

Deliberative Injustice: a Collective Failure

Moisés Barba¹

Recibido: 5-7-2021 / Aceptado: 19-1-2022 / Publicado: 30-7-2021

Abstract. The purpose of this paper is to expand the theoretical field of discursive injustice by identifying a specific kind of discursive injustice, namely, the kind we are subject to when we are unjustly prevented from exchanging reasons with others. Broadly speaking, discursive injustice is the kind of injustice we suffer when we are unjustly harmed as language users, most notably when we are prevented from using language in ways we are entitled to. The dominant approach to discursive injustice has focused on the corruption of the illocutionary force of a speech act due to the hearer's improper uptake. I claim that there is a genuinely distinct kind of discursive injustice, which I label as *deliberative injustice*, that cannot be accounted for by the dominant approach to discursive injustice. In my view, what makes deliberative injustice discursively unjust is that it amounts to denying someone a normative position that she is entitled to as a language user, which I understand in terms of being a *source of reasons*. Moreover, I explore two ways of trying to illuminate deliberative injustice by appealing to the social level of analysis, namely, by resorting to a *structural explanation* (Ayala, 2018) to explain how it comes about, and by asking what it signifies for the collective life of groups. In this regard, I show that the structural explanation has some explanatory limitations, and that deliberative injustice essentially constitutes a *collective failure*.

Keywords: deliberation; discursive injustice; social explanations; speech acts theory.

[es] Injusticia deliberativa: un fracaso colectivo

Resumen. El propósito de este artículo es expandir el campo teórico de la injusticia discursiva identificando un tipo específico de injusticia discursiva, a saber, la que sufrimos cuando se nos impide injustamente intercambiar razones con otros. Entendida de modo amplio, la injusticia discursiva es el tipo de injusticia que sufrimos cuando recibimos un daño injusto como usuarios del lenguaje, de manera más señalada cuando se nos impide emplear el lenguaje de un modo al que tenemos derecho. El enfoque dominante acerca de la injusticia discursiva se ha centrado en la corrupción de la fuerza ilocutiva del acto de un acto de habla debido a la incorrecta recepción del mismo por parte del oyente. Defenderé que hay un tipo genuinamente distinto de injusticia discursiva, que llamo *injusticia deliberativa*, que el enfoque dominante acerca de la injusticia discursiva no puede explicar. A mi ver, lo que hace discursivamente injusta a la injusticia deliberativa es que supone negarle a alguien una posición normativa a la que tiene derecho como usuario del lenguaje, y que entiendo en términos de ser una *fuerza de razones*. Además, exploraré dos maneras en que podemos intentar iluminar la injusticia deliberativa apelando al nivel social de análisis, a saber, recurriendo a una *explicación estructural* para explicar cómo tiene lugar, y preguntándonos qué significa para la vida colectiva de los grupos. A este respecto, argumentaré que la explicación estructural tiene algunas limitaciones explicativas, y que la injusticia deliberativa constituye un fracaso *colectivo*.

Palabras clave: deliberación; injusticia discursiva; explicaciones sociales; teoría de los actos de habla.

Cómo citar: Barba, M. (2022). Deliberative Injustice: a Collective Failure. *Las Torres de Lucca. Revista internacional de filosofía política*, 11(2), 347-356. <https://dx.doi.org/10.5209/ltld.77037>

My main goal in this paper is to expand the field of discursive injustice to cover a particular kind of discursive injustice, which I label as *deliberative injustice*. Over the last few years, a number of works have drawn attention to the fact that linguistic exchange is one more means through which we can be subject to, or commit, an injustice. To illustrate, here are some typical examples:

Rejecting

In a famous passage of Jane Austen's *Pride and Prejudice*, Elizabeth Bennet repeatedly rejects Mr. Collins's marriage proposal without him acknowledging the refusal. Mistaking her words for a mere display of feminine coquetry meant to motivate further gallantry on his part, Mr. Collins declares that he knows it to be "...the established custom of your [Elizabeth's] sex to reject a man on the first application" and resolves to attribute her response to her wish of "...increasing my love by suspense, according to the usual practice of elegant

¹ Universidad Autónoma de Madrid.
E-mail: moises.barba@hotmail.com.
ORCID: <https://orcid.org/0000-0002-5959-0135>

females". In the passage, Elizabeth fails time and again to dissuade Mr. Collins: the more unequivocally she asserts her intention not to marry him, the more certain Mr. Collins becomes that she is merely engaging in a sophisticated form of flirtation.

Commanding

Celia is a manager at a factory, and all her fellow workers are men. At giving an order, she resorts to the relevant linguistic conventions, e.g., she uses a voice tone usually associated with giving orders. However, her orders are met with low compliance on the part of her subordinates, as they take them to be mere requests.²

It should be clear that Elizabeth Bennet and Celia are subject to an injustice of some sort: however correct Elizabeth Bennet's and Celia's use of the relevant linguistic conventions may be, their words are not taken by the hearers in the way they intended. Thus, neither Elizabeth Bennet's rejection of a marriage proposal nor Celia's commands are acknowledged as such and consequently they are, in a sense, deprived of something they are owed. Importantly, the way Elizabeth Bennet and Celia are wronged is as language users – as subjects capable of, and entitled to, using language under the appropriate circumstances, e.g., by using the appropriate linguistic conventions. In particular, the subjects in the previous examples are prevented from giving rise to some state of affairs through linguistic interaction with others. That is to say, what they are owed, and what they are deprived of, is their capacity to *do things with words*. In the literature, this kind of injustice is labeled as *discursive injustice*.

