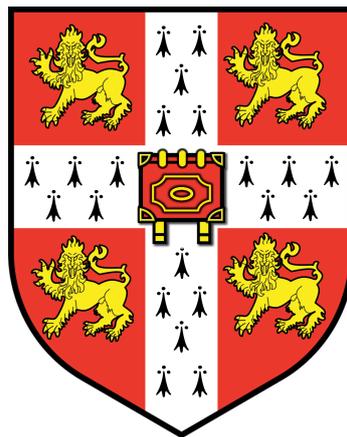


Saliency Perspectives

Ella Kate Whiteley
Gonville and Caius College,
University of Cambridge



This dissertation is submitted for the degree of
Doctor of Philosophy

February 2019

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Saliency Perspectives

Ella Whiteley

Abstract

In the philosophy of language and epistemology, debates often centre on what *content* a person is communicating, or representing in their mind. How that content is organised, along dimensions of saliency, has received relatively little attention. I argue that saliency matters. Mere change of saliency patterns, without change of content, can have dramatic implications, both epistemic and moral.

Imagine two newspaper articles that offer the same information about a subject, but differ in terms of what they headline. These articles can be said to adopt different *linguistic saliency perspectives*. Making different things salient in language is a way of making different things salient in an audience's mind: it is a way of encouraging the audience to adopt a particular *cognitive saliency perspective*. Building on Elisabeth Camp's work on *perspectives*, and Sebastian Watzl's work on attention, I suggest that one has a certain cognitive saliency perspective in virtue of better noticing, remembering, and finding cognitively accessible certain contents over others.

Drawing on psychological research into cognitive biases and framing effects, I argue that that simply making some content salient in the mind, perhaps through first making it salient in language, can be sufficient to activate substantive cultural beliefs or ideologies associated with that content. Where those beliefs and ideologies have epistemic and moral problems, we have grounds for criticising the saliency perspective that causally produced them. Besides this instrumental harm, I also suggest that certain saliency perspectives can themselves *constitute* harm. I draw on feminist work on *objectification* to argue that making the wrong thing salient about a person can constitute a way of dehumanising them. A great many factors, from physical and psychological violence, to false beliefs and credibility deficits, have already been identified as potentially harming an individual or group. What is distinctive about this argument, then, is the idea that, sometimes, mere patterns of saliency can be damaging in and of themselves.

Acknowledgments

I would like to thank the Arts and Humanities Research Council and the Faculty of Philosophy for financial support, which enabled me to complete this thesis.

I would also like sincerely to thank Prof. Rae Langton for the continuous support of my PhD research. I have found our conversations fascinating and enriching, and this thesis has benefited immeasurably from her immense knowledge and understanding, and from her encouragement. Besides my supervisor, I would also like to thank Prof. Tim Lewens, who helped a great deal by supervising me for a term. Special thanks also to Dr. Sophia Connell for her particularly illuminating and enjoyable discussions, and for her general support. I have also benefited from interesting philosophical conversations with many individuals, which have shaped the research in this thesis, including Sebastian Watzl, Rachel Fraser, Paul-Mikhail Podosky, Rachael Hadoux, Mat Simpson, and Leonie Smith. Mat in particular has also been a greatly supportive friend during my time at Cambridge. I would like to thank the staff at the Faculty of Philosophy, especially Charlie Evans, not only for being so helpful when I have made administrative queries, but for creating a warm and caring environment at the faculty.

Last but not least, I would like to thank those outside of philosophy. In particular, my Mum, Diane Hill, has helped not only with proofreading this thesis, but has generally been unconditionally supportive and encouraging. I am forever grateful for her wisdom, humour, and love. I would also like to thank my partner, Justin Tadros, not only for discussing the issues in this thesis with me, but for making me delicious meals when I was busy writing and editing, caring for me, and generally being a loving support. Thanks also to my brother, Dan Hill, for his interest in my work, and for his encouragement. Finally, I would like to thank my late Dad, Nigel Whiteley, who gave me, amongst many things, a deep curiosity about the world, which no doubt shaped my decision to embark on this research.

Table of Contents

Preface	iii
Abstract.....	v
Acknowledgements.....	vi
1. Introduction.....	1
1.1. Pinpointing elusive disagreements	2
1.2. Where is the philosophy of salience?.....	4
1.3. The significance of salience.....	7
1.4. Chapter summary.....	9
1.5. Conclusion	12
2. Salience and Language	15
2.1. Introduction.....	15
2.2. Introducing linguistic salience perspectives.....	16
2.2.1. Framing effects	16
2.2.2. Linguistic salience perspectives	19
2.3. Linguistic salience perspectives in context	21
2.3.1. Example 1: Advertising beef.....	21
2.3.2. Example 2: Oral health messages	23
2.3.3. Example 3: Genetics pedagogies.....	25
2.4. Ways of evaluating linguistic salience perspectives	28
2.4.1. Instrumental critique.....	28
2.4.2. Non-instrumental critique.....	31
2.5. How to choose linguistic salience perspectives	34
2.6. Conclusion.....	39
3. Salience and Cognition	41
3.1. Introduction.....	41
3.2. Introducing cognitive salience perspectives	42
3.2.1. Some examples	42
3.2.2. Cognitive salience perspectives	43
3.2.2.1. Dispositional nature.....	44

3.2.2.2. Scope and duration	46
3.2.2.3. Automatic and intuitive nature	47
3.3. Ways of evaluating cognitive salience perspectives	50
3.3.1. Instrumental critique.....	50
3.3.2. Non-instrumental critique.....	52
3.4. Changing cognitive salience perspectives	57
3.4.1. Individualist approach	58
3.4.2. Structural approach	60
3.5. Conclusion.....	61

4. Salience Perspectives as Instrumental Harms: The Substance Salience

Perspective	63
4.1. Introduction.....	63
4.2. Introducing the substance salience perspective	65
4.3. Resonances with psychological biases	67
4.3.1. The correspondence bias	68
4.3.2. Psychological essentialism.....	72
4.3.3. The fixed mindset bias.....	75
4.4. The substance salience perspective: instrumental harms	78
4.5. The promise of the process salience perspective.....	82
4.6. Conclusion	87

5. Salience Perspectives as Instrumental Harms: Sex/Gender..... 89

5.1. Introduction.....	89
5.2. Case study: A framing effect for sex/gender	90
5.2.1. Bruckmüller et al.'s <i>Framing Gender Differences</i>	91
5.2.2. Analysis	94
5.3. Introducing sex/gender essentialism.....	97
5.3.1. Sex/gender essentialism: the beliefs	97
5.3.1.1. Two opposing categories	98
5.3.1.2. Discrete categories, internally homogenous.....	99
5.3.1.3. Biological determinism	100
5.3.1.4. Immutability.....	102
5.3.1.5. High inductive potential.....	102

5.3.2. Sex/gender essentialism: harms	103
5.4. Choosing salience perspectives for sex/gender	107
5.4.1. Salience perspectives to avoid: instrumental harms	107
5.4.2. The novelty of this proposal.....	111
5.4.3. Salience perspectives to choose: instrumental benefits	111
5.5. Conclusion	116
6. Salience Perspectives as Constitutive Harms: Sex/Gender	119
6.1. Introduction.....	119
6.2. Case studies.....	120
6.2.1. Rape victim-survivors	120
6.2.2. Women and their bodies	127
6.3. Perspectival objectification.....	129
6.4. Perspectival sex/gender essentialism.....	133
6.5. Addressing potential objections	139
6.5.1. Perspectival harm is too broad.....	139
6.5.2. Clashing criteria for harm.....	140
6.6. Conclusion.....	144
7. Conclusion	147
7.1. Introduction.....	147
7.2. Key points.....	147
7.3. Wider implications of salience perspectives	150
7.3.1. Race and hate speech.....	150
7.3.2. Social policy.....	152
7.4. Limitations and further research	154
7.4.1. The salience patterns in this thesis	154
7.4.2. Empirical studies	155
7.4.3. The limits of salience regarding social justice	156
Bibliography	161

CHAPTER 1

INTRODUCTION

This thesis is about *salience perspectives*, which refer to patterns of salience, either in language or the mind. One might make some content salient in language by mentioning it *first*. Making this content salient in language is a way of making it salient in an audience's *mind*. Content is salient in a person's mind when they better notice and remember that content—when they better *attend* to it. The aim of this thesis is to demonstrate that simply changing the pattern of salience in some linguistic or mental contents, without making any other changes to those contents, can both *cause* and *constitute* harm. This is interesting, I suggest, because, to many intuitions, a change in the salience pattern applied to a set of contents seems inconsequential; many of us would dismiss such a change as making no difference to the truth-value of those contents, let alone the potential for those contents to harm. In the main part, my focus is on demonstrating the causal part of my claim—namely, that a mere change in salience perspective can *cause* harm, in the form of activating harmful (and often *epistemically* problematic) biases that are prevalent in our folk-cultural context. I then turn to contemplating how a pattern of salience can itself *constitute* harm. Here, I suggest that making the wrong thing salient about a person might *count as a way* of harming them, such as by disrespecting their personhood.

In this chapter, I begin by explaining how I came across the philosophy of salience in §1.1. Given that the topic of salience can sound obscure, it might be helpful to demonstrate how an everyday experience, which may well be familiar to the reader, can be illuminated by the concept of salience. In §1.2, I discuss how philosophers have tended to talk about salience, commenting on the relative lack of attention to the subject. I contrast this with the relative centrality the subject is given in certain other disciplines. I then spell out why the philosophy of salience matters in §1.3, commenting on how the very thing that might make it seem insignificant is what makes it powerful. Finally, I sketch a summary of the proceeding chapters in §1.4.

1.1 Pinpointing elusive disagreements

Some of my background is in the philosophy of biology, which led to me having discussions with other philosophers about trait development. This involved discussing how genes and environment interact to produce an organism's traits, such as eye colour. Sometimes, I would feel friction with my interlocutor, but it was hard to pinpoint precisely what that friction consisted in. This is because we seemed to agree on the facts. We agreed, for instance, that certain genes reliably correlate with certain traits across changes in the organism's environment; a particular genetic variant might consistently aid in the development of blue eyes, for example, across a number of different environments. We might also agree that, for certain other genes, a shift in the organism's environment can make that gene function in a surprisingly different way. For instance, one and the same gene might help a plant grow *taller* than other plants in high altitude, and yet help it to grow to a *smaller* height than other plants in low altitude. Ostensibly, then, we agreed on the facts. When I would hear them talk about trait development, however, I felt that something was amiss. In particular, I had a sense that they were not taking certain facts sufficiently *seriously* in some way.

The facts in question concerned how the environment can dramatically alter the function of a gene. If I pushed my interlocutor on this, perhaps by reiterating such facts, I would quickly get reassured that we were in agreement, and that there was no need to repeat myself—we shared the same beliefs.

One might be tempted to explain this friction in relatively familiar ways. Indeed, there certainly are tools in the philosopher's kit to diagnose such seemingly elusive disagreements. For instance, we might look to the distinction between explicit and implicit belief.¹ Perhaps, rather uncharitably, I might argue that my interlocutor, whilst they had an explicit, conscious belief that the environment can wildly alter the function of our genes, on some deeper and less conscious level, they operated on the basis of a quite different belief. In particular, in their actual theorising, they applied, quite unconsciously, the belief that genes function in relative isolation from the environment. Perhaps, for instance, when considering the causes of a particular trait, they unconsciously tended only to think only of the organism's genes, and not its

¹ See, for instance, Giordani (2015) and Wittenbrink & Schwarz (2007). Others have captured similar distinctions. See, for instance, Gendler's (2008) distinction between *beliefs* and *aliefs*, and Haslanger's (2006: 99) distinction between our *manifest* and *operative* concepts.

environment. Perhaps they had an unconscious association between genes and blueprints.² Alternatively, my interlocutor might propose that I myself was operating with an implicit belief that diverged from my explicit beliefs. Perhaps, whilst I explicitly believed that the both genes and environment interact to produce an organism's traits, I implicitly believed that the environment is what primarily causes our traits. Instead of a divergence in explicit belief, then, there may have been a divergence in implicit belief between us.

Increasingly, however, I felt that the friction I experienced was about something more minimal than this, and yet something with significance. I began to wonder whether it was instead down to a difference in the *salience* that we accorded to the various beliefs that we held. My interlocutor, whilst they (explicitly and implicitly) believed that environments can change the function of a gene in sometimes wildly different ways, they did not appear to give this fact sufficient salience in their mind—it was not at the forefront of their attention. For me, this fact not only *was* most salient—I felt that it *should* be most salient. It *should* be at the forefront of their mind, I thought—without, yet, a clear understanding as to why. I simply had a feeling that this fact had an epistemic and moral significance to which my interlocutor was not doing justice.

It was this experience that first led me to think about the topic of this thesis. With the notion of *salience* in my repertoire, I started to notice other instances of friction with an interlocutor that seemed like they had something to do with salience. For instance, I would hear a person, let's call them Sally, described by my interlocutor in a way that felt like they made the *wrong* attribute salient about her. They might introduce Sally as a lesbian, before mentioning her job as a politician, for example. Something felt wrong with how they gave Sally's sexual orientation such salience in mentioning it before her other identities—identities that I felt were more important in some way. In particular, I felt that it was in Sally's *interests* for her job as a politician to be made more salient than her sexual orientation—to be made more noticeable and memorable. Again, I didn't yet have the resources to fully *explain* this intuition, but the fact that I could at least pinpoint my initially inchoate feelings of unease as having something to do with *salience*, felt a step in the right direction.

Across these experiences, I sensed that the issue I was having was not with the *content* of what my interlocutor said or thought, but rather how they structured that

² There is a question as to whether implicit beliefs really are beliefs. Some have argued that they are instead non-propositional things like associations (Kihlstrom, 2004).

content so that some of it stood out more than others. I turned, therefore, to what existing theory I could find on the subject of salience. The first research that I came across that seemed helpful was psychological research into *framing effects*, which I will introduce in the next section.

1.2. Where is the philosophy of salience?

Whilst interest is growing,³ salience has generally received relatively little recognition in philosophy. There is no general entry on *salience* in the *Stanford Encyclopedia of Philosophy*, for instance. One does find mention of salience in certain areas, but this tends to occur in relatively circumscribed discussions. One of these discussions concerns *reference fixing* in language. For instance, taking the example of the proposition *the cat is in the carton*, David Lewis suggests that facts about which cat is most *salient* in a given context (e.g. it is the only one in the room) help to fix the reference for this proposition.⁴ As we will see, the sort of salience that interests me is not to do with reference fixing, but to do with the consequences for harm (as well as false beliefs and unwarranted ideologies) that come with giving relatively more salience to some linguistic or mental contents over others. Something closer to this latter notion of salience can be found in certain debates within the philosophy of *moral virtue*. It has been suggested that what one finds salient, in the sense of being *attention-grabbing*, is relevant to the evaluation of one's character as virtuous or vicious.⁵ It might be argued, for instance, that finding the needs of others insufficiently salient is partially constitutive of a moral vice. I too am interested in moral dimensions to salience, but my interest is wider than this narrow focus on virtue and vice.⁶

³ See, for instance, Archer's (forthcoming) edited book on the philosophy of salience. Introducing this book, Archer (forthcoming) also comments that salience is relatively underexplored as a philosophical topic.

⁴ Lewis (1979: 348).

⁵ Chappell & Chappell (2016), Bommarito (2013).

⁶ Further, the notion of salience discussed by these virtue ethicists is sometimes explicated in ways that imply it is more substantive than the particularly minimalist notion that interests me. Sometimes, these philosophers talk about the moral consequences of paying attention to *x* (i.e. making it salient) whilst *ignoring y* (Chappell & Chappell, 2016). In chapter 2 (§2.4.2), I describe this as a sin of *omission*. The notion of salience that interests me is more minimal. I want to show that, even when we do not *ignore y*, simply giving it *relatively less* salience in our attention can cause and constitute harm. This latter claim is more controversial.

Outside of these quite delimited discussions, however, salience is comparatively undertheorised. Sebastian Watzl also makes this point. Watzl writes about salience in the mind, which, as we will see in chapter 3, is constituted by dispositions to *attend* better to some contents over others. Commenting on this sort of salience, Watzl suggests that “In stark contrast to the intensity of public interest in attention and to the richness and detail of its scientific investigation, for a long while professional philosophers [in the West] have almost completely neglected attention [and therefore salience in the mind] as a topic of study”.⁷

Part of the neglect seems to be down to a focus on linguistic and mental *content*, at the expense of considering how that content is *structured* and *presented*. In the philosophy of language, the focus is often on *what* a person communicates (whether that content is explicitly asserted, or what philosophers call *not-at-issue* content, in the sense that it is presupposed, or otherwise implicitly communicated).⁸ In the philosophy of mind, the focus tends to be on *what* someone is mentally representing—namely, *which* beliefs, desires, feelings, and so on, that person has.⁹ Issues of how that content is structured along dimensions of salience, receive less discussion.

This is in sharp contrast to certain other fields. Salience *is* a familiar topic in the worlds of advertising and journalism, for instance. Here, there is a keen appreciation for the great variety of ways in which one can influence one’s audience in order to sell more products, or better persuade one’s readers or listeners to adopt a particular view. An advertiser for a given medicine, for instance, knows that they will find much greater success in selling that medicine if it is advertised as having a 90% success rate, instead of as failing to work in 10% of cases. In other words, consumers will respond differently to logically equivalent information, depending on whether *success* or *failure* is made salient. This is an example of a *framing effect*, something that I will discuss in chapter 2. Other framing effects occur when the *order* in which content is discussed is played with. In the next chapter, we will hear of evidence that our moral intuitions can vary significantly depending on the *order* in which we read two hypothetical scenarios. Journalists also capitalise on our susceptibility to framing effects. Journalists and editors know that which

⁷ Watzl (2017: 5). Watzl does concede that philosophical exceptions to this can be found in writings on phenomenology, as well as in Eastern philosophical traditions. He also cites the discussions of salience in moral virtue debates, as I mention above.

⁸ See, for instance, Stanley’s (2015: Ch. 4) discussion of not-at-issue content.

⁹ Camp picks up on the emphasis on content in imagination (one topic analysed in the philosophy of mind). She says that the emphasis tends to be on “*what* is represented in imagination...[as opposed] to ‘*how*’ that content is imagined” (Camp, 2017: 77).

content is made headline, and which is put front page, matters a great deal for how an audience thinks and feels about a topic.¹⁰ As I will discuss in the next chapter, framing effects can have significant implications; simply shifting the salience in some content can mean the difference between believing something false and harmful, and not.

Whilst this sort of effect of salience is much less discussed in philosophy, philosophers certainly *do* recognise the importance of salience in a different regard. Whilst not a philosophical *topic*, the notion of salience crops up in the context of *building a good philosophical argument*. We are told, for instance, to mention our key argument at the *beginning* of our paper (one way of making something salient in language, remember, is by discussing it *first*). Usually, the rationale is that using salience in this way helps us to make the best argument that we can. Students might also get told that leaving one's key argument until the very end might be inadvisable because one's examiner might miss it—their attention and interest might have waned by the time they reach it. These suggestions as to what a shift in salience can do, however, are rather conservative. Shifts in salience can help one to achieve a stronger, more persuasive argument, or can help one to ensure that one's core argument is actually noted by a sleepy examiner. This is quite different to the more radical claims, made in the framing effects literature. Here, simply shifting the pattern of salience in some content is understood to be capable of eliciting wildly different responses to that content in its audience.

One might wonder, perhaps rather uncharitably, whether this latter sort of (more radical) effect of salience is discussed less in philosophy because we in this field can have a tendency to presume ourselves to be too rational and logical to be susceptible to things like framing effects. Framing effects, as we will hear, are regularly discussed as revealing something *irrational* about human thought. Our training, a philosopher might think, guards us against such irrationalities; it allows us to reflect only on the *content* of what is being said, instead of being swayed by mere presentational shifts in salience within that content. Indeed, Louise Antony has speculated that philosophers tend to suffer from what psychologists call *over-confidence bias*, insofar as philosophers presume themselves to be too rational to suffer from the various irrational biases that psychologists suggest influence human thought. In her words, “In my experience, philosophers are quite ready to allow that other people are affected by irrelevancies like the pitch of a candidate's

¹⁰ Wahl-Jorgensen & Hanitzsch (2009).

voice, or a stereotype that links prettiness with vacuity, but not them.”¹¹ Patterns of salience might count as another so-called *irrelevancy*.¹²

Elisabeth Camp, whom I will introduce more fully in chapter 3, has also suggested that philosophers often neglect aspects of human thought that do not fit what she calls “a philosopher’s fantasy of rationality”.¹³ She explicitly mentions our susceptibility to framing effects in this regard.¹⁴ Fortunately, despite this general neglect, both Camp and Watzl’s philosophical discussions on salience in the mind provide excellent platforms for the discussions in this thesis, and their accounts influence a great deal of what is to come.

1.3. The significance of salience

Why talk about salience? Firstly, the effects of shifts in salience will likely surprise us. To many of our intuitions, a mere shift in the pattern of salience in some content seems too trivial to affect substantially how we think, feel, and act towards that content. Imagine, for instance, being asked a question about gendered leadership styles. Surely it doesn’t matter, one might think, whether one is asked *how do women differ from men as leaders?* or *how do men differ from women as leaders?* Whether women are made more salient than men, by mentioning them first, surely makes no difference to how one answers this question. Even the subtlest shifts in salience can be impactful, however. In chapter 5, we will look at evidence that asking how *women* differ from *men*, as opposed to how *men* differ from *women*, can make us more likely, amongst other things, to endorse stereotypes about sex/gender.¹⁵ If we find these stereotypes to be harmful, then we can see one way in which salience can be relevant to epistemology. Shifts in salience might make the activation of false beliefs more likely. If we find these stereotypes to be harmful, then we can also see one way in which salience is relevant to ethics.

¹¹ Antony (2012: 236).

¹² Antony does in fact mention a non-verbal counterpart to framing effects. When introducing the sorts of irrational, implicit biases to which philosophers presume themselves to be immune, she mentions “we prefer products placed on the right side of the supermarket shelf” (ibid. 236). As we have heard, some framing effects play with *order*, and this example also involves playing with order (namely, the order in which products are placed).

¹³ Camp (2015: 601).

¹⁴ Ibid. (602, fn. 8).

¹⁵ See chapter 5 for a discussion of why I use the conjunction sex/gender.

We will also hear in forthcoming chapters more intimate, and also likely surprising, ways in which salience can have relevance to epistemology and ethics. In chapters 2 and 3, I will briefly consider how patterns of salience might *constitute* an epistemic flaw. Certain patterns of salience might themselves be *unwarranted* for instance. In chapter 6, I consider in more detail the ethical counterpart to this idea. Certain patterns of salience, I will argue, can *constitute* harm. This conclusion in particular might strike one as unexpected. Intuitively, we might think that so long as one listens to what a person says, values their conversational contributions, does not harbour false beliefs, derogatory feelings, or malign desires concerning them, and acts towards them with respect, one cannot harm that person simply by finding one of their attributes more salient than another. Merely attending more to one trait of theirs rather than another, *especially if these attentional dispositions do not have any negative consequences regarding how one thinks (or feels, etc.) about or acts towards that individual*, is surely morally insignificant. This assumption, I will suggest, is mistaken; giving relative salience to one trait of a person over another in one's attention can, in itself, count as a way of harming them. We should learn about salience, then, because we might learn something new and unexpected.

Secondly, we should learn about salience because the very thing that makes salience surprising is also what makes it *powerful*. Generally, we are surprised by what a shift in salience can cause or constitute because that shift seems too minimal to matter. Importantly, the minimal nature of salience perspectives can give them an insidious strength, in two ways. Firstly, their minimal nature means that, not only are the patterns of salience in language and attention themselves rarely consciously noticed by the audience or thinker, but the *inferences* that those patterns of salience solicit us to make are also rarely consciously registered. As we will hear in forthcoming chapters, where inferences are activated under our conscious awareness, such as those to the gender stereotypes mentioned above, the beliefs and associations that we have inferred are especially likely to go on to influence our thought and behaviour. This is because, unaware of how a salience perspective is soliciting us to think and act, we do not attempt to block any problematic inferences that we might be making. The *under the radar* nature that salience perspectives tend to have, then, can give them a special power.¹⁶

¹⁶ Generally, many philosophers of language have highlighted the power that under the radar language has to shape an audience's responses. See, for instance, Stanley (2015), who discusses *not-at-issue* content, Langton & West (1999) who discuss *back-door* speech acts, and Sbisà (1999), who writes about presupposition.

In addition, the fact that salience perspectives are rarely consciously registered means that any harm that they *constitute* (as opposed to cause) can more easily continue unchecked. Further, where the harm that a salience perspective constitutes *is* noticed and articulated, it is all too easy to dismiss it as morally insignificant. Unlike harmful *beliefs* that certain races or genders are inferior, or harmful *practices* of discrimination against these social groups, harming someone simply by making the wrong thing salient about them likely will not elicit the same sympathetic responses. This, I will argue, can give salience perspectives an insidious power, in virtue of making them *effectively* unchallengeable.

Thirdly, salience is pervasive. It structures all of our language and mental content. We cannot communicate without making some contents more salient than others, not least because we must inevitably say one sentence before another. We cannot give everything equal attention in our mind. It is inevitable that we find some things more eye-catching and memorable than other things—inevitable that some things remain at the periphery of our attention. There will always be more memorable, as well as more forgettable, aspects of a person that one meets, for instance. If, as I have suggested above, these patterns of salience have implications regarding whether or not we *harm* a person or group, then we should reflect on which patterns of salience we use. Further, if, as I suggested above, salience has an *under the radar* nature that makes any harm it causes and constitutes particularly difficult to notice and challenge, then its pervasive nature is all the more concerning—we might be causing and constituting harm to others (and perhaps to ourselves) much more regularly than we think.

1.4. Chapter Summary

I begin my investigation in chapter 2 by considering the linguistic side of salience perspectives. I build this notion from psychological research into framing effects. This research provides the empirical support that undergirds many of the arguments in this thesis, namely, that merely changing the pattern of salience in some (in this case linguistic) content, without making any other adjustments to that content, can indeed have significant repercussions for how audiences respond to that content. Linguistic salience perspectives structure contents by giving some *relative* salience over others. In an utterance consisting of two contents, x and y , one might choose to give x relative salience over y (such as by talking about x *before* y), or to give y relative salience over x (such as by

talking about *y* before *x*). It is this shift in relative salience that can generate a framing effect—that can move an audience to respond in different ways to that utterance. In particular, I focus on examples of framing effects that require us to look to facts about our culture in order to explain why a given pattern of salience in language has the effect that it does. These examples set the tone for proceeding chapters; generally, I am interested in how salience interacts with prevalent beliefs and ideologies in our culture. For instance, where a belief is particularly prevalent in culture, I demonstrate that simply making salient in language content that is central to that belief can be sufficient to activate an inference to it. Where that belief is harmful, and/or has epistemic faults, we can criticise the salience perspective that leads to it. This provides an instrumental critique of linguistic salience perspectives—a way of criticising a pattern of salience on the basis of the problems with that which it *causes*. This is briefly contrasted with some ways to critique linguistic salience perspectives on non-instrumental grounds. I conclude by suggesting a *naturalised* methodology for selecting linguistic salience perspectives, which I take up in forthcoming chapters. This methodology suggests that we first review which (morally and epistemically) problematic beliefs and ideologies concerning a given subject are prevalent in our folk-cultural context, and choose a salience perspective that helps to avoid activating those beliefs and ideologies.

Chapter 3 considers the *cognitive* side of salience. I build upon Elisabeth Camp and Sebastian Watzl's work to suggest that cognitive salience perspectives consist in attentional dispositions—dispositions better to notice, remember, and find more cognitively accessible certain mental contents (such as beliefs, or one's experience of things in the world) over others. As with linguistic salience perspectives, cognitive salience perspectives are also characterised by the notion of relative salience. For instance, one might have two beliefs, *x* and *y*, but it is a further question which belief has relative salience over the other; does *x* grab one's attention more than *y*? I suggest various ways in which cognitive salience perspectives might be evaluated, along moral and epistemic dimensions. These involve instrumental and non-instrumental ways of critiquing cognitive salience perspectives, to parallel the suggestions in the previous chapter. This chapter also discusses the type of attention that cognitive salience perspectives consist in; instead of being fully under our voluntary control, for instance, a cognitive salience perspective involves our better attending to one thing rather than another in an *automatic* and *intuitive* way. This fact has repercussions for our ability to change our cognitive salience perspectives at will. Sometimes, I suggest, wider changes to

the culture in which we live will be necessary for us to be able, intuitively and automatically, to find certain things more attention-grabbing than others.

These first two chapters aim to provide a broad account of salience perspectives: what they are, the various ways in which we might critique them, and so on. Having offered this account, I then turn to using linguistic salience perspectives as tools in practice, by applying them to case studies. The case studies discussed in chapters 4 and 5 focus on bringing to life the *instrumental* critique of salience perspectives, whilst those in chapter 6 aim to make vivid the non-instrumental critique. All chapters focus on harm, instead of epistemic faults.

I turn first in chapter 4 to an example of a *generic* salience perspective. A generic salience perspective is one whose pattern of salience can be applied to multiple subjects, because the content that it makes salient is of a relatively abstract nature. The generic salience perspective on which I focus is one that makes salient the internal properties of the subject in question, and the distinctive traits that it happens to be exhibiting. I call this the *substance salience perspective*. Because of its generic nature, I suggest that this salience perspective makes salient content that is central to a number of prevalent biases in our culture: *the correspondence bias*, *psychological essentialism*, and *the fixed mindset bias*. This, I suggest, makes the substance salience perspective liable to activate these biases—biases that, according to common readings of them, are constituted by substantive *beliefs* about the contents in question. Given that many find the beliefs constituted by these biases to be harmful and false, I argue that we can take issue with the substance salience perspective insofar as it plays a causal role in bringing about their harms and falsehoods. My aim in this chapter is to demonstrate the *breadth* of impact that a single salience perspective might have. A single pattern of salience is liable to activate multiple biases.

In contrast to the broader focus of chapter 4, I turn in chapters 5 and 6 to focus on a single subject: sex/gender. In chapter 5, I pick up on one of the cognitive biases introduced in chapter 4, psychological essentialism, and explore what this bias encourages us to believe about the topic of sex/gender in particular. According to how it is commonly understood, this bias encompasses a range of harmful, and false, beliefs. In light of this survey, I suggest which patterns of salience we ought to avoid for the subject of sex/gender, in our language and attention, if we wish to prevent activating sex/gender essentialist beliefs. This includes, for instance, being wary of giving relative salience to sex/gender differences over sex/gender similarities. As part of this discussion, I consider how the idiosyncracies of *sex/gender* essentialism affect how the substance salience

perspective of the previous chapter is likely to function when applied to this specific topic.

Chapter 6 marks a break with what has been a focus on instrumental critiques of salience perspectives. In this chapter, I consider how patterns of salience might *in and of themselves* constitute sex/gender harm. (A constitutive *epistemic* critique of salience perspectives may well be possible, and indeed, I briefly detail some approaches that we might take in chapters 2 and 3, but this is not my focus.) I examine two case studies related to sex/gender that motivate the idea that salience perspectives can constitute harm. Firstly, I argue that rape victim-survivors can be harmed if others make their experience of rape their most salient feature. Secondly, I argue that women can be harmed if others make their bodies more salient than their faces, or conversational contributions. This latter discussion allows me to develop the notion of *perspectival objectification*; this refers to the idea that certain attentional dispositions can themselves *count* as a form of sexual objectification. In other words, salience perspectives count as a form of sex/gender harm. Following this idea, I also consider another sex/gender harm: sex/gender essentialism. Certain attentional dispositions might themselves *count* as form of this bias, which I call *perspectival essentialism*.

In chapter 7, I sum up the main points made in this thesis. I then suggest how these points might affect debates not yet considered in this thesis, such as those relating to racism and hate speech. I also reflect on the implications that this thesis has for social policy, especially where this regards advising how governments and corporations communicate. Further, I reflect on the limitations of this thesis, including, for instance, the lack of psychological studies testing the specific predictions that I make, and suggest related avenues for further research.

1.5. Conclusion

In this thesis, I unite discussions of salience in language with discussions of salience in the mind, to create the concept of a *salience perspective*. *Salience perspectives* refer to the patterns of salience that we use to present and structure linguistic and mental contents, so that some of those contents stand out more than others. The key insight that this thesis argues for, and builds upon, is that changing only the salience perspective applied to some contents, without making any other changes to those contents, can both cause and constitute harm. To demonstrate that salience perspectives can cause harm, I look at

case studies of certain salience patterns activating harmful beliefs and ideologies in our folk-cultural context. To demonstrate that salience perspectives can constitute harm, I demonstrate that certain salience patterns themselves *count* as sex/gender harms.

CHAPTER 2

SALIENCE AND LANGUAGE

2.1. Introduction

Ground beef tastes nicer, and is considered leaner and less greasy, when it is labelled *75% lean*, as opposed to *25% fat*.¹ At least, that's what study participants' evaluations of one and the same sample of beef, under these different labels, would lead you to think. This is an example of a *framing effect*. Framing effects occur when people respond differently to logically equivalent linguistic content when that content is presented in different ways.² More specifically, they respond differently when different aspects of one and the same linguistic content are made salient.

The guiding questions for this chapter are as follows. Firstly, what does it mean for something to be more salient in language? Secondly, how is it that a mere shift of salience in language can have the sorts of dramatic consequences identified by framing effects research? Thirdly, can we evaluate mere patterns of salience in language, and if so, how? And finally, can the answers to these questions tell us anything about how we *should* be using salience in language? The aim in this chapter and the next is to give a relatively broad overview of salience in language and salience in cognition respectively. This will allow the reader better to grasp the lay of the land in respect of salience perspectives, before I consider some applications of the concept in future chapters.

In what follows, I will begin by introducing the concept of a *linguistic salience perspective* in §2.2. This concept refers to the structuring of linguistic content, so that certain parts of it are made more salient than others. Psychological research into the phenomenon of framing effects demonstrates that simply changing the salience patterns in some linguistic content, without making any other changes to that content, can have significant implications for an audience's cognitive and behavioural response to that content. One explanation as to why some linguistic salience perspectives have these effects is, as I suggest in §2.3, that they are able to activate pre-existing cultural and cognitive biases of ours. This can occur simply by making certain content salient in language that is also central to an existing cultural or cognitive bias. Problems arise where

¹ Levin & Gaeth (1988).

² Kahneman & Tversky (1979).

that bias is false, misleading, and/or harmful. This, I suggest, provides us with a way of evaluating linguistic salience perspectives on instrumental grounds, which I make explicit in §2.4.1. In §2.4.2, I contrast this instrumental critique with some brief suggestions as to how we might critique linguistic salience perspectives on non-instrumental grounds. Finally, in §2.5 I consider the wider implications that this discussion has for how we *ought* to use linguistic salience perspectives when communicating with our audience. In particular, I suggest that where a linguistic salience perspective is liable to activate a problematic bias, we ought to adopt a counter-salience perspective that helps to avoid this. This suggestion follows the logic of naturalised epistemology.

2.2. Introducing linguistic salience perspectives

2.2.1. Framing effects.

Think back to the example of ground beef in the introduction. This sort of framing effect uses an *attribute frame*. Here, people are presented with logically equivalent information (75% *lean* is simply another way of saying 25% *fat*),³ and yet make different judgements depending on which attribute is valenced.⁴

Another type of framing effect involves *order frames*.⁵ Here, the same linguistic content is presented in different orders. Consider two common scenarios discussed in the famous *trolley problem*⁶—*Push* and *Switch*. The *Push* case involves pushing a single person in front of a train in order to prevent it from killing five others further down the track, whilst the *Switch* case involves pressing a switch to divert a train away from five

³ Some may wonder whether this sort of change in linguistic salience actually changes the truth-value of that linguistic content (Fisher, 2017). Along with many others writing about framing effects, I take for granted that this is not the case; two texts that differ only in virtue of which attribute frame they use are logically equivalent (see, for instance, Kahneman & Tversky, 1979, and Rybash and Roodin, 1989). This is consistent with the two texts having different truth-values when considered through a pragmatic lens — namely, one that takes into account how a speaker’s wider communicative context shapes the linguistic content’s meaning. In fact, I will look at an argument for a type of pragmatic inequivalence in §2.3.

⁴ Levin & Gaeth (1988).

⁵ Another type of framing effect often discussed in the literature involves *issue frames*, which involve different aspects of a single phenomenon being emphasised, such as free speech vs. public safety regarding a Ku Klux Klan rally (Nelson et al., 1997). Instead of subtle changes to the structure or presentation of a single set of contents, issue frames involve clearly substantively different linguistic contents being communicated, and, as such, are not relevant for my purposes. Others have also argued that they do not count as a real framing effect (Otieno et al., 2013).

⁶ Foot (1967), Thomson (1985).

people, instead directing it to hit one person. When asked if these cases are morally equivalent, studies have found that the professional philosophers surveyed were more likely to answer in the affirmative if the scenarios were presented (in a written text) in the *Push/Switch* order than the *Switch/Push* order.⁷ Further, telling philosophers about framing effects did little to actually minimise the impact that they had.⁸

The framing effect has consistently proven to be one of the strongest cognitive biases in decision-making,⁹ and is observed in many diverse areas, including the media, marketing, academia, and politics. There is a question as to whether people of different cultures display the same susceptibility to framing effects. One study found that Chinese participants demonstrated a stronger susceptibility to framing effects than American participants.¹⁰ Whilst the *degree* of susceptibility may alter, however, all populations studied have indeed shown to be susceptible to framing effects.¹¹ I return to the issue of cultural variation in susceptibility to cognitive biases in chapter 4. A different issue of cultural variation – namely, variation in *how* people respond to framing effects (i.e. *which* inferences they make) – will be mentioned in §2.3 below.

What can we take from the existence of framing effects? For starters, they appear to reveal something irrational about human thought. *Rationally speaking*, it is thought, we should not evaluate one and the same cut of beef differently when hearing different, but logically equivalent, descriptions of its fat content. Similarly, logically speaking, the order in which we hear the *switch* and *push* scenarios should not make a difference to their moral equivalence (or lack thereof). As Amos Tversky and Daniel Kahneman put it, “different representations of the same choice should yield the same preference. That is, the preference between options should be independent of their description.”¹²

Some have questioned whether all framing effects do indeed demonstrate irrationality. For instance, Shlomi Sher and Craig McKenzie suggest that, at least for some types of framing effects, the audience is making rational, pragmatic inferences based on their speaker’s choice of frame.¹³ For instance, whilst *half full* is a logically

⁷ Schwitzgebel & Cushman (2012). The participants were drawn from a website, as well as from 25 major research universities. Further studies concerning framing effects affecting our moral intuitions are documented in Sinnott-Armstrong (2008).

⁸ Schwitzgebel & Cushman (2012).

⁹ Kühllberger (1988).

¹⁰ Levinson & Peng (2007).

¹¹ Ibid.

¹² Tversky & Kahneman (1986: 253).

¹³ Sher & McKenzie (2008). Also see Mandel (2014) for a different defence regarding the rationality of (at least some) framing effects.

equivalent attribute frame to *half empty*, the audience might, through inferring some background knowledge of the speaker, find these to be informationally inequivalent. The speaker chose the framing of *half full*, the audience might think, because the glass had less water in it earlier; the speaker's positive framing implies, so goes the thought, a positive development regarding the amount of water in the glass.

This is not a debate that I wish to enter into here. Instead, I want to focus on a different issue: namely, the *unconscious nature* of making these ir/rational inferences. Framing effects research receives so much press precisely because our responses surprise us. Psychologically, we do not tend to think that small shifts in the presentation of content will shift our response to that content. This suggests that any inferences that we are making are largely unconsciously processed. Even the sorts of rational inferences proposed by Sher and McKenzie are, in their words, “surely *implicit*—i.e., drawn below conscious awareness. Otherwise, the non-equivalence of attribute frames would have been self-evident prior to our analysis”.¹⁴ We will return to the implications that this has in §2.3.

Whilst most of us do not tend to think that we will be susceptible to framing effects, we heard in the introduction that philosophers might be particularly reticent to accept this fact. We have already seen evidence, however, that order effects sway professional philosophers. But it is not just us. Physicians have been swayed by framing effects when deciding which treatment programme to adopt during a disease outbreak.¹⁵ Senior police officers, experienced in making criminal justice-related decisions, have been susceptible to framing effects when evaluating the accuracy of counter-terrorism techniques.¹⁶ Financial planners have been found prone to framing effects when making judgements about financial risk avoidance.¹⁷ Generally, those with high academic credentials have been found to be affected by framing effects in a similar way to those without such credentials.¹⁸ The sorts of deliberative skills and expertise knowledge that experts in these fields might develop, then, do not protect one from this bias.

¹⁴ Sher & McKenzie (2006: 489).

¹⁵ Tversky and Kahneman (1981).

¹⁶ Garcia-Retamero & Dhami (2013).

¹⁷ Roszkowski & Snelbecker (1990).

¹⁸ Loke & Tan (1992).

2.2.2. Linguistic salience perspectives

I build the notion of a linguistic salience perspective from this discussion of framing effects. Consider a journalist, making decisions about how to present the costs and benefits of immigration in a newspaper article. They are obliged, let's say, to include statistics on crime committed by immigrants, as well as information about the skills and taxes that immigrants contribute. They are obliged, then, to include a particular set of linguistic content. Even with this constraint on what content they can communicate, the journalist can decide which of these facts to *headline*—namely, which of these facts to make salient. They mock up two articles, the only difference between them being which implication of immigration is in bold at the top of the page.

We can say, in this case, that the two articles invoke different linguistic salience perspectives. A linguistic salience perspective refers to the *structuring* or *presentation* of linguistic content. (The term *structuring* looks better to capture the *order frames* discussed earlier; it captures ways of making some content more salient than another by discussing it first. The term *presenting* looks better to capture the *attribute frames*; namely, it captures ways of making some content more salient than another by choosing one or other of a logically-equivalent attribute to describe it.)

Linguistic salience perspectives structure contents by giving some *relative* salience over others.¹⁹ So, while an article might discuss both the costs and benefits of immigration, by discussing the costs first, those costs are given relative salience over the benefits. With those costs more salient in language, this linguistic salience perspective will tend to make those costs more salient in the audience's *minds*. In other words, a linguistic salience perspective will tend to produce its *counterpart cognitive* salience perspective in its audience. Having the costs of immigration relatively more salient in one's mind means better *attending* to those costs—it means better noticing and remembering them.²⁰

¹⁹ I have borrowed talk of *relative salience* from Watzl (2017: Ch. 4), who, as we will see in the next chapter, applies this idea to salience in the *mind*.

²⁰ Indeed, order effects most commonly have the effect of the participant favouring information that is presented first (a *primacy effect*). This effect occurs because the initial content becomes anchored in the person's mind. In some limited cases, however, information presented last is favoured by the participant (a *recency effect*). Where this does occur, it is usually where the topics under discussion are ones that individuals know little about, or show little interest in. By contrast, primacy effects are strongest for subject matter that subjects are highly familiar with, and those that subjects find controversial and/or interesting (see Haugtvedt & Wegener, 1994). As will become clearer in what follows, I am interested in salience patterns for content that subjects are highly familiar with—in particular, where there are culturally entrenched biases concerning that content.

Precisely what making something salient in one's mind amounts to will be addressed in the next chapter. There, I will be examining the work of Elisabeth Camp and Sebastian Watzl, who discuss patterns of salience in the mind. Indeed, it is from Camp's account of salience in the mind that I have borrowed the terminology of *perspective*.²¹ This chapter draws out the linguistic counterpart to this idea.

One question one might have is this: do the two newspaper articles above, with their different linguistic salience perspectives, present identical content? Can we say, in other words, that they have the same truth-value? Or, does changing the *order* in which various linguistic contents about immigration are communicated mean that one alters the truth-value of those contents? Whilst these are important questions, they are not ones that concern me here.²² Linguistic salience perspectives are interesting not because, where they change audience responses to some content, they do so without altering that content's truth-value. They deserve our attention because their effects tend to surprise us. As already mentioned, the framing effects literature has received so much press because we do not tend to expect simple shifts in linguistic salience to be consequential. We are surprised when we are confronted with evidence that, for instance, our responses to some content shift simply when that content is reordered.

In part, our surprise is explained by the fact that linguistic salience perspectives function in an *under the radar* manner. We are usually unaware of how a linguistic salience perspective is soliciting us to respond to some linguistic content. This is because they do not make their requests explicit, such as through the plea *you should pay more attention to the costs of immigration!* Instead, they subtly play with the presentation and structure of content, in a way that often bypasses our conscious awareness. This can give them a manipulative quality, insofar as they do not, in Cass Sunstein's words, "*sufficiently engage or appeal to [people's] capacity for reflection and deliberation* [emphasis in original]".²³ This, as we will see, can make them especially effective in shaping our thought and behaviour.

An important question that I will be asking in this chapter is as follows. Suppose that an article communicates only truths about a subject, such as UK immigration. Further, suppose (rather idealistically) that it communicates *all relevant* truths about that

When I discuss linguistic salience perspectives that play with order, then, I will presume that they will have a primacy effect.

²¹ Camp (2017).

²² See Fisher (2017) for a discussion concerning whether various changes in linguistic salience do in fact change the truth-value of the linguistic content communicated.

²³ Sunstein (2016: 82) is talking about modes of manipulation more generally here, as opposed to the specific issue of using *salience* in language.

subject. Might there nevertheless be grounds for criticising that article for giving *relative salience* to certain truths over others? Might it, by encouraging its audience to make one fact relatively more salient in their minds than another, mislead its audience? In other words, whilst we cannot criticise the article for being *untrue*, or perhaps for being misleading in virtue of *omitting* certain truths, can we nevertheless find some other epistemic and moral grounds on which to critique linguistic salience perspectives?

2.3. Linguistic salience perspectives in context

In what follows, I will look at three examples of how knowledge of cultural beliefs and ideologies is necessary to understand why a linguistic salience perspective has affected audience responses in the way that it has.

2.3.1. Example 1: Advertising beef.

Why did people evaluate one and the same cut of beef as tasting worse and being greasier when it was labelled as *25% fat*? Part of the explanation requires looking at the cultural associations with, and beliefs about, *fatness* and *leanness*.²⁴ Our culture associates leanness with health and desirability, and fatness with the opposite; all around us are adverts selling products on the basis that they include 0% fat (exploiting a prevalent belief that fat *per se* is bad for you, and that lean products *per se* are healthy), adverts depicting thin, toned bodies without a gram of cellulite as desirable, and narratives associating fatness with laziness, greed, and so on.²⁵

Inferences to these cultural associations and beliefs are *cognitively accessible*, in the sense that it takes little cognitive labour to think of them—they easily pop into our minds. In part, this is because of the sheer number of times that these beliefs and associations are repeated around us. They are also *socially licensed*, insofar as, in our culture, most people are disposed to regard inferences to these associations and beliefs as

²⁴ Various explanations as to why framing effects occur have been proposed. See, for instance Sher & McKenzie (2006), for an *information leakage* account, and Levin (1987) for an *associative* theory. My aim in this section is simply to offer *one* explanation as to why certain shifts in linguistic salience perspective result in a framing effect.

²⁵ See, for instance, Eaton (2016) on the association between fatness and laziness, amongst other negative traits. I will return to the specific issue of cultural associations with fatness in the next chapter. What precisely associations *are*, I do not address here. Minimally, we can say that A and B are associated when they are reliably correlated (Simpson, 2012: 572-3: fn. 16).

legitimate.²⁶ It takes very little to activate these associations and beliefs. Simply making salient content central to these beliefs and associations – namely, making either fatness or leanness salient – can be sufficient causally to activate them.

Part of the success of this framing effect was down to its subtlety. The experimenters did not explicitly mention any cultural associations or beliefs. They did not say *this cut of beef is 25% fat, and fat, remember, is unhealthy, greasy, and will make you unattractive*. Interestingly, if they had, it is possible that the study participants might not have as consistently responded in ways congruent with those cultural associations and beliefs (i.e. in ways that are consistent with their endorsing the belief that fat is undesirable, and so on). That is because *implicitly* activated cultural associations and beliefs are especially effective in shaping our thought and behaviour. Research shows that where stereotypes (i.e. cultural associations and/or beliefs whose content involves a social group and an attribute) are implicitly activated, more people respond in stereotype-congruent ways. Implicit activation might involve, for instance, unscrambling a word, e.g. *emotional* (from *mtoeinlta*), before evaluating the behaviour of a person who is stereotypically associated with that word, e.g. a woman.²⁷

Implicit stereotype activation even affects individuals who explicitly reject stereotypes, and who usually make efforts to respond in ways that demonstrate this rejection.²⁸ Implicitly activated stereotypes are so successful at producing stereotype-congruent responses in us because we are not consciously aware of their activation. This means that our ability consciously to monitor our thought and behavioural patterns, and to suppress any stereotypes that we might on reflection reject, is precluded. This has benefits for a speaker who wishes to activate a stereotype in their audience's mind without wanting to *argue* for that stereotype.²⁹ Whereas the audience can ask the speaker for evidence when they *explicitly* assert a stereotype, this option is precluded when the hearer is not aware that a stereotype has been activated. Advertisers are quick to take advantage of this fact.³⁰

It is worth emphasising the behavioural dimension of this stereotype activation. In the case at hand, the activation of the stereotype about fat being unhealthy does not *stay in the head*. It translates into people being more likely to buy products advertised as

²⁶ These definitions of *cognitively accessible* and *socially licensed* are borrowed from Fraser's (2018: 735-6) paper on rape metaphors.

²⁷ Banaji et al. (1993).

²⁸ Devine (1989: 194-5)

²⁹ See, for instance, Simon-Vandenberg, et al. (2007).

