Propositional attitudes, harm and public hate speech situations: towards a maieutic approach

Corrado Fumagalli

1. Introduction

Here I repeat a well-worn question: if public hate speech is harmful, what is the relevant wound? Some of the most sophisticated philosophical arguments say that public hate-speech events are harmful because they cause a discrete, traceable, and harmful change in one's propositional attitudes (e.g., Delgado 2018; Langton 2012; Matsuda 2018; Seglow 2016). I am sceptical that speech can harm, but, at least for the first part of my argument, I shall proceed as if such harm was possible. So the question is: do hate-speech events cause a traceable and discrete change in one's propositional attitudes? In this paper, I claim it is very difficult to identify a traceable and discrete change in propositional attitudes when speakers and their audience share the same common ground of propositional attitudes about the target group, which is, more or less explicitly, assumed as a necessary condition for public hate speech to accomplish something (Delgado 2018; Langton 2012; Matsuda 2018; Lawrence III 2018). As an alternative, I offer a more realistic proposal: a maieutic approach to hate speech. From this perspective, public hate-speech events do not cause changes in propositional attitudes, but rather, if successful, either such events bring a person's latent propositional attitudes into clear consciousness, or they play with propositional attitudes speakers and their audience had prior to the public hate-speech situation. Both possibilities, however, oppose the thesis that, by looking at propositional attitudes, there is sufficient grounding for identifying the harm of hate speech.

This article studies the relationship between changes in propositional attitudes and the possible effect of public hate speech. I begin by making clear the assumptions of this paper. In section 3, I construct a theoretical framework to study public hate-speech situations. In section 4, I offer a

critical discussion of prevailing philosophical arguments for the harm of hate speech. In section 5, I provide my alternative proposal: a maieutic approach to public hate speech. Section 6 concludes the paper.[1]

2. Caveats

In this paper, I maintain the distinction between speech in the ordinary sense and speech in the technical sense (Schauer 1979; Maitra and McGowan 2010). Here, "speech" is not only a vocal address delivered to a more or less wide audience. I understand the word in the technical sense so that some actions that are not speech in the ordinary sense (such as publishing a pamphlet, drawing murals, writing libels) may fall within the scope of the paper (Maitra and McGowan 2010). Moreover, my focus is on speech that is public—that is, speech and actions that address an audience and that may have a much broader resonance. More specifically, my argument is on public speech that expresses hate. The definition of "hate speech" is certainly a disputed matter, and setting the conceptual perimeter in one or another direction may imply different normative positions (Maitra and McGowan 2010; Yong 2011). By "public hate speech," I refer to actions and discourses that simultaneously deny the basic standing of individuals who belong to a certain target group (Brettschneider 2012; Delgado 2018; Lepoutre 2017; Matsuda 2018; Seglow 2016; Waldron 2014) and aim to have an effect on and beyond a certain audience (Delgado 2018; Maitra and McGowan 2012; Langton 2012; UN 1965).[2] I follow most of the literature (e.g., Brettschneider 2012; Heinze 2016; Lepoutre 2017; Waldron 2014) in situating my argument within relatively stable liberal democracies.[3] Then, in developing my argument, I employ the canonical usage of terms such as propositional attitudes and propositional contents (Cresswell 1985; Fodor 1978). Propositional attitudes are internal mental states that have a certain attitude mode, such as believing, and a traceable semantic content.[4] Propositional content is such a content. Notably, many different agents can have propositional attitudes with the same

propositional content, and the same agent may have different propositional attitudes simultaneously.

3. The architecture of public hate-speech situations

Let me begin this section with some examples. Italian politician Attilio Fontana called on his audience to defend the "white race": "We have to decide if our ethnicity, if our white race, if our society continues to exist or if it will be wiped out."

In February 2004, French humorist Dieudonne M'bala M'bala said: "'Dirty nigger, the Jews will have your skin,' I've heard this kind of slogans. They are all slave traders turned bankers, [converted to] the show-business and today to terrorist action they show supporting the policies of Ariel Sharon. Those who attack me have founded empires and fortunes on the slave trade and slavery."

Valérie Bemeriki, a presenter on the Rwandan radio station Radio Télévision Libre des Mille Collines at the time of the Rwandan genocide, used to deliver speech of the following kind: "They [the Tutsis] are all Inyenzi [cockroaches]. When our armed forces will get there, they will get what they deserve."[5]

These three examples are generally understood as public hate-speech events. In this section, my ambition is to extend the analytical perspective beyond hateful content in order to capture the five fundamental elements (*speaker* [S], *target* [T], *tolerant* [t], *intolerant* [i], *common ground* [CG]) of each state of affairs in which public hate speech unfolds. I call such states of affairs *public hate-speech situations*.[6] If we read the examples in continuity, we see that, like all the many other examples of intersubjective human communication, in public hate-speech situations there is someone or something that speaks, gives audible expressions, expresses by written or printed words, draws cartoons and manifestos, makes something publicly known, publishes something, or, in very general terms, puts into circulation a continuous piece of speech with a propositional

content. I shall call such an entity *the speaker*. In cases such as written or printed words and cartoons or manifestos, the identity of the author might be initially unknown. In such circumstances, the identity of the author does not necessarily disappear, and it remains traceable throughout production and distribution. In most cases, however, the author of the message is recognizable, and he or she is a sentient being. People sign cartoons, manifestos, articles, books, and the like. People also speak, make noise, and use body language or gestures that are distinctively their own. It is not necessarily the case that the speaker is a single individual. For instance, political and cultural groups sign libels, manifestos, pamphlets, and books. Here members of a group construct themselves as a single entity, which counts as the speaker. Generally speaking, however, it is the entity that delivers a certain propositional content that counts as the speaker in public hate-speech situations. So, when a representative speaks on behalf of his or her political party, or when someone claims to represent the interests of a third party, the speaker is not the party as a whole or the union between the represented and the representative. The speaker is the single entity that conveys the propositional content.

