ABSTRACT

Title of Dissertation:	DISADVANTAGE IN CONTEXT: FROM MICROAGGRESSIONS TO HEALTHCARE POLICY
	Javiera Maximiliana Perez Gomez, Doctor of Philosophy, 2019
Dissertation directed by:	Professor Samuel J. Kerstein, Department of Philosophy

Many dimensions of applied ethics appeal to consequentialist moral theories to evaluate the moral permissibility of an action, practice, or policy. But such an approach risks obscuring other, non-consequentialist concerns. In line with this worry, this dissertation seeks to clarify and morally examine three phenomena that may compound the disadvantages that members of historically and currently disadvantaged groups face: microaggressions, the promotion of prenatal testing for selective abortion, and the allocation of scarce medical resources.

Chapter 1, “Disadvantage in Context,” describes the notion of disadvantage that is relevant to this dissertation and explains the relation between Chapters 2-4.

Chapter 2, “Microaggressions: What’s the Big Deal?” argues that the standard view of microaggressions, which holds that microaggressions are harmful because they express devaluing messages about members of disadvantaged groups, is too underdeveloped both for identifying microaggressions and for explaining why they

are morally objectionable. I then offer an improved account of microaggressions according to which it is the content of what is expressed that determines when microaggressions are morally objectionable.

Chapter 3, “When Is the Promotion of Prenatal Testing for Selective Abortion Wrong?” addresses the imprecisions of the *expressivist objection* to prenatal testing, which maintains that when medical professionals promote the use of prenatal testing for abortion on grounds of disability, they express a harmful, devaluing message to and about extant disabled people. I then offer an improved formulation of this objection according to which the promotion of prenatal testing for selective abortion is *sometimes* wrong.

Chapter 4, “Indirect Benefits and Double Jeopardy in the Allocation of Scarce, Lifesaving Resources,” examines the question of whether or not benefits to third parties, e.g., saving their lives or improving socioeconomic conditions, should count when resources are scarce and not all can be saved. By recruiting the notion of ‘double jeopardy,’ which, as I argue, can be understood in two distinct ways, I aim to give a stronger foundation for the idea that counting indirect benefits such as social contribution would be wrong—at least given certain social conditions.

DISADVANTAGE IN CONTEXT: FROM MICROGGRESSIONS TO HEALTHCARE POLICY

by

Javiera Maximiliana Perez Gomez

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2019

Advisory Committee:
  Professor Samuel J. Kerstein, Chair
  Associate Professor Dan Moller
  Associate Professor Rachel Singpurwalla
  Professor Peter Carruthers
  Professor Ruth Zambrana

Dedication

Para Carmen Gloria, Francisco, y Francisca.

Acknowledgements

I am most indebted to Samuel Kerstein, for his guidance at every stage of this project. Sam is a patient and generous advisor whose diligent comments on many, *many* drafts of each chapter have been central to my development as a scholar. I am also greatly indebted to Dan Moller, for invaluable advice that was often also necessary for this project to move forward. I am most grateful to Rachel Singpurwalla, for being an unwavering advocate since the day we met.

I am thankful to many others who provided various kinds of support during my time writing this dissertation, and during my time at Maryland more generally. Special thanks are owed to Peter Carruthers, Louise Gilman, and Theresa Lopez. I am also grateful to Julius Schönherr and Aiden Woodcock, who, through humor (if one can call it that) and rigorous conversation, shaped this project in many ways. For valuable feedback on earlier versions of these chapters, I am also thankful to Erick Ramirez, Jennifer Morton, members of the audience at my 2018 and 2019 Philosophy Colloquium presentations at the University of Maryland, and members of the audience at my 2018 Work-In-Progress presentation in the Bioethics Department at the National Institutes of Health.

I am fortunate to also have had worthwhile discussions with, and insightful guidance from, my best friend and little sister, Francisca Perez. My deepest gratitude is to Carmen Gloria Gomez Osorio and Francisco Javier Perez Orellana, whose devotion and loving encouragement are evident even in my earliest memories. It is their literal sweat and tears (of sadness and joy!) that made this milestone possible.

Table of Contents

Chapter 1: Disadvantage in Context

Disadvantage, it may be argued, refers to unfavorable conditions that make attaining a certain end, or good, more difficult than it would be in normal conditions. For example, if an individual’s end is living a good life but she struggles with poverty, lack of proper education, poor health, or discrimination, she would be at a disadvantage if normal conditions for living a good life exclude such factors.

In the contemporary U.S., the difficulty of attaining ends involving, e.g., education, earnings, health, among others, seems to vary depending on characteristics such as gender, race, ethnicity, ability, among others. For example, according to the U.S. Department of Labor, blacks and Hispanics fare worse than whites and Asians, and people with disabilities fare worse than people without disabilities, across all of the work force characteristics examined by the report: educational attainment, occupation, earnings, and employment status.[1] Women are also reported to fare poorly in terms of earnings when compared with men: e.g., in 2014, the median weekly earnings of women who worked full time was 83 percent of men’s median weekly earnings.[2] Similar disparities exist beyond the work force: e.g., in terms of life

---

[1] That blacks and Hispanics perform worse across these factors is a trend that has been recorded by the U.S. Department of Labor since at least 2007 (U.S. Bureau of Labor Statistics 2008).

[2] Although note that “[w]omen's earnings as a percentage of men's varied by occupation. Women's median usual weekly earnings in construction and extraction occupations ($691)

expectancy and other key health outcomes. For example, while cardiovascular disease is the leading cause of death in the U.S., "[n]on-Hispanic black adults are at least 50% more likely to die of heart disease or stroke prematurely (i.e., before age 75 years) than their non-Hispanic white counterparts" (U.S. Department of Health and Human Services, Centers for Disease Control and Prevention 2013).[3] Even infants who belong to certain social groups are at a disadvantage in terms of mortality rates: among non-Hispanic blacks, the mortality rate for infants is more than double that for non-Hispanic whites.

As these data suggest, the disadvantages that members of certain social groups face in the contemporary U.S. are not isolated; rather, they tend to cluster.[4] In other words, people who are at a disadvantage in terms of, e.g., socioeconomic status, are also at a disadvantage in terms of, e.g., educational attainment and health outcomes. Furthermore, these disadvantages are not ahistorical; rather, they would seem in part to be the product of a long history of legal (and other) mistreatment of members of the relevant social groups. For example, when women weren't legally able to vote or

---

were 91 percent of the earnings of their male counterparts" (U.S. Bureau of Labor Statistics 2016).

[3] Moreover, according the World Health Organization, today, lower socioeconomic status is linked to "poor health, more stress and lower self-confidence" (The World Health Organization 2017). Given these data, it would not be unreasonable to suppose that members of socio-economically disadvantaged groups, e.g., women, blacks, Hispanics, among others, may generally be at higher risk of illness than other groups.

[4] For a defense of the view that disadvantage clusters, see Wolff and de-Shalit (2007).

own property, or to accuse their husbands of rape—because rape was defined by the law in a way that "excluded a husband's rape of his wife from the prohibited conduct"—they were clearly at a legal disadvantage (Hellman 2008, 14). When blacks were legally forced into chattel slavery and when they were not legally allowed to vote or use the same public facilities as whites, they, too, were clearly at a legal disadvantage. When disabled people were not able to get to work, and when disabled children were not able to obtain a public education because appropriate public accommodations were not legally required, people with disabilities were also clearly at a legal disadvantage.

Talk of legal and historical disadvantage may seem overblown. In the contemporary U.S., segregation is illegal, women and blacks are able to vote, people with disabilities are legally owed accommodations, etc. In other words, the law has been challenged and reformed such that a wider subset of the population is now under its protection. This fact may lend itself to the idea that people who are at a disadvantage today—e.g., in terms of attaining ends involving education, earnings, or health—are morally responsible for the conditions that prevent them from attaining such ends. For instance, one might think that individuals may be responsible for not informing themselves about the benefits of education, or about ways to find employment or more affordable healthcare. But one need not discard the importance of individual responsibility to appreciate that the current political, economic, and social conditions of the U.S. do not function independently of prior states of affairs. As the above discussion suggests, some of the same groups that were at one point at a disadvantage with respect to the law remain at a disadvantage today, although today it is with respect to a different set of factors.

In light of these ongoing patterns of disadvantage, various policies and practices have been criticized for presenting further, though often subtle, problems for members of disadvantaged groups in their interactions with individuals as well as institutions. In three standalone papers, this dissertation examines three such practices: microaggressions, the promotion of prenatal testing by medical professionals, and the allocation of scarce medical resources by governmental institutions.

Microaggressions are said to be brief, commonplace, and often unintentional slights and insults toward members of disadvantaged groups. An example might be saying to an academic who is black “You are so articulate” which, according to the standard view of microaggressions, expresses the harmful message “It is unusual for someone of your race to be intelligent” (Sue, et al. 2007). Similarly, according to supporters of *the expressivist objection* to prenatal testing, promoting prenatal testing for abortion on grounds of disability expresses a harmful message to and about extant disabled people: e.g., that “Disabled people’s lives are not worth living” or that “Disabled people’s lives are not worthy of support.” Both microaggressions and the promotion of prenatal testing by medical professionals, then, have been said to be morally objectionable not because of what they do, but because of what they express.

But as I argue in Chapters 2 and 3, the arguments for the view that these acts are morally objectionable because of what they express have thus far not been compelling. In their place, I offer a framework for refining the morally relevant expressive nature of these acts and others. This framework originates in the idea that agents often communicate much more than they explicitly state; they often communicate implied meanings—or, to use the technical term: *implicatures*. Implicatures are meanings that are not explicitly stated but are nonetheless intentionally communicated. As I argue, however, agents also often *unintentionally* communicate—through the relevant actions—information about some of the attitudes or value judgments that they hold: what I call *hyper-implicature*. In my view, it's the *content* of an agent's implicature and hyper-implicature that determines the moral relevance of her action.

Chapter 2, "Microaggressions: What's the Big Deal?" thus has two main aims. First, I offer a stronger view of what microaggressions are. I argue that a microaggression occurs when an agent *hyper-implicates*—or, unintentionally reveals—an attitude or value judgment about another in virtue of the other's membership in some social group. Then, I examine what makes microaggressions wrong. In my view, the kinds of attitudes that agents hyper-implicate in a microaggression can be categorized into (at least) two types of attitude: those that are descriptive—e.g., "In the U.S., Latinos have a higher rate of unemployment than whites"—and those that are moral or evaluative—e.g., "Blacks are morally inferior to whites." Only the latter, I argue, obviously involve a *moral error* on the part of the agent: an error that amounts to a failure to recognize the moral standing of another person by, e.g., giving inappropriate weight to the fact that she is a member of some social group.

Chapter 3, “When Is the Promotion of Prenatal Testing for Selective Abortion Wrong?” offers an improved formulation of the expressivist objection to prenatal testing: what I call *the expressivist objection**. In my view, when medical professionals engage in the action of promoting prenatal testing for selective abortion, they sometimes express implied messages that can be about extant disabled people. Moreover, sometimes expressing these messages amounts to the medical professional’s violating a duty to obtain informed consent from her patient to run tests that might present the patient with the choice of terminating the pregnancy on grounds of disability. If, for example, a medical professional expresses to her patient that “Adults with Down syndrome have no independence,” she would seem to express an attitude that is inaccurate or misleading, as research suggests that adults with Down syndrome *can* lead independent lives so long as they have the appropriate resources. Thus, insofar as this medical professional would communicate inaccurate or misleading information about, e.g., the reasons one might have for considering prenatal testing for selective abortion, to her patient, she would seem to violate a duty to obtain informed consent to run such tests and would thereby act wrongly.

Chapter 4, “Indirect Benefits and Double Jeopardy in the Allocation of Scarce, Lifesaving Resources,” turns to the policy question of how scarce medical resources, such as kidneys, should be allocated. Bioethicists generally assume that it would be wrong for a scarce resource allocation scheme to count indirect benefits such as social contribution. But the arguments for this view are surprisingly weak. By recruiting the notion of ‘double jeopardy,’ I seek to offer a stronger foundation for the view that an allocation policy that counted indirect benefits such as social contribution would be morally objectionable. As I argue, ‘double jeopardy’ can be understood in a weak sense, according to which an allocation scheme may bring about injustice, and in a strong sense, according to which an allocation scheme may compound injustice. In the contemporary U.S., for example, a policy that counted indirect benefits, e.g., in terms of economic contribution, would likely put black Americans in strong double jeopardy. This, I argue, would seem to be worse than if such a policy were to put black Americans in weak double jeopardy because the allocating entity would seem to be the same entity that brought about the original injustice on which the second injustice is based. This view would seem to have implications for how we think of other allocation schemes, including those that make assumptions about the value of life with disability.

Each of these chapters is written so that it is independent of all the others. The chapters can thus be read in any order, as they do not presuppose one another.

Chapter 2: Microaggressions: What’s the Big Deal?

Introduction

‘Microaggressions’ are said to involve brief and commonplace verbal, behavioral, or environmental indignities, whether intentional or unintentional, that communicate slights and insults toward members of disadvantaged groups (Sue, et al. 2007). Harvard psychiatrist Chester Pierce coined the term in 1970 to denote minor slights and insults directed at racial minorities; but today the term is used to denote slights and insults directed at members of various other disadvantaged groups, including women; individuals who are lesbian, gay, bisexual, transgender or transsexual; Asian-Americans; Hispanic-Americans; Muslim-Americans; as well as individuals who are obese (Lilienfeld 2017, Sue 2017). This latter use of the term originated in a 2007 paper by psychologist Derald Wing Sue and his coauthors, who are now taken to be the pioneers of the concept of microaggressions. In their view microaggressions are harmful because they send devaluing messages to individuals who are members of disadvantaged groups.

Sue et al.’s view has been widely influential yet quite controversial. Numerous businesses and college campuses have jumped on board what is sometimes called the ‘microaggression research program,’ or simply the ‘microaggression program’: real-world efforts aimed at counteracting microaggressions. Yet the program has been criticized incessantly, primarily on the basis of the negative consequences it could bring about. For instance, critics worry that it may lead to unwarranted restrictions on free speech,[5] worsen racial tensions,[6] and help create a culture of victimhood[7]—not only in college campuses but in society at large.[8]

While both the concept of microaggressions and the microaggression program have received quite a bit of attention from psychologists and from the media, philosophers have been almost entirely absent from the debate.[9] Admittedly, Regina Rini defends the moral importance of microaggressions in a *Los Angeles Times* op-ed, arguing that we should not conceive of microaggressions as part of a new "culture of

---

[5] See, for example, Lukianoff and Haidt (2015).

[6] See, for example, Haidt and Jussim (2016).

[7] See, for example, Campbell and Manning (2014).

[8] Interestingly, a parallel could be drawn between the debate on microaggressions and that of political correctness, as understood by Dan Moller (2016). I won't explore this point further, but as Moller writes, "…political correctness *is* distinctive, and a distinctively left phenomenon… Those attempting to shape discourse on the right are rarely moved by feeling sorry for some group and rarely make corresponding objections focused on avoiding offense" (2016, 92).

[9] Some psychologists have suggested that microaggressions may contribute to psychological disturbances such as anxiety and depression (Sue et al. 2007, 279). It's worth noting, though, that there is no evidence yet that the purported negative effects on individuals' mental states or mental health are actual. This lack of evidence suggests that while it may be quite intuitive to think that experiences of microaggressions may potentially damage a minority individual's mental health, to claim that microaggressions have such *causal* power has yet to be explored. For more on this point, see Lilienfeld (2017, 142).

victimhood,"[10] as critics suggest, but as signs of the work that needs to be done in order to attain a culture of solidarity: "a culture in which no one is denied full moral recognition" (Rini 2015).[11] Moreover, Emily McTernan (2017) argues that microaggressions constitute "a distinctive form of injustice," as they "form a social practice that contributes to structures of oppression and marginalization." Nevertheless, philosophical attention has been lacking both with regard to the coherence of the concept of microaggressions, as well as with regard to non-consequentialist moral concerns that microaggressions may give rise to. For example, as I will argue, microaggressions may reveal a moral failure in the *agent* who engages in the microaggression, independently of the consequences of his action, but the debate has focused only on what microaggressions mean for their victims. This paper is a call for political and moral philosophers to engage in the microaggressions debate, for, as will become clear throughout the paper, microaggressions involve a social practice that can be morally objectionable in a multitude of ways.

