

A structural explanation of injustice in conversations: It's about norms

Saray Ayala-López

Penultimate draft. Please cite published version Pacific Philosophical Quarterly 99(4) :726-748 (2018)

Abstract

In contrast to individualistic explanations of social injustice that appeal to implicit attitudes, structural explanations are unintuitive: they appeal to entities that lack clear ontological status, and the explanatory mechanism is similarly unclear. This makes structural explanations unappealing. The present work proposes a structural explanation of one type of injustice that happens in conversations, discursive injustice. This proposal meets two goals. First, it satisfactorily accounts for the specific features of this particular kind of injustice; and second, it articulates a structural explanation that overcomes their unattractiveness. The main idea is that discursive injustice is not the result of biased interlocutors, but of problematic discursive norms.

0. Intro

In the current debate on explanations of social injustice we find two important elements. First, individuals' explicit and implicit attitudes towards those who are the target of injustice; and second, social structures and the material reality beyond individuals' minds that constraint individuals' judgements, decisions and actions. Depending on the explanatory weight given to each of these elements, explanations of social injustice lean towards either the individualistic or the structural side of the spectrum. Individualistic approaches proceed by identifying causally relevant mental states in individuals' minds. Jennifer Saul (2013), for example, has argued that what explains the underrepresentation of women in philosophy is, at least in part, implicit bias, that might result, amongst other things, in unfair evaluations of women philosophers' work. Implicit attitudes have also been said to explain racial/ethnic health disparities (e.g. Chapman, et al. 2012; Green et al. 2007), and several kinds of unjust treatment of nonnative speakers, e.g. negative judgments of their competence (Boyd 2003), credibility (Giles 1973) and intelligence (Lindemann, 2002), or the glass ceiling that affects them (Huang et al. 2013). These are but a few examples of the large body of research on the topic.ⁱ

While individualistic explanations appear relatively easy to understand, structural explanations resist an easy grasp. One of the main reasons for that is that in spite of its common use in social science and

philosophy, the notion of social structure remains vague. What is the nature of social structures? Do they constitute a reality beyond individuals? The explanatory process is also unclear. How do social structures affect individuals' behavior so that they can be said to play an explanatory role in the latter? Are structural explanations causal explanations that appeal to structures as causes? What is the mechanism connecting structural factors and episodes or patterns of injustice?

Individualistic explanations have two advantages that structural explanations lack. First, the main element in the explanation, the individual, has a clear ontological status; and second, the explanatory mechanism is relatively transparent, if only because it fits into a much exploited model of our folk psychology: there are mental states (e.g. desires, beliefs) that cause us to do certain things. This model also works for implicit or unconscious attitudes.ⁱⁱ Structural explanations, however, appeal to entities that lack clear, or at least intuitive ontological status, and the explanatory mechanism is similarly unclear or unintuitive. In their appeal to individuals' attitudes, individualistic explanations do not carry the ontological vagueness and explanatory complication that structural explanations do.

The present work has two main motivations. On the one hand, to resist the recent focus on implicit bias in the literature in our explanations of social injustice.ⁱⁱⁱ On the other, to resist the vagueness and complication of appeals to structural explanations. Expanding on previous work (Ayala & Vasilyeva 2015), the present work proposes a structural explanation of discursive injustice that meets two goals. First, it satisfactorily accounts for the specific features of this particular kind of injustice; and second, it articulates a structural explanation that overcomes both the vagueness and the complication. Therefore, this work contributes to both the literature on discursive injustice, and the current debate on explanations of social injustice. With regards to the latter contribution, it is my intention to provide the debate with an articulation of structural explanation that could serve as a model to develop concrete structural explanations of other types of social injustice.

I first introduce the notion of discursive injustice as it appears in Rebecca Kukla (2014), and discuss a bias-explanation of this type of injustice in terms of prejudices in the mind of interlocutors. Then I outline ways in which structural explanations identifying higher-level social constraints on behavior address an important gap in an individualistic understanding of social injustice. I spell out a structural explanation of discursive injustice appealing to factors such as the discursive norms that operate in the communicative

exchange, and the positions that speaker and interlocutor(s) occupy in that communicative framework, and argue that a structural explanation is necessary and also sufficient to account for discursive injustice. I conclude by addressing some objections to my proposal and suggesting a possible avenue to alleviate this type of injustice.

1. Discursive injustice and implicit bias

Discursive injustice occurs when a speaker's speech act is not given the appropriate uptake, in a way that distorts its performative force and disempowers the speaker (Kukla, 2014). *Appropriate* refers here to the uptake the speech act would get in some ideal circumstances. As we will see, these ideal circumstances involve social identity. That is, discursive injustice is tied to the social identity of the speaker, as a 'social identity swap' test would reveal: if we ask 'would the wrong uptake occur if the speaker belonged to a more privileged social identity category?', the answer would typically be 'no'. In cases of discursive injustice, the action that the speaker intends, and is entitled, to bring about fails, and a different action occurs, usually one that enhances a disadvantage, creates it and/or contributes to its perpetuation.^{iv} Thus, discursive injustice undermines the speech capacity of (some) speakers, understood not as the capacity to utter certain combination of words with certain content (locutionary capacity), but rather as the capacity to do certain things with their words (illocutionary capacity).^v Let's consider the following two scenarios as illustrations of discursive injustice:

NON-NATIVE

A non-native speaker of English is attending an academic talk in an English-speaking country. The non-native speaker makes a contribution to the discussion and nobody engages with it. Eventually a native speaker raises the same point and this time people engage in a lively discussion.

JUDGES

In June 2014, the Spanish Supreme Court hosted an open house and announced the event with big banners throughout Madrid. The banners contained an image of a nearly naked woman, taken from one of the paintings covering the interior walls of the building, with the word 'conócela' (get to know her) across her pubis. A group of women judges wrote an open letter to the president of the

Supreme Court asking to remove the image. These are two of the sentences that the letter contained: ‘this is not an appropriate image to represent justice’ and ‘it is an act of discrimination to use the image of a naked woman to promote an activity’. And this is how two Spanish newspapers covered the news: ‘an ad with a naked woman upsets women judges of the supreme court’.^{vi}

