

ORDER-BASED SALIENCE PATTERNS IN LANGUAGE: WHAT THEY ARE AND WHY THEY MATTER

ELLA KATE WHITELEY

Sheffield Methods Institute, University of Sheffield

Whenever we communicate, we inevitably have to say one thing before another. This means introducing particularly subtle patterns of salience into our language. In this paper, I introduce ‘order-based salience patterns,’ referring to the *ordering* of syntactic contents where that ordering, pretheoretically, does not appear to be of consequence. For instance, if one is to describe a colourful scarf, it wouldn’t seem to matter if one were to say it is ‘orange and blue’ or ‘blue and orange.’ Despite their apparent triviality, I argue that order-based salience patterns tend to make the content positioned first more salient – in the sense of attention-grabbing – in a way that can have surprising normative implications. Giving relative salience to *gender differences* over *similarities*, for instance, can result in the activation of cognitively accessible beliefs about gender differences. Where those beliefs are epistemically and/or ethically flawed, we can critique the salience pattern that led to them, providing an instrumental way of evaluating those patterns. I suggest that order-based salience patterns can also be evaluated on constitutive grounds; talking about gender differences before similarities might constitute a subtle form of bias. Finally, I reflect on how the apparent triviality of order-based salience patterns in language gives them an insidious strength.

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 the reader 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 one’s curiosity. Does the relative salience

Contact: Ella Kate Whiteley <e.k.whiteley@sheffield.ac.uk>

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 paper, I suggest that it likely does. More specifically, I aim to show that there is a distinctively subtle and yet impactful way in which salience patterns get introduced into our language through the order in which we communicate information.

I begin by introducing the concept of *order-based patterns of salience* in §2. This concept refers to the ordering of syntactic content, so that certain parts of it are made to stand out more than others. In §3 I suggest an instrumental way of evaluating order-based patterns of salience. Drawing on psychological research into the phenomenon of order effects, I suggest that certain order-based salience patterns are liable to activate epistemically and/or ethically problematic effects – such as false, inaccurate and/or harmful associations and beliefs – and as such can be criticised on the basis of their upshots. In §4, I offer different constitutive ways of evaluating these salience patterns. Certain order-based salience patterns might themselves constitute a subtle type of epistemic or ethical flaw. Finally, in §5 I consider the wider implications that this discussion has for how we ought to communicate.

2. Order-Based Salience

There are a variety of mechanisms we can employ in language to make certain contents salient. When content is ‘salient’ it is more noticeable – more attention-grabbing – to the audience (see Watzl 2017; and Camp 2017, for a fuller discussion of attention and salience in the mind).¹

Simply mentioning something can function to make it salient. For instance, Maxime Lepoutre (2021) focuses on cases in which ‘counterspeech,’ namely, speech that acts to counter falsehoods like those contained in conspiracy theories, can inadvertently make the falsehoods they critique salient. In particular, to debunk a conspiracy theory, counterspeech usually mentions that conspiracy theory, and in so doing, inadvertently makes that conspiracy theory salient, bringing attention to it. This can be problematic for various reasons. One is that by making the conspiracy theory more salient, it becomes more familiar to us, and, according to psychological research, familiar things are more likely to be interpreted as true. If the theory is false, then this is a problem.²

1. Others also discuss how salience in language shape the *attentional* patterns of the audience, such as Fraser (2021: 4041). Salience, as it is used in sociolinguistics, also treats salient properties in language as those that ‘attract attention’ (Rác 2013).

2. Others also focus on how mentioning something functions to make it salient, such as McGowan (2022). In my experience, it is this mechanism of salience in language that has received most attention from philosophers so far.

There are several other mechanisms that introduce salience patterns into language. One is to employ phrases that function explicitly to draw attention to some content. For instance, I might say ‘Sean Bean was in *Lord of the Rings* but remember: he was also in *GoldenEye*.’ By saying ‘remember,’ I flag to the audience that they should direct their attentional resources to the content coming after that word. Another mechanism involves repeating some content more than others; here, I might discuss Bean’s various film appearances, but keep repeating *GoldenEye*. Another mechanism still is tone of voice; I might communicate the information about Bean’s performance in *GoldenEye* using a louder volume, than when I discuss the other films in which he has appeared.

In contrast to these mechanisms, I want to focus on one distinctively subtle way of making certain contents salient in language; the order in which one communicates one’s utterance can make certain contents stand out as more salient than others. Some philosophers have considered syntactic ordering in relation to salience. Susanna Siegel, for instance, talks about how a news story is “more salient, if it is on the front page of a print or digital newspaper than if it is several clicks or pages away” (2022: 239).³ It is this order-based mechanism that I wish to home in on and analyse in this paper.

Consider a psychologist like Hyde, making decisions about how to present their findings about gender comparisons in relation to various psychological traits. They have findings demonstrating differences between females and males; perhaps there are statistically significant differences regarding throwing velocity, and sexual orientation. They also have findings demonstrating striking similarity; perhaps females and males are overwhelmingly similar when it comes to negotiation style, and extraversion. Once the content of what they want to communicate in their article is settled, however, there remains a further decision: with which facts should they begin their paper? They mock up two articles, the only difference between them being whether differences or similarities are mentioned at the top of the page.

These two articles invoke different order-based salience patterns. An order-based salience pattern refers to the ordering of syntactic content in cases where that order does not, pretheoretically, appear to be of consequence informationally-speaking. Two utterances that differ only in the order-based salience pattern applied to them seem (at least pretheroetically) informationally equivalent, such that they have the same truth value. For instance, ‘the scarf is orange and blue’ or ‘the scarf is blue and orange’ seem to be two ways of communicating identical information, namely, the colour of this scarf. I use the word ‘seem’, as I am not interested in strict informational equivalence, but rather its pretheoretical

3. See too Munton’s (2022) discussion of epistemically evaluating search engines, which includes a discussion of the ordering of webpages.

appearance. Perhaps changing the order in which one mentions the scarf's colours does alter the truth value of the utterance in some subtle sense.⁴ Indeed, I will argue in §3 that such a change in order can alter the information communicated by an utterance when considered through a pragmatic lens, where one considers how a speaker's wider communicative context shapes the utterance's meaning. What is interesting about order-based salience patterns is that they change an utterance in a way that seems inconsequential—something on which I elaborate in §5. The two versions of this psychologist's article described above, I suggest, fit this bill: the alternative versions seem informationally equivalent.

Despite this, one article makes differences more salient, while the other makes similarities more salient. In other words, order-based salience patterns structure contents by giving some *relative* salience over others. Salience, in this case, is comparative. But which contents does each article make more salient? To be able to answer this, we need to consult empirical research into how audiences typically process the orders in which utterances are communicated.⁵ Salience, remember, is about content standing out to an audience.

This research – conducted primarily on English-speaking audiences – indicates that audiences often find more salient content that is presented first, in a phenomenon known as the 'primacy effect' (Haugtvedt & Wegner 1994).⁶ There are some general psychological reasons for this, which I return to throughout this paper. In part, this is down to a mundane fact about mental processing as it occurs in time; content discussed first, in part due to the simple chronological fact that it is mentally processed first, regularly becomes 'anchored' in our minds, in the sense that it becomes more heavily relied upon than content presented afterwards. This is known as the 'anchoring bias' (Tversky & Kahneman 1974: 1128–1130). This preference for what is presented first is reinforced by the nature of our attention spans; these are such that we generally focus better at the beginning of hearing some information, with our attention waning as time goes on (thus being less likely to notice and retain information communicated later).⁷ Relatedly, as speakers, we are likely (often unconsciously) to present information

4. Cf. Fisher 2017 for interesting discussion.

5. Inevitably, there will be individual differences in how we process the utterances we hear. I refer to how audiences 'typically' respond to utterances, to allow for these differences.

6. The 2015 replication crisis in psychology cast doubt on empirical findings in psychology. The order effect, however, continues to be well-replicated, even beyond the studies discussed in §3. See, e.g., (Mihailov et al. 2023: 12), (MacInnis et al. 2021), (Machery 2017: §2.6), and (Sullivan 2018).

