheless, in all cases, the emotion needs to have the structural oppression of the situation as its trigger.

To be relevant for assessing whether a microaggression has occurred, an emotion needs to be externally warranted for the victim. The victim needs to have some form of evidence that the situation is structurally and oppressively unjust. Structural accounts' frameworks provide a good background for what constitutes the structurally and oppressive unjust features of a situation. In the rest of the chapter, I will appeal to Cudd's notion of structural oppression, like Friedlaender, as Cudd offers an individualist account of structural oppression. Thus, the Emotional Account* of microaggressions builds the notion of warrant from structural accounts, as the evidence required from the victim to warrant their emotion is for them to be able to establish a link between the structure of oppression and the particular situation. For instance, the setting of the situation can inform the agent of such a thing: a Black speaker is the only person of colour at a conference. When that person is mistaken for a member of the catering team, the setting gave them some evidence that this environment could promote some subtle racism. If the conference is taking place in North America, it is often the case that labour is racialised such that most academic professionals are white, while people of colour often do service work. The white participants might simply assume that the Black speaker is a member of the catering team because of their race in that context. There are many available clues to the agent who is a target of microaggressions, to be able to warrant their emotions. However, contrary to structural definitions, the Emotional Account* explains clearly how the agent knows that there is a link between the particular situation and the structural oppression: through emotions. Indeed, emotions provide the microaggressed agent with an epistemic access to that link on a regular basis, as emotions are responding to the oppressive features of the situation. This enables the Emotional Account* to incorporate lived experience to make it relevant for

the assessment of microaggressions, while still framing microaggressions as subtle instances of structural oppression, and not as any form of harm made to someone because of their group membership. Warranted emotions provide some additional, not publicly available evidence for the link between the individual act and the structural oppression at stake. Through warranted emotions, the agent is capable of making sense of the individual act within a collective framework.

The other cognitive rationality dimensions are not relevant for the current discussion about the emotional component of the Emotional Account, as this component meant to answer the epistemological question. First, whether the emotion of indignation is fitting depends on whether the situation is linked to structural oppression. Given that the exchange must be tied to structural oppression to be a microaggression, as described by the structural accounts, the fittingness is already built in the Emotional Account. Indeed, unfitting emotions of the relevant kind will be ones directed towards things that are not unjust.[24] As the ontological part of the Emotional Account already deals with assessing whether the situation is linked to structural oppression and sorts out between cases where structural oppression is present or not, the focus can shift to emphasising the epistemic role of emotions in the recognition of individual instances of structural oppression through the notion of warrant.

Second, having a warranted emotion should not be equated with having a coherent emotional response. Indeed, emotions experienced by the agent are sometimes in contradiction with other beliefs that they either acquire themselves or are forced to

[24] This means that the account allows for the possibility of agents having warranted emotions that are unfitting. Here is an example of such a situation : imagine mishearing a slur when one of your friends tells you, an Asian descendant, that « you are acting mellow today », and you hear « you are acting yellow today » instead, which makes you react with indignation. Here, the emotion is not fitting, as the sentence « you are acting mellow today » is not unjust under any aspect, and the indignation is the result of a mishearing. Such a case would not count as a microaggression, even though the agent has a warranted emotion, because the situation is not linked to a pattern of structural oppression.

acquire through being reasoned with by other agents, thus rendering them incoherent. For example, a woman might force herself or be forced to think that a sexist joke she has just been told is funny, because all her male colleagues are laughing at it, even though she had a negative affective response to the joke, and she just cannot shake the feeling that something is wrong with that joke. Even though her colleagues might influence her beliefs and change them, the content of her emotion can conflict with the set of beliefs that she currently holds. An important aspect of microaggressions is highlighted by this last example: sometimes, it is not clear from a rational point of view that a microaggression has occurred, even for the victim. All that there is, is the feeling that something is wrong, or unjust. One is angry, but without knowing exactly why. Those cases are unfortunately the majority when it comes to microaggressions, because microaggressions are subtle and nearly invisible. Thus, how could we adapt what has been said about the need for warrant, when evidence might be hidden?

Moreover, one could object that having the presence of the structural oppression already built in the Emotional Account makes the definition sound circular. Indeed, if someone discredits the evidence showing that the situation exhibits features of the unjust, on the basis that the agent is misperceiving the nature of the evidence, the agent ends up appealing to structural oppression. However, if the objector demands evidence to prove the existence of the structural oppression, the agent seems to be going back to their starting point, using the reasons that warrant their emotions in the first place. Thus, we need a way to stop this apparent circularity, by showing that emotions reliably track the evaluative features of the situation without relying on these features to be reliable.

The solution to both those problem lies within the nature of microaggressions themselves: repetition. It is the repetition of those feelings in similar situations which helps the agent gain an epistemic advantage, through various stages of calibration. This echoes what we saw in the last chapter when I presented the Reflective Model of affective experience. In the next section, I explain how emotions triggered by

microaggressions are calibrated and how these calibrated emotions help the agent gain an epistemic advantage over agents who are not microaggressed on a daily basis.

2.2 Being Microaggressed, Calibrated Emotions and Epistemic Advantage

In order to show how being regularly microaggressed could make one gain a unique epistemic access to microaggressions through their emotions, I will first borrow from Sauer's theory of educated intuitions. Sauer argues that moral judgements are based on educated intuitions, which are automated rational responses (2012). Sauer explains that "like any other kind of education, the education of (moral) intuitions is a process of habitualization. The level of automaticity with which an intuitive judgemental response is triggered increases with the number of repetitions, a process that consists in a "migration" of controlled and effortful cognitive processes into an agent's effortless, perception-like intuitive system" (2012, 266). As intuitions are usually considered to be closer to beliefs and more intellectualised than emotions, it seems more intuitive that they can be improved through learning. However, Sauer's model can be translated to emotions, given their features of cognitive rationality. Indeed, both intuitions and emotions have a common ground in that they respond to reasons. With emotions, agents can get better at recognising the reasons and respond to them, as I explained in more details in chapter 2. For instance, Mei has a friend Yumiko who always responds to her text messages the same day. However, those last days, Yumiko has not been responding to Mei's messages, thus prompting Mei to worry about whether she might have done something to anger Yumiko. Mei then remembers that this week Yumiko has an important deadline and is behind with her work. Given this reason, Mei stops worrying and waits until Yumiko gets back to her. The following week, Mei learns that she was right, as Yumiko finally answers, explaining that she was too stressed to respond to the messages she received. In this example, Mei engages in internal calibration, where she discards her affective experience after being presented with a defeater, here the important deadline.

Some people are exposed to structural oppression on a daily basis. If the epistemology of microaggressions framed throughout this chapter is right, then it is possible that

through their emotional responses, microaggressed people can gain a unique form of understanding of the phenomenon, due to their positioning and the emotional aspect of their experience. For example, Pierce explains that racism is "a perceptual illness. In almost any black-white negotiation each participant views things differently, depending on whether he is white (the offender) or black (the offended)" (1970, 267). Following Sauer, one can adapt the model of educated intuitions to this situation: as agents become more and more exposed to structural oppression, it triggers the transition from an effortful process to one that becomes intuitive, here the emotional response. As Sauer explains, "habits are behavioural patterns whose execution is triggered in certain circumstances and becomes automatic over time" (2012, 259).

Here, the habit of having an emotional response to structural injustice is one that becomes automatic in agents who are regularly microaggressed. This makes them more likely to pick up on features of the situation that are relevant to link this particular situation to a larger pattern of oppression. This access is not infallible, but it helps microaggressed people learn about the properties of the oppressive environment better than if they were not experiencing it as victims. As Sauer highlights, educating one's intuitions is done *ex post* and *ex ante*. *Ex post* education deals with the response, while *ex ante* education is concerned with the conditions under which the response is generated (2012). In the case of microaggressions, *ex post* education is crucial to understand why microaggressed people use their emotions as a reliable method of knowledge acquisition. According to Sauer, "*ex post* education is concerned with how an intuition, once it has been generated, can and should be dealt with, and how controlled after-the-fact reflection feeds back into an agent's intuitive system" (2012, 267) and it often results in an improvement of the system after reflection (2012). This idea is built on Kennett and Fine's work, who argue that a "closer examination of the interaction between automatic and controlled reflective processes in moral judgement […] makes room for the view […] that genuine moral judgements are those that are regulated or endorsed by reflection" (2009, 78). The idea of *ex post* education fits well with the reflective mechanisms of the Reflective Model of affective experience: agents

obtain more information through reflection and this information added to the system helps improve the accuracy of the reason-tracking.

Microaggressed people feed into this loop of *post hoc* deliberation after a microaggressive situation occurred, by often asking themselves whether their anger is warranted and whether their anger tracked reasons accurately, depending on whether the situation was truly an instance of subtle racism, sexism, homophobia or any other type of oppression. More generally, as an agent notices discrepancies between their intuitions and reasons that become available to them, it becomes effective in shaping their automatic responses (Sauer 2012), and it is likely that a rational agent would stop using their anger if they realised that it stopped tracking reasons, as explained in the Reflective Model of affective experience.

Finally, those educated intuitions come in different levels of sophistication, thus explaining the differences that can be found between different agents. The more one accumulates experiences of microaggressions, the better they become at connecting separate instances together (Cherry 2012; Fatima 2020). As Fatima explains, "as time passes, one may gain epistemic certainty about the nature of their experiences either because of critical reflection, and/or because they continue to encounter other similar microaggression, and/or because they hear about similar experiences of others" (2020, 167). On the contrary, a person who doesn't know what structural oppression is[25], for instance someone privileged who never experiences oppression, will not engage in critical reflection about the phenomenon, nor develop an ability to connect various events together. This last distinction is a crucial one: privileged people will not acquire the educated emotional response to microaggressions nor the ability to track the relevant reasons in this context, unless they go the extra-mile to understand the

[25] Unfortunately, some oppressed people might not be aware of structural oppression either. Oppression can impact the flourishing and development of oppressed people, causing them to fail to understand that they are oppressed (Young 1988). This hinders the process of education of their intuitions and prevents the process of habitualisation from developing.

underlying mechanisms. This effortful process makes privileged people less likely to recognise a microaggression, as they never need to have quick and effortless responses to such phenomena. Indeed, privilege grants one the opportunity of never having to develop a reason-tracking ability that targets oppression. For people suffering from structural oppression, developing such a mechanism is automated and not a choice, as it becomes "a necessity to survive" (Fatima 2020, 166).

Thus, the emotional experience, which takes the form of educated intuitions and is a reason-tracking ability, is key for understanding microaggressions and recognising them. It puts victims of microaggressions at an epistemic advantage compared to other agents, who do not suffer from the same structural oppression. The agent has a strong internal feeling that something is wrong, and that type of feeling is the same as the feelings they got in other microaggressive situations, which had clear evidence, thus helping them recognise that this particular situation exhibits similar features. This process is one that is educated and automated, and in which past experiences feed into the loop of *post hoc* deliberation, thus improving the habitualization and helping the agent realise whether their emotion stops tracking the relevant reasons. For instance, in the case of Fatima, her emotional response, developed through her experience as an immigrant woman of colour, gave her an epistemic advantage over her colleagues, as they could not realise that the situation was microaggressive, due to a lack of understanding and educated intuitions.

However, one could object that even though emotions might become calibrated in this sense, we still need to rule out the possibility that marginalised people's emotions react to something, but not necessarily structural oppression. Indeed, Lilienfeld worries that some people might be prone to negative emotionality, a tendency to focus on the negative aspects of life, which might shape how agents perceive microaggressions, thus making emotional responses not reliable to track structural oppression (Lilienfeld 2017). Lilienfeld supports this criticism by highlighting that the high internal consistency of

microaggressions measures[26] can be best explained by the influence of personality traits such as negative emotionality. The main problem seems to be that certain minority individuals are consistently recipients of microaggressions, while others are not (Lilienfeld 2017). However, this could easily be explained by the emotional experience model: some agents have a system that is more finely calibrated than others, due to a greater sophistication after repeated microaggression exposure. It is no wonder that some agents would acquire a more fine-tuned system the more they experience microaggressions, thus enabling them to recognise microaggressions in a wider range of settings.

Moreover, the social nature of microaggressions makes it often the case that microaggressed agents are forced to reflect upon the nature of what happened to them, thus increasing the opportunity for *post hoc* education and self-reflection. Indeed, agents from oppressed groups are constantly challenged and pushed to the edge of knowing, as Fatima explains (2017). Being challenged whether it is by a member of one's group or by someone in a position of privilege, pushes the agent to reflect on the reasons that warrant their emotions. This forces the agent to engage in the social calibration of their affective experience, whether it is through extrinsic regulation or collective resistant regulation[27]. When one has the opportunity to engage in shared and collective knowledge about structural oppression, it helps them critically reflect on their own

[26] Microaggression measures are self-reported measures made by participants to various studies about whether they were microaggressed or not. High internally consistent microaggression measures refer to the idea that some participants seem to be consistently targets of microaggressions, even in a wide range of situations (Lilienfeld 2017, 156). According to Lilienfeld, "upon reflection, however, these reliability values actually give cause for concern. Microaggressions are posited to comprise an extremely diverse class of slights, insults, and snubs of various sorts emanating from a diverse array of individuals (Sue et al., 2007). Thus, it is not at all clear why microaggression measures should be internally consistent" (Lilienfeld 2017, 156).

[27] See chapter 2, section 2.3 for a more detailed explanation of the various ways to calibrate affective experience.

experience (Fatima 2020) thanks to collective resistant regulation. Indeed, communal spaces to share one's experience as a member of an oppressed group are essential to understand power relations. As Collins explains, "Black Women's experiences occurred in communities whereby individual experiences were meaningless without some sort of collective analytical frame through which to interpret them"(2019, 187-188). Engaging critically with other members of one's community to reflect together about their experience helps to educate one's responses and participates in the *post hoc* education of emotional intuitions. This mechanism drives the self-reflective monitoring and helps the agent adjust their emotional system and educate their emotional intuitions, thus improving their habit to recognise microaggressions. This provides a form of corrective feedback system, enabling marginalised agents' emotions to track structural oppression and not just any negative input.

There is second opposite worry that might appear: agents are prone to suffer from the problem of attributional ambiguity, which is when targets are unclear about whether the microaggression has occurred, because microaggression become increasingly harder to detect, thus leading to a unique form of epistemic uncertainty (Friedlaender 2021). Even more worryingly, attributional ambiguity could "distort the target's perception, potentially leaving them with the belief that a microaggression might not have occurred" (Friedlaender 2021, 240). In cases of attributional ambiguity, the Emotional Account* can explain how agents can still accurately identify individual instances of microaggressions: as Friedlaender explains, structural accounts enable the agent to identify a microaggression through the similarity of the individual action to a recognised pattern of behaviours already labelled as microaggressions (Friedlaender 2021); in the same manner, the educated emotional intuition is making the link between the pattern and the individual action, thus enabling the agent to know through their emotions that they are microaggressed. Even when the agent is later the victim of epistemic gaslighting, a form of manipulation where the victim is led to believe that their belief is wrong, they would still have had a warranted emotional response to the situation in the first place, thus making it an instance of a microaggression. As said previously, the agent

is often left with the strong internal feeling that something is wrong, even when others try to convince them otherwise, as described by Fatima (2017).

According to the whole picture, the daily experience of microaggressions helps microaggressed people develop an automated educated emotional response to microaggressions, a process calibrated through a loop of corrective feedback induced by being challenged and reflecting on one's reasons warranting one's emotion. This gives them an epistemic advantage, thus giving them a better understanding of microaggressive situations. By participating in collective resistant regulation over their individual experiences, individual agents also add to the collective resistant knowledge project about structural oppression[28]. Let us now put the Emotional Account* to the test and see if it fares better than the other two competing theories.

2.3 The Emotional Account* Tested

There were four cases that highlighted the shortcomings of Sue's experiential account and of structural accounts. Let us now explore whether the Emotional Account* does better.

Remember the Christians who were deeply offended by people saying "Happy Holidays" instead of "Merry Christmas", and thought that this was a microaggression. Can the Emotional Account* exclude cases like this? The first component is whether the Christians had an emotion, like indignation, towards what had been said to them. The answer is positive. However, is that emotion warranted towards a structural oppression? No, because there is no evidence that something structurally oppressive is happening in that case. Had the Christians been part of a religious minority that is

[28] Other theories of microaggressions, like the ones discussed in McClure and Rini (2020) or Freeman and Stewart (2020) incorporate standpoint epistemology. However, those theories are focused on group knowledge and fail to point out the role of individuals in the formation of such group knowledge, and that the mechanisms of undertaking the critical process individually and with a community form an interactive feedback loop.

persecuted, then it would have been a microaggression. But in a European or North American society, it is not the case. Thus, this case is not a microaggression, even though it might still harm those Christians.

When examining the case of the hair, the Emotional Account* doesn't seem to be able to deal with it, as the problem is the absence of indignation from the woman, which would make this case not a microaggression, but a microaggression towards other women who feel indignant when their hair is touched. Thus, relativism is present, just as with Sue's experiential account. Moreover, another case might arise, when the agent has the emotion and all the evidence to warrant it, but nonetheless, they do not think that their emotion is warranted because they are in a bad epistemic state, by for instance being conditioned to think that their indignation is unexplainable and not linked at all to structural oppression, which prevented the development of automated educated emotional responses. In that case, they don't have a warranted emotion, because oppression inhibited their flourishing, which prevents them from gaining the more thorough understanding exhibited by microaggressed people. In that case, it seems that both the ontological and epistemic questions cannot be answered satisfactorily by the Emotional Account*. This leads to a revision of the Emotional Account*:

Emotional Account: a microaggression has occurred when the agent who is the target of the brief everyday exchange could have had a warranted emotion towards the structural oppression present in that exchange.

To account for cases like the one of the hair or the oppressed person, a counterfactual is needed. Here the counterfactual is of a special kind, where it is not the facts that change, but the perception that the agent has of the same facts. Thus the conditions to explain that the agent could have a warranted emotion are the following: first, the person has an emotion this time, because structural oppression is there and this time they have a better understanding of their environment, and their affective experience responds to such an evaluative feature of the environment; second, the person is in an improved epistemic state, such that this time they do not dismiss the evidence that

enables them to warrant their emotion. I do not want to say that the agent should have had a warranted emotional response, because unfortunately, oppression works in a way that renders some agents unable to see that they are in an oppressive situation. As Frye puts it, oppression is a bird cage, in which you cannot see all the wires at the same time and thus, cannot comprehend why the bird is trapped, unless you look at the cage in its entirety (1983, 4). Some people might never understand that they are birds, trapped in a cage. Thus, a good account of microaggression should not demand that a specific agent needs to have a warranted emotional response for a microaggression to occur, but rather that an emotional response could be warranted in that situation given the evidence. A microaggression, because it derives from structural oppression, has the power to elicit an emotional response, thus making it the object of indignation or anger, and the emotion helps the agent gain knowledge about the existence of the microaggression.

Now, let us turn to the counter cases that structural accounts faced, and see if the Emotional Account can explain them. Here the Emotional Account gives a clear answer. Fatima had a warranted emotional response to what happened and therefore, the episode was a microaggression. She felt anger, and the evidence she had was that she was facing a privileged white male student full of prejudice, while she is an immigrant woman of colour, thus letting her know that the treatment she received is unjust. In the Gal Pal phenomenon example, the WDW would have an emotional response that will be warranted either by some evidence, because of the response of the interlocutor, or the general idea that girlfriends are just that, girls and friends. The Gal Pal phenomenon is a form of microaggression, with many instances, some of which are obvious, like the case of celebrities Ashley Benson and Cara Delevingne being described as having a “fun friendship ceremony” when they decided to get married in Vegas, while not signing any legal document (Kiley 2019). Another celebrity couple, with partners of different genders, would not have been described as having a friendship ceremony, had they done the same thing. It is those less ambiguous cases that enable members of the targeted group, here WDW, to learn about the pattern of structural oppression and develop calibrated emotions through calibration mechanisms. In the case I presented

about the girlfriends being said to be best friends, their automated educated emotional intuitions put them at an epistemic advantage to recognise the evidence, which might not be obvious and available to people who are not targeted by microaggressions of the Gal Pal type.

After putting the Emotional Account to the test, it appears that its advantages over structural accounts will be epistemic. The Emotional Account provides an explanation about how agents know that a particular instance is a microaggression and how they can link it to a pattern of structural oppression. This enables the Emotional Account to address both challenges to the conceptualisation of microaggressions. Moreover, by including a mechanism of *post hoc* education and calibration that interacts with other agents, the Emotional Account explains how collective discussions about oppression are nurtured by individual experiences, and vice-versa. This interaction of individual knowledge with group knowledge is explained through the lived experience of the agent and their emotions. As structural accounts distance themselves from the subjective experience of the agent, they lack this essential component to close the gap between the individual and collective knowledge of microaggressions.