In both scenarios, the speaker’s intervention is not taken for what it is, and their perceived (and most salient in that context) social identity (as a non-native speaker in the former case, and as women in the latter) seems to be driving that. More specifically, in JUDGES the speakers intend to assert something, i.e. ‘the image is discriminatory’, using a standard formula to do so, but their audience^{vii} does not give their speech act the uptake corresponding to an assertion, so the assertive act fails. Their attempted assertion is taken as an expression of feelings instead (an expressive). There are several things they could be doing with that locution, not only an act of assertion; they could, for example, perform an act of protest. The important point here is that the uptake their speech act gets betrays what they wanted to do with their words, and therefore undermines their speech capacity and disempowers them as speakers.^{viii} In the NON-NATIVE scenario, a non-native speaker tries to make a proper contribution to a conversation, but fails to do so. There is plenty of research showing that non-standard accents serve as a platform for discrimination, often masked behind appeals to unintelligibility or lack of comprehension (Lippi-Green 1997; Baugh 2016). For example, non-native speakers are given less credibility compared to native speakers (e.g. Brown, Giles & Thakerar, 1985; Giles, 1973), and they are perceived as less intelligent and competent (Boyd 2003; Bradac 1990; Lindemann 2003). Interestingly, judgements of intelligibility of non-native accented speech are affected by attitudes towards to it (Rubin 1992; Linderman 2002). For the sake of my argument, let’s say our example is one of those cases; if so, the speaker didn’t get an appropriate level of credibility from their audience because of their perceived identity as a non-standard speaker. But how does this happen? Kukla describes *entreaties* as speech acts that are performed ‘in order to be granted status as a speaker with normative standing within a discursive subspace’ (Kukla, 2014, p. 9). *Entreaties* come from outside the discursive game; they ask ‘can I play too?’. They do not count as (proper) contributions. If an intended contribution is taken as an entreaty, this prevents the audience from evaluating its content at all, for it only counts as a *request* to be part of the conversation. In NON-NATIVE, the non-native’s contribution might have been taken as an entreaty, and so

the lack of audience engagement is better seen not as a response to its content, but to its (perceived) performative force.^{ix}

Many researchers have worked on similar forms of discursive malfunctioning, proposing different frameworks to understand its nature. The purpose of introducing it here is to have a concrete example of unjust phenomenon to focus on in our inquiry into the question of different kinds of explanations of it. An arguably attractive explanation of this malfunctioning in conversations refers to the (explicit or implicit) attitudes of the audience. Referring specifically to cases of testimonial injustice, Miranda Fricker appeals to prejudices in the interlocutor that ‘will tend surreptitiously to inflate or deflate the credibility afforded the speaker’ (Fricker, 2007, 17), and specifies that ‘[T]he main type of prejudice (...) that tracks people in this way is prejudice relating to social identity’ (27). According to this, interlocutor’s prejudices about the speaker’s perceived social identity are at the basis of this type of injustice.^x This bias-explanation, tracing particular episodes or patterns of injustice to bias in individuals’ minds, is favored by a large body of psychological and philosophical literature drawing upon empirical research on implicit bias. This individualistic view with regards to explanation shapes the development of interventions that target different kinds of discriminatory behavior (e.g. studies on changing or controlling implicit biases).^{xi} In the next section I explain a possible limitation of bias-explanations using discursive injustice as a case study. Identifying that limitation paves the way for the delineation of a structural explanation.

2. Social situatedness: Off line, and online, too.

Bias-explanations might face some limitations to appropriately situate the individual in their social reality. In order to articulate this possible limitation, I draw on distributed and extended cognition research in the philosophy of mind (Clark 2008; Menary 2010; Sutton 2010) to defend a full-bodied situatedness, according to which the agent is in constant interaction with what is outside of their skin, continuously outsourcing to external (material and abstract) resources, be they material aids (Clark & Chalmers 1998) or cultural resources, including other people and social entities like institutions (e.g. Hutchins 1995, 2011; Gallaguer, 2013). In order to successfully navigate a city or recall a friend’s birthday, we might store the relevant information in our internal memory or on some external aid (e.g. smartphone), and retrieve it when needed (consulting either the internal or the external resource). Our minds are outsourcing-hungry, ‘we are

particularly adept at creating and recruiting environmental props and scaffolds (media, tools, artifacts, symbol systems) for the purpose of solving problems' (Theiner, 2013). That is, if we have the chance to outsource, we will. This makes existing external resources critically important to understand and explain our behavior.

This tight coupling between agent and external resources can be seen in different ways, and as having different implications, e.g. as a challenge to the traditional picture of the metaphysics of the mind as something contained in the head. Here I am concerned with the explanatory, rather than the metaphysical consequences of considering agents as situated and extended. According to this, when explaining a cognitive process or behavior, we need to include the external factors with which the agent is coupled, for they account for as important a part of the outcome as the internal elements. However, bias-explanations do not necessarily ignore or exclude external resources. They might, and often do, include external factors as remote causes of the mental states that are relevant in the explanation of behavior. That is, a bias-explanation might assume that social norms and other cultural dynamics at some point caused the agent to acquire implicit biases, and these implicit attitudes do the relevant explanatory job in our explanation of the behavior of interest. I argue that appealing to external factors in this way is unsatisfactory. The time parameter is crucial here. Instead of pushing the explanatory relevance of external factors back in time, as having to do with the origin of a causal chain that ends up in a particular behavior via a series of mental states, I see external factors as importantly explanatorily relevant at the time the behavior we want to explain occurs.

If bias-explanations are to account for external factors, we need to make sure they don't ignore the on-time effects of social reality at the real time of the episode of injustice. A complete account of interlocutors' situatedness in the social reality must consider not only their mental states, even if we see them as a result of their previous exposure to social factors, but also the social factors themselves that, at that moment and in that space, are guiding their interpretation of speaker's words. My hypothesis is that structural explanations can accommodate this full-bodied situatedness much easily. I emphasize that our choices about what interpretation to give to others' speech are often directly guided by external factors (cues we detect often without awareness), instead of being the end result of an internal evaluation, e.g. a judgment about which interpretation is the best one. The general idea here is that our choices in everyday situations i) are not always, and perhaps only rarely, the result of a deliberative process determined by our beliefs or other relevant mental states, and ii) offload part of the cognitive work onto existing material and cultural resources, e.g.

norms. For the first point, I draw upon Richard Holton (2006)'s analysis of choice as an act that is not always preceded by a judgment, and defend that interpretative decisions in everyday conversations are choices that more often than not are not the result of a previous judgment.^{xii} These choices respond to cues in the context of conversation that we somehow register, without being necessarily aware of them. Thus, an explanation of our interpretative choices in conversations that appeals to our beliefs or other relevant mental states is not enough, even if these mental states have been influenced by external factors, for external factors themselves might be what immediately precedes our choices. An advocate of bias-explanations might accommodate this concern, at least the part that refers to the mental process being deliberative, by simply clarifying that reference to implicit attitudes do not include reference to judgements of this sort.