The aim of this paper is twofold. First, I argue that insofar as Sue et al.'s notion of microaggressions implicitly appeals to subjective and utilitarian principles, it has hurt rather than helped the microaggression program. More specifically, insofar as it holds that microaggressions lie in the eye of the beholder and are by definition harmful, it offers little guidance for preventing microaggressions. If the concept of

---

[10] See, for example, Campbell and Manning (2014).

[11] Rini is also working on a book, *The Ethics of Microaggressions*, forthcoming, 2019. See also Tschaepe (2016), who argues that "[m]icroaggressions cause epistemic injustice and prevent human flourishing."

microaggressions is going to be useful for the microaggression program, I suggest, then it must distance itself from subjective and utilitarian principles and endorse, instead, both a more objective framework to identify when a microaggression has occurred, and a broader moral framework to determine when a microaggression is morally objectionable.

My second aim is to offer some grounding ideas for a more useful concept of microaggressions. I argue that microaggressions belong to a class of actions that communicate more than is explicitly stated or intentionally implied; they communicate certain kinds of beliefs or attitudes that the agent of a microaggression seemingly holds.[12] I refer to such beliefs or attitudes as *hyper-implicatures*. Specifically, I argue that these hyper-implicatures are about some social group to which the agent of a microaggression takes the interpreter of the microaggression to belong to. On this view, it is the *content* of a hyper-implicature that is subject to moral evaluation. In some cases, an agent who engages in a microaggression may hyper-implicate a *descriptive* attitude about another agent that is inaccurate or misleading: e.g., "Latin Americans like spicy food." In other cases, an agent who engages in a microaggression may hyper-implicate an *evaluative*, or *moral*, attitude about another that is *demeaning*: e.g., "Blacks are morally inferior to whites." A 'demeaning attitude,' as I define it here, involves a failure on the part of the agent to

---

[12] I say "seemingly" to allow for cases in which an agent acts in a way that communicates a collective attitude without actually holding that attitude, as well as cases in which the agent does hold the attitude she communicates in her action. I elaborate on this distinction a bit in footnote 36, but more in-depth analysis is a matter for another paper.

regard another appropriately.[13] The idea that I defend is that while hyper-implicating both descriptive and demeaning attitudes can raise moral concerns (e.g., if it foreseeably causes indignation, irritation, offense, etc., in the interpreter of the action), hyper-implicating a *demeaning* attitude would seem to contain a moral error and would thereby be morally objectionable even if no one was caused any pain.

The paper proceeds as follows. In Section 1, I explain Sue et al.'s notion of microaggressions in more detail. In Section 2, I elucidate two of its major shortcomings: it leads to counterintuitive claims because it is unnecessarily subjective with respect to how to identify a microaggression and unnecessarily utilitarian with respect to what makes microaggressions morally objectionable. In Section 3, I appeal to some concepts from the philosophy of language to develop a more objective framework for identifying when a microaggression has occurred. In Section 4, I develop that framework by considering the claim that it is the content of the attitudes that are hyper-implicated that makes a microaggression morally objectionable. Finally, Section 5 discusses some implications of the framework I propose, including the noteworthy implication that microaggressions are a kind of phenomenon whose target are *not* just members of disadvantaged groups.

---

[13] Indeed, as I will argue, demeaning attitudes may illustrate a failure of what Stephen Darwall calls *recognition respect*, such that if an agent hyper-implicates a demeaning attitude, she reveals a kind of morally objectionable error that is independent of the consequences that hyper-implicating that attitude may or may not bring about.

1. *Sue et al.'s view of microaggressions*

In their seminal paper on microaggressions, Sue et al. are primarily concerned with microaggressions that occur in the context of clinical practice. Their worry is that white therapists may commit microaggressions against their non-white patients, thereby preventing successful therapeutic relationships with them.[14] But the examples they offer involve a broader social context. For example, Sue et al. maintain that:

---

[14] For this reason, Sue et al. are careful to taxonomize microaggressions into three categories: microassaults, microinsults, and microinvalidations. I will not address these specific categories in this paper, but here is how Sue et al. define these terms (Sue et al. at p. 274):

> *Microassault*: "an explicit racial derogation characterized primarily by a verbal or nonverbal attack meant to hurt the intended victim through name-calling, avoidant behavior, or purposeful discriminatory actions. Referring to someone as 'colored' or 'Oriental,' using racial epithets, discouraging interracial interactions, deliberately serving a White patron before someone of color, and displaying a swastika are examples. Microassaults are most similar to what has been called "old fashioned" racism conducted on an individual level."
>
> *Microinsult*: "is characterized by communications that convey rudeness and insensitivity and demean a person's racial heritage or identity. Microinsults represent subtle snubs, frequently unknown to the perpetrator, but clearly convey a hidden insulting message to the recipient of color."
>
> *Microinvalidations*: "are characterized by communications that exclude, negate, or nullify the psychological thoughts, feelings, or experiential reality of a person of color. When Asian Americans (born and raised in the United States) are complimented for speaking good English or are repeatedly asked where they were

1. Asking a minority-individual “Where are you from?” “Where were you born?” or telling her “Your English is great” sends the message “You are not American” or “You are foreign.”
2. Saying “Everyone can succeed in this society if they work hard enough” sends the message “People of color are lazy and/or incompetent and need to work harder.”
3. Clutching on to one’s purse and crossing the street at the sight of a person of color approaching sends the message “You are a criminal” or “You are going to steal.”
4. Mistaking a person of color for a service worker sends the message “People of color are servants to whites. They couldn’t possibly occupy high-status positions.”
5. The overabundance of liquor stores in communities of color sends the message “People of color are deviant.” (Sue et al. 2007, 276-277)

Straightaway, one may be hesitant to accept that these actions actually express the messages that Sue et al. claim they express, as at least some of these messages

---

> born, the effect is to negate their U.S. American heritage and to convey that they are perpetual foreigners.”

Note that while all of these terms have some overlap, the term ‘microinsult’ seems to be closest to the phenomenon I am concerned with in this paper. ‘Microassaults’ seem to be more clearly a case of *overt* discrimination, and ‘microinvalidations’ may be a subset of microinsults that may be more adequately described as instances of what Miranda Fricker (2007) calls ‘epistemic injustice.’

seem hyperbolic or groundless.[15] One may also object to such a heterogeneous set of examples, as it may evoke a sense that too much may count as a microaggression in Sue et al.'s view. I will return to these concerns in Section 2. For now, let us grant that these examples do constitute microaggressions. In other words, let us grant that the actions described in each example do express the hurtful messages that Sue et al. claim they express. What I want to highlight is that the main characteristic of Sue et al.'s view is that it is entirely victim-focused—both in terms of identifying a microaggression and in terms of explaining why it is morally objectionable.

Consider the following passage:

> When a White employer tells a prospective candidate of color 'I believe the most qualified person should get the job, regardless of race' or when an employee of color is asked 'How did you get your job?', the underlying message from the perspective of the recipient may be twofold: (a) People of color are not qualified, and (b) as a minority group member, you must have obtained the position through some affirmative action or quota program and not because of ability. (Sue et al. 2007, 274)

In this passage, Sue et al. suggest that the messages that a microaggression expresses will depend, in part, on the context in which the microaggression occurs. The relevant contextual factors seem to include the group membership of the agent of the microaggression, the group membership of the 'victim' of the microaggression, and the conversational topic between them: e.g., one about race, merit, job

---

[15] For example, it's not clear how (5) constitutes an action.

appointments, etc.[16] Nevertheless, the passage also suggests that despite these contextual factors, microaggressions ultimately "lie in the eye of the beholder"[17]: the message that the white employer's comment expresses just is what the prospective employee of color takes it to be. Indeed, Sue et al. later write that "the [victim of the microaggression] must determine whether a microaggression has occurred" (Sue et al. 2007, 279). In other words, microaggressions involve comments that express certain messages *from the perspective of the microaggressee*. Sue et al.'s view thus appeals to the subjective judgment of the victim in order to determine whether a microaggression has occurred. That is, Sue et al.'s view is subjective in nature. The question "How did you get the job?" *becomes* a microaggression only if the recipient of the question judges that such a question sends a harmful underlying message.

Moreover, for Sue et al. microaggressions are harmful by definition: they send underlying *harmful* messages from the perspective of the microaggressee. What makes these messages harmful is that when communicated they cause various kinds of suffering in victims of microaggressions. For example, imagine that you are a person of color and a white person says to you something that suggests that "People of color are not qualified" for a certain job or that minority-individuals obtain jobs "through some affirmative action or quota program and not because of ability." Such an action would naturally cause indignation or offense (at the very least), for it says

---

[16] Interestingly, these criteria have been endorsed by the University of California Santa Cruz, which has published a list of microaggressions that are not be permitted on campus, including the statement "America is the Land of Opportunity." See Williams (2018).

[17] I borrow this terminology from Lilienfeld (2017, 141).

something negative about a social group to which you are perceived to belong to or to which you take yourself to belong to—or so Sue et al. suggest.

But Sue et al. are concerned with more than just indignation or offense. Consider asking someone who ‘looks’ Asian, Hispanic, Muslim, etc., “Where are you from?” or “Where are you *really* from?” According to Sue et al., such questions, when directed at a member of a disadvantaged social group, communicate that he does not belong to some dominant social group—e.g., with regard to race, ethnicity, religious views, etc. In the clinical context, Sue et al. suggest, the mere highlighting of this fact may be enough to cause in patients a feeling of being misunderstood, which may result in a kind of reluctance to form a meaningful relationship with their therapist. In the broader social context, the highlighting of an individual’s lack of membership in some dominant group may cause a more general feeling of exclusion, perhaps even alienation from one’s society.

Furthermore, Sue et al. indicate, microaggressions may put microaggressees in something of an internal double-bind. For example, if a microaggressee points out a microaggression, she risks being perceived as overly sensitive or paranoid, but if she does not point out the microaggression, she risks tacitly endorsing the continuance of such a microaggression. Such a predicament can, in turn, cause feelings of paranoia (“What if I’m exaggerating?”) or of being gaslighted (“She can’t possibly be serious; she must be messing with me”).[18]

---

[18] For example, a member of a minority group who is given the intended compliment “You are so articulate” may wonder whether this is a microaggression: did the complimenter say this to her because she is “so articulate” *tout court* or because she is “so articulate” *for a*

The takeaway here is that on Sue et al.'s view microaggressions are morally objectionable because, by expressing harmful messages, they cause various kinds of negative mental states in microaggressees, including indignation, offense, alienation, paranoia, anxiety, among others. We can thus think of Sue et al.'s view as sharing a feature with utilitarianism, as it locates the harm of microaggressions in the negative mental states that victims of microaggressions experience as a result of the messages that microaggressions send.

In short, because on Sue et al.'s view a microaggression occurs if the 'victim' of some action has been caused a negative mental state by the messages he takes the action to express, their view is entirely victim-focused—both in terms of how to determine when a microaggression has occurred (it is up to the victim of the microaggression), and in terms of what makes a microaggression morally problematic (it causes negative mental states in the victim of the microaggression).

2. Some shortcomings of Sue et al.'s view

As I mentioned earlier, Sue et al.'s view has served as the foundation for the 'microaggression program': efforts adopted by businesses and college campuses

---

minority individual? As Lilienfeld reports, a shortcoming of the concept of microaggressions is that they "are often or usually extremely ambiguous in nature, rendering it difficult or even impossible to ascertain whether they have actually occurred" (Lilienfeld 2017). Nevertheless, supporters of Sue et al.'s view, and McTernan, see this as one of the defining characteristics of microaggressions. In my view, however, this is not a defining feature of microaggressions, even if they are often perceived this way by minority-individuals.

aimed at counteracting microaggressions. In the business world, companies such as Airbnb, Coca-Cola, Facebook, and, most recently, Starbucks have provided training aimed at preventing implicitly prejudicial or otherwise offensive comments and actions in the workplace (Fisher 2015, Scheiber and Abrams 2018). In the academic world, colleges and universities have distributed lists of microaggressions[19] and sponsored round table discussions[20] presumably to inform potential microaggressors about what kind of language is offensive to individuals who belong to historically or currently disadvantaged groups (hereafter: disadvantaged groups). Moreover, colleges and universities have also attempted to show support to "victims" of microaggressions by encouraging them to share their experiences on "Walls of

---

[19] For example, the University of California, Berkeley has distributed lists of microaggressions to encourage faculty and students not to engage in behaviors that may be offensive to minorities (Lilienfeld 2017, 139). See also footnote 16 above.

[20] For example, the University of Maryland, College Park has sponsored "round table discussions led by trained facilitators about the impact of microaggressions on the everyday lives of Maryland's students of color": https://dsa.umd.edu/events/itooammaryland-responses-to-racial-microaggressions.

Expression" distributed across campus,[21] or by using certain 'hashtags' on social media.[22]

The microaggression program thus rests on Sue et al.'s idea that microaggressions psychologically harm minority individuals on a regular basis and in covert ways—e.g., by being ambiguous or by putting minority individuals in an internal double bind—and that these are harms that must be prevented. This is certainly a noble cause; but, as I argue in this section, Sue et al.'s notion of microaggression is currently too underdeveloped for that cause. In particular, because Sue et al.'s view is entirely victim-focused, it has had the effect of hurting rather than helping the microaggression program, as it offers little guidance both for preventing microaggressions and for explaining when they are morally objectionable.[23]

---

[21] The University of South Florida has tried to combat microaggressions via cross-campus "Walls of Expression," on which students could write "any word that struck them as offensive," no matter "if the word had been used to describe the student, if the student used it before or if the person simply viewed it as commonly out-of-context": http://news.usf.edu/article/templates/?z=38&a=5326. Presumably, words said commonly out-of-context involved words that were insulting and irrelevant to a particular conversation.

[22] For example, the University of Maryland has encouraged students to bring attention to the microaggressions they experience by using the hashtag "#ITooAmMaryland": https://dsa.umd.edu/events/itooammaryland-responses-to-racial-microaggressions.

[23] For various other potential shortcomings of the concept of microaggressions, see Lilienfeld (2017).

As I mentioned in Section 1, on Sue et al.'s view, microaggressions express harmful messages from the perspective of minority individuals. One central objection to this claim is that Sue et al. offer no support for the idea that microaggressions actually express such harmful messages (Lilienfeld 2017). This objection has two interpretations. One is that there is no empirical evidence that microaggressions actually cause the negative mental states that Sue et al. claim they cause. As Scott O. Lilienfeld (2017) argues, it may be that the negative mental states that victims of microaggressions report are the result of certain personality traits that make some individuals hypersensitive, paranoid, or neurotic. Specifically, as Lilienfeld highlights, some individuals may have what is known as the personality trait of 'negative emotionality': "a propensity toward depression and anxiety, and a tendency to react to stressful situations with unpleasant emotions."[24] Thus, Lilienfeld argues, it may be that some people are so sensitive that they may interpret interactions that need not be negative negatively. Without empirical data for the claim that microaggressions *actually* cause the harmed mental states that Sue et al. claim they cause, we can't be sure that the messages expressed by microaggressions are actually harmful—or so Lilienfeld suggests.

Sue has responded to this point that in applying the principle of skepticism to the study of microaggressions Lilienfeld negates the lived experiences of individuals who anecdotally report suffering as a result of microaggressions (Sue 2017, 171). I agree with Sue that there is plenty of anecdotal evidence of what he calls

---

[24] https://www.eurekalert.org/pub_releases/2001-09/ace-he091001.php

"microaggressive suffering" so I will set this worry aside.[25] But, I want to restate, even if there is anecdotal evidence of such suffering, Sue et al.'s view currently offers little guidance for deciphering the content of the messages that microaggressions purportedly express. Specifically, they offer no support for claiming, for example, that mistaking a person of color for a service worker actually sends the message "People of color are servants to whites. They couldn't possibly occupy high-status positions." As Lilienfeld puts it, "there is no research evidence that the microaggressions identified by Sue et al. are linked, either probabilistically or inexorably, to these negative messages, as there are no data on what proportions of minority individuals interpret each microaggression in accord with the purported message" (Lilienfeld 2017, 146). In Section 3 I will offer a way to avoid this lack-of-support objection. For now, let me elucidate the scope of this objection.

The objection involves two worries. The first is empirical: there is no evidence that minority individuals actually interpret microaggressions to express the specific messages that Sue et al. claim they express. Now, this worry again fails to account for the vast testimonial evidence of experiences of microaggressions, so I will set it aside. The second worry, however, is theoretical: even if it turns out that minority individuals do interpret microaggressions to express the messages that Sue et al. claim they express, Sue et al. offer no *theoretical* background for justifying such

---

[25] A simple YouTube search for 'things that white people say to <insert some minority-group>' reveals that the term has proved useful for many individuals who belong to disadvantaged groups to pick out all-too-familiar slights and insults that occur in their everyday interactions with individuals who belong to some dominant group.

interpretations. That is so say, minority individuals may just be wrong in their interpretations of purported microaggressions.