3. Conclusion

The Emotional Account can respond to the problematic cases encountered by currently available definitions of microaggressions. This is the result of the Emotional Account hybridising the best features of the experiential and structural accounts: first, it keeps the ties to the agent by incorporating lived experience, just like Sue's experiential account; this enables it to reflect the reality of living in an oppressive society, especially when said oppression is nearly invisible, and thus explains the epistemic access agents have to microaggressions through their emotions and their epistemic advantage provided through their lived experience and the development of automated and calibrated emotional intuitions. Secondly, it prevents sliding into relativism, by invoking the structurally oppressive nature of microaggressions, thus building on Friedlaender's and McTernan's structural accounts. However, this time, how agents know about the link between the particular exchange and the structural oppression is made clear

through the notion of warranted emotion. Moreover, the calibration mechanism of emotions explains how the collective and individual levels of knowledge about structural oppression interact with each other, and the role individual agents play into the process of building resistant epistemologies. Thus, the Emotional Account paints a cohesive picture of microaggressions, because it answers both the ontological and epistemological queries, both at the individual and collective levels, which makes it a theory that agents can use when they disagree about whether a microaggression has occurred. Therefore, the concept of microaggression should not be abandoned, as it has now been clearly framed epistemically and ontologically with an account addressing individual cases of microaggressions. Overall, the case study of microaggressions is a clear example a scenario where agents use and should use their emotions to obtain knowledge, given the importance of self-reflective monitoring through both internal and social calibration.

Chapter Four:
Source Based Epistemic Injustice

Within a hierarchical society, the norms and values that predominate tend to serve the interests of the dominant groups.[29]

Alison Jaggar

Introduction

In this chapter, I introduce an expansion to the concept of epistemic injustice. Epistemic injustice is the idea that a wrong can be done to someone in their capacity as a knower (Fricker 2007). Epistemic injustice has been widely theorised and discussed, whether it was on the scope of the concept (Anderson 2017; Congdon 2015; Fricker 2007; Kidd, Medina and Pohlhaus 2017; McKinnon 2016; Medina 2012), or the relationship between epistemic injustice and other concepts of epistemic oppression (Bailey 2018; Dotson 2011; Dotson 2012; Dotson 2014; Pohlhaus 2012; Toole 2019)[30]. I propose to expand this scholarship, introducing what I call *source based epistemic injustice* (SBEI). This is unfairness towards knowers who use epistemic methods based on sources of knowledge currently unapproved by the dominant epistemological framework. I focus my analysis on the downgrading of knowledge obtained through affective experience, like emotions or pain, to illustrate the phenomena.

I argue that SBEI arises because certain sources of knowledge are unfairly downgraded due to an insistence on intersubjective verifiability and objectivity as an epistemological methodology. This downgrading disproportionately affects marginalised groups, leading

[29] (Jaggar 1989, 165)

[30] The literature on epistemic injustice has been abundant since Fricker (2007) introduced the concept in analytic philosophy. However, it is important to note that similar power mechanics in the epistemic domain had already been theorised by Black feminist philosophers (Collins 1986; McKinnon 2016; Pohlhaus 2017). The reader should also note that the references given here are only a sample relevant to this chapter of the literature produced on this topic.

to a form of injustice. In addition, I explain how theorising epistemic injustice through the sources of knowledge enables us to provide a causal story to hermeneutical injustice, as well as offering a new insight into testimonial injustice as the focus is shifted from the individual to the epistemic methods themselves. I aim to offer a diagnosis of real-life cases of epistemic injustice, which could not be fully explained by the current framework of epistemic injustice, by analysing them through the lens of SBEI. This will lead me to sketch the direction current epistemological practices should take if we wish to start reducing instances of SBEI.

1. Sources of Knowledge and Intersubjective Verifiability

Before we start the discussion, let me introduce some useful definitions. Throughout this chapter, the notion of *epistemic resource* includes "language to formulate propositions, concepts to make sense of experience, procedures to approach the world, and standards to judge particular accounts of experience" as defined by Pohlhaus (2012, 718). The term *epistemic method* refers to the method of knowledge acquisition like perception or memory. The term *epistemological system* refers to a system of connected theories, methodologies and practices that gives all the conditions for the possibility of knowledge production and possession (Dotson 2014). Finally, *sources of knowledge* are understood in this chapter as things in the life of the knower that yield beliefs constituting knowledge (Audi 2002).

Nowadays, the accepted sources of knowledge inside the analytic tradition are perception, memory, consciousness (also called introspection), reason (also sometimes referred to as intuition), and sometimes testimony, which is included in non-basic sources of knowledge (Audi 2002). This list includes mostly sources which provide information that is objective or intersubjectively verifiable. Indeed, perception is often considered to be intersubjectively verifiable: if I see a glass on the table, another observer can also look at it themselves to confirm whether there is a glass on the table. The same applies to memory: someone else who was at the same event could confirm whether it happened or not. Reason is often considered to be objective: any other person using the same type of reasoning could obtain a similar belief. This choice of

approved sources in analytic philosophy is the result of the history of Western philosophy. It is particularly influenced by British empiricism and positivism, where reason, objectivity and intersubjective verifiability had priority over other ways of obtaining knowledge (Jaggar 1989). As Jaggar explains, "empirical testability became accepted as the hallmark of natural science; this, in turn, was viewed as the paradigm of genuine knowledge. Often epistemology was equated with the philosophy of science, and the dominant methodology of positivism prescribed that truly scientific knowledge must be capable of intersubjective verification" (Jaggar 1989, 152). Moreover, as Audi explains, genuine sources are usually seen as reliable (2002). In our current dominant epistemic practices, it seems that reliability entails that we can intersubjectively verify the truth of a belief (Audi 2002). This paradigm in epistemology has led to the exclusion of some sources like emotions and other affective experiences. Historically, emotions were seen as irrational, due to their lack of intersubjective verification, making them unworthy of being considered a source of knowledge (Jaggar 1989).

However, one might wonder whether this epistemological system relying on intersubjective verifiability is justified, and whether it should ground the exclusion of affective experience. We might question the legitimacy of the exclusion, given the similarities between affective experience and perception. Even though perception appears intersubjectively verifiable, it is nonetheless impossible to prove that it is. The problem of inverted colour qualia is a counterexample to the intersubjective verification of perception (Byrne 2020): how can we prove that we see the same colours as another agent, if a perfectly inverted spectrum of colours is possible? What if what I describe as red is what you would describe as green? If one of the main sources of knowledge, perception, is not truly intersubjectively verifiable nor objective, then it seems illegitimate to use that argument to exclude affective experience. The same argument can be made about memory: indeed, it is impossible to access the memories of another person, yet we don't exclude memory from our sources of knowledge. This raises the question of whether affective experience was assessed fairly before exclusion from the list of appropriate sources of knowledge. If epistemologists could find ways to explain why perception is reliable even if not truly intersubjectively verifiable, then maybe it is

possible to find an explanation for affective experience. However, as I have shown in previous chapters, affective experience is either excluded or, at best, ignored by epistemologists. Thus, there are preliminary grounds to think that this insistence on intersubjective verifiability leads to the downgrading of affective experience and that this might not be justified.

Nonetheless, one might argue that there are good epistemic reasons to exclude certain epistemic methods, like affective experience, and that it is appropriate to exclude affective experience given these reasons. Indeed, as I explained in chapter 1, affective experience might be too unreliable or subjective to be a source of knowledge, and knowledge truly requires intersubjective verification. Given how much false information one can get through their affective experiences, in particular their emotions, we are better off excluding it to safeguard our epistemic practices. However, we saw in chapter 2 of this thesis that it is possible to explain when affective experience should count as a way of obtaining knowledge which is sensitive to these objections against affective experience. Thus, a fair assessment of sources of knowledge shouldn't focus only on intersubjective verification and it shouldn't categorically exclude affective experience as a source of knowledge. Moreover, as we have seen in many other places in this thesis, affective experience is used by many agents to obtain information, especially about injustice. Given the widespread use of affective experience as a way of obtaining knowledge, it seems strange to exclude it without assessing it properly.

It is likely that affective experience was not given a fair chance in modern theories due to the historical association between marginalised groups and affective experience. Throughout history, women were often painted as being "emotional" (Jaggar 1989), or "angry" in the case of Black women (Collins 2000), which is used to depict them as irrational and producing at best some sort of 'subjective knowledge'. On the contrary, white men are described as being rational, because they are driven by their reason and are thus producers of objective knowledge (Jaggar 1989). In the Middle Ages, women were thought to be more embodied than men, thus compromising their mental capacities and making their testimony less reliable (Van Dyke 2018). As Alcoff explains,

certain groups have been held to have “intrinsic tendencies and limited capacities with epistemic relevance” (1999, 78). Here, the identity bias combines with the dominant assumptions within epistemology to decrease the epistemic credibility of marginalised groups associated with the unapproved sources. Indeed, associating marginalised groups with unapproved methods of knowledge acquisition like affective experience can further their oppression and serve as a tool to maintain the power of dominant knowers.

Centring objectivity and the intersubjectively verifiable in epistemology can hide political motivations, as people uncritically believe that the ideal stance of objectivity can be filled by some individuals (Toole 2022). Specifically, the individuals who are perceived to fulfil this ideal are from the dominant class (Toole 2022, 5). Consequently, the “knowledge that conflicts with operative theory or paradigms […] [is] likely to be dismissed as inaccurate and biased” (Toole 2022, 5). It is no wonder that the anger displayed by oppressed people towards their oppressors is depicted as biased: their emotion points towards a fault in the dominant system and shows that the current society is not as just as it claims. Members of marginalised groups tend to go against the assumptions of the dominant group and object the presumed universality of various statements made by the dominant group. This is one way to explain why members of marginalised groups have been associated with subjectivity and the downgraded sources, as this would add to their credibility deficit.

I will now explain how the exclusion of affective experience constitutes an epistemic injustice. Indeed, this exclusion disproportionately affects members of marginalised groups and contributes to epistemic oppression.

2. Source Based Epistemic Injustice

Before delving into how the exclusion of affective experience disproportionately affects marginalised groups, I will introduce the idea of epistemic exclusion and various types of epistemic injustice in more details.

Currently recognised types of epistemic injustice have been divided into three orders of epistemic exclusion which compromise knowledge production and the participation of

some agents in such knowledge production (Dotson 2014). First order exclusion results from some defect in the shared epistemic resources with respect to some goal or value. It includes *testimonial injustice*, where biases inside the epistemic community result in a credibility deficit in some members of the community. In general, first order epistemic exclusion is the result of relations of epistemic power, with some agents being dominated, and others dominating, thus creating an imbalance (Dotson 2014). In the second order, the epistemological system itself is flawed, as it obscures some parts of the experience of marginalised knowers. Indeed, the epistemic resources currently produced by the system are insufficient, and the system itself cannot be used to fix this injustice (Dotson 2014). A typical example of second order epistemic exclusion is *hermeneutical injustice*, the idea that an agent's experience is "obscured from collective understanding" due to a gap in collective interpretative resources (Fricker 2007, 155). Finally, in the third order of epistemic exclusion, the system is inadequate for an epistemic task and some agents are excluded when they point out the inadequacy of the dominantly shared epistemic resources. In third order epistemic exclusion, some agents are prevented from contributing to the current epistemological system because they challenge it (Dotson 2014). For instance, an immigrant person from a different culture might be the victim of third order epistemic exclusion when they try to testify to something seen as completely incompatible with the epistemological system of their interlocutor. In those cases, only radical changes within the parameters of the epistemological system would enable that testimony to be part of the knowledge production (Dotson 2014).

What happens when affective experience is excluded from our ways of knowing? Would it count as a form of epistemic exclusion? Let me provide an example. People who experience Lordean rage- the anger directed at racist actions or attitudes (Cherry 2021, 23)- are often met with scepticism when using their anger to justify that a racist act has happened. As Cherry explains, using Lordean rage to justify one's belief leads to an evidentiary burden, where one is asked to provide further evidence to prove that one is right (Cherry 2021, 45-46). In the current dominant epistemological system, emotions are seldom taken to provide immediate justification for beliefs and are not considered

a source of knowledge (Audi 2002). However, as Cherry points out, this leads to instance of racism not being recognised and makes it harder to obtain justice (Cherry 2021).

Anger at racism is not the only case where the contributions of marginalised groups are not taken seriously by the dominant group. As discussed in the last chapter on microaggressions, small phenomena like microaggressions are not recognised by external observers if the only evidence presented is subjective. There seems to be a trend to exclude the knowledge or beliefs obtained through affective experience, even though affective experience is one of the best way to know about injustice, and other subtle social phenomena. The exclusion of affective experience disproportionately impacts marginalised groups, as it prevents them from sharing knowledge about their oppression. I will now conceptualise how this constitutes a new form of epistemic injustice.

The new type of epistemic injustice that I present here falls into Dotson's second order epistemic exclusions. *Source based epistemic injustice* (SBEI) is a loss of credibility due to a prejudice towards the epistemic methods used by the knower in order to obtain knowledge, because those epistemic methods are linked to sources of knowledge currently unapproved by the dominant epistemological system. The prejudice against affective experience as a way of obtaining knowledge is an example of SBEI.

SBEI falls into Dotson's second order epistemic exclusion because it points out to a particular way in which the epistemological system is flawed: by preventing some sources of knowledge to be considered as such, the epistemic resources or beliefs produced through those sources will not be taken into account in the dominant epistemological system, thus denying the use of concepts to some people, for which those concepts would be useful to make sense and share their experience, or by downgrading the knowledge acquired through these unapproved sources. As we can see, SBEI highlights that it is not only lack of collective epistemic resources (which would amount to hermeneutical injustice), but more importantly prejudice towards epistemic methods that result in an agent's experience being obscured. I will get back to the relation between hermeneutical injustice and SBEI in section 5. For instance, the anger

of Black people has often been dismissed and the scholarship describing their experiences, acquired through that anger didn't get a lot of uptake within dominant circles. The work of Lorde (1984) about her anger and what it taught her about oppression exemplifies this. It is not that Lorde didn't have the collective epistemic resources to make sense of her oppression. Rather, it is that her anger is not taken seriously as a method of acquiring knowledge about oppression. Note that SBEI can also fall into third order epistemic oppression depending on how radical including a new source of knowledge would change our current epistemological system. If incorporating some change regarding which sources of knowledge should be approved could be done within the current epistemological system, then SBEI would only be a second order epistemic exclusion. However, adding affective experience as a source of knowledge might require radically modifying our current epistemological system, due to a revision of the importance of intersubjectively verifiable evidence within our current epistemological system. It would also require revision of the idea that there can be some sort of 'generic knower'[31] who exists regardless of the social context. This chapter does not aim to settle this question, but it is interesting that theorising SBEI might lead us to recognise that some kinds of epistemic injustice could involve elements of different orders of epistemic exclusion.

In addition, dismissing some epistemic methods might be a case of wilful hermeneutical ignorance. As Pohlhaus defines it, *wilful hermeneutical ignorance* is when "dominantly situated knowers refuse to acknowledge epistemic [resources] developed from the experienced world of those situated marginally. Such refusals allow dominantly situated knowers to misunderstand, misinterpret, and/or ignore whole parts of the world" (Pohlhaus 2012, 715). Thus, the dominant epistemological framework, by dismissing

[31] The idea of the 'generic knower' refers to the assumption in analytic philosophy that "all knowers are believed to be alike with respect both to their cognitive capacities and to their methods of achieving knowledge" (Code 1991, 6). This assumption is heavily influenced by the cartesian model of the pure inquirer (Code 1991).

affective experience as an appropriate source of knowledge, facilitates a form of wilful hermeneutical ignorance.

SBEI might initially seem redundant to the concept of wilful hermeneutical ignorance (Pohlhaus 2012). However, wilful hermeneutical ignorance only explains part of the consequences of SBEI. SBEI includes wilful hermeneutical ignorance, as dominantly situated knowers do ignore the concepts acquired from unapproved sources of knowledge. But, in addition to this phenomenon, SBEI explains why those concepts are ignored, and puts the focus onto the epistemic methods themselves, rather than on the epistemic resources developed using those sources. Moreover, wilful hermeneutical ignorance seems to be more of an attitude present in dominantly situated knowers, whereas SBEI points to a fault in our epistemic system. This opens one possibility that usually has not been discussed and cannot be accounted for within the current framework of epistemic injustice and wilful hermeneutical ignorance: the credibility deficit can affect any knower, whatever their identity, if they use an unapproved source of knowledge; and in contrast, any knower who uses currently approved sources of knowledge can gain credibility.

3. Why This Matters: The Case of Pain Testimony in Medicine

I will now show the importance of theorising SBEI, to complement the existing framework of epistemic injustice. I explore a case where it is nearly impossible to provide explanations about how one acquires information through currently approved sources of knowledge or to provide intersubjectively confirmable evidence for one's claim. Indeed, there are cases where knowledge can only be acquired through the agent's affective experience. I highlight that when this only way of knowing is downgraded, it leads to injustice and in many cases contributes to harm people.

Let us turn a practical case: the diagnosis of endometriosis. Endometriosis is a painful condition affecting 1 in 10 people with a uterus, in which tissue similar to the lining of the uterus starts to grow in other places of the body. Every month, this tissue bleeds in the same way as tissue in the womb, but unlike a typical menstrual cycle, the blood has

no way to escape, causing painful and/or heavy period (Endometriosis UK 2022). It currently takes an average of 7.5 years to get a diagnosis in the UK, after a patient sees a doctor for the first time about their symptoms (Endometriosis UK 2022). Similar diagnostic times can be observed in other countries (Nnoaham et al. 2011). Additionally, a study conducted on a group of Brazilian women highlighted that patients consulting for pelvic pain wait on average 3.4 years longer than patients consulting for subfertility to get diagnosed with endometriosis (Arruda et al. 2003).

The framework of epistemic injustice offers some good explanations as to why people who suffer from endometriosis are ignored for so long: most of them are women, and women can suffer from testimonial injustice, due to their identity as members of a marginalised group. Moreover, there are social and societal structures, as well as historical and economic realities that can influence medical practice (Blease, Carel and Geraghty 2016; Buchman, Ho and Goldberg, 2017). However, these explanations do not fully explain why there is a difference of 3.4 years in diagnostic time for the same condition, given that it is diagnosed in the same manner whether the primary symptoms are pelvic pain or subfertility.

Here the significant difference appears to be the different reasons invoked by the patient to justify that something is wrong: pain and subfertility. When testifying only through pelvic pain, the patient waits longer for a diagnosis, while testifying through a factual observation (not being able to conceive like in cases of subfertility) leads to further testing. Thus, in cases of patients with pelvic pain something appears to go wrong at the level of the testimony and the reasons invoked, and not at the level of the identity of the patient. If identity was the only important factor, both groups (subfertile patients and patients with pelvic pain) should be treated in the same manner as they are from the same demographic group. This means that the differences between both categories of patients cannot be fully explained by invoking solely testimonial injustice. Similarly, hermeneutical injustice is not the relevant variety of epistemic injustice for these cases. The patients with pelvic pain have the resources available to make sense of their experience, as they have the concept of ‘pain’. Moreover, the healthcare

practitioners (HCPs) have the concept of 'endometriosis' and know what the illness is. Thus, the lack of further investigation into the symptoms cannot be imputed to a lack of appropriate conceptual resources to investigate the situation further.

One variety of epistemic injustice that might be relevant here is *content-focused injustice*: talking about painful periods might trigger a prejudiced reaction from the HCPs, thus leading to the content of the assertion ("my period pain is abnormal") not being believed, due to some operative prejudice against a group (here women) that the hearer associates with the assertion (Dembroff and Whitcomb 2022). However, there are two reasons to pushback against such an interpretation as the ultimate explanation for the difference between cases of pelvic pain and subfertility.

First, while there is a stigma around period pain which might lead to downplaying it, it would be surprising that pain complaints would be ignored by most HCPs for many years given the prevalence of endometriosis (1 out of 10 people with a uterus). One might alternatively invoke that the relevant form of content-focused injustice is that talking about subfertility triggers a prejudiced reaction valanced positively towards testimony, as women's fertility is valued within most societies. Thus, treating subfertility appears more important than treating pain. While this is a real concern, it seems quite unlikely that most HCPs only care about their patients if they consult for subfertility.

Second, most HCPs want to help their patients but are constrained by the material reality of the healthcare system, in which time and resources are scarce. In that system, they must be efficient, and can continue investigating only if they have reasons to warrant further investigation. While the content does change between cases of subfertility and pelvic pain, as it goes from complaining about subfertility and pain to pain exclusively, we can interrogate whether it is the content itself that is the relevant dimension here or whether it is that the different content (subfertility) brings in a way of testifying that is objectively shareable. For that reason, HCPs might be more enticed to look for the cause of the symptoms when the symptoms are intersubjectively verifiable. Unfortunately, in the case of pain, the HCPs cannot know whether the cause is psychological or physical. In cases where HCPs cannot find a biological marker easily,

the lack of objectively observable evidence might lead them to abandon the search for a cause and diagnosis, and they might fall back on the idea that periods are unfortunately painful for some people, due to natural variation, rather than an illness.