For the second point, I draw upon the idea of scaffolding from the literature in extended cognition. The idea is that much of our cognitive processing is done by exploiting external resources, be them notebooks, computers or cultural practices. This exploitation can be done to specifically store information in order to ease our biological memory (e.g. sticky notes), and also as part of more general cognitive processes. The idea that cognition is extended not only onto material aids, but also on socio-cultural practices, has been explored by several researchers. For example, Ed Hutchins (2011), who talks of Enculturated Cognition, Shaun Gallagher (2013), who claims that the mind is extended onto institutions, and Richard Menary (2007, 2013), who argues that the human mind is socio-culturally extended in the sense that we 'learn how to think in richly structured cognitive niches that contain experts, skills, representations, artefacts, institutions, and so on' (28).^{xiii} I apply this same framework to what happens in our daily conversations when we interpret each other's interventions. We offload onto socio-cultural resources, e.g. norms that more or less implicitly govern and guide the conversation.^{xiv} An individualistic explanation would have a hard time, or so I think, accommodating this online explanatory role played by the external elements.

Someone might try to work out that accommodation by saying that explanations focusing on implicit attitudes do not assume a picture of individuals and their biases as isolated from their environments. Some sort of interaction between them is taken for granted. This is fair enough. The evil, however, is in the details. How does this interaction happen, and more important for my purposes, when? One way in which this interaction might happen is, as mentioned above, by external (social) factors being the cause of implicit attitudes. Another one is the reverse direction: by social norms piggy-bagging from implicit bias, that is,

implicit bias might influence the kind of social norms we get to have.^{xv} However, acknowledging this feedback loop between implicit attitudes and norms does not by itself account for one key element of the offloading process, and that is the timing. A distinction I introduce in Ayala (ms) comes handy here. This is a distinction between *origin-cause* and *maintenance-cause*, that is, between causal factors that are at the origin of a phenomenon, and causal factors that are responsible for the continued existence and maintenance of that phenomenon.^{xvi} The advocate of the individualistic view can definitely accommodate the idea that external factors are at the origin-cause of our implicit biases which in turn directly guide our discursive behavior, and so an explanation in terms of implicit biases somehow includes an indirect reference to those external factors. This saves the individualistic explanation from isolating the individual from its environment. The reverse direction is as easily accommodated: implicit biases might be part of the origin-cause of some social norms that in turn directly guides our behavior, so an explanation of the target behavior ultimately includes reference to implicit biases. In this reverse case, any reference to external factors ultimately includes a reference to implicit biases, so this secures that our explanations will have an individualistic flavor. I claim, however, that the individualist will have a harder time accommodating the details of maintenance-cause, that is, the way external factors and mental states interact in real time as to maintain certain patterns of behavior. More in particular, the on-line, as opposed to off-line, interaction that occurs at the moment of the discursive exchange (e.g. when interpreting the speaker's performative act). I defend that the relevant factors at that moment are the norms that apply to interlocutors' social identities, rather than any mental states representing the speaker as a member of a social group.^{xvii} Let me say more about the role played by external factors in our conversational exchanges.

Environmental factors not only facilitate (by providing norms as guides) but also constrain agents' behavior. When taking part in a conversation or any other social practice, we are part of a social structure, and this structure constrains our behavior in different ways. I draw on Sally Haslanger's notion of social structures as networks of social relations that are themselves constituted through practices (Haslanger, 2016). These practices locate us in different positions in the structure; Haslanger calls these positions *nodes*. From each node, a limited range of possibilities is available (i.e. possible actions to choose from). Someone occupying the node of mother/spouse in a monogamous heterosexual relationship in a community with no affordable childcare, and with a spouse earning a significantly higher salary due to a gender wage gap, has

at their disposition a very limited range of possible actions from which to choose when a baby arrives. When the child-bearer decides to quit their job, this decision does not come from an unconstrained decision-making process that straightforwardly reflects the individual's preferences.^{xviii}

Conversations, qua structured social interactions, situate interlocutors within nodes that afford a limited range of actions. The affordances of each position within a conversation (e.g., speaker, interlocutor) are adjusted in virtue of membership in intersecting social categories. For example, a homeless person has a very limited range of actions they can do with their words when addressing a passerby. In an interview to apply for a visa to enter the United States, the status of legal authority and epistemic privilege of the interviewer, and the burden of proof of a lack of malicious intent lying upon the interviewee, together reduce the range of possible things the interviewee can do with their words. Different dimensions of interlocutors' social location (race, ethnicity, sex/gender, socioeconomic status, ability, sexual orientation, immigration status, etc.) intersect in ways that stretch or shrink the range of possible speech acts. While a citizen of the US who tells an immigrant 'Speak English, you are in the United States' would be enacting discriminatory facts,^{xix} if the speaker is a homeless citizen the speech act has a different force.^{xx} The possibility to discriminate in that sense is perhaps out of that speaker's range of possible things they can do with their words.

A bias-explanation seems to leave aside the online guides and constraints that come with being situated in a structured space. In doing so, it paints a limited picture of the factors governing discursive injustice at the time of occurrence. I next argue that some of those structural factors can help delineate an explanation of discursive injustice.

3. Structural explanation

I mentioned above that the social position of an interlocutor in the broader social structure in which the conversational exchange is situated constrains the range of possible things this person can do with their words. There might be several structural factors playing an explanatory role for each particular social position, or in Haslanger's terms, for each node. To offer a structural explanation one does not have to commit to focusing on any one specific factor. In this work I don't provide a general analysis of what a structural explanation should look like in order for it to be satisfactory.^{xxi} My much more modest goal is to point out

one type of structural factor that plays an explanatory role in discursive injustice. And that is the norms or norms governing discourse. Norms can be difficult to individuate and define. For my purposes, it is enough to say that there are normal norms about what counts as a speech act of a particular type in a given context,^{xxii} together with normal ways to invoke those norms (e.g. particular words and gestures). In general we easily recognize, by interpreting words and gestures in any given conversational context, what norms an interlocutor is invoking, and we let those guide our interpretation of (the force of) their speech acts. We humans are in general very sensitive to norms; we are pretty good, from an early age, at detecting them, conforming to them right away, and punishing others for not doing so (Tomasello 2009).

My hypothesis is that our ordinary discursive interactions are governed by norms that systematically undermine the speech capacity of people perceived as occupying certain social positions, what we can call *social position-constrained norms*, C-norms for short. The existence of these norms makes discursive injustice a feature of well-functioning, although unjust, discursive exchanges. That is, this type of injustice is not the result of occasional misinterpretation of some people's speech acts, but a systematic phenomenon resulting from norms that are working more or less covertly. More important for our purposes here, it is not the result of unskilled listeners who, due to (implicit) biases against speakers' identity, do not appropriately apply the relevant discursive norms to that speaker's intervention. It is rather the unfortunate result of perfectly skilled listeners who are appropriately applying the norms operative in their communities.