There are several different takes on the nature of discursive injustice. According to Langton's approach (Langton, 1993) discursive injustice *prevents* a speech act from being fully performed, so that, for instance, Elizabeth Bennet's words do not amount to an actual rejection of Mr. Collins's proposal. On Kukla's account (Kukla, 2014) discursive injustice consists not in preventing someone's words from doing something at all, but in *distorting* them so that what they finally do differs from what the speaker intended them to do. Thus, Celia's orders are distorted to the point that they do not count any longer as orders, but as mere requests. And, on Bianchi's view (Bianchi, 2020) both Elizabeth Bennet's and Celia's words really constitute the intended speech acts, but a failure at the communicative level prevents them from doing with words what they intended to –that Mr. Collins desists from proposing marriage to Elizabeth Bennet, and that Celia's fellow workers do as she commands. Though differing in the detail, all these accounts coincide in treating discursive injustice as a matter of making a speech act fail by depriving it of its illocutionary force as a result of the hearer's undue uptake. In the remainder of this paper, I will use the expression 'standard approach' to refer to this general and dominant approach to discursive injustice that understands it in terms of the corruption of a speech act's illocutionary force through the hearer's undue uptake.

I claim that there is a distinct kind of discursive injustice that does not fit neatly in the model of the standard approach to discursive injustice, namely, deliberative injustice. This is a kind of discursive injustice that we may suffer in contexts where the linguistic interaction serves the purpose of exchanging reasons with others, most notably, in group deliberation contexts, where we come to one another to settle on a common view about some issue –e.g., when a trial jury deliberates about whether the defendant is guilty or innocent. As a discursive practice, deliberating with others provides an occasion to be unjustly treated and harmed as a language user. However, I will show that the standard approach to discursive injustice is not well suited to account for what makes deliberative injustice discursively unjust. To be clear, I do not intend to criticize the overall theoretical import of the standard approach, but merely to show that deliberative injustice does not fall under its scope but constitutes a distinct kind of discursive injustice. This is a lacuna that I aim to fill by providing an account of deliberative injustice that departs from the standard approach. Moreover, echoing recent contributions (Ayala, 2018) I will examine different ways in which appealing to the social level of analysis –i.e., social facts that are in place in linguistic interactions like that of exchanging reasons– is relevant for illuminating the phenomenon of deliberative injustice.

The paper is structured as follows: I begin in the next section by describing a case to illustrate how we can be subject to discursive injustices in reason exchange contexts. While the standard approach initially seems well fit to account for it, I will argue that other cases escape from its scope. After that, I will offer an original account of deliberative injustice as failing to acknowledge a participant in an exchange of reasons as a *source of reasons*, a notion I will cast out in terms of occupying a certain normative position that we are entitled to as language users. Thus, as I will argue, the nature of the harm inflicted in deliberative injustices differs in some relevant respects from that in standard discursive injustices, though, in both cases, the injustice at stake ultimately boils down to unjustly harming someone as a language user. Once we have a working theory of deliberative injustice, we can wonder whether resorting to the social level of analysis – i.e., appealing to social facts – helps further illuminate its nature. There are two steps in this direction. The first one is to follow Ayala's recently proposed *structural explanation* (Ayala, 2018) according to which discursive injustice in general –and by extension, deliberative injustice– calls for an explanation as an effect of cognitively outsourcing

² Adapted from Kukla (2014).

to unjust discursive norms. In the following, I present this proposal and apply it as a fitting explanation of deliberative injustice in some cases, though I also show that it has some limitations. The second step is to focus on the significance of deliberative injustice at the collective level, i.e., for the group taken as embarked on the joint endeavor of carrying out an exchange of reasons. In this regard, in the last section I will contend that deliberative injustice constitutes a collective failure of the group. To conclude, I will lay the ground for further work devoted to conceptualizing how groups can avert deliberative injustice.

The Limits of the Standard Approach to Discursive Injustice

As noted, discursive injustice is the kind of injustice we suffer when we are denied the possibility of using language in ways we are entitled to, notably when we intend to *do things with words*. The focus of the standard approach to discursive injustice has been to identify and characterize the discursive injustice consisting in the distortion of the illocutionary force of someone's speech act through the hearer's undue uptake when the speaker's use of the relevant linguistic conventions is correct. The underlying idea is that we are entitled as language users to make speech acts when we use the appropriate linguistic conventions. To elaborate, the reason for this is that linguistic interaction is a main vehicle of interaction with others, which, in turn, constitutes a great portion of what we do and who we are. Since producing speech acts has a place at the core of our linguistic interaction with others, it comes as no surprise that, at least as long as we use the appropriate linguistic conventions, we are entitled to produce speech acts and see that others respond to them appropriately. For, otherwise, we are severely harmed in our capacity to carry out our plans – say, to issue an order – and to situate ourselves in our social surroundings – say, to present ourselves as occupying the position of a boss.

