

# Beyond Speech Acts: On Hate Speech and the Ubiquity of Norm Enactment

MARY KATE MCGOWAN\*

## ABSTRACT

*This paper argues against two frameworks for thinking about how language functions. The first such framework treats language use as primarily in the business of communicating content. On this content expression view, when we say things, we are only making claims about the world and/or offering considerations for or against such claims. It is shown here that this popular and even intuitive view of language use is problematically impoverished.*

*The second framework we shall consider is considerably richer; it acknowledges Austin's insight that we can do things with words, that speech can perform actions, enact facts, exercise power, and causally impact the world around us in a myriad of powerful ways. Although these insights are necessary to properly understand the complexity of language use, a speech act framework is nevertheless ultimately insufficient. In order to fully understand how our utterances impact the social world around us, we need to go beyond the intentional and conscious world of communicated speech acts; we need to recognize the unintended and barely conscious normative impact of our words. And, doing that requires a new framework beyond speech acts.*

## TABLE OF CONTENTS

I.	INTRODUCTION . . . . .	1056
II.	THE CONTENT EXPRESSION VIEW . . . . .	1057
	A. <i>Content Expression is Complex and Inferential</i> . . . . .	1057
	B. <i>The Standard Liberal Stance: More Speech</i> . . . . .	1058
	C. <i>Criticisms of the More Speech Response</i> . . . . .	1059
III.	THE SPEECH ACT FRAMEWORK. . . . .	1061
	A. <i>Intentionalism about Speech Acts</i> . . . . .	1061

---

\* Mary Kate McGowan is the Margaret Clapp '30 Distinguished Alumna Professor of Philosophy at Wellesley College. She specializes in metaphysics, analytic feminism, philosophy of language, and philosophy of law. Her most recent book, *Just Words: On Speech and Hidden Harm* was published by Oxford University Press in 2019. I thank the participants at the Ethics of Freedom of Speech Workshop, supported by the Georgetown Institute for the Study of Markets and Ethics, for helpful comments on this paper. © 2022, Mary Kate McGowan.

B.	<i>Revisiting Public Hate</i> . . . . .	1062
C.	<i>On Verbal Norm Enactment and Authority</i> . . . . .	1064
IV.	GOING BEYOND SPEECH ACTS . . . . .	1066
A.	<i>On Conversational Norm Enactment</i> . . . . .	1066
B.	<i>Conversational Score</i> . . . . .	1067
C.	<i>Conversational Norm Enactment</i> . . . . .	1069
V.	HOW IT GENERALIZES . . . . .	1070
VI.	<i>PUBLIC HARM REVISITED</i> . . . . .	1072
VII.	CONCLUSION . . . . .	1074

## I. INTRODUCTION

Consider the following situation:

*Public Hate:* An Asian woman is walking to work when a young white man, who is unknown to her, steps toward her on the sidewalk and angrily says, “I lost my job because of your virus; go back where you belong and take your diseases with you!”

Although this particular example is fictional, the recent rise, here in the United States, in anti-Asian incidents like this one is all too real. Before we decide what to do—individually or collectively—about such incidents, we must first understand more clearly how utterances like this function—linguistically and otherwise. This paper will focus on this clarifying task.

In the process, I will argue against two frameworks for thinking about how language functions. The first such framework treats language use as exclusively communicating content. On this sort of view, which we might call the content expression view, when we say things, we are only making claims about the world and/or offering considerations for or against such claims. As we shall see, this popular and even intuitive view of language use is problematic.

The second framework we shall consider is considerably richer; it acknowledges Austin’s insight that we can do things with words, that speech can perform actions, enact facts, exercise power, and causally impact the world around us in a myriad of powerful ways.<sup>1</sup> Although these insights are necessary to properly understand the complexity of language use, a speech act framework is nevertheless

---

1. J. L. AUSTIN, *HOW TO DO THINGS WITH WORDS* (1962).

ultimately insufficient. In particular, I shall argue here that in order to understand how cases like *Public Hate* work, we need to go beyond the intentional and conscious world of communicated speech acts; we need to recognize the unintended and barely conscious normative impact of our words.

The paper proceeds as follows. In Part I, I consider and criticize the content expression view of language use. Then, in Part II, I present the speech act view and apply it to *Public Hate*. There, I raise several challenges to a speech act account of the utterance in *Public Hate*. In Part III, I argue for a hidden way that speech enacts norms and apply that mechanism to *Public Hate*. The resulting view of language use shows norm enactment to be an inescapable and prevalent aspect of participation in norm-governed practices. Finally, I identify some consequences of this view.

## II. THE CONTENT EXPRESSION VIEW

Our first order of business is to evaluate the content expression view of language use. As we see in the following section, it is considerably more complex than one might think.

### A. *Content Expression is Complex and Inferential*

The content expression view is sophisticated enough to account for the fact that there are a variety of ways to communicate or express content. It is worthwhile to briefly explore this expressive complexity so that we do not unfairly criticize this model.

One might think that when a speaker wants to communicate some proposition or claim, *p*, the speaker says something that (literally) means *p*. Then, a hearer (or the receiver or interpreter of the utterance) accesses that meaning, *p*, by simply decoding the literal meaning of what the speaker actually says. On this picture, communication is just a matter of coding and decoding the content of what the speaker wants to communicate. This picture of language use is highly intuitive. After all, we use language to communicate what we mean, and we say (or write) things that have meaning, so it makes sense that what we say (or write) should match the meaning we intend to convey.

As intuitive as this picture may be, it is wrong. It cannot deal with a wide variety of different ways in which we get content across linguistically. In fact, we rarely just come out and say what we mean. Instead, we say something else that enables the hearer to *figure out* what we mean.<sup>2</sup>

To see an example of this, consider the following. Suppose that Simon and Edi are talking about how strange Deepak has been behaving lately and Edi says, "Anyone who just got fired for cause *would* be freaking out." Simon takes Edi to be telling him that Deepak just got fired for cause and this explains why Deepak is behaving so strangely lately. Although this is precisely what Edi intends to

---

2. See, e.g., PAUL GRICE, *STUDIES IN THE WAY OF WORDS* 26–31 (1989).

communicate to Simon, it is not what Edi actually says. Edi actually says a general claim about people who get fired for cause; it is not a claim about Deepak at all. Despite this, Simon is able to infer Edi's intended meaning (the claim about Deepak) from the conventional meaning of what Edi actually says, the context, and the cooperative nature of conversation. This phenomenon is called conversational implicature and it was first identified by H.P. Grice. The point is what Edi means and what Edi says are different, so communication is not a mere matter of decoding the meaning of the language used.