Recall our scenarios JUDGES and NON-NATIVE. A C-norm in the JUDGES context could be 'when a person perceived as belonging to the category of women shows disapproval of some sort, she is expressing feelings'. This would prevent considering the speaker as making an assertion that can be true or false. A C-norm in the NON-NATIVE context could be 'when uttered in an unfamiliar way – including pronunciation and choice of words, speaker's interventions are attempts to join the conversation'. C-norms can be easily read as an adaptation of (at least some of) John Austin's felicity conditions for speech acts. Speech acts can misfire if they don't get the right uptake in the world. A condition to secure success for a speech act is that '[T]he particular person and circumstances in a given case must be appropriate for the invocation of the particular procedure invoked' (Austin 1975, 34). In Austin's own examples, if the would-be spouse in a wedding ceremony turns out to be a monkey, the attempt to perform a marriage fails, even if the right formula was uttered; if the person naming a ship doesn't have the right authority, the naming fails (Austin 1962). We can stretch the Austinian framework and say that what counts as an appropriate person for the speech act to be successful can be spelled out in terms of more nuanced social standing, e.g. race, gender/sex, class, immigration status.^{xxiii}

If I am right, we follow unjust, C-norms and we don't do anything about it. This might be due to mere moral laziness, but also to the fact that we are often wrong about the norms we follow. Mimicking the distinction between the manifest concept (the concept we take ourselves to be applying) and the operative concept (the concept we apply in practice, Haslanger 2006), we can talk of manifest and operative norms. Similar to the case of concepts, our intuitions about which discursive norms operate in our linguistic communities might be wrong, and the manifest norms might differ from the operative ones.

In at least some contexts, there exist C-norms that are covertly invoked by speaker's perceived social position, which jeopardizes their speech capacity.^{xxiv} In those contexts (C-contexts from now on), a structural explanation of discursive injustice is necessary and sufficient.^{xxv} It is necessary because a full-bodied situatedness in the social reality requires, as I defended above, considering factors outside individuals' heads that are both constraining and guiding their communicative interactions *in the moment*. How is it also sufficient? In order to argue for that I need to say something about norm-conforming behavior.

4. Norm-conforming behavior

How does the existence of C-norms help make the case for a structural explanation of discursive injustice? And in particular, how does that pave the way for the claim that a structural explanation in C-contexts is sufficient to account for discursive injustice? Someone could say that even if we grant the existence of C-norms in at least some contexts, they only work through the minds of the individuals, and so an explanation of discursive injustice will ultimately be in terms of individuals' mental states.^{xxvi} But we don't need to posit any specific mental state (in particular, implicit bias against speaker's perceived social identity) to explain the behavior of interlocutors who conform to operative C-norms at the moment of the episode. First, conforming to a norm can be an unreflective, automatic response to situational cues (Bicchieri 2006).^{xxvii} Complying to social norms is usually not a result of an evaluation of the situation and a decision to conform. This feature of norm-following is however easy to accommodate into the individualistic picture. Second, and more important for my purposes, the existence of a social norm (what for Bicchieri means that we feel expected and preferred to conform to it) gives us enough of a reason to conform to it. To clarify this power of norms I will refer to empirical work on folk attributions of mental states to others, and how they are influenced by norms. As suggested in Uttich & Lombrozo (2010), only when someone behaves in a way that violates a norm we go on to attribute to the agent some (other) reason to act that way, such as a norm-conflicting mental state. That is, the existence of a norm that prescribes a particular behavior in a given context can be seen as enough of a reason to exhibit that behavior in that context (no extra reason is required). Thus folk explanations of norm-conforming behavior do not have to invoke any mental state about the content of the norm (but instead only assume some set of cognitive processes related to norm-following). This could serve as initial support for the following hypothesis: no set of beliefs or other attitudes related to the content of the norm need to play a role in the production of norm-conforming behavior. This is a far-fetched hypothesis given that the empirical results just mentioned are about how we generally attribute mental states

to others, rather than about how mental states are actually related to norm-conforming behavior. Taking folk attributions as accurate descriptions of what is going on sounds too optimistic. Let me try a different avenue.

Let's consider two scenarios. If compliance to a particular norm required a preceding reflective process, we would expect that beliefs and other attitudes directly related to the content of the norm would be part of this process. Let's call this the mental scenario. If, on the contrary, this process is (sometimes, or even often) automatic and unreflective (let's call this the automatic scenario), we might say that those relevant beliefs and attitudes are not invoked in the production of the norm-conforming behavior. Now, in the mental scenario not all mental processes need to be conscious or explicit. Implicit bias (in this case against speaker's identity) would be one of such attitudes. However, what sets apart the mental and automatic scenarios is not whether the agent has relevant beliefs and attitudes (in some latent state), or how implicit/explicit or conscious/unconscious they are, but rather whether these beliefs and attitudes play a salient role in the particular process that leads to the norm-conforming behavior at that time. My hypothesis is that mental states that play a role in the mental scenario, play no salient role in the automatic scenario. This does not mean that in the automatic scenario the agent has no mental states at all. When following norms agents don't turn into zombies (do they?). It does not mean, either, that they do not have any mental state that is somehow related to the norm-conforming behavior. We can grant that a cognitive process related to norm-following (with detecting and exploiting existing norms in the local environment) might be involved (e.g. Stripada & Stich 2006). The agent could also use some sort of heuristic that takes situational cues as inputs, e.g. imitate the best, and this would entail having mental states that are somehow related to the existing norm, although not with its content. The idea I propose is that the application of this heuristic does not necessarily involve any or most of the relevant beliefs and attitudes that could be at work in the mental scenario. Take the example of spatial proximity to other people in a public setting. Someone familiar with the norms regulating spatial proximity in a particular context (e.g. informal settings in the US) probably has relevant (standing) beliefs and attitudes about appropriate distance for that context, but at the time of interacting with others, say at a party, the cognitive processes causally closest to their keeping a certain distance do not necessarily involve a mental state about the appropriate distance (that is, about the content of the existing norm), but a more automatic *tuning up*, perhaps a general cognitive structure dedicated to norm detection and exploitation, and the capacity to outsource onto the local environment. This story is not one in which agents are zombies, but one in which when engaged in (often rapid) social interactions, agents use their capacity to detect and exploit norms (which includes the use of simple cognitive processes, e.g. heuristics), and the structured environment, instead of mental states whose content includes the content of the existing norms.

Finding ourselves conforming to norms in the automatic scenario is pretty common, but it is not always the case. When a norm is broken, for example, the machinery of mental states related to the content of the norm is triggered and we jump from the automatic to the mental scenario, and in particular to an extremely deliberative version of it. When we are unfamiliar with operative norms in a context, we also rely on the mental scenario. Compare the example above with someone who knows the norms regulating spatial

proximity in that context but who is however not a skilled practitioner of them (e.g. they were raised in a different cultural practice and are adjusting). For this person, at least initially, the distance at which they place themselves relative to other people will be the result of a process that involves beliefs about the appropriate distance in that context (i.e. about the content of the norm). For my hypothesis to be useful in a defense of the potential of a structural explanation we don't need the automatic scenario to be always the case. Whenever we are in the automatic scenario and in a C-context (i.e. one in which C-norms are in place), a structural explanation of discursive injustice in terms of norms is sufficient.