But herein lies the strongest case against Sue et al.'s notion of microaggressions: it is too underdeveloped. Consider again the case of mistaking a person of color for a service worker. The objection is that Sue et al. offer no theoretical foundation for favoring the interpretation that this action expresses "People of color are servants to whites. They couldn't possibly occupy high-status positions" over, e.g., "The color of your shirt, not the color of your skin, suggested to me that you were a service worker." Or, consider the case of asking a minority-individual "Where were you born?" or telling her "Your English is great." According to Sue et al., this kind of action sends the (offensive) message "You are not American" or "You are foreign"; but they offer no support for that interpretation over an interpretation such as "Where in the U.S. were you born?"

The point here is this. An individual who chooses the interpretation that involves a message of disvalue over an alternative, more charitable interpretation could just be wrong, and Sue et al. offer no theoretical support for deeming the victim of the microaggression justified in interpreting an action as expressing a message of disvalue. Without a successful reply to this worry, one may be tempted to agree with Lilienfeld's call to abandon the microaggression program. After all, if there is no theoretical support in favor of microaggressees' interpretations of the messages that microaggressions purportedly express, it is unclear what the microaggression program is engaged in except for the prevention of free speech just in case someone, somewhere, might take offense to one's words—or so one might reason.

Sue et al. may respond in one of two ways. They may appeal to the idea that whites and nonwhites perceive racial relations in very different ways and claim that for this reason there is no fact of the matter regarding what microaggressions actually express. As Sue et al. put it,

> White Americans tend to believe that minorities are doing better in life, that discrimination is on the decline, that racism is no longer a significant factor in the lives of people of color, and that equality has been achieved. More important, the majority of Whites do not view themselves as racist or capable of racist behavior. (Sue et al. 2007, 277)

Sue et al.'s suggestion is that an individual's conception of race can radically affect the way she perceives the world, such that while an action may seem morally neutral to a white person who does not experience racism or discrimination on a regular basis, that same action may be perceived as an obvious instance of racism or discrimination to a person of color. But while this may be true, what Sue et al. seem to be suggesting is that there is no fact of the matter as to what is or is not racist—a point that does not seem viable. In other words, it seems wrong to say that there is no fact of the matter of whether or not an action is an instance of racism or discrimination, and to suggest that there is no fact of the matter of what an action such microaggression expresses. Surely there is some set of criteria that one might appeal to in order to determine whether an action is an instance of racism or discrimination. Similarly for what it is that a microaggression might express—a point I return to in Section 3.

Alternatively, Sue et al. may reply that microaggressions express the messages that Sue et al. claim they express only in certain contexts—contexts involving, e.g., a certain conversational topic and certain group memberships among the participants of the action. Specifically, they may suggest that some of the necessary conditions for a microaggression to occur may be that the microaggressor be a member of some dominant group and that the microaggressee be a member of some minority-group. But even if they did say this, such a set of criteria would seem to be much too vague for choosing one interpretation of a microaggression over another. Nevertheless, both of these responses that Sue et al. may offer may be supported by an appeal to the idea that microaggressions ultimately lie in the eye of the beholder: that it will ultimately be up to the victim of the microaggression to determine whether or not a microaggression occurred.

But I believe that this response to the criticism that Sue et al. offer little theoretical support in favor of choosing one interpretation of what a microaggression expresses over another highlights one of the central shortcomings of Sue et al.'s view. In claiming that microaggressions lie in the eye of the beholder, Sue et al. focus on the point of view of the victim of a microaggression to a fault. Neither the intention of the agent nor any objective criteria for an action's counting as a microaggression plays a role. The problem with this entirely subjective notion of microaggression is that it over- and under-extends—that is, too little and too much turns out to be a microaggression on Sue et al.'s view.[26]

---

[26] Lilienfeld (2017) raises this point.

Suppose a black student constantly raises her hand in response to her teacher's questions, but her teacher constantly calls on a white student first, partly because the teacher believes, explicitly or implicitly, that the black student, because she is black and likely less competent, will derail the class discussion. Now, while the teacher's action may express a message of disvalue, if the black student does not experience the action as a slight or insult, then the action is not a microaggression—or so Sue et al. would seem to suggest. But this conclusion seems counterintuitive. Many of us would agree that even though the student does not register the teacher's action as a microaggression, there is a sense in which something microaggression-like did occur. In a world in which for one reason or another no one realizes that such slights are occurring, microaggressions wouldn't exist on Sue et al.'s view. Their view thus captures too little as a microaggression.

Now, imagine that a socially conscientious teacher develops the following policy to avoid being a microaggressor: she will be careful not to automatically call on her white students first when she poses a question to her class. Suppose, further, that today she calls on a white student first because she does not see a nonwhite student raise his hand before the white student. On Sue et al.'s view, if the nonwhite student experiences the teacher's action as expressing a slight or insult, the action would indeed be a microaggression—despite the teacher's explicit intention not to engage in microaggressions. But, many of us would agree, that would be unfair to the teacher; agents don't need to be infallible in terms of foreseeing others' interpretations of what their actions express. Thus, the conclusion that the teacher's action is a microaggression is untenable. It entails that nearly any action could count as a microaggression.[27] Because Sue et al.'s view is strictly subjective, then, in addition to capturing too little, it captures too much as a microaggression.

Now, since the microaggression program appeals to Sue et al.'s notion of microaggressions, and since one of the main aims of the microaggression program is to prevent the occurrence of microaggressions, Sue et al.'s entirely subjective notion of microaggressions is unhelpful to that end, as it offers little guidance with regard to preventing microaggressions. Admittedly, Sue et al.'s notion of microaggressions does suggest that one should avoid actions that might be interpreted to express offensive messages to individuals who belong to disadvantaged groups; but it leaves the interpretation of the action entirely up to the victim of the purported

---

[27] Lilienfeld tries to illustrate this worry by citing an example that "both (a) ignoring and (b) attending to minority students in classrooms have been deemed to be microaggressions by some authors": some see ignoring an Asian-American student who raises her hand as a microaggression, whereas others see complimenting a minority student as a microaggression. Lilienfeld worries that such variance in opinion leaves potential microaggressors—in this case, teachers—in a sort of double bind: "If they ignore minority students' raised hands, they risk being accused of implicit prejudice; conversely, if they call on students and compliment them, they risk the same accusation" (Lilienfeld 2017, 150). Note, though, that it's a mistake to think that these are the only two options for teachers. Teachers could try to avoid both ignoring and complimenting their minority students. Moreover, a teacher determined to avoid microaggressions could aim to compliment a student not because he is or isn't a minority-student, but because he deserves a compliment *tout court*.

microaggression. This means that whether or not an agent engages in a microaggression may in many cases be out of her hands.

Now, recall that according to Sue et al. microaggressions are harmful by definition: they express messages that cause negative mental states in individuals who belong to disadvantaged groups. Thus, in addition to appealing to subjective principles, Sue et al.'s view also shares an important feature with utilitarianism: it suggests that because the messages expressed by microaggressions (tend to) cause suffering rather than happiness, they make the world a worse place.[28]

But this point suggests another shortcoming of Sue et al.'s view, for such a basis for the harm in microaggressions invites at least two serious worries. Suppose an individual who is tired of feeling hurt—e.g., slighted, offended, displaced, alienated, gaslighted, etc.—by the microaggressions he encounters left and right decides to adopt a stoic attitude towards such actions. Sue et al.'s view entails that if this individual does not feel hurt by the action, then the action isn't harmful (and, moreover, that it isn't a microaggression). But that again seems wrong. Many of us would agree that it would be much more appropriate to say that a morally problematic microaggression did occur, even though it did not harm the stoic individual.

Now, insofar as Sue et al.'s view attributes the harm in microaggressions entirely to the negative mental states that microaggressions, in virtue of expressing

---

[28] As Mill put it, "Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure" (Mill 1861).

harmful messages, cause, it invites a popular objection to the microaggression program: that instead of encouraging individuals who belong to disadvantaged groups to look for microaggressions left and right, the microaggression program should train such individuals to be more resilient to these supposed slights and insults.[29] I believe that this objection is misguided, but, as will become clear in the next section, it is avoidable if Sue et al. (a) move away from the idea that microaggressions are harmful by definition, and (b) broaden their understanding of what makes microaggressions morally objectionable. In short, the utilitarian feature of Sue et al.'s view captures only some of what may be objectionable about microaggressions, and thereby makes the microaggression program seem wrongheaded.

To sum up this section: insofar as Sue et al.'s view is implicitly subjective, it captures too little and too much as a microaggression. Moreover, insofar as it is utilitarian in nature, it fails to account for intuitive cases of morally objectionable microaggressions and invites an undesirable yet avoidable objection: that the microaggression program should train minority individuals to be less sensitive. For these reasons, Sue et al.'s entirely victim-focused notion of microaggressions offers little guidance for avoiding microaggressions, and thus fails to provide the support that the microaggression program needs in order to be successful.

If the concept of microaggressions is going to be taken seriously and if the microaggression program is going to have a defensible foundation, then we need to

---

[29] See, for example, Campbell and Manning (2014), Lukianoff and Haidt (2015), Haidt (2017), and Skenazy and Haidt (2017).

refine Sue et al.'s conception of microaggressions so that it is neither entirely subjective nor strictly utilitarian. In the rest of the paper, I offer some grounding ideas for such a project. Before doing so, however, let me stress that in highlighting that the concept of microaggressions must not be entirely subjective I do not intend to downplay the opinions or feelings of microaggressees. A more objective framework for identifying microaggressions is not committed to dismissing people's wounded feelings. Indeed, that many people agree that some microaggressions are hurtful is one central reason why the phenomenon is in need of further scrutiny. Rather, in offering a more objective framework for identifying microaggressions, I want to highlight that microaggressions are a distinct kind of phenomenon and that there is more to what makes them morally objectionable than has been entertained in the debate on microaggressions thus far.

3. *Refining the concept of microaggressions: How to identify a microaggression*

I mentioned in Section 2 that Sue et al. offer no theoretical support for their interpretations of the messages that they claim microaggressions express. In this section, I argue that there are concepts in the philosophy of language that we can use to try to decipher what it is that microaggressions might express. Typically, when philosophers of language talk about meaning, they talk about two kinds of meaning: literal meaning and implied meaning. What I want to suggest is that given the vagueness of Sue et al.'s view, it's easy to equate the 'messages' that they claim microaggressions *express* with the implied meanings that a speaker communicates in an utterance. But, as I will argue, this is a mistake, for microaggressions are a kind of action that communicates—or, more precisely, *reveals*—more than literal and implied meanings: it reveals some of the attitudes or value judgments that the agent of the microaggression seemingly holds.

On Grice's theory of implicature, our utterances communicate both literal meanings and implied meanings, or implicatures (Grice 1957). By the term 'implicature,' Grice refers to the act of communicating something by saying something else. In other words, an implicature is what a speaker communicates by an utterance without saying it.[30] For Grice, an implicature occurs if certain conditions are met, but for the purposes of this paper we don't need to go into detail about what those conditions require; it will suffice to note that for Grice both the speaker's intention and the hearer's understanding play a central role in what the speaker implicates.[31] Consider the following example.

---

[30] For Grice, implicatures can be conventional, or semantic, or conversational. I here focus only on the latter. See Grice (1989).

[31] Grice writes,

> A man who, by (in, when) saying (or making as if to say) that p has implicated that q, may be said to have conversationally implicated that q provided that (i) he is to be presumed to be observing the conversational maxims […]; (ii) the supposition that he is aware that, or thinks that, q is required in order to make his saying or making as if to say p (or doing so in those terms) consistent with this presumption; and (iii) the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (ii) is required. (1989, 30)

Suppose that as you are leaving the office with a couple of friends, you suddenly remember that you forgot to turn off the light in your office, so you say to your friends, “I’ll be right back—I forgot to turn off the light!” Suppose, further, that one of your friends replies “Oh, but we don’t pay for electricity!” intentionally and successfully communicating that it’s fine for you to leave without going back to turn off the light. In this example, your friend conveys both a literal meaning and an implied meaning, or implicature. His literal meaning is the content of what he explicitly states: namely, that you do not pay for electricity; and his implied meaning, or implicature, is that it is fine for you to leave without going back to turn off the light.

But notice that this is not all that your friend communicates. He also communicates—albeit unintentionally—additional information about his own set of beliefs or value judgments. This additional information is not equivalent to the meanings he communicates, but it can be inferred from them. For instance, your friend seems to also communicate that he holds a set of attitudes or value judgments that includes something along the lines of “It’s okay to unnecessarily waste resources,” “It’s okay to waste resources we don’t pay for,” “It’s not problematic to unnecessarily waste these resources,” etc.

Now, determining exactly how to calculate the precise content of the attitudes that an agent communicates through an action is a project that is beyond the scope of this paper. What I want to do instead is suggest that we might conceive of this additional, unintentionally communicated information as part of the agent’s underlying set of beliefs that best explains the literal and implied meanings that she communicates through her utterance. Accordingly, if your friend did not hold an attitude such as “It’s okay to unnecessarily waste resources” or “It’s okay to waste resources we don’t pay for,” it would be difficult to explain, in this scenario, why he would remind you that you do not pay for those resources in order to communicate that it’s okay to leave without turning off the light.

Note that the literal and implied meanings that your friend communicates also allow us to infer certain attitudes or value judgments that he must *not* hold. For example, we can infer that he seemingly does *not* hold an attitude such as “It is morally problematic to waste resources unnecessarily.” If he did hold such an attitude, it would be unclear why he would say what he said to imply what he implied.[32] Nevertheless, regardless of whether the additional, unintentionally communicated information involves attitudes that your friend holds or attitudes he does not hold, what’s important for our purposes is that these attitudes are not part of the literal or implied meanings of the action—at least not in the sense I’ve described these meanings here. In other words, even though those attitudes or value judgments can arguably be inferred from your friend’s literal and implied meanings, they are not equivalent to those meanings.

---

[32] It’s worth noting that the set of beliefs we arrive at as an explanation for the literal and implied meanings of an action may be disputed by the agent. This need not mean that the set of belief we arrive at is wrong; agents do not have transparent introspective access to their own attitudes (See Carruthers 2011). But it does suggest that there may be cases in which an agent only seemingly holds the set of beliefs that we infer from the meanings she communicates. I return to this point on footnote 36.

The point here is this: it would be a mistake to say that this additional, unintentionally communicated information is part of the meanings that an agent communicates. In other words, while we can claim that your friend communicates a literal meaning ("But we don't pay for electricity") and an implied meaning ("It is okay for you to leave without turning off the lights") it would be a stretch to say that he also communicates the *meaning* that "It's okay to waste resources we don't pay for." Rather, in addition to the literal and implied meanings that your friend communicates, he also reveals, or makes manifest, an attitude or value judgment that *he* seemingly thinks it is okay to waste resources we don't pay for.[33] That attitude is equivalent neither to his literal meaning nor to his implied meaning but can be inferred from those meanings as part of the best explanation for the pair of meanings that he communicates. In what follows, I will refer to this kind of information as *hyper-implicature*—'hyper-' because although it resembles an implicature in that it is not explicitly stated, it is beyond just implicature because it is inferred in part from the speaker's implicature.

What does this discussion about meaning and hyper-implicature have to do with microaggressions? I believe that microaggressions are the kind of action through which an agent hyper-implicates certain kinds of beliefs or attitudes that she

---

[33] I submit that microaggressions involve *unintentionally* revealed attitudes—a point that I discuss in more detail in Section 5. Moreover, note that some authors have referred to this kind of additional information as *unintentional implicature*. This is not an uncontroversial view, as most philosophers of language tend to agree that Gricean implicature requires the intention of the speaker. For a defense of unintentional implicatures, see Lassiter (2012).

seemingly holds. In other words, microaggressions may be a kind of action through which, like your friend in the example above, an agent reveals, or makes manifest, attitudes or value judgments that are separate from the literal meaning and the implied meaning that she intends to communicate. These hyper-implicated attitudes or value judgments are not explicitly stated or intentionally implied, but can nonetheless be inferred from the agent's literal and implied meanings as part of the best explanation for such meanings. Now, the kinds of attitudes or value judgments that an agent hyper-implicates in a microaggression are not about just anything; rather, they are about the *interpreter* of the agent's action, are based on some (social) group to which she belongs, and may (often for that reason) be offensive to her.