Now, I want to point out that an important part of our linguistic interaction with others consists in exchanging reasons, that is, in offering others what we take to be good considerations in favor of acting or thinking in some way and receiving similar considerations in return. Indeed, it is difficult to overstate the role that embarking on such an activity occupies in our everyday lives. We engage in exchanges of reasons about what restaurant to go to, what movie to watch, about why one made some vital choice, about why the Parliament passed some law, about whether or not, and to what extent, assume a vegetarian way of life, etc. In doing so, we may intend to carry out a broad array of different actions, say, we may want to justify ourselves before others, to flag our opposition to others' plans and ideas, to engage in an intellectually stimulating interaction, or, in cases of group deliberation, to make a joint decision. Also, through exchanging reasons we exert an influence on how others perceive us and so situate ourselves in the social surroundings, say, we can present ourselves as having a preference towards some particular cuisine or as advocates of some public policy or way of life. Since the linguistic practice of exchanging reasons with others occupies such an important place in our lives, it is plausible to think that it is something we are entitled to do, at least as long as we engage in it in the appropriate circumstances. For example, we are surely entitled to share our reasons with others and to listen to what others have to say when we engage in a debate or in a group deliberation.

That being so, it is easy to see that one way in which we can be subject to discursive injustice consists in being prevented from effectively taking part in an exchange of reasons. Now, how are we to understand such a form of discursive injustice? At first sight, it might seem that the standard approach is well-provided to that end, as on many occasions we are prevented from effectively taking part in an exchange of reasons because of our interlocutors' undue response towards the speech act whereby we intended to engage in the exchange. To illustrate this possibility, here is an example featuring a group deliberation scenario:

Ignoring

Imagine a research group deliberating about whether or not to write a paper with the empirical results they have at the current stage of their inquiry, and that one of the group members makes a contribution to the debate speaking with a strong foreign accent. After she's done speaking, there is a moment of silence, and then the debate resumes as if nothing had happened. But later on, a different group member, who is a native speaker, raises the same point and then the group debates it vividly³.

Here, an injustice has been committed: due to prejudices against foreign speakers, the foreign researcher's contribution is dismissed when it clearly should be debated, as the group implicitly acknowledges when they debate the same point later on. What is more, it is a discursive injustice, since the foreign researcher has been prevented from doing with words something she was clearly entitled to, that is, participate in the inner deliberation of a research group she is a member of. And the standard approach can explain exactly how that injustice comes about, namely, by unduly responding to the foreign researcher's rightful attempt to make her words count as a contribution to the deliberation.

³ Adapted from Ayala (2018).

Nevertheless, there are other ways in which discursive injustices can be committed in reason exchange contexts that the standard approach cannot accommodate. To see this, consider the phenomenon of the unjust self-silencing of socially disadvantaged groups, such as women, in group deliberation contexts (Gerber, 2015; Hansen, 2010). As the empirical evidence suggests, women tend to participate less than men in deliberative processes due to factors such as whether the debate revolves around political issues or whether there is a low women-men ratio inside the group (Karpowitz, Mendelberg & Lee, 2012). Plausibly, this comes as a result of social dynamics transcending the deliberative process, such as the fact that women are subject to greater critical scrutiny than men when they speak their minds, and that topics such as politics tend to be perceived as typically masculine topics over which men have more authority – all of which puts an unjust burden on women (Karpowitz & Mendelberg, 2014). In this way, when the deliberation takes place in predominantly masculine environments, women tend to self-silencing, thus refusing to intervene, and deferring to men.

The unjust self-silencing of women in group deliberation contexts constitutes an example of a discursive injustice, for, because of some widespread sexist attitudes, women see their capacity to intervene in the linguistic practice of exchanging reasons in deliberative contexts severely constrained. Thus, there is an important dimension of their lives as language users –that which has to do with exchanging reasons with others in deliberative contexts– where they have less of an opportunity to participate than they should have. But the standard approach does not explain why this amounts to a discursive injustice. The reason is, quite simply, that in self-silencing cases there is no speech act involved to begin with, so there is no hindering of any speech act's illocutionary force at all. Indeed, in such cases, the victim suffers what might be the crudest form of injustice one can be subject to as a language user, namely: she does not even get to be a speaker. Since there is no speech, there is no speech act to be hindered through the hearer's undue uptake. Thus, there appears to be a form of discursive injustice that concerns reason exchange contexts that escapes the theoretical scope of the standard approach, and the contours of which are yet to be explored. In what follows, I will use the expression "deliberative injustice" to refer to such a kind of discursive injustice, as deliberation constitutes a paradigmatic example of a linguistic practice that has the exchange of reasons at its core.

Deliberative Injustice and Being a Source of Reasons.

In this section I will put forward a view of the discursive injustice at stake in reason exchange contexts that departs from the standard approach and that successfully accounts for why cases like the one presented in the former section constitute examples of discursive injustice. It is important to bear in mind that my purpose is not to provide an alternative theory of discursive injustice *in toto*, but the humbler one of understanding a particular form of discursive injustice that escapes the theoretical scope of the dominant approach. My proposal does not contradict, but complements, current theoretical trends and purports to contribute to widen and further illuminate the theoretical territory of discursive injustice.