Although I do not develop this proposal in detail here, I hope that this outline is enough to show the promising potential of a structural explanation that focuses on norms, even if this outline alone is not enough to disprove the potential of individualistic explanations. An objection, or perhaps a follow up to this (depending on where our sympathies are) is to say that while a focus on implicit attitudes targeting others' social identity might be misplaced, focusing on mental processes is still the right way to go, just not those ones.^{xxviii}

The relationship between agents and norms in the automatic scenario might still appear, however, obscure. Somewhere else I propose a relevant analysis that might serve as a first approximation to understand this relationship (Ayala 2016). This analysis naturalizes discursive norms by turning them into structured affordances,^{xxix} what I call *speech affordances*, which interlocutors can pick up or not. According to this view, when taking part in conversations, people tune up into what each situation invites. Given that speech affordances vary across social positions, what each situation solicits to different interlocutors depends on their different social position. Having access to affordances means having access to the relation 'between an ability we have and the situation out there in which our ability can be exercised.' (10). This is a possible way to define what is involved in speech affordances:

Relevant abilities: A general ability to communicate with others according to the norms in that context and a more specific ability to perform, exploiting the norms governing discourse in that context, a particular action with our speech acts.

Situational features: These are properties of the social environment that solicit both the general ability to communicate and the specific ability to perform a specific speech act. (10)

Detecting and adequately responding to available speech affordances does not require an articulation

of a norm about how to proceed (e.g. about which formulas to use when intending to perform a particular type of speech act, or how to interpret others' speech acts). I do not offer a defense of the framework of affordances here, but rely on both the reputation that this framework has in ecological psychology as a legitimate alternative to more intellectualist approaches to perception, and the defense of speech affordances in particular that I articulate somewhere else. If we concede that the framework of (social) affordances de-intellectualizes the process of coping with the environment, and in particular the process of picking up the relevant social cues in order to behave according to accepted patterns, then we have a way of bypassing mental states when explaining how people conform to C-norms. I claim that in C-contexts we do not need to postulate implicit bias in interlocutors' minds in order to explain discursive injustice; the existence of (unjust) social norms and norms governing speech interactions is enough to explain this phenomenon. Thus, a structural explanation in terms of norms is, in C-contexts, sufficient to explain discursive injustice.

5. Weighting structural explanation and individualistic approaches

I argued for two claims. First, that a structural explanation overcomes at least one critical flaw of the bias-explanation, i.e. the lack of online situatedness; and second, the stronger claim that in C-contexts, a structural explanation is necessary and sufficient to explain discursive injustice. An additional benefit of structural explanation, pointed out by Garfinkel (1988) and Haslanger (2016), is that it is more stable, for it holds across 'inessential perturbations' of situations. There are subtleties in prejudices and implicit biases that surely vary from one individual to the other, and also within the same individual across contexts, which make a bias-explanation of any single episode of discursive injustice to be limited to that episode. On the contrary a structural explanation, even though constrained by the cultural context and by eventual variations in discursive norms, is more stable across individual episodes and individual agents. This stability makes structural explanations better at providing an understanding of discursive injustice as a social phenomenon.

This stability could be, someone could worry, a problem, rather than a virtue, in cases of subtle and unstable phenomena, e.g. discriminative behavior that only occurs once in a while, or in a very subtle way. That is, if we assume that structural explanations pick up properties that hold across small perturbations in minds and contexts, it may seem that when the phenomena we want to explain consist of exactly small perturbations, a structural explanation will miss the point, and would perhaps *overpredict* the behavior we

aim at explaining. Thus a bias-explanation attentive to minor perturbations would be more appropriate. However, this worry is based on a mistaken assumption that structural explanations imply a deterministic picture. In order to articulate my response, let's first try to understand in more detail what those unstable and subtle phenomena consist of. There are at least two scenarios that would result in unstable and/or subtle phenomena. First, one in which there is a lot of noise and the phenomenon we want to identify is unstable; and a second case is a complex system where there are many different factors affecting the phenomenon we want to explain. A structural explanation would not misrepresent either case as a deterministic and simple one. Structural constraints, whether in the former (noisy) or latter (complex) case, do not determine that certain outcome will happen, but rather make it probable, and make it probable in different degrees (Vasilyeva 2016). What a structural explanation reveals is how the object of explanation (in this case, the discriminatory behavior) is situated 'in a network of relationships within a larger whole (structure) and how these relationships modify the probability distribution over possible states of the object/individual, relative to a hypothetical isolated case not embedded in any structure, and/or relative to other positions within the structure, or relative to being a part of a different structure' (Vasilyeva, in prep, 2).^{xxx} A structural explanation of a case of discursive injustice (e.g. an interlocutor who takes a woman's intended assertion for an expressive), whether we are in a noisy or a complex system, does not necessarily overpredict the quantity or regularity of that behavior; it tells us that this outcome is more likely to occur in that kind of context for individuals situated in that particular position compared to alternative contexts and positions. If unstable and subtle phenomena were a problem for an explanation in terms of structural constraints, then they would also be a problem for a bias-explanation. In the same way that having a racist implicit bias doesn't mean either that the person always has an extreme racist behavior in all contexts, nor that they always display racist behavior along the same dimension or in the same way, the existence of unjust discursive norms regulating our conversational interactions doesn't imply always displaying extreme unjust behavior, or unjust behavior along the same dimension or in the same way (e.g. always distorting assertions and turning them into expressives when the speaker is a woman).

At this point, a seemingly appropriate question to ask is how we can distinguish, given a particular case of discursive injustice, whether it is a case of prejudices in the interlocutor's mind profiling the speaker, or a case of C-norms directly guiding interlocutor's interpretation. This question might have no clear answer,

for episodes of discursive injustice might be brought about by temporally extended complex combinations of individuals' attitudes and discursive norms, where each component reinforces the others (e.g. C-norms might become more powerful through individuals' implicit or explicit identity-prejudices; and more interestingly, individuals' attitudes might come to existence, and later be reinforced, through conforming to C-norms in their discursive practices). In any case, the question about explanatory attribution in specific cases reflects concerns that deviate from the ones that guide my proposal. The above question is a request to identify proximal causes of particular episodes of injustice. A structural explanation, however, provides a way to understand discursive injustice more generally, rather than a straightforward causal story.^{xxx} In this sense, the individualistic and the structural explanations are compatible, since they aim at satisfying two different explanatory demands.