In all public hate-speech situations, there has to be an identifiable target group, a group to which the speaker does not belong. *Targets* are the second unit of analysis of public hate-speech situations. To my knowledge, there is no account in the literature that denies the link between hate speech and the presence of references to target groups in the propositional content of a speaker's communicative act (e.g., Brettschneider 2012: 1; Langton 2012; Lawrence III et al. 2018: 1-3; Lepoutre 2017; Matsuda 1989: 2322-23; McGowan 2009; UN 1965; Waldron 2014: 27-8). Historically, target groups have been of very different kinds. It is plausible to say that, over time, possible targets of public hate speech may change together with social, political, and economic transformations. It is also reasonable to say that, in present times, possible targets of public hate speech may change together with worldviews speakers believe the targets have. So the distinctive perspective of the speaker might be very important in the construction of a target group. Let me explain. Prior to the communicative act, a speaker constructs an ideal-typical target group by associating one after the other relevant individuals because of something she takes to be a particularly prominent bad feature. In this way, the speaker can take members of

the group as synecdoche for the whole group or she can take the whole group to represent each of the members.[7]

There is at least one other way to constitute target groups. A certain set of individuals can think of themselves as members of the same group for a number of different reasons, such as common past, political ideas, cultural ties, physical traits, and the like. Such a group is not by definition a target group. It becomes a target group when a speaker constitutes it as such. In doing so, the speaker may draw upon traits that members of the group find communal, but also the speaker can reconstitute the group around an allegedly bad feature members of the group do not even identify with. For instance, Bemeriki may think all members of the Tutsi group look like cockroaches. If such a bad feature informs the speaker's propositional content, the group becomes a target. The construction of a target group is necessary but not sufficient for a public hate-speech situation to occur. Target groups impact the architecture of public hate-speech situations in two ways. First, it should be clear that target groups inform what the speaker aims to communicate. Specifically, in all public hate-speech situations, the construction of a target group influences the propositional content speakers aim to convey. Fontana constructs a variegated group of ethnicities as a menace to "whiteness." Dieudonne constructs Jews as robbers. Bemeriki stresses the cockroach metaphor to construct Tutsis as an inferior group.

Second, when we limit the perspective to speakers and targets, we offer too simplistic an analysis of public hate-speech situations. In the architecture of public hate-speech situations, third parties (neither speakers nor targets) also play an important role. I think we can identify two main groups of third parties: parties who share the speaker's beliefs about the target group; and parties who disagree with the speaker. Within such groups, there might be important variations of degrees, but these differences do not affect my analysis in any relevant way. In the following, for the sake of simplicity, I shall call one group *the tolerant* (the group of parties who disagree with the speaker). *The intolerant* is the group of parties who share more with the speaker. For instance, Bemeriki spoke to Hutu extremists, Dieudonne addressed a sympathetic public, and Fontana aimed to mobilize his voters.[8]

In the study of everyday public hate-speech situations, it is not always easy to identify the two kinds of audience. From an external point of view, for each public hate-speech situation, we can safely make the distinction between the intolerant and the tolerant only ex post, once speakers have conveyed their messages. At that stage, those parties who have received the propositional content with evident approval are recognizably intolerant. But it is by no means true that only those individuals who are recognizably intolerant are parts of the intolerant group; it might be difficult to detect an individual's intolerance. Speakers might address the wider public without explicitly galvanizing a group of sympathetic speakers (who might be electrified already). Speakers might presuppose they will have success with a limited number of third parties, and, eventually, they excite many more people than expected (or surprisingly fewer). Whether easily determinable or not, the constitution of an audience is quintessential to all public hate-speech situations, and it may affect the decision of conveying a certain propositional content rather than all possible alternatives. On this view, when speakers deliver offensive communicative acts in public, it is plausible to say that, from their own perspective, there is a reasonable expectation that someone will receive the propositional content sympathetically. At the same time, someone will react with disdain to the public hate-speech event. Someone may resort to public acts to stress how differently she believes. These acts make such people recognizable as tolerant.

Speakers, targets, the tolerant, and the intolerant can be combined to compose a framework that, however, is still partial. To have a proper analytical setup, we need also to consider public hate-speech situations for what they are: practices of speech that presuppose *certain things* as preconditions for making propositional contents of the speaker acceptable to the hearers. Within this context, it is plausible to affirm that when one delivers a hateful propositional content in public, she expects someone to adopt the hateful content. We should therefore observe that, in public hate-speech situations, a *common ground* between the speakers and at least one of the possible hearers (third parties and targets), which motivates such expectation, should be traceable. The expression "common ground" is by no means new. Over the years, a great deal of scholarship has studied the conditions for successful communication (Austin 1962) and the meaning of presuppositions (e.g., Grice 1957; Lewis 1979; Stalnaker 2002). Moreover, Rae

Langton has brought the topic to contemporary disputes about hate speech (2012). Presuppositions, as Robert Stalnaker writes, guide what speakers "choose to say and how they intend what they say to be interpreted" (2002: 701). The speaker presupposes a certain thing, as Stalnaker continues, only if she presupposes that recipients make the same presupposition (2002: 701-02). In this way, speakers and recipients act as if there was a "common ground" (Stalnaker 2002: 703-05) upon which a successful communication can be built.

For Langton (2012), scholars such as Stalnaker (2002), Grice (1967) and, possibly, Austin (1962) think about common ground as an abstract structure (such as mutually shared information, boundaries between permissible and impermissible sets of actions, a body of shared beliefs) that supports conversations. David Lewis (1979) thinks of common grounds as the attitudes of parties during the conversation: attitudes speakers can exploit, attitudes that can evolve to make sense of what is going on. In this paper, I try to keep the two perspectives together. If taken separately, each of the camps would lose explanatory power within the discourse on hate speech. For instance, by taking a too abstract account of common ground, we would neglect the fact that public hate-speech situations are not necessarily improvements in the knowledge one has. Langton is indeed right when she says that to capture the complexity of communication in today's societies, notions such as common ground and accommodation might be extended beyond beliefs. Specifically, it is generally understood that in public hate-speech situations, the principal reason for speech is not to get people to develop new true beliefs about something, but rather to make them feel something, to take in attitudes, to act (e.g., Delgado 2018; Langton 2012; Lawrence III et al.; Matsuda 1989; McGowan 2009; UN 1965; Waldron 2014). So, following Langton (2012), someone might be inclined to take a strong attitudinal perspective on common ground in public hate-speech situations. This seems more appropriate. However, with too much stress on the interacting element of public hate-speech situations, we would begin with an idea of public hate-speech situations that does not correspond to reality. In public hate-speech situations, the interaction between speakers and hearers is minimal.[9] Mostly, speakers launch hate speech because of prior presuppositions about the hearers without necessarily imagining an exchange with them. For these reasons, in the following, I shall hold a slightly different

conception of common ground. There is, I think, a common ground when the speaker and the hearer have the same stance on possible states of affairs about the target group so that both parties would be willing to ground the next steps of their conversation on such a common presupposition. By reframing the discourse in terms of propositional attitudes about the target group, I extend the discourse on common ground beyond beliefs. At the same time, this conception of common ground leaves open the possibility that such communication can continue in the future, but without postulating that this is necessarily the case. Speakers and hearers can hold propositional attitudes with different degrees of commitment. For instance, Bill may be a committed racist. Susan may consciously defend racial equality, but, at the same time, she may also harbor unconscious biases against blacks. And Peter, living without too much questioning, may just hold bad propositional attitudes as a result of his social environment. A further qualification is therefore in order: a common ground exists when the speaker and the hearer have the same stance on a possible state of affairs, and they hold such a propositional attitude *to the extent to which it is not unreasonable to expect successful uptake*.