³⁰ Till et al. (2011), Schmidt et al. (2017: 3).

lean. Other stereotypes have more complex relationships to action, as I will be discussing in future chapters. For instance, *stereotype threat* refers to scenarios where people either are, or feel, at risk of conforming to stereotypes about their social group. Consider another example of implicit stereotype activation. Research demonstrates that putting a gender tick box at the top of a maths test leads to girls performing worse in that test than in a test without such a tick box.³¹ This occurs because of how cognitively accessible and socially licensed is the stereotype that women are bad at maths; simply reminding girls of their gender in the context of a maths test is sufficient to result in the implicit activation of this stereotype.³² The point most relevant for my purposes, however, is that this stereotype does not *stay in the head*. It results in girls unconsciously living up to the stereotype by performing worse in that test. We can also flag the epistemic consequences that many have identified with stereotype threat. Whilst belief standardly aims to fit the world (my belief that the table is brown aims to fit the way the actual table in the world is), the belief that girls are worse at maths shapes its subject matter in the world to fit *it* (the girls come to fit the belief). This is often said to give the belief an epistemic flaw in virtue of having the *wrong direction of fit* with the world.³³

If linguistic salience perspectives count as one way of implicitly activating stereotypes, then, we should be mindful of the complex behavioural (and epistemic) consequences that they might engender. This should warn us that small can be powerful. The sort of change in some linguistic content that comes with altering its linguistic salience perspective is often so small that the audience does not even notice it (or, perhaps, dismiss as too trivial to be of consequence). The discussion here demonstrates that we should not underestimate the impact of small changes to salience patterns in language.

2.3.2. Example 2: Oral health messages.

Compare the following two health messages:

(A) “If you floss your teeth regularly, you will have healthier teeth and gums”!

³¹ Danaher & Crandall (2008).

³² Whilst this is an example of implicit stereotype activation, this is not an example of a *linguistic salience perspective* implicitly activating a stereotype. The difference between the maths test in the intervention condition (i.e. with the M/F tick box) and the control condition (i.e. without the M/F tick box) went beyond merely altering the structure or presentation of one and the same set of linguistic contents; instead, the difference was one of introducing *new* content (i.e. an M/F tick box) that was not present in the control.

³³ See, for instance, Langton (1993a, 2000), who is building upon Anscombe (1957), and Humberstone (1992) for a discussion of these issues.

(B) “If you don’t floss regularly, the health of your teeth and gums is at risk”!³⁴

This is a standard example of attribute framing; (A) presents information with a *gains* frame (i.e. it makes gains salient), whilst (B) presents that same information with a *loss* frame (i.e. it makes losses salient). Research demonstrates that using attribute frame (A) is more effective with White American and White British study participants; these individuals were more likely to report intentions to floss more regularly when the *positive* health impact of flossing was made salient. Things are different, however, in East Asia. East Asian participants (largely from Japan and China) are more likely to be motivated to floss when they receive the oral health message with the *negatives* made salient, as it is in (B).³⁵

Why is this the case? Again, the suggestion is that we look to culture. It has been argued that dominant cultural narratives in the U.S. and Britain focus on individual achievement, distinction, and autonomy, as well as self-promotion. Steven Heine and colleagues refer to various socio-historical factors as demonstrating this focus, including media tropes of lone cowboys, and men who are masters of their fates, as well as social practices involving regularly praising children, and calling them *special*.³⁶ This, the authors suggest, motivates White Americans and White Brits to pursue positive gains. Framing oral health care in terms of the gains that can be made, then, appeals to this demographic.

Instead, dominant narratives in Japan and China are said to focus on the relationships that an individual has, and how an individual’s own reputation is interdependent with the reputation of the groups (e.g. familial, societal) of which they are a part. Losses of an individual’s own reputation, therefore, negatively impact their group, and group reputation is easily damaged but not easily saved. Various cultural practices have developed, therefore, that focus on *avoiding* reputation losses. For instance, *Ki ga sumanai* and *hansei* are two Japanese practices that involve drawing attention to one’s deficits, and communicating one’s shame in, and one’s efforts to fix, these deficits.³⁷ The fact that avoiding losses (of reputation) is a central cultural narrative in this part of the world helps to make sense of how framing oral health care with a loss frame is particularly motivating for those with an East Asian upbringing.

³⁴ This example is cited in Brick et al. (2016).

³⁵ Ibid.

³⁶ Heine et al. (1999: 769). Their focus is on North America, but similar things have been said about Britain (see the references in Brick et al., 2016).

³⁷ Heine et al. (1999: 770-1). For similar practices in East Asia (including China) more generally, see Hamamura et al. (2009).

We can understand the cultural variation in this framing effect only by taking into account the differences in the formative cultural upbringings of the participants. In the case at hand, it looks like we need to look beyond the more simple sorts of associations and beliefs that were mentioned in the previous example. The differences between the cultures look to be more amorphous than this—they look to involve differences in *ideology*. Whilst there has been a lot of attention on the notion of ideology, Eric Swanson’s recent account provides a clear definition.³⁸ An ideology, according to Swanson, is a mutually supporting cluster of beliefs, associations, affective dispositions, codes of interaction, values, practices, and so on.³⁹ In East Asia, for instance, caught up with beliefs that one is a relational self are *affective dispositions* of shame connected to feeling oneself lacking, and *codes of interaction* that require one to communicate this feeling to others. It is this *ideology* that is cognitively accessible and socially licensed in East Asian countries, and is liable to become activated simply through making salient content central to that ideology—namely, the potential for *loss*.⁴⁰

What the examples in §2.3.1 and §2.3.2 do both demonstrate, however, is how cultural biases (whether they be associations, beliefs or ideologies) might play a role when it comes to explaining how linguistic salience perspectives function. (I will sometimes use the term *bias* as an umbrella term for these phenomena—as will become clearer later, I do not take the term *bias* to be inherently pejorative.) Both examples involve emphasising different attributes of logically equivalent linguistic content. How about linguistic salience perspectives that play with *order*? Can these also activate existing cultural associations, beliefs and ideologies?

2.3.3. Example 3: Genetics pedagogies.

Many of us will have come across the idea of *dominant* and *recessive* traits in relation to genetics, notions that have become associated with Gregor Mendel. Mendelian genetics⁴¹ is often explained through experiments with peas; the idea is that if one breeds a yellow and green pea together, the first offspring is always yellow, as the *gene for* the yellow trait

³⁸ The notion of *ideology* has a rich history. See, for instance, Althusser (1965), Gramsci (1971 [1929-1935] 323-3), West (1989: 232), Tirrell (1999), and Haslanger (2011).

³⁹ Swanson (forthcoming).

⁴⁰ This means, therefore, adding *mere patterns of salience in language* to the growing list of ways in which it has been proposed that ideologies can be activated. See Swanson (forthcoming) for other examples.

⁴¹ We should distinguish Mendelian ideas from those of Mendel himself; in addition to not using the language of *genes*, Mendel had a far more complex view of genetic qualities such as *dominance* than the common reconstructions of his ideas. See Olby (1979).

is dominant. Similarly, the *gene for* the smooth seed is dominant to its recessive wrinkled counterpart, and tall is dominant to short. Mendelian ideas – for example, of *genes for* traits and genes as relatively isolable units separate from the environment – frame the introduction and basis for classical genetics textbooks and curricula. Later, the students learn the qualification that, whilst Mendelian genetics accurately captures some instances of genetic causation, much of genetics is not as simple as this; the environment (both within and external to the organism) profoundly complicates the function of genes. (The reader may remember that these ideas were briefly introduced in the previous chapter.)

In particular, the students learn about *interactionism*, which emphasises how traits arise from the complex interaction between biology and environment.⁴² For example, interactionist research focuses on how the same gene contributing to a plant being the tallest in comparison to its competitors in one environment, e.g. high altitude, might contribute to that plant being the smallest in another environment, e.g. low altitude. A change in the environment of a so-called dominant gene can even render it *recessive*.⁴³ Interactionist research is not simply reminding us that the environment *also* contributes to the traits of an organism, i.e. *in addition* to genes. It is showing that the function of a gene can be so profoundly shaped by its environment that one and the same gene might be *for* tallness in one context whilst being *for* shortness in another. A single gene can dispose an organism to develop *opposite* traits in different environments. This means that we cannot, for many genes, refer to them as being *for* any trait in particular—at least not without specifying the environment. This interactionist research is discussed much later (i.e. in the final chapter or two) and more briefly in mainstream textbooks.

Now consider a pedagogy that *reverses* the information in the standard textbook. In other words, consider a pedagogy that starts with interactionist research, and ends with Mendel. Could this mere shift in order – in linguistic salience perspective – impact how students conceptualise genetics?

It certainly can, according to the results of Annie Jamieson and Gregory Radick's *Genetics Pedagogies Project*.⁴⁴ Retaining largely the same content as the original, so that none of the traditional views are rejected, and no completely new ideas are introduced, Jamieson and Radick altered the chronology of the traditional pedagogy in precisely this

⁴² The term *interactionism* can be used to mean a range of different things (Tabery, 2014: chapter 1). The way in which I cash out the term in the paragraph above follows what Tabery (ibid. chapter 2) calls *developmental concept of interactionism*, inspired by the work of Hogben and Lewontin.

⁴³ Jamieson and Radick (2013: 589).

⁴⁴ University of Leeds (2019).

way.⁴⁵ Their project sought, in Radick’s words, “to take what is peripheral in the existing curriculum and make it central, and to take what is central and make it peripheral”.⁴⁶ The result? Students came away with less deterministic beliefs about genes after taking the revised course.⁴⁷ *Genetic determinism* is the view that genes are the *super-causers* of most if not all of our traits, in the sense that environmental factors play little to no causal role in the creation of our traits.⁴⁸ Even for complex traits like sexuality, genetic determinists are prone to posit a single *gene for* these traits. Where a trait is deemed determined by genes in this way, it is considered to be largely unchangeable. (It should be noted that this reduction in genetically determinist beliefs in the students did not come at the expense of a proper *understanding* of genetics. One might worry that starting with the complex cases and ending with the simple cases might hinder the students’ ability to pick up knowledge about genetics. Jamieson and Radick found, however, that the intervention group received similar results in their end-of-semester exam on genetics to the control group, who were instead following the traditional pedagogy.)⁴⁹

Why did Jamieson and Radick’s revised pedagogy have the impact of reducing support for genetic determinism? Just as with the example concerning beef labelling earlier, we find the answer by examining the different curricula through the lens of our cultural context. We are surrounded by oversimplified and deterministic conceptions of genes in the media, as well as in scientific and educational communities. For instance, there is an obsession in the media with finding the *gene for* homosexuality, or the *gene for* a high IQ.⁵⁰ In fact, it has even been proposed that genetic determinism constitutes a prevalent folk psychological bias—namely, a bias that affects most people across many

⁴⁵ Now, it is true that Jamieson and Radick’s project did not achieve the status of a *pure* framing effect. Some content changes were necessary in the end due to practical limitations in the study, contrary to the authors’ original intentions. But, as the authors conclude, “[i]t is not difficult to imagine a successor study that, conditions (and funding) allowing, would enable the design of a more closely comparable Mendelian counterpart course—no different in length, delivery, content, assessments, accreditation and so on” (Jamieson & Radick, 2017: 1282). The results from the framing effects literature suggest that a project refined in this way could well still alter student attitudes.

⁴⁶ Radick (2016: 169).

⁴⁷ Jamieson and Radick (2017: 1273) tested this by asking the students how much they agreed to various statements congruent with genetic determinism, such as *cloning can produce an identical copy of an animal, so you could recreate a much-loved pet*.

⁴⁸ This definition of determinism is influenced by Kampourakis’ definition of genetic determinism (in Jamieson and Radick, 2017: 1265).

⁴⁹ *Ibid.* (1275). This paper also provides examples of power point slides and other teaching materials used by Jamieson and Radick, which demonstrate how they adroitly made the complex cases accessible (*ibid.* §3.2).

⁵⁰ Condit et al. (1998), Dar-Nimrod & Heine (2011).

different cultures.⁵¹ However widespread, the bottom line is that genetic determinism is a cognitively accessible and socially licensed bias. The traditional pedagogy, by beginning genetics education with the sort of *gene for* language that Mendelian genetics encourages, gives relative salience to instances of genetic causation that are *also* central to genetic determinism. Namely, it makes salient instances where a given gene does reliably correlate with a particular trait, across different environments. In this way, the traditional pedagogy was liable to activate existing biases that *also* make those instances of this type of genetic causation salient—namely, our existing genetic determinism biases.

2.4. Ways of evaluating linguistic salience perspectives

2.4.1. Instrumental critique.

This last example of genetics textbooks provides us with more evidence that linguistic salience perspectives can causally activate a bias (whether it is constituted by beliefs, associations, or ideologies). Most importantly, however, it provides us with evidence that a mere shift in linguistic salience perspective can mean the difference between believing something *false* and *damaging*, and not. Genetic determinism is considered to be false by contemporary (philosophy of/) biology.⁵² It is not the case that most if not all of our traits are caused by our genes, with little to no environmental influence. Further, the false beliefs of genetic determinism have a long history of being used to further social oppression.⁵³ Partly, this is because genetic determinists treat the outcomes of social inequality as instead caused by unchangeable genetic factors. For instance, unsound genetically determinist arguments have been used to justify removing educational resources for black people. Where evidence has been cited of black people having lower IQs than white people, presuming intelligence to be genetically determined instead of socially shaped has allowed individuals to argue that trying to improve the IQs of black

⁵¹ Dar-Nimrod & Heine (2011) suggest that it might be a specific strand of a more general, essentialist folk psychological bias. This more general essentialist bias will be discussed in more detail in Ch. 3. Sometimes, folk psychological biases are distinguished from cultural biases by claiming that the former are universal and innate. Increasing research into the cultural variability of our biases, including those once supposed universal and innate, is casting doubt on this distinction. For a brief discussion of these issues, see Chapter 4 (§4.3).

⁵² Lewontin (2000), Kitcher (2001), Oyama et al. (2003) and Kampourakis et al (2014).

⁵³ In addition to the references in fn. 52, see Belkhir & Duyme (1998) and Keller (2005).

people through education is a pointless task.⁵⁴ The fact that Jamieson and Radick's revised pedagogy managed to reduce acceptance of genetic determinism, then, looks to be a good thing both epistemically and socially.

We can say, then, that a linguistic salience perspective can be criticised on instrumental grounds where its *effects* are problematic in some way. We might criticise it for activating false and/or harmful beliefs, as per the traditional genetics textbooks. We might also criticise it for leading to problematic *non-doxastic* phenomena. The examples in §2.3.1 and §2.3.2 demonstrated linguistic salience perspectives activating certain *associations* and *ideologies*. Whilst associations and ideologies are not things with truth-conditions, given that they either are or include components that are non-propositional, they can be assessed on epistemic grounds. For instance, it is commonly said that an emotion can be *warranted* or not, depending on facts (facts about the situation may show that my fear of someone is unwarranted, for example).⁵⁵ As with beliefs, associations and ideologies might also be harmful, such as where black people are associated with low intelligence, or where an ideology prescribes that one treat black people as second-class citizens.

Let's return to the UK immigration newspaper article, to assess the ways in which we might critique its choice of linguistic salience perspective. Consider the version of the article that headlines the statistics on immigrant crime, and discusses how many immigrants are coming into the UK in its opening paragraphs. Much later, it discusses how the UK benefits from immigration—from the skills and revenue that immigrants contribute. By giving relative salience to immigration numbers and immigrant crime over the benefits from immigration, could this newspaper play into existing biases?

Xenophobic beliefs and ideologies – namely, those involving prejudice against people from other countries – are arguably widespread, and have a long history.⁵⁶ As for the UK in particular, it is often argued that right wing parties using xenophobic claims are increasingly entering the centre-ground in UK politics. Consider the UK Independence Party (UKIP), who surprised the nation by taking the most votes and seats

⁵⁴ See Jensen (1969) and Hernstein & Murray (1994) for examples of such an argument, and Lewontin (1970) for a rebuttal.

⁵⁵ See, for instance, Elgin (1996). Others talk of emotions being *fitting* or *appropriate* (D'Arms & Jacobson, 2000). For discussions of epistemic criteria beyond truth and falsity, more generally, see, for instance, Richard (2008), Gibbard (1990), Langton (1993a, 2004), Haslanger (2013), and Anderson (1995). These alternative criteria include, for example, accuracy, warrant, and aptness.

⁵⁶ Rydgren (2004), Wearing (2016).

in the 2014 European elections, and who is often credited with the success of the Brexit movement. UKIP has regularly been accused of taking advantage of latent xenophobic attitudes in British society. One incident that sparked controversy was their use of a poster with the words *Breaking Point* over an image of a long, winding queue of non-white immigrants. Some were concerned that it was no coincidence that this poster had uncomfortable visual similarities to Nazi propaganda of migrants post-WW1, a poster describing such migrants as “parasites”.⁵⁷ We might worry, then, that this poster represents a cultural association between immigrants and swarms of parasites. Even the main parties, however, have been criticised for pandering to Xenophobic attitudes.⁵⁸ The fact that immigrants are discriminated against in the labour and housing market could be taken as material proof of xenophobia.⁵⁹

Some might argue, on the basis of this sort of evidence, that xenophobia is a cognitively accessible and socially licensed bias. In fact, some have argued that this type of cultural xenophobia is symptomatic of a broader, more universal bias. Psychologists Catherine Cottrel and Steven Neuberg suggest that our psychology has been primed by evolution to emphasize the potential threats associated with people from other groups.⁶⁰

If these biases are cognitively accessible and socially licensed, this gives us reason to think that a linguistic salience perspective that emphasises that which is central to those biases is able to activate them. The newspaper article above does indeed make salient content central to xenophobia, namely, the *threats* and *costs* associated with immigration. It might, then, play a role in the activation of this bias.

If we can show that xenophobia is epistemically problematic and/or harmful, then we can criticise linguistic salience perspective that *leads* to the epistemic and moral problems that xenophobia encompasses. Here, we employ an *instrumental critique* of linguistic salience perspectives. A linguistic salience perspective can be evaluated on the basis of its upshots.

⁵⁷ Stewart & Mason (2016).

⁵⁸ *The Guardian* editorial (2016).

⁵⁹ See Rydgren (2004: 124) and the references therein. That xenophobia is an *ideology* would help to explain this material dimension; immigrants would likely be discriminated in this way because of, for instance, associations between immigrants and low competence, feelings of fear towards immigrants, and codes of interaction such as *treat immigrants as less trustworthy*. Others have referred to xenophobia in ways consistent with it being an ideology. For instance, it is often related to common emotional responses (such as fear and scorn) and codes of interaction (such as anti-social behavior towards those in an *out-group*). See Cottrel & Neuberg (2005).

⁶⁰ Cottrel & Neuberg (2005).

2.4.2. Non-instrumental critique.

Can a linguistic salience perspective be criticisable in its own right, namely, independent of its effects? Whilst this is not something that I will discuss in depth here, I suggest two ways in which an argument along these lines might proceed. I mention these non-instrumental critiques here for two reasons. Firstly, as mentioned in the introduction, my aim in this chapter and the next is to give a broad overview of salience perspectives in both linguistic and cognitive form. Whilst I focus on instrumental critiques in the case studies considered in chapters 4 and 5, discussing both instrumental *and* non-instrumental critiques here provides the reader with a better grasp of the lay of the land regarding salience perspectives. (I do return to non-instrumental critiques in chapter 6, considering in detail a version introduced briefly in the next chapter.) Further, the second instrumental critique that I discuss below helps us to clarify the notion of salience used in this thesis, by contrasting it with another everyday use of the term.

The first method for criticising linguistic salience perspectives on non-instrumental grounds concerns the concept of *licensing*. Language licenses a belief or ideology where it makes it seem *normatively acceptable*. How can language do this? Swanson suggests that when we use a certain word, such as *nerd*, we imply that, all things considered, it is acceptable to use the word *nerd*—that is to say, we *license* use of that word.⁶¹ Generally, Swanson suggests that where we presume a speaker to be rational, cooperative, and so on, “we will have reason to think that the speaker is putting things in an acceptable way”.⁶² Now, there is an ideology associated with the word *nerd*, including, for instance, associations between those with concern for academic achievement and social awkwardness, and codes of interaction that encourage disparaging and subjugating those deemed to be *nerds*. For use of the word *nerd* to be acceptable, Swanson argues, the ideology associated with it “would itself have to be acceptable—indeed even *good*

⁶¹ Swanson (forthcoming: §1) refers to these normative implicatures associated with use of a given word as “acceptability implicatures”. He says that uses of any word (to which I add uses of any linguistic salience perspective) are associated with a family of acceptability implicatures. These acceptability implicatures are, in Swanson’s (forthcoming: 3) words, “reflexive and higher-order: their content is a commentary on how what was said was said. In a word, they convey that it was acceptable to put it that way”. The particular term *licensing* is borrowed from Tirrell (2012), whose account has similarities with Swanson’s.

⁶² Swanson (forthcoming: 4). The speaker may try to cancel this implicature, by, for instance, saying “But I don’t mean to suggest that putting it that way was acceptable all things considered”. Cancelling the implicature is not always possible, however; for instance, it is often argued that one cannot cancel the implicatures of a slur (Swanson, forthcoming: 2-3).

[emphasis in original]”.⁶³ In an important sense, then, the speaker implies that the ideology associated with *nerd* is acceptable simply by using the word *nerd*. Swanson’s suggestion, then, is that a word can normatively license ideologies (or associations and beliefs) related to it.

The significance of this fact is that licensing is a normative, as opposed to causal, notion.⁶⁴ This means that, where the word *nerd* licenses an ideology about nerds, if we deem that ideology to be problematic, then our use of the word *nerd* bears a *constitutive* relation to the problems of that ideology. Applying this to linguistic salience perspectives, then, the idea would be that use of a linguistic salience perspective can, in addition causally to activating associations, beliefs and/or ideologies associated with that perspective, normatively *license* them. Where it does so, it bears a *constitutive* relation to any problems inherent to those biases. Where a linguistic salience perspective *constitutes* the (for example) falsehoods or harms of a bias, we can find it non-instrumentally problematic.

Relating this back to the immigration article, making salient instances of immigrants posing threats might be a way of saying *this way of using salience in language is acceptable*. One obvious way in which using that linguistic salience perspective would be acceptable is if xenophobia – an ideology associated with immigrants being threatening – is also acceptable. We might, then, want to pursue an argument that the immigration article above is a way of licensing xenophobia. If we were successful, we could say that instead of just *causing* the epistemic and moral problems constituted by xenophobia, the linguistic salience perspective employed by the immigration article *constitutes* those problems, in virtue of licensing them. In other words, we can criticise it on non-instrumental grounds.

The second way in which we might offer a non-instrumental critique of a linguistic salience perspective is if certain patterns of salience can themselves constitute a type of epistemic flaw. To see this, let’s first consider Elizabeth Anderson’s claim that a linguistic account of a phenomenon that contains only truths *and yet omits other, relevant truths* constitutes an epistemically biased account of that phenomenon. To illustrate this idea, Anderson cites a book by the controversial American Black Nationalist Louis Farrakhan. Farrakhan’s 1991 book *The Secret Relationship between Blacks and Jews* describes the role of Jews in the Atlantic slave system. Anderson lists many claims made by the

⁶³ Swanson (forthcoming: 10).

⁶⁴ Tirrell (2012), Swanson (forthcoming).

book that are true. For example, “that Jews had considerable investments in the Dutch West India Company, which played a significant role in the seventeenth century Atlantic slave trade...[and] that a larger percentage of Jews living in the U.S. South owned slaves than did Southern whites as a whole”.⁶⁵

The problem, Anderson claims, is not that Farrakhan’s book contains falsehoods. It is that it does not put the true facts that it discusses into a wider context that would allow the reader to assess the *significance* of those facts. Taken by themselves, these facts give the impression that Jewish people played a particularly special and large role in the Atlantic slave system, larger than other ethnic groups. The larger context, however, shows this impression (though not the facts themselves) to be false. As Anderson notes, “The share of the Jewish investment in the Dutch West India Company was small, and the Dutch played a significant role in the Atlantic slave trade only in the seventeenth century, when the trade was small ... [Further, a] greater proportion of U.S. Southern Jews owned slaves than other Southern whites only because they were concentrated in urban areas, where rates of slave ownership were higher”.⁶⁶

Anderson concludes that *The Secret Relationship* constitutes an epistemically biased account of the Jewish role in the slave trade in virtue of omitting significant facts that are crucial to our understanding of that topic. In this way, Anderson adds to a large literature suggesting that we need to look beyond mere truth when judging the adequacy of our language.⁶⁷ In so doing, Anderson suggests a non-instrumental critique of *The Secret Relationship* that does not rely on assessing its truth-conditions. One of the ways in which we can *directly* assess the epistemic standing of an account of a subject, in other words, is by assessing whether it includes all pertinent and significant facts.

What can this tell us about linguistic salience perspectives? Speculatively, we can suggest that making the wrong content salient might also be *constitutive* of epistemically biased language. If it is, then the linguistic salience perspective employed by the immigration article might be criticisable on non-instrumental grounds because of its emphasis on the wrong content—in this case, the wrong facts.

One thing that the discussion of Anderson helps to clarify in this chapter is the concept of salience behind this thesis. I am not concerned, in this thesis, with sins of omission. Salience, at least in the sense that interests me, is not about making an issue prominent in a way that involves omitting or otherwise ignoring pertinent content.

⁶⁵ Anderson (1995: 38).

⁶⁶ Ibid.

⁶⁷ See the references mentioned in fn. 55.

(Whilst we do sometimes use the word *salience* to refer to such a scenario, this is not an everyday usage of the term that is relevant to this thesis.)⁶⁸ The sort of salience that characterises salience perspectives is about including all pertinent contents but giving relatively more salience to the wrong content. The problems that I wish to highlight, then, look subtler than those on which Anderson focuses.

2.5. How to choose linguistic salience perspectives

Let's return to the main point. I have suggested that we examine prevalent cultural beliefs, associations and ideologies in our folk-cultural context to determine how a linguistic salience perspective is likely to function. In particular, I am interested in whether a given linguistic salience perspective, even if not false in itself, is likely to activate *false* (and/or harmful) beliefs, or *unwarranted* (and/or harmful) associations and ideologies. What can this tell us about which linguistic salience perspectives we *should* be using?

To answer this, I borrow from Jennifer Saul and Louise Antony's discussion of implicit bias. I will spend the next few pages exploring this discussion. This detour will, I hope, prove fruitful, as it will put me in better stead to answer my question.

An implicit bias is an unconscious belief, association or attitude that affects our judgements and actions. For example, whilst assessors working in a university institution think that they are not racist, research shows that they nevertheless rate a CV as less impressive if it has a stereotypically black name at the top of it, than an equivalent CV (i.e. with the same qualifications) with a stereotypically white name at the top of it.⁶⁹ Another example might be a company manager who, whilst they explicitly hold gender egalitarian beliefs, they nevertheless unconsciously associate women with the domestic setting. This leads them to avoid hiring women for professional jobs, and to trust advice and feedback from female colleagues less.⁷⁰ That these biases can *harm*, is hardly

⁶⁸ We might, for instance, talk of a newspaper making the potential benefits of Brexit salient in virtue of neglecting to mention the potential costs. Indeed, this notion of salience seems to be behind what communications researchers refer to as *issue salience* (See, for instance, Entman, 1993, and the general topic of *issue frames*, mentioned in fn. 5.)

⁶⁹ Saul (2012: 244).

⁷⁰ Brownstein (2017).

controversial. As these examples indicate, implicit biases can involve treating people unfairly.⁷¹

They can have epistemic costs, too. For instance, implicit biases can encourage us to make decisions on the basis of factors that are evidentially irrelevant to the matter at hand (e.g. a person's skin colour), which can quickly lead us to make mistakes (e.g. to rate a person's intelligence incorrectly, or to judge their credibility wrongly). Where we hire and promote individuals who are not as good as our implicit biases lead us to think, we end up with an organisation that is not as good as it could be.⁷² If our organisation publishes research, for instance, implicit biases can mean that the standard of research being published is lower than it would be without such biases influencing our hiring and promotions.

In her 2012 paper *Scepticism and Implicit Bias*, Jennifer Saul discusses which epistemic norms we should endorse in light of the fact that we are subject to these biases. She first considers the most intuitive suggestion, namely, to become *consciously aware* of the biases that are affecting our thought, and to try and overcome them by being more objective. Let's say that I am implicitly biased against people of colour, so that I consider people of colour to be less intelligent than white people. The suggested solution here is to make conscious efforts to ignore the colour of a person's skin—to really focus on assessing each individual in a neutral light.

This approach, unfortunately, does not tend to work. Ironically, trying to be more objective can work to give one a false confidence that one *has* become more objective, meaning that one is even less likely to correct for one's biased judgements.⁷³ Implicit biases simply work on a far too subconscious level for us to be able to correct for them through standard modes of reflective deliberation. This demonstrates, for Saul, the problems that we incur when we approach the issue of implicit bias with our *a priori* expectations of ourselves as neutral, rational agents—ones that can consciously *reason* their way out of a bias.⁷⁴

What *does* work, then? Many successful interventions involve flipping stereotypes on their head. For instance, spending time thinking about counter-stereotypical

⁷¹ Saul (2013: 246). Generally, there is a lot of evidence that implicit bias correlates with real-world behavior congruent with those biases (Greenwald et al., 2009).

⁷² Saul (2013: 246-50).

⁷³ Antony (2016: 159-60) makes this point. See Apfelbaum, Sommers, & Norton (2008) for evidence on this.

⁷⁴ Generally, many biases cannot be eliminated simply through making them explicit (Fischhoff, 1977).

exemplars is one effective way of suppressing the activation of implicit biases.⁷⁵ This includes, for instance, spending 5 minutes creating a mental image of a counter-stereotypic person (e.g. a strong woman, a black professor). Another effective intervention is repeating intentions such as “when I see a black face I will think ‘safe’”.⁷⁶ Saul suggests, then, that we correct for our implicit biases using “counter-intuitive mechanical techniques that draw not upon our rational agency but upon automatic and unconscious responses” [such as one’s unconscious response to one’s mental image of a counter-stereotypical person].⁷⁷ She continues “We can consciously enlist these unconscious responses, and use them to improve our epistemic responses”.⁷⁸ The epistemic norm becomes something like this: *follow those techniques that science has proved successful, whether or not they accord with our a priori intuitions about what counts as good epistemic practice.*⁷⁹

Louise Antony frames Saul’s suggestion as using a Quinean naturalised approach to the study of knowledge. This involves selecting epistemic norms by empirically investigating how humans *actually* think, qua the limited, situated creatures that we are. From this investigation, we then judge what is in fact conducive to epistemic success. In Antony’s words, this involves “[writing] a blank check for warrant”.⁸⁰ This naturalised approach, Antony suggests, demonstrates that “we should not strive to put aside all bias”.⁸¹ Whether a bias is good or bad is something to be judged on its consequences for our understanding. The effective interventions mentioned above – such as spending five minutes mentally picturing a counter-stereotypical person, and repeating intentions such as *when I see a black face, I will think ‘safe’* – count as ways of biasing one’s thought. They, in Antony’s words, “incline us to one judgement [or, we might add, association or ideology] rather than another”, for instance that black people are safe.⁸² These instances of bias, however, have been shown to help us, by counteracting our existing *bad biases*, to form a

⁷⁵ Blair (2002), Kang and Banaji (2006).

⁷⁶ Stewart & Payne (2008).

⁷⁷ Saul (2012: 260).

⁷⁸ Ibid.

⁷⁹ The type of interventions considered here are individualistic. Saul (ibid. 260-1) does suggest that individual corrective measures have limitations, and that we should be looking to make wider societal changes if we want to fully combat our implicit biases. I return to this point in the next chapter.

⁸⁰ Antony (2016: 161).

⁸¹ Ibid. Others have made similar suggestions. See for instance Richardson’s (2013) account of introducing what she calls *productive partialities* (in other words, *good bias*) into scientific method.

⁸² Antony (2016: 162).

more accurate impression of the subject in question. A bad bias is that which, in Antony's words, "incline[s] us in the wrong direction: away from the truth".⁸³ A *good bias* is that which inclines us *towards* truth. Given my earlier discussion of epistemic criteria beyond truth and falsity, we might add that bad biases incline us towards a less *accurate* or *warranted* picture of our subject, whilst good biases do the opposite.⁸⁴ We might also add a moral dimension to this: bad biases incline us towards harm, whilst good biases incline us away from harm.

What can these suggestions offer to my discussion of linguistic salience perspectives? Linguistic salience perspectives have an implicit nature. They subtly play with the presentation of linguistic content in a way that most often bypasses our conscious awareness. When these presentational shifts activate existing biases of ours, this generally occurs without our noticing. As with implicit bias generally, then, linguistic salience perspectives might be resistant to having their effects counteracted through standard modes of conscious reflection and deliberation. In fact, research does suggest that consciously reflecting on being subject to framing effects, and consciously trying to overcome them, is not a successful strategy. As I mentioned earlier, telling philosophers about framing effects does little to actually minimise the impact of framing effects in the trolley problem example discussed in §2.2.1, for instance. Generally, being given longer to reflect on how some linguistic content is presented has been found to be ineffective in reducing the impact of framing effects.⁸⁵

Mirroring the suggestions from Saul and Antony, then, I suggest that we adopt an epistemically naturalised methodology when deciding which salience perspectives to use in our communication. This involves two key steps. The first is to take stock of which (epistemically and morally) problematic biases are likely already influencing our thought, or take little to become activated in our minds. These include associations, beliefs and ideologies about a subject that are cognitively accessible and socially licensed. Secondly, we should empirically investigate which linguistic salience perspectives are

⁸³ Ibid.

⁸⁴ See Antony (ibid.) for a more detailed account of how to judge which biases are bad, and which are good. Antony (ibid. 184) also addresses the problem of how we can carry out these judgments "when we know that it is precisely such judgments that embody implicit bias". She suggests that one "mechanically [notes] all the "facts," and [makes] all inference principles as explicit as possible. Thus, for example, if we realize that we are likely to underestimate the productivity of women job applicants, we can force ourselves to explicitly count the number of publications by each applicant, and to relate them to precalibrated grades" (ibid.).

⁸⁵ LaBoeuf & Shafir (2003).

effective in neutralising, or preventing the activation of, these existing bad biases. Those that are effective may well involve using good bias.⁸⁶ (Given that I do not conduct my own empirical studies for this thesis, I instead consult existing, relevant empirical studies that can help me to judge which salience perspectives are likely to be effective.)

Just as saying *when I see a black face I will think 'safe'* is a way of *constructively* using bias in order to neutralise an existing bad bias, using a linguistic salience perspective that emphasises the opposite of what our existing bad biases make salient might also help to neutralise, or prevent the activation of, those bad biases. Indeed, what worked for Jamieson and Radick's genetics textbook was a *reversal* of the order in which they presented information about genetics. Whilst this arguably introduced a different bias into the textbook (it inclined its audience to make interactionist research most salient in their minds), this bias was one that resulted in the students having a more accurate (and less harmful) understanding of genetics. In particular, the students were less likely to endorse false (and harmful) claims about genetic causation. Adopting linguistic salience perspectives that *invert* the emphases of our bad biases looks to be a good first place to start, then.⁸⁷

Let's relate this back to the immigration article. If we have found in our survey that its linguistic salience perspective does indeed risk activating our (bad) xenophobic biases, then we can suggest reversing its pattern of salience. In particular, perhaps we might advise that we begin discussions of immigration by discussing its advantages, and mention potential costs later. As I will reiterate in future chapters, this advice is not tantamount to saying that no article (or book, or speech) is permitted to begin with a discussion of the threats and costs associated with immigration. It is simply that one will

⁸⁶ Using salience perspectives to prevent the activation of problematic biases in our audience's minds brings up issues connected to epistemic and ethical *paternalism*. In other words, we are using patterns of salience to shape our audience's inferential patterns for (what we deem is) the better, *without consulting them on the issue*. Whilst there are interesting things to be said here, I do not discuss this issue. For a defence of *epistemic* paternalism, the reader can turn to Ahlstrom-Vij (2013).

⁸⁷ The suggestion here, then, only concerns linguistic salience perspectives that are liable to lead to epistemically and morally problematic biases. When it comes to framing effects like the oral health care example in §2.3.2, there is no obvious single option that is clearly wrong (changing the framing of flossing information from a loss to a gain does not mean the difference between activating bad biases and not, arguably). Thinking about which linguistic salience perspective to use can still be important for prudential and/or practical reasons, though; if more East Asian people will form better oral healthcare intentions when information about flossing is framed as a loss instead of a gain, then this information may well rightly influence a government's choice of linguistic salience perspective.

need to justify choosing this linguistic salience perspective, given the potential *cost* of employing it (i.e. the cost of activating false and harmful xenophobic biases). In other words, the fact that a linguistic salience perspective can activate a (bad) bias of ours gives us *one* reason to avoid that perspective. This consideration must, of course, be weighed against the many other considerations that shape our overall decision-making process about how to communicate with others.

2.6. Conclusion

In this chapter, I have introduced the idea of a linguistic salience perspective, which refers to the structuring and presenting of linguistic content, so that certain parts of it are made relatively more salient than others. Citing research into framing effects, I showed that simply shifting the pattern of salience in linguistic content, without making any other changes to that content, can nevertheless result in significantly different cognitive and behavioural responses to that content.

In the introduction, I asked four questions. Firstly, what does it mean for something to be more salient in language? We might make something more salient in language, in the sense that interests me, by playing with the structure and presentation of some linguistic content in such a way that many would dismiss as inconsequential. This might involve mentioning it *before* some other content. Alternatively, where we have a choice between different logically-equivalent presentations of some content, whichever presentation we choose is that which we make salient. For instance, by presenting the fat/lean content of some beef as *25% fat*, as opposed to *75% lean*, we make its fatness more salient than its leanness.

Secondly, I asked how is it that a mere shift of salience in language can have the sorts of dramatic consequences identified by framing effects research. My explanation focussed on culture. Where there is a cognitively accessible and socially licensed bias in our culture, I argued that simply making salient content that is central to that bias is liable to activate it in the audience's mind.

Thirdly, I asked whether and how we could evaluate mere patterns of salience in language. To this end, I focused on offering an instrumental critique of linguistic salience perspectives; where a linguistic salience perspective has bad effects, such as through activating a problematic bias, it can be judged instrumentally problematic. I briefly contrasted this instrumental critique with some non-instrumental ways of critiquing

salience perspectives. Where a salience perspective *normatively licenses* a bias, or where it counts as an epistemic flaw, it can be judged on non-instrumental grounds.

Fourthly, I asked how we *should* use salience patterns in language. In answering this, I suggested following the lessons of naturalised epistemology. More specifically, I suggested first assessing which problematic cultural biases are particularly cognitively accessible, and thus take little to activate. On the basis of this assessment, I suggested using linguistic salience perspectives that make the activation of these biases less likely. Often, these counter-salience perspectives will involve making salient the opposite of that which our biases make salient. By helping to avoid the epistemic and moral problems encompassed in those biases, these counter-salience perspectives will help the audience in question to develop more accurate, and less harmful, conceptions of the subject matter in question.

We should pay more attention to how we use salience in language. All linguistic content must be presented in some way or other. Whilst we might pass off as trivial (or fail to even notice) the choices that we make when structure our communication along dimensions of salience, these choices might be shaping audience responses in important ways, with epistemic and moral implications.

CHAPTER 3

SALIENCE AND COGNITION

3.1. Introduction

Imagine a person, let's call them Charlotte. Charlotte attends to people differently depending on their weight. When a person is slender, it tends to be that person's personality that Charlotte finds to be their most salient feature, in the sense that their personality catches her attention and stays in her memory, more so than their figure. When a person is fat,¹ on the other hand, it is their size that is usually most salient to Charlotte. She does notice their personality, and she can remember certain of their character quirks when she later reflects on them, but she finds these things less striking, less memorable, than their weight.

The guiding questions for this chapter are as follows. Firstly, what does it mean for something to be more salient in a person's cognition? Secondly, can we evaluate a person's cognitive salience patterns, and if so, how? And, thirdly, if we do indeed judge that we are making the wrong thing salient, how can we change the salience patterns in our mind?

In what follows, I will introduce the notion of a *cognitive salience perspective* in §3.2, building from Elisabeth Camp and Sebastian Watzl's work. Cognitive salience perspectives are constituted by dispositions to *attend* more to certain mental states in the mind over others (these states might include beliefs, desires, perceptions of things in the world, and so on). It is this that separates them from the linguistic salience perspectives discussed in the previous chapter, which were instead constituted by ways of manipulating the structuring or presentation of linguistic content, such as by shifting the order of that content. Whether we can evaluate a person's private attentional dispositions will be addressed in §3.3. In particular, I consider the instrumental and non-instrumental critiques of linguistic salience perspectives discussed in the previous chapter, and discuss

¹ Here, I follow Anne Eaton (2016) in using the word *fat* as a value-neutral descriptive term. As Eaton points out, this parallels its usage in Fat Studies and the Fat Pride Community. She says: "The basic idea is to avoid seemingly well-intentioned euphemisms like "saftig" or "heavy" that depend on the tacit understanding that "fat" is an impolite term of derision, and also to avoid euphemisms like "overweight" and "obese" that medicalize fat as a disease." (ibid. 39).

how these might apply to cognitive salience perspectives. Finally, I consider the issue of how to *change* our cognitive salience perspectives in §3.4. In addition to individualist approaches modelled on habit formation, I also consider structural approaches, which place the burden of change on society.

3.2. Introducing Cognitive Salience Perspectives

3.2.1. Some examples

A. Music

Chun is listening to a band. She finds herself focussing on one instrument over the others. That synth is really grabbing her attention! She hears the drums and guitars, but they've melted into the background of her experience.² A week later, a friend asks Chun about the gig. The synth comes to her mind straightaway, and she mentions how great it sounded. Her friend knows the guitarist, and so asks about their performance. Chun has to try a little harder to remember the guitar; trying to jog her memory by imagining herself back in the room, the sound comes back to her, and she reports back to her friend.

B. Personal Outlook

Imran is a relatively anxious person. When he meets a new person, he focuses on any awkward moments that arise, such as when he stutters or says something embarrassing. When he reflects on his past actions more generally, the moments where he has slipped up loom largest in his mind. He does at least remember his achievements, but they don't take centre stage in his mind like his failures do.

C. Intellectual Theories

Susan and Arif come to agree on the facts about the role of genetics in trait development. They differ, though, in which fact jumps out at them. When reading academic articles about genetics, what stands out to Susan are discussions of how one and the same gene can contribute to near-opposing traits thanks to shifts in its environment. These are the articles that most easily pop into her mind when she reflects on the role of genes in trait development. For Arif, it's instead papers detailing how

² A version of this example is borrowed from Sebastian Watzl (2017: 74), whose work echoes some of the ideas in this chapter.

particular genes reliably correlate with particular traits when the environment is held constant. Arif also believes what Susan believes about how genes and environment interact in sometimes surprising ways, but he has to try a little harder to remember the relevant facts. Susan similarly must push her memory slightly more when it comes to the cases that are most salient to Arif.³

3.2.2. Cognitive salience perspectives

A cognitive salience perspective is constituted by attentional dispositions to find various mental states more salient than others. For instance, one might better attend to one belief that one has about the topic of biology over another. Alternatively, one might better attend to (one's experience of) the synth in a band over the sound of the guitar. The idea of a cognitive salience perspective finds echoes in Elisabeth Camp's account of what she calls *perspectives*. Whilst Campian perspectives are broader than what I call cognitive salience perspectives, they are partially constituted by the *structuring* of contents along dimensions of salience.⁴ For Camp, something is salient in our attention when "it sticks out relative to the background, like a bright light or a hugely bulbous nose".⁵ Sebastian Watzl's work on attention also picks up on this idea of signal-to-noise ratio. Salience in attention, he suggests, is about how that attention is *structured* so that "some things are in the foreground relative to others".⁶

As with linguistic salience perspectives, cognitive salience perspectives are individuated not in terms of content (in this case, *cognitive* content), but rather in terms of the *structuring* of that content. Again, we see this shift of focus from content to structure made explicit in Camp and Watzl's accounts. For Watzl, the study of attention means recognising that the "relations between the parts of our conscious mental life [i.e. the foregrounding and backgrounding of mental states] are as important as the intrinsic features [i.e. the *content*] of various mental states".⁷ Similarly, what interests Camp (as we

³ This example builds upon an experience of mine that I discussed in chapter 1 (§1.1).

⁴ Campian perspectives go beyond mere salience. For instance, they include dispositions to find certain properties more causal and explanatory (in her words, more "central") than others (Camp, 2017: 80). (See §3.2.2.1 below.) I use the phrase *salience perspective* to highlight how my account of perspectives is narrower than Camp's.

⁵ Ibid. 80.

⁶ Watzl (2011: 849). Watzl (2017: Ch. 4) offers what he calls *the priority structure view* of attention, which involves the activity of structuring mental states so that some are prioritized over others. He compares our mind to a newspaper, where certain stories make the headline, and others are granted less priority (ibid. 70).

⁷ Watzl (2011: 849).

heard in chapter 1) is not “*what* is represented in imagination ... [but rather] ‘*how*’ that content is imagined”.⁸

Further, as with the linguistic salience perspective of the previous chapter, cognitive salience perspectives are about giving *relative* salience to one thing over another (this time, in one’s attention). Let’s return to Charlotte. Whilst Charlotte might hold various beliefs about a subject, such as that her acquaintance Jane is fat, and that she does a lot of charity work, it is a further question which of these beliefs is given relative salience in her mind over the other. Let’s say that she gives relative salience to Jane’s fatness. When I implore Charlotte to change her salience perspective on Jane so that she attends more to her charity work, then, I am not claiming that Charlotte is *unaware* of this trait of Jane’s, or that she is *ignoring* it.⁹ Instead, I am imploring Charlotte to give it *relative salience* over her weight.

This discussion helps to make explicit one way, then, in which cognitive salience perspectives relate to linguistic salience perspectives. They are connected by the relation of similarity, insofar as the notion of relative salience applies to them both. They are also causally connected. As was mentioned in the previous chapter, making something relatively more salient in language tends to make that thing relatively more salient in an audience’s minds. If Charlotte makes Jane’s weight particularly salient when talking to Bill, such as by mentioning her weight *before* her other attributes, then this will tend to result in Bill giving relative salience to Jane’s weight in his *mind*. Conversely, when something is salient in one’s mind, this tends to result in one making that thing salient in language. Charlotte’s cognitive salience perspective on Jane will indeed make it likely that, when talking to Bill and others, she mentions Jane’s weight before her other attributes.

3.2.2.1. Dispositional Nature:

What exactly does having something relatively more salient in one’s mind amount to? When we make something salient in our minds, we are disposed better to *notice* it, *remember* it, and find it more *cognitively accessible*. As for Charlotte’s acquaintance Jane, she

⁸ Camp (2017: 77). Camp is primarily interested in how perspectives relate to our imaginative engagement with fiction, hence the emphasis on *imagination* in this quote. She also thinks that they capture how we engage with (subjects in) the real world, too, however.

⁹ In this last chapter, (§2.4.2) I clarified this notion of *relative salience* by contrasting it with another everyday usage of *salience* that I stipulated is not relevant to this thesis: namely, making *x* (e.g. Jane’s fatness) salient in a way that involves *ignoring y* (e.g. Jane’s charity work).

might better *notice* instances of others discussing Jane’s fatness than instances of others discussing her charity efforts. Because she has noticed these instances, she better *remembers* them.¹⁰ This means that Jane’s fatness is more cognitively accessible to Charlotte; it takes less cognitive labour to think of than her charitable side.¹¹

Even if we do not change our propositional attitudes (such as beliefs or desires) regarding a subject, a shift in cognitive salience perspective can affect our explanations of and behavioural responses to that subject. When Jane’s fatness is particularly salient to Charlotte, it is what pops into her mind first when asked to explain why Jane acts as she does, such as why she rarely attends social events. This can mean Charlotte is more likely to invoke Jane’s weight than her charity commitments in her explanations of Jane’s behaviour. Indeed, psychologists Larisa Hussak and Andrei Cimpian argue that where x is more retrievable in our memory than y , the accessibility of x can indirectly affect the content of our explanations. In their words, “[certain] entities are in the focus of attention when the search for an explanation is triggered, so they are likely to serve as the first, and perhaps only, retrieval cues”.¹² This “attentional spotlight”, they say, can simultaneously prevent other relevant contents from coming to mind, which lowers the probability of citing these other contents in explanations.¹³ This is one way in which cognitive salience perspectives have epistemic significance: by indirectly influencing our explanations.

The fact that a salience perspective can impact how one explains information about a subject does not mean that salience perspectives are smuggling in implicit beliefs about how explanatory, or specifically how *causal*, a subject’s properties are. When a property is more cognitively accessible, it is simply more *likely* to get invoked in explanations.¹⁴ It is here that we can see one of the more significant ways in which my

¹⁰ See Chaffee & Schleuder (1986) for a study confirming this link between attention and memory.

¹¹ The notion of cognitive accessibility was discussed in §3.1.1 of the previous chapter.

¹² Hussak and Cimpian (2017: 70).

¹³ Ibid. Siegel, who discusses the epistemic side of attentional dispositions, also confirms the epistemic significance of attention. For instance, she suggests that, “Which evidence one ends up with is a function of what one pays attention to” (Siegel, 2017: 159).

¹⁴ It should be noted that repeatedly retrieving certain properties instead of others when explaining a subject’s traits, can eventually lead to one *forming a belief* about that property’s explanatory, or specifically causal, nature; as we will hear in the next section, merely having something *salient* in one’s mind can have plenty of significant *effects*, such as the triggering of a truth-conditional belief about that property. Crucially, however, having a property cognitively accessible does not *in itself* need to involve a belief about that property’s explanatory, or specifically causal, nature.

account of salience perspectives departs from Camp's notion of perspectives. Camp talks about certain properties being more *central* than others in someone's mind, where they are considered to cause, or otherwise explain, many of the subject's other features.¹⁵ This gives the impression that there is a (perhaps implicit) belief about the property's causal profile—a belief that would clearly take us beyond the bounds of salience. (If this is not what Camp means, then she could afford to be clearer about her notion of *centrality*).