The best way to present my proposal is to explain what makes the harm inflicted in the cases of discursive injustice I am most interested in, contrasting it with the sort of harm at stake in the cases of discursive injustice that the standard approach can account for. So, let us ask: what does the harm inflicted upon Elizabeth Bennet in *Rejecting* and Celia in *Commanding* consist in? Of course, a great deal of that harm is that Elizabeth Bennet's and Celia's lives become difficult as their interlocutors put them in unpleasant situations, such as needing to deny having flirted with Mr. Collins, or burden them with the unwelcome need to carry out further courses of action to guarantee the success of their actions, such as adopting a higher tone of voice and a rude attitude so that nobody can doubt that an order is being issued. Nevertheless, there is a deeper, more general and more relevant sense in which we must understand that Elizabeth Bennet and Celia are harmed, the sense that we are mainly interested in when we study the signification of discursive injustice for socially disadvantaged groups. The harm in question is that the victim is damaged in her status as an agent: when our interlocutor impedes or distorts our actions of rejecting someone or giving orders, he is not only making our life more difficult, but also damaging our capacity to relate to the world as beings capable of possessing and pursuing projects and courses of action. More precisely, the victims of the discursive injustices explained by the standard approach are damaged in a particular, though crucial aspect of their status as agents, namely, that which has to do with acting in a social space through linguistic interaction as owners of the sense of their own words.

It should be clearer now why the standard approach cannot account for the kind of discursive injustice at play in reasons exchange contexts: in contrast with what goes on for Elizabeth Bennet and Celia, the harm inflicted in those cases does not consist in the fact that the victim is deprived of control over the sense of her own words. For the self-silencing of socially disadvantaged groups in group deliberative contexts still counts as a discursive injustice even if, should they intervene in the deliberation, their words would be taken in the intended way. So, what is the nature of the harm at play in cases of discursive injustice committed in deliberative contexts – for short, *deliberative* injustice?

My thesis is that the way in which we are harmed as language users in deliberative injustice contexts, the way in which we are prevented from fruitfully interacting with others through language, is by being deprived

of the effective exercise of a particular normative position that we hold before the others, and which I label as being a source of reasons. As I understand it, being a source of reasons for others, in the sense envisaged, consists in occupying a normative position from which our linguistic performance generates, under the appropriate circumstances, normative constrictions for others in the form of reasons for thinking or acting in some way. For example, if we take part in a debate and put forward an objection against our interlocutor's beliefs, she is bound to acknowledge our objection as a reason to change her mind, revise her stance, or provide a fitting counterargument. By doing so, she would be acknowledging us as valid sources of reasons. To the contrary, we are not taken as sources of reasons if, for example, our interlocutor just shrugs his shoulders and keeps talking as if nothing had happened. In such a case, by refusing to respond, our interlocutor fails to behave in accordance with the normative constraints that our objection puts upon him, and thus hinders the effective exercise of the normative position we occupy before him as sources of reasons.

To elaborate, I want to point out that there are two general different ways of exercising discursive injustice. The first and most obvious one is by directly blocking the exercise of one's status as a source of reasons, like when we put forward an objection that is met with disregard. But there is a different way of failing to acknowledge someone as a source of reasons that is a little – though significantly – stronger, namely, that the *person* is not taken to occupy the kind of normative position in virtue of which her linguistic conduct imposes normative constraints on her interlocutors. It is not only to block the exercise of some right, but to *exclude someone as a bearer of some linguistic entitlements*, that constitutes the harm involved in such deliberative injustice cases. To illustrate the difference, consider this example:

Online debate

Ona, a woman, is debating in an online forum using a male avatar. She is putting forward well-informed, correct and convincing arguments, and she is thus achieving to change her interlocutor's mind. Nevertheless, her interlocutor happens to be such a recalcitrant misogynist that he would never take Ona's arguments seriously, but would immediately dismiss them, if he knew about Ona's gender.

In *Online debate* there is nothing in the attitude of Ona's interlocutor preventing her from doing such things as putting forward an objection and receiving a response. However, it is still clear to me that an injustice is being committed. For, as the case goes, Ona's interlocutor would cease to respond to Ona's objections if he knew that she is a woman and, in consequence, he is not genuinely accepting *her* as his interlocutor. His conduct is not constrained by the pressure Ona ought to exert upon him in virtue of being his interlocutor, but by his attitudes towards Ona's gender. Thus, the sexist dispositions of Ona's interlocutor prevent her from actually occupying the normative position that she, as a member of the linguistic community, is entitled to bear, thus unjustly depriving her of something she is owed. This is the sense in which, in the example, Ona is the victim of a deliberative injustice even though she is still capable of interacting in a seemingly appropriate fashion.

There are three theoretical advantages to my proposal. The first one is that it allows to account for some discursive injustice cases that escape the scope of the standard approach. Thus, we can explain the self-silencing of socially disadvantaged groups as a discursive injustice by pointing out that the members of such groups are under an unjust pressure not to intervene in group deliberation processes, which hinders their capacity to make the normative position of being sources of reasons effective –or, in any cost, only at an undue cost and effort.

The second advantage is that underlining that we are entitled as language speakers to being acknowledged as sources of reasons allows us to explain the normative structure of various practices that revolve around exchanging reasons. Let us consider deliberation as a paradigmatic example of a reason exchange practice. As Corredor (2020) has pointed out, by engaging in a deliberation we situate ourselves in a net of rights and obligations towards the other participants in the exchange, such as the obligation to respond to the objections we are presented with. My proposal can account for why that is the case, namely, because deliberation is a practice founded on the exchange of reasons, and so one where the participants are bound to acknowledge each other as sources of reasons. Thus, the rights and obligations structuring the practice of deliberation stem from the entitlements that we bear as language users that engage in exchanges of reasons with others.