7. Sousa (2006: 89–95), for instance, describes research indicating that teachers should ensure the most important learning content is presented first, when it is most likely to be attended to and remembered by students. This suggests that *part* of the reason for order effects may relate to general features of human psychology. We might ask, however, if susceptibility to (at least some) order effects is culturally-specific; the fact that the studies I cite generally focus on English-speaking participants, is relevant to this question. Here, I presume that the effects I discuss are common if not universal. We might also ask if Gricean pragmatic inferences about speaker communicative

that we find most important and relevant to the case at hand first (Kesebir 2017; Formanowicz & Hansen 2022). Evidence also suggests that, as audiences, we tend to infer that content presented first is more important and relevant—something that drives us to focus more acutely on that content (Kesebir 2017).

While general psychological factors, then, contribute to the tendency to find content communicated first most salient, these ‘primacy effects’ are especially likely in certain contexts. In particular, they occur most reliably for topics with which the audience is familiar, or finds interesting or controversial (Haugtvedt & Wegner 1994). These might include topics that are discussed regularly in the news, that come up in everyday conversations, and so on.

There are exceptions, however. In some cases, the content presented last is attended to better by the audience. This latter scenario, known as a ‘recency effect’, generally happens when the content under discussion is unfamiliar, or is of little interest to the audience; here, the content discussed last is typically that which is salient to the audience (Haugtvedt & Wegner 1994). If I am discussing the various attributes of the *Helix Aspersa Maxima* snail, then, presuming that my audience knows little about snails, and/or shows little interest in this species (no shade intended to snail lovers), the information that I discuss last is liable to stand out to that audience.

As will become clearer in what follows, I am interested in order-based salience patterns as they are applied to familiar topics, which audiences often find interesting and/or controversial—in particular, I am focussing on how we communicate information about gender (something around which we structure a lot of everyday life, a subject that is routinely discussed in the news and media, and which often generates interest and controversy). Therefore, I will assume that discussing content relating to this subject first is the typical way in which one makes it salient.

Another qualification is necessary. It should be noted that while we can use this general rule of thumb for order-based salience patterns—that, the content discussed first is made more salient (especially regarding content that is familiar and/or interesting to the audience)—this picture can be complicated by the presence of other mechanisms of salience, including those mentioned earlier. Perhaps certain salience mechanisms, like the addition of the phrase ‘but interestingly’ inserted before the second-discussed content usually overrides the order-based salience mechanism that makes the initial content more salient. To keep things simple, I focus on cases where these extra mechanisms of salience are not present.⁸

intentions account for order effect findings; research from Sullivan (2018), however, suggests that this explanation does not bear out.

8. Future research could clarify how these different mechanisms interact; do certain mechanisms (e.g. repetition) reliably cancel out other mechanisms (e.g. order), for instance?

So, toying with the order in which one communicates information constitutes one mechanism for introducing patterns of salience into one's utterance. The order-based salience mechanism is particularly interesting, I suggest, because it shows that one can shift the salience pattern in one's utterance without needing to add new, or remove existing, syntactic content. Consider again Lepoutre's discussion of salience, where simply mentioning a conspiracy theory makes it salient (2021). Lepoutre implies that, to avoid making a conspiracy theory salient when debunking it, one would need to avoid mention of (i.e. remove syntactic content regarding) that theory. Of course, this would be difficult to do; to debunk a conspiracy theory, one must mention it! Hence the seeming intractability of the problem that Lepoutre identifies, whereby debunking conspiracy theories seems necessarily to come with giving that conspiracy theory undue salience.

By contrast, order-based salience patterns can shift what is salient simply through reorganising the content communicated—without needing to subtract or add new syntactic content. One can communicate exactly the same words across two utterances, and yet alter what one makes salient just by changing how those words are ordered (in a way that, pretheoretically, seems inconsequential). This feature of order-based salience patterns makes this particular salience mechanism especially subtle and difficult to detect. The consequences of this will be addressed in §5.

Hyde, then, in the title of her paper, “Gender Similarities and Differences,” is adopting a particular order-based salience pattern, which makes gender similarities more salient than differences.

3. An Instrumental Critique of Order-Based Salience Patterns

Do the order-based salience patterns we apply to language matter? Here, I expand upon the empirical evidence mentioned in the previous section, to analyse the impact that the primacy effect, generated by certain order-based salience patterns, can have. In particular, I draw on research regarding the ‘order effect.’ This shows how simply changing the order-based salience pattern in some content can lead to significantly different audience responses to that content, including the activation of substantive beliefs and associations in the audience—beliefs and associations that we might find problematic.

3.1 Empirical Evidence and a Causal Model

Some of the more well-known order effects concern pricing. The wily restaurateur, for instance, will put their most expensive wine at the top of the list. While

an otherwise identical list with the cheapest wine at the top will invite customers to stick with the cheapest, putting the most expensive bottle first makes patrons dig a little deeper in their wallets (Roller 2011). Change only the order-based salience pattern in your wine list, then, and you can make more money.

Why does this happen? As discussed in §2, the content discussed first (namely, the wine at the top of the list) is made salient in such a way that it becomes anchored in the study participants' minds (Haugtvedt & Wegener 1994). In the case at hand, this has an important effect: it sets the standard by which an individual compares the content that comes after it. Putting the most expensive wine first gets us to spend more, then, by making all the other wines look much better value.

Some order effects, through this anchoring process, have particularly interesting effects. Consider the order-based salience pattern 'men and women,' which research finds is significantly more common to hear than 'women and men' (Hegarty et al. 2011; Kesebir 2017). While there are exceptions—i.e. 'ladies and gentlemen'—making men relatively more salient than women in language is significantly more common; for instance, one hears 'men and women'; 'males and females'; 'boys and girls'; 'he or she'; 'his and hers'; 'Mr and Mrs'; 'husband and wife'; 'Romeo and Juliet'; 'Adam and Eve'; 'Kings and Queens'; and so on. Research confirms this tendency, finding this holds even after controlling for other variables such as word length (Hegarty et al. 2011; Kesebir 2017). This goes too far for nonconventional phrases. Again, controlling for other phonological factors, research finds that we are more likely to mention a man's name before a woman's name when referring to a heterosexual partnership (Hegarty et al. 2011).

There are two consequences of this order-based salience pattern that I will discuss here; audiences tend to infer that men are more agential and powerful, and they tend to infer gender stereotypes. To be able to understand why (and how), let's first turn to the initial linguistic norm that gets invoked when processing such utterances. The content discussed first, insofar as it is made more salient, is usually inferred by the audience to be more *relevant*—where relevance is the quality "of being connected, central, and important to the matter at hand" (Kesebir 2017: 264–266).⁹ As mentioned in §2, we see this norm reflected in speaker tendencies to put the more relevant content first. Consider how the conventional ordering 'men and women' is usually reversed for contexts in which women are culturally seen as more relevant, such as childcare; there, one is more likely to hear 'women and men' or 'mothers and fathers' (Kesebir 2017: 265, 266, 274).

Once the 'men' in 'men and women' have been inferred to be more relevant, this then activates¹⁰ (unconsciously, as described in §5) further inferences that

9. Cf. Oeberst & Matschke 2017: 673–4.

10. 'Activate' in this context implies that the content inferred is 'called to mind' (i.e. the brain accesses knowledge of the content).

help to make sense of—or at least are associated with—their perceived greater relevance. One way to make sense of the fact that content x has been signalled to be more relevant than content y , is to infer that content x is more agential and powerful. As Selin Kesebir says, agential, powerful parties “are typically more central and important (i.e. more relevant) than their powerless, low-status, and passive counterparts ... Power and agency afford the ability to control others’ outcomes and people attend more closely to powerful others” (2017: 264).

The English language exemplifies this rule of thumb across many of its conventions. Think of how English (and the vast majority of the world’s languages) tends to position the subject (agential) in a sentence before the object (non-agential), as in the sentence “the cat sat on the mat” (Kesebir 2017: 263). Think too of conventional binomial orderings that include agential and non-agential entities, like “living and dead” and “people and things”, “speaker and listener”, “men and machines”, “subject and object”, as well those including powerful and less powerful entities, like ‘rich and poor’, ‘kings and queens’, and ‘adults and children’ (Kesebir 2017: 264; Maass et al. 2022). That audiences are used to inferring the rule that the content discussed first is more agential can be seen in how audiences process non-conventional binomials too.¹¹ Regarding a phrase relating to an imaginary duo, audiences will generally treat Chen as more agential than Amir if they hear the phrase ‘Chen and Amir walked to the shops’ (Kesebir 2017: 264; Maass et al. 2022; Hegarty et al. 2011).¹²

The first thing to note, then, is that these quite general linguistic norms can have sexist effects when they interact with the social convention of linguistically positioning men before women. As the social convention is to mention men before women, there is good reason to think that we will regularly be coming away with the sense that men are more agential and powerful than women.