Similar things can be said about how cognitive salience perspectives relate to our behavioural responses to a subject. With Jane's weight most salient in her mind, Charlotte might be less likely to invite Jane on a charity run than if Jane's charity-giving had loomed larger in her mind. As with the point above, this does not demonstrate that cognitive salience perspectives are partially constituted by codes of interaction, such as was the case with *ideologies*, as discussed in the previous chapter. Instead, when Jane's weight is more cognitively accessible, it is simply more *likely* to influence Charlotte's decisions about how to act around Jane. What we attend to can indirectly influence our behaviour.

3.2.2.2. *Scope and duration*

Charlotte's salience perspective on Jane might pertain only to Jane. Or, as per our example in the introduction, Charlotte might apply this salience perspective to all fat people, not just Jane. These two cognitive salience perspectives represent different levels of *subject-specific* salience perspectives—where the specific subject shifts from a single person to a single social group. We can envisage a *generic* salience perspective, though, that is less confined. Perhaps Imran's cognitive salience perspective that induces him to focus on his failures is symptomatic of a broader cognitive salience perspective characterised by pessimism. This broader perspective might dispose him to better attend to negatives generally speaking, whether he is considering his own life, the economy, or the future of humankind.¹⁶ (This distinction is not confined to cognitive salience perspectives only; a linguistic salience perspective can also either be subject-specific or

¹⁵ Camp (2017: 80).

¹⁶ The distinction between generic and subject-specific cognitive salience perspectives is not entirely clear-cut. Consider a cognitive salience perspective that gives relative salience to cynical motivations over altruistic ones. Some might want to refer to this perspective as generic, insofar as one can better attend to cynicism in multiple subjects (politics, personal interactions, media, etc.). Others, however, might consider it to be subject-specific, insofar as it pertains to the specific subject of human motivation.

generic. A generic linguistic salience perspective might reliably mention negatives *before* positives, for instance, whatever the subject under consideration.)

The duration with which we hold cognitive salience perspectives can also vary. Charlotte's cognitive salience perspective on Jane might involve her reliably making Jane's fatness her most salient feature. But, if, whilst they are out together, Jane spills her drink on herself, her clumsiness might be given relative salience over her weight in Charlotte's mind, at least for a few minutes. All sorts of things might spark a temporary shift in a person's cognitive salience perspective, such as desires. If I am thirsty and desire hydration, the beverages that I pass in shop windows will have relative salience for me over the foodstuffs. This cognitive salience perspective will likely dissipate upon fulfilling my desire. Further, reading an article with a particular *linguistic* salience perspective might succeed in temporarily changing a person's cognitive salience patterns (linguistic salience perspectives, remember, tend to produce their counterpart cognitive salience perspectives in their audience's minds). I might read an article targeted at food packaging businesses that gives relative salience to the harms that our everyday plastic consumption generates over the benefits and conveniences that plastic food packaging tends to give. As a result, I start giving these harms relative salience in my attention. This effect might only be short lived; a day after reading the article, I, like many others, continue finding the conveniences of using plastics more noticeable, memorable, and cognitively accessible than the harms.

In what follows, I am interested primarily in cognitive salience perspectives that have a longer duration. More specifically, I am interested in those that have become habitual—where one reliably gives relative salience to one mental content over another.

3.2.2.3. *Automatic and intuitive nature*

Let's say that I believe that I should find the harms associated with plastics more salient than the conveniences. Does my holding this belief mean that I have successfully adopted the relevant cognitive salience perspective? To answer this, consider the duck/rabbit illusion in Figure 1 below. I might look and *know* that the lines in this image can represent a duck without successfully seeing the image *as* (depicting) a duck. It just looks like a rabbit to me. I am aware that it *can* be seen as a duck, but, try as I might, my

imagination simply doesn't arrange the image in the right way. Wittgenstein's distinction between *looking plus thinking* and *seeing-as* helps to explain this phenomenon.¹⁷



Figure 1: Duck / Rabbit Illusion

Elisabeth Camp uses examples such as these to suggest that an analogous phenomenon can occur in our attention. I might believe the proposition *the harms of plastics should have relative salience over their conveniences* – that I should find the harms more noticeable and memorable – but I can't help finding their conveniences more attention-grabbing. Or, like Imran from §2.1, I might, after speaking to a counsellor, believe the proposition *my achievements should be given relative salience over my failures*, but nevertheless persist in automatically attending more to my worries and flaws. My achievements simply don't intuitively leap out as salient in my own mind. As Elisabeth Camp phrases it, I don't get that "click".¹⁸ This demonstrates the non-doxastic nature of cognitive salience perspectives; one does not hold a particular cognitive salience perspective in virtue of believing a particular proposition.¹⁹ The non-propositional nature of cognitive salience perspectives also demonstrates that they lack truth-conditions.²⁰

¹⁷ Wittgenstein, in Camp (2017: 82). This argument – that salience perspectives are non-propositional and intuitive – is borrowed from Camp.

¹⁸ Camp (2017: 83). Sebastian Watzl (2017: 49) refers to attention that is unintentional – that is not controlled by one's will – as *arational*.

¹⁹ Indeed, Camp distinguishes her *perspectives* from belief (2017: 82), and Watzl (2011) distinguishes his account of *attention* from belief. One might wonder whether cognitive salience perspectives are instead equivalent to what Tamar Gendler (2008) calls *aliefs*. Indeed, aliefs are also distinguished from beliefs partly in virtue of their intuitive and automatic nature, as well as their tendency to come apart from one's beliefs. For instance, one might believe that the characters in a film are fictional, but *alieve* that they are real, in the sense that one feels emotions about them, and one is disposed to cry about the bad things that happen to them. It is precisely these affective and behavioural dimensions to alief, however, that distinguish them from cognitive salience perspectives.

Familiar examples of habit formation is one area where many of us experience a disconnect between our intellectual endorsement of a cognitive salience perspective and an ability to *exercise* it in an automatic and intuitive way. When we have developed a habit of focussing on particular contents – such as worries and flaws – it can take a lot of effort and practice to shift to focussing on others—such as opportunities and strengths. Those who have experience with *Cognitive Behavioural Therapy* can attest to how tricky this process of changing one’s automatic, intuitive attentional patterns can be;²¹ the challenge can be so hard that some of us never succeed in cultivating new attentional dispositions. Simply considering (or even endorsing) a salience perspective, then, is not the same as *adopting* a salience perspective (i.e. successfully cultivating the relevant attentional dispositions). I will consider the significance of this point in §4.

Further, as is also a common experience in any sort of therapy (whether formal or informal),²² our existing habits of attention are not always transparent to us. It can often take another person, or a transformative personal experience, to make clear exactly what it is we seem to be noticing, remembering, and finding cognitively accessible. It is plausible that Charlotte is unaware, therefore, that she is reliably making fat people’s weight their most salient feature. In this way, cognitive salience perspectives often bypass our conscious reflection. This does not mean, however, that it is impossible to investigate such implicit cognitive salience perspectives. In addition to techniques used in therapy,²³ other ways exist of ascertaining a person’s implicit attentional dispositions. For instance, eye tracking studies help to illuminate what is catching a person’s attention, and

Cognitive salience perspectives are more minimal. Instead of encompassing these affective and behavioural dimensions, they simply involve the structuring of our attention so that we find certain content more noticeable, memorable, and cognitively accessible, than other content.

²⁰ Camp (2017: 79) agrees that a [cognitive salience] perspective “determines no truth-conditions of its own”.

²¹ Part of Cognitive Behavioural Therapy (CBT) is about noticing and then changing what one focuses on and remembers. Padesky (1994: 268), for instance, talks about CBT changing “what we notice, attend to, and remember of our experiences”. He suggests, for instance, that someone with what he calls a *negative self-schema* “will focus on personal defects, flaws, and errors, noticing and remembering these more than strengths, positive gains, and successes” (ibid.). This way of talking is strikingly similar to the language that I use in this chapter.

²² Formal therapy might involve counselling, whilst informal therapy might involve chatting with a friend, or reading a lifestyle magazine.

²³ Cognitive behavioural therapy, for instance, has developed techniques of making a person conscious of their unconscious habits, including habits of attention. Many of these involve worksheets that get an individual to record their thoughts and feelings (Padesky, 1994).

therefore what is likely getting logged in their memory. More academic studies are making use of eye tracking technology, using it to ascertain, for instance, which parts of a person's body someone is paying more attention to,²⁴ or which sentences in a story one's eyes linger on.²⁵ Memory tests can also help to reveal which implicit cognitive salience perspective one has.²⁶ Generally, as is the case with cognition that is unconsciously processed, implicit measures (i.e. those that do not rely on an individual to report their attitudes and habits, but instead attempt directly to measure a person's unconscious responses) may be better at uncovering cognitive salience perspectives that we are not conscious of having.²⁷

3.3. Ways of evaluating cognitive salience perspectives

Can we evaluate a person's private attentional patterns as good or bad? For some, it seems that we cannot. As Susanna Siegel reflects, "Some patterns of attention seem palpably insignificant. Strolling in a meadow, one can [give relative salience to] the sky or the trees [in one's attention]."²⁸ The example of Chun listening to music might also fit this description; unless there are extenuating circumstances (such as it being part of an exam designed to test her ability to evaluate the drums, or that she is at the gig specifically to support her drummer friend perhaps), it does not seem to matter epistemically or morally whether Chun pays more attention to the synth or the drums. For other attentional dispositions, though, I think that we can more easily make sense of the idea one is treating the wrong things as salient in their attention. How?

3.3.1. Instrumental critique

Siegel suggests that we are already used to evaluating attention according to its upshots. She cites a legal decision to condemn a driver's attentional dispositions where they involve not noticing a mule that had walked onto the highway, and where this lack of

²⁴ See Karsay et al. (2017), and Gervais et al. (2013).

²⁵ Frisson & Wakefield (2012).

²⁶ See, for instance, Gaither et al.'s (2014) study into whether children better remember faces racially ambiguous or racially non-ambiguous faces.

²⁷ See, for instance, the measures developed to detect our implicit biases, such as the *Implicit Association Test*. There is a great range of implicit measures that have been developed in psychology (Rudman, 2011).

²⁸ Siegel (2017: 159).

attention led to a car accident.²⁹ This represents an instrumental critique of attention; the demerits of the driver's attention were derivative on the demerits of its effects. The form of attention that Siegel considers however, does not represent a cognitive salience perspective; the driver is not giving relatively less salience to the mule (than to something else), rather she is not attending to the mule at all. (This is what I have called a sin of *omission*.)³⁰ How can we evaluate cognitive salience perspectives in particular?

In the previous chapter, I suggested judging as instrumentally problematic linguistic salience perspectives that activate inferences to epistemically and/or morally problematic biases of ours. As we have already mentioned, linguistic salience perspectives function by first inculcating their counterpart cognitive salience perspective in their audience's mind. This means that it is actually *cognitive* salience perspectives that are more directly responsible for activating the biases discussed in the previous chapter. Following the logic in that chapter, then, we can judge as instrumentally problematic *cognitive* salience perspectives that activate inferences to epistemically and/or morally problematic associations, beliefs, or ideologies.³¹

This goes for cognitive salience perspectives whether they have been inspired by a linguistic salience perspective or otherwise. For instance, whether or not a linguistic salience perspective has inspired it, we might criticise Charlotte's cognitive salience perspective on fat people if it is activating oppressive beliefs and ideologies about fat people. These might include, for instance, beliefs that fat people are lazy and unintelligent, associations between fatness and greed, feelings of disgust, and codes of interaction that involve avoiding sitting next to a fat person on public transport.³² It has been argued that what might be called *fatist* beliefs and ideologies are cognitively accessible and socially licensed.³³ If they are, according to the conclusions reached in the previous chapter, Charlotte, simply through reliably making fat people's weight more salient than their other features, might activate those beliefs and ideologies in her mind. Where we want to condemn these beliefs and ideologies as false, unwarranted, and/or harmful, we can take issue with the attentional dispositions that led to them.

²⁹ Ibid. (160).

³⁰ Chapter 2 (§2.4.2).

³¹ Indeed, Watzl (forthcoming) suggests that we can evaluate attention (specifically attention that takes on the form of a cognitive salience perspective, i.e. where one gives *relative* attention to one thing over another) on instrumental grounds—according to its results.

³² Some of these features of fatist ideologies are described in Eaton's (2016) paper.

³³ See, for instance, Eaton (ibid.). Eaton cites various terms that refer to fat oppression, including *fatism* (ibid. 38).

Where Charlotte's cognitive salience perspective is activating these biases, it is likely to be doing so in a way that bypasses Charlotte's conscious awareness. I suggested in the previous chapter that patterns of salience activate biases in an *implicit* manner. Far from explicitly stating a belief, for instance, salience perspectives (whether linguistic or cognitive) rely on an implicit inferential process whereby we subconsciously pragmatically infer that belief. This, I suggested, can make patterns of salience especially *effective* in activating that belief in our mind, meaning that they are even more likely to influence our thought and behaviour in ways congruent with our accepting that belief. This is because their under-the-radar nature means that they do not give the individual chance consciously to reflect on whether they do indeed wish to endorse that belief that is activated in their minds.³⁴ This means that, whilst Charlotte might be disposed to *reject* an inference to a fatist belief if such a belief is explicitly articulated (if, say, an acquaintance says *Jane is fat and therefore unattractive*), she might be disposed to think and act in ways that demonstrate an endorsement of fatist beliefs if those beliefs are *implicitly* activated by her cognitive salience perspective.

As I mentioned in the previous chapter, it is this *instrumental* critique of salience perspectives that concerns me in the next two chapters. Below, however, I offer some contrasting *non-instrumental* ways of critiquing cognitive salience perspectives. Whilst I do not return to non-instrumental critiques until chapter 6, discussing them here will be helpful. My aim in these first two (substantive) chapters is to give a broad overview of salience perspectives. This allows the reader to have all the moving parts relating to salience perspectives in place before considering how they function in practice.

3.3.2. Non-instrumental critique

I briefly discussed in the previous chapter two ways in which a linguistic salience perspective might constitute, as opposed to cause, epistemic bias, and harm. One involved the normative concept of *licensing*. Licensing is a normative, as opposed to causal notion, meaning that where a cognitive salience perspective *licenses* problematic beliefs or ideologies, it bears a *constitutive* relationship to their problems. If we can show that Charlotte's cognitive salience perspective licenses fatist beliefs and ideologies, then, we can say that her perspective constitutes their epistemic and moral flaws. We can briefly sketch how this argument might look. In the previous chapter, we suggested that use of a given salience perspective can *license* beliefs and ideologies associated with that

³⁴ See Chapter 2 (§2.2.1 and §2.3.1).

perspective.³⁵ Applied to cognitive salience perspectives, we can say that, when Charlotte employs her cognitive salience perspective, she implies that *this way of attending to fat people is acceptable* (in other words, she licenses her cognitive salience perspective). There is an ideology associated with that way of attending to fat people, namely, fatism. Fatism also makes fat people's weight their most salient feature; amongst other things, fatism suggests that fat people's weight *should* be salient because fatness is a grave social problem that requires fixing, and that it is important to shame and stigmatise fat individuals.³⁶ Following the logic of the discussion in the previous chapter, for Charlotte's way of attending to fat people to be acceptable, the ideology associated with that way of attending to fat people would also have to be acceptable. After all, it is this ideology that explains *why* it is appropriate to make fat people's weight their most salient feature (fatism suggests that fatness *is* and *should be* attention-grabbing, for the reasons above). In this way, simply by employing her cognitive salience perspective on fat people, Charlotte might be *licensing* fatist ideologies. Where these ideologies are intrinsically problematic, this would mean that, given its normative as opposed to causal relationship to those ideologies, so too is Charlotte's cognitive salience perspective.

I also introduced a different way in which certain patterns of salience in language might be problematic on non-instrumental grounds in the previous chapter, namely, where they themselves count as a way of being biased. Making the wrong facts salient, I suggested, might *count* as an epistemic flaw, independently of any additional negative epistemic consequences this pattern of salience might have. We can extend this thought to patterns of salience in our attention. Indeed, Camp suggests something along these lines. For her, cognitive salience perspectives, whilst non-propositional and thus lacking in truth-conditions might be judged on the epistemic ground of warrant.³⁷ In particular, she says that a subject-specific cognitive salience perspective is (un)warranted where its assignments of salience are (in)consistent with the objective distribution of properties in the world. In her words: "features [i.e. properties in the world] that are in fact more intense [i.e. that stick out relative to the background] should be assigned higher prominence".³⁸ So, blonde hair in a room full of brunettes might warrant being made

³⁵ Chapter 3 (§3.3). This argument was inspired by Swanson (forthcoming: §1).

³⁶ Eaton (2016).

³⁷ This adds to the growing literature in epistemology highlighting the myriad ways *beyond falsehood* that our cognition can be critiqued. See fn. 55 in the previous chapter.

³⁸ Camp (2017: 83). Given that Camp's notion of a *perspective* is broader than mine, she also details other criteria for judging perspectives, which are not relevant to my more minimalist account (*ibid.*).

salient because of what Camp refers to as the “broad statistical [distribution] of properties” in that room.³⁹ Camp does suggest that other considerations mean that multiple different cognitive salience perspectives will usually be warranted for a given subject.⁴⁰ For instance, our human interests and aims (which themselves, she suggests, can be driven by practical, moral, and/or aesthetic values) can legitimately shape what we decide to make salient. For instance, I might be studying the science of Brunette hair, and this interest of mine justifies my giving relative salience to the Brunettes in the room. Further, the fact that the statistical structures in the world are complex and multi-dimensional, mean that we are often justified in choosing different salience perspectives for one and the same subject. Considering why Britain voted to leave the EU, for instance, the same data legitimates making both the influence of key figureheads salient (such as Boris Johnson), as well as the social characteristics of individual voters (i.e. their occupational status and educational backgrounds). Different levels – from individual voter, to individual politicians, to political parties – play a causal role in political decision making. These qualifications notwithstanding, Camp suggests that a cognitive salience perspective can itself count as unwarranted if it makes salient the wrong feature, meaning that she offers a non-instrumental way of critiquing cognitive salience perspectives.

Watzl and Siegel adopt a similar, if slightly different strategy. For them, attentional dispositions might *partially constitute* an epistemic bias. Their example is as follows: Sara is evaluating CVs in order to hire a person. She systematically attends more to the good parts of CVs coming from White individuals, whilst systematically attending more to the bad parts of CVs coming from BME (black and minority ethnic) individuals. Sara has no evidence that BME individuals, on average, have more negative traits that make them ill-suited to the job than White people. Against the evidence, then, Sara takes BME individuals to be generally less qualified. We can say, then, that Sara has an irrational outlook. The suggestion that Watzl and Siegel make is that a tendency to have certain attentional dispositions might be *constitutive* of this irrational outlook. In other words, it is part of what it is to harbour this irrational outlook to, for no principled reason, pay more attention to the negatives of one social group than the negatives of another social group. If the irrational outlook is non-instrumentally problematic, they suggest, then so too is the cognitive salience perspective that partially constitutes it.⁴¹

³⁹ Ibid.

⁴⁰ Ibid. (83-4).

⁴¹ Watzl (forthcoming). Siegel did not co-author this paper, but Watzl cites their joint research.

Camp, Watzl, and Siegel have identified interesting ways of epistemically evaluating attentional dispositions. In Watzl and Siegel's case, Sara's attentional dispositions also look to be problematic for moral reasons.⁴² Whilst it might be *irrational* to pay more attention to the negatives of one group than the negatives of another, it also seems *immoral*. Sara is treating people of BME individuals *unfairly*, by presuming them to have more negative traits than White people. We might say that she has a morally prejudiced outlook, and that her cognitive salience perspective partially constitutes that morally prejudiced outlook. Perhaps part of what it *is* to treat people unfairly is to attend to one social group's negatives more than the negatives of another social group.

Indeed, when it comes to evaluating Charlotte's salience perspective, it is this *moral* dimension that seems better to capture what feels wrong about her attentional dispositions. At least to my intuition, it is not so much that Charlotte's way of attending to fat people falls short of some epistemic criteria (though it may well do this)⁴³ but rather that it falls short of some moral criteria. Inspired by Camp, Watzl, and Siegel's suggestions above, then, the proposal that I would like to make is this: cognitive salience perspectives can constitute *harm*. This may be surprising. Whilst a great many factors, from physical and psychological violence, to false beliefs and credibility deficits, have already been identified as potentially harming an individual or group, facts about *salience* have not seemed particularly relevant to harm. In what follows, I will sketch an argument that certain patterns of salience can be intrinsically harmful. I will expand on these ideas in Chapter 6.

Think back to Jane. It is not hard to imagine that Jane might not want Charlotte to attend to her such that her weight is her most salient feature. She might instead prefer for her charity work to be what Charlotte and others most notice and remember about her. How can we understand her desire? Jane's charity work is something that she has chosen, that she values, that she works hard at, and so on. In particular, her charity work demonstrates her *rationality*. Jane has decided that charity work is valuable, and is rationally setting and pursuing her own charity-related goals. It demonstrates her

⁴² Whilst Siegel and Watzl (Watzl, forthcoming) do consider how attentional dispositions might partially constitute a *moral* state, such as a *virtue* (such as gratitude), they do not expand on how a constitutive moral critique might be applied to the case of the prejudiced evaluator. (This might simply be because the focus in Watzl's paper is on *epistemic* evaluations of attention.) Further, they do not consider the broader point that I make below, namely, that certain attentional dispositions might count as a way of harming someone.

⁴³ Perhaps Charlotte's surplus of attention to fat people's weight, in contrast with her lack of attention to thin people's weight, partially constitutes an irrational outlook.

autonomy—her capacity to determine what she pursues in life. Jane has freely chosen to pursue charity work. It also showcases her *personality*. Through her charity work, Jane can demonstrate her distinctive, caring personality.

All of these traits have been associated with what is sometimes called one's *humanity* or *personhood* (I will henceforth use the latter term).⁴⁴ Personhood is a concept that plays a significant role in ethical thought; it is often cited as what gives humans a special moral status, or what represents, in Immanuel Kant's words, our "inner worth".⁴⁵ In fact, it is *disrespect* of one's personhood that is regularly cited as being what makes various phenomena harmful. Consider, for instance, sexual objectification, which involves treating a person as a thing in some way.⁴⁶ When explaining why sexual objectification can be harmful, feminists regularly cite the *disrespect* that it can constitute to the personhood of the one who is objectified.⁴⁷

Jane's fatness, on the other hand, does not demonstrate her personhood. It is a superficial appearance that does not demonstrate her rationality, autonomy, or personality, in the way that her charity work does. In fact, feminists researching objectification have often suggested that the body, whether fat or thin, is commonly understood as *lacking* personhood-related qualities. Sandra Bartky, for instance, says that, historically and culturally, the body has "been regarded as less intrinsically valuable, indeed, as less inherently human, than the mind or personality".⁴⁸ This has led to a focus in feminist research on women being objectified, in a way that disrespects their personhood, *specifically through a reduction of women to their body*.⁴⁹ (I return to, and expand upon, these feminist ideas in chapter 6.)

How does this discussion help with arriving at a non-instrumental evaluation of cognitive salience perspectives? Just as our earlier discussion mentioned a way in which certain cognitive salience perspectives might *constitute* epistemic bias, the suggestion here is that certain attentional dispositions might *constitute* harm, in virtue of *counting as a way* of disrespecting an individual's personhood. It is already a common thought in feminist writings that one can disrespect an individual's personhood by reducing them to their

⁴⁴ See, for instance, Kant (*Lectures on Ethics*), Dworkin (2000: 30-1) and Bartky (1990: 130).

⁴⁵ Kant (1785).

⁴⁶ Nussbaum (1995: 257).

⁴⁷ Bartky (1990), Dworkin (2000).

⁴⁸ Bartky (1990: 35).

⁴⁹ See, for instance, Langton (2005: 246-7).

body. The suggestion here is for a more *minimal* reading of this thought.⁵⁰ Simply giving *relative salience* to an individual's body over other identities of theirs that better reflect their personhood, can count as a way of disrespecting their personhood. Crucially, disrespect of an individual's personhood is treated as harmful in and of itself. This would mean that any cognitive salience perspectives that *count* as ways of disrespecting an individual's personhood would themselves also be non-instrumentally harmful. Charlotte's cognitive salience perspective on Jane, then, might count as a way of disrespecting Jane's personhood. This would mean that it constitutes an important harm in itself, independently of any additional harmful effects this cognitive salience perspective might have.

3.4. Changing cognitive salience perspectives

I have suggested above some ways to evaluate our cognitive salience perspectives. When we have identified a problematic cognitive salience perspective, though, how do we go about changing it? Things were relatively simple in the previous chapter. When we have identified a problematic *linguistic* salience perspective, the suggested solution was practicable. For instance, if one had identified a particular pattern of salience in language that was liable to activate a particular bias of ours, the advice might simply be to reorder linguistic content. If one was writing an article, for instance, one might change the order in which one discusses information, or change which information one headlines. Or, as per the *Genetics Pedagogies Project* discussed in the previous chapter, one could swap around the chronology of chapters in a textbook.

In the context of cognitive salience perspectives, however, things are a little trickier. Suppose we have identified an epistemically and/or morally problematic bias, and we have identified a cognitive salience perspective that we ought to avoid given that it likely plays into this bias. Consider Charlotte again. Let's suppose that our investigation has indeed identified an accessible and licensed false and harmful bias that treats fat people as lazy, disgusting, less intelligent, and so on. Further, suppose that we have proved that reliably attending to fat people so that their weight is their most salient feature risks activating that bias. How, precisely, is Charlotte supposed to shift her attentional dispositions? How is she supposed to change how she intuitively and

⁵⁰ Whilst the notion of *reduction* to one's body is a common formulation in feminist writings on this subject, we will hear some feminist precedents for more minimal readings, more like the one that I mention here, in chapter 6.

automatically attends to the world, so that fat people’s weight no longer catches her attention, and sticks in her memory?

As I suggested earlier, certain habits of attention can be particularly hard to change. Simply telling ourselves that something should (not) be salient does not necessarily result in those attentional dispositions *clicking* in our mind—in them becoming intuitive and automatic. Generally, we should remember that adopting a new perspective is not a question of endorsing it intellectually, but rather of that perspective playing a functional role in one’s actual attentional patterns. Charlotte, then, might intellectually endorse a cognitive salience perspective that mirrors how she tends to attend to slender people—namely, where she reliably makes a given individual’s personality more salient than their weight. Try as she might, though, she just can’t help but find fat people’s weight their most noticeable and memorable property. What should she do?

3.4.1. Individualist approach

Even when something has become so habitual that it is largely subconscious and automatic, we can often still have some control over it. In the context of our automatic and unconsciously processed implicit biases, for instance, Jules Holroyd suggests that, whilst we may not be able to control for our attitudes and behaviours in and of the moment, we nevertheless have what she calls “long-range” control, which refers to our control over which skills, habits, *and biases* we cultivate over time.⁵¹

There is a wealth of advice, as well as consumer products, out there to help someone to stop themselves from habitually reaching for a cigarette, or unconsciously biting their nails, for instance. Even an incredibly stubborn and long-term disposition to think pessimistically can eventually be changed through following things like self help books and Cognitive Behavioural Therapy courses, for instance. Most discussions of habit change make clear that there is no *quick fix*, and that a lot of effort and time is required to successfully break a habit.⁵² We heard some suggestions from Jennifer Saul in the previous chapter regarding changing out automatic, unconscious implicit biases. Saul’s suggestions required something beyond *telling* oneself to change one’s ways—it

⁵¹ Holroyd (2012: 284).

⁵² The effort and perseverance required in Cognitive Behavioural Therapy attests to this (Robertson, 2010: 66).

involved implicit tactics that engaged one's unconscious responses. It involved spending time imagining counter-stereotypic individuals, for instance.⁵³

Camp also suggests that the fact that our (cognitive salience) perspectives are habits that are, in her words, "only partly under voluntary control" means that we must be "trained into them" through methods that do not involve standard rational deliberation.⁵⁴ Specifically, she recommends that reading narrative might be a useful tool in training oneself to adopt alternative (cognitive salience) perspectives.⁵⁵ She discusses the openness that we bring to fiction reading. Partially due to our relative freedom from practical needs and desires when reading fictions, we intentionally allow our attention to be guided by the author, and are open to cultivating a wide range of intuitive associations and attentional dispositions in order to experience aesthetic enjoyment. Reading fiction to alter one's cognitive salience perspective, then, also looks to be a way of taking advantage of one's unconscious responses. Instead of reading a reasoned, logical argument for changing one's attentional dispositions, one instead relinquishes one's conscious control to one's intuitive responses to a fictional story. Interestingly, Sally Haslanger makes similar suggestions when it comes to changing a person's intuitive perceptual, thought, and behavioural patterns more generally (to which we might add attentional patterns); narrative, she thinks, is better than standard forms of cognitive engagement (such as argument) at successfully altering these things.⁵⁶

Using these sorts of techniques, then, Charlotte might try various things. She might apply methods associated with Cognitive Behavioural Therapy. She might, for

⁵³ Interestingly, though Saul (2012) suggested that individuals might be able to control their implicit biases through these individualist approaches, she suggests that the automatic and unconscious nature of implicit bias means that individuals are not *blameworthy* for them. Holroyd (2012) disagrees with Saul, suggesting that, where certain conditions are met (including, for instance, the existence of long-range control, mentioned above), we can hold blame agents for being influenced by implicit biases. I do not discuss the issue of whether we are blameworthy for our cognitive salience perspectives, but what I say here is compatible with either of Saul's or Holroyd's accounts.

⁵⁴ Camp (2017: 92).

⁵⁵ *Ibid.* (92-3).

⁵⁶ Haslanger (2011). Further, whilst not talking about attention *per se*, Eaton says something similar in connection with the cultural distaste in fat bodies. We cannot overcome our bias against fat bodies through standard forms of education, such as through providing correct information about fatness (I take this to constitute the sort of *standard cognitive engagement* that Haslanger dismisses). Instead, Eaton suggests that we must change our *sentiments* towards fat bodies. Interestingly, it has been suggested that reading narratives can be an effective way of changing our sentiments (see Coplan, 2004, and Gernsbacher et al., 1992).

instance, start to write down personality quirks of the larger people that she meets, as a way of making these traits more memorable, ultimately training herself to better notice them the next time that she meets a larger person. She might try some of the techniques suggested by Saul's discussion of implicit bias re-training in the previous chapter. She might, for instance, repeat intentions such as *when I see a fat body I will think 'intelligent'*, in the hope of better attending to the personhood-related traits of the next fat person that she meets. Alternatively, she might seek out fictions that involve main characters who are fat, as an effort to cultivate new associations with fatness and more substantive, personhood-related traits, such as autonomy and personality.⁵⁷ These new associations might help her to reduce how much attention she pays to a fat person's weight, and to instead pay more attention to their individuality.

3.4.2. Structural approach

Are these individual corrective measures going to be enough? In the context of our implicit biases, Saul ultimately concludes that the individual-focussed corrective measures she discusses have significant limitations. In her words:

“To fully combat the influence of implicit biases, what we really need to do is to re-shape our social world. The stereotypes underlying implicit biases can only fully be broken down by creating more integrated neighborhoods and workplaces; by having women, people of colour and disabled people in positions of power; by having men in nurturing roles; and so on. The only way to be fully freed from the grip of bias-related doubt is to create a social world where the stereotypes that now warp our judgments no longer hold sway over us.”⁵⁸

Saul here is building on a point that has long been a key theme of social epistemology. A large part of our cognition is inherited from the society in which we live. This means that it may turn out that we need material changes in the society around us for us, as individuals, to successfully cultivate different associations, beliefs, ideologies, and, arguably, cognitive salience perspectives.⁵⁹

⁵⁷ For instance, *Dumplin* is a recent film portraying the life of a woman who is fat, which has been praised (see Tonic, 2018) for having a fat person as a protagonist, and whose storyline doesn't involve weight loss.

⁵⁸ Saul (2012: 260-1).

⁵⁹ See Haslanger (2015) for a similar point on how to change our implicit biases.

For Charlotte to change her attentional dispositions on fat people, for instance, we might need to increase the visibility of fat people in mainstream society. Given that, as Anne Eaton notes, “Fat bodies are rarely represented in mainstream forms of entertainment and advertising”, fatness becomes attention-grabbing—it has become noticeable for many in virtue of being unusual and striking.⁶⁰ Further, Eaton notes that the *ways* in which fat bodies are portrayed in our culture are usually quite extreme; they are depicted “as unattractive, ridiculous, contemptible, and even gross and disgusting”.⁶¹ Where something is associated with such extreme sentiments, beliefs and ideologies, it is likely to be attention-grabbing. In addition to needing more fat bodies in mainstream culture, then, we might also need cultural shifts in the *types* of narratives concerning fat bodies (i.e. to less extreme ones). For Charlotte to find fatness less salient, then, perhaps part of the change needs to come from culture.⁶²

I raise these issues here because they may well be at the back of the reader’s mind during the next few chapters. In particular, many of the cognitive salience perspectives that I will examine appear to be ones reinforced by the culture around us, and are therefore likely to be difficult to shift through individual measures alone. Given that my focus in this thesis is to demonstrate the problems that mere patterns of salience can cause and constitute, as opposed to the methods for *changing* our salience perspectives, however, I will not be exploring the suggestions in this section further. I will, however, flag where these issues are most relevant for the reader that wishes to consider them in more depth.

3.5. Conclusion

In this chapter, I have introduced the cognitive counterpart to the linguistic salience perspective of the previous chapter. In the introduction, I set out to answer three questions. Firstly, what does it mean for something to be more salient in a person’s cognition? Something is salient in a person’s cognition, in the sense that interests me,

⁶⁰ Eaton (2016: 38).

⁶¹ Ibid.

⁶² One might wonder whether there will be some attentional dispositions that cannot be shifted either through individual corrective measures or societal change. It is not uncommon to hear claims that evolution has instilled in us certain ways of thinking about, perceiving, and attending to the world. Perhaps one might think that, thanks to evolutionary forces beyond our control, we simply cannot make anything about a spider more salient than its creepy run.

when we attend to it more. We attend more to one mental state over another (whether those states are beliefs, for instance, or experiences of a properties in the world) when we better notice it, better remember it, and find it more cognitively accessible. These attentional dispositions are what constitute cognitive salience perspectives.

Secondly, I asked whether and how we might evaluate cognitive salience perspectives. Here, I suggested that we might evaluate cognitive salience perspectives on instrumental or non-instrumental grounds, as I did with the linguistic salience perspectives of the previous chapter. Where our attentional dispositions activate problematic biases, we can judge those dispositions as instrumentally problematic. Where our attentional dispositions license those biases, or where they constitute a way of being epistemically biased, or a way of being harmful, we can find those dispositions problematic in themselves.

Finally, I asked how, if we do indeed judge that we are attending to the wrong thing, we are to *change* our cognitive salience perspectives. In the previous chapter, I suggested that where we judge a linguistic salience perspective to be liable to activate a problematic bias, we ought to adopt a counter-salience perspective that helps to avoid its activation. This might mean reversing the order in which we discuss some linguistic contents, for instance. Things are not as easy with cognitive salience perspectives. Our attentional dispositions, I suggested, are like habits, which cannot be changed at the click of a finger. Instead, certain habit-breaking interventions might be required to change those dispositions. I also suggested, however, that some attentional dispositions might be particularly resilient to these individual-focussed habit-breaking measures. In particular, this might occur where our attentional dispositions are reinforced by cultural narratives and practices. For individuals successfully to break with cognitive salience perspectives that are particularly culturally entrenched, I suggested that making changes to the culture itself might be necessary.

CHAPTER 4

**SALIENCE PERSPECTIVES AS
INSTRUMENTAL HARMS: THE
SUBSTANCE SALIENCE PERSPECTIVE**

4.1. Introduction

Imagine walking down the street, and noticing a Pit Bull Terrier barking aggressively in a nearby garden. At the top of your mind are claims that you have heard that Pit Bulls have certain genes that dispose them to aggression. You have to think harder to remember claims about how breed reputation can be self-fulfilling—claims suggesting that people buy dogs of a certain breed *for* their reputation, and train them to behave in accordance with that reputation.¹ Walking away, what stands out in your memory is the dog's confrontational behaviour, which you have heard is generally taken to be distinctive of the breed. Examples of more amiable Pitt Bulls you have come across in your life, which highlight the potential for this dog's behaviour to change with the right care and attention, aren't as clear in your mind. In this chapter, I suggest that these attentional dispositions are symptomatic of a broad generic salience perspective²—one that tends to make salient the internal properties of a given subject (over features of the external context in which it exists), and the distinctive traits that it happens to be displaying (over the potential for those traits to change). For reasons that will become clear in due course, I call this the *substance salience perspective*. Further, I argue that this generic salience perspective, when applied to certain subject matter, is liable to lead us to a number of false and harmful beliefs. For instance, it is more likely to lead us to *essentialise* our subject matter—to believe that the subject in question has a deep, fixed essence that determines their identity, and causally determines their outer traits. As we will see, essentialist beliefs, especially about living organisms, are not only commonly dismissed as false; at least when applied to people from disadvantaged social groups, they are regularly criticised for helping to justify and maintain their disadvantage.

The questions guiding this chapter are as follows. Firstly, what is the substance salience perspective? Secondly, is the content made salient by the substance salience

¹ Delise (2007).

² A generic salience perspective is one that can be applied to multiple subjects (see chapter 3, §3.2.2.2).

perspective central to any cognitively accessible and socially licensed biases, and, if so, does this matter? Finally, which generic salience perspective ought we to adopt on this basis of this assessment? The aim of this chapter is to demonstrate the *breadth* of impact that a generic salience perspective can have. One and the same pattern of salience, I will argue, has the capacity causally to activate multiple biases.

The specific plan for this chapter is as follows. I introduce the substance salience perspective in §4.2, suggesting that we can in fact see it at work in the traditional genetics textbook discussed in the second chapter. With this generic salience perspective on the table, I then look for resonances it might have with existing biases in Western cultures. In §4.3, I survey three prevalent biases that focus on the same contents as the substance salience perspective: *the correspondence bias*, *psychological essentialism*, and the *fixed mindset bias*. Where these biases differ with the substance salience perspective, however, is in their doxastic nature; at least according to a common way of talking about them, these biases do not just involve a focus on internal properties and the distinctive traits a subject happens to be exhibiting, but are constituted by substantive beliefs about these contents. Many have found these beliefs to be seriously harmful when applied to certain subject matter. In §4.4, I argue that the substance salience perspective is liable causally to activate these three biases. Here, I add to existing literatures on these biases by suggesting one, particularly subtle trigger for them. If these biases are indeed harmful, then this gives us grounds for an instrumental critique of the substance salience perspective (indeed, my focus in this chapter and the next is on such instrumental critiques of salience). Finally, I suggest in §4.5 what to do on the basis of this finding. Adopting what I call *the process salience perspective*, in one's cognition and language, looks to avoid activating the various problematic folk psychological biases discussed above. In contrast to the substance salience perspective, this perspective makes salient a given subject's external context, and the potential for change *from* whichever traits that subject currently exhibits.

In this chapter, then, I follow the epistemically naturalised methodology laid out in chapter 2. Because I have already identified a potentially problematic salience perspective (namely, that which was identified in the traditional genetics textbook discussed in chapter 2), however, I begin not with a survey of existing problematic biases, as was the suggested first step of this methodology, but with the salience perspective itself. With this perspective on the table, I then carry out the methodology as detailed in chapter 2: taking stock of existing problematic biases, before consulting empirical studies

that indicate which salience perspectives are likely to be effective in preventing the activation of these biases.

4.2. Introducing the substance salience perspective

Think back to Annie Jamieson and Gregory Radick's *Genetics Pedagogies Project* discussed in chapter 2 (§2.3.3). Jamieson and Radick analysed a traditional undergraduate genetics textbook, finding that how this textbook ordered genetics information helped students to develop deterministic views of genes—to develop false and harmful beliefs that genes are the *super-causers* of most if not all of an organism's traits, in the sense that the environment plays little to no role. The environment is also believed to be unable to affect genetic function, meaning that, for the genetic determinist, genetic processes, and the traits to which they give rise, are fixed. Let's remind ourselves of this case study.

The traditional textbook started with, and therefore made most salient, simple cases of particular genes being *for* certain traits, such as being *for* a particular pea texture. The function of these genes was relatively unaffected by environmental changes; these genes would reliably correlate with a given pea texture across differences in the cellular environment around the gene, as well as differences in the environments in which the organism found itself. *Interactionist* research, which demonstrated how many of our genes can instead function in a (sometimes unexpected) variety of ways depending on their interaction with different (cellular and organismal) environments, was discussed much later, and was therefore given less salience. The revised textbook reversed this pattern of salience, beginning instead with these interactionist findings.

Can we be more specific about the different patterns of salience in these textbooks? In what follows, I suggest that these textbooks can be interpreted as emphasising two different things.

Firstly, the focus in the traditional textbook is on *internal* properties relevant to trait development—on identifying how a given gene inside the plant plays a role in the development of the plant's smooth seeds, for instance. This is contrasted with a focus on external environments in the revised textbook. We hear how non-social environments (e.g. climate and altitude) as well as social environments (e.g. social norms on diet and exercise) play a role in the development of an organism's traits, as well as in how a given gene functions.³

³ Jamieson & Radick (2013: 591).

Secondly, the traditional textbook focussed on the distinctive traits that the organisms in question happen to be displaying. More specifically, the Mendelian research made salient by this textbook focuses on contrasting the plants with smooth seeds, as distinct from those with wrinkly seeds. Finding genes that then helped to *explain* those distinctive traits was the focus: namely, Mendelian genetics identifies what it calls the *gene for* smooth peas, explaining the distinctive smooth seeds of the former plant, as well as the *gene for* wrinkly peas, explaining the distinctiveness of the wrinkly seed plant. Instead, the interactionist research made salient by the revised textbook focuses on what Jamieson and Radick call “the contingent nature of development”, and the potential that this contingency brings for a *change* in an organism’s traits, away from those it currently displays.⁴ In other words, the revised textbook makes clear early on that there are innumerable untested environments that could change genetic function in surprising, as yet unknown, ways.⁵ Plants that are currently displaying smooth seeds might instead develop wrinkly seeds if they develop in different environmental conditions (such as altitude, climate, and so on).

The traditional textbook is a subject-specific instantiation of the generic substance salience perspective, as I introduced it in the introduction. By discussing Mendelian genetics *first*, it makes salient the *internal properties* of a subject (i.e. its genes), and the distinctive traits the subject happens to display (i.e. its wrinkled or smooth seeds). Why did the substance salience perspective, as adopted by the traditional textbook, result in more students endorsing genetically determinist beliefs? Well, for the reason that I have suggested in previous chapters: where a bias is cognitively accessible and socially licensed, simply making salient content central to that bias is liable to activate it. We have already argued in chapter 2 (§2.3.3) that genetic determinism is cognitively accessible and socially licensed. Further, our discussion here clarifies the ways in which the substance salience perspective makes salient content central to this bias. Firstly, internal properties are central to the genetically determinist belief that an organism’s traits are caused by internal properties, namely, its *genes*. The actual distinctive traits that a given subject displays are central to the genetically determinist belief that our traits, and the genetic processes that cause them, are fixed. In other words, central to the belief that our genes and traits are fixed is the idea that one need only ever consult the distinctive traits that a given subject happens to be displaying, such as its wrinkly seeds, as these are

⁴ Ibid. (594).

⁵ Ibid. (589).

all it ever will display. (*There's no potential for it to develop different, smooth seeds, because its seed type is determined by its distinctive 'wrinkly seed' genes—genes whose function is also fixed.*)

The traditional textbook, then, whilst differing with the revised textbook only in virtue of the salience perspective that it employed, increased student endorsement of genetically determinist beliefs because it made salient content that is central to those beliefs.

The questions that I would like to address in the next section are as follows. Firstly, is the content made salient by the substance salience perspective central to any other cognitively accessible and socially licensed biases, outside of genetic determinism? Secondly, if it is, is this content associated with deterministic thinking elsewhere?

4.3. Resonances with psychological biases

Before I go about answering these questions, it is important to make something clear. In what follows, I draw from research into psychological bias. In this literature, the biases that I discuss below are regularly described in ways that imply that they are universal and even innate, meaning that they exist in humans from birth as opposed to having been caused by cultural environments.⁶ For instance, one of the biases that I discuss, *psychological essentialism*, has been referred to as a “basic cognitive predisposition”⁷, a “basic feature of the human mind”,⁸ as something that is not taught and is not a cultural artefact,⁹ and so on. Research is growing, however, into the cultural variability of these biases. For instance, Christopher Olivola and Edouard Machery cite evidence that susceptibility to psychological essentialism can vary wildly across cultures, from as little as 10% of people in one culture holding essentialist intuitions, compared to 70% in another country.¹⁰ As for another bias discussed below, *the correspondence bias*, research suggests that Japanese people are far less likely to display this bias than American people.¹¹ In an effort to acknowledge this cultural variability in what follows, I will make

⁶ For examples of this tendency, see Gelman, (2003), Gelman & Taylor (2000), Demoulin et al. (2006: 26), Cimpian & Saloman (2014), and Dar-Nimrod & Heine (2011). For a critical discussion of this tendency in the literature, see Olivola & Machery (2014: 499) and Nisbett & Norenzayan (2002).

⁷ Gelman & Wellman (1991: 243).

⁸ Cimpian & Saloman (2014: 473).

⁹ Leslie (2013: 115).

¹⁰ Olivola & Machery (2014: 499).

¹¹ Miyamoto & Kitayama (2002). Miller (1984) also suggests that Hindu people are far less susceptible to the bias than American people.

clear which populations the research into these biases has been drawn from. Pointing out that there is cultural variability in terms of susceptibility to these biases should not, however, lead us to conclude that they are therefore easy to shift. Within a particular culture, the cultural institutions, practices and narratives that support a given bias can themselves be very difficult to change; feminists and race equality activists fighting centuries of sexism and racism can attest to this.¹² Even when not innate, then, changing the psychological biases *within* a given culture can be easier said than done.

With these qualifications in mind, let's move on to consider which biases, if any, resonate with the substance salience perspective.

4.3.1. The correspondence bias¹³

Imagine meeting a new colleague in the staff office at work. Whilst talking to them, they drop their mug of tea, which smashes on the ground. Studies on American adults – the demographic that most psychological bias research has been tested on – show that study participants are far more likely to infer that this new acquaintance has a clumsy character, than to infer that features of the situation caused her action, such as the mug handle being unexpectedly slippery, or her being in a rush. This tendency is said to reflect a belief that individuals' *characters* causally determine the distinctive traits and behaviours that they display, so that features of the external situation play little to no role.¹⁴ *The correspondence bias* refers to this belief—a belief that Yuri Miyamoto and Shinobu Kitayama describe as “extremely robust in North America”.¹⁵

This bias is so pronounced that people in this population tend to prefer explanations of an individual's behaviour that cite their unique character even when logical reflection suggests that they should not. In one study, subjects were shown essays either supporting or opposing Fidel Castro. When they were told that the essayist *had not been free to determine which side of that issue they would argue for*, subjects nevertheless inferred strong pro- and anti-Castro attitudes respectively from the essayist's paper. Reflecting on this evidence, Daniel Gilbert and Patrick Malone concluded: “Here were perfectly

¹² Hood-Williams (1997: 43), for instance, talks about the “*longue durée*” of social relations.

¹³ The correspondence bias is regularly discussed as equivalent to what others call *the fundamental attribution error* (Harman, 1999). In what follows, I also take them to be equivalent. Some do suggest that there are differences (see e.g. Gawronski, 2004), but, for my purposes here, I need only a coarse-grained analysis and so do not engage with these finer details.

¹⁴ Gilbert & Malone (1995: 21).

¹⁵ Miyamoto & Kitayama (2002: 1239).

intelligent college students who, when exposed to the coerced political statement of another student, seemed to be saying, “Well, yes, I know he was merely completing the assignment given him by his debate coach, but to some degree I think he personally believes what he wrote”.¹⁶

Does the correspondence bias reflect the hallmarks of the substance salience perspective? Well, the bias’s focus on an individual’s inner character dispositions, instead of the situation that they are in, echoes the substance salience perspective’s emphasis on internal properties. Secondly, the bias involves focussing on however an individual happens to be behaving (whether it’s breaking a mug, or reading a pro-Castro paper), instead of considering how they might act differently in alternative situations and times. This reflects the substance salience perspective’s emphasis on the distinctive traits a subject happens to exhibit.

What we can say, though, is that the correspondence bias goes beyond merely focussing on these contents. Unlike the substance salience perspective, it makes substantive claims about these contents. In particular, it involves believing that looking to internal, character-based explanations offers most if not all of the explanatory insights into the subject in question’s traits and behaviour. These internal traits are believed causally to determine the subject’s behaviour, so that situational factors play little to no role. Instead of simply focussing on the distinctive traits the individual happens to be manifesting, the correspondence bias involves believing that these traits are representative of how that individual normally acts, across circumstance and time.¹⁷ These traits are fixed, so the belief goes, because the internal character traits causally determining those outer traits are themselves fixed. In particular, then, the correspondence bias offers *explanatory* beliefs about the contents made salient by the substance salience perspective.

The reader at this point might wonder whether I am right to explicate the correspondence bias as being *doxastic*—as consisting in beliefs. I take for granted this formulation of our psychological biases (I do briefly consider an alternative reading in chapter 6). There is certainly a precedent for this way of talking about our biases; they are regularly referred to as *implicit theories*, for instance, which are further cashed out as comprising (also often implicit) lay beliefs.¹⁸ The implicit nature of these beliefs is worth

¹⁶ Gilbert & Malone (1995: 24).