To sum up, I think all public hate-speech situations are the result of the interplay of five elements: speakers, targets, the tolerant, the intolerant, and common ground. Each of these elements participates in such situations in a distinctive way:

- *Speakers* are individuals who convey the propositional content.
- *Targets* inform speakers' propositional content and may be direct or indirect addressees.
- *The tolerant* listen to speakers without acceptance.
- *The intolerant* listen to speakers with acceptance*.
- *Common ground* enables successful communication between speakers and at least one of the other entities.

With these elements in mind, it is possible to study the harm (if any) that a certain instance of hate speech can cause to those listening, reading, or seeing it. This is what I shall do in the next sections.

4. The harms of public hate-speech events

In this section, I shall consider the two main philosophical arguments for the harm of public hate speech: public hate speech is harmful because it has effects on targets; and public hate speech is harmful because it gives third parties reasons to act against targets. In different ways, the two arguments try to demonstrate that public hate-speech events cause a discrete and traceable change in one's propositional attitudes (Langton 2012; Lawrence III 2018: 77; Matsuda et al. 2018: 15; Seglow 2016: 9).

a. The direct harm of public hate speech events

One of the two main argumentative lines about the harm of public hate speech says public hate speech directly attacks targets. In so doing, public hate speech makes victims have bad propositional attitudes about themselves, their groups, and the society as a whole. The analytical focus is on the so-called violence of the "words that wound" in order to show that public hate speech has a real, immediate, and negative effect on the victims' stance on the possible state of affairs about the target group to which they belong (Delgado 2018; Matsuda 2018; Seglow 2016; Waldron 2014).

For instance, Jonathan Seglow (2016) claims that public hate speech assaults our beliefs about self-respect (2016).[10] According to Jeremy Waldron (2014), public hate speech alters the mental state held by victims about their dignity as equal members of the political community.[11] Mary Matsuda and Richard Delgado, I think, make two sorts of claim. First, by drawing upon sociology and social psychology, they argue that public hate speech causes psychological trauma, physiological symptoms, and emotional distress "ranging from fear in the gut, rapid pulse rate and difficulty in breathing, nightmares, post-traumatic stress disorder, hypertension, psychosis, and suicide" (Matsuda 1989: 2336). I am not sure public hate speech causes all these things, but

this is an empirical claim on causality that is going to stir up debates for many more years to come, and a resolution of this disagreement is beyond the scope of this paper. Second, Matsuda and Delgado argue that targets come to feel ambivalent "about their self-worth and identity" (e.g., Delgado 2018: 91) and their personal freedom (Matsuda 2018: 24). These ideas are more relevant to my argument.

So public hate speech assaults, degrades, and damages targets in a way that makes them have different propositional attitudes towards propositions about themselves. Degrading caricatures, threats of violence, posters, signposts, literature portraying target groups in demeaning ways, leaflets, public verbal abuses, fliers advocating lynching, and graffiti have a propositional content that addresses *directly* (at least) one of the target's propositional attitudes. Specifically, as Delgado (2018) and Matsuda (2018) write, the direct and immediate negative effect of public hate speech is connected with the fact that targets already live within a demeaning context in which harassment and open and covert violence are widespread and common. In these circumstances, public hate speech inflicts harm that is "neither random nor isolated" (Delgado and Stefancic 2009: 368).

Let me try to read this argument through the five elements I presented in section 2. Speakers are individuals who convey propositional contents about a target group with the intent of causing direct harm. Speakers and intolerant people are collapsed into a single entity that expresses hate. Targets are directly affected pre-existing groups of people who recognize each other as members of the group. Tolerant people are generally left outside the main picture.[12] Such an entity speaks to targets that absorb the hateful propositional content to the point that they have harmful propositional attitudes about themselves. This is so because there is a common ground between speakers and targets.

Members of target groups, as Matsuda writes, know that stigmatizing and demeaning comments are commonplace and socially acceptable (Matsuda 2018: 48). It is on this pre-existing common ground that speakers cause harm by making targets have bad propositional attitudes about themselves. Such bad propositional attitudes are the distinctive harm of hate speech. In reverse,

for hate speech to be harmful, it has to bring about a distinctively new propositional attitude built upon shared stances on possible states of affairs about the target group. As Delgado and Matsuda argue, hate speech is harmful because prior discrete hateful speech acts rendered people in the target group vulnerable (Matsuda 2018; Delgado 2018; Delgado and Stefancic 2009). Seglow is also explicit. Speakers and writers, he says, rely upon their targets' comprehending their hateful views, "else their speech would not have its intended effects" (2016: 10). The relation between harm and common ground is sensible, but it also shows a limit of this line of argument. If we stick with the Austinian idea that saying something sometimes constitutes doing something (as the frequent appeal to the expression "Words can wound" suggests), and, at the same time, we maintain that successful communication is built upon a common ground (as many applications of Austin's theory of speech suggest), it is very difficult to prove that a specific public hate-speech act has caused a new propositional attitude in the addressees.

I recognize that, at this point of the argument, my reading of the harm of hate speech as a change in propositional attitude may sound too narrow. Literature on harm (e.g., Maitra and McGowan 2012; Matsuda 2018; McGowan 2012) and debates about racism as an ideology (e.g., Garcia 1996; Hasslanger 2017; Shelby 2002, 2003) seek to demonstrate that hate speech is one among many social behaviors that occur within the context of an already-oppressive society. As Maitra and McGowan say, harm need not be a localized phenomenon, but "harm can be due to a series of act[s] none of which is individually harmful" (2012: 23). On such a view, it does not matter whether it is hard to parcel out the distinctive contribution of public hate-speech events to the evolution of one's bad propositional attitudes about herself. For members of the target groups, public hate-speech events are harmful because they maintain and reproduce a social reality informed by prejudicial behavior as well as demeaning practices and norms.

I agree with this observation, but I do not think that it affects my critical argument. When successful, public hate speech occurs in a general environment of intolerance in which both speakers and targets are components of an oppressive network that affects their public posture

(Delgado and Lederer 1995; Maitra 2012; McGowan 2012) as well as their propositional attitudes (Langton 2012) and dynamics of mutual recognition (Whitten 2018). But, within a racist and oppressive society, it remains difficult to demonstrate that one particular instance of public hate speech should take the blame for a broader structural phenomenon that is so pervasive as to influence the content of individual mental states, various social behaviors, and expectations of both speakers and hearers.