The third and final advantage is that admitting deliberative injustice to the theoretical camp of discursive injustice opens the door to illuminating other common linguistic behaviors, which would otherwise go unnoticed, as constituting a discursive injustice. That is most clearly the case with behaviors oriented to hindering the fruitful exchange of reasons, such as when someone unjustifiably closes the exchange of reasons by stubbornly refusing to respond to his interlocutor's arguments or to change his mind in light of their strength, for instance, by claiming that he is, as it is common to put it, "entitled" to his opinion and that that is the end of it. Surely, manifesting this kind of stubbornness is a rude way to conduct oneself. But, in my view, this kind of behavior is improper at a deeper level, i.e., it constitutes a deliberative injustice. For, by closing the exchange to eschew responding to or changing his mind in accordance with the strength of the advanced arguments, he would be failing to acknowledge his interlocutor as occupying the normative position of being a source of reasons.

Deliberative Injustice and the Structural Explanation.

Once in possession of a fitting account of deliberative injustice we can proceed to examine it at a deeper level. Since deliberative injustice constitutes a social phenomenon, requiring the participation and interaction of several subjects, it is natural to begin our inquiry by considering how appealing to social facts can allow us to illuminate the nature of deliberative injustice. In the remainder of this paper I will explore two different ways of doing so. The first one—the one I will be considering in this section—consists in asking how the social level of analysis sheds light over the way in which deliberative injustice comes about. In that regard, I will resort to Ayala's structural explanation (Ayala, 2018), though I will point out that it has some important limitations. The second one consists in considering what deliberative injustice signifies for those involved in an exchange of reasons. In particular, I will argue that deliberative injustice constitutes a *collective* failure on the part of the group of participants in the exchange as a whole.

What causes deliberative injustice? We may distinguish two broad frameworks for approaching this issue: the individualist framework, where we isolate a set of facts relevant to the individual psychologies of speakers and hearers such that, when they are in place, they lead them to commit discursive injustices (Ayala & Vasilyeva, 2015) and the social framework, where the idea is that there are social facts different from what goes on in the mind of speakers and hearers that are causally responsible for their behavior. Let us begin by presenting the individualist framework. Normally, the relevant psychological facts identified by the adherents to this framework are different kinds of implicit biases towards socially disadvantaged groups. Indeed, the empirical evidence supports the existence of widespread biases that deeply affect the perception of members of socially disadvantaged groups, which systematically leads to situate them in an unfavored position as speakers (Fricker, 2007). In this way, if the individualist framework is right, the hearers commit discursive injustices due to perceiving the status of members of socially disadvantaged groups as speakers in a distorted way, and they act in consequence. For example, a common effect of such biases is to attribute women with less authority than they deserve (Karpowitz & Mendelberg, 2014). Thus, if we take the example of *Commanding*, we can explain why a discursive injustice is committed: in virtue of a sexist bias, the factory workers perceive Celia as having less authority than she deserves in virtue of her position as a manager, and, consequently, they incorrectly interpret her orders as mere requests.

As Ayala points out (Ayala, 2018), the individualist framework has two advantages: its ontological simplicity, since the existence of human individuals with individual psychologies is not a debating matter, and its explanatory clarity, since the causal chain that leads from possessing some bias to manifesting some conduct that incurs in a discursive injustice is perfectly clear. Nevertheless, Ayala (2018, p. 14) contends that there is a crucial aspect of discursive injustice that the individualist framework cannot account for, namely, our "online situatedness" as members of a society. The idea, in a nutshell, is this: typically, our situation in our social surroundings can be described in terms of a structure where we occupy the place of the nodes, and where the links between the nodes represent different ways we have to relate to one another. Thus, the connections between the nodes structure the social space by constituting the norms we are subject to—e.g., the ruling norms in a community about respect to personal space or about how to perform in a conversation—and affording the behavior we are expected to show before one another in different situations—e.g., in different conversational contexts. Importantly, we typically *rely* on our online situatedness in order to navigate our social surroundings, that is, we directly carry out the acts that are afforded by the ruling norms that structure the social space—for instance, when we are at a party, we immediately, pre-reflectively resort to the ruling conventions about how to behave towards another person's personal space. Accordingly, one way in which discursive injustice can be brought about is when the social space is structured by linguistic norms that afford discursively unjust linguistic behaviors, so that, by relying on them, the speakers commit discursive injustices. Ayala's critique on the individualist framework is that, since it focuses exclusively on implicit biases, it overlooks how common it is that discursive injustice is brought about not because the speakers have implicit biases in their psychology, but because they rely on ruling conversational norms that are unjust, e.g., conversational norms dictating taking a woman's orders as requests or, or her rejection as a display of coquetry.