There are reasons to think that looking for proximal causes of particular episodes of discursive injustice is not the most appropriate enterprise if aimed at providing an understanding of this form of injustice, and in particular, an understanding that facilitates intervention. One reason is that this sort of inquiry prioritizes identifying causes as the goal of explanation, causes of a particular kind (i.e. proximal). Causal connections do not all have the same explanatory value, they do not all illuminate the same. Obtaining understanding of the explanandum, and in particular illuminating it in a way that helps design an intervention, might require going beyond the identification of (proximal) causes. A second reason, tightly related to the previous one, is that seeking proximal causes mostly provides information about particular instantiations of the phenomenon we want to explain, rather than a general picture of why that type of phenomenon happens at all. Relevant for this latter point is Fred Dretske's distinction between *triggering* and *structuring* cause. While the triggering cause causes the behavior now, the structuring cause is what makes it so that the triggering cause produces that effect (Dretske 1988). Let's take Dretske's own example. The Queen enters the room and Clyde stands up. The question 'Why did Clyde stand up when the Queen entered the room?' can be interpreted in different ways, corresponding to different contrast spaces. If what we want to know is why Clyde stood up right at that moment, and only that, then we are asking for the triggering (or proximal) cause. But we might want to know other things, like why Clyde, as opposed to someone else, stood up, or why did he stand up as opposed to doing something else. These other questions do not ask for the cause that triggered Clyde to stand up when the Queen entered, but rather, for what makes it the case that the queen

entering the room makes Clyde (as opposed to someone else) to stand up (as opposed to sing or lie down). These other questions call for identifying structuring causes, which in Haslanger's words 'situate Clyde within a social structure: Clyde is Queen Elizabeth's subject (he is a part in a larger whole)' (Haslanger 2016, 10).

In spite of the preference for structural explanations I have defended along the paper, I see an important reason to conclude that both individualistic bias-explanations and structural explanations are (and need!) to be part of a comprehensive approach to social injustice. As articulated elsewhere, these two types of explanations respond to different questions, and so they do not measure against each other (Ayala & Vasilyeva 2015).

6. Fighting injustice in speech

An intervention against discursive injustice based on the structural explanation I have offered would start by identifying the (unjust) norm-conforming behavior as the locus of change. Changing any negative (explicit or implicit) beliefs about speakers' social identity would be, in this picture, insufficient to alleviate the injustice. Moreover, it might not be necessary. If the goal is to change the C-norms, what we need is to change our practices, and not only (and not even necessarily) our beliefs or attitudes.

Fighting discursive injustice in this way would require, first, making the operative C-norms transparent. This can be a very difficult task. These norms might be less accessible to speakers than other norms, even those that did not require explicit pre-agreement, for even after speakers make an explicit rationalization of the (manifest) norms in their discursive communities, these covert unjust norms can escape their radar. One possible way to identify C-norms is by making hypotheses of possible candidates and testing whether they hold by measuring whether people experience a sense of norm-violation when interlocutors do not conform to them. For example, we could hypothesize that in context C1, non-native speakers' assertions are taken as entreaties. Then we compare people's reactions to two scenarios: where interlocutors conform to that norm, and where they don't (i.e. assertions are generally taken as assertions no matter who the speaker is, whether native or non-native), and observe whether they have a sense of violation in the latter. One way of measuring this sense of violation can be adapted from Uttich and Lombrozo (2010), who measured how people attribute intentional states to others. If, as Uttich and Lombrozo suggest, we are more likely to judge

actions to be intentional when they violate a norm (e.g. wearing a graduation gown on the beach) than if they conform to it (e.g. wearing a graduation gown in a graduation ceremony), patterns of intentional action attribution could be used to reveal the existence of norms.

Revealing C-norms seems to require a type of intervention similar to the one Kristie Dotson identifies as needed in order to overcome what she calls contributory injustice (Dotson 2012). Contributory injustice is a breed of epistemic injustice caused by individuals holding on to a prejudiced set of hermeneutical tools when interpreting others. Overcoming this requires individuals to jump outside of their hermeneutical tool-kit and seek alternative, more appropriate tools to interpret those whose contributions escape their limited interpretative resources. Acknowledging contributory injustice means acknowledging that there are several hermeneutical resources individuals could use to interpret others. Overcoming it requires both being aware of those alternative resources and the capacity to shift amongst them. This awareness seems to be also the first step in order to reveal C-norms.

Once we reveal the operative C-norms, we can take two different avenues, depending on whether we want to fight discursive injustice as an individual enterprise, or as a political project. The former option will involve finding ways of getting around these unjust norms. This is what you would do if you are a victim of discursive injustice and want to maximize your speech capacity without necessarily questioning the communicative framework that enables discursive injustice. If, however, regardless of whether or not you are a victim of this type of injustice, you want to question and change the unjust norms producing this kind of injustice, you would go for a broader project, aimed at alleviating the existing injustice, not only avoiding it. The latter project will include establishing target norms, with the goal of turning the operative communicative framework into a just one, in which the speech capacity of speakers is not constrained as a function of their social position.

Finding out formulas to successfully do with your words what you want against discursive norms that systematically undermine your speech capacity is challenging, but possible. The fact that many people from disadvantaged groups get to succeed in different social situations in spite of discursive injustice might be the result of alternative formulas those people found to get around C-norms. If the goal is, however, pursuing a social change rather than sporadic individual successes, we have to aim at the norms themselves. That is undoubtedly a much more challenging, but also a more desirable project. It might happen that by

merely making explicit the C- norms operating in our interactions, we are already paving the way for a change.

The above reference to Dotson's work can serve an additional purpose here, and that is to wave at the idea that interventions on discursive injustice would involve both individualistic and structural changes.^{xxxii} This idea shifts the discussion over social injustice from an opposition between individualistic and structural approaches, to ways of combining those two approaches. It does so, however, for intervention purposes, which is a different conversation from the one we are mostly interested here, and that is about explanations.^{xxxiii}

4. Concluding remarks

I identified online social situatedness as an important element in discursive injustice that structural explanations can easily accommodate. I pointed out that bias-explanations would have a harder time accounting for the *in the moment* influence of social norms on individuals' conversational exchanges. I situated interlocutors in the matrix of social forces and relationships they are part of at the moment of taking part in conversations, and argued that a structural explanation that appeals to their positions in the structure and the discursive norms operative upon those positions is necessary and, in some contexts, also sufficient to account for discursive injustice. A structural explanation of conversational injustice is compatible with the existence of implicit bias in people's minds and also with their playing an explanatory role. Structural explanations might, however, respond to a more demanding explanatory enterprise than bias-explanations. Finally, I waved on several occasions at the idea that bias-explanations and structural explanations are not necessarily opposite. The introduction of the diachronic dimension, i.e. the distinction between origin-cause and maintenance-cause, allows an easy way out of a heated opposition between them: causal contributions at different times in the development of patterns of social behavior likely include implicit (and explicit) attitudes, and norms and other external factors.

References

- Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Oxford University Press.
- Ayala López, S. (in prep). Structures and Agency. *Metaphysics and Explanations*.
- 2016. Speech Affordances: A structural take on how much we can do with our words. *European Journal of Philosophy* 24(2): 879-891
- Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- Boyd, S. (2003). Foreign-Born Teachers in the Multilingual Classroom in Sweden: The role of Attitudes to Foreign Accent. *International Journal of Bilingual Education and Bilingualism*, 6: 283–95.