Language theorists now agree that language use is highly inferential. In fact, even direct literal language use involves complex inferential reasoning. Suppose, for example, I say, "Shawn is tall," and I mean what I say. I am not being ironic or insinuating something else is true of Shawn because he is tall. Even here, however, hearers must make rather complex inferences to correctly interpret my utterance. For starters, they need to figure out which Shawn I am talking about. Furthermore, they need to figure out what constitutes being tall for someone like Shawn. If he is a professional basketball player, for example, then saying he is tall probably means he is over seven feet. If, however, Shawn is a preschooler, then it means no such thing.<sup>3</sup> Even at the level of what we actually say, complex inferences are required.

There are still further ways to get across content or meaning. One such way is called presupposition.<sup>4</sup> When I say my sister just bought a house in Maryland, I presuppose (as opposed to explicitly say) I have a sister. I take it for granted. I treat it as uncontroversial. I act as if everyone already knows it or would happily accept it without controversy. As some would put it, I treat the claim that I have a sister as not-at-issue content. Still other ways to get content across include irony, insinuation, conventional implicature, and logical presupposition.

In sum, the content expression view can accommodate a wide variety of ways of expressing content. Even so, we shall now see that it is insufficiently complex to accommodate how language actually works.

### B. *The Standard Liberal Stance: More Speech*

Before moving on to consider what is missing from the content expression view, it is worth considering how that view would apply to discussions about the free speech status of hate speech.<sup>5</sup> Now, I take the utterance in *Public Hate* to be

3. There is considerable controversy in the literature regarding the role of context in fixing the conventional meaning of (or the proposition expressed by) the sentence uttered. For a sampling of positions, see ERNEST LEPORE & HERMAN CAPPELEN, *INSENSITIVE SEMANTICS: A DEFENSE OF SEMANTIC MINIMALISM AND SPEECH ACT PLURALISM* (2005); Jason Stanley, *Context and Logical Form*, 23 *LINGUISTICS & PHIL.* 391, 391–434 (2000); ROBYN CARSTON, *THOUGHTS AND UTTERANCES: THE PRAGMATICS OF EXPLICIT COMMUNICATION* (2002); John MacFarlane, *Non-indexical Contextualism*, 166 *SYNTHESE* 231, 231–50 (2009).

4. There are different kinds of presupposition (e.g. semantic and pragmatic), but we do not need to get into such technical details.

5. There is no standard satisfactory definition of hate speech; it is defined differently for different purposes, legally and otherwise. See generally ALEXANDER BROWN, *HATE SPEECH LAW: A PHILOSOPHICAL EXAMINATION* (2015). Here, I focus on the face-to-face kind of hate speech theorized by

an instance of the face-to-face type of hate speech characterized by Mari Matsuda<sup>6</sup>; notice this instance of hate speech does not involve racial slurs and the speaker might not regard what he says as hate speech. Even so, I take it to be hate speech.

Let us now look at the standard position on the free speech status of hate speech here in the United States. For the most part, hate speech is not regulated. There are some exceptions. Hate speech that constitutes a true threat is criminal; hate speech that constitutes harassment, or an intentional infliction of emotional distress, is civilly actionable; but for the most part, hate speech is considered highly protected political speech.

Let us look again at our original example:

*Public Hate:* An Asian woman is walking to work when a young white man, who is unknown to her, steps toward her on the sidewalk and angrily says, “I lost my job because of your virus; go back where you belong and take your diseases with you!”

Applying this standard position on hate speech to this case, it seems that the young man is simply expressing a political opinion. And, as a citizen of the United States, he has a right to have and to express that opinion. Applying the content expression view of language, expressing that opinion—and perhaps some considerations in favor of that opinion—is *all* the young man’s utterance is doing.

On this way of thinking about the case, the remedy is straightforward. The young man’s political position involves false and unwarranted content. What he says presupposes, for example, that COVID-19 is *her* virus, thereby communicating the claim that Asians are somehow responsible for the pandemic. This is false and should be contested. The young man also presupposes she belongs outside of the United States and—evidently—he does so simply because she appears to him to have Asian ancestry. This too is false and should be contested. On the standard take on hate speech then, the proper remedy is counter-speech. Hate speech expresses unwarranted and false content; the remedy is to take issue with that false and unwarranted content, and this has come to be known as the “more speech” response (MSR).

### C. Criticisms of the More Speech Response

The MSR has been met with considerable criticism. Here are just some of the concerns: The MSR unfairly places the burden of a remedy on those targeted by the hateful speech.<sup>7</sup> Why should it be up to the Asian woman in *Public Hate* to

---

Matsuda as opposed to the propaganda kind of hate speech (e.g., Nazi anti-Semitic propaganda circa WWII). See Mari Matsuda, *Public Response to Racist Speech: Considering the Victim’s Story*, 87 MICH. L. REV. 2320, 2320–81 (1989).

6. See *id.*

7. See Frederick Schauer, *Uncoupling Free Speech*, 92 COLUM. L. REV. 1321, 1321–61 (1992).

respond to what the young man says? Why should she be burdened with identifying the false and unwarranted content and then having to set the record straight? That seems unfair.<sup>8</sup>

The MSR also seems to assume it is always—or at least usually—safe for the addressee to respond. This assumption is dubious at best. In *Public Hate*, for example, the Asian woman might well feel physically unsafe when accosted by a stranger in this way and, even if she feels physically safe, she might well feel unsafe in other ways arguing with a complete stranger in a public space.

Moreover, this remedy seems to assume, in fact, that people *do* respond to hate speech in public places. Yet, empirical studies have shown—for a variety of reasons (e.g., concern for safety, not wanting to waste energy trying to educate the hopeless, etc.)—people do not respond to hate speech in this way.<sup>9</sup>

Additionally, the MSR appears to assume a level playing field; in particular, it seems to assume that every person is as able as any other to speak, be heard, properly understood, and be given appropriate credibility. We all already know that it's not a level playing field. Apparently then, the MSR overlooks how some people have their communicative capacities systematically undermined (in other words, silenced).<sup>10</sup>

Last but not least, the MSR seems to operate on the assumption that hate speech is merely—or even just primarily—in the business of expressing a political opinion; as we shall soon see, this is false. Speech does considerably more than express content.