But there is a further significant effect of this order-based salience pattern. Hearing the phrase ‘men and women’ is liable to activate gender stereotypes and ideologies in a way that hearing ‘women and men’ is not¹³ (Kesebir 2017; Hegarty et al. 2011; Hegarty, Mollin, & Foels 2016). Part of what happens when we attempt to make sense of why some content was signalled to be more relevant, is to draw upon background associations, beliefs, and ideologies that are particularly cognitively accessible—in the sense that it takes little cognitive labour to

11. See Oeberst & Matschke (2017) for related findings.

12. Does this imply that saying ‘ladies and gentlemen’ makes women seem more agential and powerful? Perhaps, given the results discussed here. However, what I say in the main text below, regarding how inferences to background information can interact with these general linguistic norms, should make us cautious. Where background information clashes with the notion that ‘ladies’ are more agential and powerful (as arguably it does in our culture), this can attenuate our disposition to treat the first-positioned category in this way.

13. Unless a domain is specified in which women are seen as more relevant, such as childcare (Kesebir 2017: 265).

think of them, as our minds retrieve them with ease¹⁴—and socially licensed—in the sense that, in a given culture, most people are disposed to regard inferences to that content as legitimate¹⁵—that fit with its communicated greater relevance.

As for the phrase ‘men and women’, there is a clear sense in which men really are more salient and relevant, culturally. They play more central roles in public life than do women, disproportionately inhabit positions of power, dominate main character roles in film and TV, and so on. A significant web of gender stereotypes and ideologies are on hand to explain and justify this male prominence in our culture. Some of these hark back to the issues of agency and power described above: Men are stereotyped to be assertive, decisive, and rational, among many other traits that are directly connected to notions of agency and power, which are easily marshalled to explain and justify men’s heightened prominence in public life. Feminine stereotypes generally function to rationalise women’s more peripheral social position; traits of passivity and nurturance situate them as background characters that support men’s more influential role.

An order-based salience pattern like ‘men and women’, then, can activate inferences to associations, beliefs, ideologies (and so on)¹⁶ that resonate with that pattern—namely that also treat men as more relevant and important. This includes gender stereotypes. Indeed, the centrality of gender stereotypes to our lives makes them particularly cognitively accessible and socially licensed, making opportunities for their activation especially abundant (Kesebir 2017: 262).

This suggests a basic causal model: the framing of an utterance can invoke general linguistic norms (e.g. that the content presented first is more relevant), which in turn activates other general linguistic norms (e.g. that the first-mentioned content is more relevant because it is more agential and powerful), as well as background associations and beliefs that resonate with those norms (e.g. cultural stereotypes that fit with and make sense of the first-mentioned content being more relevant, agential and powerful). Saying ‘women and men’ reduces the extent to which gender stereotypes are activated, then, because women’s linguistic positioning as more relevant, agential and powerful does not generally resonate with background associations and beliefs regarding gender

14. Srull & Wyer (1979) discuss how a key predictor of cognitive accessibility is frequency of use (both by the individual and by others in the individual’s environment).

15. ‘Socially licensed’ here implies that, in our culture, most people are disposed to regard inferences to these associations and beliefs as legitimate (Fraser 2018: 735–736).

16. They might also lead to certain problematic perspectives, in the sense described by Elizabeth Camp (2017). Some of these might be (in this case gender-) essentialising perspectives (Camp & Flores forthcoming). Investigating how order-based salience patterns in language might activate (or partially constitute, as per the discussion in §4) problematic perspectives represents a particularly interesting avenue for future research. The reason for my focus on stereotypes here is simply that it is where the empirical evidence is strongest for my causal claim regarding order-based salience patterns.

in our culture. Specifically, gender stereotypes clash with the idea of women being more relevant, agential and powerful.¹⁷

This causal model should not be taken to imply an entirely linear process of inference. Evidence suggests that the activation of background associations and beliefs can interact with the activation of the general linguistic norms described. The extent to which the linguistic norms described above—to find the content positioned first more relevant, agential, and powerful—influence an audience’s thought can be increased or lessened depending on how well those background associations and beliefs resonate (or clash) with those linguistic norms.

Taking the linguistic norm of relevance as an example, when participants were presented with a vignette of a young person commenting on their tennis practice, saying “My [mother and father OR father and mother] have met the coach last week ... [and that] My [mother and father OR father and mother] told me afterwards that they want me to work on my weaknesses...”, when asked which parent they perceived as being more relevant, participants gave strikingly different answers depending on the order in which the parents were mentioned (Kesebir 2017: Study 6). When the mother was mentioned first, 53.5% of participants said the mother was more relevant and involved. But when the father was mentioned first, 81.6% said the father was more relevant and involved. While mothers *were* seen as more relevant when positioned first, then, they were nearly 30% less likely to be treated as most relevant than were fathers when fathers were positioned first. The idea that background cultural knowledge can strengthen or attenuate the tendency to perceive the category mentioned first as more relevant provides an explanation for this, insofar as fathers are culturally seen as more relevant than mothers in the context of sport (both in the more general sense that sport is stereotyped as a masculine activity, and in the specific sense that fathers are seen as being more interested in their children’s sporting activities).

This general causal model – whereby the framing of an utterance can activate linguistic norms, which can themselves activate background associations and beliefs that resonate with those linguistic norms – is corroborated by findings concerning other framing mechanisms in language. Consider two questions: (a) ‘Do women lead differently than men in boardrooms?’; (b) ‘Do men lead differently than women in boardrooms?’. Interviewing 226 German university students, Susanne Bruckmüller and colleagues (2012) found that there was a tendency to give different responses to these two questions. Those answering the first question – namely, how *women* differ from *men* – were more likely to agree to statements indicating that men have higher social status and power in society, more likely to agree that existing inequalities between women and men are

17. As described in fn. 13, this might change if a domain is specified in which women culturally are seen as more relevant, agential and/or powerful, such as childcare.

justified, and, as I will focus on here, more likely to endorse gender stereotypes. In this latter instance, participants 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).

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 “the effect to be explained”, while (y) is “the implicit norm for the comparison” (Bruckmüller et al. 2012: 210). This tends to trigger background associations and beliefs that resonate with this framing. Men represent the norm for leadership in many cultures in a variety of senses: *statistically* (leadership positions are dominated by men); *socially* (there is a cultural stereotype associating leadership with men and masculine traits like rationality, assertiveness, lack of emotion) (Bruckmüller et al. 2012: 212–213); and *normatively* (men, in Bruckmüller et al.’s (2012: 211) words, “set the standard of culturally valued behaviour [in that domain]”, something seen in how *good* leaders tend to get described using stereotypically masculine attributes, e.g. ‘assertive’ and ‘rational’).

Women are the cultural ‘effect-to-be-explained’ in all these senses. For instance, the wide range of negative evaluations women leaders receive in their workplaces and in the media (Brescoll et al. 2010)—evaluations that routinely highlight stereotypically feminine traits like ‘emotion’ being part of the perceived problem—can be explained by women being seen as the normative abnormality in the context of leadership.