- Brown, B. L., H. Giles, and J. N. Thakerar. (1985). Speaker Evaluations as a Function of Speech Rate, Accent and Context. *Language and Communication*, 5: 207–220.
- Chapman, E., A. Kaatz, and M. Carnes, M. (2013). Physicians and Implicit Bias: How Doctors May Unwittingly Perpetuate Healthcare Disparities. *Journal of General Internal Medicine*, 28(11): 1504-1510.
- Clark, A., and D. Chalmers. (1998). The Extended Mind. *Analysis*, 58: 10-23.
- Dotson, K. (2011). A Cautionary Tale: On Limiting Epistemic Oppression. *Frontiers* 33(1): 24-47.
- Dretske, F. (1988). *Explaining Behaviour: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Gallagher, S. (2013). The Socially Extended Mind. *Cognitive Systems Research*, 25-26: 4-12.
- Gilbert, M. (2008). Social Norm Revisited. *Topoi* 27: 5–16.
- Giles, H. (1973). Communication effectiveness as a function of accented speech. *Speech Monographs*, 40: 330–331.
- Green, A., D. Carney, D. Palin, L. Ngo, K. Raymond, L. Iezzoni, and M. Banaji. (2007). Implicit Bias Among Physicians and its Prediction of Thrombolysis Decisions for Black and White Patients. *Journal of General Internal Medicine*, 22(9): 1231-1238.
- Haslanger, S. (2016). What is a (Social) Structural Explanation? *Philosophical Studies* 173(1): 113-130.
- (2015). Social Structure, Narrative, and Explanation. *Canadian Journal of Philosophy*, 45(1): 1-15.
- (2006). Philosophical Analysis and Social Kinds: What Good are Our Intuitions? *Proceedings of the Aristotelian Society*, Sup. Vol. 80: 89– 118.
- Holton, R. (2006). The Act of Choice. *The Philosophers' Imprint*, 6(3):1-15.
- Hornsby, J., and R. Langton. (1998). Free Speech and Illocution. *Legal Theory* 4(1):21-37.
- Huang, L., M. Friderger, and J. Pearce. (2013). Political Skill: Explaining the Effects of Nonnative Accent on Managerial Hiring and Entrepreneurial Investment Decisions. *Journal of Applied Psychology* 98(6): 1005–17.
- Hutchins, E. (2011). Enculturating the Supersized Mind. *Philosophical Studies*, 152(3): 437-46.
- Kukla, R. (2014). Performative Force, Norm, and Discursive Injustice. *Hypatia*, 29(2): 440-457.
- Langton, R. (1993). Speech Acts and Unspeakable Acts. *Philosophy & Public Affairs*, 22: 293-330.
- Lewis, D. (1969). *Norm: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Lindemann, S. (2002) Listening with an Attitude: A Model of Native-Speaker Comprehension of Nonnative Speakers in the United States. *Language in Society*, 31: 419–41.
- Madva, A. (2016). A Pleat for Anti-Antiindividualism: How Oversimple Psychology Mislead Social Policy. *Ergo* 3: 701-728.
- McGowan, M. K. (2004). Conversational Exercitives: Something Else We Do With Our Words. *Linguistics and Philosophy*, 27(1): 93–111.
- . (2009). Oppressive Speech. *Australasian Journal of Philosophy* 87 (3): 389-407.
- Medina, J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistance Imaginations*. Oxford: Oxford University Press.
- . (2011). The Relevance of Credibility Excess in a Proportional View of Epistemic Injustice: Differential Epistemic Authority and the Social Imaginary. *Social Epistemology*, 25(1): 15-35.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave Macmillan.
- (2010). *The Extended Mind*. Cambridge, MA: MIT Press.
- (2013). Cognitive Integration, enculturated cognition, and the socially extended mind. *Cognitive Systems Research* 25-26: 26-34.
- Saul, J. (2017). Are Generics Especially Pernicious? Inquiry DOI <http://dx.doi.org/10.1080/0020174X.2017.1285995>
- . (2013). Implicit Bias, Stereotype Threat and Women in Philosophy. In F. Jenkins and K. Hutchison (Eds) *Women in Philosophy: What Needs to Change?* Oxford: Oxford University Press.
- Stripada, C. and S. Stich. (2006) A Framework for the Psychology of Norms. In Peter Carruthers, Stephen Laurence & Stephen P. Stich (eds.), *The Innate Mind, Volume 2: Culture and Cognition*. Oxford: Oxford University Press.
- Sutton, J. (2010). Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process. In *The Extended Mind*, edited by R. Menary. MA: MIT Press.
- Theiner, G. (2013). Writing in Mind. *AVANT*, 4(2): 15-29.
- Tomasello, M. (2009). *Why We Cooperate*. Cambridge, MA: MIT Press.

Uttich, K., and T. Lombrozo. (2010). Norms inform mental state ascriptions: a rational explanation for the side-effect effect. *Cognition*, 116: 87-100.

Vasilyeva, N. (in prep.). Structural Explanations.

---- (2016). Structural Explanations: Structural Factors as Moderators and Constraints on Probable Outcomes, presented at the 42nd Meeting of The Society for Philosophy and Psychology, June 1-4 2016, Austin, TX.

Saray Ayala-López
Philosophy Department
California State University Sacramento

Acknowledgements

I'm deeply grateful to Nadya Vasilyeva for long discussions on structural explanations. I'm also grateful to two anonymous referees of this journal for their detailed and helpful comments.

ⁱ See the “Implicit Bias & Philosophy” project for relevant literature in philosophy and psychology, <http://www.biasproject.org/>

ⁱⁱ This is not to say that explanations that appeal to psychological phenomena do not encounter its own resistance, as they are sometimes seen to be difficult to combine with a materialistic understanding of the world.

ⁱⁱⁱ This explanatory focus might have as a consequence the idea that *fixing minds* is the way to go in our quest for social justice. See Ayala (in prep) for an articulation of some problems with this sort of focus in intervention. Also, the explanatory focus on implicit bias might derail us from paying attention to equally or even more important elements in the picture. A similar worry has been expressed by Jennifer Saul with regards to generic statements (Saul 2017). She writes “I view the current attention to generics as a mistake in the battle against prejudice” (1) Saul’s reasons are not that generics are not problematic, but rather that they are but one of many problematic constructions, so focusing on them is not justified. It’s worth noting that applauding the merits of bias-explanations is of course compatible with a much wider focus of intervention.

^{iv} The notion of discursive injustice is based on an important (and widely accepted) assumption according to which audience’s uptake determines what one gets to do with their words (see Austin, 1962).