Why does it matter that the symptoms are objectively shareable? I explore in the rest of this section the hypothesis that there is a link between the lack of credibility given to pain assessment and what is considered to be an appropriate source of knowledge, which leads to a slow process for the diagnosis. Whatever the identity of the patient is, there will always be a credibility deficit if the only way for the patient to prove that they should be tested for a medical condition is by testifying on the basis of their painful affective experience. For instance, when someone breaks a leg, and forms the belief that their leg is broken because it is terribly painful, there is either immediate visible evidence, or evidence can be seen through X-rays to confirm the report. However, when it comes to some illnesses with less visible causes, it can be hard to empirically prove that something is wrong.

This problem is amplified with endometriosis: periods are assumed to be painful, so how could one distinguish between a normal pain level, and an abnormal one, especially when pain assessment is subjective? As Buchman et al. explain, "such subjectivity may prompt some [Health Care Practitioners] to question the truthfulness of the pain sufferer's testimony, as symptoms may become viewed as embellished or contrived" (2017, 32). Moreover, more credibility is conferred to people who use clinical methods relying on objective factors, like results of blood tests or medical imagery, compared to methods relying on patients' testimony, like mental health screenings. This creates a hierarchy between different methods of medical assessment (Buchman et al. 2017). Due to the adoption of a form of '*mechanical objectivity*', in which medical knowledge is produced via empirical investigation and without the subjective influence of the investigator, the primary goal is to represent the scientific object just as it appears in its most natural state (Buchman et al. 2017).

But as one can already realise, it is very difficult to apply this model of knowledge acquisition to pain assessment. When it comes to pain, there can be discrepancies

between the symptoms of pain and a visible object. For the same condition, like a broken leg, one person might be in terrible pain, while another one might find it slightly inconvenient. However, in the case of endometriosis, unless one undergoes an exploratory laparoscopy[32], there is no way to prove that what is causing the pain is endometriosis and not another condition, or even just that the pain experienced is only the result of a low pain tolerance combined with some unfortunately strong period cramps. Thus, in many endometriosis cases, the patient will be faced by resistance from the medical staff, in order to prove that some tests are worth being done to investigate whether they have endometriosis or not (Kiesel 2017). This slows the process down, as the patient might not have the courage to make their voice be heard when their doctor tells them that it is just a painful period.

Current medical practice has been influenced deeply by empiricism and the need for objectivity is not surprising. Mechanical objectivity and intersubjectively verifiable evidence were crucial for medicine to be treated as a science. Indeed, without systematicity and objectivity, medicine could not have reached the status of science (Bird 2019). This explains why there is such an important focus on objectivity in medical sciences, and a rejection of the subjective. Without it, medicine would not appear to be as reliable. Nonetheless, by promoting the idea that knowledge can only be gained from the perspective of a 'generic knower', science and philosophy have both pushed forward the perspective of the world experienced from a dominant position (Pohlhaus 2012). Indeed, feminist epistemology has long tried to explain that the insistence on a generic knower and value-free knowledge is unfortunately not as neutral as we believe (Code 1991; Longino 1990). Indeed, even the conception of the generic knower reflects values and ideology of the dominant group (Code 1991). In that respect, it has erased the perspective of many individuals, and in the case of medical pain assessment indirectly

[32] A laparoscopy is "an operation in which a camera (a laparoscope) is inserted into the pelvis via a small cut near the navel" (Endometriosis UK 2022).

caused physical harm to individuals, especially women, people of colour and members of other marginalised groups.

Insofar as someone might not be treated immediately for a condition because they only know about their condition through their affective experience, then the dominant epistemic framework is unjust for dismissing their complaint, especially when that way of knowing is the only one available. If affective experience was taken seriously as a source of knowledge, then those people would have more credibility, which would lead to a better investigation of their symptoms and what could be causing such pain. It would surely reduce the time between the first appointment taken to discuss the pain and the appointment in which a correct diagnosis is offered. For instance, a preliminary diagnosis of endometriosis can be done through an MRI scanner (taking only an hour), a medical drug trial test (taking about 3 months to see if the symptoms disappear) and that preliminary diagnosis is confirmed through a laparoscopy, which is a surgical act that takes about a day (Hsu et al. 2010). As one can see, it takes nowhere near 7.5 years to be able to obtain a diagnosis. Even if the longest testing route is taken (3 months with the drug trial) then it should never be more than a few months to obtain a diagnosis.

Endometriosis is not the only medical condition where SBEI occurs. As Buchman, Ho and Goldberg explain, "people living with chronic pain have been considered systematically less credible in clinical encounters as compared to medical experts" (2017, 36). Indeed, chronic pain sufferers are less likely to be believed. Why is it so? The explanation SBEI provides is that this is the result of bias against the epistemic methods used by those patients. Indeed, if their only method of acquisition for knowledge about their illness is through the phenomenal embodied experience of that illness, i.e. the symptoms of pain, then it does not fall into the currently approved ways of acquiring knowledge, which leads to a credibility deficit. As Wendell explains, the patient's experience is not considered essential to establish a diagnosis when compared to the medical explanations that the doctor can provide, due to the authority that medicine has over medical knowledge (1997). Unfortunately, there also seems to be a tendency to minimise the patient's experiences that cannot be explained by available medical

knowledge (Wendell 1997)[33]. One important aspect to highlight is that SBEI does not have to be the only type of epistemic injustice at play in one situation. Indeed, chronic pain sufferers from marginalised groups have it even worse than patients from dominant groups, as the credibility of a patient's report is even more likely to be questioned when the patient is from a marginalised group (Buchman, Ho and Goldberg 2017). However, SBEI enables us to explain why there is a generalised credibility deficit towards chronic pain sufferers, even towards those not from marginalised groups.

By using the theoretical framework of SBEI, we can now highlight another type of epistemic injustice present in the case of some medical illnesses: the experience of patients is downgraded as a source of evidence due to the unfair downgrading of pain as a legitimate source of knowledge. Theorising SBEI shows how deeper mechanisms are at play when it comes to bias preventing groups from participating in knowledge production. To fully understand what the current problem is, we need to shift some of the focus towards epistemic methods themselves and the biases towards them. This leads me to the next important part of this discussion: what are the dynamics between the current concepts of testimonial and hermeneutical injustice and SBEI?

4. Source Based Epistemic Injustice Within the Epistemic Injustice Framework

SBEI is a form of epistemic injustice that fills two gaps within the global framework of epistemic injustice. First, SBEI ties together the dynamics of bias towards some groups with the bias towards their associated epistemic methods, and how this latter type of bias can persist even towards agents not from marginalised groups if they use those methods. SBEI is different from testimonial injustice, as the prejudice is not directly towards the knower themselves, but rather at the level of the method of knowledge acquisition. Anyone using their affective experience, like pain or emotions, to obtain knowledge could suffer from an unjust credibility deficit. Therefore, we can say that SBEI

[33] I will explore this theme in more depth in Chapter 5 of this thesis.

is closely related to the way testimonial injustice is conceptualised. However, it differs from it, because the two are happening at different levels of epistemic exclusion. Testimonial injustice is a first order epistemic exclusion, and by theorising a second order epistemic exclusion like SBEI, I have shown which flaw of the epistemological system could be fixed to prevent testimonial injustice. Indeed, if epistemic methods, like obtaining knowledge through one's affective experience, were given the same weight as methods relying on intersubjective verifiability before being dismissed, many cases of testimonial injustice might not occur.

We can also explain the paradigmatic case of Marge in the fictional movie *The Talented Mr Ripley* (Minghella 1999), presented by Fricker (2007) in a different light. In the movie, Marge's fiancé, Dickie Greenleaf is dead, and Marge thinks that Tom Ripley killed him, even though it looks like Dickie committed suicide. When she tries to explain this to the police, she is not believed. Marge appears to pick up on nuanced behavioural differences and hardly noticeable body language in Tom Ripley. Given her past experience with Dickie, she doesn't believe the others who tell her that her fiancé Dickie left her and then killed himself. The last piece of evidence she finds- that Ripley is in possession of Dickie's rings- is a confirmation of her intuitions and observations about her suspicion towards Ripley. However, when she tries to explain to Herbert Greenleaf, Dickie's father, that Ripley is responsible for Dickie's death, because he has the rings, she is only able to explain herself by saying "I don't know, I don't know, I just know it" (Minghella 1999). She then gets silenced by Herbert Greenleaf who tells her "Marge, there's female intuition, and then there are facts" (Minghella 1999). Marge cannot explicitly prove that the rings are in possession of Ripley because he killed Dickie and not because Dickie gave them to Ripley. The other clues she has obtained through interactions with Ripley are not considered a proper form of evidence, due to their subjective nature, and she is unable to prove to Herbert Greenleaf that Ripley is the killer. What happens to Marge is the result of an illegitimate dismissal of the epistemic methods she uses. If SBEI was not present, epistemic methods like acquiring knowledge through personal experience and subtle cues would have been given some attention. Given the insistence on

intersubjectively shareable evidence, Marge's testimony is dismissed as being a form of 'female intuition', even though her testimony should have been paid attention to.

The second gap filled by SBEI concerns the origin story of some occurrences of hermeneutical injustice. Hermeneutical injustice is a form of structural injustice that creates gaps in the collective hermeneutical resources, leading to difficulties in diagnosing and understanding one's experiences. To put it simply, structural injustice prevents concepts from being generated by marginalised groups and/or included in the wider epistemological system (Fricker 2007). I think that this could be the result of SBEI currently targeting members of marginalised groups who used their affective experience and emotions to better understand their world and develop related concepts. Because some important concepts generated by marginalised groups are concepts acquired through currently unapproved sources of knowledge, like affective experience, they are not taken seriously, which in turn prevents the inclusion of the developed concepts in the collective hermeneutical resources. Moreover, some of the most useful concepts developed by marginalised groups used to describe their experience as members of a marginalised group are concepts for which the application is justified through appeal to affective experience. Let us examine current concepts developed by marginalised groups like *microaggressions* and *sexual harassment* to illustrate this mechanism.

Microaggression is a term introduced in 1970 by Black psychiatrist Chester Pierce (Pierce 1970) and refers to "brief, everyday exchanges that send denigrating messages to people of colour because they belong to a racial minority group" (Sue et al. 2007, 273). The term is now used to include any member of a marginalised group. One important aspect of microaggressions is that their subtlety makes them very hard to recognise, even for the victims. However, through the repetition, the victims gain an understanding that only they can obtain. Indeed, when they try to share their experience and knowledge of the phenomenon, they are often "pushed to the edge of knowing", as Fatima explains (2017, 150), because other agents doubt the veracity of their claims. The concept of microaggression has received a lot of attention, especially from people doubting that this concept points to a real phenomenon in the world (Campbell and

Manning 2014; Lilienfeld 2017). It is interesting that the main counterargument to microaggressions is that the victims cannot provide proper intersubjectively confirmable evidence that there is an offence, thus prompting the authors to argue that it is because those offences are only perceived and not real (Campbell and Manning 2014; Lilienfeld 2017). The pertinent mechanism at play here is the following: these authors argue that perceiving an offence and knowing it through the lens of one's experience is not enough to count as knowing and use such a premise to completely discredit the concept itself, referring to it only as an instance of victimhood culture. For instance, Campbell and Manning explain that people who complain and are angry about microaggressions are "call[ing] attention to their own victimisation" and "people portray themselves as oppressed by the powerful – as damaged, disadvantaged, and needy" (Campbell and Manning 2014, 22). Moreover, they consider that reacting to microaggressions is a form of new morality, where people are "highly sensitive and easily offended" (Campbell and Manning 2015). Thus, by discrediting the way the concept is acquired and how applications of it are justified, critics like Campbell and Manning or Lilienfeld can discredit the concept itself.

Activists who represented the women's right to work in an environment without unwanted sexual advances faced a similar problem. Indeed, the term *sexual harassment* came to light in 1974, as part of Lin Farley's course at Cornell University on women and work (Siegel 2003) and made sense of the experience of many women throughout the ages. However, when time came to include sexual harassment in law, law makers decided to recognise "sexual harassment as a species of sex discrimination, it did so without acknowledging the larger social arrangements within which the practice of sexual harassment acquired dignitary meaning and distributive consequence" (Siegel 2003, 18). The problem here is that the concept was transformed to fit within the dominant framework of epistemic practice, which required that a claim of sexual harassment can be proven only if publicly available and intersubjective evidence can be pointed to. Indeed, by recognising only the practical aspects of sexual harassment, without taking into account the social dynamics linked to it, law makers rendered the concept useless to people who needed it the most. It is not surprising that in 2017,

women took over social media to highlight the high frequency of sexual harassment, with the #MeToo movement, as the concept of sexual harassment was mischaracterised since the seventies to advantage members of dominant groups.

Sexual harassment or microaggressions can be hard to prove. Therefore, victims often only rely on their own experience to prove that they are sexually harassed or microaggressed. As they often cannot produce knowledge based on intersubjectively verifiable evidence, they are dismissed because their testimony is seen as deriving only from an emotional place, thus made untrustworthy. If *sexual harassment* as a concept was not so dependent on the affective experience of the victim to come to exist in the first place, it is likely that the concept would have been theorised long before the seventies. The examples of *microaggressions* and *sexual harassment* show us that the production of concepts making sense of the experience of members of marginalised groups is itself the target of epistemic injustice, in particular SBEI. When SBEI is present, the production of concepts through unapproved sources of knowledge is impacted and the ensuing concepts are discarded from the collective epistemological resources, resulting in hermeneutical injustice.

5. Conclusion

Framing source based epistemic injustice in our current epistemological practices is vital if we want to assess bias in our knowledge production. By distrusting on a larger scale affective experience and deeming epistemic practices using such sources of knowledge as unreliable, we participate in epistemic oppression. This has consequences far beyond philosophy itself, as shown by the harm occurred by people who suffer from endometriosis, the problem of pain assessment in general and the resulting hermeneutical injustice when the concepts emerge from unapproved sources. If we want to stop such a vicious cycle, we should reconsider what we label as an appropriate source of knowledge. It would be warranted to give the current unapproved sources of knowledge a fair assessment, to see whether they should be included or not, instead of excluding some of them categorically. For instance, we could theorise about when they

are reliable, in which conditions they lead to knowledge, or when testimony based on them should get uptake.

As I showed earlier in Chapter 2 of this thesis, it is possible to theorise affective experience in a way where we can distinguish instances where it should count as a source of knowledge, and when it should not. All the cases presented in this chapter like pain towards endometriosis, or anger at microaggressions, sexual harassment or racial injustice are cases where the Reflective Model would say that these affective experiences provide the agent with a justified belief. There is no reason to exclude affective experiences from our approved ways of knowing. If at the end of the day, we end up thinking that a justification involving intersubjectively confirmable evidence is better, that does not mean that any justification appealing to affective experience or emotion is worthless or very weak. To theorise SBEI is a first step towards this enterprise, and I am hopeful that in the future there will be more theories showing the important role of emotions and affective experience in knowledge acquisition at least, and maybe how such sources deserve their place within basic sources of knowledge.

Chapter Five:
Patient Activation or Patient Exploitation ?

No one will like me for saying this, but I've often noticed that it doesn't seem to matter whether a woman goes in for yoga or reiki or dancing on hot coals at the full moon - she starts to get better when she starts to take charge.[34]

Hilary Mantel

Introduction

In recent years, there has been a shift in healthcare towards a greater empowerment of patients, by encouraging what is known as 'patient activation'. Patient activation refers to knowledge, skills, and confidence in managing one's health and healthcare. A recent study about healthcare service utilisation in England has shown that higher levels of patient activation yield better clinical indicators, as well as lower healthcare utilisation and costs (Bu and Fancourt, 2021). Another study showed that higher activation was associated with better health outcomes and lower costs two years later (Greene et al. 2015). Moreover, being an active inquirer who is informed about their condition has been highlighted in feminist medical ethics as a crucial step in being an autonomous agent (Kukla 2012). Thus, it appears that gaining knowledge about one's condition and the healthcare system comes with many positive aspects, both for patients and the healthcare system in general.

Contra this trend, I propose in this chapter to problematise some of these benefits by arguing that patient activation can lead to the exploitation of patients, in particular patients suffering from lesser-known illnesses. Indeed, there are concerns about the literature praising patient activation, as most of these studies rely on quantitative data which might omit certain factors. This could lead to a gap in the identification of barriers and challenges to implement patient activation, as both quantitative and qualitative studies are needed to identify effective patient activation interventions (Chen et al.

[34] (Mantel 2004)

2016). Additionally, the topic of patient activation has not been examined through a conceptual lens, apart from one study on the ethics of patient activation around decision-making and responsibility (Gibert, DeGrazia and Danis 2016). This opens the possibility of various unexamined problems for patient activation and the harms and downsides associated with it.

One such unexamined conceptual area that I propose to discuss in this chapter is the transmission of knowledge between patients and healthcare practitioners (HCPs) and how this relates to patient activation. HCPs are regarded as the experts to formulate a diagnosis and decide on the best care, whereas the patient's knowledge is often considered as less important. However, in the case of lesser-known illnesses, this purported asymmetry is overgeneralised; in such cases, it may well happen that the patient possesses more information and evidence than their HCPs. As I will explain, this leads to a disagreement between HCPs and patients about which ways of knowing are legitimate, especially if the patient uses their lived experience to testify and justify a further inquiry into their symptoms or a change in medical treatment.

I argue that this disagreement, sustained by an illegitimate asymmetry, leads to injustice. This results in two types of burdensome labours, which are modelled after discussions on epistemic labour and hermeneutic labour performed by women in the context of heterosexual relationships (Wilson 2021; Anderson Forthcoming). I explain how the healthcare system requires epistemic and hermeneutic labour from various patients, not just to improve some already good health outcomes, but rather to survive and avoid medical neglect and dismissal. Given the erasure of the labour performed by some patients, as well as the inescapable nature of the labour in the case of lesser-known illnesses, I argue that patient activation can be epistemically exploitative, rather than a way to empower the patient.

1. Knowledge, Expertise and Disagreement in Healthcare

In June 2022, the Guardian published a special issue on Women's health (Walker 2022). In it, various gynaecological condition like polycystic ovary syndrome (PCOS) and

endometriosis are discussed. PCOS is characterised by irregular periods, excess androgen and the ovaries become enlarged (NHS 2022b). Endometriosis is characterised by pelvic pain, painful periods and subfertility (NHS 2022a). In one of the articles (Walker 2022), the following testimony appears:

> *"I started going to the doctors about excessive bleeding and irregular periods four years ago. I was always told it was probably stress, even when I insisted it wasn't.*
>
> *I even requested my doctors' notes and found that I had been described as a "weeping woman", while one nurse had written "should have a scan" but never referred me.*
>
> *I begged for a pelvic scan in the end and had to wait six months for the appointment. I went on my own in December last year because I had been told so many times nothing was wrong, and was expecting it all to be fine, and then in an abrupt and sudden manner was told I have an 8 cm fibroid and polycystic ovary syndrome (PCOS)."*

Unfortunately, this is not an isolated instance of malpractice. As more reports and testimonies emerge, there is a pattern in healthcare systems in affluent countries of North America and Western Europe to dismiss the experience of people suffering from lesser-known illnesses (Devlin 2022; Haagaard 2022; Turner 2022; Walker 2022). In this chapter, lesser-known illnesses will be defined as illnesses where the diagnostic methods are ineffective, or where there are difficulties to diagnose the illness, or where the research to understand the illness is lacking, or rare illnesses[35]. Returning to the testimony above, we can ask ourselves: why did it take so long for this patient to get diagnosed with PCOS and the fibroid, and why were they told that nothing was wrong for four years?

One analysis of the situation would be that it stems from sexism, ableism or some other form of oppression. As Buchman, Ho and Goldberg (2017) explain, a combination of wider stigma, like gender, race, ethnicity or class inequities exists in current Western

[35] One should note that this category can include illnesses with wide prevalence, which have for various reasons been ignored by the medical establishment.

Healthcare, which leads to a downgrading of credibility, as well as gaps within knowledge, given the history of medicine as a white and male-centric discipline. Another explanation for this situation resides in the material issues facing the healthcare system. Economic realities can deeply impact the time, effort and resources provided by the healthcare system to patients. For instance, reports about the NHS England published in 2022 revealed a staffing crisis, with 10.3% of nursing posts and 5.8% of doctors' posts vacant (Campbell 2022). Thus, due to staff shortage there is ongoing pressure on healthcare practitioners to treat more patients than would be recommended. Given the historical and economic realities, some healthcare systems tend to be underfunded, which can lead to a materially ineffective system (where staff shortages and resource shortages are frequent due to the lack of resources). In turn, this materially ineffective system grounds demand for additional labour from the patient if they want to utilise the system to the best of its ability. This would explain why the patient in the example had to beg for a pelvic scan to be able to get one.