In fact, it might seem ludicrous to suggest that the young man in *Public Hate* is merely expressing a political opinion and that his utterance is just an opening move in a potential political discussion, as if what he says is an invitation to discuss a matter of mutual public concern. It seems misguided at best to suggest that the Asian woman ought to simply engage and offer counter considerations, thereby undoing everything problematic about what the young man said.

Keep in mind here that the young man's utterance—as troubling as it is—is not legally regulable—at least not currently in the United States. Although it is threatening, it is nevertheless not a true threat in the legal sense. It is also not harassment or stalking (until it is part of a documented pattern). And, although the young man's utterance could be an intentional infliction of emotional distress, good luck with that in civil court.

Here are two additional reasons to be concerned about the standard stance that hate speech is protected political expression. First, we routinely regulate

8. There are other kinds of counter-speech that do not place the burden on the target. See, e.g., KATHARINE GELBER, *SPEAKING BACK: THE FREE SPEECH VERSUS HATE SPEECH DEBATE* (2002).

9. See LAURA BETH NIELSEN, *LICENSE TO HARASS: LAW, HIERARCHY, AND OFFENSIVE PUBLIC SPEECH* (2004); Laura Beth Nielsen, *Power in Public: Reactions, Responses, and Resistance to Offensive Public Speech*, in *SPEECH AND HARM: CONTROVERSIES OVER FREE SPEECH* 148–73 (2012).

10. By now, this is an immense literature. For an overview, see Mary Kate McGowan, *On Multiple Types of Silencing*, in *BEYOND SPEECH: PORNOGRAPHY AND ANALYTIC FEMINIST PHILOSOPHY* 39–58 (2017).

panhandling without it raising any free speech concerns. So, going up to someone on the street and requesting money is regulable (on the grounds that it interferes with commerce) but going up to someone and screeching racist vitriol is not regulable (on the grounds that it merely expresses a political opinion). I find this troubling. Second, although the standard stance on the free speech status of hate speech appears to operate with a content expression view of language use, elsewhere in the law and criminal justice system it is well recognized that speech does considerably more than merely express content. Speech can hire an assassin (e.g., criminal solicitation), enact agreements to raise prices in tandem (e.g. price fixing), ask for, offer, or accept a bribe, threaten with bodily harm, enact discriminatory policies, commit blackmail, and so on. One might well wonder why this apparent double standard?

### III. THE SPEECH ACT FRAMEWORK

Let's now consider a different, and considerably less impoverished, framework for understanding how language functions: speech acts.

#### A. *Intentionalism about Speech Acts*

Ever since Austin's *How to Do Things with Words*, philosophers have paid philosophical attention to the capacity of language to perform action.<sup>11</sup> Saying "I apologize," for example, performs the act of apologizing. The utterance of "I apologize" is not a description of a separate—perhaps internal—act that is the apology; the utterance *is* the apology. Similarly, saying "I promise to cut the lawn on Saturday" performs the action of promising. Uttering that promise counts as undertaking an obligation on the part of the speaker to cut the lawn on Saturday. These examples demonstrate that saying something can sometimes constitute doing something. And, as Austin would put it, we are "doing things with words." Language use then is not exclusively in the business of making claims about the world; it is not just in the business of conveying meaning or content.

Austin called these speech actions illocutionary acts. It is important to stress that illocutionary acts are not purely physical actions, e.g., the physical action of producing sound, signs, or text. Illocutionary acts, by contrast, depend on the use of conventional language and the communication of conventional meaning. My saying "hello" to my neighbor is an illocutionary act because the sound I produce has a conventional meaning that is heard and recognized by my neighbor and, in virtue of these things, my utterance constitutes the action of greeting. So, the mere production of sound is not an illocutionary act, but a verbal greeting is an illocutionary act.

Further, Austin stressed that illocutionary acts abound. When we assert things, we commit to the truth of what we say and that act of commitment is an

---

11. See generally AUSTIN, *supra* note 1.



illocutionary act.<sup>12</sup> When we describe something, that act of description is an illocutionary act. Most of the time when we say things, then, we are also performing illocutionary acts.

There is some controversy regarding how illocutionary acts work but, on the dominant account (called intentionalism), illocutionary acts require a complex form of speaker intention and typically work via the recognition of that intention. Consider the following examples.

When Gretchen asserts that the department meeting is on Thursday, she commits to the truth of the claim that the department meeting is on Thursday. That is what it means to assert.<sup>13</sup> Gretchen intends to commit to the truth of this claim and for her addressee to recognize that intention. (Linguistic intentions are complex in this way.<sup>14</sup>) In this picture, assertion works via the recognition of the speaker's intention to assert. The same goes with promising. When Taylor says, "No worries; I swear; I'll be there on time," Taylor promises, because, when she says what she says, Taylor intends to undertake an obligation to be there on time and she also intends for her addressee to recognize this. In other words, Taylor has an illocutionary intention to promise.

There are plenty of complexities and controversies here. There is controversy over whether intentionalism is the correct account of illocution;<sup>15</sup> there is controversy over whether illocutionary acts ontologically depend on antecedent social practices; there are disputes over whether the speaker's illocutionary intention must be recognized in order for an illocutionary act to obtain.

Such complexities aside, the main point is that language use appears to be intentional, deliberate, and conscious. When we convey meaning, we intend to do so. Even when we convey more than we literally say, we intend to do so, and we are recognized as intending to do so. And that is precisely how it works. The same is true for illocutionary acts. When we perform illocutionary acts with our utterances, we are intending to do so, and we manage to succeed precisely because others recognize that intention. According to this speech act framework then, the various things we do with language we do deliberately, intentionally, and consciously.

### *B. Revisiting Public Hate*

In Part I.B, we saw that the young man's utterance in *Public Hate* does more than merely express content, but what else does it do? In particular, which illocutionary acts are performed by the young man's utterance? Unsurprisingly, this is not a straightforward question. One reason for this is that the young man's

12. This is a standard way to understand assertion, and there are others. Stalnaker, for example, treats assertion as a proposal that the asserted content be accepted by all conversational participants. See Robert Stalnaker, *Common Ground*, 25 LINGUISTICS & PHIL. 701–21 (2002); ROBERT STALNAKER, *CONTEXT AND CONTENT: ESSAYS ON INTENTIONALITY IN SPEECH AND THOUGHT* (1999).