Asking ‘how do women differ from men in the boardroom’ activates gender stereotypes more reliably than the question ‘how do men differ from women in the boardroom’, then, because the former framing, insofar as it treats women as the ‘effect-to-be-explained’ and men as the ‘implicit norm for comparison,’ activates cultural associations and beliefs that help to explain and justify, or simply resonate with, this framing—namely, associations and beliefs that also position women as the effect to-be-explained and men as the norm in our culture. An integral part of this cultural picture is gender stereotypes, as described above. While the specific mechanism discussed by Bruckmüller et al. is different than the focus of this paper—it concerns linguistic norms for comparison instead of those relating to order-salience and relevance—it corroborates the general causal model undergirding this section, wherein subtle framing devices in language activate background cultural associations, beliefs and ideologies (including gender stereotypes) which resonate with the framing.¹⁸

18. More generally, the suggestions in this section fit into a rich vein of research identifying various linguistic conventions that subtly reinforce stereotypes and other potentially harmful and/or inaccurate associations, beliefs, ideologies (and so on) about social groups, such as the use

3.2 *Instrumentally Evaluating Order-Based Salience Patterns*

How does this get us to a way of evaluating order-based salience patterns? Well, if we have reason to find these inferences false, inaccurate, and/or harmful, then we can criticise the order-based salience patterns which lead to them. Here, we employ an instrumental critique of order-based salience patterns, evaluating them on the basis of their upshots.¹⁹ I would suggest that many would find the various inferences described above problematic in these multiple ways, meaning that we can say that both ‘men and women’ and ‘do women lead differently than men in boardrooms?’ deploy instrumentally-problematic order-based salience patterns.

This is not to suggest that this order-based salience pattern is inevitably instrumentally problematic, or problematic all-things-considered. As is the case for all the order-based salience patterns identified as potentially problematic in this paper, there may be mitigating (moral, epistemic, practical, prudential) factors in certain contexts that either cancel out or override the problems highlighted. Perhaps when discussing suicide rates, for instance, it is important (morally, practically, etc.) to mention men first (and therefore make them more salient and relevant), given the troublingly high rates of suicide in that population.²⁰ It might be that, in this particular context, the problems with the inferences discussed in §3.1 are outweighed by normative and/or practical considerations in favour of using that order-based salience pattern, or even that the problematic inferences are not activated. The point I want to emphasise is that empirical

of nouns (e.g. Gelman & Heyman 1999), generics (e.g. Leslie 2014), and androcentric labels (e.g. Spender 1980).

19. There may be more long-term upshots to consider, too. Research suggests that content that is discussed first sets the foundation for how further information is accommodated and integrated, is processed more deeply, and becomes more cognitively accessible (Kesebir 2017). This might all result in longer term downstream epistemic costs, such as a less accurate understanding of the subject in question (Puddifoot 2017).

20. Someone might argue that the issue with phrases like ‘men and women’ is not order. Instead, such phrases involve parsing the world into the two purportedly relevant categories, instead of saying e.g. “people” (thank you to an anonymous reviewer for raising this point). There are three points I make in response. Firstly, the example given here, of mentioning gender differences regarding susceptibility to suicide, shows that sometimes one might need to ‘parse’ the world in this way (indeed, there are many times it is arguably useful to talk about ‘women’ and ‘men’, such as when highlighting gender inequalities). Secondly, by focussing on problematic order-based salience patterns in certain utterances, I am not claiming to have exhausted the issues with those utterances. Beyond the reviewer’s suggestion, the invisibility of non-binary identities in the phrase ‘men and women’ is another important issue. Highlighting the specific issues that *order* creates, however, is a worthwhile project, particularly because such issues are rarely philosophically analysed. Thirdly, the points I make using the ‘men and women’ example are intended to generalise outside of utterances that ‘parse’ the world into these categories in particular; they are intended to encourage us to think about how we order *all* of our communication. That might include the content in Hyde’s psychology article, a discussion of a topic that orders negative information before positive information (see the discussion in §4 below), and so on.

findings show that there is a strong tendency for certain order-based salience patterns to activate problematic inferences, meaning that those patterns come with a red flag. This ought to be factored into one's overall decision regarding whether to employ that pattern.

Consider again Hyde's paper. She has opted to avoid the order-based salience pattern 'gender differences and similarities.' Are there possible problematic inferences that this order-based salience pattern might be triggering? Consider everyday phrases like *men are from Mars, women are from Venus*, and *the opposite sex*—phrases that are materially embodied around us, such as in the pink vs. blue aisles of toyshops (Gelman & Taylor 2000). Many have suggested that this belief of extreme difference forms part of a robust and prevalent cognitive bias—gender essentialism—which puts these extreme differences between the genders down to different (usually biological) fixed essences (Meyer & Gelman 2016; Bastian & Haslam 2006: 229). In other words, there does indeed seem to be a cognitively accessible and socially licensed belief (perspective? ideology?) that paints women and men as dramatically different from one another. If this belief is indeed cognitively accessible and socially licensed, this gives us reason to think that an order-based salience pattern which makes salient what is central to this belief—namely, difference—will likely activate and reinforce it. This is because talking about gender 'differences' first signals that differences are more salient and relevant than similarities, which in turn activates the pervasive cultural beliefs described above. These beliefs resonate with, and potentially explain and justify, the heightened salience (and indicated relevance) of gender differences.

Research already indicates how little it takes for gender essentialist beliefs to become activated in audiences. For instance, Rebecca Bigler's (1995) study on American school children examined the impact of teachers giving occasional instructions that mention gender to their classes, such as 'all the girls put their bubble makers in the air.' Despite no gender essentialist stereotypes being mentioned, the mere use of gender in this functional manner led to the children in those classrooms being more likely to endorse essentialist gender stereotypes, than those in classrooms without such gendered instructions.

Hyde takes issue, on both epistemic and ethical grounds, with the essentialist belief that the genders are overwhelmingly different (2005). Examining 46 meta-analyses of psychological studies into gender comparisons, Hyde revealed that 78% of the effect sizes were small to close-to-zero, meaning that women and men range from being evaluated as overwhelmingly similar, to virtually identical, on the majority of psychological variables. She concludes, therefore, that the belief in question is false. Further, she suggests that the belief that women and men are significantly different comes with significant harms, not least because it reifies gender stereotypes, with all their associated problems (for instance, their

prescriptive nature means that those who transgress the stereotype for their gender are socially punished).

If we take the belief that women and men are overwhelmingly different to be problematic in these ways, then, as per the instrumental critique, we can criticise the order-based salience pattern that leads to the epistemic and moral problems encompassed in that belief.

What is the significance of this proposal? Well, many attempts at fixing implicitly biased language work to add ‘missing’ content. Consider job advertisements that use masculine generics, describing the ideal candidate by using phrases such as ‘He must have good organisational skills.’ Many have critiqued these advertisements as androcentric, ultimately disinclining women to apply (Kesebir 2017). A common solution has been to replace this masculine generic with the conjunction ‘he or she’, to ensure, in Kesebir’s words, “the symbolic inclusion of females in the references” (2017: 263). Consider too feminist critiques of the overwhelming focus on gender differences in science, at the expense of gender similarities (Rippon et al. 2014). Hyde herself makes this critique. One might think that a solution here is simply to add a section on gender similarities to papers that discuss only gender differences.

The novelty of the proposal here, echoing that made by Kesebir, is that this might not be enough. Even language explicitly including men *and* women, and explicitly mentioning gender differences *and* similarities (and so on), might nevertheless activate problematic biases by giving relative salience to the ‘wrong’ content.

4. A Constitutive Critique of Order-Based Salience Patterns

So far, I have suggested an instrumental critique of order-based salience patterns; this entails that the problems with a pattern are derivative of the problems with its upshots. Can order-based salience patterns instead be problematic in and of themselves? Here, I suggest that order-based patterns of salience can themselves constitute a subtle form of epistemic or moral flaw. To see this, let us take a short detour to consider a related proposal, made by Elizabeth Anderson.

In her 1995 paper “Knowledge, Human Interests, and Objectivity in Feminist Epistemology,” Anderson claims 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 book that are true. For example, “that Jews had considerable invest-

ments 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” (Anderson 1995: 38).

The problem, Anderson claims, is not that Farrakhan’s book contains falsehoods; arguably, it does not. It is that it does not put the facts it discusses into a wider context that accurately represents the *significance* of those facts. Taken by themselves, these facts give the impression that Jewish people played a particularly significant role in the Atlantic slave system, more significant than other ethnic groups. The larger context, however, shows this impression (though not the facts themselves, purportedly) 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” (Anderson 1995: 38).