Whilst it is undeniable that the political and material issues presented above constitute a systematic problem for ensuring a fair healthcare system for all, one might wonder whether such situations are only the result of political and material issues. Indeed, as Wendell points out, "the authority of medicine tends to delegitimise our experiences of our bodies as sources of knowledge about them, because the authoritative, that is, the medical and scientific, descriptions of our bodies are third-person descriptions of physical conditions" (1996, 119). As a result, the first-person phenomenological descriptions offered by patients about physical conditions are at best treated as a sign for some further truth which needs to be confirmed by further scientific and medical evidence, due to the superiority of third-person knowledge. Thus, it might be worth looking at whether there is a problem with the way in which the patient testifies and if some ways of testifying are systematically overlooked by HCPs.

In medical sciences, there has been a general adoption of *mechanical objectivity*, according to which medical knowledge should be produced via natural investigation and without the subjective influence of the investigator. Under this scientific paradigm, the

primary goal is to represent the scientific object just as it appears in its most natural state (Buchman et al. 2017).

Moreover, HCPs, in particular physicians, possess what Wendell calls 'cognitive authority' (1996, 117). That is, their expertise will always trump expertise that their patient might have. In many cases, it is true that HCPs are more knowledgeable than their patients: they underwent a long professional training and kept learning by practising throughout the years as a HCP. However, HCPs, whilst being experts in their field, do not know about every illness. Moreover, they simply cannot update their knowledge on *all* illnesses given the numerous conditions that can afflict human beings. This creates gaps in their ability to assess some new cases presented to them, as they might not be able to diagnose all illnesses or might not know how to treat some of them. Finally, they can be prone to availability bias, where their recent experience with a condition can lead them diagnose other patients with the same condition. Due to this bias, HCPs are more prone to diagnostic errors which cannot be repaired solely through adopting a reflective approach (Li, Cheng and Liu 2020). This means that patients with rarer illnesses are less likely to get diagnosed correctly than people with more common occurring conditions.

By contrast, some patients can acquire knowledge through ways that do not neatly align with the paradigm of mechanical objectivity. Many patients acquire knowledge that something is wrong with their bodies through their lived experience. Given that reliability seems to be a hallmark of knowledge acquisition and scientific practice, and that agents need to be able to test whether the deliverance of a source is confirmed, using one's subjective lived experience appears to be incompatible with the requirements of medical science. If diagnostic methods avoiding subjective experience are available, then they are preferred to the reliance on the latter. Indeed, if medical science needs evidence that is transferrable or replicable between agents, then subjective experience doesn't fit this criterion. In the case described at the start of the section, only the patient can assess whether they experience *excessive* bleeding and *irregular* periods, given that this has to be contrasted with their previous experiences.

Indeed, what may look like an irregular period for one person might be perfectly fine for another one. This tension between mechanical objectivity in the sciences and the subjective experiences featuring at the centre of some patients' testimonies creates a problem for what is considered to provide good justification for beliefs, and thus can impact whether the testimony made by the patient is taken seriously. Indeed, does lived experience make patients more knowledgeable than their HCPs? It seems that in many cases, this lived experience is an important part of the evidence that leads to further investigation and diagnosis, or to a change in care. However, due to an asymmetry in expertise, and the presupposed superiority of the expertise of the HCPs over that of the patient, the knowledge of the patient will not be taken seriously, even in cases of lesser-known illnesses where they presumably know more than their HCPs (Buchman, Ho and Goldberg 2017). Using the framework introduced in chapter 4 of this thesis, this situation can be described as a case of source based epistemic injustice.

In some cases, HCPs and patients might fall into a seemingly unsolvable disagreement due to this perceived asymmetry: the patient testifies that something is wrong with their body and that further testing is required or that their treatment should be modified. However, the HCPs, if they cannot find an answer with the tools they have, can often bypass such testimony, and declare that nothing is truly wrong, and dismiss the problem as being in the patient's head (Wendell 1996, 123). This seems to be what happened to the patient in the introductory example: they are told that their irregular periods and excessive bleeding is due to stress, even when they insist that it is not. In medical encounters, more weight is often given to the HCP's opinion and expertise, which can have adverse effects on the epistemic confidence of patients in their own bodily experiences (Wendell 1996, 120). Due to patients being less powerful than HCPs in the structure that is healthcare and medicine, they cannot convince HCPs of the reality of their experiences if the HCPs fails to see them in the first place.

More worryingly, it appears that the stance of the patients and the stance of HCPs might be so diametrically opposed that they become irreconcilable. As Leder explains, medical practice centres the body as a corpse rather than the body as lived experience (1990).

This leads to the patient being conceived as a physiological machine, where "(d)iagnosis and treatment seek to address the observed lesion, the quantified measurement, more than a person living in pain. The patient's own experience and subjective voice become inessential to the medical encounter" (Leder 1990, 147). Due to this, modern medical science misses out on many opportunities to consider the lived experience of patients and take them as a form of evidence which is part of the investigation, especially if the bodily experience of the patient cannot be explained through the mechanical objectivity framework. However, to change this practice by including subjective experience would mean to rehaul medical sciences completely. Mechanical objectivity and systematicity are seen as the legitimate method of knowledge production in medical sciences. Indeed, they are in part what provided clinical medicine with the status of science in the eighteenth century (Bird 2019).

When illnesses are not directly assessable within the medical objectivity framework, patients face their first hurdle: their experience is not going to be enough to serve as evidence, and thus, they will have to either get medical research to advance or they will need to educate their HCPs about their conditions. However, this won't be easy. Indeed, patients face the second barrier at this stage. Qua patients, their expertise is not as credible as the one of the HCPs. Thus, unless they work tirelessly to be heard, they will have to suffer the silencing of their expertise and their evidence. These difficulties stem from the presence of the presumed asymmetry of who is the expert between the HCPs and the patient. In many cases, this asymmetry is legitimate. However, as I have explained above, HCPs might bypass the testimony of the patient, even if they are in fact less knowledgeable, in cases where they do not possess the tools to assess what is happening to the patient.

Due to this, patients have to spend energy trying to bridge the gap between them and HCPs in order to get suitable care. In the next section, I explain two labours: first, through performing epistemic labour, patients acquire the tools to bypass this imbalance in knowledge and are able to deflect assumptions that they are not credible in order to be heard; second, by performing hermeneutic labour, patients gain a further understanding

of the mechanisms of power at play in interactions between patients and HCPs, to help them understand how they might be silenced and how to best to present themselves in the interaction.

2. Defining Epistemic Labour and Hermeneutic Labour

2.1 Epistemic Labour

Following Wilson, I view epistemic labour as "labour that one performs in their capacity as an epistemic agent" (2021, 114). According to her account, epistemic labour arises when one uses capacities linked to knowledge possession and production, such as planning skills, educating others, or conducting research (Wilson 2021). In general, epistemic labour is defined as the work that underwrites the acquisition of knowledge (Goldberg 2011).

Epistemic labour refers, in the context of healthcare, to the energy and time required to obtain knowledge about various aspects of one's illness or condition and share this information with others. For instance, someone with diabetes would learn what makes their sugar levels go up and down, how to check them and how to administer insulin. They would also need to explain this information in relevant contexts, like family members or staff preparing food. In addition, epistemic labour encompasses knowledge about one's condition and the medical system, but also, time spent on educating medical experts or the public about one's condition. In this case, a diabetic person might have to learn how to procure some needles, where to get insulin, or how to get it reimbursed if they have a medical insurance. It could also involve advocating to get the condition better known, or to share useful management techniques around one's condition, like what diets or activities help stabilise it. Finally, it also entails the energy and time spent on getting one's illness recognised and taken seriously. In the case of diabetes, this last form of epistemic labour is lessened thanks to the fact that this illness is well-known both by the public and HCPs, and properly recognised as an illness.

However, not all conditions benefit from this credibility. For instance, people with myalgic encephalomyelitis (also known as chronic fatigue syndrome or ME/CFS) or

Ehlers-Danlos Syndromes (EDS), a group of conditions affecting connective tissues, are often not taken seriously, as their conditions are poorly understood and their symptoms, even when severe, are not taken as being serious, or worse, can be brushed off as just being 'in the head' (see (Wendell, 1996) for ME/CFS and (Ehlers-Danlos Society 2020; Migdol 2018) for EDS). The history of the recognised symptoms for advanced multiple sclerosis offers a good illustration of this epistemic labour regarding the education of specialists: it is after patients advocated that they were experiencing severe pains in places like bones, muscles or skin that these effects were finally investigated instead of being brushed over (Wendell 1996, 124). However, whilst the research around pain in multiple sclerosis finally focused more on subjective symptoms like pain, a study from 2013 based on a meta-analysis of research findings around pain in multiple sclerosis highlighted that pain is still widely understudied for this condition (Foley et al. 2013). As these examples show, the first negative aspect of epistemic labour is that it can become burdensome. Indeed, patients need to spend a lot of time and energy performing epistemic labour, especially if they have a lesser-known illness.

One might note that epistemic labour is not necessarily bad, and that this burdensome aspect could be offset by the positive gains of performing this labour. Being an active inquirer, which involves being informed about one's medical condition, has been highlighted as an important step in being autonomous agents (Kukla 2012). Moreover, informed consent, the idea that patients can decide on their treatment and self-determine (Kukla 2005), is seen as an important step to counter medical paternalism, where HCPs are in a position of power where they decide what's best for their patients without consulting them. These concerns around countering medical paternalism have led to more patient-centred approaches. For instance, the United Kingdom has healthcare laws that centre the idea of shared decision-making between the patient and the HCPs (Gauthier-Mamaril 2022). The topic of autonomy in healthcare is widely covered in the literature, whether it is through feminist conceptions of autonomy (see Gauthier-Mamaril (2022), Lindemann (2009), McLeod and Sherwin (2000) for examples in this tradition) or traditional conceptions of autonomy (see Dive and Newson 2018, Kukla 2005, Pugh 2020 for examples of this). However, given the above discussions on

epistemic labour, it becomes clear that autonomy does not come for free. Indeed, there is a cost to becoming autonomous, as one needs to put work into educating themselves to achieve autonomy. One might still argue that epistemic labour, even if burdensome, is an acceptable cost of being autonomous.

However, this vision is flawed, as epistemic labour can be performed, yet go unrecognised, thus making it a waste of time and energy. I will now expand on this form of injustice that patients face after they have performed epistemic labour. As Wilson explained, domestic epistemic labour is performed by women, without them being recognised and acknowledged as competent epistemic agents, which constitutes an epistemic injustice (Wilson 2021). Similarly, there is already a growing body of work showing that patients are often not recognised as knowers by HCPs, and that epistemic injustice also often occurs within patient-practitioner interactions, due to the power imbalance between the two groups (Blease, Carel and Geraghty 2017; Buchman, Ho and Goldberg 2017; Carel and Kidd 2014). Indeed, HCPs are often taken to possess the knowledge and authority in the medical context, compared to their patients who are often viewed as not knowing their own bodies (Wendell 1996). Due to this presumed asymmetry about whose expertise and evidence matters, patients are silenced on epistemic grounds. This latter dimension is heightened in the case of lesser-known illnesses, as patients need to work harder to get their knowledge recognised or to acquire knowledge that isn't readily or easily available due to gaps in the current dominant framework of healthcare. The example above of multiple sclerosis and the work of patients to get more symptoms recognised and taken seriously is illustrative of this phenomenon.

To summarise, epistemic labour in the context of healthcare refers to the labour performed by patients in their capacity as epistemic agents. Patients face two problems linked to epistemic labour: first, this labour is burdensome, due to the great degree of labour that patients have to perform to secure acceptable health outcomes, especially in the case of lesser-known illnesses; second, patients are faced with an epistemic injustice, as they have to do this labour, yet are not taken seriously when they do so. As

a result, they might have to perform even more epistemic labour to overcome this injustice, or have to perform hermeneutic labour, which I will now introduce.

2.2 Hermeneutic Labour

Whilst patients with lesser-known illnesses spend a lot of time and energy performing epistemic labour, I argue in this subsection that they also have to perform hermeneutic labour. Anderson (Forthcoming) introduces the idea of hermeneutic labour, which refers to the burdensome activity of "understanding one's own feelings, desires, intentions, and motivations, and presenting them in an intelligible fashion to others when deemed appropriate"; "discerning others' feelings, desires, intentions, and motivations by interpreting their verbal and nonverbal cues, including cases when these are minimally communicative or outright avoidant"; and "comparing and contrasting these multiple sets of feelings, desires, intentions, and motivations for the purposes of conflict resolution" (Anderson Forthcoming, 8).

Anderson's concept of hermeneutic labour is useful to make sense of the other type of labour that patients have to perform in their interactions with HCPs. In the medical context, the patient is expected to put in the effort to make sure that they communicate clearly in a short amount of time to maximise the benefits during the interaction. As Wendell explains, "since modern medical science does not exercise much modesty about the extent of its knowledge […] it has a tendency to ignore, minimise the importance of, or deny outright any of my bodily experiences that it cannot explain." (Wendell 1996, 122). This forces the patient to perform hermeneutic labour, as they need to learn how to present their symptoms in a way that will be well received by their HCPs, given that their phenomenal experience is not considered essential to establish the diagnosis (Wendell 1996).

Hermeneutic labour in the context of healthcare therefore refers to the burdensome activity of:

a) *Understanding one's symptoms and what one wants for treatment and presenting them in an intelligible fashion to others when deemed appropriate;*

b) *Discerning others' intentions, and motivations by interpreting their verbal and nonverbal cues, including cases when these are minimally communicative or outright avoidant;*
c) *Comparing and contrasting these multiple sets of desires, intentions, and motivations for the purposes of advancing one's treatment and care.*

There is another reason to use Anderson's concept of hermeneutic labour in the context of healthcare. Indeed, hermeneutic labour encompasses the idea that "hermeneutic labour primarily involves patient, deliberative reflection, and is generally undertaken in solitary rumination and/or in conversations outside of the situations on which it labours, as in conversations with friends or counsellors. Hermeneutic labour reflects on social encounters after they occur, and prepares plans for future encounters." (Anderson Forthcoming, 3-4). Whilst Anderson uses intimate heterosexual relationships as her main example, the same can be said of patients who have lesser-known illnesses. Given that many lesser-known illnesses are poorly understood, patients navigate an environment where the assumptions about who possesses the best knowledge in the situation doesn't track the reality. Indeed, the HCPs, with their scientific methods of investigation might not always be the best placed to gain knowledge about the situation. For that reason, there is a burden on the patient to reflect about how to best present their needs, without bruising the ego of their HCPs, and without appearing too knowledgeable, as they otherwise risk being deemed anxious or hypochondriac. Worse, if a patient gets emotional during the interaction, their symptoms might be attributed to a mental condition, thus not being properly investigated. To avoid being branded a 'weeping woman' (Walker 2022), the patient must perform hermeneutical labour, especially when one's symptoms are painful, and one struggles to be taken seriously. Finally, many patients rely on peer networks (for instance Facebook groups or online communities) to dissect their experiences of the healthcare system. They often come up with strategies to improve their experience at the next encounter with a HCP by discussing with peers. Just like women asking their friends how to best get their partner to understand them, patients ask other patients how to best get their needs met by their HCPs. Hermeneutic labour is therefore an inherent part of the patient's work when

interacting with HCPs. Not performing such labour puts the patient at risk of being dismissed or left untreated.

2.3 Illustration: Endometriosis

Endometriosis is a painful condition affecting 1 in 10 people with a uterus, in which tissue similar to the lining of the uterus starts to grow in other places of the body. Every month, this tissue bleeds in the same way as the one in the uterus, but unlike the latter, it has no way to escape, causing painful and/or heavy period (Endometriosis UK 2022). The average time to get a diagnosis from the first-time patients consult their doctor is 7.5 years in the UK (Endometriosis UK 2022). In Germany and Austria, patients wait similar times, with 7.7 years from the first gynaecological consultation and 10.4 years from first seeking medical help (Hudelist et al. 2012).

A study conducted on a group of Brazilian women highlighted that patients consulting for pelvic pain wait on average 3.4 years longer than patients consulting for subfertility to get diagnosed with endometriosis (Arruda et al. 2003). The main difference between the two groups is the main symptom: in one case, the pain is invisible, while in the other case, there is an 'observable' symptom like infertility.

In the case of pelvic pain, the need for intersubjective verification which would satisfy the demands of mechanical objectivity is impossible to provide. For patients with pelvic pain, the only symptom that they can point to is their pain. And pain is a subjective bodily experience. Indeed, the knowledge that the patient possesses here is that this pain is different from the period pain that they used to experience, as the tissue growth happens. It is only through their subjective phenomenal experience that they can obtain this evidence. Given the disagreement about phenomenal experience being a source of evidence, patients struggle to convince HCPs to take their testimony seriously, thus leading to longer waiting times, as referrals for further examinations and tests might be refused[36]. Moreover, in some cases, HCPs might not be trained to know what

[36] See chapter 4 for a longer discussion of this issue.

endometriosis is or might have false information about it. Two of their most common shortcomings involve believing that severe period pain is normal, or that endometriosis is rare in teenagers and young women (Wood, Guidone and Hummelshoj 2016). Given that HCPs are seen as the experts, challenging these beliefs might be extremely difficult due to their perceived expertise, thus leading to longer waiting times for a diagnosis.

These difficulties result in increased epistemic and hermeneutic labour. Let us start with epistemic labour. For pre-diagnosis epistemic labour, many patient-led endometriosis organisations recommend to people seeking a diagnosis that they keep a pain diary to help with the diagnosis process (Endometriosis UK 2022). It also involves knowing about various myths around endometriosis to be able to debunk them during assessment appointments. Post-diagnosis epistemic labour might entail many different things. Patients should know about the treatment options and possess general knowledge around the illness, for example know that getting pregnant is not a cure for endometriosis. Finally, the epistemic labour post-diagnosis entails educating other patients and medical professionals about the condition and advocating for the implementation of updated guidelines. Without this epistemic labour, the care of people living with endometriosis is greatly diminished.

Regarding the hermeneutic labour done pre-diagnosis, patient with endometriosis have to carefully choose which symptoms to invoke to get treatment in order to avoid triggering the biases. For example, talking solely about pelvic pain might trigger bias about the idea that periods are supposed to be painful, thus leading to a credibility deficit. As Wendell explains, there is an incentive to make patients use descriptions that are of recognisable and definite symptoms (1996, 133). For instance, explaining that one has blood in their stool only during menstruation would be a way to avoid bias and present a scientifically definite description. Additionally, they must tread carefully in order not to appear ‘hysterical’ or to look hypochondriac, especially given medicine’s well-known tendency to attribute gynaecological problems to issues within the mind (Akbar 2021). Indeed, some clinicians assert that endometriosis is caused by poor mental health (Young et al. 2017). Thus, when interacting when HCPs, people who desire

to pursue a diagnosis of endometriosis should be careful about the way they present themselves and their symptoms, or they risk being labelled as mentally ill, which prevents any further investigation into their physical symptoms (see Turner (2022) for her story about being diagnosed with bipolar disorder instead of endometriosis). Using lived experience as their sole evidence exacerbates this issue, as they are not able to provide evidence that the HCP would accept, thus increasing their credibility deficit. All of this leads to increased hermeneutical labour, which becomes burdensome. The situation is difficult to change due to the presence of disagreement about what counts as evidence and expertise. Thus, the sole solution present for the patient is to learn how to best present themselves and their symptoms, which requires a lot of preparatory work.

One might argue that people still get diagnosed with endometriosis and therefore that disagreement between the HCPs and the patient might not be present. However, what is interesting is the difficulty for patients to get a solution, even after multiple appointments with a HCP. Indeed, even if the HCP is open to rational dialogue, they nonetheless persistently fail to take many of the patients with endometriosis seriously. Moreover, the disagreement appears persistent. In some case a partial solution might be reached, that is, a practical solution is found. For instance, the patient may finally get diagnosed. However, we can observe that in many cases the difference in belief is still there, as the next time the HCP meets with a new patient, they might still behave in the same way and dismisses pelvic pain. The disagreement is still there, even if it is practically resolved in individual cases, which contributes to the sustainability of the epistemic mechanisms that lead to burdensome epistemic and hermeneutic labour.

3. Patient Activation

Given what has been discussed around epistemic and hermeneutic labour, I would like to further discuss how these two types of labour relate to the idea of patient activation. As explained briefly in the introduction, patient activation is a measure of a person's skills, confidence and knowledge to manage their own health and higher activated individuals typically have better health outcomes, as they have better care experiences

and lower costs (Bu and Fancourt, 2021). Patient activation is often seen as a positive thing that should be encouraged in healthcare. Moreover, acquiring knowledge and confidence helps patients exercise their autonomy and participate in practices of informed consent. Both are seen as extremely positive aspects to respect patients.