13. Again, this is one standard view of assertion. See Stalnaker, *Common Ground*, *supra* note 12.

14. For details, see GRICE, *supra* note 2, at 86–137.

15. Intentions still play a crucial role in other accounts (e.g., conventionalism) of illocution.



utterance (i.e., “I lost my job because of your virus; go back where you belong and take your diseases with you!”) does many things at once.

At the very least, it accuses, blames, and instructs the addressee to go back to where she belongs. More controversially, it insults, degrades, and perhaps even essentializes, and it does so on the basis of race. Some theorists go further and argue that speech actions like this one subordinate, oppress, and even constitute acts of discrimination. These hypotheses about the illocutionary acts performed require substantial elucidation and defense, which I will not provide here. My interest lies elsewhere.

Claims like this arose in law schools. Catharine MacKinnon, for example, famously claimed that (some) pornography subordinates and silences women.<sup>16</sup> Critical race theorists, like Mari Matsuda and Charles Lawrence, claim that racist hate speech subordinates and makes people of color count as second class citizens socially.<sup>17</sup> Such claims were initially met with confusion, even incredulity.<sup>18</sup> How could speech possibly do such things? Speech can *cause* subordination, oppression, and discrimination by altering beliefs in ways that cause people to behave in subordinating, oppressive, and discriminatory ways, but how can mere words actually perform these types of actions?

Philosophers stepped in and used speech act theory to defend the coherence of these sorts of claims.<sup>19</sup> On this kind of this account, an utterance subordinates when it enacts a norm that subordinates. (Norms, in this sense, include permissions and obligations.) When an utterance enacts harmful norms, that utterance is said to *constitute*, as opposed to merely *cause*, the harm in question. Verbal harm constitution then involves verbally enacting harmful norms.

Let’s now bring this back to the young man’s utterance in *Public Hate* and let’s focus on the harm of subordination in particular. On this speech act account of harm constitution, it seems that in order for the young man’s utterance to subordinate the Asian woman, his utterance must somehow enact subordinating norms. In the following section, we explore how speech actions enact norms, at least according to the speech act framework, and then we evaluate the claim that the young man’s utterance subordinates the Asian woman in *Public Hate*.

16. See generally CATHARINE MACKINNON, ONLY WORDS (1993); CATHARINE MACKINNON, TOWARDS A FEMINIST THEORY OF THE STATE (1989); CATHARINE MACKINNON, FEMINISM UNMODIFIED: DISCOURSES ON LIFE AND LAW (1987).

17. See Mari Matsuda, *Public Response to Racist Speech: Considering the Victim’s Story*, in WORDS THAT WOUND: CRITICAL RACE THEORY, ASSAULTIVE SPEECH, AND THE FIRST AMENDMENT 17–51 (1993); Charles R. Lawrence III, *If He Hollers, Let Him Go: Regulating Racist Speech on Campus*, in WORDS THAT WOUND: CRITICAL RACE THEORY, ASSAULTIVE SPEECH, AND THE FIRST AMENDMENT 53–88 (1993).

18. See W.A. Parent, *A Second Look at Pornography and The Subordination of Women*, 87 J. PHIL. 205–11 (1990).

19. Under analytic feminism, the speech act approach started with Langton. See Rae Langton, *Speech Acts and Unspeakable Acts*, 22 PHIL. & PUB. AFFS. 293–330 (1993). Vadas offered a different elucidation and defense of MacKinnon’s claim—that pornography subordinates women. See Melinda Vadas, *A First Look at the Pornography/Civil Rights Ordinance: Could Pornography Be the Subordination of Women?*, 84 J. PHIL. 487–511 (1987).

*C. On Verbal Norm Enactment and Authority*

We are now focusing on a fairly specific speech action: norm enactment. Sometimes, an utterance enacts norms. Consider the following examples.

*Smoking Policy:* A college president enacts a new college policy when she says, “Effective October 15th, 1992, smoking is prohibited inside any college building.”

*New Bedtime:* A father says to his daughter, “Now that you are a big girl and have turned eight, you can stay up to watch the Muppet Show on Friday. Your Friday bedtime is now 8:00 p.m.”

*The Firing:* The boss says to her employee, “We have now reached that point that we discussed in our prior meetings. The measurable goals are once again unmet. You’ve been given fair warning. I am left with no viable alternative. Effective immediately, you’re fired. Clean out your desk and your locker.”

In each of these cases, the speaker’s utterance enacts changes to what others are permitted to do. In *Smoking Policy*, the college president changed college policy. After her declaration, and once the enacted college policy went into effect (on October 15th, 1992), smoking inside a college building is impermissible where it was previously permissible. The college president’s words enact changes to norms (e.g. college policy). In *New Bedtime*, the father made it permissible for his daughter to stay up until 8:00 pm on Friday evenings. His saying so made it so; he enacted a new household rule and widened the permissible options for his daughter. Finally, in *The Firing*, the boss fired the employee, which enacts changes to various obligations and permissions for that employee. Once fired, the employee is neither permitted to remain in the workplace nor entitled to demand a paycheck for work done after the firing.

As these examples illustrate, speech can change norms. So, enacting new norms is something that can be done with speech; it is a kind of illocutionary or linguistic action. And, it is worth pointing out that each of these examples have several features in common. First, the speaker intends to be changing norms. The college president, the father, and the boss are each intending to change what is permissible for others. Second, others recognize the speaker’s intention to do this. The members of the college community, the 8-year-old daughter, and the employee each recognize what the speaker is intending to do. Third, and finally, the speaker has—and is exercising—the authority to change the norms in question. Only the college president can enact new college policy; only a parent (or legal guardian) can change a child’s bedtime; and only a boss can fire an employee. These permission-changing speech actions appear to require an exercise of speaker authority.

On this speech act account then, verbal norm enactment requires an exercise of speaker authority. It also seems to work the way communication generally does—namely, via the recognition of the speaker’s conscious intention to enact the

norms in question. But these conditions raise problems for a speech act account of harm constitution in cases like *Public Hate*.