Against this backdrop, we might read public hate-speech events as contributions to a public hateful environment. Harm would be a long-term cumulative harm that accretes discriminatory attitudes and behaviors. As Sumner also argues (2010: 210), it is reasonable to say that public hate speech makes some contribution to the constitution of an unequal social context. However, it remains difficult to identify the extent of this contribution. For instance, in many public hate-speech situations, speakers do not have the formal authority to ensure successful uptake of their speech acts. Speakers might have, as Maitra (2012) argues, practical authority because their audience accommodates their presumptions. The constitution of such a common ground, then, may cause effects on one's propositional attitudes. This argument, however, relies upon the idea that certain beliefs and practices are so widespread that public hate speech, which draws upon this milieu, can be accommodated. Seen in this way, it is not public hate speech that is subordinating and harmful. The relevant public hate speech event comes about within an oppressive context that has already inculcated the relevant propositional attitudes in the target group or, at least, the predisposition to accommodate speakers' presumptions.

If we, therefore, want to accept the thesis of the cumulative harm of public hate speech within racist or hateful societies, we also have to accept that, as I argue, these acts do not necessarily cause an immediate and recognizable change in one's propositional attitudes, but rather they maintain an oppressive discourse that, then, informs the mental states of targets and third parties. By doing so, I think public hate-speech events would be understood as more or less subtle forms of advocacy, rather than as acts with an immediate effect on targets. This brings us to the second line of argument.

b. The indirect harm of public hate-speech events

Another prevailing line of argument begins with the observation that a great deal of speech aims at getting people to want (to feel) things they did not want (or feel) before (e.g., Delgado and Stefancic 1994; Hornsby and Langton 1998; Langton 2012; Lawrence III 2018; Maitra 2009). These views direct attention towards a conception of public hate speech as outspoken advocacy of hatred and discrimination of any form (UN 1965). According to this description, public hate speech causes social harm by promoting, disseminating, and upholding ideas based on racial superiority, discrimination, and negative stereotypes.

Here the harm occurs through the mental actions of the audience. The content of public hate speech persuades the audience to believe things or develop attitudes that, via their mental mediation, translate into harmful conduct. For instance, public hate speech convinces hearers to believe ideas based on racial hatred and discrimination. Public hate speech persuades hearers to engage in harmful conduct. Public hate speech normalizes certain conditions in a way that legitimizes violence against target groups. According to Charles Lawrence III (1992), public (racist) hate speech marks people of colour as socially subordinate and, in so doing, justifies discriminatory racist behaviour. Delgado and Stefancic also think public hate speech has effects on hearers' attitudes. For example, "before launching their wave of deadly attacks on the Tutsis in Rwanda," they write, "Hutus in government and the media disseminated a drumbeat of messages casting their ethnic rivals as despicable. The Third Reich did much the same with the Jews during the period leading up to the Holocaust" (2009: 363).

At the heart of these proposals, we find a more (Hornsby 1994; Hornsby and Langton 1998; Langton 2012) or less (Delgado and Stefancic 2009) explicit idea: speech-act theory is a helpful framework to study the effects of speech on the listeners.[13] According to Langton (2012), public hate speech has effects on hearers. Such effects are there because the incitement of hatred is a

characteristic of public hate speech. Much as the utterance "You are fired" constitutes a firing, hate speech constitutes an action of a particular kind—namely, intense hate (Maitra and McGowan 2010: 350). Hate speech, then, provided that the intended audience recognizes the action, and provided also that recipients perform a mental act (that leads them to take up a different attitude from the attitude they had before the utterance), may cause hearers to do certain things (Langton 2012). Actually, the effects are twofold. The first kind of effect is linked with the very presence of public hate speech in the society. The second kind of effect connects with what recipients do as a result of intense hate being present in the society. But how do these effects come about? Langton gives us an admittedly tentative explanation of this process.

The attitudes of hearers change, she says, because speakers act as if the recipients had the relevant attitude (2012). On this ground, speakers utter a series of fact-like and normative propositions. The repeated statement of such propositions alters attitudes so that recipients end up sharing the relevant attitude about the target group. In so doing, then, as Langton continues (2012), recipients absorb the relevant propositional attitude, and on this basis they may perform certain actions. For instance, as Langton (2012) writes, by telling a story, a speaker might present it as a fiction about the target group, but as a fiction that says something about the world through a series of fact-claiming and normative propositions. These propositions may change hearers' cognitive (because of the fact-claiming propositions) and emotional (because of the normative propositions) attitudes towards the object of the story. In this way, speakers succeed in enabling the effect they wanted to have.

In this way, I have pictured a particularly successful communication. However, to explain a public hate-speech situation, I should also think about what can go wrong with hateful statements. To do so, as many scholars have noted (e.g., Levin 2010; Maitra 2009), I cannot just concentrate on the individual speaker. It is important to extend my gaze to encompass the total situation. In this vein, let us try to read this situation through my five elements. Speakers are individual entities who utter sentences with a propositional content that degrades a target group. Here the propositional content is about targets, but targets are not direct addressees. Public hate speech,

for instance, depicts targets as socially subordinated as a way to motivate discrimination against them.[14] Then, the tolerant and intolerant are collapsed into a single entity as the direct addressees of public hate speech. The intolerant shared the same propositional attitude about the target group's states of affairs before the utterance. The tolerant, as I said before, are brought to have such a propositional attitude, or in Langton's own words, "Speakers invite hearers not only to join in a shared belief world, but also a shared desire world, and a shared hate world" (2012: 86). Public hate speech, then, can have effects on the tolerant and intolerant because there is a common ground between them and speakers. Therefore, a common ground between the speaker and her intended audience of tolerant and intolerant people is required for Langton's discourse on public hate speech to be successful.

It is generally understood that uptake requires an understanding of the propositional content and recognition of the intentions behind the performance of a certain act (Austin 1962; Maitra 2009: 313). Langton suggests that the phenomenon of accommodation between third parties and speakers could explain the effects of public hate speech on targets and the society as a whole through the audience. Some of the third parties had bad propositional attitudes before the public hate-speech event. Others adapt to the presuppositions of the speaker and, because of this newfound common ground, happen to have bad propositional attitudes towards the target groups that they did not have before. Since public hate speech intensifies pre-existing bad propositional attitudes and legitimizes new bad propositional attitudes about the target group, it is harmful. Here, increases in intensity and the addition of a new propositional attitude to one's set of propositional attitudes are the distinctive harms of a public hate-speech event.