However, we have now seen that the healthcare system requires epistemic and hermeneutic labour from various patients, not just to improve some already good health outcomes, but rather simply to survive, avoid medical neglect and dismissal and improve diagnostic times, like in the case of endometriosis. To be able to reach high levels of patient activation, patients require knowledge, skills and confidence, yet these can be achieved only through epistemic and hermeneutic labour. If epistemic and hermeneutic labour is burdensome and unfair, as patients are not taken seriously once they do the labour, then patient activation becomes the opposite of empowerment. Even worse, selling patient activation as a form of empowerment might be a way for the healthcare system to hide some of its epistemically exploitative nature. Let me now flesh out this idea, by first introducing some of the literature around epistemic exploitation and then explaining how it applies to healthcare.

Lorde introduces ideas around epistemic exploitation in her paper *Age, Race, Class, and Sex: Women Redefining Difference,* delivered in 1980. In it, she highlights three main facets of this labour done by members of oppressed groups to educate people in dominant positions. First, members of oppressed groups are "expected to stretch out and bridge the gap between the actualities of our lives and the consciousness of our oppressor" (Lorde 2019, 114). She talks about how oppressed people have to become familiar with the manners and language of the people who are in a position of power, echoing the idea of hermeneutic labour. Second, the responsibility falls on the oppressed to highlight failures in the system and educate members of the dominant group. In contrast, members of the dominant group can evade responsibility and reinforce their position of privilege. These dynamics mirror the one discussed around epistemic labour in healthcare, as patients work tirelessly to improve knowledge around their illnesses and debunk myths. Third and finally, this labour is "a constant drain of

energy" which might be used more productively towards projects for oneself (Lorde 2019, 115). This final aspect parallels the idea that epistemic labour and hermeneutic labour are burdensome and take out a lot of energy from patients, as they are necessary to secure good care.

Berenstain proposed a more recent precise definition of epistemic exploitation. She defines it as "a variety of epistemic oppression marked by unrecognised, uncompensated, emotionally taxing, coerced epistemic labour" (2016, 570). In general, epistemic exploitation often refers to the forced production of education by marginalised individuals for the benefit of privileged people, echoing Lorde's proposal. For instance, a Black person who is forced to educate a white person about the harms of blackface is epistemically exploited into providing knowledge to the person in a position of power (Berenstain 2016). An interesting insight from Berenstain's work are the three core features of epistemic exploitation that she highlights. First, epistemic exploitation is constituted by labour which is costly, unpaid and often unacknowledged (Berenstain 2016, 572). Second, the labour is expected to come from the marginalised, which creates a double bind: "marginalised persons often do not have the option to simply disengage from an epistemically exploitative situation without being subjected to harm as a result of their perceived affront". (Berenstain 2016, 576). And in cases where they decide to disengage and not produce the required knowledge, members of marginalised groups cannot criticise the system, because they are seen as failing to actively work to change it (Berenstain 2016, 578). Finally, when marginalised people share their lived experience and offer their knowledge, it is met with scepticism. For instance, their experience might be doubted, or it might be questioned whether it falls within a larger pattern of oppression (Berenstain 2016, 578).

The experiences of people with lesser-known illnesses are very similar to the one of marginalised people in the context of oppressive structures that Berenstain describes. Indeed, the epistemic and hermeneutic labour produced by people with lesser-known illnesses fulfil Berenstain's core criteria of epistemic exploitation. Regarding the first one, as we have discussed, their labour is unpaid, often unacknowledged, and mostly

importantly, costly in terms of time and energy spent. Second, these patients are in a double bind: either they educate themselves and try to educate HCPs, or nothing will change. If they decide not to educate themselves and others, they risk being dismissed and misdiagnosed. Worse, if they don't appear knowledgeable enough, their poor care might be explained as a failure to reach high enough levels of patient activation. Finally, many people with lesser-known illnesses are doubted on two fronts: their symptoms might be doubted or ascribed to some mental illness instead of a physical problem; and/or they might be doubted when recounting a bad interaction with a HCP or producing a complaint about how they were treated. The double bind stems from the authority of HCPs over patients, and their failure of practitioners to see their own shortcomings when it comes to diagnosing and caring for certain patients.

One might object that epistemic exploitation in healthcare doesn't truly count as a form of exploitation. In the traditional literature on exploitation, exploitation is characterised by cases where the exploiters wrong their victims to maximise their own gain. In short, the exploiter must benefit from the wrongdoing for it to count as exploitation (Malmqvist and Szigeti 2019). In the case of HCPs in healthcare, it seems inaccurate to describe them as benefitting directly from the exploitation of patients. However, this traditional model doesn't account for cases where a third-party benefits from the exploitation. As Malmqvist and Szigeti (2019) explain, exploitation can have third-party beneficiaries, and these beneficiaries might be unaware of the exploitative conduct. In healthcare, HCPs themselves don't directly benefit from the epistemic and hermeneutic labour done by some patients. However, other patients and the general healthcare system do. Indeed, epistemic and hermeneutic labour reduce healthcare costs and frees up some resources which can be used for other patients or infrastructure.

Given the erasure of the labour performed by people with lesser-known illnesses and the presence of the disagreement about expertise and evidence within healthcare, epistemic exploitation can arise. Unfortunately, as these exploitative epistemic mechanisms are not talked about within healthcare, there is a risk that championing patient activation might lead to patient exploitation. Currently, patient activation

seemingly promotes autonomy and empowerment in patients. In reality, it hides the epistemically exploitative nature of the current healthcare system when the responsibility is put on the patients to activate themselves. Unfortunately, it is beyond the scope of this chapter to discuss where exactly patient activation becomes a trade-off that leads to patient exploitation. It would be interesting to explore when it becomes a loss not to empower the patient and when it becomes a benefit in order to avoid exploitation. Alternatively, if changes could be enacted so that HCPs are the one supporting patients to be activated and patient activation is scaffolded, these negative effects might be avoided.

4. Conclusion

In this chapter, I have argued that patients, in particular patients with lesser-known illnesses, perform epistemic and hermeneutic labour to get good care. This labour is burdensome, and patients are forced to perform this labour to achieve acceptable health outcomes. Furthermore, they face an injustice if they are silenced after performing this labour, as their expertise and evidence are downgraded. Whilst the political and material issues facing healthcare do certainly contribute to epistemic and hermeneutic labour, I have argued that we need to attend more closely to the epistemic mechanisms behind some of the interactions between patients and HCPs. In addition, I have shown that highlighting the presence of disagreement about expertise and evidence between HCPs and patients helps us make sense of cases like the one of diagnostic delays for patients with pelvic pain who have endometriosis. I finished by raising concerns about patient activation, as it might foster a climate of epistemic exploitation if it is the responsibility of the patients to be activated. Patients are encouraged to perform the labour necessary to exchange with their HCPs and navigate the healthcare system. Yet, without addressing how epistemic and hermeneutic labours can be mitigated and made optional, the healthcare system risks turning patient activation into patient exploitation.

Chapter Six:
Dialectical White Scepticism

But the problem with white sceptics is not their scepticism, per se.
It is the selectiveness of their scepticism.[37]

Myisha Cherry

Introduction

In this chapter, I introduce a new subtype of white scepticism called *dialectical white scepticism*. White scepticism, as a general concept, is defined as the phenomenon of people exhibiting resistance towards testimony about racial injustice, particularly made by members of racially oppressed groups (Cherry 2021; Mills 2007). As I will explain in more detail, dialectical white scepticism, as a subtype of white scepticism, is characterised by the argumentative concerns of the person who resists the testimony. These concerns appear to be motivated by epistemic and dialectical reasons rather than white privilege, thus making it difficult to recognise it as a form of white scepticism. Indeed, these concerns appear to be initially warranted: if a challenge is raised because the justification is judged to be insufficient or not satisfactory, our traditional argumentative practices recommend that one has to defend the assertion in the face of the challenge (Rescorla 2009a, 2009b, 2009c). To explain why dialectical white scepticism puts unsatisfiable demands on racially oppressed people, I first detail in more depth what dialectical white scepticism is, how it relates to ordinary scepticism about common knowledge, and why it appears warranted for most agents. I then show why some beliefs about racial oppression cannot be justified inferentially, in a similar fashion to ordinary scepticism, which leads to unsatisfiable demands for arguments defending the assertion, as it is not possible to give an answer that wouldn't beg the question against the sceptic. Based on this structural parallel, I explore ways in which we could use this parallel to further argue that the demands of the white sceptic are

[37] (Cherry 2021, 44)

unreasonable, just like the demands of the ordinary sceptic, despite differences between the two cases. I conclude by sketching avenues for further research that could be explored to respond to the white sceptic.

1. Racism

Members of racially oppressed groups often need overwhelming evidence to convince other people that they have been racially oppressed or that a particular situation is racially oppressive. This need for overwhelming evidence is a direct consequence of the nature of racism itself. Racism, understood as racial oppression, can be defined in the following way, where oppression is defined following Cudd's definition of oppression: it is a harm where the group of racialised people is "systematically and unfairly or unjustly constrained, burdened, or reduced by any of several forces" (2006, 23). As Cudd explains, members of privileged groups do not need to see themselves as oppressors nor intend to gain unjustly from their actions or omissions to benefit from their privilege (2006). With respect to racism in particular, Pierce explains that actions to maintain oppression are often "subtle and stunning", where an individual oppressive action is seen as "innocuous", but the cumulative effect to the victim is immense (1970, 266). As Pierce explains, racism is akin to an American Football match, where the offence is most effective when deceptive: "the best offenders can cheat on the call, so that they are into action, illegally, with such speed and precision that a referee can't disqualify them even if he is suspicious of the action" (1970, 270). Racism, in its current form, is not most effective when obvious, as most people, including white people who benefit from it, would call out the obvious occurrences of racism as unjust. However, when subtle, it can permeate everyday life in ways that maintain white privilege, while going unnoticed by the majority. It is this type of racism, with its subtle everyday instances that uphold the systematic oppression of racialised people, that this chapter will focus on.

2. Dialectical White Scepticism

Given the subtle nature of some oppressive actions that maintain racial oppression, testimony about this form of racial injustice is hard to justify with evidence that will be obvious to all agents, often leading to argumentative challenges. When people don't

accept testimony about racial injustice and exhibit dialectical white scepticism, they don't think that they do so because they are racist. Rather, they purport that the reasons offered to support the assertion about racial injustice are insufficient and need to be defended further. In their eyes, the grounds for rejection are epistemic and dialectical rather than motivated by white privilege, and therefore they are not engaging in white scepticism. Indeed, one commonly accepted norm of argumentation and dialectical exchange is that agents should be ready to defend any assertion they make when challenged, or they should retract it if they are unable to meet the challenge (Rescorla 2009b). Dialectical white scepticism thus has two defining features: first, the agent who engages in dialectical white scepticism appears well-intentioned towards racially oppressed people and open to a fair dialectical exchange, where both agents engage in reasoned discourse[38]; second, the agent presents their resistance as being motivated purely by epistemic reasons, rather than by a form of white privilege which would benefit them. Their resistance appears rooted in dialectical concerns: when people engage in reasoned discourse, they should be open to being challenged and must meet the challenge when presented with it or retract their assertion (Rescorla 2009b). In short, their resistance appears to be neutral from a moral and political point of view, which validates this resistance as the result of their good epistemic agency and as agents who abide by the norms of good dialectical exchange.

Let us consider a fictional case illustrating this dynamic: John, a Black man living in the UK, was followed at the local supermarket by an employee during his most recent weekly shop. When John recounts his shopping trip to his friend Kate, he explains that he thinks that this incident was an instance of racial profiling. Kate doesn't think that this is necessarily the case and asks John for more information. John explains that in the area where they live, racial profiling is quite prevalent. Upon hearing this justification, Kate asks John what his justification for this particular case is. Indeed, she thinks that

[38] Here, reasoned discourse refers to the activity through which agents engages with each other by providing arguments and counterarguments (Rescorla 2009b).

this could have been an innocent mistake, where the employee is new and overeager (and where they might follow all customers in the shop), thus not making it necessarily a case of racial profiling. By making this remark, Kate presents a dialectical challenge to John: he now must either defend his claim that it was racial profiling, or he needs to retract his assertion. John takes up the challenge and starts explaining to Kate that given the statistics about racial profiling and the nature of racist implicit biases, he has no reason to think that this was an innocent mistake, contrary to what Kate thinks. This last argument still does not satisfy Kate, who asks John how he knows that the statistics are reliable. She continues to challenge John every time he brings out more evidence to support that this was an instance of racial profiling, making further demands for John to defend his assertion which he cannot fulfil, as any further defence he could provide would seem to beg the question. Kate, like Lewis Carroll's tortoise (Carroll 1895), will never be satisfied by John's explanation and will continue to ask for further explanation indefinitely, under the guise that she is open to being convinced and that a way to convince her exists.

Let us now turn to a useful comparison. Indeed, the problem of unsatisfiable dialectical demands is a traditional one in epistemology (Rescorla 2009b). In these core cases, it is not the existence of oppression that the agent tries to prove, but rather the existence of the external world. For instance, they try to defend the knowledge that we have two hands. However, these beliefs are challenged by the sceptic who asserts that we cannot know that we have two hands if we cannot rule out an alternative explanation, like the possibility that an evil demon is deceiving us into thinking that we have two hands. However, most people would assert that we know that we have two hands (Moore 1939). Whilst the alternative explanation cannot be fully ruled out, we often accept that given that we don't currently have any positive counterevidence to the belief that we have hands, this belief is justified and does not need further defence in a dialectical exchange (Rescorla 2009b). Let us compare the following two assertions:

A. "I have two hands."

B. "I was racially profiled by the employee when shopping this morning."

In case A, the unreasonable demands for defence often take the following form: first, the assertion is challenged, by invoking an alternative explanation, like the possible existence of an evil demon: if you cannot rule out the possibility of an evil demon, then how can you know that you have two hands? Any additional arguments given to favour the initial hypothesis that one has two hands will be rejected in a similar fashion. For example, one could try to argue that they have not seen any evil demon or don't have any evidence in favour of its existence. However, here this additional defence is refuted, by explaining that if the evil demon was here, then one would not be able to see it anyway or would not have evidence to believe that it was there, as that is one of the core features in the evil demon case. Thus, any additional argument will be rejected by the sceptic as it seems to beg the question.

In case B, the demands for defence take a similar form: the assertion B is challenged by invoking an alternative hypothesis, like the evil demon hypothesis. Here the idea that the individual incident (being followed by the employee at the supermarket) is not linked to the wider pattern of oppression (racism). For instance, if one cannot rule out the possibility that the employee was just new and overeager, then one cannot know that this incident was racially motivated; then, if an additional layer of argumentation is added, it is immediately challenged on the basis that it does not provide sufficient evidence to reject the alternative hypothesis. For example, John tries to say that he has no evidence in favour of it being an honest mistake, and on the contrary has more reasons to believe that it is a case of racial profiling. This argument could be challenged by saying that if the employee was new and overeager, it would not appear to be any different from the racially motivated case, and John would struggle to point to the difference between the two cases to a third party. Therefore, he would still not have any way to prove that it was not racially motivated, even if it is, given the subtle nature of many racist incidents. Indeed, the subtlety of some racist incidents makes them hard to prove, as the victim is only capable of recognising them due to their suitable position (Cherry 2021). These incidents are difficult to distinguish from innocent cases for agents who do not endure racism themselves. Just like the evil demon scenario is designed to prevent agents from invoking tangible shareable proof that the real world exists, racism

in its subtle forms follows the same design. In both cases, the assertions “I have two hands” or “I have been followed for racially motivated reasons” are hard to justify inferentially, as support for them would beg the question against the sceptic. For instance, trying to say that one knows that they have two hands because their vision has been working just fine lately would be a form of justification which would beg the question, as it uses the same method which has been put into question. In the case of racism, appealing to the statistical data would not help John, as the way the data was collected is also based on testimony made by racially oppressed people, thus begging the question against the dialectical white sceptic and opening the door to more challenges.

If we summarise both cases, the main idea is that if one cannot defend themselves against the challenging hypothesis (the evil demon/the over eagerness of the new shop employee) for why one has an experience (having two hands/being followed by the shop employee), then one cannot defend the assertion that they made (I have two hands/I was followed due to racial profiling). In the ordinary knowledge case, the evidence doesn’t enable one to distinguish between the hypothesis that one has real hands and the hypothesis that the hands are an illusion caused by an evil demon, so one isn’t justified in defending the first hypothesis over the second. In the case of racism, the evidence available doesn’t enable one to distinguish between the hypothesis of it being racially motivated and the hypothesis that the shop employee is new and over eager, so one isn’t justified in defending the first hypothesis over the second. In both cases, there is an initially plausible epistemic closure principle, where if p entails q, and one cannot know that q, then one cannot know that p: having hands entails not being deceived into thinking you have hands by an evil demon, and if one cannot know that they are not deceived, then they cannot know that they have hands. Being followed for racially motivated reasons entails that the shop employee is not just overeager. As one cannot know that the shop keeper is not just overeager, then one cannot know that they are followed for racially motivated reasons. Dialectical demands to defend assertions about common knowledge and knowledge of racism are thus difficult to satisfy without relying on the claim or method of acquiring the justification which is doubted. By using this

plausible epistemic closure principle and the unavailability of non-circular justification in these cases, both types of sceptics have the perfect recipe for an unanswerable challenge.

When trying to defend assertions about subtle individual occurrences of racial oppression, no inferential justification can be offered, as it is impossible to justify some higher-level belief detailing the many ways in which racism can occur without appealing to other beliefs about these individual occurrences themselves. Consider a contrasting case: in a different shop, the employee shouts a racist slur at John. In this case, the incident is clearly racist, and it is much easier for John to prove that this incident is part of the larger pattern of racism, and another external agent could observe it too. The argument to follow excludes clear instances like the racist slur incident. In many cases of everyday racism where the incidents are subtle, it isn't as easy to show that the incident is part of the wider pattern. These last types of cases are the ones where dialectical white scepticism will be the most resistant. As the explanation of the type (racism) rests on the token (individual occurrences of racism), and the token cannot be fully explained without the type, agents struggle to make sense of their experience and to justify their beliefs related to racial oppression in a way that would not appear circular to the dialectical white sceptic. One should note that many people accept that the existence of racism as a wider pattern can be justified through empirical facts, like statistics about racial violence, higher rates of health issues, poverty indexes, etc. (Cherry 2021; Lagewaard 2021) However, proving that a certain individual occurrence is linked to this wider pattern is the crux of the issue here. Indeed, recognising that a particular incident is not isolated is not the same as recognising that there are incidents in general, and it requires some specific epistemic standing that only people who are the targets of this pattern can recognise. For instance, the frequency of everyday racial violence can only be appreciated if one is a person of colour, as a white person might not realise the extent of the phenomenon if they are not the target (Cherry 2021). Demanding of agents to prove this link in a way that doesn't appear to beg the question against the dialectical white sceptic is impossible, as explained earlier.

3. Solutions

Given the strong parallels between ordinary scepticism and dialectical white scepticism (it is impossible to satisfy both types of sceptics without begging the question), I will now argue that we can extend the structural strategy of defence against the ordinary sceptic to the case of dialectical white scepticism. When confronted with the ordinary sceptic, one common strategy is to say that we don't have to defend ourselves against the sceptic in some relevant context (Lewis 1996). What truly matters is that the beliefs were formed in an epistemologically appropriate manner. For instance, beliefs formed through a reliable process, or based on some basic beliefs don't have to be justified against the ordinary sceptic, unless they present some positive evidence that legitimates the challenge (Rescorla 2009b, 2009c). Without such positive evidence, the challenge of the sceptic can be labelled as unreasonable. I will now present some ways members of racially oppressed groups form beliefs about everyday racism in an epistemologically appropriate manner. This legitimates the use of the same response strategy between the two cases, given this relevant structural similarity between beliefs about the existence of the world and the existence of everyday racism. This leads me to argue that the demands of the dialectical white sceptic should also be considered unreasonable unless they can present positive evidence to legitimate the challenge.

In the following two subsections, I explore two options about how racially oppressed form their beliefs in an epistemically appropriate manner: first by discussing epistemic conservatism; second by discussing three ways racialised people acquire beliefs about racism through reliable methods. These ways are inductive reasoning, statistical reasoning and the Reflective Model of affective experience. I contrast the case of a racially oppressed person using these methods with the case of agents using the same methods to participate in racial oppression. I explain why the two cases differ, based on considerations about risk and the difference in the application of recognition-primed decision and use of affective experience, which enables me to argue that the belief of the racially oppressed person is supported whilst the one of the oppressor is not.

3.1 Epistemic conservatism

The first common strategy to stop an ordinary sceptic is to ascribe to epistemic conservatism. Epistemic conservatism is the idea that some beliefs are *prima facie* justified because they are held by agents. Moreover, these beliefs constitute the background for our other beliefs (Koons 2006). For instance, the belief that the external world exists is *prima facie* justified. A modern version of epistemic conservatism has been presented by Koons (2006). In his version, epistemic conservatism is defined in the following way: "basic beliefs are justified by their history, and so an epistemic agent is entitled to these beliefs even if she is not able to provide any inferential justification for them" (Koons 2006, 203). In this account, basic beliefs are defined as beliefs for which no inferential justification can be offered. Moreover, the justification has two main components: it is diachronic and social in nature. Indeed, the justification of basic beliefs relies on their ability to have faced the tribunal of inquiry and survived, through a historical process that takes place for generations within an epistemic community. Additionally, even though justification is conservative, it doesn't mean that beliefs are immune to revision: indeed, for this form of epistemic conservatism to work, no belief can be immune to *de jure* revision (Koons 2006, 205-206).