Consider the paradigm microaggression of asking a minority individual "Where are you from?" Now, suppose that the question is asked by a white American to an Asian-American and that, given the context of their interaction, the white American intentionally implies "What *part of Asia* are you from?" In other words, the white American's literal meaning is what he explicitly states—namely, "Where are you from?"—and his implicature is what he intentionally implies—namely, "What *part of Asia* are you from?" If the preceding is on the right track, then we can infer from these meanings a hyper-implicature—that is, information involving the set of attitudes or value judgments that the white American holds (or does not hold). Suppose, then, that given the context of the action, we can infer that the content of her hyper-implicature approximates "This person must not be American; people who look Asian are not American" or "This person must not be American; people who look Asian must have been born somewhere in Asia." Thus, in asking "Where are you from," the white American not only communicates a literal and an implied meaning, but also reveals, or makes manifest, a hyper-implicature: a set of false beliefs that he seemingly holds, including, in this scenario, that people who 'look' Asian are not American or that they must have been born outside of the U.S. Absent some alternative, better explanation, her question "Where are you from?" accompanied by her implicature "What *part of Asia* are you from?" would not seem to make sense.

Note that the hyper-implicature in this example—namely, that "This person must not be American; people who look Asian are not American"—is based on the interpreter of the action's membership in the (social) group 'Asian-American' (or: '*looks* Asian-American'). Specifically, the white American's hyper-implicature is an inference about the interpreter of the action from his membership in that (social) group. Note, also, that it would not seem unreasonable to think that such a hyper-implicature might be offensive to the interpreter of this action *from the point of view of the agent of the microaggression*. In other words, it would not be unreasonable to think that the content of such a hyper-implicature might, e.g., undermine the interpreter's sense of belonging or make her feel abnormal. Now, I do not mean to suggest that the agent of the microaggression must actually foresee that such a hyper-implicature may be offensive to the interpreter; rather, I think that given, e.g., the agent's socio-historical situation, we can say that she could, practically speaking, foresee that the interpreter of her action might be caused pain or offense by her hyper-implicature. Furthermore, I do not mean to suggest that the interpreter of the action must actually be offended by such a hyper-implicature; all I want to suggest is that given the context in which the action occurs (e.g., a certain sociohistorical context), it would not be unreasonable for the interpreter of the action to take offense.

We can sum up my proposal for how to identify a microaggression as follows. A microaggression takes place if, and only if, the following three conditions are met:

1. An agent hyper-implicates to another an attitude or value judgment about the other,
2. The attitude or value judgment that the agent hyper-implicates is based on the other's membership in some (social) group, and
3. The agent can—although she need not actually—foresee that such an attitude or value judgment may be insulting to the agent about whom it is.

Note, that although these are necessary and jointly sufficient conditions for a microaggression to take place, they are not jointly sufficient for a *morally objectionable* microaggression to take place. I will discuss the necessary and sufficient conditions for a morally objectionable microaggression to take place in the next section.

The framework I'm proposing for identifying when a microaggression takes place is more objective than Sue et al.'s current framework. By appealing to the concepts of literal meaning, implicature, and hyper-implicature, it takes into account not just the understanding of the interpreter of the action (i.e., the "victim" of the microaggression), but also the intention of the speaker, as well as the context in which the action occurs, in order to determine what the agent of a microaggression reveals or makes manifest about himself by engaging in the microaggression. Consequently, my framework is open to faring better than Sue et al.'s current framework in terms of accounting for some intuitive cases of microaggressions. Admittedly, a more sophisticated account for calculating the exact content of a hyper-implicature will be needed, but it’s worth noting that if such an account is developed and adopted, we will be in a good position to avoid two counterintuitive entailments of Sue et al.’s view.

The first involves the case, from Section 2, of the white teacher who constantly calls on his white students first because he believes that calling on his black student, because she is black and likely less competent, will derail the class discussion. On Sue et al.’s current view, if the student does not deem the teacher’s action a microaggression, no microaggression has occurred. But according to my proposal, if the teacher communicates certain literal and implied meanings and thereby hyper-implicates a certain kind of attitude—e.g., that “A black student is less competent because she is black and blacks are less competent than whites”—then her action *is* a microaggression, for it reveals a certain kind of potentially offensive belief that the teacher holds, even if the student does not register it. This is because in this example the teacher’s hyper-implicature involves a value judgment about his student in virtue of the student’s membership in some social group—a value judgment that the teacher could, it would be reasonable to believe, foresee causing the student pain or suffering.

The second counterintuitive entailment of Sue et al.’s view that we can avoid by adopting the framework I am proposing involves the teacher, also from Section 2, who has consciously decided that she will not be a microaggressor. In that case, after failing to see a nonwhite student raise his hand first, the teacher calls on a white student and the nonwhite student deems the action a microaggression. On Sue et al.'s current view, the teacher's action is a microaggression because, insofar as microaggressions lie in the eye of the beholder, the teacher's intentions do not play a role. But according to my proposal, it's not straightaway clear that the teacher's action is a microaggression. Again, more needs to be said with regard to how to calculate the content of the set of attitudes that the teacher hyper-implicates in such a scenario; but her intention will play a role at the very least indirectly: e.g., in terms of the meanings she intentionally communicates, whether explicitly or implicitly. This means that my framework offers more guidance for preventing microaggressions than Sue et al.'s framework, for the agent's intention will play at least an indirect role in what she reveals through a microaggression.

4. *Refining the concept of microaggressions: When microaggressions are morally objectionable*

With this more objective framework for identifying microaggressions in hand, we can now ask: When are microaggressions morally objectionable? In what follows, I argue that the answer to this question lies in the kind of attitude that an agent hyper-implicates in a microaggression. In some cases, the attitude will be a *descriptive* attitude that, when hyper-implicated, may bring about consequences that may be morally problematic: e.g., foreseeably causing pain or suffering in another. In other cases, the attitude will be an *evaluative* or *moral* attitude that may be morally objectionable despite the consequences of hyper-implicating it. Microaggressions of this latter kind, I will suggest, may involve a failure of what Stephen Darwall calls *recognition respect*: a failure to regard another person appropriately.

Imagine *A* and *B*, both American citizens, having a conversation about certain kinds of food. *A* is of Chilean descent and *B* is of European descent, and each knows about the other's descent and citizenship status. Imagine, further, that *A* tells *B* about her love for a certain spicy Korean dish and that *B* responds—in a manner that makes it obvious that she is not joking—"Oh, right, you guys like spicy."[34] Suppose *B* is intentionally implying something along the lines of "You, Latin Americans, like spicy food."[35] Given *B*'s implicature and literal meaning, it would be reasonable to infer, she hyper-implicates not only that she believes something along the lines of, "Of course *A* likes spicy; Latin Americans like spicy," but also that she does not believe something along the lines of, "Chili powder is used in Mexican cuisine, but that does not mean that chili powder is also used in the cuisine of every other Latin American country." If she did believe this, it would be unclear why she would say "Oh, right, you guys like spicy" and implicate "You Latin Americans like spicy food." Notice that the attitude that *B* hyper-implicates is about *A*'s membership in a certain (social) group (i.e., 'Latin American'), and that it would not seem unreasonable to think that, given *B*'s socio-historical situation, she could foresee that a hyper-implicature such as

---

[34] It's important to note that while *B* is not joking when she says this, humor can raise complications by altering the implied meanings and thereby the attitudes or value judgments that we can infer from those meanings.

[35] I don't take this to be controversial, especially given the background against which this kind of comment occurs, at least in the U.S.: namely, one that involves many people's belief that Mexican food is spicy and that whatever may be true of Mexico will be true of the rest of Latin America.

"Of course *A* likes spicy; Latin Americans like spicy" may be hurtful or offensive to *A*—after all, in the contemporary U.S., there is a tendency of making generalizations about Latin Americans as if Latin American was a homogenous region of the world.

The idea that I wish to defend with this example is that *B* hyper-implicates a misleading descriptive attitude that need not, in virtue of being hyper-implicated, make her action morally objectionable. For example, in some cases, it would not be unreasonable to expect *A* to be unfazed by what *B*'s comment communicates, or to investigate whether *B* does in fact believe what she hyper-implicates and correct her if she does. In other words, there may be cases in which a microaggression's interpreter may unreasonably experience pain or suffering: e.g., if the hyper-implicature is not one that the interpreter of the action experiences on a regular basis. But it would be a different story if *A* regularly undergoes hyper-implicatures involving inferences about her in virtue of inaccurate or misleading stereotypes about a group to which she is taken to belong. In that case it would be understandable if *A* experienced negative mental states as a result of such hyper-implicatures: negative mental states such as annoyance, indignation, or offense. And it would then seem unreasonable to expect *A* to be unfazed by such experiences. Of course, one could argue that if *A* had such experiences on a regular basis, then it would be in her best interest to grow thicker skin. But without getting too far from the point, let me note that there's reason to think that this is an unviable option. Imagine regularly undergoing hyper-implicatures that suggest that others regard you not as an individual, but as a member of some group about which some (often negative) stereotype is true. Being unfazed by, or growing thicker skin towards, this kind of treatment may come at the risk of consequences that may be worse than irritation: e.g., becoming detached from or resentful towards one's community or society. Nevertheless, regardless of whether or not detachment from or resentment towards one's community or society is in fact worse than plain irritation, the point here is just that some microaggressions may occur often enough that each additional microaggression may understandably become harmful even if a single such experience should cause no harm. Kelly Burns puts this point nicely when she writes, "While one paper cut is not a big deal, if you are continually being cut in the same place, it can become quite painful" (Burns 2014, 140).

To put the point differently, we might say that if an agent engages in a microaggression against another agent and thereby foreseeably causes pain or suffering in the other agent, then her action amounts to a morally problematic microaggression. Importantly, as above, by 'foreseeably' I do not mean that the agent must *actually* foresee that the content of her hyper-implicature will cause pain in the interpreter of her action, but rather that, given, e.g., her socio-historical situation and cognitive abilities, the agent could, practically speaking, foresee that the interpreter would be caused pain.

Note that this point that microaggressions may or may not cause harmful mental states is not one that Sue et al.'s current view seems able to account for. In their view, microaggressions are harmful by definition: they are harmful because the messages they express cause pain or suffering. On the framework I'm proposing, however, a microaggression can occur without being morally objectionable. This is an advantage of my proposal in part because it entails the intuitive conclusion that in the case of the stoic individual who is unfazed by a microaggression, a microaggression still occurs; it just does not cause the stoic individual pain—something Sue et al. do not seem in a good position to claim.

Nevertheless, against the framework I am proposing, an opponent may highlight that in cases such as the stoic's, in which the microaggression does not cause any pain or suffering, my framework leads to the conclusion that the microaggression isn't morally objectionable. In other words, if no one is irritated, offended, etc., what's the problem? The opponent may then conclude that my framework, just like Sue et al.'s, fails to justify the microaggression program, as we could instead train potential victims to be like the stoic and refrain from reacting negatively to microaggressions. But I think that this objection relies on a very narrow understanding of what makes an action morally objectionable: that an action is morally objectionable only if it causes negative mental states—or perhaps more charitably: only if it leads to bad consequences. Indeed, that seems to be Sue et al.'s own take of what makes microaggressions morally objectionable. But as I have been arguing, that approach is misguided, for it excludes other types of wrongs. What I want to suggest is that wrongs come in various forms, only some of which involve consequences. As I will now argue, agents who engage in microaggressions sometimes hyper-implicate attitudes that are morally questionable independently of whether or not anyone is caused any pain or suffering.

Imagine an agent who hyper-implicates an attitude the content of which is something along the lines of "Blacks are inferior to whites." Note that there is a morally significant difference between these attitudes and *descriptive* attitudes such as the one hyper-implicated in the example above: namely, “Latin Americans like spicy.” Whereas attitudes such as “Latin Americans like spicy” are arguably descriptive, attitudes such as “Blacks are inferior to whites” are evaluative or moral; they involve a value judgment: in this case, one about the moral standing of another agent (or group of agents). What I want to suggest is that an agent who holds an erroneous descriptive attitude such as “Latin Americans like spicy” makes a factual error about Latin Americans as a group, but that an agent who holds an attitude such as “Blacks are inferior to whites” makes a *moral* error about blacks as a group: one that involves holding an attitude that is demeaning.

Examining the precise nature of a demeaning attitude is beyond the scope of this paper, but Stephen Darwall’s (1977) notion of *recognition respect* is useful to that end. As I understand Darwall, recognition respect involves “a disposition to weigh appropriately some feature or fact in one’s deliberations,” such that that feature or fact limits our conduct (Darwall 1977). With this notion of recognition respect in hand, we might contend that recognition respect for another person might also involve a disposition to weigh in one’s moral deliberations the fact that the other is a person, such that the other’s status restricts not just an agent’s conduct toward another, but also the set of *attitudes* that it is permissible for him to hold about another agent. We might then hold that a set of permissible attitudes toward another agent would seem to exclude attitudes that fail to recognize the status of another agent as a person. A demeaning attitude might then be understood as an attitude that involves a failure of recognition respect for another person. A clear example of such an attitude would seem to be an attitude that gets the moral standing of another agent wrong, such as “Blacks are morally inferior to whites.”

The idea, then, is that some moral or evaluative attitudes may be demeaning: they may involve an error about the moral standing of another agent. In the context of microaggressions, we might contend that when an agent engages in a microaggression that involves an attitude that is demeaning—e.g., “Blacks are inferior to whites”—he reveals, through the microaggression, a failure of recognition respect for the other agent. In other words, he reveals that he does not give appropriate weight to the fact that the other agent is a person. To be precise, the agent of the microaggression would seem to have made an error, in his moral deliberations, regarding the set of attitudes that it is permissible for him to hold about another agent.[36]

Now, in line with the sufficient condition for a morally problematic microaggression to take place examined above, this kind of microaggression would seem to clearly be morally problematic if the agent foreseeably causes pain or suffering in the interpreter of her action as a result of engaging in the microaggression. But the suggestion I am after is that an agent who engages in a microaggression that involves a *demeaning* attitude about another would, through such microaggression, seem to reveal something morally objectionable regardless of whether or not she foreseeably causes pain or suffering in the interpreter of her

---

[36] Now, one might think that just because an agent has a certain belief about another does not mean that the belief enters into the agent’s deliberation regarding what to do or of how to regard another. But it’s worth noting that such a belief may nonetheless enter into the agent’s deliberations as an underlying assumption.

action. This is because in *holding* such an attitude the agent of the microaggression would be making an error about the moral status of the other agent. More specifically, an agent who engages in a microaggression that involves a demeaning attitude would seem to have made an error, in his moral deliberations, regarding the set of attitudes that it is permissible for him to hold about another agent.

Put differently, the idea is that a morally problematic microaggression takes place in at least two ways: if, in virtue of a microaggression, an agent foreseeably causes pain or suffering in another, and if the attitude that she hyper-implicates is demeaning. And this latter instance of a microaggression would seem to be morally objectionable regardless of whether or not the other is foreseeably caused any pain.[37]

---

[37] Earlier I mentioned that there may be cases in which an agent only seemingly holds the attitudes or value judgments that she hyper-implicates, such that some microaggressions may involve an agent's revealing that she holds a certain group-based attitude, while others may involve an agent's *seemingly* revealing that she holds a certain group-based attitude when she does not in fact hold it. In my view, a microaggression would take place so long as the hyper-implicature obtains—that is, independently of whether or not the agent holds the attitude that she hyper-implicates. This suggests that, in my view, an agent can hyper-implicate a certain kind of attitude or value judgment without actually holding the attitude that she hyper-implicates. Imagine an agent who hyper-implicates the attitude "Blacks are morally inferior to whites" to another agent who is black, but that the agent who hyper-implicates this attitude does not actually hold this attitude. The question is: assuming this is a microaggression, is it a morally objectionable one? I take it that although the attitude that this agent hyper-implicates is demeaning, if the agent does not foreseeably cause anyone any pain, then his action would

An interesting question arises here: Does a microaggressor who hyper-implicates a demeaning attitude, thereby revealing a moral error but without harming anyone, *act* wrongly? I cannot begin to answer this question here, but let me note that how we answer this question might turn on whether or not we think that there is such a thing as doxastic wronging. Doxastic wrongs are said to be wrongs that occur in virtue of what an agent *believes* about another agent—not in virtue of how an agent *acts* towards another. The idea is that an agent can wrong another by believing, e.g., that she can treat him as a mere means, even if she never so treats him.[38] In the context of microaggressions, it seems plausible that even if we agree that an agent who hyper-implicates a demeaning attitude but does not harm anyone does not *act* wrongly, there may still be reason to think that something morally objectionable takes

---

not be morally objectionable. Moreover, if the agent does foreseeably cause pain or suffering in another, his action would be morally problematic regardless of whether or not he actually held the attitude he hyper-implicated. In other words, in this latter case, the agent's action would meet one of the sufficient conditions for a morally problematic microaggression to take place. Interestingly, though, insofar as the agent does not actually believe that "Blacks are morally inferior to whites," there is a question of whether or not he can be held responsible or blameworthy for hyper-implicating such an attitude. Although an appropriate answer to this question is not a project I can tackle here, I take it that a relevant consideration in answering this question will be whether or not the agent could have foreseen that his action would hyper-implicate the relevant attitude or that hyper-implicating such an attitude would cause pain or suffering in another.