I think this view requires too much of a public hate-speech event. Here my strategy is to understand public hate speech as an instance of a speaker meaning something (about targets) by her utterances. From this perspective, through the performance of a public speech act, speakers intend to make an audience respond in certain ways. While speakers can form reasonable expectations about certain intolerant groups (those groups that have responded as expected in prior public hate-speech events), such expectations, when directed at ostensibly tolerant third

parties, would not be grounded in the same way unless speakers already know they have genuine (and thus far hidden) bad propositional attitudes about targets. However, if such bad propositional attitudes already exist in the minds of otherwise ostensibly tolerant parties, it becomes very difficult to prove that a public hate-speech act causes changes in propositional attitudes and therefore harm.

It is not true that speakers will always make all third parties respond as they would expect. To me, Langton underestimates the observation that not all communications are effective.[15] When speakers say things in public, they may address people they know share the relevant propositional attitudes, but they may also fail to communicate successfully. As simple as the observation is, for a public hate speech to be successful as Langton (and many others) imagines, a common ground between speakers and third parties (tolerant and intolerant) has to exist. Yet, if such common ground already exists, it is very difficult to prove that a specific public hate-speech event is harmful.

Someone may object that my conception of common ground is too demanding: for meaningful conversations, it suffices to share some minimal elements that enable mutual understanding. This objection ignores two analytical elements. First, it overlooks the observation that public hate-speech events are not like quick and informal exchanges between two or more parties. Second, this objection neglects that a too minimal account of common ground would also undermine the thesis I am criticizing. For proponents of the view that hate speech harms, it is not important to demonstrate that hate speech makes sense to the audience. They also have to prove that hate speech is harmful. As said before, many scholars think hate speech is harmful because it occurs in a context in which some people have bad and deep-seated propositional attitudes about targets.

Moreover, as an anonymous reviewer has argued, my view takes the holding of the relevant bad propositional attitudes as 'rather more "black and white" than is plausible'. If we consider variations in intensity more closely, through exposure to hate speech an audience may come to hold the relevant propositional attitude with a higher degree of commitment. In such a case, the

account of hate speech being harmful on account of its changing propositional attitudes can hold. However, I believe that, when members of the audience have, either consciously or unconsciously, the relevant propositional attitude to the extent to which it is not unreasonable to expect successful uptake, public hate speech does not change existing propositional attitudes, but rather it makes existing ones more visible or relevant to the eyes of haters. If this is the case, I read this development as denoting a process bringing latent or weak propositional attitudes into clear consciousness. As I shall say in the next section, there is nothing intrinsically harmful about a propositional attitude coming to be more visible.

To sum up, in this section I have argued that when we take the full hate-speech situation into account, it is theoretically possible but very difficult to identify the specific harm of a public hate-speech event. If we presuppose that a common ground between speakers and targets exists, it is very difficult to prove that a public hate-speech event has caused a new propositional attitude in the targets. If we presuppose that a common ground between speakers and third parties exists, it is also hard to identify an increase in intensity or change in propositional attitudes. Alternatively, if we presuppose that common ground is not necessary, we overlook one of the conditions for speech to have effects. On these grounds, in the next section, by using the same explanatory tools, I offer an alternative proposal to study public hate-speech situations and the effects of public hate-speech events on targets and on tolerant and intolerant third parties.

5. A maieutic approach to public hate-speech situations

In this section, I offer an alternative approach to the study of public hate-speech situations and the effects of public hate-speech events that remains coherent with the theoretical framework of the arguments of section 4. So far, I have tried to remain consistent with the philosophical assumptions of scholars who argue for the direct/indirect harm of public hate speech. In the following, my aim is to think of an alternative way to assess the effects of public hate-speech

events. As I shall demonstrate, public hate-speech events do have effects on targets and third parties, but these effects are not necessarily harmful.

At the heart of my analysis, there are two observations: sentient beings have both conscious and unconscious propositional attitudes, and the combination of third parties' conscious and unconscious propositional attitudes influences what effect a public hate-speech event has on them. While the latter observation is relatively commonplace,[16] the former may be a source of many philosophical concerns. For instance, John Searle (1992) argues there is only conscious intentionality. Other philosophers would disagree about the propositional content of unconscious attitudes (e.g., Chalmers 2004; Crane 2015; Tye 1995). Here I follow Tim Crane (2016): unconscious mental states, such as unconscious propositional attitudes, are parts of one's entire attitude towards reality that are not as specific, determinate, and individuated as conscious attitudes, but are still very relevant for one's mind-to-world relation (2016). Mental states, then, may constitute networks of attitudes that, despite influencing our overall disposition towards the external world, are not fully intelligible to us unless an external phenomenon triggers a process of self-reflection.

The notion of unconscious mental states provides me with a plausible beginning of a different approach to public hate-speech situations. Targets, tolerant and intolerant people may have both conscious and unconscious attitudes towards propositional contents about the same state of affairs. These attitudes shape their perspectives on social and political realities, often causing the appeal to the same kind of tacit assumptions in different contexts and domains. By bringing unconscious attitudes into the discourse about public hate speech, we also multiply the possible avenues to establish common ground between speakers and third parties. Now a common ground exists even if one of the poles is not fully aware of all nodes of his or her network of attitudes. For this reason, in the picture we should include the fact that speakers can play with conscious and unconscious propositional attitudes to set the common ground that is necessary for successful communications. On closer inspection, this picture goes along with the view that

speakers themselves may be aware of the mind-to-world relations some addressees have, but it is fair to say they cannot know exactly what all possible addresses have in their minds.

In light of all this, I believe an alternative approach may give better accounts of public hate-speech situations and the effect of public hate-speech events. I call this perspective a maieutic approach to public hate-speech situations. In my view, unconscious and conscious mental states give purchase to the idea that public hate-speech events do not add new propositional attitudes, but rather they bring a person's latent dispositions into clear consciousness. From the third-person perspectives of targets, speakers, and genuinely tolerant people, public hate-speech events are junctures to make sense of the society in general, especially when certain dispositions are otherwise-implicit but widespread. At the same time, from the first-person perspective of the addressees, realizing their dispositions through public hate-speech events is a matter of self-knowledge. In some cases, through exposure to public hate-speech events, some third parties may gain greater clarity that their self-conception as nonracist individuals is not an accurate representation of their mental states. However, as an anonymous reviewer has suggested, people who are racist or homophobic often resent being told they are racist or homophobic. Mental states may inform their political preferences, their choice of living in certain neighborhoods, or their choice of renting rooms and apartments only to certain people, but they may honestly not even realize it. Arguing that, in a context of normalized racism and widespread oppressive practices, public hate-speech events enable an audience to admit to being racist would be a very large claim. However, public hate-speech events involve looking inwards and, vis-à-vis an expression of hatred, turning to our mental states, which may result in the self-knowledge of otherwise-unchecked mental states.