For instance, the belief that humans ordinarily have two hands can be justified the following way: humans are entitled to the belief that they have two hands, as this belief has not been proven wrong for generations. Even though people like sceptics have tried to show that this belief should be abandoned, the epistemic community has nonetheless kept it, as they reached the conclusion that this challenge was not strong enough to get rid of the belief. Given that the belief 'humans ordinarily have two hands' has survived a generations-long process and multiple challenges, it is justified through its history. Contrast this case with the hypothesis of a geocentric universe. As Koons explains, it was in the 1920s that Edwin Hubble discovered that galaxies are receding from the earth (2006, 204). Prior to this discovery, the hypothesis that the earth was at the centre of the universe was not discussed. It was rather used as a hypothesis against which other hypotheses were tested. However, when presented with Hubble's positive evidence,

scientists decided to discard the geocentric universe hypothesis (Koons 2006)[39]. In this case, humans were entitled to the belief that the universe is geocentric until they were proven wrong. As Koons would say, the geocentric universe hypothesis did not withstand the "test of time and evidence", whereas the hypothesis that the external world exists has yet to be proven wrong (Koons 2006, 205).

Similarly, the higher order belief about how racism is constituted and how individual occurrences are linked to the wider pattern is a good candidate for this type of justification. This general belief has a content of the following sort: "situations with characteristics x, y, z are instances of racism". The characteristics x, y, z are determined contextually by the group of racially oppressed people, and shift throughout the historical context. For instance, such characteristics could be 'causes harm to an agent because they are a person of colour' or 'unjustly burdens people of colour'. This general belief containing these characteristics is one that generations of racially oppressed people have held, and it has withstood the external pressure made by the white supremacist society to modify it (Pierce 1970; Cherry 2021). Moreover, when trying to justify this belief, no inferential justification can be offered, as it is impossible to justify the belief whose content is detailing the many characteristics of racism without appealing to other beliefs about individual occurrences of racism. Finally, this belief is not immune to *de jure* revision, as it could have been modified by the relevant epistemic community, here racially oppressed people. For instance, we could discover that some of the characteristics which seemingly participate to the oppression of racialised people are not oppressive and can be explained by a phenomenon other than racism. The basic belief about how racism is constituted and the link between the wider pattern and the characteristics of individual occurrences is justified non-inferentially according to epistemic conservatism. If this belief is not revised despite being subjected to a generation's long scrutiny by the community- here racially oppressed people, then it is

[39] Note that Copernicus argued for heliocentrism and against geocentrism long before the discovery by Hubble. Yet Copernicus was not taken seriously due to the religious beliefs held at the time.

justified. Thus, beliefs about individual occurrences of racism are basic in the relevant sense and so the regress of justification can be stopped.

3.2 Reliable Methods

In the main case presented earlier in this chapter, John, a Black man, believes that he was racially profiled by a supermarket employee. I argue that John has a justified belief, even if he cannot fully defend it in the face of the dialectical challenges brought by his friend Kate, who thinks that it is just an innocent mistake on the part of the supermarket employee. In the first part of this subsection I explore how John seems to use a form of inductive reasoning to justify his belief that he was racially profiled by the shop employee. However, to do so, I need to show how his inductive reasoning differs from the one used by the shop employee, and how it differs from statistical reasoning. At the end of the sub-section, I offer to use the Reflective Model of affective experience to provide an alternate explanation to show why John's belief is formed reliably.

To start, let us look at various cases showcasing the various types of reasoning used.

<u>Inductive reasoning cases:</u>

> **Case 1.i:** John, a Black man, is followed by a shop employee. In the previous experiences, e1, e2 and e3, John has been followed by shop employees in the past. In experience e1, e2, and e3, he has clear evidence that it was racial profiling, as his white friends went into the same shops right after him but were never followed by the same shop employees. Based on the previous times where he was followed, which were clearly racially motivated, he infers inductively that he is followed for racially motivated reasons in this new case, even though he cannot confirm his hypothesis as his white friends are not there with him. He forms the belief that he was racially profiled by the shop employee.

Case 2.i: Andy, the shop employee, has caught shoplifters in occasions f1, f2 and f3. In all three occasions, the thieves were people of colour. Based on inductive reasoning, he forms the belief that it is likely that the next person to shoplift will be a person of colour. For this reason, he follows customers of colour around the shop to prevent further shoplifting.

Statistical evidence cases:

Case 1.s: John, a Black man, is followed by the shop employee. He knows that racial profiling is extremely prevalent in the neighbourhood where he is shopping. Based on this statistical knowledge, he forms the belief that he was racially profiled by the shop employee.

Case 2.s: Andy, the shop employee, has been told during his training that people of colour are more likely to shoplift. Based on this statistical knowledge, he follows customers of colour around the shop to prevent shoplifting.

What are the relevant differences between Andy and John in each pair of cases? When comparing both inductive cases, it appears that John and Andy use the knowledge they acquired in previous situations. In both cases, John and Andy seem to be justified in their belief given that inductive reasoning is accepted in many cases. However, could we say that both of them are justified in believing to the same degree? Wouldn't we want to say that Andy should have a lower degree of credence compared to John about the idea that the next person who will shoplift will be a person of colour?

There are two reasons to think Andy should lower his credence compared to John. First, Andy might be prey to a form of confirmation bias in this situation. If he only follows customers of colour, he is less likely to catch shoplifters who do not fit this description. Hence, he is less likely to update his belief to reflect that not only people of colour are shoplifters. Thus, his behaviour serves to reinforce the evidence available to him as he

will not be confronted to any defeating evidence. On the contrary, we have more grounds to believe that John is not in the same situation. If he were not followed, he would not form the belief that he is being racially profiled. Moreover, if he was acting suspiciously, he would probably notice and form the counter belief that he is followed because he is acting suspiciously rather than because he is Black. Thus, John does not suffer from the same confirmation bias as Andy.

Moreover, there are reasons to believe that John's inductive reasoning can be likened to a form of expertise. Kahneman and Klein (2009) and Kahneman (2011) explain that in skilled situations, the interaction of system 1 and system 2 enable agent to make recognition-primed decisions (RPD). According to this model, the intuition present here is a form of pattern recognition. Kahneman (2011, 237) uses Herbert Simon's definition of expert intuition: "the situation has provided a cue; this cue has given the expert access to information stored in memory, and the information provides the answer" (Simon 1992, 155). In John's case, the new situation he is presented with provides him with the cue that he might be followed. He accesses information about previous situation that he has stored in his memory, which in turn provides him with further justification that he is being racially profiled. However, what would make John an expert? Kahneman and Klein explain that for an expert intuition to be reliable, two conditions need to obtain: "first, the environment must provide adequately valid cues to the nature of the situation. Second, people must have an opportunity to learn the relevant cues" (2009, 520). In short, the environment must be regular enough to be predictable and there needs to be some form of prolonged practice in that environment.

How would John's case fit these two conditions? For the first condition, the presence of racial structures and subsequent racial oppression will be homogenous in a certain geographic area unless there is sufficient demographic renewal or strong political initiatives to modify them. If most people in John's neighbourhood have been raised in the same manner, it is likely that they will exhibit the same level of racial bias. Thus, the level of racial oppression experienced by John daily will be reliable. To satisfy the second condition, John needs to be able to learn the relevant cues. Let us consider the scenario

in which John is wrong because he misevaluates that the situation is not racially oppressive when it is. In this scenario, John is likely to learn fast that he was mistaken, as he might be on the receiving end of racial violence or might be prevented from accessing some opportunities due to racial oppression. For instance, if he doesn't consider that he might be stopped by the police due to racial profiling and might get shot if he appears threatening in the slightest way, he is more likely to get injured during an encounter with the police. For racially oppressed people, to be able to understand the cues from the environment and to be able to react to them in a quick manner is a question of survival (Cherry 2021). Even though the stakes seem a lot higher in the police case than in a supermarket environment, the underlying learning mechanism is the same. As a racialised person living in a racially oppressive society, John had to become an expert at picking up cues and information about racial oppression. Therefore, he has expertise and displays expert intuition about being followed at the supermarket. This enables him to recognise the pattern of racial profiling and to form the belief that he is being followed due to racial profiling. Therefore, the inductive reasoning displayed by John is not just that: it is actually a form of skilled intuition. As explained in chapter 2, ways of obtaining knowledge do not need to be infallible to count as legitimate. Indeed, most sources of knowledge are fallible, yet still provide justification. In this case, John's skilled intuition might fail in some situations, but that does not mean that it does not provide him with justification. John's belief that he was racially profiled is justified, even though it is defeasible if presented with positive evidence against it.

In contrast, Andy does not exhibit such skilled intuition. He is not responding to cues from the environment. Rather he bases his belief on information stored in his memory, without responding to a link between current cues and this background information. His reasoning is not dynamic like John's and amounts more to a statistical inference: if most customers he catches shoplifting are people of colour, then he is more likely to catch a shoplifter by following people of colour. However, if he only applies this reasoning, he is clearly not attuned to the relevant features of the situation: that the race of a shoplifter has no causal relation to shoplifting and might only be correlated to it. Thus, the way Andy forms his belief in the inductive reasoning case is not far off from his

reasoning displayed in the statistical evidence case. One might wonder whether Andy's belief formation process is just more prone to failure compared to John. Given that fallibility does not ground the exclusion of a way of obtaining knowledge, we need to explain why Andy's reasoning should not justify his belief.

To do this, we can move out of the domain of the purely epistemic to take into account moral and practical stakes by bringing in some ideas from moral encroachment (Bolinger 2020; Gardiner 2019). As Bolinger explains, reasoning based on racial generalisation is bad, because it "disproportionately expose[s] group members to the risks associated with mistakenly assuming stereotypical propositions, and so magnify the wrong involved in relying on such inferences without adequate justification" (2020, 2415). In case involving racial generalisations (like the one that most shoplifters are people of colour), Bolinger argues that one must abstain from accepting a certain belief, given the risk in the event of a mistake. Here, mistakes wrongfully harm the group of racially oppressed people, given that many people make these mistakes, which collectively create a strong hindrance to people of colour by maintaining structural oppression. Given that one has a strong moral reason to avoid contributing to the harm, the risk incurred by the mistake is too high to justify immediate accept of certain beliefs based on racial generalisations. Bolinger argues that the belief p can be accepted only if there is sufficient evidence to justify ignoring the possibility that non-p (2020). For instance, if Andy has sufficient evidence to justify why he ignores the possibility that John is not a shoplifter, then he can ignore it. For instance, if he sees John entering with an empty bag, but then leaving with a seemingly full bag, without paying. Or if John acts suspiciously and seems to hide things under his coat. All these signs would warrant Andy to believe that John is a shoplifter. But as one might notice, this new evidence clearly makes it the case that Andy is not relying on a racial generalisation, and rather uses perceptual evidence and reasoning to form his belief that John is a shoplifter and thus that he should follow him around the shop.

In the case of John believing that he is being racially profiled, the same worry about risk does not apply. Indeed, Andy is not harmed by John believing and acting as if he believed

that he was racially profiled. There are two reasons for this. First, agents are usually harmed by repeated exposure to the same risk. Therefore, one-time risk exposure does not have the same effect as the repeated exposure of risk (Bolinger 2020, 2424). Andy is not repeatedly exposed to a possibility of error that has detrimental effects to him. In his case, it's even likely that he would not even realise that John thinks that he is acting in a racially motivated manner. Thus, Andy does not suffer as a result of John's belief and he is not wronged. Second, John, unlike Andy, has sufficiently strong evidence to justify ignoring the possibility that Andy is an overeager shop employee given that Andy is not giving him any counterevidence. Indeed, Andy *is* following John. And given that it would be worse for John to assume that Andy's behaviour is not racially motivated if it is than to assume that the behaviour is racially motivated when it isn't, John has pragmatic reasons to maintain his belief that the behaviour is racially motivated in the face of the lack of evidence to suggest the contrary. The risk associated in the case of Andy is not strong enough to require John to avoid making any statistical generalisations. One should note that this response only works if we accept that there is genuine moral and pragmatic encroachment on the epistemic. However, this type of position remains controversial, and some readers might feel unsatisfied by this response.

For this reason, I would like to present an alternative option which does not rely on accepting moral encroachment. In chapter 2, I introduced my own model of affective experience, the Reflective Model. Under this conception, affective experience is a source of knowledge if an agent has rational control over using it, and they would stop using it if they didn't consider it to be rational in the relevant conditions. If an agent satisfies this, then a belief formed based on their affective experience is justified. Let me analyse at the case of John once more. It is likely that he felt anger or annoyance at being followed around the shop. John's emotion towards being followed counts as a source of knowledge because John is rational when using it. Indeed, he is responsive to the relevant defeaters and would stop using his anger to acquire beliefs if he presented with evidence that it fails to track. Moreover, he knows that he can usually trust his anger in racist situations because he went through various internal and social calibration

mechanisms[40]. This means that he is likely to notice the occasions in which his anger fails to track reasons and would stop using it in the relevant scenarios.

Which forms of calibration ensure that John engages in self-reflection? First, John has learnt whether his anger tracks racism on other occasions and he can engage in immediate intrinsic calibration or metacognitive intrinsic calibration. For instance, when he went to the shop with his white friends, he obtained evidence that his anger on this occasion tracked racist features accurately. Moreover, John has likely been confronted with social types of calibration before this incident. For example, by being asked by Kate whether he is sure that he was racially profiled or whether it might have been an over eager new shop employee, he engages in extrinsic regulation. When asked to engage in extrinsic regulation, John will question whether his anger reliably tracks racist features, and he will adapt and improve this ability if needed. Finally, by interacting with other Black people, John participates in collective resistant regulation, and he will learn more about characteristics of racism. This will enable him to fine-tune his emotional reactions and will give him more information about whether his anger accurately tracks racism or not. Given all of this, John would stop using his anger if it failed to track racism. We can thus say that his affective experience is a source of knowledge in these conditions and thus justifies his belief that he was racially profiled.

John's emotion was triggered by being followed around the shop. His anger alerted him to the racist features of the situation and enabled him to link the situation to the wider pattern of oppression. Just as we have seen in chapter 3 with the phenomenon of microaggressions, emotions play a useful role to link subtle instances of oppression to a wider pattern. Moreover, as Cherry (2021) explains, racialised people living a racially oppressive society pick up on cues that remain invisible to other agents due to the repetition of small racist actions. Through their anger, they learn about the racial reality. Andy doesn't learn about who is a shoplifter in the same way, given that he doesn't have

[40] See chapter 2 for more details on each type of calibration.

an affective experience which would count as justifying his belief. Indeed, even if Andy is afraid of shoplifters, he would not stop using his fear were it to stop tracking. As I have explained above, Andy is not responsive to the evidence available to him. If he were responsive, he would know that his fear is not reliable to identify shoplifters, given that it is rooted in racial bias and is prone to confirmation bias.

To summarise, John is justified in believing that Andy is racially profiling him, because he is either: using expert intuition in a reliable context; or basing his belief upon statistical data that doesn't risk harming the target of his belief; or using his anger as a reliable source of knowledge to justify his belief. On the contrary, Andy doesn't use expert intuition, is prone to confirmation bias and his fear of shoplifters would not count as a source of knowledge. Moreover, Andy's use of statistical data and use of racial generalisations doesn't justify his belief. Given the risk of harm and wronging, he should abstain from forming the belief that John is a shoplifter unless he can provide positive evidence to justify that John is a shoplifter. Given that he doesn't provide such evidence, his belief is not justified. For all of these reasons, Andy is not justified in maintaining the belief that John will shoplift and shouldn't follow John around the shop.

4. Conclusion

One important lesson from traditional epistemology and the literature on dialectical regress (Rescorla 2009a, 2009b, 2009c) is that the burden of proof is usually on the agent who challenges an ordinarily accepted belief, rather than on the agent who holds the belief. For instance, it is on the sceptic to provide positive evidence for the hypothesis of the evil demon as a good alternative explanation as to why we have hands, rather than on the person with the belief to prove that the evil demon doesn't exist. When it comes to beliefs about racism, that is not the case. Indeed, in our current society and epistemic practices, it seems that the burden of proof has been put on people from racially oppressed groups to prove that their beliefs about racism are justified, due to the political forces at play (Cherry 2021). Given the widely accepted existence of racism and the generally accepted hypothesis that people who are racially oppressed are best placed to recognise instances of racism (Cherry 2021; Lagewaard 2021), it appears odd

that the burden of proof is still on racially oppressed people. Indeed, given the strong similarities between the case of common knowledge and the case of knowledge about racism, it follows that both cases should be treated similarly. If we are to adopt a consistent set of epistemic rules and dialectical practices, the burden of proof should be on the dialectical white sceptic to provide positive evidence in favour of the alternative hypothesis. However, it appears that this discrepancy reinforces racial oppression by forcing racially oppressed people to spend time and energy in hopeless dialectical exchanges, where their objector appears to want to be convinced but will never be satisfied, just like Carroll's tortoise (Carroll 1895). This discrepancy has been widely accepted due to historical and social biases, and it contributes to the epistemic burden carried by racially oppressed people (Cherry 2021).

Given that dialectical white scepticism puts the bar too high for what it will accept as a good defence, it should not be accepted as *de facto* posing a legitimate challenge. Rather, the burden of proof should be on the agents who engage in dialectical white scepticism to provide positive evidence showing why their challenge is legitimate, in a similar vein to the challenges made by the sceptic about ordinary knowledge (Rescorla 2009c). The demands for defence in the case of dialectical white scepticism appeared warranted on a surface level, but as this chapter has shown, these demands are unreasonable and unsatisfiable. If it were the case that dialectical white scepticism is motivated by purely epistemic and dialectical reasons, then the agent would recognise that they are putting the challenged in an inescapable epistemic and dialectical situation, where their demands could never be satisfied. Maintaining the discrepancy between how we respond to different types of scepticism only serves to reinforce white privilege and furthers the epistemic burden carried by racially oppressed people.

Conclusion

In this thesis, my goal has been to argue that our epistemological discussions should include affective experience in our sources of knowledge. This goes against the current status quo in analytic philosophy, as affective experience is traditionally excluded from what analytic epistemologists consider an appropriate source of knowledge and justification. However, the legacy of feminist philosophy gives us strong reasons to question whether this exclusion is warranted, given the important role emotions play in recognising oppression. My main task was thus to build a model of affective experience as a source of knowledge which would withstand the scrutiny of analytic epistemologists, whilst serving the goals of feminist and resistant theory.

I started by assessing in **Chapter 1** the main arguments against using affective experience as a source of knowledge. First, affective experience tends to provide us with false information and seems in general to have quite a low reliability. Indeed, between cases of phobias, emotions being elicited by implicit biases, or affective experience leading to confirmation bias, affective experience seems particularly prone to distortions. However, as I argued, this reasoning would sound quite pessimistic if applied to other commonly accepted sources like perception. I then proceeded to show that the main issue of affective experience is that it is not as sensitive to defeaters as the other sources of knowledge. Indeed, the personal nature and phenomenal component of affective experience complicate how an agent responds to defeaters. This in turn leads to the possibility of conflict without contradiction, where an agent can have a conflict between the content represented in their affective experience and their belief. The possibility of conflict without contradiction explains why affective experience is not as sensitive to defeaters, as their affective experience might persist even when presented with defeating counterevidence. Affective experience does not fare well within the traditional model of knowledge acquisition, as agents can be rational yet unresponsive to defeaters for their beliefs acquired through affective experience. Thus, this combination of theoretical issues supports the exclusion of affective experience from our approved sources of knowledge. I complicated the situation by introducing two

applied cases which support the inclusion of affective experience in practical contexts. Without pain being included as a source of knowledge, patients face an important diagnostic delay when trying to get diagnosed with endometriosis. When it comes to microaggressions, agents rely heavily on their emotions to link individual incidents to wider structures of oppression through pattern recognition. This led to a tension between the theoretical arguments against affective experience and the practical concerns in favour of its inclusion.

In **Chapter 2**, I dissolved this tension by providing a conception of affective experience as a source of knowledge which accounts for the arguments presented in the previous chapter. After discussing Döring and Brady's theories and rejecting them because they failed to explain convincingly how affective experience provides immediate justification for beliefs, I proposed to build an alternative model. Taking stock from Jones, I started by highlighting the differences between reason-tracking and reason-responding abilities and argued that affective experience is a reason-tracking ability. Under this conception, I introduced the *Reflective Model of Affective Experience*: a rational agent is one who exercises reflective self-monitoring and would stop relying on affective experience if they realised that it failed to track. I argued that agents using affective experience can engage in reflective self-monitoring, through four main ways of calibrating their affective experience: immediate intrinsic regulation, metacognitive intrinsic regulation, extrinsic regulation, and collective resistant regulation. If an agent engages in these types of calibration, they would stop using their affective experience if it fails to track reasons. Calibration plays a dual role: first, through calibration, the agent realises that they failed to track; second, they can work to improve how their affective experience tracks or discard any belief formed based on a type of affective experience which fails to track. This new model centres the idea that we regulate our affective experience both internally and socially, given the deeply social world we live in.