Anderson concludes that ‘The Secret Relationship’ constitutes a biased representation 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 simply truth when judging the adequacy of our language (cf. Elgin 1996; Richard 2008; Gibbard 1990).²¹ In so doing, Anderson suggests a constitutive epistemic critique of ‘The Secret Relationship’ that does not rely on merely assessing its truth-conditions. One way in which we can directly assess the epistemic standing of a linguistic account of a subject, in other words, is by assessing whether it includes all pertinent and significant facts. A subtle way in which one’s account of a phenomenon can be biased, then, is by omitting relevant, significant facts about that phenomenon.²²

What about order-based salience patterns? Making the ‘wrong’ content relatively more salient than the ‘right’ content might also be constitutive of a subtle bias. Imagine an amended version of Farrakhan’s book, which included all significant truths, but ordered them so that all the negative truths about the Jewish role in the slave trade were discussed first. The complaint here would be that, *organising* these (significant) facts in a way that makes salient the negative

21. Alternative criteria to truth/falsehood include, for example, accuracy, warrant, and aptness. Importantly, these criteria can capture ways in which non-propositional entities (like desires, emotions, hopes, and, I suggest, order-based salience patterns) can be epistemically assessed on constitutive grounds.

22. Others have made similar claims. Outside of Siegel (2022), mentioned earlier, Hermann and Chomsky’s (2008) critique of media propaganda for related ideas, insofar as the selection of certain truths over others can mislead an audience. Thank you to an anonymous reviewer for pointing out this latter connection.

features of the Jewish role in the slave trade counts as a subtle way in which one's account is biased.²³ One subtle way in which a representation can be biased, in other words, consists in treating the 'wrong' facts as more salient than the 'right' facts. The order-based salience pattern employed by this hypothetical version of Farrakhan's book might also count as a subtle way of being anti-Semitic. Where these states, of bias and anti-Semitism, are constitutively epistemically and/or ethically flawed, then the order-based salience pattern that counts as a version of these phenomena is also constitutively flawed.

How more specifically are we to understand this constitutive claim?²⁴ One way is non-instrumentally. Imagine the sorts of order-based salience patterns described in this piece never finding an audience, and thus not having problematic effects (presume too they don't encourage the speaker to develop or embolden their prejudices,²⁵ or some other negative effect relating to the speaker's later conduct). The patterns might structure syntactic contents in a private journal entry, an academic paper that never sees the light of day, a person's

23. An analogy can be made here to an attempt at a real-life version of this hypothetical, instead regarding a textbook. Jamieson & Radick (2013) intended to create a modified version of a classic undergraduate textbook on genetics, the only change being the order in which its contents were discussed. In particular, they wanted the 'interactionist' research, which emphasises how traits arise from the complex interaction between genes and environment, discussed first instead of last. This interactionist research shows, for instance, how some genes can function in surprisingly different ways depending on which environment they interact with. Mendelian genetics, which is prone to 'gene for' talk (i.e. the 'gene for' blue eyes), was to be moved from the beginning chapters of standard textbooks to the end of Jamieson and Radick's revised textbook. For various practical reasons, however, Jamieson and Radick (2017) could not ensure that these were the only changes made in their revised textbook, meaning that, in practice, more substantive content changes were necessary in their study. It is interesting, however, that students taking their revised course (and its revised textbook) emerged with less genetically determinist attitudes regarding genetics than those taking the standard course. If similar findings could be replicated in future studies that ensure no substantive content changes to the original genetics textbook, these findings would speak primarily to the instrumental evaluation of order-based salience patterns; if we find (epistemic and/or ethical) problems with genetic determinism, we can critique the salience pattern (in this case, the chapter order in the standard textbook) that cultivates this attitude. Perhaps, though, a constitutive critique is also possible; we might say the salience pattern of the standard textbook *constitutes* a subtle form of genetic determinism.

24. Thank you to an anonymous reviewer for helping to clarify my thoughts about these constitutive critiques in this section.

25. See DiFranco & Morgan (2023: 5–6) for further discussion of this qualification. One point from this discussion is worth mentioning here. Critics might insist that for speech to be morally problematic, it must have bad effects. Consider slurring utterances, as is the focus for DiFranco and Morgan. They point out one particularly counter-intuitive consequence for such an effects-focussed theorist; such "theorists must admit that, *ceteris paribus*, solitary slurring by nonracists is morally reprehensible (because it can have measurable negative effects [namely, emboldening the prejudices of those nonracists], while solitary slurring by extreme racists is not wrong (because it cannot make them any worse)" (2023: 6). Given that the effects-focussed theorist might believe their view to be the more intuitive option, it is worth pointing out this plausibly counter-intuitive outcome of their view.

utterances when speaking to themselves, and so on.²⁶ To take one example, imagine someone who keeps a private diary in which they discuss their thoughts about a range of people. Whenever they mention people who are women, they describe their appearance before they describe their personalities, behaviours, careers, and so on. The opposite is true for when this individual writes about men. I think that these patterns are problematic in themselves, despite their lack of problematic upshots.

One way of justifying this is to suggest that the order-based salience pattern constitutes a failure of a moral duty, such as treating people fairly. This strategy has been employed to defend the constitutive evaluation of things like certain beliefs and utterances that lack negative effects. For instance, Rima Basu (2023: 5) suggests that certain of our beliefs, whether or not they generate negative effects, can wrong others by constituting a failure in our duty to adopt what P. F. Strawson calls the ‘participant stance’ with other people—referring to the need to treat others, in deed and thought, as persons with agency, instead of as mere things. Discussing the broader category of derogatory attitudes, Ralph DiFranco and Andrew Morgan (2023: 2) suggest that these “involve failing in one’s moral duty to be affectively open, that is, in the duty to avoid reacting to others with hostility.” I find these arguments plausible, despite their likely demanding implications. Instead of defending them here, though, I will focus on two other ways of developing a constitutive critique of order-based salience patterns that I believe will be less controversial.

Consider one response to the journal entry above. A critic might suggest that any epistemic or moral problem here would be derivative on a prior mental attitude influencing that order-based salience pattern, meaning that it’s the prior mental attitude where the problems reside; perhaps the salience pattern of talking about women’s appearances before more substantive traits is bad only in virtue of the objectifying outlook of the individual uttering it.²⁷ Critically, I want to suggest that even if we concede the relevance of a prior mental state—insofar as the prior mental attitude is necessary for making a principled judgment between order-based salience patterns that are acceptable versus those that are not—that needn’t mean that the order-based salience pattern itself loses its normative quality.

Consider DiFranco and Morgan’s (2023) argument regarding slurs that are uttered privately and thereby generate no negative consequences. The slurring utterance will, they concede, likely be influenced by a derogatory attitude had by the slurrer. So, they ask, why not find the moral fault in that attitude instead of the slur? Their response is this: “if a particular derogatory attitude is morally

26. Thank you to an anonymous reviewer for raising these examples.

27. Siegel (2017: Ch. 9) tends to frame things in this way. Where salience patterns in attention are problematic, this is usually because of the prior mental attitude influencing them, making them *derivatively*, not constitutively, problematic.

defective, then it is wrong to express that attitude as well". Here the utterance of a slur is considered an expressive manifestation of that attitude, *endorsing* and *manifesting* that attitude in a way that inherits its badness. Consider a person who (privately) says 'I hate gay people,' in a way that expresses (endorses) their homophobic attitude. Compare this with an individual who, at the end of a therapy session aimed at changing their objectionable attitudes, says the same thing. DiFranco and Morgan (2023: 7) suggest that this latter instance is simply "a neutral and well-informed evidential self-report on her mental life." I take it that we have different moral evaluations of these utterances. We can make sense of those different evaluations by acknowledging that some utterances simply report an attitude of hatred, while others action-express them. Utterances that action-express hatred, for instance, endorse and manifest that hatred in a way that inherits the moral badness of that attitude.²⁸

In a recent *Ethics* blogpost, Zachary Irving (2023) makes a related claim, this time about attentional patterns.²⁹ Irving gives an example of someone who has prejudicial biases, like sexist, objectifying beliefs. In one scenario, this influences their attentional patterns so that they find women's bodies more salient than their conversational contributions. In another scenario in which they still have those biases, the individual successfully exercises attentional control to suppress the usual attentional patterns above. The former is worse than the latter, Irving suggests—an intuition we can make sense of by adopting what he calls the 'process view' where the guiding state *and* the resulting salience pattern in attention are morally bad. In other words, instead of insisting that the mental state guiding the attentional patterns is the locus of evaluation, one might say it's the whole process of *the mental state guiding a salience-based structure in attention*. We might extend this idea to order-based salience patterns in language; the locus of evaluation can be the whole process, whereby the *mental state guides the order-based salience pattern applied to the individual's utterance*. This gives us a constitutive evaluation of order-based salience patterns in language, albeit restricted to cases where those patterns are influenced (in the right way) by a prior defective attitude.