An advantage of this proposal is that it gives a more realistic account of individual dispositions and the interplay of different dispositions in a public hate-speech situation without resorting to disputable empirical claims. My description accepts the complexity of individual dispositions and, in this way, gives a nuanced portrait of the actors in a public hate-speech situation. The

observation that a great deal of what comprises our worldview is incomplete, maybe obscure, and not necessarily conscious is not new (Crane 2016; Freud 2012; Wilson 2002). I think such an observation should play an important role in how we think of realistic communicative interactions on complex political and social issues, such as those interactions characterizing public hate-speech situations. Another advantage of my proposal is that it gives a more nuanced account of the effect of public hate-speech events. Public hate-speech events do not directly cause any immediately traceable discrete change in one's conscious and unconscious propositional attitudes.

Let me explain. Comparisons with other accounts may help to clarify my view on the effect of hate speech. First, let me contrast my picture with the idea that public hate-speech events directly harm members of the target group. As seen, such a picture implies a common ground of bad propositional attitudes about the target groups shared between targets and the speaker. On that ground, because of what the speaker has uttered, members of the target group have new harmful (conscious and unconscious) propositional attitudes about themselves. The new attitudes constitute the distinctive harm of a public hate-speech event.

The alternative I am proposing is that public hate-speech events do not add any propositional attitude to the worldview of targets that is immediately traceable from the first-person perspective of the addressee. In other words, public hate-speech events do not cause any new conscious propositional attitude. A speaker plays with what she knows about the targets, but she also plays with what she thinks the target knows about herself. It is plausible to say this material cannot exist without the speaker's and targets' realizing it. If so, such a common ground of conscious propositional attitudes existed prior to the public hate-speech event. But it was also recognized as such by both speakers and targets. This observation undermines the thesis that a public hate-speech event could add a new conscious propositional attitude to the worldview of targets. In principle, it might add an unconscious propositional attitude, but, in an environment in which targets already have a number of other conscious bad propositional attitudes,[17] assuming

such dispositions arose without the targets' realizing, it is odd to argue retrospectively that it was such-and-such public speech event that caused the relevant propositional attitude, which, then, grew without the targets' realizing.

I think a maieutic approach gives a more accurate picture of public hate-speech situations. An analysis of public hate-speech situations should account for the possibility of communicative success, but it should have the instruments to study conditions for communicative unsuccess. A maieutic approach to hate speech helps us do so by looking at the set of propositional attitudes speakers and targets actually have. From this perspective, it remains very difficult (but not impossible) to say a speaker caused a traceable change in propositional attitudes. At the same time, according to a maieutic approach, speakers may fail to successfully communicate hatred when there is no common ground, neither conscious nor unconscious, between them and targets. Moreover, for a maieutic approach, speakers may also make targets realize something concerning the propositional attitudes they have about themselves. However, such self-discovery does not necessarily have an adverse effect on targets of the kind we need to argue that, through public hate speech, speakers directly harm their targets.

Let me now contrast my picture with the idea that speakers indirectly harm targets because public hate speech causes transformations in third parties' propositional attitudes. Such a picture implies that in public hate-speech situations, there is a common ground between speakers and third parties (the tolerant and intolerant). On such ground, tolerant and intolerant parties have new conscious harmful propositional attitudes about the target groups because of a public hate-speech event. These propositional attitudes are conscious because, at least according to the proponents of this view, they cause actions. As said before, if propositional attitudes against the target group are to be considered as conscious before the public hate-speech event, it is difficult to refute the claim that third parties had prior knowledge of their disposition towards the target. So it is also very difficult to detect the traceable effect of public hate speech on third parties. If

such bad propositional attitudes are to be considered as unconscious before the public hate-speech event,[18] it is hard to deny that such attitudes were already there waiting to be discovered. Here a maieutic approach to hate speech argues there is no addition of propositional attitudes through public hate-speech events, but rather these events make tolerant and intolerant people resolve part of the indeterminacy in their worldviews about the target group. In this way, we have a neutral perspective that is less deterministic on the possible effects of public hate speech. From such a standpoint, it is better to say public hate-speech events may involve realizing strong sympathy with the speakers. In this way, we maintain that public hate speech may have certain effects, but we are also able to hold that, despite speakers' presuppositions, public hate-speech events can involve realizing the absence of any ground of common shared propositional attitudes between oneself and the speaker.

We should not consider unconscious propositional attitudes as empty boxes. They have a propositional content that, at a certain moment of one's development, has been accepted as true. They are unconscious because, when a public hate-speech event occurs, we do not know about them and we are not aware of the ways we accepted such propositional content as true (Freud and Breuer 2004: 128-137). For instance, taboo breaking through public hate-speech events, as a reviewer has noticed, may make some bad propositional attitudes so socialized as to be consciously and unconsciously accepted as true. On this view, a public hate-speech event is one of many practices that make certain propositional attitudes socially acceptable (or even normalized) up to the point to which a person, in his or her development, may be influenced without even realizing that this is the case. Nevertheless, my account accepts also a second hypothesis. Exposure to taboo-breaking is not necessarily a bad thing. These events may trigger a process through which members of the audience realize the content of some propositional attitudes that they do not know themselves to have.[19] Following these observations, it is plausible to say a maieutic approach offers an original standpoint to study the effects of public hate-speech events without necessarily assuming such effects are harmful. On the one hand, my view says that, speakers can have an effect on targets when a common ground of shared bad

propositional attitudes about the target groups exists between them. If this is the case, public hate speech events make targets realize something about themselves. Otherwise, according to a maieutic approach, speakers are individual entities who utter sentences with a propositional content about a target group. Here the propositional content is about targets, but targets are not direct addressees. On such a view, tolerant and intolerant remain two separate entities. Speakers postulate two different common grounds: one with some of the tolerant people, one with the intolerant people. Therefore, speakers also address tolerant third parties on the presupposition that some of them might have the same unconscious propositional attitude about the target group. If common ground exists, there is no addition of new propositional attitudes, because, prior to the public hate-speech event, third parties already had all the necessary bad propositional attitudes towards targets. If such a common ground does not in fact exist, there is no basis for arguing that such a public hate-speech event has caused people to do something.