^v The phenomenon that Rae Langton identified as *illocutionary disablement* (Langton, 1993) would be an extreme case of this. Discursive injustice is closely related to *testimonial injustice* (Dotson, 2011; Fricker, 2007). In cases of testimonial injustice, a speaker is granted a deflated level of credibility; their testimony is undermined and it doesn’t get to count for what it is. While not all cases of discursive injustice involve deflation of speaker’s credibility, in all cases of testimonial injustice some sort of distortion of the speaker’s speech act occurs. In both discursive and testimonial injustice speaker’s social identity trumps the force of their speech. An advantage of the notion of discursive injustice is that it not only identifies a type of conversational malfunction, but could also provide a possible mechanism for testimonial injustice. When a speaker’s credibility has been deflated, what is exactly what the speaker gets to do with their words? What is what the judges end up doing with their attempted assertion, and what does the non-native speaker get to do with their attempted contribution to the discussion? By analyzing cases of testimonial injustice with the tools of discursive injustice we can specify what happens to the speaker’s words.

^{vi} El Pais and El Economista, June 11, 2014

http://politica.elpais.com/politica/2014/06/11/actualidad/1402491426_341616.html ;

<http://www.eleconomista.es/espana/noticias/5856766/06/14/El-vicepresidente-del-Supremo-destaca-la-sensibilidad-de-las-disenadoras-del-cartel-con-una-mujer-desnuda.html>

^{vii} For the sake of simplicity, I take the two aforementioned newspapers to be the audience.

^{viii} If we take the intended act of the judges to be one of protest, the unfortunate result of discursive injustice in this case is even more upsetting: by having their protest taken as an expression of feelings, the discrimination they are trying to protest gets reinforced.

^{ix} Epistemic trustworthiness has two components, competence and sincerity (Fricker, 2007). In a case of testimonial injustice, any of those, or both, can be impugned. Given the aforementioned research showing that non-native accent undermines perceived competence of the speaker, it seems that competence, but not

sincerity, is what gets impugned when a non-native speaker is granted deficient credibility. And competence, more clearly than sincerity, seems to be directly related to the performative force of the speaker's speech acts. The defective attribution of competence to a non-native speaker might be happening via a recognition of their speech act as an entreaty, or in other words, via a lack of recognition of their speech act as an intervention within the discursive game.

^x José Medina (2013) expands on Miranda Fricker's account of testimonial injustice and seems to point out the need of a less individualistic approach when he says that 'in order to address the injustice we have to go well beyond the individual involved in the exchange: we have to go to the *social roots* of the problem' (ibid. p.80).

^{xi} My work does not depend on whether a bias-explanation of discursive injustice an explanation has been articulated and defended as one that is in opposition to a structural explanation. My goal is to outline how a structural explanation could be articulated, and to use discursive injustice as a case study to show some important elements that are easily capture by such a structural explanation.

^{xii} One relevant comment here is that the processes leading to the choice might not consist of a conscious judgment, but perhaps an unconscious one, guided by unconscious beliefs. Holton is willing to limit his proposal to conscious judgments, that is, that there is often choice in the absence of conscious judgment (Holton 2006, 9). This modified version is enough for my purposes here.

^{xiii} Krueger (2010) proposes an account of social cognition as partially constituted by environmental scaffolding.

^{xiv} Davidson & Kelly (ms) propose an account that aims at bridging the gap between individualistic and structural approaches, and that also appeals to social norms.

^{xv} I thank one anonymous reviewer for inviting me to make this explicit.

^{xvi} This distinction is inspired by the notions of *morphogenesis* and *morphostasis* (Buckley 1967; Archer 1979).

^{xvii} A relevant consideration here is: are there attitudes that might be relevant at the moment of interpreting others' speech, and which are not implicit biases against the speaker's social identity? In his sophisticated defense of the individualistic approach, Alex Madva (2016) emphasizes the significance of other attitudes (e.g. attitudes towards the malleability of social systems). These other attitudes might be crucial for a successful intervention against injustice, as suggested by studies like Johnson and Fujita (2012). Madva uses this as part of his strategy against those who criticize the individualistic approach by merely pointing out the problems with implicit biases.

^{xviii} See Haslanger (2015a) for a detailed analysis of this example.

^{xix} See Mary Kate McGowan (2004; 2009) for a framework of how enactment of oppressive facts can occur in casual conversations.

^{xx} I take this example from Ayala (2016, 4).

^{xxi} See Vasilyeva (2016; in prep) for an analysis of structural explanation that influenced this work.

^{xxii} Very roughly, norms are regularities in behaviour, which we deploy either consciously or unconsciously, and through which we organize and coordinate our behaviour. Contrary to Lewis' traditional account (Lewis, 1969), and like many others, I don't take norms to require common knowledge or to be advantageous on all occasions (see Gilbert, 2008).

^{xxiii} Thanks to an anonymous referee of this journal for pointing out this way of accommodating C-norms into Austin's framework.

^{xxiv} A detailed analysis of these norms is needed, one that would specify, among other things, (1) whether there is a norm for each and every type of speech act, or they are rather general and distort any kind of speech by toning it down – e.g. command-to-request, assertion to expressive; (2) how speaker's intersectional identities affect the operation and detection of these norms; and (3) how these norms regulating conversations relate to other kinds of social injustice. Echoing what Medina calls credibility excess (Medina 2011), which is the opposite of the credibility deficit that happens in testimonial injustice, we can ask whether there are also norms that somehow boost speech capacity by toning up speech acts' performative force, e.g. request to command. I leave this in-depth exploration for a future work.

^{xxv} What about contexts in which we do not know whether these unjust norms are in place? Given the lack of information, it would be reasonable not to discard either factor and consider both individualistic and structural considerations as potentially explanatorily adequate.

^{xxvi} Haslanger (2016) offers a possible response to this concern that draws on a distinction between *triggering* and *structuring* cause (Dretske, 1988). I briefly introduce this towards the end of the paper.

^{xxvii} See Bicchieri (2006) for a detailed distinction between norms, descriptive norms, and social norms. For my purposes here we don't need to distinguish among them.

^{xxviii} This is a possibility that I mentioned in footnote xviii, and which Madva (2016) emphasizes.

^{xxix} The notion of affordances, defined as possibilities for actions, was introduced in psychology by Gibson (1979) in his account of perception. 'The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*' (127, emphasis in the original). The notion has been elaborated and refined by many others.

^{xxx} See Ayala-López (in prep) for an alternative understanding of structural explanations in modal terms.

^{xxxi} As suggested in Vasililyeva (2006; ms), structural explanations might not be causal explanations after all.

^{xxxii} Dotson writes "Both the structurally prejudiced or biased hermeneutical resources and the agent's situated ignorance are catalysts for contributory injustice. As such, it is located within the gray area between agential and structural perpetuation of epistemic injustice." (2012, 31).

^{xxxiii} See Ayala López (in prep) for a defense of the importance to keep these different projects separate.