¹⁷ Harman (1999: 317).

¹⁸ Haslam et al. (2000: 114) refer to a “growing recognition that [folk psychological] concepts are embedded in theories rather than simply being clusters of covarying

emphasising; an individual who holds the various beliefs discussed in this chapter is often not conscious of holding them. This means that, instead of directly asking a person if they are biased in the ways detailed in this chapter, indirect measures are often required.¹⁹

The reader might also wonder whether the beliefs of the correspondence bias are *partially constituted* by the substance salience perspective in its cognitive form. Part of what it is to believe that a person's character determines their behaviour, one might think, is to find the internal properties of a subject more salient in one's attention than features of their external situation. Indeed, Gilbert and Malone sometimes refer to the correspondence bias as partially involving a "tendency to overlook or ignore situational forces", which sounds like a claim about attention.²⁰ Whether one does want to say that the correspondence bias is partially constituted by the substance salience perspective depends on how one cashes out the type of attention characteristic of the correspondence bias. Cognitive salience perspectives, remember, are constituted by dispositions to give *relative salience* to x over y ; both x and y exist in one's attention, but x is foregrounded and y backgrounded. Attending to x so that one *ignores* y , however, does not count as employing a cognitive salience perspective. Here, I will take at face value Gilbert and Malone's phraseology of *ignoring* situational factors, and presume that the correspondence bias is not partially constituted by the substance salience perspective.²¹

For the purposes of this chapter, then, I am taking the correspondence bias to be constituted by beliefs. Further, I am presuming that the correspondence bias is not partially constituted by the substance salience perspective, but simply that this

properties". For instance, they discuss "laypeople hold[ing] essentialist theories" (ibid.), consisting of "essentialist beliefs" (ibid. 116). For others who fit this trend, see Heyman & Giles (2006) Haslam et al. (2006), Gelman (2005), Stanley & Blanchard-Fields (2011), Dweck (2000), Medin & Ortony (1989), and Wilson & English (2017). In chapter 1, I flagged that some do not consider implicit beliefs really to be beliefs. Instead, it has been proposed that they are, for instance, non-propositional things like *associations* (Kihlstrom, 2004). In this thesis, I take the use of the phrase *implicit belief*, used by the theorists that I cite here and elsewhere, on face value. Implicit beliefs, I presume in this thesis, really are beliefs.

¹⁹ Given that many of the beliefs that interest psychologists are held implicitly, there is a great range of techniques to measure them that have been developed (Rudman, 2011).

²⁰ Gilbert & Malone (1995: 27).

²¹ If the correspondence bias (and the other biases discussed in this chapter) was (were) partially constituted by the substance salience perspective, then we might be able to develop a non-instrumental critique of the substance salience perspective. If we can show that these biases are intrinsically (epistemically and/or ethically) problematic, then, insofar as the substance salience partially constitutes them, it inherits their intrinsic problems. See this method for evaluating salience perspectives on non-instrumental grounds discussed in chapter 3 (§3.2.2).

perspective makes salient content central to the constituent beliefs of this bias. (Further, I apply these assumptions to the other biases discussed in this chapter.) With this in mind, I can move onto the next question.

Are there problems associated with the correspondence bias? There certainly seem to be epistemic costs, according to many who discuss this bias. Generally speaking, the belief that our behaviour is mostly accounted for by character traits is considered to be false, meaning that the correspondence bias leads us, in Betram Gawronski's words, "to systematic errors in first impressions of other individuals".²² The Castro example gave a clear example of how illogical the bias can be.

Given that this thesis is focussed on harm, however, we can instead turn to a different question: do the false beliefs of the correspondence bias have ethical significance? Many suggest that they do. For instance, thanks to the false belief that a person's behaviour is almost entirely caused by their internal character, the correspondence bias is often criticised for leading one wrongly to blame a person for their behaviour when in fact they were constrained by external factors outside of their control.²³ This might involve unfairly believing that a person's poverty, for instance, is down to faults in their character, instead of recognising the social situations, structures and practices that have constrained them. These latter factors might include the relative lack of opportunities that they had growing up in a deprived area, the racism that they experience when looking for jobs, the lack of role models in positions of power with their skin colour, gender, accent, and so on. Or, it might involve unfairly believing that a person's success, such as their confidence and powerful job, is down to merits in their character, instead of down to the social factors, such as their wealth and the plentitude of positive role models with their social attributes, that have unfairly privileged them.²⁴

Generally, feminists and other social theorists have pointed to the harms that can be generated by believing situational factors to play little role in shaping an individual.²⁵

²² Gawronski (2007: 194). Choi, Nisbett, and Norenzay, (1999: 47) also emphasise that the core belief behind the correspondence bias is "One of the greatest and most remarkable misunderstandings we have about people". Gilbert and Malone (1995: 21) concur, referring to the correspondence bias as a "mistake". Further, there are many studies documenting how our behaviours are highly responsive to the situations in which we find ourselves, meaning that the correspondence bias involves us regularly making mistaken predictions about the consistency of an individual's behaviour. See, for instance, Isen & Levin (1972).

²³ Gilovich & Eibach (2001: 25-6).

²⁴ Scopelliti et al. (2017: 27).

²⁵ See Willett (2016) and the references therein.

As indicated above, the concern is that explaining people's traits in terms of internal dispositions means dismissing the broader structures and practices of power that in fact oppress, or privilege certain individuals qua their social group. (Very broadly, we can say that oppression refers to the unjust social institutions and practices that disadvantage certain social groups, and advantage others, in a variety of ways.)²⁶ This in turn means that the oppression can continue; why change social structures and practices (and therefore the oppression that they in fact cause and constitute) when they have nothing to do with why an individual has ended up the way that they have? Only that individual is to be blamed (or praised) for their undesirable (or desirable) traits and behaviour.²⁷ The correspondence bias, then, is criticised for justifying unjust inequalities, including those connected with oppression.

4.3.2. Psychological essentialism

- A. A woman receiving a heart transplant from a man will gain some *male energy*.
- B. The scarcity of women in science is down to male-female differences in intrinsic aptitude.
- C. Evolutionary theory is false, because one species cannot transform into another.²⁸

These are examples of psychologically essentialist claims that feel intuitively true to many of us, particularly in Western cultures.²⁹ Psychological essentialism, another widespread and compelling psychological bias in these cultures, represents a tendency to ascribe deep, intrinsic and fixed essences to certain categories, both social and natural.³⁰

Essences are believed to be substances, such as genes or souls, which give individuals their category identity, and generate any outward similarities between members of the same category.³¹ XX chromosomes might be thought of as the essence of being a woman for instance. They purportedly give relevant individuals their identity as women, and are causally responsible for the various traits women supposedly share,

²⁶ This is influenced by Bohmer & Briggs' definition of *oppression* (1991: 155)

²⁷ For these reasons and others, Harman (1999: 330) goes as far as to suggest that the correspondence bias "has deplorable results, leading to massive misunderstanding of other people, promoting unnecessary hostility between individuals and groups, distorting discussions of law and public policy, and preventing the implementation of situational changes that could have been useful".

²⁸ These examples are borrowed from Gelman (2005).

²⁹ See Olivola and Machery's (2014) study mentioned at the beginning of §4.3, regarding cultural variability in susceptibility to psychological essentialism. See also O'Connor & Joffe (2014: 498).

³⁰ Medin and Ortony (1989), Gelman (2005).

³¹ Gelman, (2003: 61).

such as being nurturing, and being underrepresented in STEM fields. Essences are conceived as existing inside individuals from birth, and are considered to be largely unaffected by any environmental changes the organism goes through. This means that the traits that essences give rise to are largely fixed from birth.³²

This does not mean that the individual in which the essence resides must also not change, but rather that any changes that the individual goes through, such as developing and aging, come pre-specified by the essence. In this sense, the essence prescribes what Susan Gelman calls the “innate potential” of the individual.³³ It is this innate potential that is believed to be unalterable by the environment. For example, studies show that American children who learn that a baby rabbit was raised by goats or monkeys since birth predict that its upbringing will have no impact on its developmental stages and behaviour; when it is old enough, the belief is that it will grow up acting like a rabbit, hopping and eating carrots.³⁴ Similar views have been found in Menominee Indian children, Yucatec Mayan children, and urban Brazilian children.³⁵ As children get older, more concessions to the role of the environment are made, but internal properties are still favoured in explanations. As Sarah Jane Leslie suggests, these essentialist beliefs become more implicit; whilst adults become less inclined to state them explicitly (though, many essentialist claims *are* explicitly stated by adults),³⁶ measures that aim to uncover our *implicitly* held beliefs indicate that essentialism is alive and well into adulthood.³⁷ (These issues and others will be considered in more depth in the next chapter, where I discuss psychological essentialism about sex/gender.)

Of particular interest to me are the resonances between psychological essentialism and the substance salience perspective. The focus on genes or immaterial souls echoes the substance salience perspective’s focus on internal properties. The focus on (what are judged to be) the distinctive traits that particular gender or race groups

³² See Gelman (2005).

³³ Ibid.

³⁴ Gelman & Hirschfeld (1999). When the rabbit is old enough to eat carrots, American children believe, in Gelman’s (2005) words, that “this property [of carrot-eating] will eventually be expressed”, even if the rabbit’s adoptive goat parents never eat carrots. See also Gelman & Wellman (1991) for similar findings.

³⁵ See Leslie (2013: 117) and the references therein.

³⁶ Gelman (2005) makes this point. Consider, for instance, how common it is to find newspaper (and academic) articles suggesting that complex traits from male aggression to homosexuality are *in the genes* (Dar-Nimrod & Heine, 2011).

³⁷ Leslie (2013).

exhibit, for instance, resonates with the substance salience perspective's emphasis on a subject's distinctive, actual traits.

As with the correspondence bias, however, psychological essentialism does not just involve a focus on these things. It is also constituted by truth-conditional, explanatory beliefs about them. It involves a belief that *there are internal properties*, namely, *essences*, that give individuals their identities. There is a belief that these internal properties *causally determine* the innate potential of an organism. One need only look inside an individual, so goes the thought, to find out about how an organism will develop—the external environment is believed to tell you little to nothing about an organism's development. There is also a belief that the distinctive traits and behaviours that certain categories tend to exhibit are indicative of the category's essence, and that, thanks to the unchanging nature of this essence, these traits and behaviours will reliably be manifested across time and circumstance. The essentialist might propose that there will always be more men in science, for instance, given their intrinsic aptitude for the subject. Even the concept of innate potential, which describes developmental *change*, is believed to be fixed; potential for variation *from* this developmental path is considered largely impossible.

As with the correspondence bias, there are multiple problems with psychological essentialism. Epistemically speaking, there is overwhelming consensus that essentialism is false both as a metaphysical³⁸ and biological³⁹ doctrine. There are no essences, as the psychological essentialist conceives them. Again, though, my focus is on the ethical side of things.

False essentialist beliefs are associated with a great deal of harm. For instance, psychological essentialism comes with a higher endorsement of stereotypes.⁴⁰ One problem associated with stereotypes is *stereotype threat*; as mentioned in chapter 2 (§2.3.1), stereotypes can have a self-fulfilling potential, meaning that individuals either are, or feel at risk of, conforming to stereotypes about their social group. Stereotypes that build on essentialist premises are certainly effective in shaping an individual's behaviour; for instance, women who read an essay saying that men are better at maths due to an innate biological difference (read: essence) then perform worse in a maths test than women who

³⁸ Mayr (1991), Medin & Ortney (1989).

³⁹ Dupré (1993).

⁴⁰ See Haslam et al. (2006) and Cimpain & Salomon (2014: 1298). The discussion above discusses the harms of stereotypes, but they are also plagued by epistemic problems (Puddifoot, 2017).

had read an essay saying that men's increased ability is due to their increased effort.⁴¹ There are significant social and economic implications that come with any group, especially one that already suffers structurally systemic disadvantage, performing under their potential.⁴²

Further, psychological essentialism receives criticism for helping to foster a greater psychological and moral acceptance of racial and gender inequalities.⁴³ For instance, if one believes that the differences between white and black people are deep and immutable, then one is likely to see any racial inequalities in education, wealth, employment, and so on, as natural and insurmountable—that they're simply the inevitable outcome of different racial essences. The problems for such a belief deepen when one considers research demonstrating that many of us are quick to move from descriptive premises concerning the inevitability of a given phenomenon, to normative premises concerning the justifiability of that phenomenon. In other words, where one believes that an inequality is the result of natural differences between races, one is liable to infer that those inequalities are therefore morally acceptable.⁴⁴ We mentioned in chapter 2 how reports of racial differences in IQ tests have been used in arguments that attempt to justify removing educational resources for black people, for instance.⁴⁵ Psychologically essentialist beliefs are regularly used in arguments aiming to justify this unjust disadvantaging of certain social groups. They are, in other words, often used to justify relations of oppression.

4.3.3. The fixed mindset bias

Alvin takes a physics test at school. He feels relatively confident, but when he gets his results back, he finds that he has failed. He concludes that he is obviously not good at physics, and decides not to pursue that subject the next year; no point putting in more

⁴¹ Dar-Nimrod & Heine (2006).

⁴² Walton & Cohen (2007). There many also be epistemic issues that are generated here, insofar as these essentialist beliefs end up with the wrong *direction of fit* with the world (see the brief mention of this in chapter 2 §2.3.1).

⁴³ Heyman & Giles (2006), Williams & Eberhardt (2008).

⁴⁴ See Williams & Eberhardt (2008) for a study confirming that the inference from natural-to-justifiable (or good) is a common one for people who have essentialist beliefs about race. Generally, the *appeal to nature* fallacy, in which something is considered good if it is natural, or bad because it is unnatural, is particularly common in human thought (Twarek & Cimpian, 2016). These authors also explain the epistemic and ethical problems with this type of inference.

⁴⁵ Chapter 2 (§2.4.1). See Jensen (1969) for an example of this argument.

effort, he thinks, if he just doesn't have the brain for the subject. Alvin's mindset can be characterised as *fixed*; he presumes that his abilities are not changeable.

According to Carol Dweck, I have a fixed mindset when I see my ability (e.g. intelligence) as "an entity that dwells within [me] and that [I] can't change".⁴⁶ I might be able to learn new things, but I can't change how intelligent I am. With this mindset, I would only seek opportunities to *confirm* my intelligence, instead of trying to develop it through work and effort. If I fail a test, I take this to indicate that I simply do not have the relevant ability. Even if there are clear situational factors that have contributed to my failure (such as an illness, or family grievance), I nevertheless prefer to put my result down to my intrinsic aptitude. My belief that personal abilities are fixed corroborates this interpretation of my failure; I need only one example of my failure to feel justified in predicting great cross-situational consistency of my failing behaviour.⁴⁷ Research indicates that many of us, especially in America and Europe, have this fixed mindset about our abilities. Dweck argues, for instance, that school cultures in many parts of Europe and America are premised on fixed mindset beliefs and values.⁴⁸

The fixed mindset has crossovers with the correspondence bias and psychological essentialism. All suggest that our inner traits are fixed, and determine our (fixed) outer traits. Unlike the wider scope of the correspondence bias, however, the fixed mindset focuses on traits associated with *intelligence*, and ability more generally.⁴⁹ Further, a fixed mindset is concerned with interpretations of one's *own* abilities, whereas the correspondence bias looks only at the traits of *others*. As for psychological essentialism, the essentialist posits inner entities (read: essences) that individuals possess in virtue of being members of a particular *category* (such as the category *women*), whereas someone with a fixed mindset will tend to postulate inner entities that are idiosyncratic to oneself, qua individual. The fixed mindset, like the correspondence bias, also comes with less metaphysical baggage than psychological essentialism; unlike for the essentialist, there is no specific claim that one's inner traits determine one's identity.

⁴⁶ Dweck (2000: 2).

⁴⁷ Dweck, Chiu & Hong (1995: 277).

⁴⁸ Dweck, (2006a). As we will hear in §4.5 below, this does not tend to extend to East Asian cultures.

⁴⁹ The fixed mindset bias is overwhelmingly discussed in relation to the trait of intelligence in particular, but Dweck has also considered fixed mindsets in relation to moral character (Dweck, Chiu & Hong, 1995).

Whilst there are differences, then, there are certainly similarities between these three biases.⁵⁰ Not least because of the parallels, finding the footprint of the substance salience perspective is not difficult. Demonstrating the substance salience perspective's emphasis on internal properties, the fixed mindset focuses on entities that *dwell within* the individual. Further, having a fixed mindset involves focussing on whichever test results one happens to get, which echoes the substance salience perspective's focus on the distinctive traits one happens to be displaying.

The fixed mindset bias certainly goes beyond the substance salience perspective insofar as it also comprises explanatory beliefs about these two contents. Instead of simply focussing on internal properties, a fixed mindset involves believing that one's intelligence is caused by an entity that is found *inside* of one—an entity that is believed impervious to outside influence (such as educational resources). Instead of simply emphasising the distinctive traits one happens to be displaying, a person with a fixed mindset believes that their actual performance (on, perhaps, a test) is evidence of an unchangeable level of ability; potential for a *different* level of ability simply doesn't exist, so there is no need to consider the chance for one to *change* and develop one's abilities.

Does the fixed mindset come with problems? The accumulating evidence of brain plasticity, and the malleability of things like IQ results given educational opportunities and the level of wealth in one's family, has led many to conclude that the fixed mindset is wrong about our abilities.⁵¹

Unsurprisingly, the fixed mindset can be prudentially very costly; faced with failure, one does not engage in opportunities to learn and improve, and thus one's personal development can become stunted.⁵² Evidence also suggests that those with fixed mindsets are particularly susceptible to conforming to stereotypes. I mentioned earlier that women who are reminded of the stereotype that men are better at maths underperform in a subsequent maths test. Research also shows that women who hear the fixed mindset message – namely, that maths ability is a fixed, inner trait – in conjunction with this stereotype are *especially* likely to underperform men.⁵³ Further, holding a fixed mindset makes it more likely that these women then go on to pull out of maths classes.⁵⁴

⁵⁰ Dar-Nimrod & Heine (2011) and Norenzayan et al. (2002) have also noticed similarities between these biases.

⁵¹ Fine et al. (2013); Ramsden et al. (2011)

⁵² Good, Rattan & Dweck (2012: 701-2).

⁵³ Good, Rattan & Dweck (2007).

⁵⁴ Ibid. Again, similar epistemic issues concerning *direction of fit* apply here (see a brief discussion of these issues in chapter 2 §2.3.1). Dropping out of maths classes will end up

Research also shows that, where women read a company mission statement that endorses the fixed mindset approach to ability, those women perform worse in business-related tasks, disengage more from those tasks, and worry more about being negatively stereotyped.⁵⁵ Similar effects occur for Black and Latino individuals, upon reading that a given company endorses the fixed mindset.⁵⁶ It is through these sorts of consequences that we can start to see how the false beliefs constituted by the fixed mindset have not just a moral dimension, but one related to oppression. A fixed mindset makes things worse for those from oppressed social groups, by compounding problems they already face, such as stereotype threat. Further, if one is already of an oppressed group, who suffers systematic disadvantages in society, then a fixed mindset will encourage one falsely to believe that the shortcomings in one's abilities are inherent and fixed, as opposed to caused by oppressive features of one's social context (such as a lack of resources for, or explicit and implicit biases against, one's social group). This helps to protect those oppressive arrangements; indeed, it makes one less likely to call for the oppressive institutions and practices, which have *in fact* impeded the development of one's abilities, to be *changed*.

4.4. The substance salience perspective: instrumental harms

Whilst the substance salience perspective simply involves giving relative salience in one's attention to internal properties of a subject, and to whichever distinctive traits it happens to be displaying, then, we have seen that the three biases discussed here instead involve substantive *beliefs* about these contents.

Further, the sorts of beliefs that comprise these biases are principally of an explanatory nature. They purport to explain an individual's or group's traits, by suggesting what has *caused* those traits. (Psychological essentialism also involves a metaphysical belief, insofar as it posits that essences are responsible for making individuals what they fundamentally are.)

More specific still, the type of explanation that these biases engage in is *deterministic*. In chapter 2, I defined genetic determinism as the belief that genes are the

affecting these women's maths ability. The stagnation of their ability looks to these women like *confirmation* that their maths ability is unchangeable, which ultimately looks to confirm their fixed mindset. In fact, their belief might have the wrong direction of fit with the world.

⁵⁵ Emerson & Murphy (2015).

⁵⁶ Emerson & Murphy (2013).

super-causers of an organism's traits, so that environmental factors play little to no role in that organism's trait development. Instead of genes *per se*, the biases at hand treat a range of specifically *internal* properties as super-causers, whether these are character dispositions, essences, or internal abilities. These internal properties are thought to *determine* one's behaviour and observable traits so fully that situational and environmental factors play little to no role.

In addition to epistemic problems, we have seen that research indicates that this internal property determinism comes with a range of harms when applied to *people*. (I have said nothing here to indicate that internal-property determinism would have epistemic or ethical problems when applied to other subjects, such as mechanics or physical laws.)⁵⁷ More specific still, the potential for this internal-property determinism to cause *harm* in particular is most significant when the subject matter includes people from oppressed social groups. The harms associated with the correspondence bias were clearest when one made internal, dispositional attributions for behaviour and traits that were instead caused by inequalities in society. Those who condemn essentialist beliefs focus on how they are often used to justify societal inequalities that hurt women, ethnic minorities, homosexuals, those with mental disorders, and other oppressed groups the most.⁵⁸ Research suggested that the fixed mindset disproportionately hurts individuals from groups who suffer from negative stereotypes about their ability, such as women in the context of maths (in other words, it tends to hurt individuals from oppressed groups). In chapter 2, we heard that genetic determinism has a long history of being used to justify societal conditions that hurt oppressed groups; Arthur Jensen, for instance, suggested curtailing educational resources for black people in the 1970s because their (purported) lower average intelligence levels were caused by fixed, genetic differences.⁵⁹ The internal property determinism in these folk psychological biases generates these harms for those in oppressed groups because it takes existing differences between social groups, including those caused by unjust social arrangements, and encourages us to explain those differences in terms of the fixed, internal properties of the individuals from those groups.

⁵⁷ There is reason, however, to think that internal property determinism might harbor epistemic (though perhaps not ethical) problems when applied to biology generally speaking; genetic determinism is considered to be an epistemically flawed way of approaching genetics generally, even for non-human organisms, such as plants (Lewontin, 2000).

⁵⁸ Haslam & Whelan (2008).

⁵⁹ Jensen (1969).

It is for this reason that we might in fact want to consider extending the harms of internal property determinism to certain other subjects. Think back to the dog mentioned in the introduction. Pit Bull Terriers are, due to the *Dangerous Dogs Act*, banned in the UK, as the breed is deemed to be aggressive. There are many vocal critics about this act. The sort of criticism levied against it tends to sound like the one we have just rehearsed against internal property determinism. For instance, *The Kennel Club* suggests that the act focuses only on (what they claim to be false) claims about the *genetics* of dog breeds, and, in their words, “ignores the most important factors that contribute to biting incidents – primarily anti-social behaviour by people who train their dogs to be aggressive and irresponsible dog ownership”.⁶⁰ In addition to harming the public (by failing to prevent biting incidents), they suggest that the act also has costs for the dogs themselves, “[resulting] in the unnecessary euthanasia of dogs simply because of their breed or type”.⁶¹ Perhaps we can pursue an argument, then, that the *Dangerous Dogs Act* is an example of an internal property determinism that harms dogs, by unnecessarily euthanizing them. Existing trends of Pit Bulls acting aggressively are presumed to be caused entirely by the internal properties of those dogs (i.e. Pit Bull genetics), instead of by environment factors, such as irresponsible dog ownership. This might give us pause for reflection on whether to extend the harms of internal property determinism beyond groups of human people. I will not pursue this thought further here, however.

What can all of this tell us about the substance salience perspective? We heard in the previous chapters that having some content salient in one’s mind (perhaps through first having it made salient in language, as with the traditional genetics textbook) can be sufficient to activate an inference to a (set of) substantive, *belief(s)* about that content. This occurs where a salience perspective makes salient content that is central to that (set of) belief(s), and where that (set of) belief(s) is cognitively accessible and socially licensed.

Indeed, I have shown that the biases above share a focus on the same content as that made salient by the substance salience perspective. Further, as discussed, research demonstrates that these biases are cognitively accessible and socially licensed for many of us, at least in many Western cultures. This means that, in these cultures, use of the substance salience perspective risks activating those biases. Once activated, these biases go to work on interpreting the explanatory (and in certain cases the metaphysical) *significance* of the contents made salient by the substance salience perspective. The internal

⁶⁰ The Kennel Club (2019).

⁶¹ Ibid.

properties of an individual might go from simply being salient, for instance, to being deemed entirely causally responsible for an individual's outer traits.

Where it does activate these biases, it would most likely do so in an under-the-radar manner. We have already heard in previous chapters how patterns of salience subtly shape our cognition and action, mostly instigating inferential processes that bypass our conscious awareness. This, I suggested, can make an individual even more likely to think and act in ways that demonstrate an endorsement of that bias (unaware of that bias's activation, they have not had chance consciously to reflect on whether they wish to *reject* inferences to that bias, and its associated influences on our behaviour). The implicit nature of how the substance salience perspective would activate the folk psychological biases discussed in this chapter, then, suggests that it might have a special power in successfully inculcating them in the individual in question.

Indeed, there is already evidence that subtle triggers can activate these biases. For instance, the use of generics in language has been suggested as a trigger for psychological essentialism.⁶² A generic is a generalization that omits quantifiers such as *some* or *most*, such as the phrase *girls like dolls*. Simply omitting quantifiers in this way has been shown to lead to a higher endorsement of essentialist statements. The discussion in this chapter, then, adds to existing discussions concerning the activation of the three cognitive biases discussed in this chapter by identifying another, particularly subtle and insidious trigger for those biases: the substance salience perspective.

If we can show both that the substance salience perspective does indeed help to activate the correspondence bias, psychological essentialism, and the fixed mindset bias, and that the existing critiques of these biases stand, we can criticise the substance salience perspective by criticising its effects: by pointing to the numerous problems that many have identified with these biases. Given that the problems with these biases look most serious when they are applied to the domain of people, qua social individuals, the instrumental critique of the substance salience perspective is similarly limited to this domain.

Criticising the substance salience perspective in this way is not equivalent to suggesting that one can never employ it. There are many considerations to take into account when deciding which generic attentional dispositions to cultivate.⁶³ Avoiding the

⁶² Leslie (2007).

⁶³ I mentioned in the previous chapter that the idea of deciding which generic attentional dispositions to cultivate is not an unfamiliar one. For instance, resolving to focus less on

risk of activating epistemically problematic and harmful biases is just one such consideration, and may be outweighed in favour of another.⁶⁴ That these biases can, in addition to lead us to false beliefs, engender significant harm, however, should remind us to take the consideration noted here seriously. Further, the fact that this single salience perspective is liable to activate not one but *three* especially problematic biases, should give us serious pause for thought.

4.5. The promise of the process salience perspective

I suggested in chapter 2 that we adopt an epistemically naturalised methodology when deciding which salience perspectives to use. This involved first taking stock of which epistemically and morally problematic biases are particularly cognitively accessible and socially licensed. In this chapter, I identified the correspondence bias, psychological essentialism, and fixed mindsets as problematic biases, which take little to become activated, and which feel warranted to many of us. Secondly, I suggested that we empirically investigate which salience perspectives are effective in neutralising, or preventing the activation of, these existing biases. It is this second question to which I now turn my attention.

Fortunately, as discussed in chapter 2, we already have empirical evidence of a successful intervention, collected by Jamieson and Radick in their *Genetics Pedagogies Project*. In an attempt to prevent the activation of our genetic determinism biases, Jamieson and Radick reversed the order in which the traditional curriculum discussed information about genetics. In so doing, they therefore adopted the antithesis of the substance salience perspective. Instead of starting with Mendelian genetics, their revised pedagogy began with interactionist research. This, as I suggested in §4.2, meant making salient two key things: the external context in which the subject exists; and the potential for change *from* whichever traits it currently exhibits. It involved adopting something that I am calling the *process salience perspective*.

As we heard in chapter 2, this counter-salience perspective came with epistemic and moral benefits, insofar as it reduced endorsement of genetic determinism among students. Jamieson and Radick's research gives us reason to endorse the process salience perspective, then, at least in the context of genetics. Just as I criticised the substance

the *negatives*, whether one is attending to one's own life, another's life, contemporary affairs, and so on, is an idea that many of us come across in daily life.

⁶⁴ I mentioned this point in relation to linguistic salience perspectives in chapter 2 (§2.5).

salience perspective on the basis of its negative upshots, then, we can praise the process salience perspective on the basis of its positive upshots.

Is there any reason to think that these epistemic and moral successes of the process salience perspective might generalise out beyond the case of genetics textbooks? To see how they might, think back to Alvin. Imagine that, instead of a fixed mindset, Alvin has a *growth* mindset. He takes a maths test at school, and finds that he has failed. Instead of concluding that maths is simply not his strong suit, Alvin resolves to put in more effort. He asks for extra homework, and seeks out short tests on the internet so that he can get used to exam-style questions. He also reflects on how, unlike the situation for many other students in his class, the fact that his parents cannot afford private tuition might have negatively impacted his abilities.

Alvin's growth mindset is characterised by his belief that ability is, in Dweck and colleagues' words, a "malleable quality that can be changed and developed".⁶⁵ Failure does not indicate to him that he lacks intelligence; it instead invites him to *develop* his abilities by studying harder, and by engaging in more educational opportunities. His belief in the incremental, contextual nature of his skill means that, in Dweck and colleagues' words, at any given point he is "simply judging a momentary level of ability".⁶⁶ He cannot describe his personal qualities once and for all, since they are continuing processes whose potential he cannot prejudge. It is this belief that leads Dweck and colleagues to suggest that, for an individual with the growth mindset, "reality can never be known with any finality".⁶⁷

We can find clear parallels between this growth mindset and the process salience perspective. Firstly, instead of a focus on internal (perhaps genetic) determinants of one's performance, the focus is much more on features of the external context—on educational resources, exam condition pressures, and so on. Secondly, the potential for change, specifically, for a different, hopefully improved, level of ability to what one currently exhibits, takes centre stage.

Instead of just focussing on these things, however, the growth mindset makes substantive claims about these contents. Firstly, the growth mindset involves the belief that one's performance is primarily *caused* by factors in one's external context, such as one's school teacher quality, how many revision sessions one attends, and which environmental stresses that one is under. Secondly, whatever level of ability one currently

⁶⁵ Dweck, Chiu & Hong (1995: 267).

⁶⁶ Dweck & Leggett (1988: 263).

⁶⁷ Dweck, Chiu & Hong (1995: 269).

exhibits is believed to be changeable; with more effort, there is a belief that one *can* improve one's intelligence. Further, the extent to which one might be able to change one's intelligence is unknown; perhaps, one day, one might be able to beat a previously unbeatable competitor.

The difference between the beliefs constituted by the growth mindset and those constituted by the fixed mindset is that the former are better supported by the aforementioned empirical research into brain plasticity. There are also prudential advantages to the growth mindset. For instance, one can develop one's skills further; unlike with a fixed mindset, the growth mindset encourages one to stick with studying after encountering setbacks and failures.⁶⁸ Studies suggest that this translates into higher levels of achievement; for instance, having a growth mindset has been found to predict higher grades in university courses.⁶⁹ Research also demonstrates that levels of self-esteem, as well as enjoyment of the subject under study, are higher with a growth mindset; the suggestion is that it is easier to enjoy learning when one's anxieties about failure are lower.⁷⁰ Further, instead of the social problems engendered by the fixed mindset, the growth mindset beliefs look to have social *advantages*. For instance, having a growth mindset has been shown to help protect negatively stereotyped individuals from stereotype threat.⁷¹ In sharp contrast to the negative consequences that a fixed mindset can incur when adopted by those from negatively stereotyped groups, research is building that American women and African American students obtain higher grades when they have been encouraged to see intelligence as malleable in the ways suggested by the growth mindset.⁷² Individuals from these groups with a growth mindset have also been found to be more likely to maintain their intention to pursue maths courses in the future. In other words, having a fixed mindset can, in Catherine Good and Carol Dweck

⁶⁸ There has been some very recent scepticism concerning the efficacy of the growth mindset in relation to student achievement (Sisk et al., 2018). Others, however, suggest that this scepticism is misplaced; according to Yeager et al. (2013), the only scenarios in which the growth mindset does not lead to the benefits above is when it has been taught poorly or incorrectly. Even those who are skeptical about the benefits of the growth mindset, however, concede that students with low socioeconomic status, or who are academically at risk, benefit from its adoption (Sisk et al., 2018).

⁶⁹ Grant & Dweck (2003). See also Blackwell et al. (2007), who found that over two years, students who believed growth was possible reported higher math grades than students who believed that their abilities were fixed.

⁷⁰ Dweck (2006a: 48).

⁷¹ Good, Rattan & Dweck (2007).

⁷² Aronson, Fried & Good (2001), Good, Aronson & Inzlicht (2003), Good, Rattan & Dweck (2012).

words, “have real [and positive] consequences for [the] career aspirations and achievement” of members from oppressed social groups.⁷³

Research indicates that the growth mindset is cognitively accessible and socially licensed in certain other parts of the world, especially in East Asian cultures.⁷⁴ Here, then, the process salience perspective would be well poised to activate this prevalent set of beliefs. Given the epistemic, prudential, and moral benefits of these beliefs, we can praise the process salience perspective, insofar as it is likely to *activate* inferences to those beliefs, for its good effects.

The fact that the growth mindset has a purchase in this demographic is significant. Some have argued that humans may struggle to pay as much attention to contextual factors (such as social norms) as they do to persons and their internal properties. For instance, contextual factors are, according to Larisa Hussak and Andrei Cimpian, “seldom apparent to an observer”.⁷⁵ We can’t, for instance, *point* to a social norm. Indeed, Gilbert and Malone suggest something similar, under the term *the invisibility problem*. They say: “Actors can be weighed and behaviours can be filmed, but when one tries to point to a situation, one often stabs empty air”.⁷⁶ This would make the process salience perspective very difficult to adopt. Perhaps humans naturally don’t find contextual content more noticeable, memorable and cognitively accessible than intrinsic information! The fact that research suggests that referring to contextual information comes naturally in East Asian cultures suggests that humans can indeed find at least some of the properties made salient by the process salience perspective intuitively salient in their attention.⁷⁷

⁷³ Good, Rattan & Dweck (2012: 713).

⁷⁴ Chen & Wong (2015). Generally, East Asian countries place more emphasis on effort and the social factors that contribute to achievement (Yip, 2018; Heine et al, 1999; Uttal, 1997). This coheres with other differences that I discussed in chapter 2 (§2.3). I cited research suggesting that cultural understandings of the self in Japan, for instance, tend to highlight one’s interdependence with one’s wider social networks, which also demonstrates a focus on external context instead of internal properties (Markus & Kitayama, 1991: 227). Focussing more on external context is characteristic of the growth mindset.

⁷⁵ Hussak & Cimpian (2017: 70).

⁷⁶ Gilbert & Malone (1995: 25).

⁷⁷ The discussion concerning cognitive salience perspective change in §3.4 of the previous chapter is relevant here. If we find that the process salience perspective is proving difficult for those in Western countries to adopt, wider changes to Western culture might be necessary to allow those in this demographic successfully to shift their attentional dispositions.

As for Western cultures, there is reason to think that the prevalence and acceptability of the growth mindset is, to embrace the pun, growing. The growth mindset is being taught more and more in classrooms across America and the UK.⁷⁸ Discussions of the growth mindset are receiving more airtime in newspapers and other news sources.⁷⁹ This suggests that the growth mindset is becoming increasingly cognitively accessible and socially licensed in these parts of the world. It is therefore becoming more likely that simply making salient that which is central to the growth mindset – the external context in which the subject in question exists, and the potential for its traits to change – might be sufficient to activate the growth mindset in these demographics.

Even where the process salience perspective fails to activate the growth mindset (perhaps in conditions where the growth mindset is insufficiently cognitively accessible and socially licensed) it is at least likely to decrease the chance of activating the false and harmful beliefs contained within the correspondence bias, psychological essentialism, and the fixed mindset. Indeed, it successfully decreased the activation of another bias relating to internal property determinism, namely, genetic determinism, in Jamieson and Radick's study.

It is also worth noting that the growth mindset does not represent an obscure, one-off set of beliefs. The divide between the fixed and growth mindsets looks to be representative of a broader division in the way in which people conceptualise all sorts of phenomena. Carol Dweck and Ellen Legget refer to this as a difference between “thinking in terms of static, reified entities versus thinking in terms of dynamic, malleable processes”.⁸⁰ They cite a wide range of philosophers, anthropologists, historians of science, linguists and psychologists, who document these different ways of thinking: Whitehead, observing the contrasting paradigms of static substances vs. dynamic processes in scientific and philosophical thought; Heller, contrasting the static, pre-Renaissance conception of persons with the dynamic, developmental post-Renaissance conception; Bloom, analysing the linguistic divide between the English and Chinese languages, with only the former persistently *entifying* the properties of people and things, by turning adjectives and verbs into nouns; and Piaget, distinguishing static conceptions

⁷⁸ Hennessey (2018), Busch (2018)

⁷⁹ See, for instance, Busch (2018) Kirp (2016). There are even resources on the BBC for parents to teach their children the growth mindset (BBC, 2019).

⁸⁰ Dweck and Leggett (1988: 267).

of the world in cognitive development from conceptions that emphasise process and becoming.⁸¹

Within these contrasting paradigms, there is a focus on the internal and the actual on the one hand, and the external and the possible on the other. For instance, substance metaphysics focuses on substances, which are believed to be ontologically basic things, whose identities are *internal* to them (their identities do not depend on anything external to them).⁸² Process metaphysics instead focuses on processes, which are considered to be ontologically basic. Processes are taken to influence, and be influenced by, their surroundings, meaning that one must look at a process's wider *context* to understand it.⁸³ Further, substance metaphysics is characterised as using what Johanna Seibt refers to as a “snapshot view of reality”, which captures an *unchanging* reality of static substances that “must be what they are at any instant in time”⁸⁴. By contrast, process metaphysics is characterised as believing reality to be “continuously changing”.⁸⁵ It is because of these resonances that I have chosen the terms *substance* and *process* to refer to the two salience perspectives in this chapter.

It looks, then, that the substance vs. process division is a socio-historically important one. It might characterise substantive metaphysical commitments, such as in the substance and process metaphysics traditions mentioned above or, as I argue in this chapter, it can manifest itself in subtle shifts of salience, via the minimalist notion of a salience perspective. It is not just through the lens of the three folk psychological biases discussed in this chapter, then, that we should see the substance salience perspective. We may need to look to an even wider folk-cultural context to fully understand how the substance salience perspective is likely to function—to understand the various beliefs and ideologies it might be capable of activating. This demonstrates the potential breadth of the consequences that a single generic salience perspective can have.

4.6. Conclusion

I started this chapter with three questions. Firstly, what is the substance salience perspective? The substance salience perspective is a generic salience perspective that

⁸¹ Ibid.

⁸² Robinson (2018). See in particular §1 and §2.5.2.

⁸³ Dupré & Nicholson (2018).

⁸⁴ Seibt (2018).

⁸⁵ Ibid.

emphasises the following two contents: a subject's internal properties; and the distinctive behaviour that subject happens to be manifesting.

Secondly, I asked if the content made salient by the substance salience perspective is central to any cognitively accessible and socially licensed biases, and, if it is, whether this matters. As for the first part of this question, I have demonstrated that it is. The correspondence bias, psychological essentialism, and the fixed mindset bias all make salient the same contents as the substance salience perspective. This means that use of the substance salience perspective, in linguistic and/or cognitive form, is liable to activate these biases. As for the second part of this question, I answered that this fact does indeed matter. All of these biases are associated with what I called *internal property determinism* which, at least when applied to people, can be particularly harmful. Ultimately, this led me to criticise the substance salience perspective on instrumental grounds—on the basis of its effects.

Thirdly, I asked which salience perspective we ought to adopt in light of this discussion. It was here that I examined the promise of the process salience perspective. This perspective makes salient the opposite of that which the substance salience perspective emphasises, namely, the external context and the potential for change. It was the process salience perspective that was employed by the revised genetics textbook, and reduced student endorsement of false and harmful beliefs associated with genetic determinism. This provided empirical evidence that the process salience perspective is able to reduce support for internal property determinism. Further, the content made salient by the process salience perspective is central to the antithesis of the fixed mindset: the growth mindset. Where the growth mindset is cognitively accessible and socially licensed, using the process salience perspective might be liable to activate this mindset. Where it is, given that the growth mindset is associated with epistemic and moral advantages, we can judge the process salience perspective that activates this mindset to be instrumentally good.

Generally, this chapter aimed to demonstrate the wide reach of generic salience perspectives. The fact that a single generic salience perspective is liable to activate a number of particularly troublesome biases reveals the potential power that generic patterns of salience can have.

CHAPTER 5

**SALIENCE PERSPECTIVES AS
INSTRUMENTAL HARMS:
SEX / GENDER**

5.1. Introduction

Gender Similarities and Differences is a 2014 paper written by Janet Shibley Hyde. In it, Hyde consults meta-analyses to investigate whether the genders are more similar than they are different. Instead of discussing her findings, however, let's consider her choice of title. There's a chance it might strike us as unusual; we are so used to hearing about gender differences in our culture that hearing about gender similarities *before* differences (or indeed hearing the phrase *gender similarities* at all), might pique our curiosity. Does the relative salience that Hyde has given to gender similarities over differences, in virtue of mentioning them first, have any significant consequences for how her readers are likely to think about gender? In this chapter, I suggest that it does.

In the previous chapter, I looked at one generic salience perspective, *the substance salience perspective*, which I said risked activating various problematic cognitive biases. This salience perspective, I suggested, is most likely to cause problems when applied to subject matter relating to social categories of people. In contrast to the breadth of focus in the previous chapter, I turn in this chapter to consider a specific subject, namely, the specific social category of *sex/gender*¹ (though it should be kept in mind that what I say has implications for other topics, such as race). Narrowing the focus still, I examine a single cognitively accessible and socially licensed bias regarding this subject, namely, *psychological essentialism about sex/gender*.

Three questions drive my investigation. Firstly, is there any empirical evidence from which I can draw demonstrating that a mere change of salience perspective can activate problematic biases for the specific subject of sex/gender? Secondly, what does a survey of the prevalent sex/gender essentialism bias tell us about which patterns of

¹ In what follows, I follow Rippon et al. (2014) in adopting the conjunction sex/gender. Briefly, it is to highlight how *gender* differences of socialisation can show up as *sex* differences in our biology, meaning that it is often not possible cleanly to separate the two. When discussing another person's research, however, I stick with their terminology, hence my use of the singular word *gender* when discussing Hyde's research above.

salience we ought to avoid for attending to, and communicating about, sex/gender? Finally, what can we hope to achieve by making salient content that is the *opposite* to that which is central to sex/gender essentialism?

To answer these questions, I begin in §5.2 by examining a case study of a framing effect in the context of sex/gender; this tells us that a mere change in the order in which content is discussed, without making any further changes to that content, can mean the difference between activating existing androcentric biases, and not. (Androcentrism is a bias that involves treating maleness and masculinity as the norm in various ways.) Demonstrating that we should indeed pay attention to how patterns of salience interact with prevalent biases concerning sex/gender, this case study hopefully motivates the reader to take seriously the particular project of this chapter. In §5.3, I begin this project, by surveying a different sex/gender bias—*sex/gender psychological essentialism*. Here, I follow the epistemologically naturalised methodology introduced in the second chapter, which suggests that our decisions about which salience perspective to adopt for a subject should first take account of which problematic biases are cognitively accessible and socially licensed for that subject. Having detailed the features of sex/gender psychological essentialism, I move in §5.4 to the second step of this naturalised methodology—suggesting, on the basis of this survey which patterns of salience we should and should not use for sex/gender. Whilst there is a not insubstantial feminist literature giving advice about how to avoid activating sex/gender essentialist beliefs, I make the novel proposal in this chapter that salience perspectives might count as a particularly subtle potential trigger for our essentialist beliefs—a trigger whose subtlety can make it powerful.

5.2. Case study: A framing effect for sex/gender

Given that I am looking into salience perspectives that engender instrumental harms in the context of sex/gender, it would be apt to begin with a sex/gender-specific framing effect. Whilst the effect described here does not demonstrate the involvement of psychological essentialism, it does provide evidence that a mere change in linguistic salience perspective can be sufficient to activate substantive cognitively accessible and socially licensed biases about sex/gender in particular.

5.2.1. Bruckmüller et al.'s Framing Gender Differences

Do women lead differently than men in boardrooms? Various things might pop into our minds, concerning style, tone, and so on. What if we phrase the question slightly differently: do men lead differently than women in boardrooms? Surely this change in word order is too trivial to evoke a different response from us?

Interviewing 226 German university students, Susanne Bruckmüller and colleagues found that there was indeed a tendency to give different responses to these two questions.² Firstly, those answering the first question – namely, how *women* differ from *men* – attributed a higher level of social status and power to men in society. For instance, they were more likely to agree to various statements indicating that men had more privileges and opportunities for achievement than women. Secondly, those answering the first question were more likely to agree that existing inequalities between women and men are justified. For instance, they were more likely to endorse statements claiming that men had qualities making them better suited for these privileges and opportunities, and that the wage gap is acceptable because women and men are doing different jobs. Thirdly, participants answering the question as to how *women* differ from *men* were more likely to endorse gender stereotypes.³ In particular, they were more likely to attribute a greater number of stereotypically masculine traits to men (such as self confidence, independence, and decisiveness), and a greater number of feminine stereotypic traits to women (such as being emotional, compassionate, and warm).

All this, prompted by a simple swap of the words *women* and *men*? The authors suggest that an explanation lies in how linguistic norms intersect with background cultural associations and beliefs. By asking how x differs from y in respect of z , we invoke a linguistic norm that positions the former group (x) as, in Bruckmüller and colleagues' words, "the effect to be explained", with the latter (y) as "the implicit norm for the comparison".⁴ Which type of norm we end up invoking depends on the background cultural associations and beliefs that we have about the groups being compared. Asking how red squirrels differ from grey squirrels, for instance, might invoke a simple statistical norm; my limited background beliefs about squirrels tells me that grey

² Bruckmüller et al. (2012). This was the total number of individuals interviewed across three separate studies that helped to deliver the conclusions detailed in this section.

³ I remind the reader that, when referring to another's research, I stick with their terminology, instead of using my preferred conjunction *sex/gender*.

⁴ Bruckmüller et al. (2012: 210).

squirrels are more numerous in the UK, meaning that red squirrels are the group that needs explaining in virtue of their statistical abnormality.

The sort of cultural associations and beliefs about gender in the context of leadership, however, are more complex. Men do indeed represent the statistical norm for leadership—there are, after all, more male CEOs named *John* than women CEOs.⁵ This might go some way to explaining why, in Bruckmüller et al.’s pre-test, men were called to mind much more easily when participants were asked to imagine a leader.⁶ Men also represent the social norm of a leader, in the sense that there is a cultural stereotype associating leadership with men and masculine traits. The stereotype of a manager, consultant, professor, and so on, is man. Research suggests that the qualities that we tend to associate with leaders, and/or believe that leaders possess, are stereotypically masculine traits—rationality, assertiveness, lack of emotion.⁷

Crucially, stereotypes do not just describe—they commonly *prescribe*. Where one believes that the stereotype of a leader is a man, and specifically a man with masculine traits, this often goes hand in hand with believing that leaders *should* be men, and have masculine traits.⁸ This prescriptive side of male leader stereotypes goes some way to explaining various facts that I will return to in §5.2.2 below. For instance, studies show that *good* managers tend to get described largely by masculine attributes, and masculine attributes are often considered essential to be a successful leader.⁹ Feminine traits are considered a poor fit with leadership.¹⁰ This goes some way to explaining studies demonstrating that women who are leaders are often seen to be defective in some way; for instance, they tend to receive more negative evaluations of their performance than men who are leaders.¹¹ This would suggest that men are also the *normative* norm in the context of leadership, in the sense that they, in Bruckmüller et al.’s words, “set the standard of culturally valued behaviour [in that domain]”.¹² In other words, leadership represents an androcentric domain in our culture. Women are not just a statistical abnormality in the context of leadership, then, but a *normative* abnormality.

⁵ Miller et al. (2018). There are also many other top jobs where men named *John* (and sometimes *David*) outnumber women, as Miller et al.’s article discusses.

⁶ Bruckmüller et al. (2012: 213).

⁷ See *ibid.* (212) and the references therein. Also see Bruckmüller et al. (2013: 457).

⁸ Burgess & Borgida (1999), Terborg (1977), and Heilman (2001, 2012).

⁹ See Heilman (2001: 659-60) and the references therein.

¹⁰ *Ibid.* (660-1).

¹¹ Brescoll et al. (2010).

¹² Bruckmüller et al. (2012: 211).

By linguistically positioning *women* as the effect and *men* as the norm, cultural associations and beliefs that resonate with this positioning – namely, which *also* position women as the effect (i.e. the statistical and/or normative abnormality) and men as the (statistical and/or normative) norm *in our culture* – are pragmatically inferred by the audience. This helps to explain the results that we began with. The bundle of associations and beliefs associated with men being the norm for leaders in our culture involves recognising that men do indeed have more power and status in society. In virtue of being statistically much more likely to be leaders than women, men are much more likely to have the power and social status that being a leader involves. This helps to explain the first result: namely, the tendency to attribute more status and power to men when men are linguistically positioned as the norm. The cultural norm of men leaders also suggests that any such gendered power and status imbalances are *justified*; the stereotype of a leader suggests that leaders *should* be men, and therefore that the power and status benefits that leadership brings to men are warranted. This helps to explain the second result: namely, the tendency to think that existing gendered inequalities in status and power are *justified*, when men are linguistically positioned as the norm. Finally, the cultural norm of men leaders involves identifying leadership traits with stereotypically *masculine* traits. Women are assumed to be a poor fit for leadership in large part because they are assumed to possess stereotypically feminine traits—traits that the ideal leader should *not* possess. The relevance of gender stereotypes to the cultural norm of men leaders helps to explain the higher endorsement of gender stereotypes, when men are linguistically positioned as the norm.