I believe that a maieutic approach to public hate-speech situations allows us to sharpen our perspective on the possible effects of public hate-speech events. By stressing the making conscious of otherwise-unconscious bad propositional attitudes, my account takes public hate-speech events as the first steps of processes through which the spread of hatred can be recognized and, maybe, challenged. On this view, some intolerant parties may realize (or they might be pushed to admit) the true character of their propositional attitudes and start working out how to revise them. Meanwhile, tolerant parties may become conscious that certain bad propositional attitudes are silent but widespread in their society. I know that such a picture is very optimistic. Nevertheless, it can indicate new avenues for future research in the area of normative responses to public hate speech. When inspired by a maieutic approach to public hate-speech situations, a normative justification against bans would begin with the idea that public hate-speech events occur when speakers know that someone will successfully uptake their expressions of hatred. Hate-speech events, therefore, would be crucial junctures to access those propositional attitudes that tolerant and intolerant individuals actually have. With this in mind, the goal of an argument for counter-speech would not be that of challenging speakers on their

own grounds. Counter-speech would be an instrument to work with third parties, both intolerant and tolerant, who, more or less consciously, contribute to the reproduction of discriminatory discourses and practices.

6. Conclusion

In this paper, I have provided an argument against the harm of public hate speech. To do so, I have identified the essential conceptual architecture of public hate-speech situations, I have assessed existing arguments for the direct and indirect harm of public hate speech, and I have proposed a novel way to approach public hate-speech situations: a maieutic approach. On this perspective, public hate speech brings a person's latent dispositions into clear consciousness. Overall, if my account holds, it is another proof that, in targets and third parties, public hate speech does not cause any discrete and traceable addition of harmful propositional attitudes. Therefore, a strong-enough reason to restrict public hate speech should be found elsewhere.[20]

Notes

1. Please note that I share the same concern as post-colonial, radical feminist, critical legal, queer, black empowerment, and other movements. It is not true that all residents in a liberal democracy can speak with equal voices. I also think that unequal clusters of power, wealth and influence affect political and social life in contemporary liberal democracies. The intent of this paper is to scrutinize the widespread idea that public hate speech causes a discrete and traceable change in one's propositional attitude. To my mind, this claim is not intuitive. I shall assess the descriptive and normative

observations that seem to support this thesis without questioning the idea that power differentials permeate social life. From this perspective, the maieutic approach can be seen as a starting point for a critical theory of hate speech situations.

2. This excludes face-to-face intimidations, targeted threats, and harassments from my picture. This is a common position in the debate about the harm of hate speech (e.g., Langton 2012; Lepoutre 2017; Maitra and McGowan 2012; Waldron 2014).

3. Therefore, the argument is not necessarily valid for all cultures at all times. Yet, I do not think that liberalism and democracy are perfectly coextensive. Specifically, the normative basis of a liberal argument against hate speech can be very different from the grounding of a democratic argument against hate speech. On this issue, see Heinze (2013 2016a 2016b).

4. For an example, "I believe that god exists" is a propositional attitude in which "I believe" expresses an attitude mode and "god exists" expresses a semantic content, which may be true or false.

5. I borrow this example from Rae Langton (2012).

6. In this paper, I maintain that public hate-speech events and public hate-speech situations are two connected but separate units of analysis. A public hate-speech event is what, like uttering a sentence or publishing a libel, triggers the interplay among the elements of a public hate-speech situation.

7. I think this framework helps to identify cases that seem to be hate-speech situations, but in fact are not. For instance, it is not rare to find Jewish black comedians offering impertinent descriptions of the Jewish community, or black hip-hop stars presenting demeaning representations of black people and ghetto life. Are these public hate-speech situations? It is reasonable not to think so. My analysis offers a simple way to substantiate this feeling: if speakers and addressees are both of the same target group, even if the content might be very offensive, there is no public hate-speech situation. I wish to thank [left empty for blind review] for pushing me on this point.

8. It is important to notice that these roles are not static: the same individual may perform multiple identities under different circumstances. For instance, a tolerant can act as an intolerant when targets change or when otherwise-tolerant people find themselves in new public hate speech situations. Equally, individuals who are members of the target group in a certain public hate speech situation may act as tolerant or intolerant in a new public hate speech situation with a different target.

9. According to Richard Delgado, one of the most significant characteristics of hate speech is that it is not intended to invite conversation (2018: 108).

10. Beliefs and propositional attitudes are not one and the same thing. However, it is commonplace in the literature on propositional attitudes to accept that a change in propositional attitudes implies a change in beliefs. Therefore, by saying a person changes her beliefs about self-respect, I am also saying a person changes her propositional attitudes connected with such beliefs.

11. Waldron (2012) says public hate speech attacks dignity as a civic status. Here I am bringing this position down to the fundamental units of thoughts and contents. On my reading of Waldron, a victim has a propositional attitude towards propositions such as "I am an equal member of the political community."

12. I do not mean to say all such scholars neglect the role of the tolerant and intolerant. Scholars such as Matsuda (1989) and Delgado (2009 2018) demonstrate that there could be a direct impact on third parties, but they support this claim with a different argument.

13. Here I focus only on Langton (2012). In this way, I do not account for all the subtle differences (and different degrees of philosophical complexity) within this line of argument. Nevertheless, I think Langton's way of approaching the problem makes explicit the fundamental conceptual structure to support the thesis that harm can occur through the mental actions of the audience.

14. Besides the work of Rae Langton, expressions such as "oppressive speech" (McGowan 2009) and "silencing speech" (Maitra 2009) reveal the same conceptual architecture. In these cases, speech is "oppressive" or "silencing" because it induces third parties to do something that reduces targets' room for action.

15. I think Ishani Maitra has criticized Hornsby and Langton's account of silencing with an analogous observation in mind. On her view, targets are silenced because their intended audience does not recognize their informative intention or because their intended

audience does not have reasons to produce the expected responses (2009: 324-27). Despite a strikingly different intent, Maitra also stresses the observation that not all communicative acts are successful.

16. At least, it could be seen as an elaboration of the idea of common ground as presented in section 3.

17. These pre-existing bad propositional attitudes are the ground upon which the very possibility of a successful communication can be built.

18. I think this is the combination of conscious and unconscious propositional attitudes that proponents of the indirect-harm thesis have in mind. For instance, see Lawrence III (2018: 58).

19. In this paper, I focus on public hate-speech events as possible new beginnings for self-acknowledgement. In this way, I do not want to deny the plausibility of public hate-speech events as contributing to the social acceptability of otherwise-unacceptable taboos.

20. These considerations do not exclude the possibility that public hate speech can be harmful or can contribute to the constitution of a harmful environment. My idea, however, is that it remains difficult to demonstrate that public hate speech is so recognizably harmful as to fall beyond the threshold of the harm test. Moreover, I see the burden of proof as resting on scholars who think public hate speech is harmful. My view draws upon an interpretation of the harm test in which the default presumption of

liberty assigns the burden of proof to those who want to subject speakers to coercive

interference (e.g., Sumner 2010). A longer discussion of this position, however, goes

outside the feasible reach of this paper.

Bibliography

Austin JL (1962) *How to Do Things with Words.* Oxford: Clarendon Press.