First, consider the issue of speaker authority. In *Public Hate*, the speaker, i.e., the young man, does not appear to have the authority required to enact subordinating norms. This is called the authority problem.<sup>20</sup> In response, theorists have pointed out that not all authority is official;<sup>21</sup> they have also offered a variety of reasons to regard the young man's utterance as an authoritative one. Perhaps the young man has an unofficial kind of authority by virtue of having a long history of institutionalized racism behind him;<sup>22</sup> perhaps the young man has the authority of the government behind him since his utterance is treated as highly protected political speech;<sup>23</sup> perhaps the young man has authority conferred on him by bystanders' failure to object.<sup>24</sup> As ingenious as these strategies are, it's ultimately unclear that any of them succeed in establishing the authority of ordinary speakers of hate speech. But, even if any of these strategies were to succeed, other crucially important felicity conditions of verbal norm enactment remain unsatisfied, at least on a speech act account.<sup>25</sup> In *Public Hate*, the young man does not appear to be intending to enact any norm, never mind a subordinating one; the Asian woman does not take him to be doing so, and his utterance does not express the content of a subordinating norm. All of these are reasons to reject a speech act account of harm constitution in this case.

Of course, I acknowledge that these considerations are also a reason to reject the claim that the utterance in *Public Hate* involves harm constitution at all. Why assume that it does or even that it might?

Although a full defense of this claim is beyond the scope of the present paper, a brief sketch of the theoretical machinery behind it is both possible and prudent. For starters, I work within a school of thought that regards language use as central to social hierarchy and social constriction. To help to see how this might be, notice first that there is very good reason to believe that everyday utterances play a crucial role in social construction generally. Social constructions involve what John Searle calls a constitutive rule, i.e., X counts as Y in C.<sup>26</sup> To illustrate, let's consider one of his favorite examples of a socially constructed thing: money. This particular piece of paper (X) counts as a ten-dollar bill (Y) here in the U.S. (C), and this means that the possessor of the piece of paper is entitled to certain things as a result. In other words, social construction imposes a socially shared status on things. And, if that status is not collectively recognized—if, for

20. See Ishani Maitra, *Subordinating Speech*, in *SPEECH AND HARM: CONTROVERSIES OVER FREE SPEECH* 94, 95 (2012).

21. See Rae Langton, *Subordination, Silence and Pornography's Authority*, in *CENSORSHIP AND SILENCING: PRACTICES OF CULTURAL REGULATION* 261, 270–73 (1998).

22. See generally Matsuda, *supra* note 5, at 2332–34.

23. See *id.*

24. See generally Maitra, *supra* note 20.

25. For a more detailed exploration, see Mary Kate McGowan, *Conversational Exercitives and the Force of Pornography*, 31 *PHIL. & PUB. AFFS.* 155, 155–89 (2003).

26. See JOHN R. SEARLE, *THE CONSTRUCTION OF SOCIAL REALITY* 27–29 (1995).

example, all of a sudden, we ceased to regard that piece of paper as valuable—then it would cease to function as currency; in short, it would cease to *be* money. In this way, social constructions ontologically depend on this (complex) kind of collective recognition. The central point here is that communicative acts, and thus language use, are crucial to both the imposition of the status function and the continued and shared recognition of it.

Now, if, as many theorists argue, things like race and gender are socially constructed, then language also plays a crucial constitutive role in social ontology—that is, in who counts as what socially and thus in who gets treated how. And this suggests that ordinary everyday speech actions are playing a crucial constitutive role in all of this—in coordinating the behavior that undergirds it and in signaling the shared recognition of who counts as what socially. So, on this sort of picture, we have good reason to suspect—and thus to want to further investigate—the possibility that the young man’s utterance in *Public Hate* is constitutively harmful after all. It’s just that the enacting of harmful norms isn’t conscious, intentional, and deliberate the way speech acts are; it’s sneaky and rendered invisible by its ubiquity the way social constructions are.

#### IV. GOING BEYOND SPEECH ACTS

If we go beyond speech acts, beyond what speakers are aware of doing and intending to do, and think of language use as an embedded social practice, we shall see that there is another way that speech enacts norms. This other way does a much better job explaining how the young man’s utterance might enact subordinating norms. And, even if one is unsympathetic to the claim that ordinary instances of hate speech can constitute harm, we still have good reason to attend to this other way that speech enacts norms. After all, it’s an important part of what speech does.

First, I will argue for this other mechanism of norm enactment within conversation. Then, in Part IV, I will argue that the phenomenon is generalizable.

##### *A. On Conversational Norm Enactment*

In this section, I will argue for a subtle mechanism of norm enactment that is operative in conversations. Although we all participate in conversations, when we do so, we are doing considerably more than we either realize or consciously intend. In fact, we are routinely enacting conversation-specific norms.

To see this, notice first that conversations are norm-governed activities. This just means that there are shared normative expectations regarding what one ought to do when one contributes to a conversation. Let’s look at a few examples.

One such norm is the requirement that a person speak in ways that others are able to understand. If Madeleine and Trayvon are discussing local parks, and Madeleine suddenly starts speaking French, which Trayvon does not speak, then Madeleine is violating this norm. She is not being communicatively cooperative; she is not behaving in ways that enable Trayvon to understand what she means to get across.

Another conversational norm is that one ought to have sufficient evidence for what one says. How much evidence is sufficient can, of course, depend on the context, but sometimes it is perfectly clear that a conversational contribution is unwarranted. To see this, suppose that when Madeleine and Trayvon are discussing local parks. Madeleine says, “that park is going to close for good next week,” but Madeleine has no evidence whatsoever for making this claim. In such a case, her conversational contribution is legitimately subject to critique.

Yet another conversational norm is that one’s contributions ought to be relevant to the conversation at hand. Blurting out irrelevant things, even if what is said is easily understood and well supported, is nevertheless conversationally inappropriate. When contributing to conversations, participants ought to make their contributions relevant to the conversation.

These examples are enough to show that conversations are governed by prescriptive norms. I call these norms g-norms; they govern conversations generally. That conversation is a g-norm-governed activity means that contributions to them can be inappropriate and even entirely out of bounds. It also means that conversational contributions can be more or less appropriate, and they can be more or less appropriate along several different dimensions of appropriateness. A contribution can, for example, be more or less relevant, more or less warranted, and more or less intelligible.

### B. Conversational Score

Philosophers of language and linguists have different ways to keep track of what is happening in a conversation; in other words, there are different ways to specify the conversational context.