Further, Bruckmüller and colleagues cite evidence that when we ourselves are the ones asking questions, or talking more generally, about contrasts between social groups, cultural associations and beliefs concerning those groups influence how we phrase our comparisons.¹³ In the case at hand, this would mean that people are more likely spontaneously to phrase a question about gender and leadership as one about how *women* lead differently to *men*, as opposed to the other way around. This goes for academics as well as laypeople. Consider, for instance, another form of androcentrism. There is a cultural convention that presents the male body as the norm, something we see, for instance, in portrayals of the human body that overwhelmingly depict the male anatomy. (For instance, in popular Western anatomy texts used in pedagogical contexts, one study found that, where there were sections dealing with non-sex-specific anatomy, male

¹³ See for instance, Miller et al. (1991), cited in Bruckmüller et al. (2012: 216).

subjects were depicted in 64% of illustrations in which a sex was discernable, whilst females appeared in only 11% of images.)¹⁴ Research demonstrates that, consistent with this cultural trend of depicting the male body as the norm, scientific publications are more likely linguistically to position maleness as the implicit norm when asking questions about sex/gender differences in biology and psychology.¹⁵ One is more likely to hear contrasts, for instance, about how *women* differ from *men* in average body size, or in mental rotation ability.

Other cultural associations and beliefs instead position women as the norm. The research cited here suggests that we would be more likely to ask how *men* differ from *women* when it comes to childcare, for instance. Partially, this tendency to choose linguistic framing that reflects cultural norms is down to our desire to facilitate easy communication. As Bruckmüller et al. put it, “Scientists [and others] usually strive to communicate findings in a way that is most easily understandable and relying on culturally shared habits of linguistic framing facilitates understanding”.¹⁶ The implications of choosing linguistic positioning of comparative terms that goes *against* the grain of cultural norms will be discussed in §5.4.

5.2.2. Analysis

The terminology from this thesis can be used to capture what is going on in Bruckmüller et al.’s study. The respective linguistic positioning of *men* and *women* counts as one of the presentational devices that constitute a linguistic salience perspective. Linguistic salience perspectives, remember, involved manipulating the presentation of linguistic contents in subtle ways so as to make certain features of that content more salient than others. This might involve changing the *order* in which that content is discussed, for instance. As for the two questions about gender and leadership, the only difference between them is indeed the *order* in which the words *men* and *women* appear in the sentence. The fact that this change in linguistic salience perspective produced different responses from study participants, then, demonstrates a framing effect.

This linguistic salience perspective led to a two-step (pragmatic) inferential process. Firstly, it activated a pragmatic inference to the *linguistic* norm suggesting that whichever group is positioned first in such a comparison is to be treated as the effect to

¹⁴ Giacomini et al. (1986). Sex-neutral representations made up the remaining 25%.

¹⁵ See Hegarty & Buechel (2006), discussed in Bruckmüller et al. (2012: 210).

¹⁶ Bruckmüller et al. (2012: 216).

be explained, and whichever group is last is the implicit norm for comparison.¹⁷ With this linguistic norm having been activated in the participants' minds, and with it having rendered men the implicit norm for comparison, and women the effect to be explained, pragmatic inferences to *cultural* norms that resonate with this structure were activated. The cultural, androcentric norm that paints male leaders and masculine traits as the ideal (i.e. the *normative* norm), and women leaders and feminine traits as abnormal and deviant became activated in this process. Again, this norm was readily inferred given its cognitively accessible and socially licensed nature. If we find this cultural norm problematic, we can find the linguistic salience perspective that led to it problematic. In other words, we can judge the linguistic salience perspective that made *women* salient to be instrumentally problematic on the basis of its upshots.

The inferential process described above did not occur on a conscious level; far from being explicitly described in the questions, these linguistic and cultural norms were communicated implicitly.¹⁸ As we have learnt from previous chapters, the fact that this inferential process bypassed conscious awareness plausibly helps to explain why study participants answering the *how do women differ from men?* question so consistently ended up endorsing the cultural associations and beliefs described above. Unable to monitor their inferences, these participants were likely unable to suppress those that they might, if given chance to deliberate, have rejected.

What is the significance of Bruckmüller and colleagues' study? In a rather general suggestion, the authors themselves warn, "the ways that we habitually frame group differences appear to be part of power processes that reproduce stereotypes about differences between real groups that differ in status".¹⁹ What, though, is so bad about this stereotype reproduction? Whilst the authors do not go into detail, we can make some more specific suggestions on the basis of what we have already learnt about the action-engendering nature of stereotypes.

Activating inferences to stereotypes can mean also activating various forms of behaviour that are consistent with these stereotypes.²⁰ For instance, research suggests that, thanks to the stereotype that leaders are men, we are more likely to hire a man as a

¹⁷ Bruckmüller et al. (2012).

¹⁸ See in particular the discussion in chapter 2 (§2.3.1).

¹⁹ Bruckmüller et al. (2012: 216).

²⁰ Tirrell (2013: 165). See also Wheeler & Petty (2001: 820) for a discussion of how stereotypes activate behavioural codes associated with them.

manager.²¹ Research has shown that men are also more likely to get promoted over women, and receive larger bonuses and variable pay, even when their performance evaluations are identical.²²

Further, research demonstrates that where women display assertiveness, a trait stereotypically associated with leadership, and a trait that Laurie Rudman and Julie Phelan describe as “necessary for success in the business world”, they tend to get viewed and evaluated more negatively, and this comes with significant financial setbacks for women leaders.²³ It has been suggested that these factors combine to explain why fewer women put in for leadership roles.²⁴ Not only do stereotypes tell women that they do not fit the bill of a leader, but the prescriptive side of these stereotypes means that they will likely be penalised if they *do* try to lead, which dissuades them from pursuing leadership roles.²⁵ Where women *do* pursue leadership roles, research suggests that negative women leader stereotypes can engender stereotype threat effects, affecting how well women leaders perform. For instance, research conducted on American adults found that when negative woman leader stereotypes are implicitly activated (i.e. where feminine-coded traits, such as empathy, are linked to poor negotiation outcomes), women tend to underperform men in leadership tasks.²⁶

This brief outline of some of the issues that women face due to gender stereotypes about leadership vividly demonstrate ways in which implicitly activating a bias, such as through using a certain linguistic salience perspective, can ultimately play a role in entrenching existing power and status inequalities between men and women. It is not just our conscious discounting of the bias itself that is precluded by the implicit nature of its activation. Its associated codes for behaviour are also not consciously interrogated, and therefore not suppressed.

The lesson to be taken from Bruckmüller et al.’s study is that we should pay careful attention to which linguistic salience perspectives we employ to discuss gender differences; making women salient in gender comparisons in androcentric domains, such

²¹ Heilman (2012, 2001).

²² Castilla (2008).

²³ See Rudman & Phelan (2008: 65-6) and the references therein. See also Heilman & Okimoto (2007), and Lyness & Heilman (2006).

²⁴ Hoyt & Murphy (2016).

²⁵ Ibid. (388).

²⁶ Kray et al. (2001). As mentioned in chapter 2 (§2.3.1), this can generate epistemic issues, insofar as the stereotype ends up with the wrong direction of fit with the world. The effect of stereotype threat here can look retroactively to confirm the stereotype that women are inferior leaders.

as those connected to leadership, can help to activate inferences to harmful, action-engendering cultural associations and beliefs (and their associated behaviours).²⁷ Bruckmüller et al.'s study gives us one reason to avoid employing salience perspectives that make women salient in the context of sex/gender comparisons in androcentric domains.

5.3. Introducing sex/gender essentialism

With this evidence in mind, of a mere salience perspective activating problematic androcentric biases about sex/gender, I turn to the particular project of this chapter. Instead of looking at androcentric biases, I will instead be focussing on a different sex/gender bias—psychological essentialism. The hope is that, just as Bruckmüller and colleagues examined cognitively accessible and socially licensed cultural norms and stereotypes about leaders in order to reach their conclusion about which salience perspectives to utilise for sex/gender comparisons (they suggested that our androcentric biases counsel against making *women* salient in gender comparisons for leadership), this survey of our *essentialist* bias will illuminate which salience perspectives for sex/gender are to be avoided. I will be relatively brief in my discussion of the sex/gender essentialist bias, but refer the reader to the references for a deeper examination of the (extensive) psychological evidence.

5.3.1. Sex/gender essentialism: the beliefs

The previous chapter suggested that psychological essentialism involves the belief that certain categories have an underlying, internal, and fixed nature – an *essence* – that gives category members their kind identity. Precisely how psychological essentialism functions when applied to different categories (such as race, class, and sexual orientation), however, varies.²⁸ When it comes to sex/gender, sexed/gendered categories tend to be seen as akin to natural kinds, in the sense that they are considered to be *species-like*—at least according to the folk understanding of species. This means thinking that there are different sexes/genders, and that their differences emanate from their distinct, fixed,

²⁷ My focus on harm in this thesis means that I have not addressed how these cultural associations and beliefs might be flawed on epistemic grounds. See Blum (2004) and Puddifoot (2017) for general discussions of how stereotypes in particular can be critiqued on such grounds.

²⁸ Medin & Ortony (1989), Haslam and Whelan (2008).

biological natures.²⁹ In fact, in a study on American adults' representations of 40 social categories, it was found that sex/gender categories are the most robustly and potently essentialised in this way—more so than other social categories, such as race and sexual orientation.³⁰

Below, I discuss the various beliefs associated with sex/gender essentialism that have been proposed in psychological research, alongside the cultural narratives and practices, often also reflected in sex/gender science, that reinforce them. These are of course not the only cultural messages that exist about sex/gender. In addition to the essentialist messages detailed in what follows, we also come across some anti-essentialist messages in our culture, such as *girls can do anything boys can do*. Whilst these latter sorts of messages do exist, it is, as we will see, essentialist conceptions of sex/gender that are particularly prevalent and mainstream.

Following Sarah-Jane Leslie, Susan Gelman and Marianne Taylor, I suggest that we treat essentialism as a cluster concept; one can be more or less essentialist depending on how many of the various beliefs, detailed below, that one holds.³¹ An individual who holds essentialist beliefs is not always conscious of holding them. Instead, these beliefs are often held only *implicitly*. As mentioned in the previous chapter, this means that, instead of directly asking a person if they hold a given essentialist belief, indirect measures can be required.³²

Further, it should be borne in mind in what follows that there is little research into the cultural variation of psychological essentialism about sex/gender in particular;³³ the vast majority of research discussed below has been carried out on North American and European populations. The cultural narratives that I discuss have therefore also been drawn from these Western demographics.

5.3.1.1. *Two opposing categories*

At least in Western cultures, we tend to think in terms of a *woman/man* binary. For reasons that will become clearer in §5.3.1.3, these gender terms are treated as synonymous with the sex terms of *female* and *male*, which are also treated as a binary. Further, the differences that these alternative kind essences give rise to are considered to

²⁹ Rothbart and Taylor (1992), Bastian and Haslam (2006: 229).

³⁰ Haslam et al. (2000).

³¹ Leslie (2013: 116), Gelman & Taylor (2000).

³² See back to chapter 4 (§4.3.1).

³³ A notable exception is Mahalingam (2003).

be extreme. For instance, Carol Martin and colleagues found that the American children they studied tended to use their knowledge about what girls do to predict the *opposite* for boys.³⁴ This goes for adult folk theories, too; for instance, many comment on a tendency in scientific publications to exaggerate sex/gender differences, suggesting that this is symptomatic of how essentialist beliefs continue into adulthood.³⁵

Cultural narratives reinforce these messages. Infamous phrases, such as *the opposite sex*, and *men are from Mars, women are from Venus*, help us to think that there are just two genders, and that they are as different as chalk and cheese.³⁶ These ideas are materially embodied around us, whether it is in toy shops that are separated into two aisles of pink vs. blue, or in different dress codes (dresses vs. trousers) and hairstyles (long flowing locks vs. short or shaved).³⁷ When it comes to sex science, the trend to talk as though there are indeed just two sexes, which helps to make various intersex conditions invisible, is overwhelming.³⁸ Further, in addition to the trend of *exaggerating* sex/gender differences mentioned above, Janet-Shibley Hyde has shown that sex/gender *similarities* are rarely investigated; psychological research has an obsession with documenting sex/gender *differences*.³⁹

5.3.1.2. Discrete categories, internally homogeneous

Related to the binary above, the two categories of *women* and *men* are often treated as discrete—as having sharp category boundaries, with no overlap. Studies of American children indicate an assumption that a woman, for instance, cannot have gender roles stereotypically associated with two gender categories; she cannot, for example, be both a mother and doctor, or a woman and a firefighter.⁴⁰ The *normative* counterpart of this belief tends to remain into adulthood—whilst adults tend to accept that some men and

³⁴ Martin et al. (1995).

³⁵ See, for instance, Hyde (2005), Hyde et al. (1990), Caplan, MacPherson & Tobin (1985).

³⁶ One of the best-selling nonfiction hard-back books of all time is called *Men are from Mars, Women are from Venus* (Gray, 1992), which runs with spurious popular science claims that men and women are virtually different species. (One might also wonder whether this phrase plays into androcentric biases, in virtue of mentioning men *before* women.)

³⁷ Gelman & Taylor (2000).

³⁸ Fausto-Sterling (1993).

³⁹ Hyde (2005). See also Rippon et al. (2014).

⁴⁰ See the studies cited in Gelman & Taylor (2000). Some children deny that such individuals (e.g. a female firefighter) exist, whilst others misremember them (e.g. as a male firefighter).

women *do* have traits that overlap gender boundaries, there is a prevalent belief that this *should not* be the case. For instance, a recent UK study found that 49% of the men interviewed believed there was a social stigma associated with taking shared parental leave, given that childcare is considered a *feminine* job.⁴¹ This normative belief in discreteness also applies to sexed biology; research tends to find that intersex individuals are highly stigmatised in society.⁴²

When it comes to thinking about the members of a given sex/gender group, the sex/gender essentialist believes that within each sex/gender group, members are considered to be largely homogeneous. For instance, Martin and colleagues found that the American children in their study tend to presume that “what a person of one sex likes, other people of the same sex also will like”.⁴³ The fact that Kiera likes dolls is taken as a good indicator that other girls will have the same preference, for instance. Research conducted on American undergraduate students indicates that this trend persists into adulthood; these students tended to presume that same-sex peers were much more likely than other-sex peers to share their own liking of a given object, such as a sculpture.⁴⁴

The idea that women/females and men/males form two internally homogeneous but discrete categories is particularly strong in commercial culture. It is hard to find a shop online or on the high street that doesn’t have a *gifts for her* vs. *gifts for him* section, for instance—a practice that presumes men generally like the same things as other men, that women like the same things as other women, and that these preferences are non-overlapping. *Men are from Mars*, and *women are from Venus*, after all! Sex/gender science has also been criticised for neglecting both within-sex variability, and between-sex overlap.⁴⁵

5.3.1.3. *Biological determinism*

Studies indicate that the influence of the environment on sex/gender development is often either downplayed or disregarded altogether by the sex/gender essentialist. Marianne Taylor’s study on American children under the age of nine demonstrated that these children “viewed gender-stereotyped properties as relatively impervious to environmental influences”.⁴⁶ For instance, many of her participants thought that a boy

⁴¹ TSB (2018). For other examples of this trend, see Brescoll et al. (2010), as well as Bruckmüller et al.’s (2012) study, mentioned earlier.

⁴² Fausto-Sterling (2000: 8), Preves (1999).

⁴³ Martin, Eisenbud & Rose (1995: 1468).

⁴⁴ Ibid (1460).

⁴⁵ Nelson (2016), Vanwesenbeeck (2009), Hyde (2005).

⁴⁶ Taylor (1996: 1568).

raised only with girls and women, surrounded by feminine-stereotyped games of dolls and princesses, will inevitably end up developing masculine-stereotyped properties such as an interest in football. When asked to explain why a boy raised with girls would nevertheless develop masculine stereotyped properties, these children tend to refer to biological properties, such as *brains* and *instinct*.⁴⁷ This reflects the idea of *innate potential* discussed in the previous chapter, which referred to the idea that an organism's development is entirely pre-specified by the organism's in-born (in this case, biological) essence. Studies suggest that similar beliefs have been found alive and well in adults. In one study of American young adults, for instance, it was found that study participants more frequently suggested that a boy raised by women would nevertheless develop masculine-stereotyped behavioural traits.⁴⁸

Turning to popular culture in Western countries, we see a clear emphasis on *biological* properties when explaining sexed/gendered traits.⁴⁹ It is not uncommon in newspapers and magazines to hear the claim that men's brains are hardwired differently to women, a finding that explains why men are more logical, better at reading maps, unempathetic, and so on.⁵⁰ Pick up a self-professed *women's magazine* like *Cosmopolitan*, and find the claim that men are messy and unempathetic because it's been written in their genes since they were cavemen.⁵¹ These ideas are not confined to popular culture. It is common to find scientific studies discussing hardwired brain differences between the sexes/genders, as well as fixed differences in genetic code, and for these differences to be understood as causing sexed/gendered differences in behaviour.⁵²

⁴⁷ Taylor & Gelman (1999).

⁴⁸ Eidson & Coley (2014). Some earlier studies had suggested that adults instead consider environmental properties as having more causal power over our sexed/gendered traits (Taylor, 1996). Eidson and Coley, however, have found these studies to be flawed, given their use of explicit self-report measures. I mentioned earlier that many of our essentialist beliefs are implicit (i.e. the individual with them is unconscious of having them), meaning that implicit measures are required to uncover them.

⁴⁹ See, for instance, Arthurs & Zacharias (2006), and O'Connor & Joffe (2014).

⁵⁰ See, for instance, *The Guardian's* 2013 article *Male and Female Brains Wired Differently, Scans Reveal* (Sample, 2013), and Langton and Dupré's (2013) response.

⁵¹ Hasinoff (2009).

⁵² For critical discussions of such studies, see e.g. Rippon et al. (2014), Fine (2010), and Dupré (2001). Rippon et al. (2014: 2), for instance, talk about a trend in sex science to take "a person's biological sex [to be] a good proxy for gendered behavior".

5.3.1.4. *Immutability*

The results of the study mentioned earlier, where children judged that a boy's masculine traits would remain constant over significant changes in his environment (such as where he is brought up exclusively with women and girls), led Taylor to suggest that American children "view gender categories as relatively immutable".⁵³ Whilst it is hard to find research explicitly testing for beliefs that gender categories are immutable in adults, the fact that it is so common to come across narratives of *hardwired* brains and *fixed* genetic coding in Western cultures *and in scientific studies* of sex/gender (see §5.3.1.3 above) suggests that sex/gender is regularly interpreted as fixed by this demographic.

Generally, where a sex/gender essence is presumed to be *biological* (as it regularly is, according to the discussion in §5.3.1.3), ideas of immutability follow close behind. Western cultural conceptions of biology tend to treat it as fixed and isolated from social influence.⁵⁴ This idea can also be seen in countless cultural artefacts around us. Consider, for instance, the *women's magazines* mentioned earlier. After suggesting that men's behaviour is down to their genetics, one representative article of *Cosmopolitan* implores "Look, the dude's not going to change his weirdo ways, but at least now you have a clue".⁵⁵ We also see the ideas of fixity in social responses to trans identities. Rachel McKinnon notes how trans women regularly come up against the sex/gender essentialist trope that they "are still, and forever will be, men, no matter how indistinguishable they may be from cisgender women".⁵⁶

5.3.1.5. *High inductive potential*

Because of this causally potent essence, sex/gender categories are treated as having high inductive potential, in the sense that once you learn a person's sex/gender, you can infer a great many other facts about them.⁵⁷ On their study on American children, Thomas Berndt and Kirby Heller found that those who had learnt gender stereotypes (i.e., above kindergarten age) tended happily to infer a person's toy preferences solely on the basis of

⁵³ Taylor (1996: 1568). See Bem (1989) for similar results. Some studies have suggested that children do not consider sex/gender identity to be stable. American preschool children, for instance, have been found to assume that a boy who wears a dress will become a girl, for instance (Marcus & Overton, 1978). Other studies have questioned results such as these, however, finding flaws in their methodology (Ruble & Martin, 1998).

⁵⁴ Keller (2005), Fine (2010), Jordan-Young (2010).

⁵⁵ Bodnar, in Hasinoff (2009).

⁵⁶ McKinnon (2014: 859).

⁵⁷ Gelman and Taylor (2000), Gelman (2003), Rothbart and Taylor (1992).

their perceived gender category.⁵⁸ For instance, they tended consistently to suggest that a boy, of whom they know nothing other than that he is a boy, would rather fix a bicycle than bake brownies. Kay Deaux and Laurie Lewis found that similar results extend to American adults.⁵⁹ When no other information about a person is given beyond whether they are female or male, these adult participants nevertheless tended to make a wide range of gender-stereotypical inferences about that person's physical characteristics (e.g. strong vs. slender), role behaviours (e.g. head of household, vs. child-carer), occupations (e.g. doctor vs. nurse), and sexual orientation (e.g. heterosexual, homosexual).

The sheer salience of sex/gender in our culture reinforces the impression that it *is* indeed a treasure trove of inductive inferences.⁶⁰ Whether it is the centrality that gendered personal pronouns (e.g. *he/she*) play in our language,⁶¹ administrative forms requiring knowledge of our sex/gender, the reliability with which new parents are asked whether they are having a boy or a girl before anything else, the regularity that shop assistants ask whether one is buying for a boy or a girl when their advice is solicited on gift purchases, Western societies, in McKinnon's words, "are structured such that one's gender *identity* is always relevant".⁶² This gives the impression that sex/gender is a deeply important cultural vector—one from which we *should* be making lots of inferences. The fact that the variable of *sex/gender* is so regularly tested in science without a good rationale as to whether it is indeed relevant, suggests that the belief that gender has high inductive potential seeps into science.⁶³

5.3.2. Sex/Gender Essentialism: Harms

Sex/gender essentialism has grave problems from an epistemic standpoint. Generally, the belief that there are biological gendered essences is largely rejected by contemporary

⁵⁸ Berndt and Heller (1986).

⁵⁹ Deaux & Lewis (1984). See Prentice & Miller (2006) for similar findings.

⁶⁰ Gelman & Taylor (2000: 186), Bem (1993), Freed (1996).

⁶¹ Frye (1983) famously draws attention to this issue, referring to this grammatical phenomenon as *sex marking*.

⁶² McKinnon (2014: 865-6).

⁶³ For instance, the research from Hyde (2005) that was mentioned above, which suggested that psychological research is obsessed with documenting sex/gender differences in human psychology, even when the evidence is that these differences are usually either non-existent or small, speaks to this point. Scientists so regularly test for differences in psychological traits using the variable of gender, whether or not there is good reason to do so, and whether or not there is good reason to test for *this variable in particular* over others (such as race, class, occupation, and so on). It is important to flag that an androcentric bias also pulls the other way, however, insofar as sex/gender differences are not investigated where they should be (Heidari et al., 2016).

science.⁶⁴ In fact, contemporary science shows how each constituent belief of sex/gender essentialism, as detailed above, has serious problems from an epistemic standpoint.⁶⁵ Given that the focus in this thesis is on harm, however, I concentrate in what follows on the *ethical* repercussions of these false beliefs.

A whole host of harms associated with sex/gender essentialist beliefs has been found. For instance, these beliefs have been shown to generate stereotype threat for disadvantaged sex/gender groups, such as women and trans men and women.⁶⁶ Trans men and women also face a range of other harms due to sex/gender essentialist beliefs, including social stigmatisation. We heard earlier that the line between a descriptive and normative belief is blurred for many of the components of sex/gender essentialism. Sometimes, the belief that there *is* a binary of women/females and men/males becomes a normative belief that everyone *should* fit into this binary. This is associated with a belief that trans identities are deviant and punishable. It is this essentialist belief that has been suggested to help account for the terrible statistic that more than a third of trans people in the UK suffered hate crimes in 2017, for instance.⁶⁷

Here, though, I will focus on just one harm of sex/gender essentialism—one that I mentioned in the previous chapter, and that I will return to in chapter 6. There is an abundance of research into how false sex/gender essentialist beliefs are regularly used to justify, and therefore *maintain*, the oppression of disadvantaged sex/gender groups.⁶⁸ Take women, for instance. The thought goes something like this: it is in women's nature to be carers, and it goes against their nature to have power and authority; what is natural is good; therefore, women should stay at home and be discouraged from holding positions of power. Studies have found that essentialist explanations that follow this structure are often given by couples to justify a sexist division of labour in the household.

⁶⁴ See, for instance, Fine (2010), Leslie (2013), Dupré (1986), Jordan-Young (2010), Richardson (2013), and Rippon et al. (2014). For more general criticisms of the folk understanding of biology, see Griffiths (2002) and Medin & Atran (1990).

⁶⁵ In addition to the references in the footnote above, see Hyde (2005), Fausto-Sterling (2000), O'Connor (2010), Richards, Bouman & Barker (2017), Hyde & Mezulis (2002: 556), Mikkola (2017), Feng et al. (2007), and Haier et al. (2009).

⁶⁶ See Heyman & Giles (2006), who detail some of the issues for girls and cis-women (namely, women who have female genitalia and whose chosen social identity is *woman*). See McKinnon (2014) for examples of the problems of sex/gender essentialist stereotype threat for trans women.

⁶⁷ Butcher (2018). See Fausto-Sterling (2000) and Minto et al. (2003) for a discussion of how these normative beliefs have negatively affected intersexuals.

⁶⁸ Skewes et al. (2018), Bem (1993), Fausto-Sterling (1985), Fuss (1989), Eagly et al. (2000), Hoffman & Hurst (1990), Yzerbyt et al. (1998).

For instance, in his interviews with 20 American dual earner couples, Scott Coltrane found that couples for whom the woman was primarily responsible for the household and childcare were likely to invoke the essentialist belief that “maternal instincts”, and *natural* differences between the sexes/genders (citing, for instance, the mother’s genes setting her up to be nurturing) justified their domestic arrangement.⁶⁹ By contrast, couples who adopted a more equal division of household labour and childcare tended to cite a belief that the sexes/genders were similar. Other studies corroborate these findings. Ruth Gaunt interviewed 209 Israeli couples, finding that “the more the father believed that women have an essence that makes them more suitable for child care, the more hours the mother spent as the sole care provider for the child”.⁷⁰

This arguably oppresses women in at least two ways. (Here, I borrow from Iris Marion Young’s influential account of oppression, in which she takes it to include “five faces”: exploitation, powerlessness, marginalization, cultural domination, and violence.)⁷¹ Firstly, it can be said to *exploit* them, insofar as women unfairly perform the majority of domestic labour and childcare.⁷² This is widely recognised to be culturally undervalued work, for which they are not remunerated, and from which men benefit.⁷³ Secondly, it arguably helps to make women powerless. Essentialist beliefs often trap women in domestic roles, the burdens of which restrict their options regarding jobs and activities outside of the home. This limiting of options is itself a form of powerlessness, but it also has knock on effects in terms of women’s ability to get *positions* of power in society (such as leadership positions); the burdens of domestic activities affect their ability to progress in the workplace.

Essentialist beliefs have been found to justify oppression outside of the home, too. In Thomas Morton and colleagues’ study involving 552 adults across the globe, half of study participants read an article arguing that sex/gender differences in thought and behaviour were down to biological differences (read: essences) between females and males, whilst the other half read the same article alongside another that refuted each of

⁶⁹ Coltrane (1989).

⁷⁰ Gaunt (2006: 532).

⁷¹ Young (1990).

⁷² A 2016 ONS (cited in Burkeman, 2018) study found that UK women do 60% more of the housework than men.

⁷³ For an argument that domestic work is culturally undervalued, see NWLC (2017).

the first paper's arguments.⁷⁴ Men reading the first paper only were more likely to agree to various statements saying that they would support discriminatory practices that disadvantaged women (such as *If I would work in a company where my manager preferred hiring men to women, I would privately support him*, and *If I were a manager, I would believe that more often than not, promoting men is a better investment in the future of the company than promoting women*). These examples suggest that gender essentialist beliefs move beyond an intellectual justification of sex/gender inequalities in the status quo; the action-engendering consequences of these beliefs mean that they, in Haslam's words, "promote the *maintenance* of current social divisions" [my emphasis].⁷⁵ These social divisions are oppressive to women in various ways, not least insofar as they help to make women powerless; in addition to restricting their options more generally, they also restrict their access to positions of power in particular.

Essentialist beliefs have also been said to oppress other disadvantaged sex/gender groups, namely, those that are non-binary. The belief that there is, or should be, a sex/gender binary is oppressive to non-binary groups in virtue of *marginalising* them.⁷⁶ (Marginalisation, as we noted above, is one face of oppression, according to Young.) The descriptive counterpart of this belief can help to make non-binary groups *invisible*, by denying their existence. The normative counterpart of this belief accepts these groups' existences, but stigmatises them as deviant. Both beliefs help to relegate non-binary groups to the edges of society. This is marginalising, in the sense that these groups are less heard, less represented, less respected, less able to influence decision-making, and so on.⁷⁷

It is widely recognised, then, that false essentialist beliefs are regularly used to justify various practices that oppress disadvantaged sex/gender groups.

⁷⁴ Morton et al. (2009). See Kray et al. (2017) for a similar study, instead showing how exposure to essentialist ideas that gender roles are fixed increases support for gendered inequalities society.

⁷⁵ Haslam (2011). Also see Keller (2005), Yzerbyt, Rocher, & Schadron (1997) and Fuss (1989).

⁷⁶ Bettcher (2016), Preves (1999).

⁷⁷ Honneth (1995) for instance, talks about social recognition being a vital part of self-respect, and therefore that social marginalisation constitutes a harm. See also Darwall (1977) and Bartky (1990).

5.4. Choosing Salience Perspectives for Sex/Gender

5.4.1. Salience perspectives to avoid: Instrumental harms

The discussion in §5.3.1 suggests that sex/gender essentialism is cognitively accessible and socially licensed, at least in Western cultures. This fact helps to explain research indicating that it takes very little for our existing essentialist biases to become activated. For instance, Rebecca Bigler's study on American school children demonstrated that where teachers gave occasional instructions that mention gender, such as *all the girls put their bubble makers in the air*, the children in that classroom were more likely to endorse essentialist sex/gender stereotypes than those in classrooms without such gendered instructions.⁷⁸

More minimal still, Susan Gelman and Gail Heyman's study of 115 American children found that simply conveying information in noun form could trigger their essentialist thinking.⁷⁹ These children were told about a hypothetical person, Rose, who liked carrots. When this information was conveyed in noun-form, i.e. *Rose is a carrot-eater*, it generated more essentialist thinking about Rose than when it was conveyed in verb-form, i.e. *Rose eats a lot of carrots when she can*. In particular, hearing the information in noun-form led children to judge that Rose's predilection for carrots would be significantly more stable over time and over different situations. There is evidence, then, that we are used to inferring substantive essentialist beliefs (i.e. that the genders are essentially different, or that our behaviours are immutable) from the subtlest of inputs. For instance, the fact that a change in the grammatical structure of a sentence would no doubt be something that bypasses an audience's conscious awareness plausibly makes sense of the result above. We have already heard about how language that has an implicit, under-the-radar quality can be especially effective in shaping an audience's response to the linguistic content in question. Where an inference is made below an individual's conscious awareness, it is especially likely to go on to shape that individual's thought and behavioural patterns. This is because the individual cannot consciously reflect on whether they do indeed wish to endorse that inference, which prevents them from blocking inferences that they might otherwise wish to reject.

Salience perspectives count as another more minimal, implicit input from which we are liable to infer essentialist beliefs. As we have learnt from the previous chapter, as

⁷⁸ Bigler (1995).

⁷⁹ Gelman & Heyman (1999).

well as from Bruckmüller's study from this chapter, we should be wary of using salience perspectives that make salient content that is central to our biases; these biases are already socially licensed and cognitively accessible, and so simply having contents central to them at the forefront of one's mind (perhaps through first having them made salient in language) can be sufficient to activate inference to those biases' *beliefs* about those contents. Further, the activation of these biases would most likely occur in an under-the-radar manner, which can give salience perspectives the sort of special power mentioned above.⁸⁰

This chapter suggests that, simply by making the following contents *salient*, we risk activating our sex/gender essentialist beliefs, and their associated practices of oppression.

Firstly, *differences* between *women/females* and *men/males* over their similarities: this mere pattern of salience might result in activating problematic essentialist beliefs, such as that women/females and males/men are overwhelmingly different on most dimensions. It might also generate an inference to a different essentialist belief. We heard in the previous chapter that an emphasis on the distinctive traits a given subject happens to be manifesting can activate the essentialist belief that those traits are *fixed* (this essentialist belief was discussed in §5.3.1.4). We can arguably see this inferential pattern in cultural discussions of sex/gender differences. Consider, for instance, how media reports that women/females and men/males have *distinctive* brains, or even simply *distinctive* behavioural traits, regularly elicits the conclusion that those differences are *hardwired*, or otherwise immutable.⁸¹ We might wonder, then, whether giving relative salience to descriptions of the distinctive traits that women/females and men/males happen to exhibit, over the potential for their sexed/gendered traits to *change*, might elicit inferences to sex/gender essentialist beliefs that any existing differences are *fixed*.

Secondly, the sex/gender binary over the larger sex/gender space: this salience pattern risks activating the belief that the sex/gender binary either does or *should* exhaust the sex/gender space, so that alternative categories (such as intersex, queer) are rendered either invisible or deviant.

⁸⁰ See chapter 2 (§2.2.2 and §2.3.1), and chapter 3 (§3.3.1) for a discussion of the under-the-radar nature of salience perspectives and their activation of biases.

⁸¹ See, for instance, Maney's (2016) paper reviewing the *Perils and Pitfalls of Reporting Sex Differences*. See Sample (2013) for an example of this poor reporting. Also see Greenberg et al (2018) for an example of the inference from finding distinctive behavioural traits (e.g. empathy vs. an interest in systems) in women/ females and men/males in particular to the conclusion that those differences are fixed.

Thirdly, ways in which the sex/gender categories are *discrete* over ways in which they overlap: giving relative salience to the former over the latter risks activating the problematic belief that an individual either cannot or *should* not have a combination of female(/feminine) and male(/masculine) traits.

Fourthly, ways in which members of the same sex/gender category are similar over ways in which they are diverse: this pattern of salience might activate problematic beliefs, including that there either is or should be little diversity *within* each sex/gender category.

Fifthly, biological insides of sexed/gendered individuals over their social environments: this salience pattern may help us to infer biologically determinist beliefs that treat our behaviour and traits qua sexed/gendered individuals as entirely caused by sexed/gendered biological essences (such as genes, hormones, or brains).

Finally, we should be aware of the fact that simply making sex/gender salient in contexts where it is not strictly relevant can activate essentialist beliefs about sex/gender.⁸² This issue is connected less to cognitive salience perspectives, and more to a different type of salience. Instead of giving *relatively* less salience to some contents over another, the suggestion here is that we should *omit* mention of sex/gender where its mention is not strictly relevant. This is because our essentialist biases mean that we tend to imbue sex/gender categories with great inductive potential, and that mere mention of a person's sex/gender can activate a great host of essentialist inferences about them.

The research so far suggests that emphasising any one of these contents (differences, binaries, and so on) may be sufficient to activate an essentialist belief about that content. Further, epistemic and social problems have been identified with each component belief.⁸³ This means that it is not just when a salience perspective activates all five beliefs that problems can occur. Remember, however, that essentialism is best thought of a cluster concept; one is more or less essentialist depending on how many of essentialism's constituent beliefs one holds. A text, or set of attentional dispositions, then, that makes *all* of these contents salient risks activating the strongest form of sex/gender essentialism, by activating inferences to *all* of its component beliefs. Where a

⁸² What counts as relevant will of course be quite a task to cash out. We can think of examples that lend sense to this idea, however. For instance, it could be helpful to remove sex/gender boxes for application forms for jobs that do not require a particular sex/gender; removing a person's sex/gender would help to prevent the assessor making lots of implicit inferences about their skills and traits, for instance. (See the related advice on guarding against implicit bias, such as that discussed in Chapter 1 §4.)

⁸³ The references in §5.3.1 demonstrate this.

salience perspective is likely to activate these beliefs, we can deem it instrumentally problematic—on the basis of the epistemic and/or moral problems with its upshots.

This point, namely, that different salience perspectives might activate varying degrees of sex/gender essentialism, helps us to see how the substance salience perspective fits into the discussion. We heard in the previous chapter that the substance salience perspective involved making salient the internal properties of a subject, and the distinctive traits it happens to be displaying. These contents are indeed central to the constituent beliefs of sex/gender essentialism. There is a focus on internal (specifically biological) properties of women/females and men/males, which were considered to be relatively impervious from environmental influence, as well as on the distinctive traits that different sexes/genders happen to display (there is a focus on the actual traits that make women/females and men/males *different*). This gives us reason to suspect that where we use a substance salience perspective in either our writing about sex/gender, or in our mind, we risk inducing a form (though not the strongest form) of sex/gender essentialism.

The list above goes *beyond* the contents highlighted by the substance salience perspective, however. We are warned against making salient contents such as the *binary* categories of women/female vs. men/male, and the ways in which members of a given sex/gender group are similar. This demonstrates the importance of considering subject-specific biases when making judgements about the likely impact of generic salience perspectives. The peculiarities of *sex/gender* essentialism mean that we must consider the impact of making a wider range of contents salient than those central to the substance salience perspective.

Echoing qualifications made in previous chapters, the suggestion here is not that we are *never* permitted to make salient the contents listed above in the context of sex/gender. There might be occasions, perhaps in a discussion of how certain medicines tend to affect differently sexed/gendered bodies in different ways, that require giving relatively more salience to internal, biological differences between women/females and men/males.⁸⁴ The point is rather that there is a risk that by doing so we activate our essentialist beliefs and their associated problems, and that this risk must be weighed into consideration. Is adopting this salience perspective worth the potential costs?

⁸⁴ Even here, though, there may sometimes be cases where giving relative salience to biological sex/gender differences over similarities is flawed and harmful (Risberg et al., 2009).

5.4.2. The novelty of this proposal

The suggestions discussed above might well sound familiar to some readers. Many feminist philosophers, especially those writing about sex/gender science, have given advice that sounds similar to that given here. *Be wary of studies that only look for sex/gender differences; look out for research treating members of the same gender group as homogeneous; be sceptical of studies painting our biology as fixed*: these are all lessons that someone acquainted with this literature has heard before.⁸⁵ The novelty of the suggestion being made here is that salience perspectives give us a particularly subtle way in which the communication of linguistic content concerning sex/gender, or attentional dispositions for this content or for sexed/gendered persons, might perpetuate bias. It is not just language that *avoids* talking about sex/gender similarities,⁸⁶ or language that uses terms like *hardwired* brains, that can lead the audience to essentialist conceptions of sex/gender. Language that explicitly discusses sex/gender similarities, or that explicitly mentions brain plasticity, can nevertheless activate problematic sex/gender biases. It can do this simply by giving less salience to these features of sex/gender. Bruckmüller et al.'s study should remind us to take the suggestion that mere salience can activate sex/gender biases seriously; this study has already demonstrated that where a cognitively accessible and socially licensed bias concerning sex/gender exists, it can take the subtlest patterns of salience, in language and attention, to activate it. Further, I have many times reiterated the point that subtle, and thus mostly unconsciously processed, mechanisms can have a special power in shaping our inferential patterns.

5.4.3. Salience perspectives to choose: Instrumental benefits

Do men lead differently than women in boardrooms? Bruckmüller and colleagues, in the paper mentioned at the beginning, refer to this as “an unconventional framing” (henceforth *salience perspective*) for a question about sex/gender and leadership.⁸⁷ It sounds odd to us, in our culture, for the reasons expressed in §5.2; we are used to men being treated as the implicit norm for androcentric domains. Because of this, Bruckmüller et al. commented on our tendency spontaneously to phrase gender comparisons in androcentric domains by asking how *women* compare to *men*.

⁸⁵ See, for instance, Rippon et al. (2014) and Maney (2016).

⁸⁶ This would instead count as a *sin of omission* (see the discussion of Anderson in chapter 2, §2.4.2).

⁸⁷ Bruckmüller et al. (2012: 212).

What are of interest to us here, however, are the ethical (and epistemic) gains that can be made through selecting *unconventional* (linguistic and cognitive) salience perspectives for sex/gender. When Bruckmüller et al.'s question about gender and leadership was phrased as men differing from women, study participants attributed less status and power to men, perceived any gendered status and power inequalities as significantly less legitimate, and attributed fewer stereotypic traits to men and women. True, they did not find women to have more status and power than men, despite women being positioned as the linguistic norm; the strength of the cultural stereotype of men as leaders is too strong for that. Successfully lessening the *degree* of status and power attributed to men remains a noteworthy effect, however, especially considering that this is the result of a particularly small linguistic alteration.

Why did this happen? We've heard in previous chapters that adopting salience perspectives that make salient the opposite content to that which is central to our biases often helps to prevent the activation of those biases. Bruckmüller and colleagues suggest a reason as to why this is the case. When it comes to linguistic salience perspectives, adopting those that are unconventional functions to impair communication, as the audience spends time trying to understand why the speaker used an unconventional salience perspective for their utterance. This can result in the arousal of the audience's cognitive attention.⁸⁸ The results of Bruckmüller et al.'s unconventional linguistic salience perspective for the leadership question demonstrate the benefits of this arousal. In their words, disrupting the usual conversational flow in this way was able to disrupt "the subtle reinforcement of beliefs about gendered status and power [and gender stereotypes] resonating with a conventional [salience perspective]".⁸⁹ They continue, suggesting "an unconventional [salience perspective] might lead participants to examine these status differences more critically and to question their legitimacy".⁹⁰

Whilst Bruckmüller et al.'s choice of words here might imply that unconventional salience perspectives move their audiences *consciously* critically to examine and question their dispositions to invoke prevalent cultural beliefs and stereotypes, other research that they cite suggests otherwise. The research into how unconventional linguistic salience perspectives arouse the audience's cognitive attention tends to find that its effects occur below the audience's conscious awareness. Consider, for instance, Allyson Holbrook et al.'s study, in which they asked participants a question with simple, dichotomous

⁸⁸ Holbrook et al. (2000), Roese et al. (1998).

⁸⁹ Bruckmüller et al. (2012: 212).

⁹⁰ Ibid.

answers. For instance, one question was: *if you could vote on requiring all drivers to have accident insurance, would you vote for or against this?* This question represents a conventional linguistic salience perspective; usually, we ask if someone is *for* or *against* something, in that order. Holbrook et al. found that study participants both took longer to respond, and gave unpredictable answers, when this question was phrased unconventionally—namely, where participants were asked if they were *against* or *for* something. When asked to explain the thoughts that went through their mind whilst answering the unconventionally framed question, the fact that these participants never mentioned the unconventional order of response choices lead Holbrook et al. to conclude that an unconventional question order (or, what I would call salience perspective) “[has] its effects completely outside of consciousness”.⁹¹

Bruckmüller et al.’s unconventional salience perspective likely engaged the study participants’ *unconscious* brain, then, when it helped them to, in the authors’ words, *examine status differences more critically*. This mirrors the discussion of effective implicit bias countering techniques mentioned in chapter 1; those that were successful used, in Jennifer Saul’s words, “counter-intuitive mechanical techniques that draw not upon our rational agency but upon automatic and unconscious responses”.⁹² Just as the implicit, unconscious mental processing of stereotypes and biases can have a special power in shaping our mental and behavioural responses *for the worse*, then, the discussion here affirms that our implicit, unconscious processing can also shape our mental and behavioural responses *for the better*.

With this in mind, we can suggest various ways to expand Bruckmüller and colleagues’ study in line with the discussion in this chapter. The conventional question posed to participants by Bruckmüller and colleagues was: *do women lead differently than men in boardrooms?* They played around with the order of *women* and *men* to elicit different responses from study participants, but the discussion in this chapter suggests that they might also play around with other aspects of the question. Firstly, given that a belief that the sexes/genders are overwhelmingly different to each other is a component of sex/gender essentialism, we might suggest adopting an unconventional linguistic salience perspective that instead asks about the *similarities* between the sexes/genders in their leadership style. What we have learnt about the impact of unconventional linguistic salience perspectives suggests that this small change might be sufficient to alter audience

⁹¹ Holbrook et al. (2000: 484).

⁹² Saul (2013: 260).

responses for the better—in other words, to reduce endorsements of essentialist beliefs about sex/gender difference. This brings us back to Hyde’s paper discussed at the very beginning of this chapter. Indeed, Hyde is concerned about sex/gender psychology research pandering to various problematic sex/gender biases, so it is no surprise that she has opted for the unconventional linguistic salience perspective *Gender similarities and differences*.⁹³ We can conceive of a further study, along the lines of that by Bruckmüller and colleagues, which involves a control paper discussing the same linguistic content as Hyde’s paper, but with sex/gender *differences* mentioned before *similarities*.

Secondly, we might also suggest playing around with which sex/gender categories are mentioned in Bruckmüller et al.’s study. The assessment of sex/gender essentialism in §5.3.1 revealed a belief that the sexes/genders form a binary, in the sense that either there are, or should be, only women/females and men/males (there are no intersex sexes or queer identities, or where they exist they are deviant). We might, then, ask something like this: *how do the various sexes/genders lead in ways that are different (similar) to each other?* Might this shift in linguistic salience perspective steer the audience to less essentialist responses? We might also consider going a little further than strictly what counts as a change in linguistic salience perspective, by introducing new content—namely, by *naming* an example of a gender identity beyond the binary. We might ask, for instance: *how do genderqueer individuals, women and men lead in ways that are different (similar) to each other?* The research discussed here suggests that these unconventional salience perspectives and phrasings might be instrumentally beneficial, insofar as they reduce our essentialist beliefs about sex/gender binaries.

Generally, this chapter suggests that making salient (in both our language and mind) the opposite to that which our sex/gender essentialist biases make salient, might help to avoid the activation of our sex/gender essentialist beliefs. Instead of binary categories and difference, we might consider giving relative salience to a wider range of sex/gender categories, as well as the similarities between them; instead of discreteness of sex/gender categories and the homogeneity within each category, we could give more attention to the overlap between categories, and the variability within a category; instead of internal biological properties, we might give relative salience to external properties such as social norms and stimuli; instead of actual distinctive traits of sex/gender groups, we could consider attending more to the potential for change in group attributes. (The

⁹³ See Hyde’s (2005) paper for a more explicit discussion about her concern about sex/gender psychology research playing into problematic sex/gender biases.

contents in this list overlap with what we introduced as *the process salience perspective* in the previous chapter. This perspective made salient the external context in which a subject exists, and the potential for change *from* whichever traits that subject exhibits. There is a case, then, that the process salience perspective could be of some help in avoiding the activation of sex/gender essentialism.)

Just as with Bruckmüller et al.'s unconventional salience perspective for androcentric domains, the scope of the benefits to be gained from using a salience perspective that counters our sex/gender essentialist biases will likely be similarly circumscribed. Just like our androcentric biases, sex/gender essentialist beliefs are too strong (i.e. too cognitively accessible and socially licensed) to expect the use of a mere counter-salience perspective to be able to activate thorough-going *anti-essentialist* beliefs about sex/gender in its audience, such as that sex/gender is a continually evolving set of traits that might alter unrecognisably in the future. (Further, anti-essentialist conceptions of sex/gender are plausibly *not* cognitively accessible and socially licensed, which means that salience perspectives, given that they rely on a bias being so accessible and licensed, are unlikely to activate them.) Whilst we cannot expect our suggested counter-salience perspective for sex/gender to help us to conceptualise sex/gender in radically different, anti-essentialist ways, we can expect a lessening of the *degree* to which we endorse essentialist biases. Given that the suggestions made here are for particularly small, minimal-effort linguistic changes, the fact that we might achieve a lesser degree of sex/gender essentialism remains a significant result.⁹⁴

What about adopting this counter-salience perspective for sex/gender in our *cognition*? Changing patterns of salience in language is one thing, but how does one develop long-term habits of focussing on sex/gender categories beyond the binary, on the similarities between the sex/gender categories, on the potential for different

⁹⁴ One way in which this counter salience perspective might have benefits beyond the mere lessening of our essentialist biases is if it is able to activate a *good* bias. We heard in the previous chapter about the *growth mindset*, which involved a person believing their intelligence to be something that can develop over time in response to effort. I suggested that the growth mindset was becoming increasingly cognitively accessible and socially licensed in Western cultures. Now, the growth mindset is associated with benefits in the context of sex/gender. For instance, where one believes that maths ability grows with effort, grades improve and, in Fine's (2010: 185) words, "gender gaps diminish". It was also said better to protect women against essentialist stereotypes about maths ability (Dweck, 2006b). Given that this mindset makes salient at least two of the contents emphasised by our counter-salience perspective for sex/gender – external context and the potential for change – it might be liable to become activated through use of this perspective.

gendered traits in the future, and so on? Further, how do we do this when the society around us is functioning to make sex/gender binaries, sex/gender differences, and so on, attention grabbing, memorable, and cognitively accessible? In chapter 2, I suggested that individualist approaches to changing one's cognitive salience perspective will often be insufficient in this sort of case. Instead, we need to make structural changes within the culture itself. For instance, we might want to look to changing the toyshops that divide gifts along stereotypically-gendered lines, the media reporting of hardwired brain differences between the sexes/genders, the administrative forms that require indicating one's sex/gender where this is not strictly relevant, and so on.