Brettschneider C (2012) *When the State Speaks, What Should It Say? How Democracies Can Protect Expression and Promote Equality*. Princeton: Princeton University Press.

Chalmers D (2004) The Representational Character of Experience. In: Leiter B (ed.) *The Future for Philosophy*. Oxford: Oxford University Press, pp. 153-81.

Cresswell MJ (1985) *Structured Meanings: The Semantics of Propositional Attitudes*. Oxford: Oxford University Press.

Crane T (2015) The Mental States of Persons and their Brains. In: O'Hear A (ed.) *Mind, Self and Person, Royal Institute of Philosophy Supplement 76*. Cambridge: Cambridge University Press, pp. 253 - 270.

Crane T (2016) The Unity of Unconsciousness. *Proceedings of the Aristotelian Society* 117 (1): 1-21.

Delgado R (2018) Words that Wound: A Tort Action for Racial Insults, Epithets, and Name Calling. In: Matsuda M (ed.) *Words That Wound: Critical Race Theory, Assaultive Speech, And The First Amendment*. New York: Routledge, pp. 89-110.

Delgado R and Stefancic J (1994) Hateful Speech, Loving Communities: Why Our Notion of a Just Balance Changes so Slowly. *California Law Review* 82: 851-69.

Delgado R and Lederer L (1995) *The Price we pay: the case against racist speech, hate propaganda, and pornography*. New York: Hill and Wang.

Delgado R and Stefancic J (2009) Four Observations About Hate Speech. *Wake Forest Law Review* 44: 353-70.

Fodor JA (1978) Propositional Attitudes. *The Monist* 61: 501-23.

Freud S (2012) *A General Introduction to Psychoanalysis*. Ware: Wordsworth Editions.

Freud S and Breuer S (2004) *Studies in Hysteria*. London: Penguin.

Garcia JLA (1996) The Heart of Racism. *Journal of Social Philosophy* 27 (1): 5–45.

Grice HP (1957) Meaning. *The Philosophical Review* 66 (3): 377-88.

Hasslanger S (2017) Racism, Ideology, and Social Movements. *Res Philosophica* 94 (1): 1-22.

Heinze E (2013) Hate Speech and the Normative Foundations of Regulation. *International Journal of Law in Context* 9 (4): 590-617.

Heinze E (2016a) *Hate Speech and Democratic Citizenship*. Oxford: Oxford University Press.

Heinze E (2016b) Towards a post-liberal theory of free expression. *Critical Legal Thinking*. http://criticallegalthinking.com/author/eric-heinze/

Hornsby J (1994) Illocution and its Significance. In: Tsohatzidis SL (ed.) *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*: New York-London: Routledge, pp. 187–207.

Hornsby J and Langton R (1998) Free Speech and Illocution. *Legal Theory* 4(1): 21-37.

Langton R (2012) Beyond Belief: Pragmatics in Hate Speech and Pornography. In: Maitra I and McGowan MK (eds.). *Speech and Harm: Controversies Over Free Speech*. Oxford: Oxford University Press, pp. 72-93.

Lawrence III CR (2018) If He Hollers Let Him Go: Regulating Racist Speech on Campus. In: Matsuda M (ed.) *Words That Wound: Critical Race Theory, Assaultive Speech, And The First Amendment*. New York: Routledge, pp. 53-88.

Lawrence III CR et al. (2018) Introduction. In Matsuda M (ed.) *Words That Wound: Critical Race Theory, Assaultive Speech, And The First Amendment*. New York: Routledge, pp. 1-16.

Lepoutre M (2017) Hate Speech in Public Discourse: A Pessimistic Defense of Counterspeech. *Social Theory and Practice* 43(4): 851-83.

Levin A (2010) *The Cost of Free Speech: Pornography, Hate Speech, and their Challenge to Liberalism*. Basingstoke: Palgrave Macmillan.

Lewis D (1979) Scorekeeping in Language Game. *Journal of Philosophical Logic* 8 (3): 339-59.

Maitra I (2009) Silencing Speech. *Canadian Journal of Philosophy* 39 (2): 309-38.

Maitra I and McGowan MK (2010) On Racist Hate Speech and the Scope of a Free Speech Principle. *Canadian Journal of Law and Jurisprudence* 23 (2): 343-72.

Maitra I and McGowan MK (2012) Introduction and Overview. In : Maitra I and McGowan MK (eds.). *Speech and Harm: Controversies Over Free Speech*. Oxford: Oxford University Press, pp. 1-23.

Matsuda M (2018) Public Response to Racist Speech: Considering the Victim's Story. In: Matsuda M (ed.) *Words That Wound: Critical Race Theory, Assaultive Speech, And The First Amendment*. New York: Routledge, pp. 17-52.

McGowan MK (2009) Oppressive Speech. *Australasian Journal of Philosophy* 87: 389-407.

McGowan MK (2012) On "whites only" signs and racist hate speech: Verbal acts of racial discrimination. In: Maitra I and McGowan MK (eds.). *Speech and Harm: Controversies over Free Speech*. Oxford: Oxford University Press, pp. 121-147.

Schauer F (1979) Speech and Speech - Obscenity and Obscenity: An Exercise in the Interpretation of Constitutional Language. *Faculty Publications.* 883.

Seglow J (2016) Hate Speech, Dignity and Self-Respect. *Ethical Theory and Moral Practice* 19 (5): 1103-16.

Shelby T (2002) Is Racism in the "Hearth"?. *Journal of Social Philosophy* 33 (3): 411-20.

Shelby T (2003) Ideology, Racism, and Critical Social Theory. *The Philosophical Forum* 34 (2): 153-88.

Stalnaker R (2002) Common Ground. *Linguistics and Philosophy* 25: 701-21.

Sumner LW (2010) Incitement and the Regulation of Hate Speech in Canada: A Philosophical Analysis. In Hare I and Weinstein J (eds.) *Extreme Speech and Democracy*. Oxford: Oxford University Press, pp. 204-20.

Tye, M (1995) *Ten Problems of Consciousness*. Cambridge MA: MIT Press.

UN (1965) General Assembly Resolution 2106A(XX), 21 December 1965, entered into force 4 January 1969.

Waldron J (2014) *The Harm in Hate Speech*. Cambridge MA: Harvard University Press.

Whitten S (online first) A Recognition-sensitive Phenomenology of Hate Speech. *Critical Review of International Social and Political Philosophy* https://doi.org/10.1080/13698230.2018.1511170

Wilson T (2002) *Strangers to Ourselves*. Cambridge, MA: Harvard University Press.

Yong C (2011) Does Freedom of Speech Include Hate Speech?. *Res Publica* 17 (4): 385-403.