I moved from the theory to the practice in **Chapter 3**. I showed how this idea of affective experience being both social and individual helps us characterise how agents detect individual microaggressions. The existence of microaggressions is often questioned due

to a lack of theories about how agents know that an incident is a microaggression. I offered to fill this gap with the *Emotional Account* of microaggressions: a microaggression has occurred when the agent who is the target of the brief everyday exchange could have had a warranted emotion towards the structural oppression present in that exchange. On this new account, agents know that a microaggression has occurred through their emotions, as their affective experience enables them to link the particular incident to a wider pattern of structural oppression. This provides an advantage over existing accounts of microaggressions, as they failed to explain the link between the individual perspective of the victim and the wider group they belong to. The Emotional Account includes a mechanism of *post hoc* education where the agent interacts with other group members, explaining how collective discussions about oppression are nurtured by individual experiences and vice-versa. Through the example of microaggressions, I showed that including affective experience as a source of knowledge makes it easier to make sense of the epistemic contributions of members from marginalised groups, such as how they recognise and conceptualise subtle forms of oppression, like microaggressions.

Chapters 1, 2 and 3 of this thesis constituted my positive argument for the inclusion of affective experience in our sources of knowledge. In the remaining chapters, I focused my attention on some of the practical and ethical consequences of the exclusion of affective experience. This enabled me to build a negative argument against the exclusion of affective experience: without the inclusion of affective experience in our sources of knowledge, we risk perpetuating injustice. Indeed, by upholding the epistemic standards that benefit the privileged and dominant group, we fail to recognise some of the oppressive features of our society. I developed this idea in chapters 4, 5 and 6.

In **Chapter 4**, I built the theoretical grounds to support this argument and introduced a new form of epistemic injustice, *source based epistemic injustice* (SBEI). SBEI is the idea that a knower can incur a loss of credibility due to an unfair prejudice towards unapproved sources of knowledge. I argued that SBEI stems from an over-reliance on intersubjective verifiability within the analytic tradition and disproportionately affects

marginalised groups, making it a type of epistemic injustice. I used this new framework to reconceptualise what goes wrong in the diagnostic delay experienced by patients with endometriosis who testify about their pelvic pain. Due to SBEI, their pain is not considered to justify further medical investigation. I then explained how SBEI fits within the wider framework of epistemic injustice. First, it is closely related to testimonial injustice, as it explains why some agents from dominant groups can suffer a credibility deficit if they use an unapproved source like affective experience, as these unapproved sources are often associated with members of marginalised groups. Second, SBEI provides a causal story for some cases of hermeneutical injustice. I showed that affective experience is often key to develop concepts, and to justify the application of said concept to various situations. *Microaggression* and *sexual harassment* are paradigmatic cases of this phenomenon, as members of marginalised groups developed these concepts to explain their experience. Yet these phenomena are doubted because they are hard to prove without relying on affective experience. This leads to hermeneutical injustice as these concepts developed through affective experience lack uptake in the dominant epistemological system.

At the core of SBEI is an insistence on intersubjectively verifiable evidence, which has been used to exclude affective experience as a source of knowledge. In **Chapter 5**, I explored the ramification of this idea within Western healthcare systems, and how it shapes the interactions of healthcare practitioners and patients with lesser-known illnesses. I problematised the idea of *patient activation*, arguing that encouraging greater levels of patient activation can lead to the epistemic exploitation of patients. Indeed, if agents are encouraged to know more about their illnesses, yet their knowledge can only be justified through their lived experience, they are likely to experience silencing. Given the focus of Western healthcare on mechanical objectivity and intersubjectively verifiable evidence, testimonies relying on subjective experience will be downgraded. I argued that this puts patients in a position where they are forced to perform epistemic and hermeneutic labour in order to bypass the credibility deficit and present their evidence in a way that will be well received by their healthcare provider. This need for epistemic and hermeneutic labour is unfair and burdensome, as

patients can perform such labours, yet not be taken seriously. I showed that the performance of epistemic and hermeneutic labour by some patients can fit the category of epistemic exploitation. This led me to conclude that championing patient activation can lead to the exploitation of patients, given that the labour is not optional to obtain good health outcomes, and that patients who perform said labour are silenced.

Testimonies made by patients are not the only testimonies being downgraded due to their reliance on affective experience. Some testimonies about everyday racism encounter the same type of pushback, due to the difficulty in providing intersubjectively verifiable evidence. In **Chapter 6**, I developed this idea by conceptualising *dialectical white scepticism*, a subtype of white scepticism motivated by argumentative concerns. I argued that dialectical white scepticism is structurally similar to ordinary scepticism, as racially oppressed people cannot offer inferential justification about why an individual occurrence of everyday racism is racist without relying on a form of justification that doesn't beg the question against the dialectical white sceptic. Indeed, the way to get out of this would be to recognise that affective experience provides a form of justification which would stop the regress of justification. However, the justification derived from affective experience will not satisfy the dialectical white sceptic given that affective experience is not taken seriously as a source of knowledge. I provided two ways to respond to the dialectical white sceptic and show that the burden of proof should fall on them rather than racially oppressed people. First, I argued that racially oppressed people could argue that their beliefs about racism are basic and justified by a form of epistemic conservatism. Second, I discussed the idea that racially oppressed people use reliable methods of knowledge acquisition, like their lived experience of racism which can be likened to a form of expert intuition or the use of statistical evidence which doesn't harm the target of the belief. The case of dialectical white scepticism highlighted once again that our epistemic standards are not applied uniformly. Indeed, if ordinary knowledge and knowledge about racism were treated the same way, the burden of proof would fall on the dialectical white sceptic the same way it falls on the ordinary sceptic. However, the current discrepancy within these standards reinforces racial oppression and contributes to the epistemic burden carried by racially oppressed

people. If affective experience were to be recognised as a proper source of knowledge, there is hope that these double standards would disappear, and the burden on racially oppressed people would be alleviated.

As I have shown throughout this dissertation, affective experience is a way of obtaining knowledge, especially about subtle phenomena that only some people will notice due to their membership to a specific group. Whether it is noticing sexism as a woman, racism as a person of colour, microaggressions as a member of a marginalised group, or changes in pain for someone with endometriosis, affective experience plays a crucial role in linking small instances with a wider pattern. These cases are paradigmatic of affective experience accurately tracking the world, because these agents will experience a variety of internal and social calibration on multiple occasions. Affective experience should be considered a source of knowledge in these contexts, as the agents have plenty of occasions to assess whether their affective experience tracks, and they can recalibrate it if needed.

It should not come as a surprise that affective experience plays an important role in acquiring knowledge about social phenomena like oppression. Indeed, this point was already raised by feminist philosophers who explained that our knowledge is shaped by our social standpoint (see for instance Alcoff 1999; Code 1991; Harding 2004; Longino 1990, 1993; Toole 2019, 2022). As Alcoff explains "social identity is relevant to epistemic judgement not because identity determines judgement but because identity can in some instances yield access to perceptual facts that themselves may be relevant to the formulation of various knowledge claims" (1999, 83). The idea that social identity influences what we can know about the world is often what has set apart standpoint epistemologists from contemporary analytic epistemology, given the tension between objectivity and the feminist idea that all knowledge is situated and influenced by social features (Toole 2022).

I have explained throughout this thesis that affective experience can help us bridge the perspective of the individual agent and the collective view of the group, due to the regulative mechanisms embedded in the use of affective experience as a source of

knowledge. More broadly, I think that such a theory can also be a step towards reconciliating standpoint epistemologists and contemporary analytic epistemology. Indeed, analytic epistemology focuses on the individual perspective whereas standpoint epistemology highlights social features. With the *Reflective Model* of affective experience, both these aspects are addressed. Indeed, we now have a way to satisfy the requirements of contemporary analytic epistemology whilst explaining why not all knowledge is available to all knowers.

Indeed, agents can first fail to use affective experience as a source of knowledge. This means that some relevant information will be inaccessible or obscure to them. Everyday racism, microaggressions, and other forms of oppression are a good example of this, as many people who don't belong to the oppressed group will lack the emotional response to the oppression, as they are not the relevant target. Without the emotion being elicited, they might fail to grasp the evidence necessary to form the relevant belief.

Second, agents might be unable to participate in certain types of social regulation necessary to adjust their affective experience so that it tracks the relevant features of the world. For instance, someone might feel anger towards some microaggressions but not others, as they don't have access to the group knowledge used to fine-tune their affective experience. In these cases, the social positioning is key to be able to engage in certain forms of regulation and obtain certain type of evidence. Affective experience is a source of knowledge which includes social aspects and embodied features in our knowledge production. By excluding it from the sources of knowledge, analytic epistemology has greatly diminished its ability to explain why certain beliefs and evidence are available to some agents and not others due to their social positioning. And this is exactly the fault that feminist standpoint epistemology has highlighted in this dominant model.

By introducing affective experience in the approved sources of knowledge of analytic epistemology, we are able to propose a modified version of the dominant model of epistemology. This new model takes into account the knowledge production of members from marginalised groups and includes it within our mainstream

epistemological discussions. Hopefully, this modification will be a step towards a more equitable treatment of the contributions made from the margin, as well as a solution to various forms of epistemic silencing currently experienced by members of marginalised groups.

Bibliography

Akbar, Arifa. 2021. 'Pain on the Page: Is This the End of the Hysterical, Ill Woman of Literature?' *The Guardian*, 19 June 2021. https://www.theguardian.com/books/2021/jun/19/pain-on-the-page-is-this-the-end-of-the-hysterical-ill-woman-of-literature.

Alcoff, Linda Martín. 1999. 'On Judging Epistemic Credibility: Is Social Identity Relevant?' *Philosophic Exchange* 29 (1): 73–93.

———. 2010. 'Epistemic Identities'. *Episteme* 7 (2): 128–37.

Anderson, Derek Egan. 2017. 'Conceptual Competence Injustice'. *Social Epistemology* 31 (2): 210–23.

Anderson, Ellie. Forthcoming. 'Hermeneutic Labor: The Gendered Burden of Interpretation in Intimate Relationships Between Women and Men'. *Hypatia*.

Archer, Alfred, and Georgina Mills. 2019. 'Anger, Affective Injustice, and Emotion Regulation'. *Philosophical Topics* 47 (2): 75–94.

Arruda, M. S., C. A. Petta, M. S. Abrão, and C. L. Benetti-Pinto. 2003. 'Time Elapsed from Onset of Symptoms to Diagnosis of Endometriosis in a Cohort Study of Brazilian Women'. *Human Reproduction* 18 (4): 756–59.

Audi, Robert. 2002. 'The Sources of Knowledge'. In *The Oxford Handbook of Epistemology*, edited by Paul K. Moser, 71–94. Oxford University Press.

———. 2013. *Moral Perception*. Princeton University Press.

Bailey, Alison. 2018. 'On Anger, Silence, and Epistemic Injustice'. *Royal Institute of Philosophy Supplements* 84 (November): 93–115.

Bain, David, and Michael Brady. 2014. 'Pain, Pleasure, and Unpleasure'. *Review of Philosophy and Psychology* 5 (1): 1–14.

Battaly, Heather. 2008. 'Virtue Epistemology'. *Philosophy Compass* 3 (4): 639–63.

Ben-Ze'ev, Aaron. 2000. *The Subtlety of Emotions*. Cambridge, MA, US: The MIT Press.

Berenstain, Nora. 2016. 'Epistemic Exploitation'. *Ergo, an Open Access Journal of Philosophy* 3 (22): 569–90.

Bergqvist, Anna, and Robert Cowan, eds. 2018. *Evaluative Perception*. Oxford University Press.

Bird, Alexander. 2019. 'Systematicity, Knowledge, and Bias. How Systematicity Made Clinical Medicine a Science'. *Synthese* 196 (3): 863–79

Blease, Charlotte, Havi Carel, and Keith Geraghty. 2017. 'Epistemic Injustice in Healthcare Encounters: Evidence from Chronic Fatigue Syndrome'. *Journal of Medical Ethics* 43 (8): 549–57.

Bolinger, Renée Jorgensen. 2020. 'The Rational Impermissibility of Accepting (Some) Racial Generalizations'. *Synthese* 197 (6): 2415–31.

Brady, Michael S. 2009. 'The Irrationality of Recalcitrant Emotions'. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 145 (3): 413–30.

———. 2013. *Emotional Insight: The Epistemic Role of Emotional Experience*. Oxford University Press.

Brogaard, Berit. 2018. 'Phenomenal Dogmatism, Seeming Evidentialism and Inferential Justification'. In *Believing in Accordance with the Evidence: New Essays on Evidentialism*, edited by Kevin McCain, 53–67. Synthese Library. Cham: Springer International Publishing.

Brown, Jessica. 2007. 'Externalism in Mind and Epistemology'. In *Internalism and Externalism in Semantics and Epistemology*, edited by Sanford C. Goldberg, 13–34. Oxford University Press.

Brownstein, Michael. 2019. 'Implicit Bias'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2019 Edition. <https://plato.stanford.edu/archives/fall2019/entries/implicit-bias/>.

Brownstein, Michael, and Jennifer Mather Saul, eds. 2016. *Implicit Bias and Philosophy*. First edition. Oxford, United Kingdom: Oxford University Press.

Brun, Georg, Ulvi Doğuoğlu, and Dominique Kuenzle, eds. 2008. *Epistemology and Emotions*. Ashgate Epistemology and Mind Series. Aldershot, Hants, England ; Burlington, VT: Ashgate.

Bu, Feifei, and Daisy Fancourt. 2021. 'How Is Patient Activation Related to Healthcare Service Utilisation? Evidence from Electronic Patient Records in England'. *BMC Health Services Research* 21 (1): 1196.

Buchman, Daniel Z., Anita Ho, and Daniel S. Goldberg. 2017. 'Investigating Trust, Expertise, and Epistemic Injustice in Chronic Pain'. *Journal of Bioethical Inquiry* 14 (1): 31–42.

Burki, Talha. 2020. 'The Online Anti-Vaccine Movement in the Age of COVID-19'. *The Lancet Digital Health* 2 (10): 504–5.

Byrne, Alex. 2020. 'Inverted Qualia'. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2020 Edition. https://plato.stanford.edu/entries/qualia-inverted/.

Campbell, Bradley, and Jason Manning. 2014. 'Microaggression and Moral Cultures'. *Comparative Sociology* 13 (6): 692–726.

———. 2015. ‘The New Millennial “Morality”: Highly Sensitive and Easily Offended.’ *Time*, 17 November 2015. https://time.com/4115439/student-protests-microaggressions/.

———. 2016. ‘Campus Culture Wars and the Sociology of Morality’. *Comparative Sociology* 15 (2): 147–78.

Campbell, Dennis. 2022. ‘Staffing Crisis Deepens in NHS England with 110,000 Posts Unfilled’. *The Guardian*, 3 March 2022. https://www.theguardian.com/society/2022/mar/03/staffing-crisis-deepens-in-nhs-england-with-110000-posts-unfilled.

Candiotto, Laura, ed. 2019. *The Value of Emotions for Knowledge*. Cham: Springer International Publishing.

Carel, Havi, and Ian James Kidd. 2014. ‘Epistemic Injustice in Healthcare: A Philosophial Analysis’. *Medicine, Health Care and Philosophy* 17 (4): 529–40.

Carroll, Lewis. 1895. ‘What the Tortoise Said to Achilles’. *Mind* 4 (14): 278–80.

Chen, Jie, C. Daniel Mullins, Priscilla Novak, and Stephen B. Thomas. 2016. ‘Personalized Strategies to Activate and Empower Patients in Health Care and Reduce Health Disparities’. *Health Education & Behavior: The Official Publication of the Society for Public Health Education* 43 (1): 25–34.

Cherry, Myisha. 2021. *The Case for Rage: Why Anger Is Essential to Anti-Racist Struggle*. Oxford, New York: Oxford University Press.

Code, Lorraine. 1991. *What Can She Know?: Feminist Theory and the Construction of Knowledge*. Cornell University Press.

Collins, Patricia Hill. 1986. ‘Learning from the Outsider Within: The Sociological Significance of Black Feminist Thought’. *Social Problems* 33 (6): S14–32.

———. 2000. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. Rev. 10th anniversary ed. New York: Routledge.

———. 2019. *Intersectionality as Critical Social Theory*. Durham: Duke University Press.

Colombetti, Giovanna. 2014. *The Feeling Body: Affective Science Meets the Enactive Mind*. MIT Press.

Congdon, Matthew. 2015. ‘Epistemic Injustice in the Space of Reasons’. *Episteme* 12 (1): 75–93.

Correll, Joshua, Bernadette Park, Charles M. Judd, and Bernd Wittenbrink. 2002. ‘The Police Officer’s Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals.’ *Journal of Personality and Social Psychology* 83 (6): 1314–29.

Cowan, Robert. 2018. 'Epistemic Sentimentalism and Epistemic Reason-Responsiveness'. In *Evaluative Perception*, edited by Anna Bergqvist and Robert Cowan, 219–36. Oxford University Press.

Craig, Edward. 1999. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.

Cudd, Ann E. 2006. *Analyzing Oppression*. Oxford University Press.

Cut, dir. 2016. *Can I Touch Your Hair ? / Black Women/ One Word/ Cut*. Cut. https://www.youtube.com/watch?v=qaU67Ps7TWw.

D'Arms, Justin, and Daniel Jacobson. 2000. 'The Moralistic Fallacy: On the "Appropriateness" of Emotions'. *Philosophy and Phenomenological Research* 61 (1): 65.

Dasgupta, Nilanjana, and Anthony G. Greenwald. 2001. 'On the Malleability of Automatic Attitudes: Combating Automatic Prejudice with Images of Admired and Disliked Individuals.' *Journal of Personality and Social Psychology* 81 (5): 800–814.

Dembroff, Robin, and Dennis Whitcomb. 2022. 'Content Focused Epistemic Injustice'. In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler, John Hawthorne, and Julianne Chung, 7:48–70. Oxford University Press.

Deonna, Julien, and Fabrice Teroni. 2012. *The Emotions: A Philosophical Introduction*. Routledge.

Devlin, Hannah. 2022. 'Menopause, Endometriosis and More: Four Ways England Is Failing Women'. *The Guardian*, 2 June 2022. https://www.theguardian.com/society/2022/jun/02/menopause-endometriosis-and-more-four-ways-england-is-failing-women.

Döring, Sabine A. 2007. 'Seeing What to Do: Affective Perception and Rational Motivation'. *Dialectica* 61 (3): 363–94.

———. 2009. 'The Logic of Emotional Experience: Noninferentiality and the Problem of Conflict Without Contradiction'. *Emotion Review* 1: 240–47.

Dotson, Kristie. 2011. 'Tracking Epistemic Violence, Tracking Practices of Silencing'. *Hypatia* 26 (2): 236–57.

———. 2012. 'A Cautionary Tale: On Limiting Epistemic Oppression'. *Frontiers: A Journal of Women Studies* 33 (1): 24–47.

———. 2014. 'Conceptualizing Epistemic Oppression'. *Social Epistemology* 28 (2): 115–38.

Dover, Michael A. 2016. 'The Moment of Microaggression: The Experience of Acts of Oppression, Dehumanization, and Exploitation'. *Journal of Human Behavior in the Social Environment* 26 (7–8): 575–86.

Endometriosis UK. 2022. 'Information | Endometriosis UK'. 2022. https://www.endometriosis-uk.org/information.

Enlow, Courtney. 2015. 'The Way the Media Reports on Kristen Stewart and Her Girlfriend Is Hilarious'. Pajiba. 21 April 2015. https://www.pajiba.com/celebrities_are_better_than_you/the-way-the-media-reports-on-kristen-stewart-and-her-girlfriend-is-hilarious.php.

Fatima, Saba. 2017. 'On the Edge of Knowing: Microaggression and Epistemic Uncertainty as a Woman of Color'. In *Surviving Sexism in Academia: Feminist Strategies for Leadership*, edited by Kirsti Cole and Holly Hassel, 147–57. Routledge.

———. 2020. 'I Know What Happened to Me'. In *Microaggressions and Philosophy*, edited by Lauren Freeman and Jeanine Weekes Schroer, 163–83. Routledge Studies in Contemporary Philosophy 137. New York, NY: Routledge.

Feldman, Fred. 2004. *Pleasure and the Good Life: Concerning the Nature Varieties and Plausibility of Hedonism*. Oxford University Press.