5.5. Conclusion

Bruckmüller and colleagues' study gave us evidence of salience perspectives activating inferences to substantive, truth-conditional beliefs concerning sex/gender. These beliefs were inferred in part because they were particularly cognitively accessible and socially licensed. Finding these beliefs harmful, Bruckmüller et al. recommended that we adopt a linguistic salience perspective (or, *linguistic framing* in their words) that has a better chance at thwarting the easy retrieval of these beliefs. This answered the first question that we had at the beginning of this chapter; there is indeed empirical evidence that a mere change in salience perspective for the specific subject of sex/gender can mean the difference between activating problematic (androcentric) biases, and not.

At least in Western cultures, another set of particularly cognitively accessible and socially licensed beliefs in the context of sex/gender is sex/gender essentialism. This chapter has laid out the components of this bias, on the (common) assumption that it does indeed have a doxastic structure. These beliefs, I argued, are both false and harmful. I then suggested that linguistic and cognitive salience perspectives that simply made the contents central to these beliefs *salient* risked activating inferences to these beliefs. In particular, I said that, when talking about sex/gender, or attending to sexed/gendered subjects, this means being wary of making salient: sex/gender binaries; sex/gender differences (i.e. the idiosyncrasies sex/gender groups currently display); ways in which sex/gender categories are discrete; ways in which members of a given sex/gender category are similar; the biological insides of sexed/gendered persons. (Whilst not relevant to cognitive salience perspectives *per se*, I also suggested that we ought to think twice before mentioning a person's sex/gender in situations where it is not strictly relevant.) This answered the second question posed in the introduction—namely, it

helped to answer which sex/gender contents we should avoid making salient in our language and thought. The substance salience perspective was identified as making salient many, though not all, of the contents central to sex/gender essentialism, and thus was considered liable to activate a form (though not the strongest form) of sex/gender essentialism.

Finally, I asked what we can hope to achieve by making salient content that is the *opposite* to that which is central to sex/gender essentialism. Here, I suggested that the potency of the sex/gender essentialist bias in Western cultures means that the benefits of a counter-salience perspective will no doubt be limited. It is unlikely to help to activate radically anti-essentialist beliefs about sex/gender. The results from using unconventional salience perspectives in Bruckmüller et al.'s study, however, gives us hope that it will at least thwart the easy activation of sex/gender essentialist beliefs.

The discussion in this chapter should make us all (from lay people, to journalists, to academics) carefully consider how we present information about sex/gender, and which salience perspectives we decide to cultivate. This is especially important given the implicit, under-the-radar way in which salience perspectives can activate inferences to our biases—this can give them *more* power to shape our cognitive and behavioural responses. The fact that these responses can help to *oppress* people from disadvantaged social groups, as was the case with sex/gender essentialism, further demonstrates the importance of employing the right salience perspective.

CHAPTER 6

**SALIENCE PERSPECTIVES AS
CONSTITUTIVE HARMS:
SEX / GENDER**

6.1. Introduction

In previous chapters, I have spent most of the time focussing on the harms, as well as some of the epistemic faults, that salience perspectives can *cause*, through their activation of problematic biases. In this chapter, I expand upon one of the ways mentioned in chapter 3, in which we might offer a *constitutive* critique of salience. Here, I build upon the suggestion that salience can constitute harm in itself, independently of any harmful effects it may have. Focussing again on the topic of sex/gender, I argue that we can harm someone in this way simply by making certain things salient about them.¹

A woman can be harmed, for instance, when what she is wearing is more salient in her interlocutor's mind than her conversational abilities. A philosopher can be harmed when what others find most salient about her is her identity as a woman, as opposed to her philosophical expertise. For an athlete who has been raped, it can be in her best interests for her athletic achievement to be her most salient feature, and not the fact that she was raped.

My aim is to clarify and expand upon the suggestion made about Charlotte's cognitive salience perspective on Jane, discussed in chapter 3. Namely, making the wrong thing salient about a person can harm them by constituting a way of disrespecting their personhood. The questions driving this chapter are three-fold. My primary question is as follows: is the idea that mere salience can constitute harm plausible? Secondly, if it is, what can the notion of salience as constitutive harm offer to the subject of sex/gender

¹ The type of harm being considered in this chapter is what Feinberg (1987: Ch. 1) calls "setbacks to interests". I am asking, therefore, if patterns of salience can themselves hinder our interests. Our interests, Feinberg suggests, are those things in which we have a stake. They are components of our well-being, insofar as, in Feinberg's words, one "flourishes or languishes as [one's interests] flourish or languish" (ibid. 34). A person can set back that interest, and thus harm me, by making it difficult for me to achieve that fulfilment, for instance. Patterns of salience might count as harms in *another* sense identified by Feinberg, too. Making the wrong thing salient about a person might involve *wronging* them, in the sense of violating their *rights* (ibid. 34-5). Whilst this may be the case, it is not something that I consider in this chapter.

harm? Thirdly, what, if anything, makes the type of harm constituted by patterns of salience important to consider?

In what follows, I begin in §6.2 by consulting two case studies relating to the topic of sex/gender harms. These are designed to motivate and clarify the idea that certain patterns of salience can constitute harm. The first case study in §6.2.1 looks at salience-based harms in relation to women who have experienced rape. The second case study in §6.2.2 considers how making women's bodies their most salient feature can harm them. In §6.3, I build upon this latter case study and suggest that attending primarily to a woman so that her body is her most salient feature *counts* as a form of sexual objectification—what I call *perspectival objectification*. In other words, salience perspectives count as one aspect of sex/gender harm. This perspectival sex/gender harm is important to address. Its *subtlety*, I suggest, gives it an insidious power. In particular, its subtlety makes it difficult to challenge, both because it makes it hard to notice, and because it makes it easy to dismiss as trivial where it is noticed. In §6.4, I consider the topic of the previous chapter, *sex/gender essentialism*, with the preceding discussion in mind. In particular, I suggest another way in which salience perspectives might count as an aspect this sex/gender harm. Certain attentional dispositions, I suggest, might themselves count as a *form* of sex/gender essentialism: *perspectival essentialism*. In §6.5, I clarify the ideas discussed so far by defending them against two objections.

6.2. Case studies

6.2.1. Rape victim-survivors²

Consider a situation that too many women find themselves in, though it is a difficult one to talk about. Many women are survivors of rape: 85,000 women are raped in England and Wales alone every year.³ Whilst both men and women experience rape, it is women who disproportionately suffer from this terrible crime, and men who disproportionately

² I use the term *victim-survivor* as the generic term for someone who has experienced rape. I have chosen this expression in an effort to acknowledge the issues that come with choosing one or other term (*victim* or *survivor*). According to Jean-Charles (2014: 39), *victim-survivor* better “grapples with the multiple responses to and experiences with rape, as well as the different ways they figure on cultural production”.

³ ONS (2013: 6).

are the assailants.⁴ Further, research demonstrates how gendered norms and practices help to explain this gendered differential. For instance, many have argued that gender norms for men tend to encourage dominance and control, as well as the view that women are sexual conquests.⁵ It is with this research in mind that I assume, in what follows, that rape is a sex/gender issue.

Despite the gravity of the crime, many rape cases are not brought to trial.⁶ There are multiple reasons why women who have been raped might want to avoid testifying in the courts. The exceedingly low conviction rate is one.⁷ The emotional distress of having to recount a horrific attack to a room full of people is another.⁸ A different reason, however, might be that these women do not want the fact that they were raped to become the most salient thing about them. They do not want it to be what others find most noticeable, memorable, and cognitively accessible.

Part of their concern may well be about the harmful beliefs and ideologies that this cognitive salience perspective is liable to activate. The salience of their experience of rape might activate harmful victim-blaming beliefs and ideologies, in both the minds of others, as well as in these victim-survivors' own minds. Victim-blaming beliefs and ideologies are arguably culturally pervasive and socially licensed.⁹

There looks to be something more, however. To understand this, let's look to rape victim-survivor statements. A common concern raised by those who have been raped is that one's experience of rape ends up masking other aspects of one's identity. Monika Korra, who was kidnapped and raped when she was out on a run, has said that she wants to be known as a runner, not a rape victim. Calling running her passion, Korra describes it as "the thing that gave me identity in life", and that the fact that she was raped threatened the primacy of that identity.¹⁰ Similarly, Simone Biles, 4 times Olympic

⁴ Over 2017-8, for instance, 93% of users of Rape Crisis Centres in England and Wales were women (Rape Crisis, 2019a). Perpetrators are predominantly male (see, for instance, Black et al, 2011: 3).

⁵ See Brownmiller (1975), Edwards (1987), Cahill (2001), MacKinnon (2016) and Stotzer & MacCartney (2016) for further discussions of this point.

⁶ For instance, the crime survey for England and Wales found that around 5 in every 6 victims did not report their experiences to the police (ONS, 2018).

⁷ It is estimated that only 5.7% of reported rape cases result in a conviction for the perpetrator (Kelly et al, 2005).

⁸ See these and other issues eloquently discussed in an open letter by the victim in the Stanford University rape case (Anonymous, reprinted in Osborne, 2016).

⁹ Suarez & Gadalla (2010), Whatley (1996). These authors also explain the harm of victim-blaming beliefs and ideologies.

¹⁰ Korra, in Lopez (2016).

gymnastics champion, recently cited similar concerns when she announced that she had been sexually abused by her USA Olympic national team doctor. In a statement published on social media, Biles stressed that “this horrific experience does not define me. I am much more than this. I am unique, smart, talented, motivated, and passionate. I have promised myself that my story will be much greater than this...[emphasis in original]”.¹¹

Many complex issues around identity, and the *victim* status in particular,¹² no doubt play a role in these women’s thoughts here. One way of thinking about their statements, though, is by invoking cognitive salience perspectives. Let’s think back to what cognitive salience perspectives entail. What would it mean for Korra, for instance, to have her experience of rape taken to be her most prominent feature? Amongst other things, this would involve people *noticing* properties connected to the fact that she was raped more than others; for instance, they might notice others discussing her experience of rape more than they would notice others discussing her other properties, such as her skills, her interests, and so on. Further, the fact that Korra was raped would be the most remembered feature of her life by others. It would involve the fact that she was raped being particularly cognitively accessible to people when they think about Korra’s other of traits.

This is a problem when one benefits from having *different* features of one’s person most salient in the minds of others (and one’s own mind). Korra and Biles would do better having *runner* and *Olympic gymnast* respectively more salient than *victim of rape* in people’s minds. Firstly, *it matters to us how we are attended to*. Korra and Biles *want* to be known primarily as athletes: for their athletic successes to be noticed and remembered more easily than reports of their rape; for their determination and commitment to their sport to be at the top of people’s minds when they think of them, contemplate their behaviour, and act towards them. People can be harmed when they fail to be attended to in the way in which they wish.¹³

¹¹ Biles’ Twitter statement, reprinted in Lutz (2018). Other examples that echo Biles’ and Korra’s concern include Ashley MacDonald. MacDonald says “I definitely don’t want to be known as the rape victim. That’s not my goal in life”. She goes on: “I hope that other people will acknowledge that there is much more going on in my life than that” (MacDonald, in *CBC News*, 2017).

¹² See Jean-Charles (2014) and Kelly et al. (1996) for discussions of the issues surrounding both words, *victim* and *survivor*, in the context of rape.

¹³ Many philosophers have made similar points. For instance, one might look to Korsgaard’s notion of *practical identity* to explain this type of harm. An agent’s practical identity is their sense of self or, in Korsgaard’s (1996: 101) words, “a description under

This idea is a familiar one. Artists from minority backgrounds commonly complain about being seen primarily in terms of their ethnicity; they might get referred to as (for example) an *Indian artist*, whilst another, (white) artist will simply be referred to as an *artist*.¹⁴ Using the prefix *Indian* is one way of making these artists' ethnicity their most salient feature. Further, women scientists not infrequently object to being seen first and foremost as a woman instead of as a scientist. Take, for instance, Monica Esopi, a doctoral candidate in chemical engineering, who talks about the relief that she felt moving into a department that openly and earnestly discussed diversity-related issues. She says "I no longer feel like I'm seen as a woman *first* [emphasis added]; I am just a researcher, a scientist, an engineer".¹⁵ The issue that these individuals raise is not (always) that others see them *only* as a minority ethnic group or woman. Rather, it is that their identity as a minority ethnic group or woman seems to take precedence over their other identities—identities that they wish for others to focus on.

I want to consider another, stronger reason, however, for Korra and Biles having their identity as *athlete* more prominent than that of *rape victim-survivor*—a reason that does not depend on their particular desires. Consider sexual objectification, which involves treating a person as a thing in some way. (I return to this topic in §6.2.2 below.) A common feminist explanation as to why phenomena like objectification harms women, have referenced the idea that we deserve to be recognised as agents with personhood—with, amongst other things, rationality, a capacity to set and pursue our own ends, integrity and personality.¹⁶ We heard about this briefly in chapter 2. Objectification is harmful because it involves *disrespecting* an individual's personhood. Rachel Fraser has noted how this idea crops up in feminist writing on rape. She observes how, in order to

which you value yourself, a description under which you find your life to be worth living and your actions to be worth undertaking". This can include one's status as *father*, *Brit*, or indeed *athlete*. Others have a duty to respect our practical identity, according to Korsgaard. Honneth's (1995) ideas on *social recognition* might also help. According to Honneth, we cannot experience self-realisation, or the development of our capabilities, unless the others around us respect our identity. One might be able to expand upon these philosophers' ideas, then, to demonstrate that one way in which our chosen (practical) identities can harmfully be disrespected is by others not *attending* to that identity in the right way—namely, by failing to make our chosen identity our most salient one.

¹⁴ Pollock & Parker (1981: xix).

¹⁵ Esopi, in Science Careers Staff (2018).

¹⁶ The first two features of this list are borrowed from Kant (*Lectures on Ethics*), who is cited by many theorists writing about objectification (see Papadaki, 2015: §1). The latter features come from Dworkin (2000: 30-1) and Bartky (1990:130), who expand upon Kant.

avoid this harm, “feminists have long argued for the importance of recognising the complex personhood and agency of those who have experienced rape”.¹⁷ This is especially important given that, in our culture, rape victim-survivors are commonly understood and portrayed as passive and lifeless. In Susan Brison’s powerful philosophical paper detailing her experience of rape, she talks about how crime novels and detective films portray the rapist as agentic, whilst the victim is treated as “a merely passive pretext for our entertainment”.¹⁸ Fraser herself cites rape metaphors as evidence for this way of seeing rape victim-survivors. Metaphors such as *Germany is raping Brazil in the football right now* are so prevalent and successful in part *because* of how common it is to think of those who are raped as powerless and passive.¹⁹ In other words, rape victim-survivors are commonly portrayed as *lacking* personhood, and this harms them.

How can we apply the common thought that we can harm a person by disrespecting their personhood to the context of salience perspectives? Well, as was the brief suggestion in chapter 3, attending to a person so that their personhood-related traits are their most salient attribute looks to be one way of *respecting* their personhood. (A personhood-related trait is a trait reflecting one’s individuality, rationality, agency, and so on. For instance, it might be a trait over which one has been able to exercise some agency, whose cultivation involved exercising one’s rational faculties, and/or whose existence demonstrates one’s personality, and so on. Whilst this notion will need clarification, I think that we can make some good intuitive sense of it.) For instance, it seems clear that being a rape victim-survivor, on its own, is not a trait that one has exercised agency over, whereas being an athlete usually does require setting one’s own goals and pursuing them. Having *athlete* at the top of our minds when thinking about Korra and Biles helps us to notice and remember an identity that they have autonomously chosen, displays their individuality, reflects their ability to set their own goals (thereby reflecting their rationality and agency), and so on. It is in their interests for their personhood-related traits to be made most salient. Conversely, attending to a person so that their *non*-personhood-related traits are their most salient feature is a way of disrespecting their personhood—namely, of *harming* them.²⁰ Having *person who was raped* most prominent in our minds directs our attention to a passive status as someone

¹⁷ Fraser (2018: 747).

¹⁸ Brison (1993: 11).

¹⁹ Fraser (2018: 745).

²⁰ See back to fn. 1., where Feinberg suggests that we can be harmed when we experience a setback to our interests.

who has been attacked—a status that they did not choose, and so does not reflect their individuality, their agency, and so on. (I consider the case of the individual who finds their status as rape victim-survivor to be powerfully agentic, and wishes for others to attend primarily to that identity, in §6.5.)

Disrespecting an individual’s personhood is harmful not just instrumentally, but in itself, independently of its further effects. This means that, where a cognitive salience perspective is a way of disrespecting someone’s personhood, it too can be considered non-instrumentally harmful. It can *constitute*, as opposed to cause, harm. (We might also suggest that a *linguistic* salience perspective can constitute harm in this way by making an individual’s non-personhood-related traits most salient, such as by mentioning Korra’s experience of rape *before* her athletic career.)

This idea might helpfully be applied to the other examples that were mentioned earlier, namely, of the artist complaining that their ethnicity is given relative salience over their identity as an artist, and the scientist objecting to her sex/gender taking precedence over her identity as a scientist. Understanding that we can be harmed when a non-personhood-related feature of ours is made our most salient attribute might help to identify one of the harms occurring in these cases. Indeed, the professional identity of these individuals (artist, scientist) seems a better candidate for a personhood-related trait than the ethnicities or genders of these individuals.²¹

Before moving on, it may be helpful to note that the claim being made here – namely, that a given phenomenon might involve not just causal but also *constitutive* harms – is a familiar one. For instance, Rae Langton discusses how common this approach is in the context of pornography. She says [in-text citations are omitted for clarity of presentation]:

“Considerable work has been done to unpack the idea that pornography might harm women in a constitutive manner: for example, the idea that it discriminates against women; it subordinates women; it enacts facts about what is permissible and not permissible; it alters conventions governing women’s speech acts; it is comparable to hate speech and group libel. Agreement might well be reached that certain forms of pornography—for example, an infamous Hustler image of a headless naked woman being fed into a meat grinder—harm women in the

²¹ For those who do wish to have their ethnicity or sex/gender as their most salient feature, we may be able to give responses that parallel those in §6.5 below.

way that hate speech and racial insults are understood to be constitutively harming their targets.”²²

As Langton suggests, then, there are a whole host of ways to argue that something, in this case pornography, might involve constitutive harms. Her own defence of this idea – or rather, her defence of the conceptual grounds of Catherine MacKinnon’s argument that pornography constitutes harm –²³ borrows from speech act theory. This theory begins from the observation that, as J. L. Austin would put it, we use language to *do* things. For instance, saying *I do* in an official marriage ceremony does not just *cause* two people to become married: the utterance itself *constitutes* the act of marrying. Or, the utterance *Run! There’s a fire!* said by someone in a burning building, itself *constitutes* the act of warning.

Sometimes, the act constituted by the utterance is harmful. Developing examples closer to the case of pornography, Langton gives the example of apartheid law. She says that the utterance *Black people are not permitted to vote* uttered in certain contexts, *constitutes* subordination.²⁴ In particular, it constitutes the acts of (unfairly) *ranking* black people as inferior, (unjustly) *depriving* them of rights and powers, and *legitimizing* discrimination against them.²⁵ We might want to put it this way: even if (for some hard-to-fathom reason) this utterance didn’t have the pernicious *effects* we associate with apartheid law, it would nevertheless constitute harm, by itself *counting* as an act of subordination. (I take it for granted that Langton deems subordination to be a harm in itself.) In other words, even if there is a possible world in which this utterance does not lead to black people staying away from polling booths (and so on), we would nevertheless want to say that the utterance *Black people are not permitted to vote* is harmful. In particular, it would be harmful in this possible world because the utterance is, in and of itself, a way of enacting the subordination of black people.

It is with these examples in mind, built on the scaffold of speech act theory, that Langton suggests understanding MacKinnon’s proposal that pornography constitutes

²² Langton (2008: 1).

²³ MacKinnon (1987).

²⁴ Langton (1993b: 302). In this work, Langton (ibid. 304) argues that, for the utterance in question to count as a given act, the utterer must have *authority* in the relevant domain (in the Apartheid case, she suggests that the utterer might be a legislator in Pretoria). Many others since have questioned this condition, suggesting that, often, one does not require any special authority to nevertheless perform various acts through one’s utterances (see, for instance, Bauer, 2005).

²⁵ Langton (1993b: 303).

subordination. We can understand MacKinnon, according to Langton, to be suggesting that pornography is a speech act of subordination. In and of itself, pornography (unfairly) *ranks* women as inferior, (unjustly) *deprives* them of power, and *legitimizes* sexual violence against women.²⁶

The point that I would like to dwell on from the preceding discussion is this: there is an important precedent for arguing that a phenomenon might not just cause harm, but also *constitute* an act of harm. Instead of suggesting that certain *utterances* might constitute harm, my claim is that certain *patterns of attention* might constitute harm. Further, instead of focussing on the specific harm of *subordination*, as per Langton's argument above, the type of harm that I identify is *disrespect for an individual's personhood*. Generally, though, I follow the spirit of these other arguments: just as pornography itself might *constitute* a way of subordinating women, I want to suggest that certain patterns of attention might *constitute* a way of disrespecting an individual's personhood.

6.2.2. Women and their bodies

Let's consider another case study that will help to clarify and expand upon this idea. Imagine a man, let's call him Terry, who, when interacting with other men, tends to notice their voice and face more than their body. He easily remembers what they say, as well as their facial expressions. When he considers an individual man that he has met, he tends to find these aspects of them more cognitively accessible than memories of build or biceps. By contrast, when Terry interacts with other women, he often notices and remembers a woman's figure or bust more than their face and voice, finding these attributes most accessible in his mind when he later reflects on what they were like. Now, Terry still notices and remembers other aspects of the women he meets, including their face and voice. He remembers and values their conversational contributions and personality quirks, but they are simply less prominent in his mind than these other features.

Terry is systematically attending to women differently to men. Is this a problem? Well, you might think it is if Terry's cognitive salience perspective on women is activating objectificatory beliefs and ideologies, and their associated behaviour. One type of sexual objectification involves reducing a person to their body, or body parts.²⁷ Indeed, many have suggested that our culture is saturated in objectificatory images and

²⁶ See Langton (ibid. 307) for these specific suggestions as to how pornography might subordinate women.

²⁷ Langton (2005: 246-7).

narratives that encourage us to think of women in particular as reducible to their bodies and appearance more generally.²⁸ This suggests that objectificatory beliefs and ideologies about women are plausibly cognitively accessible and socially licensed. Perhaps Terry's cognitive salience perspective, by making women's body parts their most salient feature, is helping to activate and license these beliefs and ideologies. It might be resulting in him thinking things like *women are reducible to their bodies*, or acting in objectifying ways, for example by touching women inappropriately.²⁹

Is there a problem with Terry's cognitive salience perspective independently of these effects, though? Let's follow the logic of the previous section. How we communicate through speech is a paradigmatic way in which we express our agency, personality, rationality, and so on. In other words, it is a way in which we express and exercise our personhood. Attending to men so that their conversational contributions are their most salient feature, then, is a way in which Terry *respects* men's personhood. By contrast, women's bodies in particular have taken on an especially object-like quality in contemporary (at least Western) culture. For instance, it is particularly common in adverts to crop photos of women so that their body, but not head, is in the shot.³⁰ Men, on the other hand, are depicted with a greater focus on the face in the media generally speaking.³¹ At least in Western cultures, the head and face is commonly associated with one's personhood; we tend to think of the mind, with its thoughts, feelings, desires and so on, as located in the head. Our faces are where we usually express ourselves and communicate, through speech and facial expressions. Indeed, research finds that those depicted with greater focus on the face in media and advertising images are evaluated as more intelligent (i.e. rational), assertive, and ambitious (i.e. agentic, autonomous) than those depicted with emphasis on the body.³² In other words, they are evaluated as having more personhood. Removing the head from photos of women takes away these things, reducing a woman to a feeling-less, desire-less body—something more like an object than a person. Indeed, historically and culturally, Sandra Bartky suggests that the body has “been regarded as less intrinsically valuable, indeed, as less inherently human, than the mind or personality”.³³ Attending to women so that their body parts are their most

²⁸ Goh-Mah (2013), Belsky (2019).

²⁹ Another form of objectification involves treating women as though they do not have boundary integrity (Nussbaum, 1995: 257).

³⁰ Belsky (2019).

³¹ Archer et al (1983), Szillis & Stahlberg (2007).

³² Archer et al (1983) and Szillis & Stahlberg (2007).

³³ Bartky (1990: 35).

salient feature, then, is one way in which Terry *disrespects* women's personhood. This means that Terry's cognitive salience perspective on women *constitutes* a way of harming them.³⁴

6.3. Perspectival objectification

The particular type of harm that Terry seems to be responsible for looks to be connected even more intimately to objectification than our case of rape victim-survivors. The way in which Terry is attending to women looks not just to be a way of disrespecting their personhood (a notion that I have borrowed from the literature on objectification), but specifically to count as a way of disrespecting women's personhood *by treating them as a thing*. Martha Nussbaum's influential account of objectification details seven ways of treating a person as a thing, including, for instance, denying a person's autonomy, and treating them as fungible.³⁵ Other ways have been added since; Rae Langton has also suggested that silencing a person, as well as reducing them to their appearance or to their body count as ways in which we treat a person as a thing.³⁶ The suggestion being made here is that attending to a person so that their thing-like properties are their most salient feature is another way of treating them as a thing. In other words, it is itself a form of objectification.

This idea is arguably in the background of some women's complaints regarding how others, often men, look them at. The notorious quip *I'm up here*, said to a man who is looking more at his interlocutor's chest than her face, could be construed as a claim about attention. It is not uncommon to come across personal stories from women on this theme.³⁷ This complaint is not (always) about men entirely *ignoring* a woman's face, speech, and other personhood-related traits, rather it is about eyes lingering on certain body parts more than one's face, for instance. Studies indicate this sort of gaze can have

³⁴ The fact that Western culture is saturated in objectifying depictions of women such as those suggested above suggests that cultural factors may make it difficult for Terry to adopt attentional dispositions that do *not* involve giving relative salience to women's bodies. As per my discussion about *changing* entrenched cognitive salience perspectives in chapter 3 (§3.4), we might find that cultural changes are necessary in order to allow individuals (such as Terry) to shift cognitive salience perspective—for individuals automatically and intuitively to find women's faces and voices more salient than their bodies.

³⁵ Nussbaum (1995: 257).

³⁶ Langton (2005: 246-7).

³⁷ See, for instance, Adebisi (2018), Burriss (2017), Roberts (2002: 327).

pernicious effects. For instance, one Australian study suggests that adolescent girls' feelings that they are stared at in public swimming pools inhibits the frequency and quality of their participation in physical activities.³⁸ The suggestion here, though, is that these attentional dispositions can be harmful even notwithstanding these effects. One can harmfully objectify simply in virtue of attending primarily to a person's thing-like traits.

This idea is also implicit in existing feminist writings on objectification. Sandra Bartky, for instance, claims that women can be objectified in virtue of being “too closely identified with [their body]”.³⁹ Paraphrasing Bartky's words, Evangelina Papadaki writes in a way quite consonant with the proposal at hand: “All the focus is placed on a woman's body, in a way that her mind or personality are not adequately acknowledged”.⁴⁰ We could read Papadaki here as suggesting that instead of requiring a wholesale *reduction* of a woman to her body to count as objectification (perhaps meaning that a woman is wholly *identified* with her body, so that her mind and personality is not even registered), objectification can be a matter of the *degree of focus* on a woman's body.

Despite these resonances, the particularly *minimalist* dimension of mere salience perspectives is never made entirely explicit. In addition, other aspects of Papadaki's phrasing, such as where she states that *all of the focus* is placed on a woman's body, could undermine my reading of focus coming in degrees. Further, Bartky herself goes on to expand her view in ways that indicate that something beyond mere salience is at issue. For instance, she uses words like *infatuation* synonymously with her idea of a *focus* on the body;⁴¹ *infatuation* implies an intense and all-encompassing focus on something, as opposed to the subtler notion of giving *relative* salience to one thing over another in one's attention. She also uses phrases such as “[being objectified] is to have one's entire being identified with the body”.⁴² The sort of objectification that I suggest – let's call it *perspectival objectification* – is not as strong as this.⁴³ Aspects of one's person beyond one's body (such as one's personality and autonomy) are recognised by the objectifier; it is simply that the objectifier *better* attends to one's body. They give it *relative salience* over one's personality. They find it more noticeable, memorable, and cognitively accessible.

³⁸ James (2000).

³⁹ Bartky (1990: 130).

⁴⁰ Papadaki (2018: §3).

⁴¹ Bartky (1990: 131-2).

⁴² Ibid. (35).

⁴³ I introduce *perspectival objectification*, and the notion of *perspectival harm* more generally, in Whiteley (forthcoming).

(The example of Charlotte’s cognitive salience perspective on Jane, discussed in chapter 3, looks like an instance of perspectival objectification. Charlotte is giving relative salience to Jane’s thing-like properties, namely, her fat body, over her personhood-related properties, namely, her charity work.)

The concept of perspectival objectification can be clarified by contrasting it with what it is not. In particular, it can be contrasted with *doxastic* objectification, which involves objectificatory *beliefs*, such as that women are submissive and object-like. Philosophers such as Sally Haslanger and Rae Langton discuss this form of objectification.⁴⁴ A distinctive way in which these beliefs cause harm is in being false, or perhaps in being unjustified given their wrong direction-of-fit with the world—in this latter case, objectificatory beliefs are often true only in virtue of shaping women’s behaviour so that they *become* submissive and object-like.

Perspectival objectification is *non-doxastic*; instead of being constituted by beliefs, it is constituted by attentional dispositions. Now, other forms of non-doxastic objectification have been proposed; Susan Bordo talks about objectifying *associations* (between women and bodies), for instance, and Alison Assiter discuss objectifying *desires*.⁴⁵ Further, the language of *treating* a person as a thing, adopted by Nussbaum, suggests that objectification might also consist in *behaviours*.⁴⁶ A distinctive way in which these non-doxastic forms of objectification are harmful is in their being *unwarranted* (they cannot be true or false, given their non-propositional nature).⁴⁷ Indeed, we might want to argue that perspectival objectification constitutes harm through its attentional dispositions being unwarranted, perhaps by following the suggestions made by Elisabeth Camp, Sebastian Watzl and Susanna Siegel in chapter 3 (§3.3.2).

What makes perspectival objectification unique, even from these other non-doxastic forms of objectification, however, is its minimal nature. Perspectival objectification is a distinctively subtle way of disrespecting another’s personhood.

Why offer this minimalist form of objectification to the table? Well, we often do not notice, let alone articulate the ways in which we attend to people and social groups.

⁴⁴ Haslanger (1993), Langton (1993a, 2004).

⁴⁵ Bordo (1993), Assiter (1988: 68). See also Langton (2000) and Kant (*Lectures on Ethics*: 165) for discussions of objectifying desires.

⁴⁶ Nussbaum (1995: 257). See also Halwani (2010: 187) who defines objectification as comprising behaviours. Those who advance doxastic accounts of objectification also mention how those beliefs *cause* harmful behavior (see, for instance, MacKinnon, 1987: 138).

⁴⁷ In chapter 2 (§2.4.1), I suggested that non-doxastic phenomena are commonly evaluated on epistemic grounds of warrant.

If our attentional dispositions really can harmfully objectify, then their subtle, rarely noticed or articulated nature gives them an insidious power, not least insofar as it makes them difficult to challenge. (Indeed, we cannot challenge a harm that we do not notice occurring, meaning that this harm can continue unchecked.)

When we *do* notice and articulate our attentional dispositions, it is easy to dismiss them as trivial. *What does it matter if someone notices one thing better than another! Shouldn't we be focussing on what they believe, and how they act? These things have real consequences!* Those who feel that they are being harmed by another's attentional dispositions risk being told that they are overreacting if they express this concern. This functions to silence them.⁴⁸ Either they pre-empt negative reactions and so silence themselves by not speaking, or they *do* articulate the harm that they have experienced, but are silenced insofar as their complaint is not respected.⁴⁹ As for this latter form of silencing, a complaint about perspectival objectification risks being dismissed in the following way: *Oh, come on, you're being irrational! What does it matter if he found your figure especially memorable! It's not like he ignored your face or what you said; you were listened to!* This also contributes to the insidious power of cognitive salience perspectives; they remain *effectively* unchallengeable, even when they are noticed and articulated.

We see a similar type of harm in what are termed *microaggressions*, which are defined as subtle and often brief everyday events that denigrate individuals on the basis of their group membership.⁵⁰ For instance, it is often suggested that a person of colour living in a primarily Caucasian country faces a microaggression when they are asked *where are you (really) from?* This is a subtle way of reinforcing the idea that this person is not really from that country—that they are a foreigner in their own land. Whilst different, harmful salience perspectives share with microaggressions a minimal nature, making them all too often overlooked or dismissed as trivial. This, according to Catherine Wells, creates a special type of harm. She says: “a microaggression does not just bring injury, but also brings the practical need to pretend that the aggression never happened. If one

⁴⁸ Wing Sue (2010: 66) makes this point in relation to microaggressions, which we will hear about briefly below.

⁴⁹ The former scenario counts as silencing in the sense of what Dotson (2011:244) refers to as *testimonial smothering*. Testimonial smothering is a kind of coerced self-silencing that occurs when, in Dotson's words, “the speaker perceives one's immediate audience as unwilling or unable to gain the appropriate uptake of proffered testimony” (ibid.). The latter scenario instead counts as what Langton (1993a) refers to as *perlocutionary silencing*. Here, one does speak, but one's objection, whilst heard, does not have the intended *effect*; instead of being registered as a valid complaint, one's complaint is dismissed.

⁵⁰ Wing Sue (2010: xvi).

is left angry, speechless, or hurt, one must hide that fact as best one can. Better to be seen as stumbling and inarticulate than to be seen as sensitive in irrational ways”.⁵¹ This, I suggest, also goes for harmful salience perspectives, such as perspectival objectification. The distinctiveness of perspectival harm, then, lies in its minimal nature; this gives it an insidious strength, insofar as it both difficult to notice and easy to dismiss as trivial.

6.4. Perspectival sex/gender essentialism

The discussion above tells us that simply making the wrong thing salient about a person can harm them. It can count as a way of disrespecting their personhood. It also tells us that salience perspectives might count as one aspect of sex/gender harm: making a woman’s thing-like properties more salient than her personhood-related properties can count as a form of objectification—perspectival objectification. Can these conclusions illuminate another way of approaching a different sex/gender harm discussed in this thesis, namely, sex/gender essentialism?

As we saw in the previous chapter, psychological essentialism about sex/gender is the idea that men and women have different underlying biological essences that determine their respective traits, behaviours and identities. Sex/gender essentialism is regularly cashed out as a set of *beliefs*, most commonly implicit, which is how we treated it in the previous chapter. For instance, Lea Skewes, Cordelia Fine and Nick Haslam’s recent paper on the topic refer to sex/gender essentialism as involving beliefs, as does Gail Heyman and Jessica Giles’ influential article *Gender and psychological essentialism*.⁵²

The discussion in this chapter suggests that more subtle forms of psychological essentialism might also exist—namely, ones that involve attentional dispositions, instead of beliefs, and that these might involve constitutive harm. What would this look like? Consider Pedro, a sex/gender researcher. He believes that there are more than just two sexes/genders—namely, that intersex and genderqueer people, for instance, exist. He believes that social norms play an important causal role in many of our sexed/gendered traits and behaviours. He also believes that we do not know how sex/gender might change in the future, as environments, and our biologies, alter. And yet, when he comes across information about sex/gender – in academic papers, newspapers, his interaction with people in the world, and so on – it is females/women and males/men (or

⁵¹ Wells (2013: 329).

⁵² Skewes et al (2018), Heyman & Giles (2006). See further instances of essentialism being discussed as though it is comprised of beliefs in chapter 4 (fn. 18).

information *about* females/women and males/men) that better catch Paul's attention. It's their *differences* (or facts *about* their differences) that better stick in his memory. The information about their *biology*, as opposed to social environments, is what he finds most cognitively accessible when later thinking about sex/gender. Generally, then, his attentional dispositions track content central to sex/gender essentialism.

We have already heard in chapter 5 that these sorts of attentional dispositions might be sufficient to activate sex/gender essentialist beliefs, making Pedro's cognitive salience perspective instrumentally harmful (and instrumentally epistemically problematic).

Brief discussions in previous chapters have also suggested some ways in which we might want to develop a *constitutive* critique of Pedro's attentional dispositions. Firstly, we might want to go down the *licensing* route, and argue that Pedro's cognitive salience perspective *licenses* psychological essentialism. The fact that licensing is a normative, as opposed to causal, notion would mean that Pedro's cognitive salience perspective, insofar as it does license sex/gender essentialism, bears a constitutive relation to this bias's harms (and epistemic faults).⁵³ Secondly, we might want to suggest that Pedro's attentional dispositions are *partially constitutive* of sex/gender essentialism. Perhaps partially constitutive of the belief that there is a sex/gender binary, for instance, is the attentional disposition better to notice, remember, and find more cognitively accessible binary sexes/genders over nonbinary sexes/genders. If we can show that sex/gender essentialism is non-instrumentally problematic, then, where Pedro's cognitive salience perspective partially constitutes this bias, his attentional dispositions can also be deemed non-instrumentally problematic.⁵⁴

Instead of investigating these lines of thought, let us instead consider how the discussion in this chapter might shed light on another type of constitutive critique of Pedro's cognitive salience perspective on sex/gender. Can his attentional dispositions themselves *count* as a way of harming sex/gender groups, such as women, or those who are non-binary?

Consider the relative salience given to actual biological idiosyncracies between women/females and men/males in Pedro's attention, over social context and the

⁵³ See the discussion in chapter 2 (§2.4.2). In particular, one might look to Tirrell (2012) and Swanson (forthcoming) for this approach.

⁵⁴ Perhaps, inspired by Watzl and Siegel's (Watzl, unpublished) approach discussed in chapter 3 (§3.3.2), one might try to demonstrate that sex/gender essentialism is an irrational outlook, which is non-instrumentally problematic.

potential for change in sexed/gendered traits. Pedro makes current findings that there are average differences in brain structure between men and women, for instance, more salient than the social structures and practices that shape men and women's traits. These attentional dispositions might be a problem if it is in the interests of women for these social features, and the potential for change in the sexed/gendered status quo, to be more noticeable, memorable, and salient.⁵⁵ Can we sketch an argument that it is?

To be able to do this, it will be helpful to begin by returning to the harms of *doxastic* sex/gender essentialism. We heard in the previous chapter that, where sex/gender essentialist beliefs are condemned as harmful, the harm cited is often one connected to oppression—to unjust social relations that disadvantage disadvantaged sex/gender groups.⁵⁶ Consider, for instance, essentialist beliefs that women have a submissive and nurturing biological essence (which might be located in their brains) that makes women well suited for domestic work and childcare. This, I suggested, helps to further inequalities in the domestic sphere, which oppress women not least by exploiting them and rendering them powerless. Oppression, I take for granted, is a harm in and of itself.⁵⁷

Inspired by the ideas in §6.2 and §6.3, we might try to show that *one way of oppressing women* is by making the social structures and practices that contribute to their oppression, and the potential for change *from* the traits that they exhibit in the current (unequal) sexed/gendered status quo, less salient. We might argue, for instance, that making the distinctiveness of women's brain structures and hormone levels more salient than patriarchal, capitalist institutions and practices that benefit from women doing most of the unpaid domestic labour, counts as a way of upholding those social institutions and practices that disadvantage women. Remembering that one *form* of oppression is powerlessness, we might also suggest that having these latter social features less noticeable, memorable and cognitively accessible is one way of limiting women's power to change the conditions of their oppressed status.⁵⁸ If oppression is a harm in itself, then

⁵⁵ In fn. 1 of this chapter, I mentioned Feinberg's account of harm as suggesting that one can harm a person by setting back their interests. If we can show that Pedro is setting back women's interests through his attentional dispositions, we can have a reason to think that he is harming them.

⁵⁶ Chapter 5 (§5.3.2).

⁵⁷ Others also treat it as a harm as itself, such as Cudd (2005: 23).

⁵⁸ We heard in the previous chapter about Young's (1990) influential account of oppression comprising of five *faces*: exploitation, powerlessness, marginalization, cultural domination, and violence.

so too are these attentional dispositions, insofar as they count as a way of oppressing women.

This also demonstrates a way in which the substance salience perspective, discussed in chapter 4, is relevant to this discussion. The substance salience perspective was a *generic* salience perspective, which could be applied to multiple subjects (including but not limited to sex/gender). This generic perspective *also* involved making internal properties more salient than external context, and the actual traits a given subject exhibits more salient than the potential for those traits to change. Instead of simply *activating* beliefs associated with oppression, which is what we argued for in chapter 4, the suggestion here is that these attentional dispositions might, when applied to certain subject matter, *constitute* a form of oppression. One way of oppressing disadvantaged individuals (whether they are women, disadvantaged ethnic groups, disadvantaged social classes, and so on) is, arguably, by making the social factors that oppress them, and the potential for their traits to change, less salient.

Returning to the example of Pedro the sex/gender researcher, though, we might examine another avenue for a constitutive critique. Consider the relative salience that Pedro gives to sex/gender binaries over non-binary groups, such as intersex and genderqueer. For Pedro, non-binary groups are less noticeable, memorable, and cognitively accessible to him than those falling within the sex/gender binary. This is a problem if it is in the interests of members of non-binary categories for those categories to be more salient in people's minds. Can we show that this is the case? (Note that this is a slightly different issue to that discussed so far. Instead of harm occurring through a person or group having the wrong *feature* made salient about them, a members of a group are harmed in virtue of how salient that group is *per se*.)

Again, let's return to the doxastic counterpart to the attentional disposition to find sex/gender binaries more salient than non-binary groups. As we suggested in the previous chapter, an obvious way of explaining what is harmful about the essentialist *belief* that sex/gender is or should be a binary (of women/females and men/males), is by pointing to the *marginalisation* of non-binary groups that this belief entails. This essentialist belief makes those from non-binary groups less heard, less able to influence decision making, less respected, and so on. As we heard in the previous chapter, this counts as a way of oppressing non-binary groups (marginalisation, as we heard in the previous chapter, marginalisation is one form of oppression, according to Iris Marion

Young).⁵⁹ Applying this thought to the context of Pedro's salience perspective, we might want to suggest that *one way of marginalising, and thus oppressing, members of non-binary groups is by making them less noticeable, memorable, and cognitively accessible than those from binary groups.*⁶⁰ If oppression is a non-instrumental harm, then so too are the attentional dispositions described above if they count as a way of oppressing members of non-binary groups.

I have only very briefly sketched how an argument, that attentional dispositions like Pedro's constitute harm, might proceed. If this argument can be defended, though, there would be grounds to suggest a new form of sex/gender essentialism. In parallel to the suggestion of perspectival objectification earlier, we might speculatively suggest another, particularly minimalist form of sex/gender essentialism—*perspectival essentialism*. Instead of being constituted by beliefs, one can be perspectivally essentialist in virtue of what one pays attention to, remembers, and finds cognitively accessible. (Indeed, whilst I will not consider this here, we might also pursue a similar avenue with the substance salience perspective. Instead simply of *activating* the correspondence bias, psychological essentialism, and the fixed mindset bias, the substance salience perspective might *constitute* particularly minimalist, perspectival versions of these biases. Perhaps one way of being subject to the correspondence bias, for instance, is by paying relatively more attention to the internal properties of a person, and the actual traits they happen to exhibit, than to the person's situation, and the potential for their traits to change.)⁶¹

Whilst there do not seem to be the same precedents for this minimalist, perspectival essentialism in that literature as there were for perspectival objectification in its respective literature, there are precedents for finding *non-doxastic* versions of essentialism. For instance, whilst Susan Gelman regularly talks about essentialist *beliefs* and *theories*, she at times adds ambiguity to this, saying at one point essentialism “can be

⁵⁹ Ibid.

⁶⁰ A critic might object to this, suggesting that we are not warranted in making non-binary identities and conditions more salient than their binary counterparts because of what Camp would call the *broad statistical distribution of properties* in the world (see Chapter 3, §3.3.2). Binary identities and conditions, this critic continues, are simply more numerous than their non-binary counterparts, which warrants them having relative salience! A possible response to this critic also makes use of Camp's account. One of the considerations that shape what we legitimately can make salient, Camp suggested, includes *our human interests and aims*. She explicitly mentions our *moral values* as legitimately shaping these interests and aims. Perhaps the moral value relating to treating everyone equally, and giving all equal visibility and voice, is relevant here; it justifies our suggestion that no one sex/gender group should be given relative salience over any other.

⁶¹ This perspectival version of the correspondence bias would involve constitutive harms if it *counts* as a way of oppressing members of disadvantaged social groups, as suggested above.

considered an unarticulated heuristic rather than a detailed, well-worked-out theory”,⁶² and at another times that it is “a skeletal framework ... rather than a detailed set of beliefs, scientific or otherwise”.⁶³ Adopting phraseology that sounds decidedly non-doxastic, Frank Keil sometimes refers to essentialism as a “stance”, “construal” or “heuristic”.⁶⁴ In a slightly different vein, Maykel Verkuyten suggests moving away from the idea that essentialism is about an individual’s inner cognitions, and instead embracing the idea that it is a social practice.⁶⁵ These various ways of capturing essentialism, whilst not mentioning attentional dispositions *per se*, do capture the desire to widen what counts as essentialism from simply the domain of belief. Perspectival essentialism is offered in this spirit.

Perspectival essentialism, then, is contrasted with the *doxastic* essentialism of the previous chapter. Overwhelmingly, sex/gender essentialism (and psychological essentialism more generally) is described as a set of *beliefs*. As mentioned above, Skewes and colleagues and Heyman and Giles talk of essentialist *beliefs* about sex/gender, as do Susan Gelman, and Sarah-Jane Leslie.⁶⁶ The distinctive harm in doxastic sex/gender essentialism is to be found in falsehoods (falsehoods that, as we have seen in the previous chapter, are used in arguments to justify various oppressive practices). The issues of stereotype threat in the previous chapter also suggest ways in which essentialist beliefs risk the wrong direction of fit with the world.⁶⁷

Perspectival essentialism is instead non-doxastic. Other forms of non-doxastic essentialism might exist; these might involve associations (such as between women, oestrogen, and nurturing personalities), or perhaps the social practices that are the focus of Verkuyten’s research, mentioned above (these might involve *behavioural codes* to treat women as domestic labourers). Whilst the non-propositional nature of these non-doxastic phenomena would mean that they lack the truth-conditions, they might be criticised on grounds of warrant.

What makes perspectival essentialism different even from these other non-doxastic forms of essentialism, however, is its minimal nature. Perspectival essentialism is a distinctively subtle way of oppressing underprivileged sex/gender groups. As discussed in relation to perspectival objectification in §6.3, this subtlety gives it a special strength;

⁶² Gelman (2003: 21).

⁶³ Gelman (2005:266).

⁶⁴ Keil (1995).

⁶⁵ Verkuten (2003)

⁶⁶ Skewes et al (2018), Heyman & Giles (2006), Gelman (2005), Leslie (2013).

⁶⁷ See chapter 5 (§5.3.2).

the rarely articulated nature of our attentional dispositions, and the fact that they are regularly dismissed as trivial where they *are* noticed, would give perspectival essentialism an insidious power.

6.5. Addressing potential objections

So far, then, I have introduced the general notion of perspectival harm, and suggested how it might relate to sex/gender harm. In particular, I have introduced perspectival versions of two familiar sex/gender harms: objectification, and essentialism. Here, I clarify these ideas by responding to two possible objections.

6.5.1. Perspectival harm is too broad

The notion of perspectival harm looks too broad. If one perspectivally harms someone by making their non-personhood-related traits more salient than their personhood-related traits, then this seems to include a great many scenarios that, to many intuitions, are morally benign. A doctor, for instance, will no doubt pay more attention to a person's *body parts* than their personality, as well as to their *biological insides* rather than their social context, but this hardly seems problematic. Far from subtly objectifying or essentialising the patient, the patient surely *benefits* from these attentional dispositions of the doctor's.

I agree. Indeed, this is a qualification that many philosophers make when talking about objectification in particular. Within this literature, diagnosing whether harm results from the various ways in which one can treat a person as a thing often requires examining the particular case and context.⁶⁸ A doctor arguably *reduces* her patient to his body, but various considerations concerning the patient's goal of health, and the role of a doctor, mean that harm does not occur.⁶⁹ Conversely, as highlighted by Langton, one can harmfully objectify another in a way that does not involve denying their personhood; sadistic rape, for instance, can involve a recognition of someone's personhood, with the aim of *suppressing* it.⁷⁰ The same goes for cognitive salience perspectives. We need

⁶⁸ Nussbaum (1995), Langton (2005).

⁶⁹ Some ways in which a doctor treats her patient as an object, however, *do* seem to be harmful. It has been argued that patients can be harmfully objectified by their doctors when their emotions, and personal perspectives on their health, are ignored, for instance (see Berglund et al, 2012).

⁷⁰ Langton (2005: 249). Cases like this suggest that the *manner* with which one attends to a given property will also help to decide whether harm occurs (see Watzl, forthcoming, for a related point).

knowledge of the particular case and context to decide whether *attending* to a person in a way that makes their non-personhood-related traits (or their internal biological properties, for instance) their most salient feature does in fact harm them. My claim is simply that one *can* perspectivally harm someone in the ways that I have described.

Whether one reserves the terms *perspectival objectification* or *perspectival essentialism* for instances where harm does occur, therefore defining these terms as inherently pejorative, depends on one's project. Consider, for instance, the different approaches taken within the objectification literature. Nussbaum's approach is to offer a cluster definition of *objectification* that aims to capture the varied ways in which the term is employed in ordinary usage, including what Nussbaum views as benign or even positive instances of the phenomenon.⁷¹ By contrast, Catherine MacKinnon's approach is to begin from the observation that there is a particular morally problematic phenomenon in the world, and to use the term *objectification* to refer to it.⁷² In what follows, I borrow MacKinnon's approach, and define *perspectival objectification* and *perspectival essentialism* as necessarily morally bad, restricting their uses accordingly. Little of philosophical importance rides on this decision, however, and the reader can choose broader definitions, which will involve specifying which instances of perspectival objectification and essentialism in fact constitute harm.