[38] See Basu and Schroeder (2019), Basu (2019), and Schroeder (2018).

place, insofar as the agent would seem to engage in a moral error with respect to the status of another agent.

To sum up the framework I gave been proposing: A microaggression takes place if, and only if, an agent hyper-implicates to another an attitude or value judgment that is based on the other's membership in some social group, and she can—although she need not actually—foresee that such an attitude or value judgment may—although it need not actually—be insulting to the agent about whom it is. But these necessary and jointly sufficient conditions for a microaggression to take place are not jointly sufficient for a *morally objectionable* microaggression to take place. Rather, a morally objectionable microaggression takes place if an agent who engages in a microaggression foreseeably causes pain or suffering in the victim of the microaggression. But a morally objectionable microaggression can take place even if the agent does not foreseeably cause anyone any pain or suffering: for example, a morally objectionable microaggression may occur if an agent of a microaggression hyper-implicates a demeaning attitude, such as "Blacks are morally inferior to whites." In such cases, the agent reveals that he has made a moral error about the moral standing of another agent.

In a nutshell, there are at least two sufficient conditions for a morally objectionable microaggression to take place:

1. If an agent foreseeably causes pain or suffering in another as a result of a microaggression, and
2. If the attitude that an agent hyper-implicates in a microaggression is demeaning.

Notice that whereas Sue et al.'s view currently appeals strictly to the pain or suffering that microaggressions may cause in their victims to explain why microaggressions may be morally objectionable, my proposal highlights both that not all microaggressions are morally objectionable and that microaggressions may be morally objectionable for reasons other than the harm they may cause in their victims. Moreover, recall that Sue et al.'s view, insofar as it is entirely victim-focused, invites the objection that instead of encouraging minority-individuals to look for microaggressions left and right, the microaggression program should train such individuals to be more resilient to microaggressions. My proposal need not deny that there is virtue in resilience, but, in holding that some microaggressions may involve moral errors on the part of their agents, it sheds light on a different aspect of when microaggressions are morally objectionable. My proposal for amending Sue et al.'s view of microaggressions is thus more action-guiding than Sue et al.'s current view, as it makes avoiding some microaggressions a matter that is strictly up to the microaggressor.

5. *Some lingering questions*

A few questions may have lingered in the background of this discussion. Let me start by restating that on the framework I am proposing, committing a microaggression is not always wrong. Some microaggressions may neither have bad consequences nor hyper-implicate demeaning attitudes, but they still are microaggressions if the agent of the microaggression hyper-implicates a certain kind of group-related attitude that might foreseeably cause pain in the interpreter of the action.

Now, I mentioned earlier that on Sue et al.'s current view microaggressions are slights directed at members of disadvantaged groups. Does an action that communicates an attitude about an individual in virtue of his membership in some *non*-disadvantaged group count as a microaggression, in my view? Imagine *C*, a white waitress who regularly experiences covert slights from her coworkers—slights having to do with their perception of her as a "Becky" (a term sometimes used to refer to blonde women who are said to have predictable preferences and to lack common sense). I think that such slights would indeed count as microaggressions if they communicated certain (negative or misleading) attitudes about *C* in virtue of her being a "Becky." What's central to microaggressions is not that the group about which the hyper-implicated judgment is made be a disadvantaged group, but that the judgment strips the individual of her individuality by judging her in virtue of a (misinformed or erroneous) belief about a group to which she is taken to belong. This means that microaggressions are a kind of phenomenon that occurs not only to members of disadvantaged groups, but to anyone. Nevertheless, this point that microaggressions need not involve judgments about disadvantaged groups should be taken with care for two reasons. First, microaggressions that don't involve disadvantaged groups may be less pervasive than those that do; after all, microaggressions involving disadvantaged groups may show up not just in one's workplace, for instance, but also in other spheres of life, including, for example, education, healthcare, customer service, etc. And second, microaggressions that don't involve disadvantaged groups may be less likely to be demeaning, as they are not accompanied by a history of arguably unjust treatment of the members of those groups.

Another worry that may have arisen throughout the paper may be that since I have argued that a microaggression is the kind of action that *unintentionally* reveals information about the belief set of the agent of the microaggression, it follows that actions that *intentionally* reveal information about the belief set of the agent of the microaggression may not count as microaggressions on my proposal. It seems to follow, then, that when Roseanne Barr, a public figure who is white, compared Valerie Jarrett, a former government official who is African-American, to an ape, she did not engage in a microaggression.[39] And, when Donald Trump compares illegal immigrants to "animals," or when he refers to sanctuary cities as places where immigrants "breed," he also does not engage in microaggressions.[40] I welcome these conclusions because even if such actions have some of the same negative effects of microaggressions (e.g., cause negative mental states), and even if they evince the same kind of moral failure in the agent that a microaggression that reveals a demeaning attitude evinces (e.g., a failure to regard another appropriately), they seem to be actions of a different kind: they seem to involve *intentionally implied* racist and dehumanizing meanings, or, to use the technical term, implicatures—and not attitudes that can only be inferred from those meanings. Moreover, they seem to amount to

---

[39] See Heil (2018).

[40] See Wolf (2018). For more on Trump's overtly racist actions, including the comments "Black guys counting my money! I hate it. … I think that the guy is lazy. And it's probably not his fault, because laziness is a trait in blacks," see Leonhardt and Philbrick (2018).

overtly aggressive and hostile behavior toward members of certain social groups, not covert, inferable attitudes. Such actions may instead amount to behavior that is primarily morally objectionable because it is intended to harm.

Finally, if microaggressions are unintentional, potentially offensive ways to communicate certain kinds of attitudes, and not overtly aggressive behavior, one may worry that the term 'micro*aggression*' is a misnomer. After all, the word *aggression* seems to fundamentally mischaracterize the phenomenon as a violent action. A better term, one might suggest, would be *micro-insult* or *covert-insult*. This is a question worthy of a much deeper analysis than I can offer here. But it's worth noting in favor of this worry that insofar as the phenomenon under investigation here is unintentional, it seems counterintuitive to refer to it as an "unintentional aggression."[41] Yet it's also worth noting that regular experiences of expressions of otherwise covert demeaning attitudes, in particular, may have a *feel* of hostile attacks, or aggressions, to a person's dignity or moral standing.

6. *Concluding remarks*

I have argued that Sue et al.'s notion of microaggression, insofar as it has subjective and utilitarian features, offers little guidance for identifying and preventing microaggressions, and explains only some of the ways in which microaggressions may be morally objectionable. I have also offered a way to improve Sue et al.'s

---

[41] Although it's worth noting that people do sometimes unintentionally and without any self-awareness act aggressively: e.g., if they are having a bad day and, unintentionally and without even realizing it, slam a door.

notion of microaggressions: by adopting a framework that appeals to concepts from the philosophy of language in order to identify microaggressions. On this framework, we can morally evaluate the content of the attitudes that an agent hyper-implicates in a microaggression in order to determine when a microaggression is morally objectionable. This proposal generates a notion of microaggressions that is neither subjective nor strictly utilitarian and thus does a better job than Sue et al.'s current view in terms of how to identify microaggressions, and in terms of how to explain when they are morally objectionable. Yet, if this proposal is going to be successful, a more detailed account of how to calculate the content of the attitudes that are hyper-implicated in a microaggression, as well as of how to determine the precise nature of a demeaning attitude will have to be developed.

One final point. I mentioned earlier that critics of the microaggression program generally argue that supporting the microaggression program is a bad idea because it may have negative consequences not only for college campuses but for society at large. They argue, for instance, that it may lead to unwarranted restrictions on free speech, worsen racial tensions, and help create a culture of victimhood. While these consequentialist worries face problems of their own,[42] if my proposal is on the

---

[42] For instance, Nathan J. Robinson writes that "Haidt has devoted himself to trying to understand the psychological roots of political ideology, and yet he doesn't seem interested in listening to the perspectives of the people whose motivations he wants to diagnose" (Robinson 2018). There are grounds for this response to Haidt, for he mistakenly claims that the microaggression program *trains* "diverse students, who arrive from all over the world with very different values and habits" to "react with pain and anger to ever-smaller specks

right track, these critics will be advised to reconsider their arguments at least with respect to microaggressions that may contain demeaning attitudes, as those microaggressions may turn out to constitute a particularly morally objectionable social practice for which the victims of microaggressions are not to blame.

---

that they learn to see in each other's eyes" (Haidt 2017). To the contrary, see footnote 25 above.

Chapter 3: When Is the Promotion of Prenatal Testing for Selective Abortion Wrong?

Introduction

Medical professionals routinely offer prenatal genetic testing services to their expecting patients. In theory, this testing helps prospective parents better prepare for the birth of their child, e.g., if the child will have a disability. In practice, however, such testing often leads to the termination of pregnancies that would produce a child who has a disability. In light of our society's history of poor treatment of people who have disabilities, some bioethicists believe that when medical professionals promote the use of this testing for abortion on grounds of disability, they express a devaluing message to and about extant disabled people: that their lives are not worth living or that their lives are not worthy of support (Parens and Asch 2000).[43] Supporters of this *expressivist objection* further maintain that in expressing such a message, medical professionals not only cause wounded feelings in extant people who are disabled; they also reinforce negative, misinformed attitudes about them and thereby further stigmatize them. While the expressivist objection has received quite a bit of support from disability rights theorists—in part because of its intuitive appeal—its current

---

[43] The expressivist objection has been raised both against parents who *use* prenatal testing for selective abortion and against medical professionals who *promote* or *routinize* such testing. While this paper focuses only on medical professionals' promotion of such testing, it has implications for how we think about prospective parents' use of such testing, as well as medical professionals' mere *offer* of such testing.

formulation suffers, I believe, from a lack of precision that renders it implausible. A careful examination of this imprecision, along with an appeal to some tools from the philosophy of language, might reveal a more adequate formulation of the objection that preserves its core insight.

The paper develops as follows. Section 1 examines in more detail the expressivist objection's current formulation along with some of its shortcomings. Section 2 recruits two concepts from the philosophy of language—*illocutionary acts* and *implicatures*—in order to offer an improved formulation of the expressivist objection. According to this *expressivist objection**, the promotion of prenatal testing for selective abortion is a kind of illocutionary act that expresses precise information. With this formulation in hand, Section 3 examines when this illocutionary act is morally objectionable. Finally, Section 4 considers the possibility that sometimes the promotion of prenatal testing for selective abortion involves a morally objectionable *belief*. I conclude that the promotion of prenatal testing for selective abortion is at least sometimes wrong.

1. The expressivist objection's current formulation

According to the broadest formulation of the expressivist objection, the promotion of prenatal testing for selective abortion expresses the message that the lives of disabled people are less worth living than the lives of nondisabled people (Asch 2000, Asch and Wasserman 2015, DeGrazia 2012).[44] The idea is that "[i]f we abort a fetus on the

---

[44] Now, the objection can be raised at the mere promotion of prenatal testing—that is, even if it doesn't lead to terminating pregnancies on grounds of disability; but I here examine what I

basis of a single trait, this symbolically suggests that not only this fetus's life but the life of *anyone* who has this trait is not worth living or preserving” (Kukla and Wayne 2018). In this section, I argue that when formulated in more detail the expressivist objection is imprecise, both with respect to the content of the messages that are purportedly expressed by the promotion of prenatal testing for selective abortion, and with respect to how such an action brings about harms for extant disabled people. Moreover, I argue that this imprecision renders implausible the expressivist objection’s claim that the promotion of prenatal testing for selective abortion expresses devaluing messages, and that it obscures non-consequentialist wrongs that this action may contain.

Supporters of the expressivist objection have described the messages purportedly expressed by the promotion of prenatal testing for selective abortion as “disparaging,” “disrespectful” (Nelson 1998), “hurtful” (Parens and Asch 2000), “demeaning” (Saxton 1997), “discriminatory” (Asch 2003), or just “negative” (Holm 2008). Moreover, some have claimed that the content of such messages is that “disabled people’s lives are not worth living” or that disabled people’s lives are not “worthy of support” (Parens and Asch 2000, 8), while others have claimed that it is that disabled people are morally inferior to nondisabled people, or that disabled people should not be—or should not have been—born (Scott 2005). But the heterogeneity of these claims raises concerns. To clarify, the claim that the messages

---

take to be the strongest form of the objection, namely, that against the promotion of prenatal testing for selective abortion.

expressed by prenatal testing are “hurtful” or “demeaning” suggests that the moral concern with such messages is that they harm extant disabled people psychologically: e.g., that they attack disabled people’s sense of self-worth or cause them offense. But the claim that the messages are “discriminatory” suggests that the moral concern with the expression of such messages lies in the kind of *treatment* that extant disabled people may encounter in our society. This latter concern is about wrongful discrimination and can be independent of the psychological states of extant disabled people.[45] Nevertheless, all of these concerns are unified by the idea that it is the consequences of an action that determine its moral permissibility. After all, supporters of the expressivist objection are concerned with the *expressive power* of the promotion of prenatal testing for selective abortion: i.e., what expressing negative or devaluing attitudes may cause for disabled people—including, e.g., psychological harms such as attacking their sense of self-worth, civic harms such as further stigmatizing them (e.g., by promoting or reinforcing false or denigrating attitudes about them), and material harms such as reducing the amount of resources that disability communities receive.

However, as critics highlight, supporters of the expressivist objection offer no account for how to determine the kind and content of the messages or meanings that the objection claims the promotion of prenatal testing expresses. Admittedly, supporters of the expressivist objection have suggested that what our actions express

---

[45] For a different interpretation of the variance in these claims, see Hofmann (2017), who argues that these are various *elements* of the expressivist objection. I find that interpretation of the debate much too charitable.

is necessarily connected with our society's history—a history that includes various forms of unjust treatment toward people who have disabilities. But the way in which such a history determines the meanings or messages that an action such as promoting prenatal testing for selective abortion expresses has not been addressed.[46]

David DeGrazia (2012) develops this kind of objection in some detail. While he focuses explicitly on the *use* of prenatal testing for selective abortion by expecting parents, his argument is also applicable to the *promotion* of prenatal testing for selective abortion by medical professionals. According to DeGrazia, while nonverbal actions that belong to agreed-upon systems of signs express precise meanings or messages, nonverbal actions that do not belong to agreed-upon systems of signs do not express precise meanings or messages. For example, if a driver gives another the middle finger after the other cuts him off on the highway, DeGrazia maintains, the driver couldn't then claim that he used this sign to mean "Top of the morning!" to the other driver. Giving someone the finger in the context of someone cutting you off is part of an agreed-upon system of signs and carries a specific meaning ("F-you"). Using prenatal testing for selective abortion, by contrast, is not part of an agreed-upon system of signs, and so does not express a specific, agreed-upon meaning. Thus, the nature and content of the message it expresses is underdetermined. A message that "disabled people's lives are not worth living" is no more likely than a message such as "It is acceptable to prefer not to have a child with a particular substantial disability and to make reproductive decisions accordingly" (DeGrazia 2012, 105). In the latter

---

[46] For support for this claim, see Buchanan (1996) and Nelson (1998).

case, no negative message about extant disabled people, e.g., that they have lower moral status, is expressed. The expressivist objection against the use of prenatal testing fails, DeGrazia concludes, because the use of prenatal testing for selective abortion cannot express messages or meanings that are precise enough for the claims of the expressivist objection.[47] In other words, using prenatal testing for selective abortion, unlike giving someone the finger, does not *necessarily* express a message of disvalue.

Now, in the case of the promotion of prenatal testing for selective abortion DeGrazia's objection would be that promoting prenatal testing for selective abortion is a nonverbal action that, unlike giving someone the finger, does not have agreed-upon meanings or messages.[48] Thus, it cannot express messages or meanings that are precise enough for the claims of the expressivist objection.