One such way is the conversational score.<sup>27</sup> David Lewis, for example, worked within a conversational score framework.<sup>28</sup> On Lewis’s rather inclusive notion of conversational score, it tracks *everything* that is relevant to the proper development and assessment of the conversation. This includes, among other things, the presuppositions, the appropriate standards of accuracy, and the relevant topics.<sup>29</sup> Since the various components of conversational score affect such a wide variety of linguistic phenomena (which may not be familiar to some readers), it is worthwhile to consider some examples.

---

27. Common ground is another way, developed by Robert Stalnaker. See Robert Stalnaker, *Presuppositions*, 2 J. PHIL. LOGIC 447 (1973); Robert Stalnaker, *Pragmatic Presuppositions*, in SEMANTICS AND PHILOSOPHY 197 (Milton K. Munitz & Peter K. Unger eds., 1974); Robert Stalnaker, *On the Representation of Context*. 7 J. LOGIC, LANGUAGE & INFO. 3 (1997).

28. See David Lewis, *Scorekeeping in a Language Game*, J. PHIL. LOGIC 339 (1979).

29. *Id.* at 345.

Definite descriptions are one such linguistic phenomenon; they are descriptions that purport to uniquely refer or refer to exactly one thing.<sup>30</sup> Examples include ‘the tallest student in my logic class,’ ‘Deirdre McGowan Galway’s fourth child,’ and ‘my favorite color.’ As is well known, however, many definite descriptions appear to succeed in uniquely referring even though these descriptions fail to uniquely describe their referent. The expression ‘the desk,’ for example, may pick out a particular desk even though there are many desks in the universe, and there may even be several desks in the room in question. Salience appears to account for this.<sup>31</sup> On this account, a definite description refers to the most salient satisfier of the description.<sup>32</sup>

Let’s look at an example. Suppose that Edilia mentions that her car has just been to the repair shop, and I ask if the car is now functioning properly. Edilia’s car is certainly not the only car in the universe and may not even be the only car present (since I have one too), but I have nevertheless managed to refer to her car with the expression ‘the car.’ This is because Edilia’s car is the most salient car in the context of this particular conversation. Salience is a component of the conversational score, and this salience component of the conversational score helps to settle the appropriate use of definite descriptions by helping to fix the unique referent of such descriptions.

Consider now another linguistic phenomenon that draws our attention to a different component of the conversational score, the scope of quantifiers. When we use words like ‘all,’ ‘some,’ ‘every,’ or ‘any,’ we are making claims about groups of things. We might be saying that all of the things in the group have some property or that some of them do. The scope of these quantificational terms is a technical way of specifying the group of objects in question. To see that the scope of such terms is a component of conversational score, consider the following. Suppose that Kiya, while talking to her son about her shopping list, asks him whether there is any cream for her coffee. In this conversational context, the group of objects in question is the collection of things in their possession right now. Kiya is not asking whether there is any cream anywhere in the universe, and she is not asking whether there is any cream at the store in Memphis (where she normally buys it). She is asking, of the things in their possession right now, is any of it cream. As one can see then, the scope of quantificational terms is a component of conversational score.

Other components of conversational score include standards of accuracy, presupposition, and relevance. This list is not meant to be exhaustive. In fact, this notion of score is highly inclusive and includes, by definition, whatever is relevant to the assessment and proper development of the conversation. This is an

30. For complexities, see Bertrand Russell, *On Denoting*, 14 MIND 479 (1905), and Zoltán Szabó, *Descriptions and Uniqueness*, 101 PHIL. STUD. 29 (1999).

31. There are other ways to account for this. See, e.g., WILLIAM G. LYCAN, *PHILOSOPHY OF LANGUAGE: A CONTEMPORARY INTRODUCTION* 24–25 (1999).

32. Many theorists appeal to salience. See generally Lewis, *supra* note 28; HERBERT H. CLARK, *USING LANGUAGE* (1996).



especially inclusive notion of score, and not all ways of specifying the conversational score are as inclusive.<sup>33</sup>

### C. *Conversational Norm Enactment*

Since conversational contributions enact changes about what is true about that conversation, and since the conversational score tracks everything that is relevant to the proper development of a conversation, conversational contributions enact changes to the conversational score. Since conversations are g-norm-governed activities, and what counts as appropriate in any conversation depends in part on what has happened so far in that conversation, what counts as fair play in a conversation depends on its score. As a result, changing the score thereby changes the bounds of conversational permissibility. This means that any conversational contribution enacts changes to what is subsequently permissible in that very conversation.<sup>34</sup>

That was pretty quick and quite abstract. So, going through an example will help to illustrate this fairly hidden but no less real conversational norm-enacting phenomenon. Suppose that I am talking to my neighbor Bobby about his dog, Logan, and at a certain point in our conversation, I say, “Braun is still afraid to be left alone in the house, so we have to leave him outside.” By bringing up our dog, I enact a change to the score (regarding which dog is most salient in the conversation), and this affects the appropriate use of the expression ‘the dog.’ Now that Braun is the most salient dog, it is conversationally inappropriate to try to refer to any other dog with that expression. Thus, my conversational contribution enacted a conversational permissibility fact concerning the appropriate use of the expression ‘the dog.’ This conversational permissibility fact is a norm; it’s what I call an s-norm. S-norms are specific to particular conversations and are enacted by particular conversational contributions.

Several things are worth noticing about this for our purposes. First, enacting this conversational s-norm is not my aim; in fact, I’m hardly aware of doing it. Even though I am not consciously intending to enact this s-norm, I definitely am. Thus, verbally enacting a norm like this does not seem to require that the speaker

33. There are many ways to specify the score—in terms of what it does and does not track, in terms of how it works, and in terms of what it is ontologically. Conceptions of score that are narrower than mine include: ERNIE LEPORE & MATTHEW STONE, *IMAGINATION AND CONVENTION: DISTINGUISHING GRAMMAR AND INFERENCE IN LANGUAGE* (2015); Richmond H. Thomason, *Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics*, in *INTENTIONS IN COMMUNICATION* 325, 325–63 (Philip R. Cohen et al. eds., 1990); Elisabeth Camp, *Insinuation, Common Ground, and the Conversational Record*, in *NEW WORK ON SPEECH ACTS* 40–66 (Daniel Fogal et al. ed., 2018).