28. Drawing on other accounts, DiFranco and Morgan suggest that when we find the attitude that an utterance expresses morally problematic, we have a good reason not to speak it. See Zheng and Stear's (2023: 399) related argument regarding private imaginings (e.g. imagining a 'ghetto thug'); certain private imaginings can be intrinsically unethical in virtue of realising—in the sense of expressing or instantiating—certain meanings that an oppressive social context makes available to it. Here, the expression is not of an attitude had by the agent, but something external to the agent—meanings in the agent's wider social context.

29. In his 2023 co-authored paper "Catch 22 of Forgetfulness," Irving and his colleagues allude to this argument. They consider cases of distraction in which people implicitly judge a distracted individual for the whole process (prior mental attitude influencing their distracted state), since it matters to the normative evaluation of the distraction whether its harms are under control or not, whether they issue from the individual's character or not, and so on.

There are other ways of developing a constitutive evaluation of these order-based salience patterns. A rich vein in the philosophy of language has used Austinian speech act theory to argue that certain utterances, beyond causing problematic effects, can themselves constitute actions (e.g. Langton 1993; Maitra & McGowan 2012). Saying “I pronounce you married” in a marriage ceremony does not just cause two people to become married; it *constitutes* the act of marrying. Saying ‘Run – there’s a fire!’, said by a person in a burning building, *constitutes* the act of warning. Sometimes, the act constituted by the utterance is harmful, or otherwise problematic. Rae Langton (1993: 302–303) gives the example of Apartheid law. She says that the utterance ‘Black people are not permitted to vote’ uttered in certain contexts, *constitutes* subordination; it constitutes the acts of (unfairly) *ranking* black people as inferior, (unjustly) *depriving* them of rights and powers, and *legitimizing* discrimination against them. Where those acts are harmful, that makes the utterance, insofar as it *is* those acts, constitutively harmful. This constitutive route, unlike those described above, is not non-instrumental; it does not attempt entirely to screen off the *effects* of the utterances. Indeed, for the speech act to be successful, certain upshots are required, including hearer uptake. What it does do, however, is draw a *norm-driven*, as opposed to causal, connection between the utterance and its effects.³⁰

If we extended this idea to order-based salience patterns, this would mean restricting the subset of order-based salience patterns considered constitutively problematic to those with problematic effects—specifically, those that do actually contribute to relevant problematic social activities (thus excluding the private utterances described above). Still, it provides a way of finding those order-based salience patterns constitutively problematic.

We have at least three separate ways of arriving at a constitutive critique for assessing order-based salience patterns, then. These involve locating the problem: solely in the salience pattern itself; in the combination of the salience pattern and a guiding prior attitude; and/or in the combination of the salience pattern and its effects. Different readers might be persuaded by different routes to a constitutive evaluation.

With these routes sketched, let’s return to the particularly common linguistic structuring of mentioning men before women. I considered an instrumental critique of this in §3, referring to the gender stereotypes this is liable to activate. The suggestion here is that, in addition to these potential problematic upshots, this common order-based salience pattern might itself count as a subtle form of androcentrism in one or all the ways described in this section. Perhaps, for instance, it constitutes the (speech) act of unfairly placing men at the centre of things.

30. Other accounts that develop constitutive critiques of language in this non-instrumental sense include those focusing on normative notions like ‘licensing’, e.g. Tirrell (2012).

Return too to Hyde. Perhaps the title that Hyde avoided, which made gender differences more salient than similarities, would itself constitute a subtle form of gender essentialism. It might, for example, be a way of expressing a problematic essentialist outlook.

As has been suggested in work on speech act theory, offering a way of arguing that certain speech implicates epistemic or ethical problems in a way that does not rely on establishing causal relationships can be very helpful. For instance, writing about the debate concerning pornographic speech and its purported effects on women's subordination, McGowan recommends further exploring pornography's constitutive harms (in the third sense considered above), saying "As is well known, it is notoriously difficult to establish the truth of...complex [causal] claims" (2005: 28). Indeed, it takes a significant amount of empirical work to establish the sorts of findings discussed in §3.1. Further, while I believe this would be to underestimate the robustness of the tendencies identified in that section, the importance of context means that it is always possible for the critic to say, 'maybe features of *this* context mean that this order-based salience pattern *doesn't* activate the sorts of problematic inferences you are concerned about.' Indeed, it is implausible to think we can empirically test the effects of an order-based salience pattern across every conceivable communicative context (across different types of speaker, either pro- and preceded by different utterances, delivered with different tones of voice, etc.). If we can instead suggest that a given order-based salience pattern *constitutes* an epistemic or moral bias, then one can condemn the utterance in question without needing to prove those causal claims.

5. The Fallout: Paying Attention to Order-Based Salience Patterns

What is the fallout of this discussion? Consider first how these instrumental and constitutive critiques of order-based salience patterns might help to shed light on various existing complaints. To expand on an example mentioned earlier, it is common to hear frustrations raised about how routinely news articles, novels, obituaries, and so on, describe women by their appearance before their personalities and careers—an order-based salience pattern that is much less often applied to men (Buchanan & Crucchiola 2018). Relatedly, many from marginalised ethnic backgrounds complain about utterances that give undue salience to their demographic properties; for instance, Zoe Kravitz, referring to a famous quote by Jean-Michel Basquiat, in which he claims that he is not a *black* artist, but rather an artist, says this: "Happy to be black. Just don't need to say it in front of

everything” (Kravitz, in Willis 2017).³¹ The tools in this paper provide ways of criticising these order-based salience patterns on instrumental and constitutive grounds. We might find that the former activates objectificatory beliefs about women, for instance, while the latter constitutes a way of expressing disrespectful, othering attitudes.

What, practically speaking, are we to do in light of these findings? Note that salience is pervasive. All language must be structured along dimensions of salience. We cannot communicate without making some contents more salient than others, not least because we must inevitably say one sentence before another. If, as I have suggested, these order-based salience patterns have implications regarding whether we activate a false, inaccurate, and/or harmful belief, or whether they constitute some epistemically or ethically problematic state, then we should reflect on which patterns of salience we use, and which are being communicated to us. Do we tend to say ‘men and women’ when discussing gender in everyday conversations, in teaching scenarios, in the papers that we write? Do we tend to hear gender differences mentioned before similarities in the studies that we read?

Suggesting that we reflect on the order-based salience patterns that we use and hear is a modest proposal, but it has significance. We tend not to reflect on the salience patterns that we and others employ in language, in part because we are prone to treat them as insignificant. There is a sense in which philosophers, for instance, are aware of the importance of ordering their arguments so that they build nicely, and ultimately are more persuasive to the reader. But the more radical effects of order-based salience patterns discussed here—where a change in order can mean the difference between communicating something false or biased, and not—are rarely acknowledged.

Indeed, order effects—and the wider research into ‘framing effects’³²—receive so much press precisely because our responses surprise us. In part, this is because not only are these patterns of salience themselves rarely consciously noticed by the audience, but any inferences that those patterns solicit us to make are also rarely consciously registered (Sher & McKenzie 2006; Holbrook et al. 2000). This has real significance. Research demonstrates that where inferences are activated under our conscious awareness, such as those to gender stereotypes as discussed in §3, the beliefs and associations that we have inferred are

31. See Whiteley (2023) for a discussion of this case. There, I focussed on the moral problems with the attentional patterns these utterances likely invite in their audience. Here, I focus on the problems with the utterances themselves.