Mary Ann Baily (2000) also argues that it's not clear why we should accept the interpretation that using prenatal testing necessarily or even generally expresses the devaluing meanings or messages that supporters of the expressivist objection claim it expresses. As she writes, "Acceptance of prenatal diagnosis and abortion certainly *can* express rejection of people with disabilities. But *must* it? Or can one

---

[47] See also McMahan (2005).

[48] DeGrazia suggests this when he writes that "If use or institutional promotion of PGD [prenatal genetic diagnosis] communicates anything, it does so nonverbally. Verbal communication, by contrast, involves a public language: a highly complex, rule-governed system of signs whose conventions of meaning are broadly accepted by the linguistic community in question… Words have public meanings" (at pp. 104-105).

look at it another way?" Baily's point is that an agent's reasons for acting surely play a role in what his action expresses. For instance, a prospective mother may choose to have prenatal testing for selective abortion for various reasons, including, for example, a preference not to have a child who, in having a disability, would make her life more difficult than she would want. Thus, Baily writes, if using prenatal testing for selective abortion expresses anything:

> It does not express my opinion on whether people with disabilities can have worthwhile, fulfilling lives (they can), whether parenting a child with a disability can be rewarding (it can), whether other people should be pressured into aborting fetuses with disabilities (they shouldn't), or whether societal resources should be devoted to improving the lives of people with disabilities (they should).
>
> It also does not express my belief that I'm 'entitled to a perfect child.' […]
>
> It expresses only the fact that given a choice, I would rather my child did not have a disability. That's all. (Baily 2000)

In terms of the promotion of prenatal testing for selective abortion, Bailey's objection might be that medical professionals who promote prenatal testing for selective abortion may do so for many reasons, including, e.g., informing their expecting patients' of their choices and thereby respecting their patients' ability to make their own informed decisions. In other words, they need not express disvalue for the lives of extant disabled people. According to this objection, then, it's not clear why we should accept the interpretation that promoting prenatal testing necessarily or even generally expresses the devaluing message or meaning that supporters of the expressivist objection claim it expresses.

DeGrazia's and Baily's objections appeal to the lack of necessity of an act's expressing a certain specific meaning. In light of these objections, supporters of the expressivist objection respond that regardless of whether or not using or promoting prenatal testing for selective abortion *necessarily* expresses a devaluing message to and about extant disabled people, and regardless of whether or not the agent in either action *intends* to express a devaluing message, many extant disabled individuals nonetheless receive that message (Hofmann 2017, Edwards 2004).[49] But note that this response suggests a sufficient condition for what using or promoting prenatal testing expresses which renders the expressivist objection implausible—or so I think many of us would agree. In other words, the idea seems to be that that disabled individuals receive such a message is *sufficient* for the action of using or promoting prenatal testing for selective abortion to have sent that message. Consider the following case.

> ***A, B,* and *C.*** A promotes prenatal testing to her patient, *B*, because *B* is a carrier of Tay-Sachs disease (TSD).[50] *A* does this because he would like for *B* to know that she can use prenatal testing to learn whether or not her future child will have TSD and, if so, terminate the pregnancy in order to prevent her future child's suffering—as well as her own. *A* does not believe that disabled people are not worth saving, or that children with TSD are morally inferior to children without TSD. All he wants is for *B* to be informed about her options. Now, *C*, a disabled person, learns about *A*'s action and interprets it as sending

---

[49] See also Parens and Asch (2000).

[50] TSD involves the loss of a child's ability to see, hear, and move, along with intense pain and seizures, which result in death in early childhood.

a devaluing message such as “Disabled people should not be born,” or, equivalently in *C*’s opinion, “All disabled people should be killed.”

If all that’s required for us to say that prenatal testing expresses a general or specific devaluing message is the judgment of extant disabled people, then we would have to say that *A*’s action expresses that “All disabled people should be killed.” But that seems wrong. There must be more to what an action expresses than the interpreter’s judgment of what it expresses—or so many of us would agree.

Put differently, the point is that if all that’s required for an action to express a message of disvalue to and about a certain group of people is the judgment of members of the relevant group, then *any* action could in theory meet this criterion. A person born with spina bifida could interpret the promotion of taking folic acid to prevent spina bifida as expressing a devaluing message about extant people with spina bifida. Or, a person with osteoporosis could interpret the action of taking calcium supplements to avoid developing severe osteoporosis as expressing a message of disvalue about extant people with osteoporosis. But to claim that now the promotion of taking folic acid is morally objectionable because it expresses a devaluing message to and about extant people with spina bifida, or that the promotion of taking calcium supplements is now morally objectionable because it expresses a devaluing message to and about extant people with osteoporosis, seems extreme.[51]

---

[51] See Hofmann (2017) and Edwards (2004) put this point slightly differently. Hofmann writes, “Offence is different from harm, and persons with disabilities, like all others, must accept offences” (at p. 510). Furthermore, he writes, “offence has to be balanced against other issues: ‘one may choose to have an abortion in the knowledge that this will offend prolife

Again: there must be more to what an action expresses than what its interpreter takes it to express. The current formulation of the expressivist objection thus suggests an implausible condition for how to determine the kind of message that is expressed by an action.

I don't mean to claim that we now have reason to ignore that some or many disabled people do receive devaluing messages from the use or promotion of prenatal testing for selective abortion. My point is just that if prenatal testing does express a devaluing message about extant disabled people, this claim has not been adequately defended on a theoretical level. Now, I believe that the expressivist objection's core insight—that there are expressive grounds for objecting to the promotion of prenatal testing for selective abortion—can be preserved by recruiting some tools from the philosophy of language and devising a stricter set of criteria for what our actions express, and I embark on that project in the next section. But before turning to that project, let me highlight another shortcoming of the current formulation of the expressivist objection.

Recall that supporters of the expressivist objection agree that prenatal testing for selective abortion is morally objectionable only because of what it brings about for extant disabled people—e.g., it may further stigmatize them, reduce the amount of resources allocated to disability communities, etc. In response to this claim, opponents of the expressivist objection highlight that there is no evidence that the

---

groups, but it does not follow that the offence caused to these groups outweighs one's right to make such a choice'" (quoting Edwards).

promotion of prenatal testing for selective abortion actually expresses the attitudes that its supporters claim it expresses, so there is no evidence that the promotion of prenatal testing for selective abortion actually brings about the consequences that supporters of the expressivist objection claim it brings about. Moreover, critics suggest, even if the promotion of prenatal testing for selective abortion did express a negative attitude about the quality of life of disabled people, it does not follow that it also expresses an attitude about how disabled people should be treated. Attitudes about the quality of life of disabled people are conceptually distinct from attitudes about how such individuals should be treated (Shakespeare 2006, Glover 2008, Buchanan 1996, McMahan 2005).[52] Therefore, even if the promotion of prenatal testing for selective abortion did express an attitude such as "Life with disability is not worth living," it does not follow that it also expresses an attitude such as "People who have disabilities should be treated poorly." And if this latter attitude is not expressed, then it is doubtful that the promotion of prenatal testing for selective abortion would promote the poor treatment of extant disabled people—or so opponents of the expressivist objection maintain.

Supporters of the expressivist objection respond that even if there is no evidence that prenatal testing for selective abortion has negative consequences for extant disabled people, this does not amount to evidence that prenatal testing for selective abortion does *not* have such bad consequences. But, while true, this point offers no further support for the claim that promoting prenatal testing for selective

---

[52] Cf. Reinders (2000) and Asch (2003).

abortion, in virtue of expressing devaluing messages, brings about the negative consequences that supporters of the expressivist objection claim it does.

This highlights what I take to be another shortcoming of the current formulation of the expressivist objection. Suppose, for the sake of argument, that supporters of the expressivist objection are wrong that the promotion of prenatal testing for selective abortion further stigmatizes extant disabled people or diminishes the resources that disability communities receive. Would the promotion of prenatal testing for selective abortion, absent such bad consequences, cease to be morally objectionable? I don't believe so. Imagine the following case:

> **D**. *D* is a medical professional who firmly believes that people who are born with any kind of disability are morally inferior to people who are not born with a disability. For example, she believes that the value of the life of people who are born with a disability is lower than the value of the life of people who are not born with a disability, and that the interests of people who are born with a disability are owed less consideration than the interests of people who are not born with a disability. For this reason, she makes sure that all of her patients understand that they should use prenatal testing and terminate their pregnancy if their future child will be disabled. Suppose that *D*'s action expresses an attitude such as "Disabled people have lower moral worth and should not be born," and that in expressing this attitude, her action has no negative consequences for extant disabled people. That is, it does not cause any psychological harms (e.g., no disabled person learns about it), civic harms (e.g., it does not further stigmatize people who are disabled), or material

harms (e.g., it has no effect on the number of social resources that extant disabled people receive).

Now, insofar as supporters of the expressivist objection maintain that what's morally objectionable about prenatal testing for selective abortion amounts solely to its negative consequences for extant disabled people, they would have to maintain that *D*'s action is not morally objectionable. After all, it causes no psychological, civic, or material harms to such individuals. But many of us would still find *D*'s promotion of prenatal testing morally problematic; it's just that what makes it morally problematic may not lie in the consequences of her action. I will return to this point in Section 3; for now, all I want to suggest is that even if the promotion of prenatal testing for selective abortion does not have any negative consequences for disabled people, it does not follow that it is not morally objectionable. But the current formulation of the expressivist objection, because it focuses solely on the consequences of this action, obscures the idea that actions can be morally objectionable on non-consequentialist grounds, too.

To sum up, the current formulation of the expressivist objection is imprecise, both with respect to what the promotion of prenatal testing for selective abortion purportedly expresses, and with respect to the reasons why promoting prenatal testing for selective abortion is morally objectionable. Moreover, as a result of this imprecision, the current formulation of the expressivist objection adopts an implausible sufficient condition for what an action such as the promotion of prenatal testing for selective abortion expresses, and obscures other, non-consequentialist wrongs that such an action may involve.

A more sophisticated formulation of the expressivist objection will offer a stricter set of criteria for what an action such as promoting prenatal testing for selective abortion might express. It will also have a wider scope regarding what's relevant for morally evaluating such an action. The rest of the paper will lay out the basics for such an improved formulation of the expressivist objection. Before proceeding, however, let me clarify that in engaging in such a project, I do not intend to imply that disabled people's judgments about what the promotion (or use) of prenatal testing for selective abortion expresses are insignificant. Nor do I intend to imply that the potential consequences of expressing such a message are not serious. Rather, my goal is to offer an improved formulation of the expressivist objection that both better explains what an action such as promoting prenatal testing for selective abortion expresses, and better captures when such an action is morally objectionable.

2. What our actions express

I believe that we can begin to offer a more adequate defense of what an action such as promoting prenatal testing for selective abortion expresses by appealing to two concepts from the philosophy of language: *illocutionary acts* and *implicatures*. In what follows, I argue that the promotion of prenatal testing for selective abortion can be conceived of as an illocutionary act that can occur over the course of a clinical interaction. Moreover, I argue that when a medical professional engages in such an act, she sometimes expresses precise meanings (or messages) about extant people who are disabled. This improved formulation of the expressivist objection—which I will call *the expressivist objection**—sidesteps the DeGrazia-Baily objection discussed in Section 1 and lays the groundwork for morally evaluating the promotion of prenatal testing for selective abortion more thoroughly.

Illocutionary acts, or *illocutions*, are verbal or written utterances which themselves constitute some intended action (Bach 1998, Austin 1962, Searle 1969).[53] For example, if I utter "I promise I won't eat the rest of the marionberry pie," I engage in the act of promising. My utterance itself constitutes my intended action. Similarly, if a priest says, "I now pronounce you husband and wife," his action itself constitutes the act of marrying.

Now, an utterance need not be explicit in order to count as an illocution. If I utter "Yes" when my daughter asks me "Do you promise you'll take me to Lizzy's once I finish my homework?" my utterance counts as an act of promising even if I don't state "I promise I will take you to Lizzy's once you finish your homework" explicitly. Furthermore, the authority of the speaker makes a difference in whether or not an action counts as an illocutionary act; we would not say of someone who utters "I hereby pronounce you husband and wife," but does not have the authority to marry two individuals, that she has married them.

---

[53] Importantly, for Austin, illocutionary acts refer to what an agent does *in* uttering what he utters—an action that does not require an effect on the interpreter. For example, stating, asserting, or suggesting are all illocutionary acts. These acts are, for Austin, distinct from *perlocutionary* acts. Perlocutionary acts refer to what an agent does *by* uttering what he utters—e.g., persuading, convincing, and scaring are all examples of perlocutionary acts. Note that perlocutionary acts *do* require an effect on the interpreter.

Suppose, then, that a nurse practitioner utters the words “I hereby promote the use of prenatal testing for selective abortion.” I take it that this utterance would, quite straightforwardly, constitute his intended action of promoting the use of prenatal testing for selective abortion. But notice that the nurse does not have to utter these words explicitly in order for his utterance to count as an illocutionary act of promoting prenatal testing for selective abortion. Similarly for the following case.

> ***E and F.*** During her first prenatal visit, *E*’s doctor, *F*, says something along the lines of:
>
> > Thanks to various recent technological advances, at this point in the pregnancy you have the option of doing some tests to learn more about the fetus. This testing is not high-risk at all, and the benefits of doing it outweigh the drawbacks, particularly because it gives you the option of using the test results to make a decision about whether or not to continue the pregnancy—say, if the fetus has some kind of chromosomal abnormality, such as Down syndrome, and you don’t want to have a child who won’t even be able to go to the bathroom on his own.

Here, *F* does not utter the words “I hereby promote prenatal testing for selective abortion,” yet his action arguably does have the illocutionary force of promoting the use of prenatal testing for selective abortion.[54] My goal here is not to

---

[54] Note that whether or not I’m right will depend in part on what kind of action *promoting* is. One might think that the act of promoting something is best described as an act that requires an effect in the interpreter—a perlocutionary act such as persuading, convincing, or scaring.

offer a set of necessary and sufficient conditions for why *F*'s action amounts to the illocutionary act of promoting prenatal testing for selective abortion, although it's worth noting two promising points of departure to that end: that *F* has medical authority in this context (i.e., he is a prenatal medical professional), and that he presents prenatal testing for selective abortion in a positive light: e.g., "thanks to recent technological advances," "the benefits of prenatal testing outweigh its drawbacks," etc. Rather, my point is that we can conceive of the act of promoting prenatal testing for selective abortion as an illocutionary act—one that can occur during a clinical interaction between a medical professional and a prospective parent.

What I want to highlight next is that in addition to presenting the option of using prenatal testing in a positive light, *F* communicates other kinds of information,

---

One might then say that *in* uttering what he utters, *N states* (illocutionary act) that there are some tests *O* can opt for, and that such tests are not dangerous. Moreover, one might say that *by* uttering what he utters (perlocutionary act), *N* devalues or, e.g., expresses disvalue, for extant disabled people. I'm not sure about the virtues of this route, but it seems to me that the act of promoting something is not one that requires an effect in the interpreter. That is, I think that promoting something is an illocutionary act—an action that resembles more the act of stating rather than the act of, say, persuading. Nevertheless, what seems at least a little bit plausible is that *by* uttering what she utters, *N* devalues extant disabled people. But here the question becomes a question of what kind of act *devaluing* is. And crucial in answering that question is whether or not someone can devalue another without having an effect on him. I think so, but if I'm wrong, then the act of devaluing another might best be described as a perlocutionary act.

too. For example, he communicates that *E* “has the option of doing some tests to learn more about the fetus,” that “this testing is not high risk at all,” and that “the benefits of doing [this testing] outweigh the drawbacks.” Importantly, however, he also communicates implied meanings or messages—that is, information that he does not state explicitly but that he nonetheless communicates. Philosophers of language call this kind of information *implicature*. On a Gricean picture of meaning, speakers convey two kinds of meaning in their utterances: a literal meaning and an implied meaning, or *implicature* (Grice 1957). Whereas a speaker’s literal meaning amounts to what he explicitly states, ‘implicature’ refers to what he communicates without explicitly stating it (Grice 1989).[55] How, then, are implicatures deciphered? According to Grice, implicatures are calculated from the intention of the speaker, the understanding of the interpreter, and the context in which the utterance occurs.[56] The following example will prove clarificatory:

---

[55] For Grice, implicatures can be conventional, or semantic, or conversational. I here focus only on the latter.