As an alternative to the individualist framework, Ayala offers her structural explanation. The first step is to notice that, as contemporary cognitive science and epistemology has it, cognitive processes do not take place uniquely within the confines of our heads, but they are often *outsourced*, i.e., partly realized by external factors on which we cognitively rely (Clark & Chalmers, 1998). Artifacts provide the paradigm of cognitive outsourcing, such as GPS systems on which we rely to navigate an unknown city, the notebook where we write phone numbers, or even the pen we use to carry out a mathematical operation. But recent works have convincingly defended that we can also cognitively outsource to social factors, such as another person with whom we stand on a certain relationship, e.g., a spouse on whom we always rely at storing beliefs about our experiences together (Gallagher, 2013; Menary, 2013; Krueger, 2011). Institutions, such as the norms that regulate social practices, are also among the social factors that can operate as external cognitive resources, and the second step in Ayala's argumentation is that, in discursive contexts, we typically outsource to the discursive norms that regulate the conversation. That is, instead of embarking on internal deliberations to

decide what to do, we simply rely on being guided by the conventions of the practice. Just as we immediately rely on the guidance of GPS systems without paying it too much of a thought, so we rely on the guidance of conversational norms to conduct ourselves in a conversation. In a final step, Ayala argues that unjust discursive behavior results from cognitively relying on, and thus acting in accordance to, the unjust discursive norms that articulate unjust social structures. In other words, what explains discursive injustice is the existence of certain unjust norms that govern our behavior in a conversation, such as the norm to interpret a woman's rejection as a mere display of coquetry. In this way, discursive injustice does not result from the hearer's incompetence at carrying out a discursive enterprise, but from its competence at abiding by discursive norms that structure the social space in an unjust way. This is, in essence, the structural explanation of discursive injustice.

Now, it is easy to see how the structural explanation, as an account of discursive injustice in general, can account for deliberative injustice in particular: it would be sufficient to identify the ruling norms about exchanging reasons in a given linguistic community such that, by relying on them, the subjects involved in an exchange of reasons find themselves committing deliberative injustices. For example, if there is a norm of such a kind to the effect that men are to be given prevalence when discussing about politics, women that rely on it would be prone to self-silencing in contexts where the exchange of reasons concerns political issues –which, as we saw, constitutes a deliberative injustice, for, in such cases, women are not effectively acknowledged as sources of reasons.

Nevertheless, I also want to point out that Ayala's structural explanation has some limitations, which should give us pause before dismissing the individualist framework *in toto* as a possible alternative explanation of deliberative injustice and, more generally, of discursive injustice. To be sure, Ayala (2018, p. 16) explicitly acknowledges that the structural explanation may not replace the individualist framework but complement it but she does not deepen into the limitations of her proposal. By making such limitations clear, I aim to make progress in the way of clarifying the respective theoretical scope of each of the two different approaches. There are two such limitations that I want to point out. First, *appealing to the hearers' implicit biases can be sufficiently explanatory in itself*, without the need to consider whether the hearers are also abiding by unjust discursive norms or not. Indeed, what the empirical evidence shows is that there is a strong relation between having some biases and manifesting an unjust discursive conduct. For instance, the evidence shows that, when a partaker in a conversation expresses herself in a non-native accent, the native speakers tend to perceive her in such a way that they attribute her less authority than deserved (Boyd, 2003; Lindemann, 2003; Lippi-Green, 1997). Thus, the individualist framework suffices to explain some typical cases of discursive injustice, which makes it unnecessary to seek for a structural explanation in such cases.

And second, I suggest that *the structural explanation does not possess sufficient explanatory power in a broad range of cases*. The simplest of them is the case in which a member of a socially disadvantaged group commits a discursive injustice against a member of a socially advantaged group. In such a case, a discursive injustice would be in place even if no unjust discursive norms are being followed, so the structural explanation is unable to explain where the injustice comes from. For instance, we can imagine an alternative scenario of *Rejecting* with the tables turned, where it is Mr. Collins who fails to make his words count as a rejection of Elizabeth Bennet's marriage proposal due to her unduly disregarding his speech acts as such rejection. I think it is clear that, in such an example, Mr. Collins would be wronged in his status as a language user, as he would be prevented from making his words count in the way he wants due to the hearer's inappropriate uptake. Of course, discursive injustices committed against socially disadvantaged groups are far more serious, as they stem from unjust prejudices and reproduce unjust social structures. But it is still the case that a discursive injustice has been committed in our alternative scenario, one that the structural explanation has trouble accounting for. A different kind of case, perhaps more difficult to imagine, is that where the discursive norms regulating conversational practices are not unjust – because, say, we have achieved social justice and have gotten rid of all the structures that build social advantages and disadvantages. However, in such a world it would still be possible to corrupt the illocutionary force of someone's speech act: there is simply nothing to prevent people from, for example, unduly interpreting someone's orders as requests because of a *personal* prejudice against her. Since the discursive norms in such a socially just world would not be unjust, it is difficult for the structural explanation to explain what would be causing the injustice.

Deliberative Injustice as Collective Failure

In this final section, I will explore an alternative way of appealing to the social level of analysis to shed light upon the nature of discursive injustice as a general phenomenon. To that end, I will be focusing on deliberative injustice, leaving open the possibility to expand the same account to other forms of discursive injustice. My proposal is that we consider the *significance* of deliberative injustice for those that are involved in the joint enterprise of carrying out an exchange of reasons, and who, accordingly, constitute a *collective agent*, in the sense of a group of subjects that come together to carry out a common enterprise. To make things simpler, I will once again focus on group deliberation as the paradigmatic example of a practice that revolves around an

exchange of reasons. Thus, what I suggest is that we attend to what committing a deliberative injustice involves for deliberation taken as a collective practice, i.e., as a kind of enterprise that we carry out by joining efforts and coordinating with others. Let us focus on what deliberative injustice constitutes at the collective level.