6.5.2. Clashing criteria for harm

A second objection is as follows. Earlier, I suggested that people can be harmed when they are not attended to in the way in which they wish. I then proceeded to offer different, non-subjective criteria for harm, such as that individuals can be harmed when a non-personhood-related trait of theirs is made more salient than a personhood-related trait. These criteria might clash. For instance, a person might desire for a non-personhood-related trait of theirs to be their most salient feature. Consider a woman, let's call her Chun, who desires for her appearance to be her most salient feature. She might find others attending primarily to her figure to be empowering, and desires for her body to be given relative salience over her personality.⁷³ In other words, Chun wishes for (what I am calling) a non-personhood-related trait of hers to be most salient in the minds

⁷¹ Nussbaum (1995).

⁷² MacKinnon (1987). The readings of both MacKinnon's and Nussbaum's projects here have been borrowed from Stock (2015).

⁷³ This is indeed the message that many so-called *women's magazines* push. See *Glamour* magazine's article on beauty products that supposedly empower women for an example of this (Kay, 2018).

of others. Although we are attending primarily to a non-personhood-related trait of Chun's, does the fact that we are respecting her wishes mean that we avoid perspectively objectifying, and therefore harming, her? My account, so goes this objection, looks incapable of adjudicating such clashes.

Whilst my account needs clarifying to be able to resolve such cases, various options are, I contend, open to me. In particular, how one adjudicates these sorts of cases depends on to what extent one thinks that agents can be *wrong* about what contributes to their well-being.⁷⁴ Can we say that Chun is wrong to say that having others attend to her body is empowering? An important point to note in this context is that a great many feminists have argued that our choices, preferences and desires are socially constructed, in the sense that society shapes them.⁷⁵ Where society is sexist, women can internalise sexist ideologies and end up choosing and desiring things that are congruent with sexism. For instance, due to internalising sexist body ideals for women, many women desire to be so thin that a medical professional would consider them seriously underweight.⁷⁶ Although these women's weight is often a result of their choices and desires, we can see in this instance how society might have perniciously influenced these women, so that they end up choosing and desiring something unhealthy and harmful.

We must be alive to the fact, then, that the woman who desires for her appearance to be her most salient attribute may not be aware of how society has perniciously shaped her desires in a way that harms her. Whilst an individual can arguably be harmed when she is not attended to in the way in which she wishes,⁷⁷ we may do that individual *greater* harm by heeding her wish, if that wish is for us to attend to her in a way that disrespects her personhood—namely, if that wish is for us perspectively to objectify her. This sort of conclusion is not an uncommon one. Many feminists recognise that a balancing act is necessary regarding respecting an agent's desires, choices, and preferences on the one hand, whilst, on the other hand, being critical of what has *caused* those desires, choices, and preferences.⁷⁸ Whilst being critical of a woman's choice might

⁷⁴ We might look to prevalent distinctions in contemporary ethics when deciding this issue, such as between *objectivist*, *informed desire*, and *subjectivist* accounts of well-being (see Crisp, 2017: §4).

⁷⁵ See, for example, Hirschmann (1996) and Hirshman (2006).

⁷⁶ McCarthy (1990).

⁷⁷ See the references in fn. 13.

⁷⁸ Thwaites (2017), for instance, thinks that we should not equally support all choices made by women. For her, those that act to extend inequality and maintain the patriarchal status quo deserve less support. Snyder-Hall (2010: 259) agrees, suggesting that we

look like a way of disrespecting her personhood – of disrespecting her ability to exercise her rationality, agency, and so on – the point here is that pernicious social forces might have already compromised these things, meaning that she has in fact exercised little rationality, autonomy, and so on.

Similar things could be said for those who wish to be perspectively *essentialised*. A woman, for instance, might wish for facts about her biology to be more salient than facts about the sexist social practices that influence her. In deciding whether to overrule these individuals' desires, we might investigate whether there are social narratives that have malignly shaped these individuals' choices. Perhaps, for instance, one might suggest that women can internalise essentialist social narratives that allegedly *celebrate* men and women's biological differences—narratives that suggest that women are especially valuable in society because of how their distinctive hormones make them nurturing and empathetic. One might claim that these narratives are not in women's interests.⁷⁹

Things are a little more complicated, I think, when it comes to applying this discussion to the first case study, involving rape victim-survivors. What are we to conclude about the rape victim-survivor who wishes for her experience of rape to be her most salient feature, finding her status as rape victim-survivor to be powerfully connected to her personhood? Should we respect her desire and make her experience of rape her most salient feature? Again, considering the specificities of each case will be important.

Consider, for instance, the individual whose experience of rape has allowed her to create support networks for other victim-survivors, and to spread awareness of her experience as a victim-survivor. In this case, I would suggest that the apparent conflict between wanting to respect another's personhood on the one hand, and respecting their desires on the other, might well be apparent instead of real. I wonder whether this individual is benefitting from others attending primarily to her status as an *informed political activist*, as opposed to rape victim-survivor. Making her experience of rape salient is simply *instrumental* in highlighting the political and/or social goals that she wishes to achieve, such as to improve the welfare of rape victim-survivors. Indeed, it is her informed activism oriented towards these goals that demonstrates her personhood. For instance, her activism is a career that, to a certain extent, she has chosen, demonstrating

should not celebrate all women's choices equally—only those made by women who recognise the politics of their choices.

⁷⁹ See, for instance, Giora (2002) for a discussion of these narratives, and for an argument that they are not in women's interests.

her agency and rationality.⁸⁰ Through her deciding to talk about her experience, she has been able to showcase her activist *interests* and resilient and charitable *personality*. Considered by itself, the status of rape victim-survivor, does not demonstrate her personhood in these ways. In fact, it is an identity that these activists are working hard to erase from the world, precisely *because* of the harm it does to the personhood of the one who is raped.⁸¹

As for the individual who wishes to make their experience of rape their most salient feature in the absence of these other factors, we might respond in different ways. We might, as with the example of Chun above, want to investigate whether there are exploitative ideologies that the individual might have internalised, meaning that her choice might not be a free one, made in her best interest. Perhaps we could envisage a scenario where a rape victim-survivor has internalised the messages about rape victim-survivors given in the various crime television shows and films mentioned by Brison in §6.2.1. In these shows and films, rape victim-survivors tend to be portrayed not just as a *passive pretext for our entertainment*⁸² they are often depicted in troublingly sexualised ways.⁸³ Given the complexity of the rape victim-survivor experience, however, and the special importance of giving rape victim-survivors the chance, in the charity *Rape Crisis's* words, to “feel in charge of their own lives again”, it is no doubt more important to take the rape victim-survivor’s choices and desires at face-value, and simply to support their decisions regarding how they wish to be attended to.⁸⁴

The account of perspectival harm that I have offered, then, leaves various options open regarding how to resolve clashes between one criterion of harm, namely, that an individual can be harmed when they are not attended to as they wish, and another—namely, the more objective criteria of harm offered above (which suggests that an individual can be harmed when a non-personhood-related trait of theirs is made more salient than one that does relate to their personhood, or that an individual can be harmed

⁸⁰ I add the qualifier *to a certain extent* as this individual may never have considered the career unless she had experienced rape, and she certainly did not choose to experience rape.

⁸¹ As we heard in §6.2.1, the harm in rape is regularly discussed in terms of the harm to an individual’s personhood. Cahill (2001: 13), for instance, says that a comprehensive account of the harm in rape must focus primarily on its attack of the “personhood of a woman”.

⁸² Brison (1993: 11).

⁸³ Vanstone (2016).

⁸⁴ Rape Crisis (2019b).

when their internal properties are made more salient than external factors contributing to their oppression).

6.6. Conclusion

In this chapter, I have expanded upon an idea that was first explored in chapter 3. This was that cognitive salience perspectives can *constitute* harm, in virtue of counting as a way of disrespecting an individual's personhood. Given that disrespect of one's personhood is a non-instrumental harm, attentional dispositions that count as a *form* of this harm can also be judged as non-instrumentally harmful. Harm can extend beyond the material, behavioural, or doxastic level, then, to include mere salience patterns in our attention.

I began this chapter with three questions. Firstly, I asked if the idea that mere salience can constitute harm is plausible. Here, I have tried to show that it is, by considering in depth how this idea relates to two case studies. By consulting rape-victim statements, I have shown that there appears to be a common thread through certain rape victim-survivor responses to rape, which the notion of salience perspective as constitutive harm can help to explain. In particular, the concern that one's experience of rape ends up masking other aspects of one's identity can be understood as a claim about attention—about which aspects of one's identity are most noticeable, memorable, and cognitively accessible to others. Further, by relating cognitive salience perspectives to the topic of objectification, I have highlighted that something *like* attention is already in the background of some women's complaints about how they are attended to, as well as in feminist discussions of subtle types of objectification.

Secondly, I asked what the notion of salience as constitutive harm offers to the subject of sex/gender harm. Here, I answered that salience perspectives might constitute one aspect of sex/gender harm. Firstly, certain attentional dispositions might count as a form of objectification. Simply making more salient a person's thing-like properties (such as their body) over their personhood-related properties (such as their voice), can *count* as a form of harmful objectification. I called this *perspectival objectification*. Secondly, I suggested that attentional dispositions might count as a form of sex/gender essentialism. Simply making women's biological features more salient than the social structures that oppress them, for instance, might count as a way of harmfully essentialising women. I called this *perspectival essentialism*.

Thirdly, I asked what, if anything, makes the notion of perspectival harm (and therefore these perspectival versions of sex/gender harms) important. The significance of perspectival harm, I answered, is in its minimalist nature. This makes it difficult to challenge, insofar as it is difficult to notice this harm occurring. Further, where it is noticed, the fact that it can be so easily dismissed as trivial gives it a further insidious power, insofar as it remains *effectively* unchallengeable.

CHAPTER 7

CONCLUSION

7.1. Introduction

In this thesis, I have brought together psychological research into framing effects and philosophical research into the notion of salience in the mind to develop the concept of a *salience perspective*. This refers to the structuring and/or presentation of some linguistic or mental contents, so that some contents are more prominent than others. In order to demonstrate more concretely why salience *matters*, I have applied the notion of a salience perspective to research into various folk-cultural biases, as well as feminist research into topics such as sexual objectification. These latter discussions have shown the breadth and depth of the types of problems that mere patterns of salience can cause and constitute.

In this final chapter, I begin in §7.2 by discussing the key points that have been made in this thesis. I move in §7.3 to consider the wider implications of these points, looking in particular at the consequences for other topics in philosophy not so far discussed, as well as those for social policy. In §7.4, I discuss the limitations of this thesis, and avenues for future research. In particular, I address: the limitations to my own self-reflection regarding the salience patterns used within this thesis; the possibility of future empirical studies testing the predictions that I have made; and the utility of future discussions that clarify the place that salience has in a broader movement for social justice.

7.2. Key points

This thesis can be understood as making seven key points.

1. When we communicate some linguistic contents, or when we consider some mental contents in our minds, we inevitably employ patterns of salience for those contents. With regard to linguistic contents, we must, for instance, say one sentence before another. With regard to mental contents, we invariably find some mental states, such as beliefs or experiences of properties in the world, more attention-grabbing than others. To refer to these patterns of salience, I have introduced the concept of a *salience*

perspective. Linguistic salience perspectives relate to cognitive salience perspectives in different ways. Firstly, they are connected by the relation of similarity. They are similar insofar as the notion of *relative salience* applies to them both; both are constituted by patterns of salience that involve giving relative salience to some (either linguistic or mental) contents over others. Secondly, they are causally connected. Making something relatively more salient in language tends to make that thing relatively more salient in its audience's minds. Conversely, having something salient in one's mind tends to result in one making that thing salient in language.

2. Merely changing the pattern of salience in some contents, without making any other changes to those contents, can result in the activation of a belief or ideology. This can occur, I argued, when one makes salient that which is central to that belief or ideology, and when that belief or ideology is cognitively accessible and socially licensed. This has epistemic repercussions when the belief is false, or the ideology unwarranted in some way. This also has ethical repercussions where the belief or ideology is harmful. Where a salience perspective is liable to lead to these negative consequences, we can criticise it on instrumental grounds.

3. Sometimes, we must cast a wide net when considering which problematic cultural beliefs and ideologies risk being activated by a given salience perspective. Instead of activating just one belief or ideology, certain patterns of salience might be capable of activating multiple sets of biases. In particular, *the substance salience perspective*, which is a single generic salience perspective emphasising a subject's internal properties, and the distinctive traits it happens to exhibit, is liable to activate at least three cognitively accessible and socially licensed problematic biases in Western cultures: the correspondence bias; psychological essentialism; and the fixed mindset. This adds to the literatures on these respective cognitive biases, by suggesting a new, subtle trigger for them. Taking a common reading of these biases, which suggests that they are composed of false and harmful beliefs, we can say that the substance salience perspective is instrumentally problematic on epistemic and ethical grounds.

4. Sometimes, it is instead helpful to take a narrower approach, and look in depth at a particular subject, and a specific cognitively accessible and socially licensed bias associated with it, when judging how to use salience for that subject. In particular, a

review of the particular topic of sex/gender, and the specific bias of sex/gender essentialism, suggests that giving relative salience to various contents, such as sex/gender binaries over non-binary conditions and identities, and sex/gender differences over similarities, can be sufficient to activate sex/gender essentialism. This proposal adds to existing literature on sex/gender essentialism by suggesting a new, subtle trigger for these beliefs. Using a common reading of this bias, which suggests that it is composed of false and harmful beliefs about sex/gender, we can judge that the salience perspective identified in this investigation is instrumentally problematic for epistemic and ethical reasons.

5. Salience perspectives can have epistemic and ethical significance insofar as they *constitute* an epistemic flaw, or harm. This was the most controversial and novel claim in this thesis. Focussing on harm, I suggested that making the wrong thing relatively more salient about a person can, for instance, count as a way of disrespecting their personhood. Given that disrespecting someone's personhood is usually considered to be a harm in itself, salience perspectives that count as a way of disrespecting someone's personhood also count as intrinsically harmful.

6. Perspectival versions of existing sex/gender harms might exist, such as perspectival objectification and perspectival essentialism. This suggestion marked a contrast with how these phenomena tend to be treated. Often, objectification is understood principally in terms of beliefs and behaviours (see, for instance, Sally Haslanger and Martha Nussbaum), or sometimes in terms of desires and associations (see, for example, Susan Bordo and Alison Assiter).¹ Essentialism, on the other hand, is overwhelmingly discussed as comprising beliefs (see, for instance, Gail Heyman and Jessica Giles, and Susan Gelman).² The findings here added to these existing accounts of objectification and essentialism, by suggesting new, particularly subtle and minimalist forms of these sex/gender harms. Objectification and sex/gender essentialism might instead consist, I contend, in certain patterns of salience, and these patterns can have intrinsic harms.

¹ Haslanger (1993), Nussbaum (1995: 257), Bordo (1993), Assiter (1988: 68).

² Heyman & Giles (2006), Gelman (2005).

7. Salience perspectives are important not least because their minimalist nature gives them an insidious power. Firstly, where a salience perspective *activates* a belief or ideology, it likely does so in an *under the radar* manner. This means that an individual has less power to block inferences to beliefs and ideologies that they might, if conscious of those inferences, reject. This makes salience perspectives that activate beliefs and ideologies especially *effective* in doing so. Secondly, cognitive salience perspectives that *constitute* harm (which includes perspectival objectification and essentialism) have an insidious power insofar as they are difficult to challenge, both because they are difficult to notice, and because they are easy to dismiss as trivial, making them *effectively* unchallengeable.

7.3. Wider implications of salience perspectives

Salience is pervasive, insofar as it structures all of our language and mental content. If patterns of salience have the powers that I have suggested, then this has at least one general implication: it is important for all of us, whether philosophers or otherwise, to reflect on the salience patterns that we inevitably employ. Below, I suggest more specific implications that salience perspectives have: those for certain debates within philosophy so far not addressed in this thesis, and those for social policy.

7.3.1. Race and hate speech

Whilst the case studies in this thesis have largely focussed on sex/gender, we could examine other topics, such as race, class, sexuality, and disability, with the suggestion that salience can cause and constitute harm in mind. Take, for instance, the subject of race. I briefly suggested some examples of people in chapter 6 whose race was made more salient than their profession, which I suggested could harm them. We might take this idea further, and consider whether the phenomenon of racism involves not just things like false and harmful beliefs,³ unwarranted and harmful feelings and emotions,⁴ and discriminatory behaviours,⁵ but wrongful salience patterns.

³ For instance, Appiah (1990) suggests that racism is partially constituted by false beliefs about the science of race (combined with certain evaluative or moral beliefs about moral standing and status).

⁴ For instance, Blum (2002) suggests that one form of racism is defined by feelings of antipathy, such as hostility and hatred.

⁵ Flew (1990) argues that racism is a matter of racially discriminatory behavior, instead of belief.

The finding that salience can constitute harm in particular would also affect certain other debates in ethics. Consider, for example, the topic of hate speech, which I will consider in a little more depth. Hate speech tends to receive criticism on two grounds. One is its liability to inculcate false and harmful beliefs in its audience. Jeremy Waldron, for instance, suggests understanding hate speech as group libel, whose harm consists in defaming members of a group through making *false* statements about them.⁶ Waldron considers a leaflet published in 1950s Chicago which urges people to protect the white race from being “mongrelized” and terrorised by the “rapes, robberies, guns, knives, and marijuana of the negro”.⁷ The harm of this hate speech, he suggests, is primarily in its *falsity*, and the damage that false assertion does to the reputation of black people.

Alternatively, hate speech is often criticised on the basis of the violence that it causes and licenses (i.e. legitimises). Lynne Tirrell, for instance, focuses on how hate speech can “[open] the door to previously prohibited [and violent] actions”.⁸ Looking at the Rwandan genocide of 1994, Tirrell considers hate speech targeted at the Tutsi people. She notes that Tutsi people are regularly called *inyenzi* (Kinyarwanda for *cockroach*). Tirrell suggests that, partly because of the social meanings associated with the word *cockroach* (e.g. that they are dirty and disgusting), uses of this word ended up licensing actions consistent with the Tutsi people *actually being* cockroaches, such as the mutilation of their bodies. Tirrell locates one harm of hate speech, then, in its ability to license violent, even genocidal, actions. Hate speech, can, in an important sense, *kill*.⁹

The discussion in chapter 6 suggests that we should consider the possibility of another harm. Hate speech might be harmful simply in virtue of inculcating the wrong pattern of attention in its audience. The Chicago leaflet, for instance, arguably helps its audience to find any instances of crimes committed by black people particularly cognitively accessible—more so than, say, any good deeds they might do. The Rwandan hate speech arguably helps its audience better to notice and remember traits associated with cockroaches in Tutsi people—more so than traits associated with their personhood. For instance, cockroaches are taken to be ubiquitous, dirty and disease-ridden. Calling Tutsi people *cockroaches* can serve simply to help one to notice crowds of Tutsi people,

⁶ Waldron (2012: 48).

⁷ Ibid.

⁸ Tirrell (2012: 175).

⁹ For this phraseology, see Tirrell’s participation in the podcast *Can Speech Kill?* (Philosophy Talk, 2017).

and to help one remember instances of unclean or unwell Tutsi people. In other words, hate speech might succeed simply in making certain traits of an individual, qua their group membership, particularly *salient*. This, as I have suggested, can *constitute* a way of harming them. It might, for instance, count as a way of disrespecting their personhood.

7.3.2. Social policy

Outside of philosophy, the ideas in this thesis also have implications for arenas like social policy. If mere changes in salience can both cause and constitute harm, then this is something that should influence how governments and corporations (for instance) communicate.

Fortunately, some organisations that aim to influence social policy have already picked up on the importance of salience. For instance, *The FrameWorks Institute* in Washington DC is a think tank that aims to improve the quality of communications in the non-profit sector, in a way that aims to “further public understanding of specific social issues”, such as mental health, climate change, and immigration.¹⁰ Their focus is on how, in their words, “the subtle selection of certain aspects of an issue [can] cue a specific response” in an audience.¹¹ Researching how shifts in patterns of salience can generate framing effects is a part of their research programme.¹² Further, as has been the focus in this thesis, The FrameWorks Institute examines how salience interacts with what they call “the public’s deeply held worldviews and widely held assumptions”, which might include the sort of cultural beliefs and ideologies mentioned in previous chapters.¹³ Echoing my discussion of instrumental critiques of salience perspectives, The FrameWorks Institute says, for instance, “without knowing the existing contours of public thinking and how unproductive cultural models are activated, policy experts and advocates often inadvertently trigger them [such as by using the wrong pattern of salience] in their public education and messaging efforts”.¹⁴

¹⁰ FrameWorks Institute (2017).

¹¹ FrameWorks Institute (2019a).

¹² FrameWorks Institute (2019b). They also consider topics beyond salience, such as how certain metaphors, narratives, visuals, tones, and so on, shape audience responses to some content.

¹³ Ibid.

¹⁴ Manuel & Arvizu (2010: 6).

The Behavioural Insights Team in the UK is a similar outfit.¹⁵ Part of their aim is to, in their words, “improve equality and tackle discrimination across society...[encouraging] change in workplaces, education establishments, households, public services and public spaces”.¹⁶ Like The FrameWorks Institute, the Behavioural Insights Team researches what it calls “simple changes to tackle major policy problems”.¹⁷ In other words, BIT markets itself as offering relatively minor, low-cost interventions that can nevertheless have powerful impacts.¹⁸ For instance, they look at the impact that making a retirement plan *opt out* instead of *opt in* has on participation rates, and what colour-coding nutritional information on food packets does to influence consumers to make healthier choices.¹⁹ Shifts in the salience perspective used to communicate some contents counts as one of these minor, low-cost interventions.

The research in this thesis adds to the call for more organisations like these. It also might add to these existing organisations, by clarifying which biases are likely to become triggered by making the wrong things salient. Consider, for instance, the suggestions in this thesis as to how the correspondence bias, psychological essentialism (including psychological essentialism about sex/gender in particular), and the fixed mindset can become activated through patterns of salience in language and the mind. Given their epistemic and ethical implications, these biases might count as the sort of *unproductive cultural models* that these institutions aim to prevent communicators activating through their speech.

Further, my finding that salience can *constitute harm* represents one way of expanding the focus of these institutions, whose current emphasis seems instead to be on the *epistemic* issues that salience can *cause*. For instance, The FrameWorks Institute focuses on furthering “the public *understanding* [emphasis added]” of various issues,²⁰ which seems to be an epistemic project. Further, they use *causal* language of *cueing*, *activating*, and *triggering*, when discussing how shifts in salience relate to the *unproductive cultural models* mentioned earlier.²¹ Why expand their focus in this way? As has been argued in other

¹⁵ BIT was created in 2010, partly in response to advice from Nobel-prize winning economist Richard Thaler, whose book *Nudge* (2008), coauthored with Cass Sunstein, brought *nudge theory* to prominence. Nudge theory is defined in §7.4.3 below.

¹⁶ BIT (2019a).

¹⁷ BIT (2019b).

¹⁸ They advertise their services as generating “impact quickly and cost-effectively” (BIT, 2019c).

¹⁹ Service et al. (n.d.).

²⁰ FrameWorks Institute (2017).

²¹ Manuel & Arvizu (2010: 6), FrameWorks Institute (2019a).

philosophical debates, offering a way of arguing that certain speech implicates harm in a way that does not rely on establishing *causal* relationships, can be very helpful. Take, for instance, the immigration article discussed in chapter 2, and the suggestion it might, by mentioning crimes that immigrants have committed *before* the benefits associated with immigration, causally activate xenophobic biases. Demonstrating this might be difficult. In Mary Kate McGowan's words, "As is well known, it is notoriously difficult to establish the truth of...complex [causal] claims".²² If we can instead suggest that this article *constitutes* a way of being xenophobic (i.e. that it *constitutes* a harm and/or epistemic bias), then one can condemn this article without needing to prove complex causal claims.

Generally, the various philosophical distinctions in this thesis add clarity to the critical approaches taken by these organisations. For instance, contrasts between epistemic and ethical, as well as instrumental and non-instrumental critiques of salience, as well as the contrasts between doxastic, and non-doxastic (including *perspectival*) forms of biases and harms suggested by this thesis, add precision and rigour.

7.4 Limitations and further research

7.4.1. The salience patterns in this thesis

In the introduction to this thesis, I commented on how the pervasive nature of salience will inevitably affect my own writing. I have spent a lot of time reflecting on the ways in which I myself am employing salience patterns in this thesis. For instance, I spent time editing my linguistic salience patterns in chapter 5, on sex/gender essentialism. In this editing process, I noticed that, when talking about sex/gender binaries, I would automatically and unreflectively make men/males more salient by mentioning them before women/females. I would talk, for instance, about similarities and differences between men/males and women/females.

Once I became conscious of this salience pattern in my writing, I reflected both on why it felt intuitive to employ it, and which biases it risks activating. We have already heard in the very chapter in question (§5.2) about the cognitive accessibility and social license of our androcentric biases, which treat men/males as the norm in our culture. I realised that my linguistic salience perspective made salient content central to this bias (namely, by talking about men/males *first*, it made them more salient than

²² McGowan (2005: 28). McGowan makes this point in relation to MacKinnon and Langton's discussions of the harms caused and constituted by pornography.

women/females). In addition to helping me make sense of why giving relative salience to men/males felt intuitive to me, recognising that my salience pattern made salient content central to androcentrism demonstrated the potential for my writing to activate androcentric biases in the reader's mind (and my own). Realising this, I changed the pattern of salience that I was using by *reversing* it. In other words, I mentioned women/females *before* men/males. One limitation of this research, therefore, is that to my own self-reflection. If I can uncritically use patterns of salience in the very chapter that explains why these salience patterns are harmful, it is inevitable that, despite my efforts, I will have employed other linguistic salience patterns that unhelpfully and harmfully pander to certain biases of ours.

I admit to this in part to demonstrate the difficulty of becoming aware of the harmful (and potentially unwarranted) salience patterns that we use—even those transparently written on the page in front of us. I have talked quite a bit about the subtle nature of salience perspectives, which does indeed make them difficult to notice. Further research would do well to illuminate effective ways of becoming *aware* of the salience patterns that we use in our language and attention. In chapter 3, I mentioned the potential utility of Cognitive Behavioural Therapy in the context of individual measures for *changing* entrenched cognitive salience perspectives of ours that we have deemed to be problematic. We might also look to this and related therapies for suggestions regarding how to first *notice* the salience patterns that we use (before we can begin to think about changing them). Indeed, research in this area has developed many techniques for *spotting* unhelpful patterns of thought, which we might apply to salience perspectives.²³ Often, however, as is the case in therapy, we will need someone else to point out the problematic salience perspectives that we use (when it comes to *linguistic* salience perspectives, an editor for our work could play this role).

7.4.2. Empirical studies

Another key limitation of this research is the lack of empirical studies testing the various hypotheses that I have made, such as those regarding how certain salience perspectives could activate sex/gender essentialist beliefs, or how the substance salience perspective

²³ Chapter 3 (§3.4.1). Cognitive Behavioural Therapy has developed techniques for making a person conscious of their unconscious habits, including habits of attention (Padesky, 1994). I also mentioned eye tracking studies and memory tests as potentially being of use when it comes to identifying our *cognitive* salience perspectives. Further research might also look to these techniques.

might activate various biases such as the correspondence bias. When making these predictions, I instead extrapolated from existing psychology studies into how mere shifts in salience perspective can mean the difference between activating certain cultural beliefs and ideologies and not. Ideally, further research would conduct the studies needed to support the predictions made in this thesis.

For instance, the potential studies that I discussed in §5.4.3 of chapter 5, inspired by that of Susan Bruckmüller and colleagues, would be carried out.²⁴ These aimed to test how salience patterns affect support for sex/gender essentialism. For instance, does mentioning women/females and men/males *before* non-binary identities and conditions affect how likely study participants are to endorse statements that intersex or genderqueer individuals are odd or aberrant in some way? We might also conduct studies that test for the claims in chapter 4, and present information about a person or group either by giving relative salience to the properties highlighted by the substance salience perspective, or to those highlighted by the process salience perspective. We might, for instance, talk about a person's character before their situational constraints, and vice versa, and see whether these different salience patterns affect the participants' endorsement of the correspondence bias, psychological essentialism, or fixed mindsets. Given the implicit nature of these various biases, we would need to make use of the various *implicit* measures for testing a person's attitudes that have been developed in psychology.²⁵

7.4.3. The limits of salience regarding social justice

Finally, further research could reflect on the *place* that salience perspective change has in a broader movement for social justice. Whilst learning about salience perspectives is helpful for a variety of reasons, including solely *epistemic* goals of using patterns of salience that are warranted, and that avoid the activation of false beliefs, many might find the subject of this thesis appealing for its contribution to ethical goals. Indeed, I have focussed on the *harms* that patterns of salience cause and contribute, given my own interest in minimising the harm that we do to others—especially disadvantaged social groups.

One problem that could be raised about the topic of linguistic salience perspectives in connection to social justice goals is its potential to be used by apologists for

²⁴ Bruckmüller et al. (2012).

²⁵ Rudman (2011).

conservative approaches to institutional and social reform. Consider *nudge theory*, which uses insights from behavioural economics to suggest that simple, low-cost interventions can influence people to think and act differently.²⁶ For instance, simply etching an image of a fly onto urinals has been found to reduce the spillage of urine around the urinals by 80%.²⁷ Framing effects (and therefore, what I call salience perspectives), are discussed as one type of nudge.²⁸ Chief political commentator for *The Observer* Andrew Rawnsley suggests that nudge theory has particular successes during recession periods, which is no coincidence given the “big appetite for low-cost solutions to public policy challenges”.²⁹ He continues: “Nudging appeared to offer easy ways of reforming society without committing to large spending programmes”.³⁰ Eduardo Porter of the *New York Times* also suggests that nudging “fosters a belief that tweaks based on an understanding of people’s psychology could lead to a vastly improved society at little to no cost to taxpayers”.³¹ Porter finds this belief inaccurate, insofar as it is unduly optimistic. He suggests that governments push this belief about nudging not necessarily because of a sincere belief in its power, but because nudging offers a way of saving money, and of pandering to the lack of political will to make more serious and far-reaching changes to how society is run. Changing linguistic salience perspectives offers a similar appeal; interventions that involve simply altering salience patterns in language are simple and low cost. The risk in this is that they detract from larger, more radical social reform that is needed. The utility of salience-based interventions could be used as an excuse by governments and corporations for avoiding more substantive measures, which require more effort, and cost more money.

Consider, for instance, the Behavioural Insights Team’s intervention in the Avon and Somerset Constabulary, aimed at increasing its racial diversity. This involved adding a particular sentence to a standard email sent to applicants during the application process—a sentence that attempted to create a friendlier tone in that email. (Whilst not an intervention involving salience perspectives, it was one of a similarly minimal nature.) For reasons that we will not go into here, this intervention managed to increase the probability that a Black or minority ethnic (BME) individual passed a certain stage in the

²⁶ Thaler & Sunstein (2008).

²⁷ Ibid. (3-4).

²⁸ Ibid. (36).

²⁹ Rawnsley (2017).

³⁰ Ibid.

³¹ Porter (2016).

application process by 50%.³² (Data has not yet been collected regarding whether this translated into more BME recruits, however.)

This certainly seems to be an admirable result. But the appeal of this intervention, namely, its simplicity and low-cost nature, might be a double-edged sword. These minimalist, headline-grabbing interventions might be used to distract from the lack of resources and effort the British Police Force is directing to fix the deeper problems that it faces. For instance, many have suggested that the British Police Force suffers from institutional racism, which disadvantages BME individuals in a great many ways.³³ Nifty, cheap prompts placed in emails should not be seen as *substitutes* for the large scale, difficult, and costly moves the police force would need to make to address these deeper, more complex ethical issues. To fix institutional racism, for instance, all sorts of different interventions would be needed beyond mere nudges (including those that play with salience), such as explicit and implicit bias training, the introduction of more counter-stereotypical exemplars into positions of power, the funding of outreach programmes, the overhauling of certain (e.g. stop and search) policies, and so on.³⁴

Further research, then, would clarify the precise role that salience perspectives can play in a wider movement for justice. In particular, it would suggest how best to communicate what makes the research in this thesis regarding linguistic salience perspectives *attractive*, without giving the impression that linguistic salience perspective change is a miracle cure. So long as it is understood that salience-based interventions are but *one* tool among many for achieving ethical (and prudential, and epistemic) goals, then they can play an important role in achieving those goals.

This is not to undermine the more radical nature of the claims made in this thesis. There *is* something importantly radical about the idea that simply changing the pattern of salience in some linguistic or mental content, without changing anything else about that content, can mean the difference between causing or constituting harm to a person or group, and not. This confronts conventional models of thought, such as those that treat individuals as rational and unaffected by mere presentational and structural shifts in some linguistic or mental content.³⁵ It confronts conventional models of harm, which instead locate harm in phenomena such as false beliefs, malign desires, or physical

³² BIT (2015: 29).

³³ See, for instance, Dearden (2018), Gayle (2018) and MacPherson (1999).

³⁴ See MacPherson (1999) for some of these suggestions and others.

³⁵ See chapter 1 (§1.2) and chapter 2 (§2.2.1).

violence.³⁶ Indeed, the fact that, even when one's beliefs, desires, feelings, and actions are ethically sound, one can *still* harm a person or group simply through the patterns of salience that one applies to them, is hardly a fact that apologists for conservative reform would like to dwell on. Neither would they like to dwell on how the pervasive nature of salience, in conjunction with its *under the radar* nature, means that it is likely that we are causing and constituting harm to people and social groups far more than we might have thought.

Whilst we might want to ensure that the attractiveness of salience-based interventions is not co-opted into cynical interests, then, the take-home conclusion from this thesis stands. Small can be powerful.

³⁶ See chapter 5, and the discussion of hate speech and racism in §7.3.1 in this chapter.

BIBLIOGRAPHY

- Adebisi., 2018. Staring at women is not a compliment, *African Feminism*, 11 Jun [online] Available at: <<http://africanfeminism.com/staring-at-women-is-not-a-compliment/>> [Accessed 20 January 2019].
- Ahlstrom-Vij, K., 2013. *Epistemic paternalism: A defence*, Basingstoke: Palgrave MacMillan.
- Althusser, L. 1965. Theory, Theoretical Practice and Theoretical Formation: Ideology and Ideological Struggle, Translated by J. H. Kavanagh, In: G. Elliott, ed, *Philosophy and the Spontaneous Philosophy of the Scientists*. New York, NY: Verso, 3–42.
- Anderson, E., 1995. Knowledge, Human Interests, and Objectivity in Feminist Epistemology, *Philosophical Topics*, 23(2), 27–58.
- Anscombe, E., 1957. *Intention*, Oxford: Blackwell.
- Antony, L., 2016. Bias: Friend or foe? Reflections on Saulish skepticism. In: M. Brownstein, M., & J. Saul, eds., *Implicit Bias and Philosophy: Metaphysics and Epistemology* (Vol. 1), Oxford: Oxford University Press, 157–190.
- 2012. Different voices or perfect storm: why are there so few women in philosophy?, *Journal of Social Philosophy*, 43(3), 227-55.
- Apfelbaum, E. P., Sommers, S. R., & Norton, M. I., 2008. Seeing race and seeming racist? Evaluating strategic colorblindness in social interaction, *Journal of Personality and Social Psychology Research*, 95(4), 918-32.
- Appiah, A., 1990. Racisms. In David T. Goldberg, ed., *The anatomy of racism*, Minneapolis, M.N.: University of Minnesota Press.
- Archer, D., Iritani, B., Kimes, D. D., & Barrios, M., 1983. Face-ism: Five studies of sex differences in facial prominence, *Journal of Personality and Social Psychology*, 45(4), 725-735
- Archer, S., forthcoming. *Salience: A Philosophical Inquiry*, Routledge.
- Aronson, J., Fried, C. B., & Good, C., 2001. Reducing the effects of stereotype threat on African American college students by shaping theories of intelligence, *Journal of Experimental Social Psychology*, 38, 113-25.
- Arthurs, J., & Zacharias, U., 2006. Introduction: Scientific/biological determinism: Media models of genetics and gender, *Feminist Media Studies*, 6(2), 221-30.

- Assiter, A., 1988. Autonomy and pornography. In: M. Griffiths & M. Whitford, eds., *Feminist perspectives in philosophy*, Bloomington: Indiana University Press.
- Banaji, M. R., Hardin, C., & Rothman, A., J., 1993. Implicit stereotyping in person judgment, *Journal of Personality and Social Psychology*, 65, 272-81.
- Bartky, S., 1990. *Femininity and domination: Studies in the phenomenology of oppression*, New York, NY: Routledge.
- Bastian, B., & Haslam, N., 2006. Psychological essentialism and stereotype endorsement, *Journal of Experimental Social Psychology*, 42, 228–235.
- Bauer, N., 2005. How to do things with pornography. In: S. Shieh & A. Crary, eds., *Reading Cavell*, London: Routledge, 69-97.
- BBC, 2019. *Growth mindsets: Help your child try new things* [online] Available at: <<https://www.bbc.co.uk/cbeebies/grownups/help-your-child-try-new-things>> [Accessed 1 February 2019]
- BIT (Behavioural Insights Team), 2019a. *Equality and diversity* [online] Available at: <<https://www.bi.team/what-we-do/policy-areas/equality-and-diversity/>> [Accessed 15 November 2018].
- 2019b. *About us* [online] Available at: <<https://www.bi.team/about-us/>> [Accessed 15 November 2018].
- 2019c. *Designing and testing interventions* [online] Available at: <<https://www.bi.team/what-we-do/our-services/designing-and-testing-behavioural-insights-interventions/>> [Accessed 15 November 2018].
- 2015. *Update Report 2013-5* [online] Available at: <http://www.behaviouralinsights.co.uk/wp-content/uploads/2015/07/BIT_Update-Report-Final-2013-2015.pdf> [Accessed 3 February 2019].
- Belkhir, J. A., & Duyme, M., 1998. Intelligence and race, gender, class: The fallacy of genetic determinism: Rethinking intelligence from the position of the oppressed, *Race, Gender and Class*, 5(3), 136-78.
- Belsky, M., 2019. *The headless women of Hollywood* [online] Available at: <<http://headlesswomenofhollywood.com/>> [Accessed 20 November 2018].
- Bem, S. L. 1993. *The lenses of gender: Transforming the debate on sexual inequality*, New Haven, CT: Yale University Press.
- 1989. Genital knowledge and gender constancy in preschool children, *Child Development*, 60(3), 649-62.

- Berglund, M., Westin, L., Svanström, R., & Sundler, A., 2012. Suffering caused by care—Patients' experiences from hospital settings, *International Journal of Qualitative Studies on Health and Well-Being*, 7, 1-9.
- Berndt, T. J., & Heller, H. K., 1986. Gender stereotypes and social inferences: A developmental study, *Journal of Personality and Social Psychology*, 50(5), 889-898.
- Bettcher, T. M., 2016. Intersexuality, transgender, and transexuality. In: L. J. Disch, & M. E. Hawesworth, eds., *The Oxford handbook of feminist theory*, New York, NY: Oxford University Press.
- Bigler, R. S., 1995. The role of classification skill in moderating environmental influences on children's gender stereotyping: A study of the functional use of gender in the classroom, *Child Development*, 66(4), 1072-87.
- Black, M.C., Basile, K.C., Breiding, M.J., Smith, S.G., Walters, M.L., Merrick, M.T., Chen, J., & Stevens, M.R., 2011. The National Intimate Partner and Sexual Violence Survey (NISVS): 2010 summary report, *National Center for Injury Prevention and Control, Centers for Disease Control and Prevention*. Atlanta, GA.
- Blackwell, L. S., Trzesniewski, K. H., & Dweck, C. S., 2007. Implicit theories of intelligence predict achievement across an adolescent transition: A longitudinal study and an intervention, *Child Development*, 78(1), 246-63.
- Blair, I. 2002. The Malleability of Automatic Stereotypes and Prejudice, *Personality and Social Psychology Review*, 3: 242-261.
- Blum, L., 2004. Stereotypes and stereotyping: A moral analysis, *Philosophical Papers*, 33(3), 251-89.
- 2002. What it is and what it isn't, *Studies in Philosophy and Education*, 21(3), 203-18.
- Bohmer, S., & Briggs, J. L., 1991. Teaching privileged students about gender, race, and class oppression, *Teaching Sociology*, 19(2), 154-63.
- Bommarito, N., 2013. Modesty as a virtue of attention, *Philosophical Review*, 122(1), 93-117.
- Bordo, S., 1993. *Unbearable weight*, Berkeley, CA: University of California Press.
- Brescoll, V. L., Dawson, E., Uhlmann, E. L., 2010. Hard won and easily lost: The fragile status of leaders in gender-stereotype-incongruent occupations, *Psychological Science*, 21(11), 1640-2.

- Brick, C., McNully, S. N., Updegraff, J.A., Ehret, P.J., Areguin, M.A., & Sherman, D.K., 2016. The impact of cultural exposure and message framing on oral health behaviour: exploring the role of message memory, *Medical Decision Making*, 36(7), 834-43.
- Brison, S., 1993. Surviving sexual violence: A philosophical perspective, *Journal of Social Philosophy*, 24(1), 5-22.
- Brownmiller, S., 1975. *Against our will: Men, women and rape*, New York, NY: Simon and Schuster.
- Brownstein, M., 2017. Implicit bias. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), [online] Available at: <<https://plato.stanford.edu/archives/spr2017/entries/implicit-bias/>> [Accessed 20 November 2017].
- Bruckmüller, S., Hegarty, P., & Abele, A., 2012. Framing gender differences: Linguistic normativity affects perceptions of power and gender stereotypes, *European Journal of Social Psychology*, 42, 210-18.
- Bruckmüller, S., Ryan, M., Haslam, S. A., Peters, K., 2013. Ceilings, cliffs and labyrinths: exploring metaphors for workplace gender discrimination. In: C. Harland, M. Ryan, G. Nassimbeni, N. Branscombe, & E. Schneller, eds., *The SAGE Handbook of Gender and Psychology*. London: SAGE publications.
- Burgess, D., Borgida, E., 1999. Who women are, who women should be: Descriptive and prescriptive gender stereotyping in sex discrimination, *Psychology, Public Policy, and Law*, 5(3), 665-92.
- Burkeman, O., 2018. Dirty secret: Why is there still a housework gender gap? *The Guardian*, 17 Feb [online] Available at: <<https://www.theguardian.com/inequality/2018/feb/17/dirty-secret-why-housework-gender-gap>> [Accessed 20 December 2018].
- Burriss, R., 2017. Eye-gaze and attraction, *Medium*, 9 Oct [online] Available at: <<https://medium.com/@RobertBurriss/eye-gaze-and-attraction-97955ddd217c>> [Accessed 20 December 2018].
- Busch, B., 2018. Research every teacher should know: Growth mindset, *The Guardian*, 4 Jan [online] Available at: <<https://www.theguardian.com/teacher-network/2018/jan/04/research-every-teacher-should-know-growth-mindset>> [Accessed 1 February 2019].
- Butcher, R., 2018. More than a third of all trans people suffered hate crimes in 2017, research suggests. *The Independent* [online] Available at:

<<https://www.independent.co.uk/news/uk/crime/trans-hate-crimes-stats-figures-2017-more-than-third-a8166826.html>> [Accessed 20 January 2019].

Cahill, A., 2001, *Rethinking rape*, Ithaca, NY: Cornell University Press.

Camp, E., 2017. Perspectives in imaginative engagement with fiction, *Philosophical Perspectives*, 31, 73-102.

Caplan, P. J., MacPherson, G. M., & Tobin, P., 1985. Do sex-related differences in spatial abilities exist? A multilevel critique with new data, *American Psychologist*, 40(7), 786-99.

Castilla, E. 2008. Gender, Race, and Meritocracy in Organisational Careers, *American Journal of Sociology*, 113(6), 1479-526.

CBC News, 2017. *'That Doesn't Define Me': What Life is Like After Speaking Out About Sexual Assault*. Canada, 18 Jun [online]. Available at:
<<https://www.cbc.ca/news/canada/newfoundland-labrador/life-after-publicly-speaking-about-sexual-assault-1.4163264>> [Accessed 3 August 2018].

Chaffee, S., & Schleuder, J., 1986. Measurement and effects of attention to media news, *Human Communication Research*, 13(1), 76-107.

Chappell, R. Y., & Yetter-Chappell, H. 2016. Virtue and salience, *Australasian Journal of Philosophy*, 94(3), 449-63.

Chen, W. & Wong, Y., 2015. Chinese mindset: Theories of intelligence, goal orientation and academic achievement in Hong Kong students, *Educational Psychology*, 35(6), 714-25.

Choi, I., Nisbett, R. E., & Norenzayan, A., 1999. Causal attribution across cultures: variation and universality, *Psychological Bulletin*, 125(1), 47-63.

Cimpian, A., & Saloman, E., 2014. The inherence heuristic: An intuitive means of making sense of the world, and a potential precursor to psychological essentialism, *Behavioural and Brain Sciences*, 37, 461-527.

Coltrane, S., 1989. Household labor and the routine production of gender, *Social Problems*, 36(5), 473-90.

Condit, C., Ofulue, N., Sheedy, K., M., 1998. Determinism and mass-media portrayals of genetics, *American Journal of Human Genetics*, 62(4), 979-84.

Coplan, A., 2004. Empathetic engagement with narrative fictions, *The Journal of Aesthetics and Art Criticism*, 62(2), 141-52.

Cottrel C. A., & Neuberg, S. L., 2005. Different emotional reactions to different groups: A sociofunctional threat-based approach to “prejudice”, *Journal of Personality and Social Psychology*, 88(5), 770-89.

Crisp, R., 2017. Well-being. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), [online] Available at: <<https://plato.stanford.edu/entries/well-being/>> [Accessed 15 December 2018].

Cudd, A., 2005. How to explain oppression: Criteria of adequacy for normative explanatory theories, *Philosophy of the Social Sciences*, 35(1), 20-49.

D’Arms, J., & Jacobson, D., 2000. The moralistic fallacy: on the ‘appropriateness’ of emotions, *Philosophy and Phenomenological Research*, 61(1), 65-90.

Danaher, K., & Crandall, C. S., 2008. Stereotype threat in applied settings re-examined, *Journal of Applied Social Psychology*, 38(6), 1639-1655.

Dar-Nimrod, I., & Heine, S. J., 2011. Genetic essentialism: On the deceptive determinism of DNA. *Psychological Bulletin*, 137(5), 800–18.

--- 2006. Exposure to scientific theories affects women’s math performance, *Science*, 314(5798), 435.

Darwall, L., 1977. Two kinds of respect, *Ethics*, 88: 36–49.

Dearden, S., 2018. British police ‘still institutionally racist’, senior officer admits as new recruitment strategy launched, *The Independent*, 12 Oct [online] <<https://www.independent.co.uk/news/uk/crime/police-racist-uk-recruitment-bame-institutionally-npcc-strategy-women-disability-a8581646.html>> [Accessed 12 December 2018].

Deaux, K., & Lewis, L. L., 1984. Structure of gender stereotypes: Interrelationships among components and gender label, *Journal of Personality and Social Psychology*, 46(5), 991-1004.

Delise, K., 2007. *The Pit Bull placebo: The media, myths and politics of canine aggression*, Manonville, NY: Anubis Press.

Demoulin, S., Leyens, J-P., & Yzerbyt, V., 2006. Lay theories of essentialism, *Group Processes and Intergroup Relations*, 9(1), 25-42.

Devine, P. G., 1989. Stereotypes and prejudice: their automatic and controlled components, *Journal of Personality and Social Psychology*, 56, 5-18.

Dotson, K., 2011. Tracking epistemic violence, tracking practices of silencing, *Hypatia*, 26(2), 236-57.

- Dupré, J., 2001. *Human nature and the limits of science*, Oxford: Oxford University Press.
- 1993. *The disorder of things: Metaphysical foundation of the disunity of science*, Cambridge, MA: Harvard University Press.
- 1986. Sex, gender, and essence, *Midwest Studies in Philosophy*, 11(1), 441-57.
- Dupré, J., & Nicholson D. J., 2018. A manifesto for a processual philosophy of biology. In: D. J. Nicholson, & J. Dupré, eds., *Everything flows: Towards a processual philosophy of biology*, Oxford: Oxford University Press, Ch. 1.
- Dweck, C. S., 2006a. *Mindset: the new psychology of success*. New York, NY: Random House.
- 2006b. Is maths a gift? Beliefs that put females at risk. In: S.J. Ceci & W. Williams, eds., *Why aren't more women in science? Top researchers debate the evidence*, Washington, D.C.: American Psychological Association, 47-55.
- 2000. *Self theories: Their role in motivation, personality, and development*, Philadelphia, PA: Taylor & Francis.
- Dweck, C., Chiu, C., Hong, Y., 1995. Implicit theories and their role in judgements and reactions: A world from two perspectives. *Psychological Inquiry*, 6(4), 267-85.
- Dweck, C., & Leggett, E., 1988. A social-cognitive approach to motivation and personality, *Psychological Review*, 95(2), 256-73.
- Dworkin, A., 2000. Against the male Flood: Censorship, pornography, and equality. In: D. Cornell, (ed.), *Feminism: Feminism and Pornography*, Oxford: Oxford University Press, 19-44.
- Eagly, A. H., Wood, W., & Diekmann, A. B., 2000. *Social role theory of sex differences and similarities: A current appraisal*. In T. Eckes, & H. M. Trautner, eds., *The developmental social psychology of gender*, Mahwah, New Jersey: Erlbaum, 123-174.
- Eaton, A., 2016. Bodily taste and fat oppression. In: S. Irvin, *Body Aesthetics*, Oxford: Oxford University Press.
- Edwards, A., 1987. Male violence in feminist theory: An analysis of the changing conceptions of sex/gender violence and male dominance. In: J. Hanmer, & M. Maynard, eds., *Women, violence and social control: Explorations in sociology*, London: Palgrave Macmillan, 13-29.
- Eidson, R., & Coley, J., 2014. Not so Fast: reassessing Gender essentialism in young adults, *Journal of Cognition and Development*. 0(0), 1-11.
- Elgin, C., 1996. *Considered Judgement*, Princeton: Princeton University Press.