[56] This is an uncontroversial understanding of Grice. Grice himself writes of conversational implicature that:

> A man who, by (in, when) saying (or making as if to say) that p has implicated that q, may be said to have conversationally implicated that q provided that (i) he is to be presumed to be observing the conversational maxims […]; (ii) the supposition that he is aware that, or thinks that, q is required in order to make his saying or making as if to say p (or doing so in those terms) consistent with this presumption; and (iii) the speaker thinks (and would expect the hearer to think that the speaker thinks) that it is

> ***G and H.*** Carrying with him an empty plastic water bottle, *G* arrives at *H*'s home and asks her "Where do you keep your recycling?"; *H* replies "Ugh, I'm too lazy for that! The trash is under the sink."

In this example, each agent communicates a literal meaning and an implied meaning, or implicature. The literal meaning that *G* communicates is the content of the question "Where do you keep your recycling?" Moreover, in part because he is holding a plastic bottle, the implicature he communicates is something along the lines of, "I wish to dispose of this plastic bottle." *H* also communicates a literal meaning: namely that she is too lazy for a recycling bin and that the trash is under the sink; and, in part because she understands that *G* is trying to dispose of something, she also communicates an implicature: something along the lines of "I don't recycle; just use the trash bin."

What is the relevance of this discussion for the promotion of prenatal testing for selective abortion? Let us return to the case of *E and F* and reexamine the second half of what *F* says. Here it is reproduced:

> [Running these prenatal tests] gives you the option of using the test results to make a decision about whether or not to continue the pregnancy—say, if the fetus has some

---

> within the competence of the hearer to work out, or grasp intuitively, that the supposition mentioned in (ii) is required. (Grice 1989, 30)

For a more controversial understanding of Grice, according to which implicature can be unintentional—which I discuss in Section 4—see. Lassiter (2012).

> kind of chromosomal abnormality, such as Down syndrome, and you don’t want to have a child who won’t be able to even go to the bathroom on his own.

Note that *F* arguably implies—or, to use the technical term, *implicates*—something along the lines of “People with Down syndrome have no independence.” Now, I don’t intend to offer a set of necessary and sufficient conditions for how to determine exactly what *F* implicates here. But it’s worth noting that what *F* communicates by implicature will depend at least in part on the context surrounding this interaction—namely, a clinical interaction about prenatal testing. The point just is that clinical interaction—like many other verbal interactions—can involve communicating both explicit and implicit meanings. In the *E and F* example, *F* engages in the action of promoting prenatal testing for selective abortion and thereby communicates both a literal meaning and an implied meaning or message. Moreover, at least some of the implied meanings or messages that he communicates are about extant disabled people—namely, something along the lines of “People with Down syndrome have no independence.”

We can now begin to reformulate the expressivist objection—at least with respect to what the action of promoting prenatal testing for selective abortion might express:

> **The expressivist objection***: The promotion of prenatal testing for selective abortion is an illocutionary act that can occur over the course of clinical interaction. When medical professionals engage in this action, they sometimes

express implied messages or meanings that can be about extant disabled people.

If this is right, then the expressivist objection* sidesteps the DeGrazia-Baily objection discussed in Section 1. According to the DeGrazia-Baily objection, you'll recall, promoting prenatal testing for selective abortion is a nonverbal action that does not have any agreed-upon meanings. It thus need not express the messages the expressivist objection claims it expresses. But according to the expressivist objection*, the promotion of prenatal testing for selective abortion *is* a verbal action and *does* express precise messages or meanings. It thus can—at least sometimes—express precise enough messages or meanings about extant disabled people.

With this framework in hand, we can now ask: Do medical professionals who promote prenatal testing for selective abortion act wrongly? This is the question I now turn to.

3. *When our actions express vague or misleading attitudes*

You'll recall from Section 1 that supporters of the expressivist objection differ with respect to the exact content of the messages that promoting prenatal testing for selective abortion purportedly expresses. Some claim it is roughly that "Disabled people's lives are not worthy of support"; others claim it is that disabled people should not be, or should not have been, born. Now, given the framework developed in Section 2, I think it's plausible to suppose that medical professionals may sometimes express a meaning or message such as "Disabled people's lives are not worthy or support" or "Disabled people should not be born" when they promote prenatal testing for selective abortion—especially when they recommend testing for various disabilities with the aim of selective abortion—so I will assume this in what follows. What I want to examine is whether and when the promotion of prenatal testing for selective abortion is wrong in virtue of expressing these attitudes and others. As I will argue, promoting prenatal testing for selective abortion is wrong when it involves violating a duty to obtain informed consent. It thus can be wrong even if it does not bring about any negative consequences for extant disabled people.

Bioethicists tend to agree that medical professionals have a duty to obtain informed consent from their patients before running certain tests on them. One of the main goals of obtaining informed consent is to ensure that the patient makes a voluntary and informed decision regarding her medical treatment, with a reasonable understanding of the consequences of her options. In prenatal care settings involving the promotion of prenatal testing for selective abortion, obtaining informed consent would thus seem to involve ensuring that the expecting parents understand that if they opt for this testing and their future child receives a positive diagnosis for some disability, they might face the tough moral choice of whether or not to continue the pregnancy.[57] Thus, it is important that the information that the medical professional communicates while promoting prenatal testing for selective abortion is not false or misleading, e.g., by conveying unreflective stereotypes. If such information is false or misleading, then the expecting parents would not be able to make an informed decision about whether or not to use prenatal testing and the medical professional

---

[57] For an interesting discussion on abortion and moral fallibilism, see Moller (2011).

would not be acting in accordance with a duty to obtain informed consent. For example, imagine that a medical professional says to her patient, “You can’t legally opt out of this testing.” In this case, we would agree that this medical professional would *not* be acting in accordance with a duty to obtain informed consent from her patient to run prenatal tests that might present her with a choice of whether or not to terminate the pregnancy; she would be coercing her patient by appealing to the authority of the law.

Now, I think that many of us would also agree that the medical professional would not act in accordance with a duty to obtain informed consent if the information she communicates *by implicature* is false or misleading. For example, suppose that a medical professional promotes prenatal testing for selective abortion and thereby expresses, or communicates by implicature, the meaning or message “You can’t legally opt out of this testing.” Although this medical professional does not *state* this information explicitly, she still communicates it—albeit by implicature. Moreover, insofar as this information is false and misleading, this medical professional violates a duty to obtain informed consent from her patient. And she does this because even if she only expresses this information, she nonetheless communicates it—or so I think most of us would agree.

What I’m suggesting, then, is the uncontroversial idea that medical professionals can violate a duty to obtain informed consent not only by explicitly stating false or misleading information, but also by implicating it. For the purposes of the expressivist objection*, however, it’s what the medical professional *expresses*—again: what he communicates by implicature—that’s relevant for morally evaluating his action. Thus, we can determine whether or not a medical professional acts wrongly by looking at whether or not what he expresses, or implicates, while promoting prenatal testing for selective abortion is inaccurate or misleading. If it is, then in expressing such information, he violates a duty to obtain informed consent from his patient.

Let us reexamine the case of *E and F* from Section 2. In that example, *F* expresses or implicates the meaning or message that “People who have DS have no independence.” Thus the question is: Does *F* violate a duty to obtain informed consent and thereby act wrongly when he promotes prenatal testing for selective abortion to *E*? Now, qualitative data on the daily lives of adults who have DS are scarce, but a recent study surveyed family and caregivers of adults with DS, most of whom lived in the Midwest, in order to gather such data (Matthews, et al. 2018). According to this study, over eighty percent of adults with DS had “no difficulty”—that is, needed no assistance—when engaging in everyday tasks such as moving around in or out of the house, eating meals independently, and dressing and undressing themselves. Additionally, the majority of these adults were reported to have “no difficulty” using the toilet independently, grooming themselves, and washing themselves, and around sixty-percent were reported to have “no difficulty” carrying out domestic activities, such as making their beds and light housecleaning, or being left home alone for a period greater than two hours.

I cite this study not to suggest that the message “People who have DS have no independence” is obviously false, or that in general people with DS *are* independent. After all, according to the same study, seventy-eight percent of adults with DS were reported to live with family, and only nine percent were reported to live in an independent home. Furthermore, over half of these adults were reported to have "a lot of difficulty" going out alone, and over seventy-three percent had "a lot of difficulty" managing daily finances.[58] Rather, I cite this study to suggest that a message such as "People who have DS have no independence" is vague or misleading at least in part because what counts as independent behavior is not clear.[59]

Returning to the example of *E and F*, though, the idea is that if a meaning or message such as "People who have DS have no independence" is inaccurate or misleading, in expressing it, *F* violates a duty to obtain informed consent from *E*, insofar as *E* would be making a decision to opt for or against prenatal testing based, in part, on the idea that a lack of independence might constitute a reason for terminating her pregnancy. In other words, *F* prevents *E* from making an informed decision about whether or not to opt for prenatal testing.

If I'm right, we can summarize this section by offering a more complete formulation of the expressivist objection*—one that, on the one hand, addresses what the promotion of prenatal testing for selective abortion might express, and, on the

---

[58] "A lot of difficulty" means that they were unable to.

[59] There is also evidence that two main factors associated with greater employment access and higher self-sufficiency rates among people who have DS are familial and social support. In particular, studies suggest that one reason we don't see more people who have DS having full-time jobs may be more due to their lack of social support than to their cognitive ability. See Sloper and Turner (1996) and Eggleton et al. (1999).

other, explains why such an action can, in virtue of what it expresses, be morally objectionable:

> **The expressivist objection***: The promotion of prenatal testing for selective abortion is an illocutionary act that can occur over the course of clinical interaction. When medical professionals engage in this action, they sometimes express implied messages or meanings that can be about extant disabled people. When these expressed messages or meanings are inaccurate or misleading, medical professionals act wrongly, as they violate a duty to obtain informed consent to do prenatal, disability-related testing on their expecting patients.

Note another important advantage of the expressivist objection* over the original expressivist objection. Whereas the original expressivist objection maintains that the promotion of prenatal testing for selective abortion is morally objectionable because, in expressing a devaluing message, it brings about harmful consequences for extant disabled people, the expressivist objection* maintains that the wrong lies in the *speech act* that the medical professional engages in. In other words, according to the expressivist objection*, the promotion of prenatal testing for selective abortion can be wrong in virtue of what it expresses even if no bad consequences obtain for extant disabled people. This does not mean that on the expressivist objection* the consequences that the promotion of prenatal testing for selective abortion, in virtue of what it expresses, brings about are irrelevant to whether or not this action is morally objectionable. The expressivist objection* can grant that bringing about harmful consequences for extant disabled people would be bad. But the point is that the expressivist objection* does not need to appeal to those consequences of promoting prenatal testing for selective abortion in order to account for the moral wrongness of this action. All it needs to appeal to is a duty-based, agent-centered consideration that remains even if no bad consequences obtain for extant disabled people.[60]

---

[60] Interestingly, expressing a generic attitude such as "People who have DS have no independence" may be morally objectionable independently of whether or not it's false or misleading. For example, the generic "Tigers are striped" often communicates the essentialist idea that members of the kind 'tiger' share a 'distinctive, non-obvious and persistent property or underlying nature that causally grounds' their stripedness. Recent findings also suggest that while generics are a significant mechanism by which we transmit social information, they may be insidious in more ways than one might think: for example, they may be implicated in the foundations of social prejudice involving, for example, race, gender or sexuality, or a combination of such categories (Leslie 2014, 209). And this may seem especially alarming given experimental evidence which suggests that there may be a causal link between essentialized beliefs and prejudice (Keller 2005). I don't mean to suggest that this research gives us reason to claim that medical professionals who communicate, by implicature, generics about people with DS or about people with disabilities more generally, contribute to social prejudice against people with DS or against people with disabilities more generally. After all, this research is about explicit generic *statements*, not about *implicit* generic attitudes or messages. But I do want to suggest that it is at least plausible to think that a *message* such as "People who have DS have no independence" might suggest the essentialist generic: namely, that people who have DS, *in virtue of having DS*—and not, for example, in virtue of other health conditions or environmental factors—have no independence.

4. *Implicature and hyper-implicature*

Thus far, this paper has examined and rejected the original formulation of the expressivist objection—a view that locates the wrong of promoting prenatal testing for selective abortion strictly in the victim of the action. In its place, I offered the expressivist objection*—a view that locates the wrong of promoting prenatal testing for selective abortion in the *speech act* that medical professionals who promote prenatal testing for selective abortion engage in. An interesting puzzle arises when we consider the possibility that the wrong lies neither in the victims of the action nor in the speech act that medical professionals engage in, but, rather, in the agents themselves: specifically, in their *belief sets*. In this section, I examine what I call *hyper-implicature*—that is, information that agents *unintentionally* communicate via their utterances—in order to argue that medical professionals who promote prenatal testing for selective abortion may sometimes unintentionally reveal information about their own disability-related beliefs. As I will argue, such cases suggest that medical professionals may sometimes engage in a moral wrong even if they don't, in virtue of what they reveal, bring about bad consequences or violate a duty to obtain informed consent.

Recall from Section 2 the case of *G and H*, the aim of which was to highlight that on a Gricean picture of meaning, agents communicate at least two kinds of meanings: literal meanings and implicatures. Here it is again:

> ***G and H.*** Carrying with him an empty plastic water bottle, *G* arrives at *H*'s home and asks her "Where do you keep your recycling?"; *H* replies "Ugh, I'm too lazy for that! The trash is under the sink."

The idea was that *G* communicates the literal meaning "Where do you keep your recycling?" and implicates, roughly, that he wishes to dispose of the plastic bottle; and that *H* communicates the literal meaning that she is too lazy for a recycling bin, that the trash is under the sink, and implicates, roughly, that she does not recycle and that *G* should just use the trash bin. Notice, though, that this is not all that *G and H* communicate in this interaction. Given the literal and implied meanings that they communicate, we can infer some additional information about each agent's belief set—or so I believe.

For instance, given the literal and implied meanings that *G* communicates, we can infer that he probably believes that having recycling bins is the norm, or that he thinks of *H* as someone who recycles. If this were not the case, it would not be clear why he would ask "Where do you keep your recycling?" to communicate that he wishes to dispose of the plastic bottle—instead of, say, something like "Do you have a recycling bin?" or "Do you recycle?" Moreover, in the case of *H*, we can infer that recycling is not something that she is concerned with—or, at the very least, that recycling is not a priority for her. If recycling did concern her, or if it was a priority for her, it would be odd for her to communicate to *G* that he should just use the trash bin because she is too lazy to recycle. In short, the idea is that sometimes we can infer information about an agent's set of beliefs or value judgments from the meanings she communicates. In the case of *G and H*, we can infer that *G* believes that *H* has a recycling bin, and that *H* does not care much about recycling. In what follows I will refer to this revealable kind of information as *hyper-implicature*—‘hyper-’ because although it is implicature-like, it is beyond just implicature in that it can be inferred, in part, from what is implicated.[61]

Note that what an agent communicates by hyper-implicature is distinct from the *meanings* that she communicates—at least according to the Gricean picture of meaning I’ve been considering. Philosophers of language tend to agree that, on a Gricean picture of meaning, the intention of the speaker is an important factor in the meanings that he communicates. It would thus be a mistake to say that *G* expresses the *meaning* or *message* “Recycling is the norm” or “*H* recycles,” or that *H* expresses the *meaning* or *message* “Recycling is not something that concerns me.” A more appropriate way to describe what is communicated, I submit, is that each agent *reveals* (although not intentionally) a belief or value judgment that he or she holds. Thus, *G reveals* (although not intentionally) that he believes that *H* recycles, and *H reveals* (although not intentionally) that she does not value recycling all that much.

If this is right, it’s conceivable that, in the same way that *G* and *H* reveal information about their own set of beliefs or value judgments, medical professionals also sometimes reveal information about *their* own set of beliefs or value judgments. It also seems plausible that, in contexts involving prenatal testing and selective

---

[61] Note that although philosophers of language tend to agree that implicatures are intentional, at least some philosophers think that implicatures can also be unintentional. See, for example, Lassiter (2012).

abortion, these hyper-implicatures can be about extant disabled people or about life with disability more generally.