My thesis in this regard is that *deliberative injustice constitutes a collective failure*, i.e., a failure of those involved in a deliberation *qua* deliberative group. Now, of course, there is a sense in which saying that deliberative injustice is a failure is no news. Since not being acknowledged as a source of reasons plausibly amounts to a moral harm, inflicting a deliberative injustice or, in any case, failing to prevent it or to amend it involves a failure at fulfilling some moral standard. But there is a different, more interesting sense in which deliberative injustice constitutes a failure, namely, the sense in which it is a failure of the deliberation itself - or, to put it otherwise, a failure of the deliberative group *qua* deliberative group. For the success of collective endeavors in general requires not only bringing about the goal we set out to achieve but achieving it *together*. This, in turn, requires that those partaking in the joint effort abide by the rules of the collective practice in question. Since deliberative injustice involves ignoring the central deliberative rule of acknowledging the other members of the deliberative group as sources of reasons, it follows that it involves the failure of the deliberation as a collective endeavor, or, in other words, it involves the failure of the deliberative group.

Allow me to illustrate with an example. Say that you and I take part in a team sport, such as a rowing race, and say that we win the race. Certainly, we have achieved the goal we were seeking in our common effort, namely, winning the race. But imagine that we win only because you have rowed vigorously, while I have done nothing but relax and enjoy the trip. Have we succeeded in our joint endeavor? It seems to me that we have not, even though we may have won the race. To the contrary, I believe that the verdict that comes to mind is that it has been a failure. The reason for this is that, since rowing is a collective practice – something we carry out through joining others in a common effort – succeeding at rowing involves not only achieving the goal in question but achieving it *through* our common effort and coordination. And, more precisely, the way in which such a common effort should be expressed in the practice of rowing is through abiding by the rule that we both row as hard as we can. So, yes, the race may have been won, but there is an important sense in which *we* have not won it: the sense in which winning the race has not been the result of a common effort. That being so, our rowing, *qua* joint endeavor, has been a failure. In order for it to amount to a genuine success, I need to do what is expected of me as a partaker in the activity of rowing and cooperate properly.

Let us return to deliberation. As a collective enterprise, the success of deliberation involves not only fulfilling some purported goal – such as arriving at a consensus, or true belief, or what have you – but doing so through the proper cooperation of those that engage in it. Now, as we know, deliberation is a collective practice founded on the exchange of reasons, and so proper cooperation in a deliberation requires showing some form of respect to one another and, in particular, the kind of respect that is expressed in treating each other as sources of reasons. Indeed, as we saw above, treating each other as sources of reasons occupies a central, structural place within the rules of deliberation, as rules such as the rule to respond to the objections that one's interlocutor may raise can be seen as stemming from the more fundamental requirement to acknowledge your interlocutor as someone who occupies the normative position of being a source of reasons. That being so, we can now appreciate the full significance of deliberative injustice: when someone commits a deliberative injustice, he does not only incur a moral harm towards others, but also fails to properly cooperate as a partaker in an exchange of reasons. In consequence, deliberative injustice involves the failure of the exchange of reasons as a joint enterprise, or, more exactly, it involves the failure of the group as a group of people engaged in the collective enterprise of carrying out an exchange of reasons.

Allow me to elaborate on the collective significance of deliberative injustice. I have argued that deliberative injustice constitutes a failure of the deliberative group because it impedes the full success of the practice that such a group is engaged in, as succeeding in a deliberation – and, indeed, in any practice that revolves around exchanging reasons – involves that those that partake in it treat each other as sources of reasons. But there is a second sense in which injustice constitutes a collective failure, namely, that the deliberative group is *responsible* for such a failure. I do not mean to say that the deliberative group, taken as a whole, is always *causally* responsible for giving rise to deliberative injustices. That may be the case on many occasions, but we need to make room for cases where it is simply one group member, acting on his own, who commits the injustice. The thesis in question is weaker than that, namely that, since deliberative groups are committed to the success of their deliberation, they are accordingly committed to ensuring that the deliberation is not thwarted by deliberative injustice. That being so, when deliberative injustice occurs, it is not only the individual group members causally responsible for it that are accountable, but the deliberative group as a whole.

At this point, further theoretical work must be devoted to conceptualizing the different ways in which deliberative groups can be up to the responsibility of countering deliberative injustice. Though I do not have the space to develop a full proposal, I want to advance the following sketch: first, we can distinguish two different goals such that, by fulfilling them, deliberative justice is effectively countered, and second, we can accordingly distinguish two different kinds of collective abilities such that, by exercising them, deliberative groups achieve to fulfill those goals. The first of the goals in question is the *prevention* of deliberative injustice, understood as giving rise to the conditions under which deliberative injustice does not occur. For example, deliberative

groups can institute specific inner rules and institutions to that effect. The second goal is the *emendation* of deliberative injustice wherever it takes place, which can be fulfilled through, for example, reproving whoever interrupts and inviting the aggrieved group member to repeat her contribution. As for the relevant abilities, we can expect deliberative groups capable of fulfilling those goals to manifest both practical abilities, related to acting to prevent and emend deliberative injustice, and epistemic abilities, related to the identification of the different forms that deliberative injustice can take and of the different ways to prevent and emend it. In this regard, those deliberative groups capable of countering deliberative injustice must be sensitive both to the general social context –so that they are sensitive to, for example, the difficult situation of members of socially disadvantaged group in deliberative contexts– and to the social context inner to the group itself –so that they are sensitive to, for example, the possible conflictive personal relationships among the group members. This concludes my proposal.