34. Elsewhere I call them conversational exercitives. See McGowan, *supra* note 25. See also Mary Kate McGowan, *Conversational Exercitives: Something Else We Do With Our Words*, 27 *LINGUISTICS & PHIL.* 93, 93–111 (2004); Mary Kate McGowan, *Oppressive Speech*, 87 *AUSTRALASIAN J. PHIL.* 389, 389–407 (2009); Mary Kate McGowan, *On ‘Whites Only’ Signs and Racist Hate Speech: Verbal Acts of Racial Discrimination*, in *SPEECH AND HARM* 121–47 (Ishani Maitra & Mary Kate McGowan eds., 2012); MARY KATE MCGOWAN, *JUST WORDS: ON SPEECH AND HIDDEN HARM* (2019). ‘Exercitive’ is Austin’s term for speech acts that enact facts about what is permissible or appropriate in some realm. See generally AUSTIN, *supra* note 1, at 151.



has a conscious intention to do so. This kind of verbal norm enactment appears to work quite differently.

Second, by enacting this conversational s-norm, I am not exercising any special authority in doing so. Instead, I am making an ordinary contribution to a conversation. Although one might think that enacting norms requires the exercise of authority, the routine way that conversational contributions enact s-norms tells against this. Verbally enacting s-norms for conversations appears to be a mere consequence of the g-norm-governed nature of conversation; it does not depend (as examples of verbal enactment in Part II.C appear to do) on a conscious intention to do so or on an exercise of speaker authority.

One final point before moving on. It is important to stress that there are two different sorts of norms at play. First, there are the norms that guide all conversations (e.g., norms of cooperation, syntax, and grammar). These are the general norms or g-norms. Second, there are the mini conversation-specific norms that are enacted by particular conversational contributions. These conversation-specific norms are s-norms. Conversational contributions routinely enact s-norms. My assertion about Braun, for instance, enacted an s-norm about the appropriate use of ‘the dog.’

In sum, when we speak, we are typically contributing to a conversation. And, because of the g-norm-governed nature of conversation, adding to a conversation also enacts s-norms that we don’t realize we are enacting, don’t consciously intend to enact, and don’t require any special exercise of authority in order to enact. This brand of verbal s-norm enactment is a ubiquitous feature of participating in the g-norm-governed activity of conversation.

## V. HOW IT GENERALIZES

This s-norm-enacting phenomenon generalizes. After all, many activities—not just conversations—are governed by g-norms. In fact, any move in a g-norm-governed activity thereby enacts s-norms for that token instance of that activity. Moreover, speech can be a move in these other g-norm-governed activities and, when it is, the utterance in question enacts s-norms in those activities too.

Recall that it doesn’t take much for an activity to be g-norm-governed.<sup>35</sup> It requires only that some actions count as inappropriate with respect to that activity. Conversations, dancing with a partner, informal social interactions, playing improvisational jazz, chess, checkers, and baseball are all g-norm-governed activities.

Any move in a g-norm-governed activity enacts s-norms for that activity. To see this, notice first that moves enact score changes. This is because moves are relevant to the proper development and assessment of that activity. And, since the score, by definition, tracks all such relevant factors, the score tracks all moves. Consequently, any move enacts a change to the score.

---

35. The norms in question need not be explicit, formal, exception-less, or even consciously recognized. I am also agnostic about their ontology.

Furthermore, because what is permissible at any point in a g-norm-governed activity depends on what has happened so far in that activity (along with the g-norms governing it), and because what has happened so far is captured by the score, enacting a change to the score thereby enacts s-norms for the activity in question. In this way, we see that a move in a g-norm-governed activity enacts s-norms for that activity.<sup>36</sup>

Examples abound. When my opponent moves her checker, she makes it permissible for me to move mine. As soon as Simsy mentions the big stain on the front of her blouse, it's ok for others to comment on it. Clearly, our actions are often contributions to g-norm-governed activities and, when they are, they enact changes to what is subsequently appropriate in those activities.

Here is another example. Suppose that Xian arrives at her high-school reunion and sees her old frenemy, Rory, for the first time in decades. (Rory and Xian are friendly despite disliking each other.) Back in the day, Xian grew tired of Rory's tendency to treat people like mere props in her never-ending personal drama, so Xian ended their friendship, and this split caused big problems for their friend group. Xian goes over to several people from that friend group and says, "I see Rory is over there. I haven't seen her in ages! She looks amazing; I can't wait to hear how she's been and give her a huge squeeze!" Xian's utterance is a move in the conversation; it thereby enacts s-norms for the conversation. In particular, it makes Rory the most salient woman in the context of this conversation, thereby enacting changes to the proper use of words like 'she.' In addition, however, what Xian says is also a move in the g-norm-governed activity of social interaction. By expressing joy at seeing Rory and an intention to hug her, Xian takes the high road and treats Rory well, even though she might well be warranted in resenting Rory. This grown-up and generous action sets the tone for the friend group and thereby has normative consequences. It encourages taking the high road, and it discourages pettiness. By treating Rory nicely, Xian affects what sorts of moves ought to be made by others in the friend group. Xian's action makes it less permissible for others to be uncharitable and petty (towards Rory or anyone else). Our speech actions contribute to broader (extra-conversational) social practices and thereby enact s-norms for those practices, too.

Again, just as with conversational contributions, our contributions to these broader (extra-conversational) social practices routinely enact s-norms in the token instances of those practices to which we contribute. This is so even though we are not aiming to do so, we are hardly aware of doing so, and we are not exercising any peculiar authority in doing so. The g-norm-governed nature of these practices makes it the case that contributions to them enact s-norms in them.

Granted, this has been a very quick and simplifying summary of a considerably more complex and nuanced phenomenon. That said, I hope it is clear that s-norm enactment is a ubiquitous feature of social practices. Since harm constitution

---

36. Although all moves enact score changes, not everything that enacts a score change is a move. See MCGOWAN, *JUST WORDS*, *supra* note 34, at 86–90.

requires norm enactment, highlighting a previously overlooked (and widespread) mechanism of s-norm enactment can bring further instances of harm constitution to light.

Let's now return to consider how these insights might apply to the utterance in *Public Harm*.