32. Framing effects include order effects, but they also concern other types of effect, like those relating to attribute choice (Kahneman & Tversky 1979).

especially likely to go on to influence our thought and behaviour (Banaji et al. 1993). This is because, unaware of how a salience pattern is soliciting us to think and act, we do not attempt to block any problematic inferences that we might be making. The study participants in Kesebir's and Bruckmüller and colleagues' studies did not simply call to mind gender stereotypes. They then thought in ways congruent with those stereotypes, going on to attribute more masculine stereotypes to men, and more feminine stereotypes to women (as opposed to rejecting the stereotypes that had been called to mind). The *under the radar* nature that order-based salience patterns (and comparative question framings of the sort Bruckmüller et al. examined) have, then, means that any problems that they cause and constitute can more easily continue unchecked.³³

This arguably gives order-based salience patterns an insidious power. The very thing that makes the impact of order-based salience patterns surprising is also what makes them powerful. This concern about the insidious strength of order-based salience patterns might motivate us to go further than the modest suggestion above. Beyond simply reflecting on which order-based salience patterns we employ and encounter, we might have reason to employ salience patterns that actively *counter* problematic ones. Perhaps we have a directed epistemic duty to our audience, such that we *owe* truth- or accuracy-conducive utterances to our audiences. This might mean that we owe to our audience order-based salience patterns that avoid the likely activation of false beliefs, and so on, as well as patterns that do not constitute a type of epistemic flaw. Perhaps more intuitively we have a duty to avoid wronging, disrespecting, or harming others. This might give us a reason to avoid order-based salience patterns that are liable to trigger wrongful, disrespectful or harmful beliefs (and so on) as well as patterns that constitute a moral wrong. Hyde has, of course, opted to *reverse* the standard order-based salience pattern, which makes gender differences more salient than similarities. We might see her as fulfilling epistemic and/or ethical duties in doing so, by disrupting the likely activation of cognitively accessibly false and harmful beliefs that the order 'gender differences and similarities' is liable to generate.

So, ought we adopt order-based salience patterns that *counter* problematic conventional ones? There are many things to discuss regarding this proposal. One thing to reiterate is that context matters. In some contexts, talking about gender differences before similarities (perhaps even talking *only* about gender differences), for instance, might not cause or constitute problems. Consider, for instance, the need to correct for androcentric biases in medicine, by emphasising how a medicine tested on male subjects is in fact much less effective on average

33. 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 and Langton & West (1999) who discuss *back-door* speech acts.

in female bodies (Criado-Perez 2019); emphasising sex differences here seems epistemically and ethically important. In other contexts, even if we do think that this salience pattern is liable to activate or constitute gender essentialism, any concerns we might have in this regard may be outweighed by other (moral, political, epistemic, practical) considerations. A second point of contention for this latter proposal relates to epistemic paternalism. The proposal at hand recommends using patterns of salience to shape our audience's inferential patterns for (what we deem is) the better, without consulting them on the issue. One would need to show that this is acceptable (see Jackson 2022: 134, 144, for a persuasive defence of epistemic paternalism in relation to order effects). These issues deserve further attention, which I do not have the space for here. They do, however, demonstrate the relevance of this discussion to the lively debates about nudge theory (Niker 2018; Noggle 2017)—a theory using insights from behavioural economics to suggest that simple, low-cost interventions can influence people to think and act better (Thaler & Sunstein 2008).

6. Conclusion

Suppose that an article communicates only truths about a subject, such as gender psychology. Further, suppose (rather idealistically) that it communicates all relevant truths about that subject. Might there nevertheless be grounds for criticising that article for giving relative salience to certain truths over others? In other words, while we cannot criticise the article for being untrue, or perhaps for being misleading in virtue of omitting certain truths, can we nevertheless criticise it for its order-based salience patterns? This paper has suggested that we can. We might evaluate an order-based salience pattern on instrumental grounds, for its liability to activate an epistemically and/or ethically problematic belief, ideology, and so on. Or we might evaluate it on constitutive grounds, finding the salience pattern itself constitutive of some epistemic and/or ethical flaw. In light of this, I suggest that we pay more attention to the order-based salience patterns that shape all communication, given their potential epistemic and ethical significance.

Acknowledgements

This paper has benefitted significantly from insightful reviewer comments. For helpful discussion, I would also like to thank my former colleagues at the LSE, particularly Lewis Ross, Liam Kofi Bright, and Johanna Thoma, as well as my partner, Justin Tadros.

References

- Anderson, Elizabeth (1995). Knowledge, Human Interests, and Objectivity in Feminist Epistemology. *Philosophical Topics*, 23(2), 27–58.
- Banaji, Mahzarin R., Curtis Hardin, and Alexander J. Rothman (1993). Implicit Stereotyping in Person Judgment. *Journal of Personality and Social Psychology*, 65(2), 272–281.
- Bastian, Brock, and Nick Haslam 2006. Psychological essentialism and stereotype endorsement. *Journal of Experimental Social Psychology*, 42, 228–235.
- Basu, Rima (2023). The Morality of Belief I: How Beliefs Wrong. *Philosophy Compass*, 18(7), 1–10.
- Bigler, Rebecca 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), 1072–1087.
- Brescoll, Victoria. L., Erica Dawson, and Eric L. Uhlmann (2010). Hard won and easily lost: The fragile status of leaders in gender-stereotype-incongruent occupations. *Psychological Science*, 21(11), 1640–1642.
- Bruckmüller, Susanne, Peter Hegarty, and Andrea E. Abele (2012). Framing gender differences: Linguistic normativity affects perceptions of power and gender stereotypes. *European Journal of Social Psychology*, 42(2), 210–218.
- Buchanan, Kyle and Jordan Crucchiola (2018, April 4). How 50 Famous Female Characters Were Described in their Screenplays. *Vulture*. Retrieved from <https://www.vulture.com/2018/04/how-50-female-characters-were-described-in-their-screenplays.html>
- Camp, Elisabeth (2017). Perspectives in Imaginative Engagement with Fiction. *Philosophical Perspectives*, 31(1), 73–102.
- DiFranco, Ralph and Andrew Morgan (2023). No Harm, Still Foul: On the Effect-Independent Wrongness of Slurring. *Journal of the American Philosophical Association*, 9(3), 1–19.
- Hermann, Edward and Noam Chomsky (2008). *Manufacturing Consent: The Political Economy of the Mass Media* (2nd ed.). Random House.
- Criado-Perez, Caroline (2019). *Invisible Women: Data Bias in a World Designed for Men*. Abrams Press.
- Elgin, Catherine (1996). *Considered Judgement*. Princeton University Press.
- Fisher, Sarah (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.
- Flores, Carolina and Elisabeth Camp (forthcoming). “That’s All You Really Are”: Centering Social Identities Without Essentialist Beliefs. In Sally Haslanger (Ed.), *Mind, Language, and Social Hierarchy*. Oxford University Press. Draft retrieved from https://static1.squarespace.com/static/5c17834dc258b4dac072be23/t/601605be8d4d172b15e30596/1612055998642/Camp_Flores.ThatsAllYouReallyAre.pdf
- Formanowicz, Magdalena and Karolina Hansen (2022). Subtle Linguistic Cues Affecting Gender In(Equality). *Journal of Language and Social Psychology*, 41(2), 127–147.
- Fraser, Rachel (2018). The Ethics of Metaphor. *Ethics*, 128(4), 728–55.
- Fraser, Rachel (2021). Narrative testimony. *Philosophical Studies*, 178, 4025–4052.
- Gelman, Susan. A. and Gail D. Heyman (1999). Carrot-Eaters and Creature-Believers: The Effects of Lexicalization on Children’s Inferences About Social Categories. *Psychological Science*, 10(6), 489–493.