Conclusions

I have argued for distinguishing a particular form of discursive injustice, *deliberative* injustice, which we are subject to when we are prevented from making effective a normative position that we are entitled to hold before others as language users, namely, the position of being a ‘source of reasons’ for others in reason exchange contexts. I have shown that deliberative injustice cannot be accounted for in the terms of the standard approach to discursive injustice, as there are cases where we are subject to deliberative injustice without thereby seeing the illocutionary force of our speech acts corrupted in any way. Accordingly, deliberative injustice constitutes a distinct kind of discursive injustice, different in nature from the other types of discursive injustice considered in the existing literature.

Moreover, I have also considered two different ways in which resorting to the social level of analysis is important for illuminating the phenomenon of deliberative injustice. First, I have considered Ayala’s structural explanation, according to which discursive injustice in general comes as an effect of hearers cognitively outsourcing to and abiding by socially unjust discursive norms. As I have shown, we can appeal to the structural explanation to explain how deliberative injustice comes about in many cases; nevertheless, the structural explanation also has some limitations, and there will be other cases where we should not resort to it but to the more standard approach of considering the speakers’ implicit biases. Second, and finally, I have suggested attending to the significance of deliberative injustice at the collective level, i.e., for the group that exchanges reasons taken as embarked on a collective endeavor. In this regard, I have argued that deliberative injustice constitutes a *collective failure* for the group *qua* group.

On a final note, I have laid the ground for further work regarding how groups can counter deliberative injustice. In particular, I have suggested that they can do so by developing and exercising collective epistemic and practical abilities related to identifying the conditions under which deliberative injustices are committed and to actively countering them by fulfilling two goals, namely, preventing and emending them.

Bibliographic references

- Ayala, Saray (2018). A structural explanation of injustice in conversations: it’s about norms. *Pacific Philosophical Quarterly* 99(4), 726-748. <https://doi.org/10.1111/papq.12244>
- Ayala, Saray & Vasilyeva, Nadya (2015). Explaining speech injustice: individualistic vs. structural explanation. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, and P. P. Maglio (eds) *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Bianchi, Claudia (2020). Discursive injustice: the role of uptake. *Topoi* 40(1), 181-190. <https://doi.org/10.1007/s11245-020-09699-x>
- Boyd, Sally. (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, 283-295. <https://doi.org/10.1080/13670050308667786>
- Clark, Andy and Chalmers, David (1998). The extended mind. *Analysis* 58(1), 7-19. <https://doi.org/10.1093/analys/58.1.7>
- Corredor, Cristina (2020). Deliberative speech acts: an interactional approach. *Language & Communication*, 71, 136-148. <https://doi.org/10.1016/j.langcom.2020.01.005>
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>
- Gallagher, Shaun (2013). The socially extended mind. *Cognitive Systems Research* 25 26, 4–12. <https://doi.org/10.1016/j.cogsys.2013.03.008>
- Gerber, Marlène (2015). Equal partners in dialogue? Participation equality in a transnational Deliberative Poll (Europolis). *Political Studies*, 63(1), 110-130. <https://doi.org/10.1111/1467-9248.12183>

- Hansen, Kasper Moller (2010). The equality paradox of deliberative democracy: a review of theoretical concepts and empirical findings. In M. R. Wolf, L. Morales, K. Ikeda (Eds.) *Political Discussions in Modern Democracies* (pp.26-43). Routledge.
- Karpowitz, Christopher, & Mendelberg, Tali. (2014). *The silent sex: Gender, deliberation, and institutions*. Princeton University. <https://doi.org/10.1515/9781400852697>
- Karpowitz, Christopher, Mendelberg, Tali, & Shaker, Lee (2012). Gender inequality in deliberative participation. *American Political Science Review* 106(3), 533-547. <https://doi.org/10.1017/S0003055412000329>
- Krueger, Joel. (2011). Extended cognition and the space of social interaction. *Consciousness and Cognition* 20(3), 643–657. <https://doi.org/10.1016/j.concog.2010.09.022>
- Kukla, Rebecca (2014). Performative force, convention, and discursive injustice. *Hypatia* 29(2), 440–457. <https://doi.org/10.1111/j.1527-2001.2012.01316.x>
- Langton, Rae (1993). Speech acts and unspeakable acts. *Philosophy & Public Affairs* 22, pp. 293–330. [10.1093/acprof:oso/9780199247066.003.0002](https://doi.org/10.1093/acprof:oso/9780199247066.003.0002)
- Lindemann, Stephanie (2003). Koreans, Chinese or Indians? Attitudes and ideologies about non-native English speakers in the United States. *Journal of Sociolinguistics* 7, 348-364. <https://doi.org/10.1111/1467-9481.00228>
- Lippi-Green, Rosina (1997). *English with an accent: language, ideology, and discrimination in the United States*. Routledge.
- Menary, Richard (2013). Cognitive integration, enculturated cognition, and the socially extended mind. *Cognitive Systems Research* 25–26, 26–34. <https://doi.org/10.1016/j.cogsys.2013.05.002>