Frankish, Keith. 2016. 'Playing Double: Implicit Bias, Dual Levels, and Self-Control'. In *Implicit Bias and Philosophy, Volume 1: Metaphysics and Epistemology*, edited by Michael Brownstein and Jennifer Saul, 23-46. Oxford University Press.

Freeman, Lauren, and Heather Stewart. 2018. 'Microaggressions in Clinical Medicine'. *Kennedy Institute of Ethics Journal* 28 (4): 411–49.

———. 2020. 'Sticks and Stones Can Break Your Bones and Words Can Really Hurt You : A Standpoint Epistemological Reply to Critics of the Microaggression Research Program'. In *Microaggressions and Philosophy*, edited by Lauren Freeman and Jeanine Weekes Schroer, 36–66. New York: Routledge.

Freeman, Lauren, and Jeanine Weekes Schroer, eds. 2020. *Microaggressions and Philosophy*. Routledge Studies in Contemporary Philosophy. New York: Routledge.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Fridja, Nico. 1986. *The Emotions*. Cambridge University Press.

Friedlaender, Christina. 2018. 'On Microaggressions: Cumulative Harm and Individual Responsibility'. *Hypatia* 33 (1): 5–21.

———. 2021. 'Microaggressions and the Problem of Attributional Ambiguity'. In *The Routledge Handbook of Social and Political Philosophy of Language*, edited by Justin Khoo and Rachel Katharine Sterken, 1st Edition, 232–46. New York: Routledge.

Frye, Marilyn. 1983. *The Politics of Reality: Essays in Feminist Theory*. The Crossing Press Feminist Series. Trumansburg, N.Y: Crossing Press.

Furtak, Rick Anthony. 2018. *Knowing Emotions: Truthfulness and Recognition in Affective Experience*. New York, NY : Oxford University Press.

Gardiner, Georgi. 2019. 'Legal Burdens of Proof and Statistical Evidence'. In *The Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase, 179–95. Routledge Handbooks in Philosophy. New York: Routledge.

Gibert, Sophia H., David DeGrazia, and Marion Danis. 2017. 'Ethics of Patient Activation: Exploring Its Relation to Personal Responsibility, Autonomy and Health Disparities'. *Journal of Medical Ethics* 43 (10): 670–75.

Gigerenzer, Gerd. 2008. *Gut Feelings: Short Cuts to Better Decision Making*. Penguin Books.

Goldberg, Sandy. 2011. 'The Division of Epistemic Labor'. *Episteme* 8 (1): 112–25.

Goldman, Alvin I. 2016. 'Reliabilism'. In *Routledge Encyclopedia of Philosophy*, 1st ed. London: Routledge.

Greene, Jessica, Judith H. Hibbard, Rebecca Sacks, Valerie Overton, and Carmen D. Parrotta. 2015. 'When Patient Activation Levels Change, Health Outcomes and Costs Change, Too'. *Health Affairs (Project Hope)* 34 (3): 431–37.

Gross, James J., ed. 2014. *Handbook of Emotion Regulation*. 2. ed. New York, NY: Guilford Press.

Haagaard, Alex. 2020. 'What It's Like Living and Working With a Chronic Illness'. *The Muse* (blog). Accessed 12 June 2022. https://www.themuse.com/advice/working-living-with-a-chronic-illness.

Harding, Sandra G., ed. 2004. *The Feminist Standpoint Theory Reader: Intellectual and Political Controversies*. New York: Routledge.

Harrison, Eilidh. 2021. 'The Prospects of Emotional Dogmatism'. *Philosophical Studies* 178 (8): 2535–55.

Haslanger, Sally. 1999. 'What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology'. *Philosophical Perspectives* 13: 459–80.

hooks, bell. 1996. *Killing Rage: Ending Racism*. An Owl Book. New York: Holt.

Hudelist, G., N. Fritzer, A. Thomas, C. Niehues, P. Oppelt, D. Haas, A. Tammaa, and H. Salzer. 2012. 'Diagnostic Delay for Endometriosis in Austria and Germany: Causes and Possible Consequences'. *Human Reproduction* 27 (12): 3412–16.

Huemer, Michael. 2005. *Ethical Intuitionism*. Palgrave Macmillan.

Jaggar, Alison M. 1989. 'Love and Knowledge: Emotion in Feminist Epistemology'. *Inquiry* 32 (2): 151–76.

James, William. 1890. *The Principles of Psychology*. New York: Holt.

Jarry, Jonathan. 2021. 'Don't Fall for the "VAERS Scare" Tactic: Anti-Vaxxers Show How a Precious Vaccine Adverse Event Reporting Database Can Be Used to Scare the Public'. *McGill Office for Science and Society* (blog). 18 June 2021. https://www.mcgill.ca/oss/article/covid-19-critical-thinking-health/dont-fall-vaers-scare-tactic.

Johnston, Mark. 2001. 'The Authority of Affect'. *Philosophy and Phenomenological Research* 63 (1): 181–214.

Jones, James M. 1997. *Prejudice and Racism*. 2nd ed. McGraw-Hill Series in Social Psychology. New York: McGraw-Hill Companies.

Jones, Karen. 2003. 'Emotion, Weakness of Will, and the Normative Conception of Agency'. *Royal Institute of Philosophy Supplement* 52 (March): 181–200.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. Toronto: Doubleday Canada.

Kahneman, Daniel, and Gary Klein. 2009. 'Conditions for Intuitive Expertise: A Failure to Disagree.' *American Psychologist* 64 (6): 515–26.

Keltner, Dacher, and Robert J. Robinson. 1996. 'Extremism, Power, and the Imagined Basis of Social Conflict.' *Current Directions in Psychological Science* 5: 101–5.

Kennett, Jeanette, and Cordelia Fine. 2009. 'Will the Real Moral Judgment Please Stand up? The Implications of Social Intuitionist Models of Cognition for Meta-Ethics and Moral Psychology'. *Ethical Theory and Moral Practice* 12 (1): 77–96.

Kidd, Ian James, José Medina, and Gaile Pohlhaus. 2017. *The Routledge Handbook of Epistemic Injustice*. Routledge.

Kiesel, Laura. 2017. 'Chronic Pain: The "Invisible" Disability'. Harvard Health Blog. 28 April 2017. https://www.health.harvard.edu/blog/chronic-pain-the-invisible-disability-2017042811360.

Kiley, Ashley. 2019. 'Did Cara Delevingne and Ashley Benson Get Married?' 4 August 2019. https://www.pride.com/celebrities/2019/8/04/did-cara-delevingne-and-ashley-benson-get-married.

Kornblith, Hilary. 2012. *On Reflection*. 1st ed. Oxford: Oxford University Press.

Kukla, Rebecca. 2005. 'Conscientious Autonomy: Displacing Decisions in Health Care'. *The Hastings Center Report* 35 (2): 34–44.

———. 2007. 'How Do Patients Know?' *Hastings Center Report* 37 (5): 27–35.

Lagewaard, Thirza J. 2021. ‘Epistemic Injustice and Deepened Disagreement’. *Philosophical Studies* 178 (5): 1571–92.

Lange, Carl Georg. 1922. *Om Sindsbevægelser: Et Psyko-Fysiologisk Studie. Translated as The Emotions (along with William James “What Is an Emotion?”)*. Translated by A. Haupt. Baltimore: Williams and Wilkins.

Leder, Drew. 1990. *The Absent Body*. Chicago, IL: University of Chicago Press.

Lewis, David. 1996. ‘Elusive Knowledge’. *Australasian Journal of Philosophy* 74 (4): 549–67.

Li, Ping, Zi yan Cheng, and Gui lin Liu. 2020. ‘Availability Bias Causes Misdiagnoses by Physicians: Direct Evidence from a Randomized Controlled Trial’. *Internal Medicine* 59 (24): 3141–46.

Lilienfeld, Scott O. 2017. ‘Microaggressions: Strong Claims, Inadequate Evidence’. *Perspectives on Psychological Science* 12 (1): 138–69.

Lindemann, Hilde. 2009. ‘Holding on to Edmund: The Relational Work of Identity’. In *Naturalized Bioethics: Toward Responsible Knowing and Practice*, edited by Hilde Lindemann, Marian Verkerk, and Margaret Urban Walker, 65–79. Cambridge University Press.

Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.

———. 1993. ‘Feminist Standpoint Theory and the Problems of Knowledge’. by Dorothy Smith, Liz Stanley, Susan Hekman, and Sandra Harding. *Signs* 19 (1): 201–12.

Lorde, Audre. 1984. *Sister Outsider*. New York: Penguin Books.

Lorde, Audre, and Claudia Tate. 1983. ‘Audre Lorde’. In *Black Women Writers at Work*, edited by Claudia Tate, 100–116. New York: Continuum.

Lyons, Jack C. 2018. ‘Perception and Intuition of Evaluative Properties’. In *Evaluative Perception*, edited by Anna Bergqvist and Robert Cowan, 183–99. Oxford University Press.

Malmqvist, Erik, and András Szigeti. 2019. ‘Exploitation and Joint Action’. *Journal of Social Philosophy* 50 (3): 280–300.

Mantel, Hilary. 2004. ‘“Every Part of My Body Hurt”’. *The Guardian*, 7 June 2004. https://www.theguardian.com/society/2004/jun/07/health.genderissues.

Mason, Elinor. 2007. ‘The Nature of Pleasure: A Critique of Feldman’. *Utilitas* 19 (3): 379–87.

McBean, Sam. 2016. ‘The “Gal Pal Epidemic”’. *Celebrity Studies* 7 (2): 282–84.

———. 2018. '"We Fuck and Friends Don't Fuck": BFFs, Lesbian Desire, and Queer Narratives'. *Textual Practice* 32 (6): 957–72.

McClure, Emma. 2019. 'Theorizing a Spectrum of Aggression: Microaggressions, Creepiness, and Sexual Assault'. *The Pluralist* 14 (1): 91–101.

McClure, Emma, and Regina Rini. 2020. 'Microaggression: Conceptual and Scientific Issues'. *Philosophy Compass* 15 (4): e12659.

McGrath, Sarah. 2004. 'Moral Knowledge by Perception'. *Philosophical Perspectives* 18 (1): 209–28.

McKinnon, Rachel. 2016. 'Epistemic Injustice'. *Philosophy Compass* 11 (8): 437–46.

McLeod, Carolyn, and Susan Sherwin. 2000. 'Relational Autonomy, Self-Trust, and Health Care for Patients Who Are Oppressed'. In *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*, edited by Catriona Mackenzie and Natalie Stoljar, 259–79. Oxford University Press.

McTernan, Emily. 2018. 'Microaggressions, Equality, and Social Practices'. *Journal of Political Philosophy* 26 (3): 261–81.

Medina, José. 2012. 'Hermeneutical Injustice and Polyphonic Contextualism: Social Silences and Shared Hermeneutical Responsibilities'. *Social Epistemology* 26 (2): 201–20.

Mills, Charles. 2007. 'White Ignorance'. In *Race and Epistemologies of Ignorance*, edited by Shannon Sullivan Nancy Tuana, 11–38. State Univ of New York Pr.

Milona, Michael. 2015. 'Review: Michael Brady, Emotional Insight: The Epistemic Role of Emotional Experience'. *Ethics* 125 (2): 567–71.

Minghella, Anthony. 1999. *The Talented Mr. Ripley: A Screenplay*. Miramax.

Moore, George Edward. 1939. 'Proof of an External World'. *Proceedings of the British Academy* 25: 273–300.

Moretti, Luca. 2015. 'Phenomenal Conservatism'. *Analysis* 75 (2): 296–309.

Morton, Adam. 2009. 'Epistemic Emotions'. In *The Oxford Handbook of Philosophy of Emotion*, edited by Peter Goldie, 385-400. Oxford University Press.

Neblett, William. 1979. 'Indignation: A Case Study in the Role of Feelings in Morals'. *Metaphilosophy* 10 (2): 139–52.

NHS. 2022a. 'Endometriosis'. 5 September 2022. https://www.nhs.uk/conditions/endometriosis/.

———. 2022b. ‘Polycystic Ovary Syndrome’. 11 October 2022. https://www.nhs.uk/conditions/polycystic-ovary-syndrome-pcos/.

Nnoaham, Kelechi E., Lone Hummelshoj, Premila Webster, Thomas d’Hooghe, Fiorenzo de Cicco Nardone, Carlo de Cicco Nardone, Crispin Jenkinson, Stephen H. Kennedy, and Krina T. Zondervan. 2011. ‘Impact of Endometriosis on Quality of Life and Work Productivity: A Multicenter Study across Ten Countries’. *Fertility and Sterility* 96 (2): 366-373.

Nussbaum, Martha C. 2004. ‘Emotions as Judgments of Value and Importance.’ In *Thinking about Feeling: Contemporary Philosophers on Emotions.*, edited by Robert Solomon, 183–99. Series in Affective Science. New York, NY, US: Oxford University Press.

———. 2016. *Anger and Forgiveness: Resentment, Generosity, and Justice*. Oxford, New York: Oxford University Press.

Page, Allison, and Jacquelyn Arcy. 2020. ‘#MeToo and the Politics of Collective Healing: Emotional Connection as Contestation’. *Communication, Culture and Critique* 13 (3): 333–48.

Palmer, Raymond. 2022. ‘Covid19 Vaccines and the Misinterpretation of Perceived Side Effects’. *BioMedicine* 12 (3): 1–4.

Pelser, Adam C. 2014. ‘Emotion, Evaluative Perception, and Epistemic Justification’. In *Emotion and Value*, edited by Sabine Roeser and Cain Todd, 107–23. Oxford University Press.

Pierce, Chester. 1970. ‘Offensive Mechanisms’. In *The Black Seventies*, edited by Floyd Barbour, 265–82. Boston, MA: Porter Sargent.

Pohlhaus, Gaile. 2002. ‘Knowing Communities: An Investigation of Harding’s Standpoint Epistemology’. *Social Epistemology* 16 (3): 283–93.

———. 2012. ‘Relational Knowing and Epistemic Injustice: Toward a Theory of Willful Hermeneutical Ignorance’. *Hypatia* 27 (4): 715–35.

Pollock, John L., and Joseph Cruz. 1999. *Contemporary Theories of Knowledge*. 2nd ed. Studies in Epistemology and Cognitive Theory. Lanham, Md: Rowman & Littlefield Publishers.

Prinz, Jesse. 2006. ‘The Emotional Basis of Moral Judgments’. *Philosophical Explorations* 9 (1): 29–43.

Prinz, Jesse J. 2006. *Gut Reactions: A Perceptual Theory of Emotion*. Philosophy of Mind. Oxford, New York: Oxford University Press.

Proust, Joëlle. 2013. *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. First edition. Oxford ; New York: Oxford University Press.

Rescorla, Michael. 2009a. ‘Assertion and Its Constitutive Norms’. *Philosophy and Phenomenological Research* 79 (1): 98–130.

———. 2009b. ‘Shifting the Burden of Proof?’ *The Philosophical Quarterly* 59 (234): 86–109.

———. 2009c. ‘Epistemic and Dialectical Regress’. *Australasian Journal of Philosophy* 87 (1): 43–60.

Rini, Regina. 2018. ‘How to Take Offense: Responding to Microaggression’. *Journal of the American Philosophical Association* 4 (3): 332–51.

———. 2021. *The Ethics of Microaggression*. Abingdon, Oxon; New York, NY: Routledge.

Roberts, Robert C. 2003. *Emotion: An Essay in Aid of Moral Psychology*. Cambridge University Press.

Sauer, Hanno. 2012. ‘Educated Intuitions. Automaticity and Rationality in Moral Judgement’. *Philosophical Explorations* 15 (3): 255–75.

Scarantino, Andrea. 2014. ‘The Motivational Theory of Emotions’. In *Moral Psychology and Human Agency*, edited by Justin D’Arms and Daniel Jacobson, 156–85. Oxford University Press.

———. 2016. ‘The Philosophy of Emotions and Its Impact on Affective Science’. In *Handbook of Emotions*, edited by Lisa Feldman Barrett, Michael Lewis, and Jeannette M. Haviland-Jones, 4th ed. New York (N.Y.): Guilford press.

Scarantino, Andrea, and Ronald de Sousa. 2018. ‘Emotion’. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2018. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2018/entries/emotion/.

Setiya, Kieran. 2012. *Knowing Right From Wrong*. Oxford, New York: Oxford University Press.

Siegel, Reva B. 2003. ‘Introduction: A Short History of Sexual Harassment’. In *Directions in Sexual Harassment Law*, edited by Catharine A. MacKinnon and Reva B. Siegel, 1–28. Yale University Press.

Simon, Herbert A. 1992. ‘What Is an “Explanation” of Behavior?’ *Psychological Science* 3 (3): 150–61.

Sinnott-Armstrong, Walter. 2008. ‘Framing Moral Intuitions’. In *Moral Psychology, Vol 2: The Cognitive Science of Morality: Intuition and Diversity*, 47–76. Cambridge, MA, US: MIT Press.

Skuse, Alanna. 2021. *Surgery and Selfhood in Early Modern England: Altered Bodies and Contexts of Identity*. Cambridge, United Kingdom ; New York: Cambridge University Press.

Slote, Michael. 2014. *A Sentimentalist Theory of the Mind*. Oxford University Press.

Solomon, Robert. 1993. *The Passions: Emotions and the Meaning of Life*. Indianapolis: Hackett Pub. Co.

———, ed. 2004. *Thinking about Feeling: Contemporary Philosophers on Emotions*. Series in Affective Science. Oxford, New York: Oxford University Press.

Sousa, Ronald B. de. 1987. *The Rationality of Emotion*. MIT Press.

Sue, Derald Wing. 2010. *Microaggressions in Everyday Life: Race, Gender, and Sexual Orientation*. John Wiley & Sons.

Sue, Derald Wing, Christina M. Capodilupo, Gina C. Torino, Jennifer M. Bucceri, Aisha M. B. Holder, Kevin L. Nadal, and Marta Esquilin. 2007. 'Racial Microaggressions in Everyday Life: Implications for Clinical Practice'. *The American Psychologist* 62 (4): 271–86.

Tappolet, Christine. 2000. *Emotions et Valeurs*. Paris: Presses Universitaires de France.

Thagard, Paul. 2006. *Hot Thought: Mechanisms and Applications of Emotional Cognition*. Cambridge, Mass: MIT Press.

The Ehlers-Danlos Society. 2020. '4 Years after My Diagnosis, I Am Still Struggling to Be Taken Seriously'. 30 April 2020. https://www.ehlers-danlos.com/story/kirsty-l/.

Toole, Briana. 2019. 'From Standpoint Epistemology to Epistemic Oppression'. *Hypatia* 34 (4): 598–618.

———. 2022. 'Objectivity in Feminist Epistemology'. *Philosophy Compass* 17 (11).

Turner, Hannah. 2022. 'They Said It Was Bipolar II. Now I Know It Was Endometriosis'. *Refinery29*, 14 March 2022. https://www.refinery29.com/en-gb/endometriosis-diagnosis-mental-health?utm_source=email&utm_medium=editorial&utm_content=uk&utm_campaign=220314-how-to-get-salon-quality-hair-at-home&fbclid=IwAR3BTjxD1wxgSq2MtzxamguCNl3h9Ff9mzsq0lgEK0UR0SWhTlODE4O8XiM.

Dyke, Christina Van. 2018. 'What Has History to Do with Philosophy? Insights from the Medieval Contemplative Tradition'. In *Philosophy and the Historical Perspective*, edited by Marcel van Ackeren, Lee Klein, and Marcel van Ackeren, 155–70. British Academy.

Walker, Amy. 2022. '"The Pain Is Inhumane": How NHS Gynaecology Delays Affect Women's Health'. *The Guardian*, 2 June 2022. https://www.theguardian.com/society/2022/jun/02/the-pain-is-inhumane-how-nhs-gynaecology-delays-affect-womens-health?fbclid=IwAR3zo3b0F9lMIX7SKe8LuvWAlyKmUPvPF_ERrApw0fVS1NxrwYvxFrzHGoU.

Wendell, Susan. 1996. *The Rejected Body: Feminist Philosophical Reflections on Disability*. New York: Routledge.

Wilson, Emilia L. 2021. 'The Dual Erasure of Domestic Epistemic Labour'. *Proceedings of the Aristotelian Society* 121 (1): 111–25.

Wood, Ros, Heather Guidone, and Lone Hummelshoj. 2016. 'Myths and Misconceptions in Endometriosis'. *Endometriosis.Org* (blog). 20 November 2016. https://endometriosis.org/resources/articles/myths/.

Woolf, Virginia. 1929. *A Room of One's Own*. London: Hogarth Press.

Young, Iris M. 1988. 'Five Faces of Oppression'. *Philosophical Forum* 19 (4): 270.

Young, Kate, Jane Fisher, and Maggie Kirkman. 2017. 'Clinicians' Perceptions of Women's Experiences of Endometriosis and of Psychosocial Care for Endometriosis'. *Australian and New Zealand Journal of Obstetrics and Gynaecology* 57 (1): 87–92.

Zagzebski, Linda Trinkaus. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.