- Gelman, Susan, and Marianne Taylor (2000). Gender Essentialism in Cognitive Development. In Patricia H. Miller and Ellin Scholnick (Eds.), *Towards a Feminist Developmental Psychology* (169–90). Routledge.
- Gibbard, Alan (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.
- Haugtvedt, Curtis. O. and Duane T. Wegener (1994). Message Order Effects in Persuasion: An Attitude Strength Perspective. *Journal of Consumer Research*, 21(1), 205–18.
- Hegarty, Peter, Sandra Mollin and Rob Foels (2016). Binomial word order and social status. In Howard Giles and Anne Maass (Eds.), *Advances in Intergroup Communication: Language as Social Action*, VII(21), (119–135). Peter Lang Publishing.
- Hegarty, Peter, Nila Watson, Laura Fletcher, and Grant McQueen (2011). When gentlemen are first and ladies last? Effects of gender stereotypes on the order of romantic partners' names. *British Journal of Social Psychology*, 50(Pt. 1), 21–35.
- Holbrook, Allyson L., Jon A. Krosnick, Richard T. Carson, and Robert C. Mitchell (2000). Violating Conversational Conventions Disrupts Cognitive Processing of Attitude Questions. *Journal of Experimental Social Psychology*, 36(5), 465–94.
- Hyde, Janet S. (2005). The Gender Similarities Hypothesis. *American Psychologist*, 60(6), 581–92.
- Hyde, Janet S. (2014). Gender Similarities and Differences. *Annual Review of Psychology*, 65, 373–98.
- Irving, Zachary (2023, July 25–26). Discussion of Ella Whiteley's paper 'A Woman First and a Philosopher Second'. *PEASoup Blog*. Retrieved from <https://peasoupblog.com/2023/07/ella-whiteley-a-woman-first-and-a-philosopher-second-precis-by-susanna-siegel/>
- Irving, Zachary, Samuel Murray, Aaron Glasser, and Kristina Krasich (2023). The Catch-22 of Forgetfulness: Responsibility for Mental Mistake. *Australasian Journal of Philosophy*, 75(1): 100–118.
- Jackson, Elisabeth (2022). What's Epistemic About Epistemic Paternalism? In Jonathan Matheson and Kirk Lougheed (Eds.), *Epistemic Autonomy* (132–50). Routledge.
- Jamieson, Annie, and Gregory Radick (2013). Putting Mendel Back in His Place: How Curriculum Reform in Genetics and Counterfactual History of Science Can Work Together. In Kostas Kampourakis (Ed.), *The Philosophy of Biology: A Companion for Educators* (577–95). Springer.
- Jamieson, Annie and Gregory Radick (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.
- Kahneman, Daniel, and Amos Tversky (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47(2), 263–91.
- Kesebir, Selin (2017). Word Order Denotes Relevance Differences: The Case of Conjoined Phrases with Lexical Gender. *Journal of Personality and Social Psychology*, 113(2), 262–279.
- Langton, Rae (1993). Speech Acts and Unspeakable Acts. *Philosophy and Public Affairs*, 22(4), 293–330.
- Langton, Rae, and Caroline West (1999). Scorekeeping in a Pornographic Language Game. *Australasian Journal of Philosophy*, 77(3), 303–19.
- Lepoutre, Maxime (2021, July 12). Conspiracies and 'Counterspeech', *Cardiff University Blogs: Open for Debate*. Retrieved from <https://blogs.cardiff.ac.uk/openfordebate/2021/07/12/conspiracies-and-counterspeech/>

- Leslie, Sarah-Jane (2014). Carving Up the Social World With Generics. In Tania Lombrozo, Joshua Knobe, and Shaun Nichols (Eds.), *Oxford Studies in Experimental Philosophy, Vol. 1*. (208–231). Oxford University Press.
- Maass, Anne, Carmen Cervone and Illayda Ozdemir (2022). Language and Social Cognition. In Oliver Braddick (Ed.), *Oxford Research Encyclopedia of Psychology*. Oxford University Press.
- Machery, Edouard (2017). *Philosophy Within its Proper Bounds*. Oxford University Press.
- MacInnis, Bo, Joanne M. Miller, Jon A. Krosnick, Clifton Below, and Miriam Lindner (2021). Candidate name order effects in New Hampshire: Evidence from primaries and from general elections with party column ballots. *PLoS ONE*, 16(3): e0248049.
- Maitra, Ishani, and Mary Kate McGowan (Eds.) (2012). *Speech & Harm: Controversies Over Free Speech*. Oxford University Press.
- McGowan, Mary K. (2005). On Pornography: MacKinnon, Speech Acts, and “False” Construction. *Hypatia*, 20(3), 22–49.
- McGowan, Mary K. (2022). On Saliency and Sneakiness. In Sophie Archer (Ed.), *Saliency: A Philosophical Inquiry (177–192)*. Routledge.
- Meyer, Meredith and Susan A. Gelman (2016). Gender Essentialism in Children and Parents: Implications for the Development of Gender Stereotyping and Gender-Typed Preferences. *Sex Roles*, 75(9), 409–421.
- Mihailov, Emilian, Ivar R. Hannikainen, and Alex Wiegmann (2023). Sacrificing Objects Instead of Persons: Order Effects Without Emotional Engagement. *Philosophical Psychology*, 10(1), 1–20.
- Munton, Jessie (2022). Answering Machines: How to (Epistemically) Evaluate a Search Engine. *Inquiry*. 40(1), 1–29.
- Niker, Fay (2018). Policy-Led Virtue Cultivation: Can We Nudge Citizens Towards Developing Virtues? In Tom Harrison and David Walker (Eds.), *The Theory and Practice of Virtue Education (153–167)*. Routledge.
- Noggle, Robert (2017). Manipulation, Saliency, and Nudges. *Bioethics*, 32(3), 164–170.
- Oeberst, Aileen and Christina Matschke (2017). Word Order and World Order. *Journal of Experimental Psychology: General*, 146(5), 672–690.
- Puddifoot, Katherine (2017). Dissolving the Epistemic/Ethical Dilemma Over Implicit Bias. *Philosophical Explorations*, 20(Sup. Vol. 1), 73–93.
- Rácz, Péter (2013). *Saliency in Sociolinguistics: A Quantitative Approach*. De Gruyter Mouton.
- Richard, Mark (2008). *When Truth Gives Out*. Oxford University Press.
- Rippon, Gina, Rebecca Jordan-Young, Anelis Kaiser, and Cordelia Fine (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.
- Roller, Colleen (2011, March 7). How Anchoring, Ordering, Framing, and Loss Aversion Affect Decision Making. *UX Matters*. Retrieved from <https://www.uxmatters.com/mt/archives/2011/03/how-anchoring-ordering-framing-and-loss-aversion-affect-decision-making.php>
- Sher, Shlomi and Craig R. M. McKenzie (2006). Information Leakage from Logically Equivalent Frames. *Cognition*, 101(3), 467–94.
- Siegel, Susanna (2017). *The Rationality of Perception*. Oxford University Press.
- Siegel, Susanna (2022). Saliency Principles for Democracy. In Sophie Archer (Ed.), *Saliency: A Philosophical Inquiry (235–266)*. Routledge.

- Sousa, David. A. (2006). *How the Brain Learns* (Third Edition). Corwin Press.
- Spender, Dale (1980). *Man Made Language*. Routledge.
- Strull, Thomas K. and Robert S. Wyer (1979). The Role of Category Accessibility in the Interpretation of Information About Persons: Some Determinants and Implications. *Journal of Personality and Social Psychology*, 37(10), 1660–1672.
- Stanley, Jason S. (2015). *How Propaganda Works*. Princeton University Press.
- Sullivan, Jessica (2018). The Primacy Effect in Impression Formation: Some Replications and Extensions. *Social Psychological and Personality Science*, 10(4), 432–439.
- Thaler, Richard and Cass Sunstein (2008). *Nudge: Improving Decisions About Health, Wealth and Happiness*. Yale University Press.
- Tirrell, Lynne (2012). Genocidal Language Games. In Ishani Maitra and Mary K. McGowan (Eds.), *Speech and Harm: Controversies Over Free Speech* (174–221). Oxford University Press.
- Tversky, Amos and Daniel Kahneman (1974). Judgment Under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–31.
- Watzl, Sebastian (2017). *Structuring Mind: The Nature of Attention and How it Shapes Consciousness*. Oxford University Press.
- Whiteley, Ella (2023). A Woman First and a Philosopher Second: Relative Attentional Surplus on the Wrong Property. *Ethics*, 133(4), 473–528.
- Willis, Kiersten (2017, July 16). Zoe Kravitz Responds to Criticism Over ‘I Am Not a Black Artist’ Image by Proclaiming ‘#ArtIsArt’. *Atlanta Black Star*. Retrieved from <https://atlantablackstar.com/2017/07/16/zoe-kravitz-responds-criticism-not-black-artist-image-proclaiming-artisart/>
- Zheng, Robin and Nils-Hennes Stear (2023). Imagining in Oppressive Contexts, or What’s Wrong with Blackface? *Ethics*, 133(3), 329–472.