#### VI. PUBLIC HARM REVISITED

Recall that the speech act approach seemed unable to account for how the young man's utterance might constitute an act of subordination by enacting subordinating norms. This is because the felicity conditions for verbal norm enactment, at least according to the speech act framework, appear to be unmet in this case: the speaker does not have the authority required to enact norms, the speaker does not appear to be intending to enact norms, what the speaker says does not seem to express the content of a subordinating norm, and hearers do not take the young man to be enacting subordinating norms.

How might the phenomenon of s-norm enactment fare? A lot better than the speech act framework, in my estimation. When the young man says what he says, he enacts s-norms for any potential conversation that might ensue between him and the Asian woman. In particular, by asserting that he lost his job, the topic of job loss is made relevant, and this affects what sorts of responses are conversationally apt. Although conversational s-norms could, in principle, be harmful, the real action is extra-linguistic.

When the young man approaches the Asian woman and speaks to her, his action is a move in the g-norm-governed activity of social interaction.<sup>37</sup> Which particular s-norm is enacted by the young man's (speech) action will depend on three things: the score (for this particular social interaction) at the time of the utterance, the particular move he makes, and the g-norms governing social interactions.<sup>38</sup> Here are some reasons to believe that the young man's utterance might

---

37. We are now focusing on a fairly specific speech action: One might regard the young white man's action as so disruptive that it essentially terminates the social interaction and consequently cannot constitute harm as a contribution to that interaction. First, the young white man's action does not *terminate* the social interaction even if it shortens it (by preventing the Asian woman from engaging). Second, his action still enacts an s-norm and a harmful one at that. For example, imagine how unsafe that Asian woman might feel and how she might thereby be denied equal access to that public street. Elsewhere I have argued that similar speech actions constitute a hostile environment and ought to be civilly actionable as such. See generally McGowan, *JUST WORDS*, *supra* note 34; McGowan, *On 'Whites Only' Signs*, *supra* note 34, at 121–47. Third and finally, I operate with simplifying assumptions about the individuation of both types of activities and token instances of them. It is more complex, and the right way to individuate is interest-relative. I thank Brian Leiter for raising this set of concerns.

38. The social world is a messy place; g-norms are both complex and contested. We need social science to settle some of the relevant issues. What the g-norms are, what the harms are, how speech brings those harms about, and what the remedies ought to be are all questions that require multiple types of expertise to satisfactorily settle. As a philosopher of language, my main contribution is the identification of this hidden mechanism of norm enactment as well as establishing its relevance to these issues. In short, I am a firm believer in the division of intellectual labor. I thank Bill English for raising this question.

enact subordinating norms. First, the move that he makes mistreats the Asian woman, and it does so because she is Asian. Second, the social context in which this takes place, the contemporary United States, is one in which non-white persons are systematically disadvantaged because of their race. This means, among other things, that there are operative g-norms that socially disadvantage non-whites relative to whites. If this is correct, then we have some reason to believe that the young man's utterance in *Public Hate* enacts subordinating s-norms. In thinking through this, keep in mind that the g-norm-governed nature of social interaction is sufficient to show that the young man's utterance enacts s-norms in this case. The only remaining question is whether those s-norms are subordinating. Given the social context, the presence of subordinating g-norms, and the fact that the young man's utterance mistreats the Asian woman in virtue of her race, we have fairly good reason to believe that they are subordinating.<sup>39</sup>

According to this framework, s-norm enactment is ubiquitous and unavoidable in both conversation and social life. Whenever we do and say things, we are also enacting s-norms, both in the conversation to which we are contributing and in the particular social interactions in which we take part. S-norm enactment is an unavoidable aspect of social life. We cannot opt out. Whatever we do in the social world has normative consequences.<sup>40</sup>

On this picture, when ordinary (i.e., non-authoritative) utterances subordinate, they do so because of the subordinating g-norms operative in the broader culture. The power to subordinate is not located in the individual speaker, but in the social norms operative in the surrounding culture.

On this picture, the ordinary actions of ordinary people routinely enact the s-norms constitutive of the broader structure. This makes good on the insight that we are each complicit in these systems; it makes good on how the everyday actions of individuals constitutively contribute to the social constructions in our midst.

We began our discussion by considering the standard liberal line on hate speech which would have it that the young man in *Public Hate* is merely

---

39. That harms are constituted, as opposed to caused, is important for some legal strategies. See CATHARINE MACKINNON, *ONLY WORDS* (1993). Although I have elsewhere argued for the civil actionability of some ordinary hate speech, at the end of the day, it is unclear how much the cause-constitute distinction matters for hostile environment claims. See *id.*; Ishani Maitra, *Hateful Speech and Hostile Environments*, AUSTRALASIAN PHIL. REV. (forthcoming in 2022). The enacting of norms is crucially important, however, in making good on the claim (to which I am committed) that ordinary language use enacts and perpetuates social hierarchy. I thank Erin Miller for raising these issues.

40. S-norm enactment is ubiquitous but not all s-norms are harmful. Moreover, not all cases of harmful norm enactment are on a par. I am particularly interested in *group-based* harms that occur in social contexts in which that group is *systematically disadvantaged* relative to other groups. Although, I am not *here* focused on conversational harms, elsewhere I explore how this framework can account for the phenomenon of hepeating. See Mary Kate McGowan, *New Applications, Hepeating, and Discrimination: Response to Anderson, Horisk, and Watson*, 98 RES PHILOSOPHICA 537, 537–44 (2021); Claire Horisk, *Can McGowan Explain Hepeating?*, 98 RES PHILOSOPHICA 519, 519–27 (2021). I thank Brad Jackson for drawing attention to conversational harms and to how this framework might account for them.

expressing a political opinion. We now know better. Utterances do far more than merely express content. They also do far more than we either realize or intend. When we verbally disparage others, we do considerably more than express our thoughts. We enact norms. And when the social context is one in which that group is marginalized, we might well constitute harm.

## VII. CONCLUSION

I have here argued that in order to understand what hate speech does, we need to go beyond speech acts. First, I argued that this sneaky mechanism of norm enactment does a better job of explaining how non-authoritative hate speakers can verbally subordinate. But, even if one is not sympathetic to the idea that ordinary hate speakers can enact harmful norms, we nevertheless have plenty of reason to recognize the ubiquity of s-norm enactment. It is an inescapable aspect of social life. Whatever we do in the social world, we are changing norms. We might as well face it, understand it, co-opt it, and deliberately act in ways that will shape our normative surroundings to be the way we want them to be.