Emerson, K. T., Murphy, M. C. 2015. A company I can trust? Organisational lay theories moderate stereotype threat for women, *Personality and Social Psychology Bulletin*, 41(2), 295-307.

--- 2013. *How organizational lay theories of intelligence impact the motivation and performance of Black and Latino individuals*, unpublished manuscript, Indiana University, Bloomington

Entman, R., 1993. Framing: Toward clarification of a fractured paradigm, *Journal of Communication*, 43(4), 51-8.

Fausto-Sterling, A., 2000. *Sexing the body: Gender politics and the construction of sexuality*, New York, NY: Basic Books.

--- 1993. The five sexes: Why male and female are not enough, *The Sciences*, 33, 20-24.

--- 1985. *Myths of gender: Biological theories about women and men*, New York, NY: Basic Books.

Feinberg, J., 1987. *The moral limits of the criminal law: Volume 1: Harm to Others*, Oxford: Oxford University Press.

Feng, J., I. Spence, and J. Pratt., 2007. Playing an action video game reduces gender differences in spatial cognition, *Psychological Science*, 18, 850–855.

Fine, C., 2010. *Delusions of gender*, New York, NY: Norton and Company.

Fine, C., Jordan-Young, R., Kaiser, A., & Rippon, G., 2013. Plasticity, plasticity, plasticity . . . and the rigid problem of sex, *Trends in cognitive sciences*, 17(11), 550-551.

Fischhoff, B., 1977. Perceived informativeness of facts, *Journal of Experimental Psychology: Human Perception and Performance*, 3(2), 349-58.

Fisher, 2017. *Philosophical perspectives on framing: Investigating inequivalence* [Poster] Exhibited at: 25th Annual Meeting of the European Society for Philosophy and Psychology, *University of Hertfordshire*, 14th-17th August 2017.

Flew, A. 1990. Three concepts of racism, *Encounter* 75, 63–66.

Foot, P. 1967, The problem of abortion and the doctrine of double effect, *Oxford Review*, 5, 5–15.

FrameWorks Institute, 2019a. *About Strategic Frame Analysis®* [online] Available at: <<https://www.frameworksinstitute.org/sfa-overview.html>> [Accessed 24 December 2018].

--- 2019b. *Frame Effects* [online] Available at: <<https://www.frameworksinstitute.org/frame-effects.html>> [Accessed 24 December 2018].

--- 2017. *Mission* (3 June 2017 update) [online] Available at:
<<https://www.frameworksinstitute.org/mission.html>> [Accessed 23 December 2018].

Fraser, R., 2018. The ethics of metaphor, *Ethics*, 128(4), 728-55.

Freed, A. F., 1996. Language and gender research in an experimental setting. In: V. L. Bergvall, J. M. Bing, & A. F. Freed, eds., *Rethinking language and gender research: Theory and practice* (2nd ed.), London: Longman, 54-76.

Frisson, S., & Wakefield, M., 2012. Psychological essentialist reasoning and perspective taking during reading: a donkey is not a zebra, but a plate can be a clock, *Memory and Cognition*, 40(2), 297-310.

Frye, M., 1983. *The politics of reality: Essays in feminist theory, freedom*, Trumansburg, NY: The Crossing Press.

Fuss, D. 1989. *Essentially speaking: Feminism, nature, and difference*, New York, NY: Routledge.

Gaither, S. E., Schultz, J. R., Pauker, K., Sommers, S., R., Maddox, K. B., & Ambady, N., 2014. Essentialist thinking predicts decrements in children's memory for racially-ambiguous faces, *Developmental Psychology*, 50(2), 482-88.

Garcia-Retamero, R., & Dhami, K. D., 2013. On avoiding framing effects in experienced decision makers, *Quarterly Journal of Experimental Psychology*, 66(4), 829-42.

Gaunt, R., 2006. Biological essentialism, gender ideologies, and role attitudes: What determines parents' involvement in child care, *Sex Roles*, 55, 223-533.

Gawronski, B., 2007. Correspondence bias. In: R. F. Baumeister, & K. D. Vohs, eds., *Encyclopedia of social psychology*, Thousand Oaks, CA: SAGE Publications, 194-5.

--- 2004. Theory-based bias correction in dispositional inference: The fundamental attribution error is dead, long live the correspondence bias, *European Review of Social Psychology*, 15(1), 183-217.

Gayle, D., 2018. Structural racism at heart of British society, UN human rights panel says, *The Guardian*, 27 Apr [online] Available at:
<<https://www.theguardian.com/world/2018/apr/27/racism-british-society-minority-ethnic-people-dying-excessive-force>> [Accessed 12 December].

Gelman, S., 2005. Essentialism in everyday thought. *Psychological Science Agenda* (APA Online), 19(5), [online] Available At:
<<http://www.apa.org/science/about/psa/2005/05/gelman.aspx>> [Accessed 28 December 2017].

- 2003. *The essential child: Origins of essentialism in everyday thought*. Oxford: Oxford University Press.
- Gelman, S. A., & Heyman, G. D., 1999. Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories, *Psychological Science*, 10, 489-493.
- Gelman, S., & Hirschfeld, L., 1999. How biological is essentialism? In: D. L. Medin & S. Atran, eds., *Folkbiology*, Cambridge, MA: MIT Press, 403-46.
- Gelman, S., & Taylor, M., 2000. Gender essentialism in cognitive development. In: P. H. Miller, & E. Kofsky Scholnick, eds., *Towards a Feminist Developmental Psychology*. New York, NY: Routledge, 169-90.
- Gelman, S., & Wellman, H., 1991. Insides and essences, *Cognition*, 38, 214-44.
- Gendler, T., 2008. Alief and belief, *Journal of Philosophy*, 105(10), 634-63.
- Gernsbacher, M. A., Hill Goldsmith, H., & Robertson, R. R. W., 1992. Do readers mentally represent characters' emotional states? *Cognition and Emotion*, 6, 89-111.
- Gervais, S., Holland, A., & Dodd, M., 2013. My eyes are up here: The nature of the objectifying gaze toward women, *Sex Roles*, 69, 557-70.
- Giacomini, M., Rosée-Koker, P., Pepitone-Arreola-Rockwell, F., 1986. Gender bias in human anatomy textbook illustrations, *Psychology of Women Quarterly Banner*, 10(4), 413-20.
- Gibbard, A., 1990. *Wise choices, Apt feelings*, Cambridge, MA: Harvard University Press.
- Gilbert, D. T., & Malone. P. S., 1995. The correspondence bias, *Psychological Bulletin*, 117(1), 21-38.
- Gilovich, T., & Eibach, R., 2001. The fundamental attribution error: Where it really counts, *Psychological Inquiry*, 12(1), 23-6.
- Giora, R., 2002. Theorizing gender: Feminist awareness and language change. In: B. Baron, & H. Kotthoff, eds., *Perspectives on femininity and masculinity in ethnography and discourse*, Amsterdam: John Benjamins Publishing.
- Giordani, A., 2015. A suitable semantics for implicit and explicit belief, *Logique et Analyse* 58(231).
- Goh-Mah, J., 2013. The objectification of women goes further than sexy pictures, *HuffPost*, 09 June [online] Available at: <<https://www.huffingtonpost.co.uk/joy-goh->

mah/objectification-women-sexy-pictures_b_3403251.html> [Accessed 28 August 2018].

Good, C., Aronson, J., Inzlicht, M., 2003. Improving adolescents' standardized test performance: An intervention to reduce the effects of stereotype threat, *Applied Developmental Psychology*, 24, 645-62.

Good, C., Rattan, A., & Dweck, C. S., 2012. Why do women opt out? Sense of belonging and women's representation in mathematics, *Journal of Personality and Social Psychology*, 102(4), 700-17.

--- 2007. *Theories of intelligence influence females' sense of belonging, intent to continue, and achievement in math*, unpublished manuscript, Columbia University, New York.

Gramsci, A., 1971 [1929-1935]. *Selections from the prison notebooks*. Translated and edited by Q. Hoare & G. Nowell Smith. New York, NY: International Publishers.

Grant, H., & Dweck, C. S., 2003. Clarifying achievement goals and their impact, *Journal of Personality and Social Psychology*, 85(3), 541-53.

Gray, J., 1992. *Men are from Mars, women are from Venus: A practical guide for improving communication and getting what you want in your relationships*, New York, NY: Harper Collins.

Greenberg, D. M., Warrier, V., Allison, C., & Baron-Cohen, S., 2018. Testing the Empathizing–Systemizing Theory of sex differences and the Extreme Male Brain theory of autism in half a million people, *PNAS Proceedings of the National Academy of Sciences of the United States of America*, 115(48), 12152-12157.

Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R., 2009. Understanding and using the implicit association test: III. Meta-analysis of predictive validity, *Journal of Personality and Social Psychology*, 97(1), 17-41.

Griffiths, P. E., 2002, What is innateness? *The Monist*, 85(1), 70-85.

The Guardian, editorial, 2018. The Guardian view on a Conservative crisis: One of the Tories' own making, *The Guardian*, 30 Sep [online] Available at: <<https://www.theguardian.com/commentisfree/2018/sep/30/the-guardian-view-on-a-conservative-crisis-one-of-the-tories-own-making>> [Accessed 10 December 2017].

Haier, R., Karama, S., Leyba, L., & Jung, R. 2009. MRI assessment of cortical thickness and functional activity changes in adolescent girls following three months of practice on a visual-spatial task, *BMC Research Notes*, 2,174.

Halwani, R., 2010. *Philosophy of love, sex, and marriage: An introduction*, New York, NY: Routledge.

- Hamamura, T., Meijer, Z., Heine, S. J., Kamaya, K., & Hori, I., 2009. Approach—avoidance motivation and information processing: A cross-cultural analysis, *Personality and Social Psychology Bulletin*, 35(4), 454-62.
- Harman, G. 1999. Moral philosophy meets social psychology: Virtue ethics and the fundamental attribution error, *Proceedings of the Aristotelian Society*, 99, 315–31.
- Hasinoff, A., 2009. It's sociobiology, hon!: Genetic gender determinism in *Cosmopolitan Magazine*, *Feminist Media Studies*, 9(3), 267-83.
- Haslam, N., 2011. Genetic essentialism, neuroessentialism, and stigma: Commentary on Dar-Nimrod and Heine (2011), *Psychological Bulletin*, 137, 819–824.
- Haslam, N., Bastian, B., Bain, P., & Kashima, Y., 2006. Psychological essentialism, implicit theories, and intergroup relations, *Group Processes & Intergroup Relations*, 9(1), 63-76.
- Haslam, N., & Rothschild, L., & Ernst, D. 2000. Essentialist beliefs about social categories, *British Journal of Social Psychology*, 39, 113-27.
- Haslam, N., & Whelan, J., 2008. Human natures: Psychological essentialism in thinking about differences between people, *Social and Personality Psychology Compass*, 2(3), 1297-312.
- Haslanger, S., 2015. Distinguished lecture: Social structure, narrative and explanation. *Canadian Journal of Philosophy*, 45(1), 1-15.
- 2013. Social meaning and philosophical method, *American Philosophical Association*, 110th Division Annual Meeting, (December 2013).
- 2011. Ideology, generics, and common ground. In: C. Witt, ed., *Feminist metaphysics: Explorations in the ontology of sex, gender, and the self*, Dordrecht: Springer, Ch. 11.
- 2006. Philosophical analysis and social kinds: What good are our intuitions?, *Proceedings of the Aristotelian Society* (Sup. Vol. 80), 89-119.
- 1993. On being objective and being objectified. In: L. M. Antony & C. Witt, eds., *A mind of one's own: Feminist essays on reason and objectivity*, Oxford: Westview Press.
- Haugtvedt, C. O., & Wegener, D. T., 1994. Message order effects in persuasion: an attitude strength perspective, *Journal of Consumer Research*, 21, 205-18.
- Hegarty, P., & Buechel, C. 2006. Androcentric reporting of gender differences in APA Journals: 1965–2004, *Review of General Psychology*, 10, 377–389.
- Heidari, S., Babor, T. F., de Castro, P., Tort, S., & Curno, M., 2016. Sex and gender equity in research: Rationale for the SAGER guidelines and recommended use, *Research Integrity and Peer Review*, 1(2), doi: 10.1186/s41073-016-0007-6.

- Heilman, M., 2012. Gender stereotypes and workplace bias, *Research in Organizational Behaviour*, 32, 113-35.
- 2001. Description and prescription: How gender stereotypes prevent women's ascent up the organizational ladder, *Journal of Social Issues*, 57(4), 657-74.
- Heilman, M. E., & Okimoto, T. G. 2007. Why are women penalized for success at male tasks?: The implied communality deficit, *Journal of Applied Psychology*, 92, 81-92.
- Heine, S.J., Lehman, D., R., Markus, H. R., & Kitayama, S., 1999. Is there a universal need for a positive self-regard? *Psychological Review*, 106(4), 766-94.
- Hennessey, J., 2018. Mindsets and the learning environment: How do classroom environments influence underserved students' mindsets and academic performance? *Mindset Scholars Network*, 21 Feb [online] Available at: <<http://mindsetscholarsnetwork.org/mindsets-learning-environment-classroom-environments-influence-underserved-students-mindsets-academic-performance/>> [Accessed 6 February 2019].
- Hernstein, R. J., & Murray, C., 1994. *The bell curve: Intelligence and class structure in American life*, New York, NY: Free Press Paperbacks.
- Heyman, G. D., & Giles, J. W., 2006. Gender and psychological essentialism, *Enfance; Psychologie, Pédagogie, Neuropsychiatrie, Sociologie*, 58(3), 293-310.
- Hirschmann, N., 1996. Toward a feminist theory of freedom, *Political Theory*, 24(1), 46-67.
- Hirshman, L., 2006. *Get to work: A manifesto for women of the world*, New York, NY: Viking.
- Hoffman, C., & Hurst, N. 1990. Gender stereotypes: Perception or rationalization? *Journal of Personality and Social Psychology*, 58, 197-208.
- Holbrook, A. L., & Krosnick, J. A., Carson, R. T., & Mitchell, R. C., 2000. Violating conversational conventions disrupts cognitive processing of attitude questions, *Journal of Experimental Social Psychology*, 36, 465-94.
- Holroyd, J., 2012. Responsibility for implicit bias, *Journal of Social Philosophy*, 43(3), 247-306.
- Honneth, A. 1995. *The struggle for recognition: The grammar of social conflicts*, Cambridge: Polity Press.
- Hood-Williams, J., 1997. Real sex/ fake gender: A reply to Robert Willmott, *The Sociological Review*, 45(1), 42-58.

- Hoyt, C. L., & Murphy, S. E., 2016. Managing to clear the air: Stereotype threat, women, and leadership, *The Leadership Quarterly*, 27, 387-99.
- Humberstone, I. L. 1992. Direction of fit, *Mind*, 101(401), 59-84.
- Hussak, L., & Cimpian, A., 2017. Memory accessibility shapes explanation: Testing key claims of the inherence heuristic account, *Memory and Cognition*, 46(1), 68-88.
- Hyde, J. S., 2014. Gender similarities and differences, *Annual Review of Psychology*, 65, 373-98.
- 2005. The gender similarities hypothesis, *American Psychologist*, 60(6), 581-92.
- Hyde, J. S., Fennema, E., & Lamon, S. J., 1990. Gender differences in mathematics performance: A meta-analysis, *Psychological Bulletin*, 107, 139-155.
- Hyde, J. S., & Mezulis, A. H., 2002. Gender difference research: Issues and critique. In: J. Worrell, ed., *Encyclopedia of women and gender: Sex similarities and differences and the impact of society on gender*, New York, NY: Academic Press, 551-60.
- Isen, A. M., & Levin, P. F., 1972. Effects of feeling good on helping: Cookies and kindness, *Journal of Personality and Social Psychology*, 21, 384-88.
- James, K., 2000. "You can *feel* them looking at you": The experiences of adolescent girls at swimming pools, *Journal of Leisure Research*, 32(2), 262-80.
- Jamieson, A., & Radick, G., 2017. Genetic determinism in the genetics curriculum: An explanatory study of the effects of Mendelian and Weldonian emphases, *Science and Education*, 26(10), 1261-90.
- 2013. Putting Mendel back in his place: How curriculum reform in genetics and counterfactual history of science can work together. In: L. Kampourakis, ed. *The philosophy of biology*, Netherlands: Springer, 577-95.
- Jean-Charles, R. M., 2014. Toward a victim-survivor narrative: Rape and form in Yvonne Vera's *Under the Tongue* and Calixthe Bayala's *Tu t'appelleras Tanga*, *Research in African Literatures*, 45(1), 39-62.
- Jensen, A. R., 1969. How much can we boost IQ and scholastic achievement? *Harvard Educational Review*, 39, 1-123.
- Jordan-Young, R. 2010. *Brain storm: The flaws in the science of sex differences*, Cambridge, MA: Harvard University Press.
- Kahneman, D., & Tversky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-91.

Kampourakis, K., Reydon, T., A. C., Patrinos, G. P., & Strasser, B. J., 2014. Genetics and society—Educating scientifically literate citizens: Introduction to the thematic issues, *Science and Education*, 23, 251-8.

Kang, J. & Banaji, M. 2006. Fair measures: A behavioral realist revision of ‘affirmative action’, *California Law Review*, 94, 1063-1118.

Kant, I. *Lectures on Ethics*, Translated by L. Infield, New York, NY: Harper and Row, 1963.

--- 1785. *Groundwork of the metaphysics of morals*, Translated by M. Gregor, Cambridge: Cambridge University Press, 1997.

Karsay, K., Matthes, J., Platzer, P., & Plinke, M., 2017. Adopting the objectifying gaze: exposure to sexually objectifying music videos and subsequent gazing behaviour, *Media Psychology*, 21(1), 27-49.

Kay, C., 2018. Our favourite beauty product heroes that give back to empower women, *Glamour Magazine*, 8 March [online] Available at: <<https://www.glamourmagazine.co.uk/article/empowering-beauty-products>> [Accessed 10 September].

Keil, F. C., 1995. The growth of causal understandings of natural kinds. In: D. Sperber, D. Premack, & A. Premack, eds., *Causal cognition: A multidisciplinary debate*, Oxford: Oxford University Press, 234-62.

Keller, J., 2005. In genes we trust: The biological component of psychological essentialism and its relationship to mechanisms of motivated social cognition, *Journal of Personality and Social Psychology*, 88(4), 686-702.

Kelly, L., Burton, S., & Regan, L., 1996. Beyond victim or survivor: Sexual violence, identity and feminist theory and practice. In: L. Adkins, & V. Merchant, eds., *Sexualising the social: Power and the organisation of sexuality*, London: Palgrave Macmillan, Ch. 4.

Kelly, L. Lovett, J., & Regan, L., 2005. Gap or chasm?: Attrition in reported rape cases, *Home office research study no. 293*. London: Home Office [online] Available at: <<http://webarchive.nationalarchives.gov.uk/20110218141141/http://rds.homeoffice.gov.uk/rds/pdfs05/hors293.pdf>> [Accessed 20 August 2018].

The Kennel Club, 2019. “*Dangerous dogs*”: *Deal with the deed, not the breed*. [online] Available at: <<https://www.thekennelclub.org.uk/our-resources/kennel-club-campaigns/dangerous-dogs/>> [Accessed 6 February 2019].

Kihlstrom, J. F., 2004. Implicit methods in social psychology. In: C. Sansone, C. C. Morf, & A. T. Panter, eds., *The Sage handbook of methods in social psychology*, Thousand Oaks, CA: SAGE Publications, 195-212.

Kirp, D. L., 2016. Nudges that help struggling students succeed, *The New York Times*, 29 Oct [online] Available at: <<https://www.nytimes.com/2016/10/30/opinion/nudges-that-help-struggling-students-succeed.html>> [Accessed 1 July 2017].

Kitcher, P., 2001. Battling the undead: how (and how not) to resist genetic determinism. In: R. Singh, K. Krimbas, D. Paul & J. Beatty, eds., *Thinking about evolution: Historical, philosophical, and political perspectives (Festschrift for Richard Lewontin)*, Cambridge: Cambridge University Press, 396-414.

Korsgaard, C., 1996. *The sources of normativity*, Cambridge: Cambridge University Press.

Kray, L. J., Howland, L., Russell, A. G., & Jackman, L. M., 2017. Theories on gender system justification: Fixed beliefs strengthen masculinity to preserve the status quo, *Journal of Personality and Social Psychology*, 112(1), 98–115.

Kray, L. J., Thompson, L., & Galinsky, A. 2001. Battle of the sexes: Gender stereotype confirmation and reactance in negotiations, *Journal of Personality and Social Psychology*, 80, 942–958.

Kühlberger, A., 1998. The influence of framing on risky decisions: a meta-analysis, *Organizational Behaviour and Human Decision Making Processes*, 75, 355-73.

Langton, R., 2008. Comments on A.W. Eaton's 'A Sensible Antiporn Feminism', *Symposia on Gender, Race, and Philosophy*, 4(2), 1-5.

--- 2005. Feminism in philosophy. In: F. Jackson & M. Smith, eds., *The Oxford handbook of contemporary philosophy*, Oxford: Oxford University Press, Ch. 9.

--- 2004. Projection and objectification. In: B. Leiter, ed., *The future for philosophy*, Oxford: Oxford University Press, 285-303.

--- 2000. Feminism in epistemology: exclusion and objectification. In: M. Fricker & J. Hornsby, eds., *The Cambridge companion to feminism in philosophy*, Cambridge: Cambridge University Press, 127-45.

--- 1993a. Beyond a pragmatic critique of reason, *Australasian Journal of Philosophy*, 71, 364-84.

--- 1993b. Speech acts and unspeakable acts, *Philosophy and Public Affairs*, 22(4), 293-330.

Langton, R., & Dupré, J., 2013. Gender differences all in the mind [Letter to the editor] *The Guardian*, 5 Dec [online] Available at: <<https://www.theguardian.com/world/2013/dec/05/gender-difference-all-in-mind>> [Accessed 30 April 2018].

Langton, R., & West, C., 1999. Scorekeeping in a pornographic language game,

Australasian Journal of Philosophy, 77, 303-19.

LeBoeuf, R. A., & Shafir, E., 2003. Deep thoughts and shallow frames: On the susceptibility to framing effects, *Journal of Behavioural Decision Making*, 16(2), 77-92.

Leslie, S. J., 2013. Essence and natural kinds: When science meets preschooler intuition. In: T. S. Gendler & J. Hawthorne, *Oxford studies in epistemology* (Vol. 4), Oxford: Oxford University Press, 108-165.

--- 2007. Generics: Cognition and acquisition, *Philosophical Review*, 117(1), 1-47.

Levin, I., 1987. Associative effects of information framing, *Bulletin of Psychonomic Society*, 25(2), 85-6.

Levin, I., & Gaeth, G., 1988. How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, 15(3), 374-8.

Levinson, J., & Peng, K., 2007. Valuing cultural differences in behavioural economics, *ICFAI Journal of Behavioural Finance*, 4, 32-47.

Lewis, D., 1979. Scorekeeping in a language game, *Journal of Philosophical Logic*, 8(1): 339-359.

Lewontin, R., 2000. *The triple helix: Gene, organism, and environment*, Cambridge, MA: Harvard University Press.

--- 1970. Race and Intelligence, *Bulletin of the Atomic Scientists*, 26(3), 2-8.

Loke, W. H., & Tan, K. F., 1992. Effects of framing and missing information in expert and novice judgment, *Bulletin of the Psychonomic Society*, 30(3), 187-90.

Lopez, T., 2016. Following rape, 'runner' rather than 'victim' defines survivor's identity, *U.S. Army*, 21 April [online]. Available at:
<https://www.army.mil/article/166506/following_rape_runner_rather_than_victim_defines_survivors_identity> [Accessed 2 August 2018].

Lutz, T., 2018. Simone Biles: I was sexually abused by USA gymnastics doctor Larry Nassar. *The Guardian*, 15 Jan [online]. Available at:
<<https://www.theguardian.com/sport/2018/jan/15/simone-biles-larry-nassar-sexual-abuse-allegations-usa-gymnastics>> [Accessed 3 August 2018].

Lyness, K., & Heilman, M., 2006. When fit is fundamental: Performance evaluations and promotions of upper-level female and male managers, *Journal of Applied Psychology*, 91(4), 777-85.

MacKinnon, C. A., 2016. Rape redefined, *Harvard Law and Policy Review*, 10(2), 431-77.

--- 1987. *Feminism unmodified: Discourses on life and law*, Cambridge, MA: Harvard University Press.

MacPherson, W. 1999. *The Stephen Lawrence Inquiry*. London: Home Office. CM 4262-I. Available at: <<https://www.gov.uk/government/publications/the-stephen-lawrence-inquiry>> [Accessed 12 January 2018].

Mahalingam, R., 2003. Essentialism, culture, and beliefs about gender among the Aravanis of Tamil Nadu, India. *Sex Roles*, 49 (9-10), 489-496.

Mandel, D., 2014. Do framing effects reveal irrational choice? *Journal of Experimental Psychology: General*, 143(3), 1185-98.

Maney, D.L. 2016. Perils and pitfalls of reporting sex differences, *Philosophical Transactions of the Royal Society, Series B*, 371, 20150119.

Manuel, T., & Arvizu, S., 2010. *Helping immigration experts change the public conversation: A study circle evaluation brief*, Washington, D.C.: FrameWorks Institute.

Marcus, D. E., & Overton, W. F. 1978. The development of cognitive gender constancy and sex role preferences, *Child Development*, 49, 434-444.

Markus, H. R., & Kitayama, S. 1991. Culture and the self: implications for cognition, emotion, and motivation, *Psychological Review*, 98(2), 224-53.

Martin, C. L., Eisenbud, L., & Rose, H., 1995. Children's gender-based reasoning about toys, *Child Development*, 66(5), 1453-71.

Mayr, E., 1991. *One long argument: Charles Darwin and the genesis of modern evolutionary thought*, Cambridge, MA: Harvard University Press.

McCarthy, M. 1990. The thin ideal, depression, and eating disorders in women, *Behavioural Research and Therapy*, 28(3), 205-15.

McGowan, M. K., 2005. On Pornography: MacKinnon, speech acts, and "false" construction, *Hypatia*, 20(3), 22-49.

McKinnon, R., 2014. Stereotype threat and attributional ambiguity for trans women, *Hypatia*, 29(4), 857-72.

Medin, D., & Atran, S., eds., 1990. *Folkbiology*, Cambridge, MA: MIT Press.

Medin, D. L., & Ortony, A. 1989. Psychological essentialism. In: S. Vosniadou & A. Ortony, eds., *Similarity and analogical reasoning*, Cambridge: Cambridge University Press, 179-195.

- Mikkola, M., 2017. Feminist perspectives on sex and gender. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), [online] Available at: <<https://plato.stanford.edu/entries/feminism-gender/>> [Accessed 16 January 2018].
- Miller, C. C., Quealy, K., & Sanger-Katz, M., 2018. The top jobs where women are outnumbered by men named John, *The New York Times*, 24 April [online] Available at: <<https://www.nytimes.com/interactive/2018/04/24/upshot/women-and-men-named-john.html>> [Accessed 20 December 2017].
- Miller, D., Taylor, B., & Buck, M. L. 1991. Gender gaps: Who needs to be explained? *Journal of Personality and Social Psychology*, 61, 5–12.
- Miller, J. G., 1984. Culture and the development of everyday social explanation, *Journal of Personality and Social Psychology*, 46(5), 961-78.
- Minto, C. L., Liao, L. M., Woodhouse, C. R., Ransley, P. G., & Creighton, S. M., 2003. The effects of clitoral surgery on sexual outcome in individuals who have intersex conditions with ambiguous genitalia: A cross-sectional study, *Lancet*, 12(361), 1252-7.
- Miyamoto, Y., & Kitayama, S., 2002. Cultural variation in correspondence bias: The critical role of attitude diagnosticity of socially constrained behaviour, *Journal of Personality and Social Psychology*, 83(5), 1239-48.
- Morton, T., Postmes, T., Haslam S., & Hornsey, M., 2009. Theorising gender in the face of social change: Is there anything essential about essentialism? *Journal of Personality and Social Psychology*, 96(3), 653-64.
- Nelson, J. A., 2016. Not-so-strong evidence for gender differences in risk taking, *Feminist Economics*, 22(2), 114-42.
- Nelson, T. E., Clawson, R. A., & Oxley, Z. M. 1997. Media framing of a civil liberty conflict and its effects on tolerance, *American Political Science Review*, 91, 567-83.
- Nisbett, R. E., & Norenzayan, A. 2002. Culture and cognition. In: D. L. Medin, ed., *Stevens' handbook of experimental psychology: Cognition* (3rd ed.), New York, NY: Wiley, 561–97.
- Norenzayan, A., Choi, I., & Nisbett, R. E., 2002. Cultural similarities and differences in social inference: Evidence from behavioral predictions and lay theories of behaviour, *Personality and Social Psychology Bulletin*, 28(1), 109-20.
- Nussbaum, M., 1995. Objectification, *Philosophy and Public Affairs*, 24(4), 249-91.
- NWLC (National Women's Law Center), 2017. *Undervalued: A brief history of women's care*

work and child care policy in the United States [online] Available at: <https://nwlc.org/wp-content/uploads/2017/12/final_nwlc_Undervalued2017.pdf> [Accessed 2 February 2019].

O'Connor, C., & Joffe, H., 2014. Gender on the brain: A case study of science communication in the new media environment, *PLoS ONE*, 9(10), 1-15.

O'Connor, D. 2010. This is what happens when you forget about gender, *American Journal of Bioethics*, 10(7), 27-9.

ONS (Office for National Statistics), 2018. *Sexual offences in England and Wales: Year ending March 2017*. UK [online].

Available at:

<https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/sexualoffencesinenglandandwales/yearendingmarch2017>

[Accessed 1 September 2018].

--- 2013. *An overview of sexual offending in England and Wales*. UK [online].

Available at:

<<http://webarchive.nationalarchives.gov.uk/20140711232238/https://www.gov.uk/government/publications/an-overview-of-sexual-offending-in-england-and-wales>>

[Accessed 1 September 2018].

Olby, R. 1979. Mendel no Mendelian? *History of Science*, 17, 55–72.

Olivola, C., & Machery, E., 2014. Is psychological essentialism an inherent feature of human cognition? *Behavioural and Brain Sciences*, 35(5), 499.

Osborne, S., 2016. Stanford university rape case: Victim's letter in full, *The Independent*, 06 Jun [online] <Available at:

<https://www.independent.co.uk/news/world/americas/stanford-university-rape-case-the-victims-letter-in-full-a7067146.html>> [Accessed 2 August 2018].

Otieno, C., Spada, H., Renkl, A., 2013. Effects of news frames on perceived risk, emotions, and learning, *PLoS ONE*, 8(11), 1-12.

Oyama S., Griffith P.E. & Gray R.D., eds. 2003. *Cycles of contingency: Developmental systems and evolution*. Cambridge, MA: MIT Press.

Padesky, C. A., 1994. Schema change processes in cognitive therapy, *Clinical Psychology and Psychotherapy*, 1(5), 267-78.

Papadaki, E., 2018. Feminist perspectives on objectification. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), [online] Available at:

<<https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=feminism-objectification>> [Accessed 15 July 2018].

- Philosophy Talk, 2017. *Can speech kill?* 10 December [podcast] Available at: <<https://www.philosophytalk.org/shows/can-speech-kill>> [Accessed 15 August 2018].
- Pollock, P., & Parker, R., 1981. *Old mistresses: Women, art, and ideology*, London: Routledge & Kegan.
- Porter, E., 2016. Nudges aren't enough for problems like retirement savings, *The New York Times*, 23 Feb [online] Available at: <<https://www.nytimes.com/2016/02/24/business/economy/nudges-arent-enough-to-solve-societys-problems.html>> [Accessed 15 Jan 2019].
- Prentice, D., & Miller, D., 2006. Essentializing differences between women and men, *Psychological Science*, 17(2), 129-35.
- Preves, S., 1999. For the sake of the children: Destigmatizing intersexuality. In: A. D. Dreger., ed., *Intersex in the age of ethics*, Hagerstown: University Publishing Group, 51-8.
- Puddifoot, K., 2017. Dissolving the epistemic/ethical dilemma over implicit bias, *Philosophical Explorations*, 20 (Sup. Vol. 1), 73-93.
- Radick, G., 2016. Presidential address: Experimenting with the scientific past, *British Journal for the History of Science*, 49(2), 153-72.
- Ramsden, S., Richardson, F. M., Josse, G., Thomas, M. S., Ellis, C., Shakeshaft, C., Seghier, M. L., & Price, C. J., 2011. Verbal and non-verbal intelligence changes in the teenage brain, *Nature*, 479(7371), 113-16.
- Rape Crisis, 2019a. *Rape crisis England and Wales headline statistics 2017-8* [online] Available at: <<https://rapecrisis.org.uk/get-informed/rcew-statistics/>> [Accessed 1 September 2018].
- 2019b. *Supporting a survivor* [online] Available at: <<https://rapecrisis.org.uk/get-help/looking-for-information/supporting-a-survivor/>> [Accessed 1 September 2018].
- Rawnsley, A., 2017. Why nudge theory works until a kick in the backside is needed, *The Guardian*, 22 Oct [online] Available at: <<https://www.theguardian.com/commentisfree/2017/oct/21/nobody-minds-a-gentle-nudge-except-in-the-wrong-direction>> [Accessed 3 January 2019].
- Richard, M., 2008. *When truth gives out*, Oxford: Oxford University Press.
- Richards, C., Bouman, W. P., & Barker, M-J., 2017. *Genderqueer and non-binary genders*, London: Palgrave MacMillan.

- Richardson, S., 2013. *Sex itself: The search for male and female in the human genome*, Chicago: University of Chicago Press.
- Rippon, G., Jordan-Young, R., Kaiser, A., & Fine, C., 2014. Recommendations for sex/gender neuroimaging research: Key principles and implications for research design, analysis, and interpretation, *Frontiers in Human Neuroscience*, 8(650), 1-13.
- Risberg G., Johansson E.E., Hamberg K., 2009. A theoretical model for analysing gender bias in medicine, *International Journal for Equity in Health*, 8:28.
- Roberts, T-A., 2002. III. The woman in the body, *Feminism & Psychology*, 12(3), 324–29.
- Robertson, D., 2010. *The philosophy of Cognitive-Behavioural Therapy (CBT): Stoic philosophy as rational and cognitive psychotherapy*, London: Karnac Books.
- Robinson, H. 2018. Substance. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), [online] Available at: <<https://plato.stanford.edu/archives/win2018/entries/substance/>> [Accessed 3 January 2019].
- Roese, N., Sherman, J. W., & Hur, T., 1998. Direction of comparison asymmetries in relational judgment: The role of linguistic norms, *Social Cognition*, 16(3), 353-62.
- Roszkowski, M. J., & Snelbecker, G. E., 1990. Effects of framing on measures of risk tolerance: Financial planners are not immune, *Journal of Behavioural Economics*, 19, 237-46.
- Rothbart, M., & Taylor, M., 1992. Category labels and social reality: Do we view social categories as natural kinds? In: G. R. Semin & K. Fiedler, eds., *Language and social cognition*. London: SAGE Publications, 11-36.
- Ruble, D. N., & Martin, C. L. 1998. Gender development. In: W. Damon & N. Eisenberg, eds., *Handbook of child psychology: Social, emotional, and personality development*, Hoboken, NJ, US: John Wiley & Sons Inc. 933-1016.
- Rudman, L. 2011. *Implicit measures for social and personality psychology*. London: SAGE Publications.
- Rudman, L., & Phelan, J. E., 2008. Backlash effects for disconfirming gender stereotypes in organizations, *Research in Organizational Behaviour*, 28, 61-79.
- Rybash, J. M., & Roodin, P., 1989. The framing heuristic influences judgments about younger and older adults' decision to refuse medical treatment. *Applied Cognitive Psychology*, 3, 171-80.
- Rydgren, J., 2004. The logic of xenophobia, *Rationality and Society*, 16(2), 123-48.

Sample, I., 2013, Male and female brains wired differently, scan reveals, *The Guardian*, 2 Dec [online] Available at: <<https://www.theguardian.com/science/2013/dec/02/men-women-brains-wired-differently>> [Accessed 30 April 2018].

Saul, J. 2013. Implicit bias, stereotype threat, and women in philosophy. In: F. Jenkins & K. Hutchison, eds., *Women in philosophy: What needs to change?*, Oxford: Oxford University Press.

--- 2012. Scepticism and implicit bias, *Disputatio*, 5(37), 243-63.

Sbisà, M. 1999. Ideology and the persuasive use of presupposition. In: J. Verschueren, ed., *Language and ideology: Selected papers from the 6th International Pragmatics Conference* (Vol. 1), Antwerp: International Pragmatics Association, 492-509.

Schmidt, S., Langner, S., Hennigs, N., Wiedmann, K. P., Karampournioti, E., & Lischka, G., 2017. The green brand: Explicit and implicit framing effects of ecolabelling on brand knowledge, *Cogent Psychology*, 4(1), 1329191.

Schwitzgebel, E., & Cushman, F., 2012. Expertise in moral reasoning? Order effects on moral judgements in professional philosophers and non-philosophers, *Mind and Language*, 27(2), 135-53.

Science Careers Staff, 2018. Celebrating women in science, *Science*, 9 Feb [online] Available at: <<https://www.sciencemag.org/careers/2018/02/celebrating-women-science>> [Accessed 3 January 2019].

Scopelliti, H., Min, L., McCormick, E., Kassam, K. S., & Morewedge, C. K., 2017. Individual differences in correspondence bias: Measurement, consequences, and correction of biased interpersonal attributions, *Management Science*, 1-32.

Seibt, J., 2018. Process philosophy. In: E. N. Zalta, ed. *The Stanford Encyclopedia of Philosophy* (Winter 2018 Edition), [online] Available at: <<http://plato.stanford.edu/archives/fall2013/entries/process-philosophy/>>. [Accessed 3 January 2019].

Service, O., Hallsworth, M., Halpern, D., Algate, F., Gallagher, R., Nyugen, S., Ruda, S., Sanders, M., Pelenur, M., Gyani, A., Harper, H., Reinhard, J., & Kirkman, E., n.d. *EAST: Four simple ways to apply behavioural insights*, London: Behavioural Insights Team/Cabinet office.

Sher, S., & McKenzie, C., 2008. Framing effects and rationality. In: N. Chater & M. Oaksford, eds., *The probabilistic mind*, New York, NY: Oxford University Press, Ch. 4.
--- 2006. Information leakage from logically equivalent frames, *Cognition*, 101, 467-94.

Siegel, S., 2017. *The rationality of perception*, New York, NY: Oxford University Press.

- Simon-Vandenberg, A-M., White, P. R. R., & Aijmer, K., 2007. Presupposition and 'taking-for-granted' in mass communicated political argument: An illustration from British, Flemish and Swedish political colloquy. In: A. Fetzer, & G. E. Lauerbach, eds., *Discourse in the media: Cross-cultural perspectives*, Amsterdam: John Benjamins Publishing.
- Simpson, R., 2012. Un-ringing the bell: McGowan on oppressive speech and the asymmetric pliability of conversations, *Australasian Journal of Philosophy*, 91(3), 555-75.
- Sinnott-Amstrong, W. 2008. Framing moral intuitions. In: W. Sinnott-Amstrong, ed., *Moral Psychology: The cognitive science of morality: intuition and diversity* (Vol. 2), Cambridge, MA: MIT Press. 107-44.
- Sisk, V. F., Burgoyne, A. P., Sun, J., Butler, J. L., & Macnamara, B. N., 2018. To what extent and under which circumstances are growth mind-sets important to academic achievement? Two meta-analyses, *Psychological Science*, 29(4), 549-71.
- Skewes, L., Fine, C., & Haslam, N., 2018. Beyond Mars and Venus: The role of gender essentialism in support for gender inequality and backlash, *PLoS ONE*, 13(7), e0200921.
- Snyder-Hall, C., 2010. Third-wave feminism and the defence of "choice", *Perspectives on Politics*, 8(1), 255-61.
- Stanley, J. S., 2015. *How propaganda works*, Princeton, N. J.: Princeton University Press.
- Stanley, J., S., & Blanchard-Fields, F., 2011. Beliefs about behaviour account for Age differences in the correspondence bias, *Journal of Gerontology Series B: Psychological Sciences and Social Sciences*, 66B(2), 169-76.
- Stewart, B. D., & Payne, B. K., 2008. Bringing automatic stereotyping under control: Implementation intentions as efficient means of thought control, *Personality and Social Psychology Bulletin*, 34,1332-1345.
- Stewart, H., & Mason, R., 2016. Nigel Farage's anti-migrant poster reported to police, *The Guardian*, 16 June [online] Available at: <<https://www.theguardian.com/politics/2016/jun/16/nigel-farage-defends-ukip-breaking-point-poster-queue-of-migrants>> [Accessed 1 February 2019].
- Stock, K., 2015. Sexual objectification, *Analysis*, 75(2).191-195.
- Stotzer, R. L., & MacCartney, D., 2016. The role of institutional factors on on-campus reported rape prevalence, *Journal of Interpersonal Violence*, 31(16), 2687-707.
- Suarez, E., & Gadalla, T. M., 2010. Stop blaming the victim: A meta-analysis on rape myths, *Journal of Interpersonal Violence*, 25(11), 2010-35.

- Sunstein, C., 2016. *The ethics of influence*, New York, NY: Cambridge University Press.
- Swanson, E., forthcoming. Slurs and ideologies. In: R. Celikates, S. Haslanger, and J. Stanley, (eds.). *Ideology*. Oxford: Oxford University Press.
- Szillis, U., & Stahlberg, D., 2007. The face-ism effect in the Internet differences in facial prominence of women and men, *International Journal of Internet Science*, 2(1), 3-11.
- Tabery, J., 2014. *Beyond versus: The struggle to understand the interaction of nature and nurture*, Cambridge, MA: MIT Press.
- Taylor, M. G., 1996. The development of children's beliefs about social and biological aspects of gender differences, *Child Development*, 67, 1555-71.
- Taylor, M. G., & Gelman, S. A., 1999. Children's beliefs about animal and human development: The role of category-membership and environment, Unpublished manuscript, University of Puget Sound, Washington.
- Terborg, J. R. 1977. Women in management: A research review, *Journal of Applied Psychology*, 62(6), 647-64.
- Thaler, R., & Sunstein, C., 2008. *Nudge: Improving decisions about health, wealth and happiness*, New Haven, C. T.: Yale University Press.
- Thomson, J. J., 1985. The Trolley Problem, *Yale Law Journal*, 94, 1395-415.
- Thwaites, R. 2017. Making a choice or taking a stand? Choice feminism, political engagement and the contemporary feminist movement, *Feminist Theory*, 18(1), 55–68.
- Till, B. D., Baack, D., & Waterman, B., 2011. Strategic brand association maps: Developing brand insight, *Journal of Product and Brand Management*, 20(2), 92-100.
- Tirrell, L., 2013. Studying genocide: A pragmatist approach to action-engendering discourse. In: G. Hubbs and D. Lind., eds., *Pragmatism, law, and language*, London: Routledge.
- 2012. Genocidal language games. In: I. Maitra & M. K. McGowan, eds., *Speech and harm: Controversies over free speech*, Oxford: Oxford University Press, 174-221.
- 1999. Racism, sexism, and the Inferential Role Theory of meaning. In: C. Hendricks, & K. Oliver, eds., *Language and liberation: Feminism, philosophy, and language*, Albany, NY: State University of New York Press.
- Tonic, G., 2018. 8 Plus-size bloggers rate Netflix's *Dumplin'*, *Refinery 29*, 10 Dec [online] Available at: <<https://www.refinery29.com/en-gb/plus-size-women-netflix-dumplin>> [Accessed 1 February 2019].

TSB, 2018. *Research by TSB finds societal stigma stops new dads from taking shared parental leave*, 12 Feb [online] Available at: <<https://www.tsb.co.uk/news-releases/new-fathers-fear-societal-stigmata-so-dont-take-shared-parental-leave/>> [Accessed 10 June 2018].

Tversky., & Kahneman, D., 1986. Rational choice and the framing of decisions, *The Journal of Business*, 59(4), pt. 2, 251-78.

--- 1981. The framing of decisions and the psychology of choice, *Science*, 211, 453-58.

Tworek, C. M., Cimpian, A., 2016. Why do people tend to infer “ought” from “is”? The role of biases in explanation, *Psychological Science*, 27(8), 1109-22.

University of Leeds, 2019. *Genetics Pedagogies Project*. [online] Available at: <<https://arts.leeds.ac.uk/geneticspedagogiesproject/>> [Accessed 1 February 2019].

Uttal, D. H., 1997. Beliefs about genetic influences on mathematics achievement: a cross-cultural comparison, *Genetica*, 99, 165-72.

Vanstone, E., 2016. How has rape become such a common trope of television drama?, *The Globe and Mail*, 16 Sep [online] Available at:

<<https://www.theglobeandmail.com/arts/television/how-has-rape-become-such-a-common-trope-of-television-drama/article31931181/>> [Accessed 20 May 2017].

Vanwesenbeeck, I., 2009. Doing gender in sex and sex research, *Archives of Sexual Behaviour*, 38(6), 883-98

Verkuten, M., 2003. Discourses about ethnic group (de-)essentialism: Oppressive and progressive aspects, *British Journal of Social Psychology*, 42, 371-91.

Wahl-Jorgensen, K. & Hanitzsch, T., 2009. *The handbook of journalism studies*, New York, NY: Routledge.

Waldron, J., 2012. *The harm in hate speech*, Cambridge, MA: Harvard University Press.

Walton, G. M., & Cohen, G. L., 2007. A question of belonging: race, social fit, and achievement, *Journal of Personality and Social Psychology*, 92(1), 82-96.

Watzl, S., forthcoming. Norms of attention: A classification. In: A. Appiah, J. Davis, & J. Ganari, eds., *Virtues of Attention. Cosmopolitan Perspectives*.

--- 2017. *Structuring mind: The nature of attention and how it shapes consciousness*, Oxford: Oxford University Press.

--- 2011. The nature of attention, *Philosophy Compass*, 6(11), 842-53.

Wearing, D., 2016. Racism and xenophobia are resurgent in the UK, and the centre-left is partly to blame, *Open Democracy UK*, 11 July [online] Available at:

<<https://www.opendemocracy.net/uk/david-wearing/racism-and-xenophobia-are-resurgent-in-uk-and-centre-left-is-partly-to-blame>> [Accessed 15 February 2019].

Wells, C., 2013. Microaggressions in the context of academic communities, *Seattle Journal for Social Justice*, 12(2), 319-48.

West, C., 1989. *The American evasion of philosophy*. New York, NY: University of Wisconsin Press.

Whatley, M. A., 1996. Victim characteristics influencing attributions of responsibility to rape victims: A meta-analysis, *Aggression and Violent Behavior*, 1(2), 81-95.

Wheeler, S. C., & Petty, R. E., 2001. The effects of stereotype activation on behaviour: A review of possible mechanisms. *Psychological Bulletin*, 127(6), 797-826.

Whiteley, E. K., forthcoming. Harmful salience perspectives, In: S. Archer, ed., *Salience: A Philosophical Inquiry*, Routledge.

Willett, C., 2016. Feminism and the self. In: E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), [online] Available at: <<https://plato.stanford.edu/cgi-bin/encyclopedia/archinfo.cgi?entry=feminism-self>> [Accessed 30 November 2017].

Williams, M., & Eberhardt, J. L., 2008. Biological conceptions of race and the motivation to cross racial boundaries, *Journal of Personality and Social Psychology*, 94(6), 1033-47.

Wilson, A. E., & English, J. A., 2017. The motivated fluidity of lay theories of change. In: C. M. Zedelius, Müller, B. C. N., & Schooler, J. W., 2017. *The Science of Lay Theories: How Beliefs Shape Our Cognition, Behavior, and Health*, Cham, Switzerland: Springer, 17-43.

Wing Sue, D., 2010. *Microaggressions in everyday life: Race, gender, and sexual orientation*, Hoboken, NJ: John Wiley and Sons.

Wittenbrink, B., & Schwarz, N., eds., 2007. *Implicit measures of attitudes*, New York, NY: Guilford.

Yeager, D., Walton, G., & Cohen, G. L., 2013. Addressing achievement gaps with psychological interventions, *Phi Delta Kappan*, 94(5), 62-5.

Yip, M. C., 2018. *Cognition, metacognition and academic performance: An East Asian perspective*, Oxon: Routledge.

Young, I. M., 1990. *Justice and the politics of difference*, Princeton, N. J.: Princeton University Press.

Yzerbyt, V., Rocher, S., & Schadron, G., 1997. Stereotypes as explanations: A subjective essentialistic view of group perception. In R. Spears, P. J. Oakes, N. Ellemers, & S. A. Haslam, eds., *The social psychology of stereotyping and group life*, Malden, MA: Blackwell Publishing, 20-50.

Yzerbyt, V. Y., Rogier, A., & Fiske, S. T. 1998. Group entitativity and social attribution: On translating situational constraints into stereotypes, *Personality and Social Psychology Bulletin*, 24, 1089-1103.