Now, one might highlight that if these hyper-implicatures are inaccurate or misleading, a medical professional who reveals them while promoting prenatal testing for selective abortion may violate a duty to obtain informed consent from his patient in the same way that he would violate a duty to obtain informed consent from his patient if he *expressed*—or, to use the term I'm proposing: *hyper-implicated*—inaccurate or misleading information. For example, one might worry that a medical professional who, in promoting prenatal testing for selective abortion, hyper-implicates "People who have DS have no independence" would violate a duty to obtain informed consent from his patient to run a prenatal test. I think this possibility is worth exploring—especially if we consider that patients tend to think that medical professionals' training and qualifications render their beliefs (and recommendations) authoritative. But I will not argue for this here. What I want to examine instead is that even if medical professionals who promote prenatal testing for selective abortion and thereby hyper-implicate inaccurate or misleading beliefs or value judgments about extant disabled people do not, in virtue of revealing such information, violate a duty to obtain informed consent from their patients, and even if they do not, in virtue of what they reveal, cause any harms to extant disabled people, such cases may still contain a moral wrong.

Consider the following example:

> ***J and K.*** *J* is a medical professional who works at a reproductive clinic where prenatal testing for selective abortion is often promoted. In part as a result of working at this reproductive clinic, *J* has come to believe that “People who have disabilities, in virtue of having disabilities, have less moral standing than individuals who lack disabilities.”[62] For example, *J* has come to believe that the interests of people who have disabilities are owed less consideration than the interests of people who lack disabilities. Today, *J* is promoting prenatal testing for selective abortion to his patient, *K*, who is disabled. In so doing, he unintentionally reveals that he believes, or hyper-implicates, that “People who have disabilities, in virtue of having disabilities, have less moral standing than individuals who lack disabilities.” No harmful consequences for extant disabled people obtain.

In this example, while *J* does not intend to reveal that he believes that “People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities,” he nonetheless reveals that he believes this (given what he intentionally communicates). Nevertheless, *J* does not, in virtue of what he reveals, bring about any harmful consequences for extant disabled people. Now, remember that we are setting aside the possibility that *J* violates a duty to obtain informed consent in virtue of what he reveals. Does a wrong occur even though no harms occur, and even though no duty to obtain informed consent is violated? I believe that if it does, then at least part of the explanation lies in that *J and K* contains a doxastic

---

[62] I’ve based this attitude on one of DeGrazia’s examples, at p. 103.

wronging: a kind of wrong that occurs in virtue of what one agent believes about another.

Philosophers tend to agree that it is actions—not beliefs—that wrong. If I treat another person merely as a means, I wrong her, but if I merely *believe* that I can treat her merely as a means, without ever so treating her, I do not wrong her. Contrary to this idea, however, Rima Basu and Mark Schroeder (2019) have recently offered a prima facie case for the existence of doxastic wrongs.[63] As Basu and Schroeder put it, a doxastic wronging occurs "if one person wrongs another in virtue of what she believes about him." Consider their own example:

> [S]uppose that you have struggled with an alcohol problem for many years, but have been sober for eight months. Tonight you attend a departmental reception for a visiting colloquium speaker, and are proud of withstanding the temptation to have a drink. But when you get home, your spouse smells the wine that the colloquium speaker spilled on your sleeve while gesticulating to make a point, and you can see from her eyes that that she thinks you have fallen off of the wagon. If you are like us, then you will be prone to feel wounded by this. Yes, you have a long history of falling off of the wagon, and yes, there is some evidence that this time is another. You can see how it could be reasonable for someone to draw this conclusion. But it

---

[63] Basu and Schroeder's aim, though, is to defend the existence of doxastic wrongs from some likely objections, and argue that if they are right, then the existence of doxastic wrongs offers an argument for a kind of pragmatic encroachment on epistemic rationality.

> still hurts – not least because in your eyes, tonight was an achievement to stay on the wagon despite adverse circumstances. (Basu and Schroeder 2019)[64]

Basu and Schroeder extract three characteristics of doxastic wrongs from this example. First, they indicate, that you feel wounded by your spouse's having the belief that you have fallen off the wagon suggests that your spouse isn't just wrong in believing this, but that she has wronged *you*.[65] Second, the wrong that occurs is not a matter of what led your spouse to have the belief that you fell off the wagon—i.e., that you have this history and that tonight you came home smelling like wine. Nor is it a matter of what actions follow from her belief—e.g., if she were to falsely accuse you. Rather, it is a matter of her *having the belief*.[66] Finally, the wronging occurs in

---

[64] In support for the existence of doxastic wrongs, Basu and Schroeder also cite the Book of Common Prayer's Rite II version of the Eucharist, in which the congregation confesses that "we have sinned against you in thought, word, and deed."

[65] Note that this point is not convincing. Although people may feel wounded when they are wronged, that one might feel wounded when one is wronged need not entail that one has been wronged. In other words, feeling wounded need not track wrongings. Nevertheless, one need not agree with this point of Basu and Schroeder's analysis to consider the possibility that doxastic wrongs exist.

[66] As Basu and Schroeder put it, "For example, it would feel insincere and unsatisfying if [your spouse] apologized for the upstream act of not investigating more carefully before forming this belief, but continued to believe it anyway, or if she apologized for the downstream act of revealing her belief to you by the look in her eyes, but not for the belief itself."

virtue of *what* your spouse believes: she believes something false about you. In short, Basu and Schroeder believe that doxastic wrongs are directed, that they occur in virtue of holding a certain belief, and that they are wrongs in virtue of what is believed (something false).

Notice that *J and K* seems to fit these criteria. In believing that "People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities," *J* seems to *wrong K* (and people like *K*). Moreover, what seems morally objectionable about this case is not what leads *J* to have this belief—e.g., his work environment. Nor is it what follows from his belief, since no bad consequences obtain in virtue of his having or revealing this belief. Rather, it is *that J believes that* "*People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities*" that seems morally objectionable. Finally, the wrong that occurs is a wrong in virtue of what is believed: that *K* (and people like *K*), in virtue of their disabilities, have less moral standing than people who lack disabilities.

What I'm proposing, then, is that if we find the case of *J and K* morally disconcerting, it might be because *J believes* that "People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities"—and not because *J* brings about negative consequences for extant disabled people in virtue of what he reveals (he doesn't), nor because he violates a duty to obtain informed consent in virtue of what he reveals (since we are setting this possibility aside). But what is it about *J*'s belief that wrongs *K*?

Basu and Schroeder would suggest that part of the wrong lies in that *J*'s belief is false.[67] But I want to consider a different possibility that I find plausible: that *J* wrongs *K* because his belief involves a moral error. Others have considered the idea that the moral error involved in a doxastic wrong lies in the *process* by which a morally questionable belief is formed. For example, Basu (2019) argues that beliefs wrong when they violate an obligation of "what we epistemically owe to each other."[68] What we epistemically owe to each other, she maintains, is to treat each other, in *thought*, not as things that can be studied and predicted but as persons. When we fail at this, we make a moral error.

---

[67] In line with this point, one might think that *J*'s belief involves an *epistemic* error. See (Schroeder 2018). To clarify, the idea is that if doxastic wrongs exist, then there may be beliefs that are morally wrong, and this would seem to impact what we take to be *epistemically* justified. Schroeder appeals to stakes-related reasons in order to argue that beliefs wrong only when they falsely diminish and that they falsely diminish when there are moral reasons that make it harder to justify a belief. Such moral reasons lie in the risks of having a belief that's false, such that "believing it can be unacceptably risky, even in cases in which it happens to be true." Examining Schroeder's account appropriately is well beyond the scope of this paper, but it's worth noting that it would seem to suggest that the wrong in *J and K* stems from the unacceptable risk of *J*'s believing that "People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities," such that if *J* is wrong, his belief would "guarantee an *epistemic* fault."

[68] Basu invokes Sherlock Holmes as an example: Holmes "observes everyone as objects to be studied, predicted, and managed" rather than as objects that are owed a certain kind of consideration.

What I want to examine, though, is that the moral error in a doxastic wrong lies not (or not only) in an obligation that is violated in the *process* in which the belief is formed, but (also) in the *kind* of belief that is held. Consider the following case:

> ***L and M.*** While promoting prenatal testing for selective abortion, *L* reveals to his patient *M*, who has DS, that he believes that "People who have DS can't live alone." No harmful consequences occur as a result of *L*'s revealing that he holds this belief, and *L* does not violate a duty to obtain informed consent from *M*.

To clarify, *L* does not explicitly say that "People who have DS can't live alone," nor does he implicate it. Rather, he communicates it by hyper-implicature. Nevertheless, I think that many of us would agree that unlike *J* in the example above—who seems to *wrong K* in virtue of revealing that he believes that "People who have disabilities, in virtue of their disabilities, have less moral standing than people who lack disabilities"—*L* does not seem to *wrong M* in virtue revealing that he believes that "People who have DS can't live alone"—even if this belief is false.[69]

---

[69] It's worth noting that one might worry that given his profession, *L* has an obligation to learn basic pertinent facts about people who have DS, such that in believing that "People who have DS can't live alone," *L* might fail to fulfill this duty and might thereby fall short of giving sufficient concern to his patients. But I will set this worry aside in part because according to the study from above, the majority of adults with DS (seventy-eight percent) were reported to live with family, and only nine percent were reported to live in an independent home.

This difference in judgment, I believe, stems from a familiar distinction between evaluative judgments and descriptive judgments.[70] To clarify, *J*'s belief is quite straightforwardly an evaluative judgment: it is a moral judgment about the moral standing of another person or group of people in virtue of their having a disability. But *L*'s belief is not (or not clearly) an evaluative judgment. *L*'s belief seems better described as a descriptive judgment about what life with DS involves. My point is not that there is a clear-cut distinction between an evaluative judgment such as *J*'s and a descriptive judgment such as *L*'s. After all, one might think that an attitude such as "People with DS can't live alone" at the very least suggests a judgment about the value of living alone. My point is simply that if this distinction holds up, one might think that the reason *J* wrongs *K*, and *L* does not wrong *M*, is that *J* makes a *moral error*—not just a descriptive or epistemic error such as *L*'s.

An adequate defense of this position is beyond the scope of this paper, but let me note one plausible way we might proceed. We might claim that *J*'s belief, but not *L*'s belief, involves a kind of category mistake about what we are morally permitted or not permitted to believe about other agents. Stephen Darwall's (1977) notion of

---

[70] Hare (1952) makes this kind of distinction, but he is concerned with the relationship *between* descriptive and evaluative judgments. As he argues, descriptive remarks are often given as a reason for making an evaluative remark: for example, "This strawberry is sweet" (a descriptive judgment) is often given as a reason for "This is a good strawberry" (at p. 111). For the purposes of this paper, however, this relationship is not relevant; all I want to highlight is that there is a distinction between factual or descriptive judgments and evaluative judgments.

recognition respect offers some guidance here. As I read Darwall, recognition respect involves a disposition to weigh appropriately in one's deliberations the fact that another agent is a person, such that that person's status restricts the set of *actions* that are permissible to an agent. Invoking this notion, we might contend that recognition respect might also involve a disposition to weigh appropriately in one's moral deliberations the fact that another agent is a person, such that that person's status restricts the set of *attitudes* or *judgments* that are permissible to hold about another person. Doxastic wrongs, then, could be understood as involving a failure of recognition respect. In other words, they could be understood as involving a disposition to give inappropriate weight to some irrelevant fact about another agent in one's moral deliberations. On this proposal, *J* might wrong *K* insofar as in believing that *K* has lower moral standing in virtue of her disability, he gives inappropriate weight to some irrelevant fact about *K* (i.e., her disability), thereby failing to respect her in the 'recognition respect' sense of the term. And in so failing, he makes a moral error. By contrast, in the example of *L and M*, it's not straightaway clear that *L* fails to 'recognition respect' *M* by believing that people like *M*, in virtue of having DS, can't live alone. This belief may be somewhat inaccurate or misguided, but it need not involve giving inappropriate weight to the fact that *M* is a person. In other words, it need not amount to wronging her.

If I'm right, we can summarize this section by offering a final formulation of the expressivist objection* that leaves room for the existence of doxastic wrongs:

> **The expressivist objection***: The promotion of prenatal testing for selective abortion is an illocutionary act that can occur over the course of clinical

> interaction. When medical professionals engage in this action, they sometimes express implied messages or meanings that can be about extant disabled people. When these expressed messages or meanings are inaccurate or misleading, medical professionals act wrongly, as they violate a duty to obtain informed consent to do prenatal, disability-related testing on their expecting patients. But medical professionals may also hyper-implicate morally erroneous attitudes about extant disabled people. In such cases, medical professionals might wrong extant disabled people even if they don't, in virtue of what they reveal, violate a duty to obtain informed consent from their patients.

Now, an opponent might ask: given the training that medical professionals in general and genetic counselors in particular receive, how likely is it that they will promote prenatal testing for selective abortion in the morally objectionable ways that I have described? If not at all likely, then how is the position I've presented interesting? I don't know how likely it is that medical professionals will express or reveal the morally problematic attitudes I have offered as examples throughout this paper. What I do know is that there is reason to be hesitant to accept the idea that because medical professionals are trained professionals, they are immune to moral error. Without even going into detail about our history of morally questionable medical research, consider that there is reason to worry that race[71] and gender[72] may play a role in the treatment that patients receive in healthcare settings, and that some medical professionals have reportedly refused to treat LGBT patients.[73] More relevant to the topic at hand, however, consider that some disabled people report questionable interactions with medical professionals, and that research suggests that older people who have a disability are more likely than those who don't have a disability to experience discrimination in healthcare settings.[74] While these data do not directly suggest that medical professionals who promote prenatal testing for selective abortion do so in the morally objectionable ways I have described, they do give us a reason not to assume that they don't.

In this section, my goal was to establish the possibility that medical professionals may engage in something that's morally wrong not only if they engage in the speech act of promoting prenatal testing for selective abortion when so doing violates a duty to obtain informed consent, but also if they 'merely' hold a certain kind of belief about people who have disabilities. On this view, the promotion of prenatal testing for selective abortion is sometimes a *manifestation* of a doxastic

---

[71] See, for example, Tello (2017).

[72] For instance, women may have a hard time finding a doctor who takes them seriously in their request for tubal ligation. See Mattoni (2017).

[73] See, for example, Green (2016).

[74] According to Rogers et al. (2015), "[o]ne out of five adults over the age of 50 years experiences discrimination in healthcare settings. One in 17 experience frequent healthcare discrimination, and this is associated with new or worsened disability by 4 years."

wrong. In other words, the promotion of prenatal testing for selective abortion is not, on this view, wrong in itself. Rather, it contains a moral wrong when it reveals that the agent holds a morally erroneous belief about another agent. Importantly, though, it's not the *revealing* of such a belief that's morally objectionable (although one might think that revealing such beliefs betrays basic decency); rather, it's the holding of such a belief that's morally objectionable. In other words, a medical professional may be engaged in something that's morally objectionable even if she never reveals that she holds a morally erroneous belief about another agent.

If I'm right, then it seems to follow that medical professionals in general—that is, not just medical professionals who promote prenatal testing for selective abortion—may wrong some of their patients merely by holding morally objectionable beliefs about them. While I will not argue for this here, I welcome this consequence of my view and hope that I have said enough to motivate an interesting line of thought that may help us better understand why a case like *J and K*, but not a case like *L and M,* seems to be morally objectionable.

5. Concluding Remarks

In this paper, I first argued that the current understanding of the expressivist objection faces various problems. Specifically, it is imprecise, both with respect to what the promotion of prenatal testing for selective abortion purportedly expresses and with regard to why promoting prenatal testing for selective abortion is morally objectionable. Moreover, as a result of this imprecision, the current formulation of the expressivist objection adopts an implausible sufficient condition for what an action such as the promotion of prenatal testing for selective abortion expresses, and obscures other, non-consequentialist wrongs that such an action may involve. I then offered an improved formulation of the expressivist objection: the expressivist objection*. This improved formulation conceives of the act of promoting prenatal testing for abortion on grounds of disability as an illocutionary act that involves both literal meanings and implied meanings, or implicatures. According to this expressivist objection*, when medical professionals express implied meanings that are inaccurate or misleading, they violate a duty to obtain informed consent from their patients and thereby act wrongly. Nevertheless, sometimes medical professionals hyper-implicate, or reveal (although not intentionally), that they hold inaccurate or misleading attitudes about people who have disabilities. Such cases seem to manifest a different kind of wrong: a doxastic wrong that is independent both of any consequences the medical professional may bring about in virtue of revealing a morally erroneous attitude, and of any violations of a duty to obtain informed consent that may occur as a result of what the medical professional reveals. In short, according to the expressivist objection*, medical professionals may engage in a moral wrong not only when they express inaccurate or misleading attitudes while promoting prenatal testing for selective abortion, thereby violating a duty to obtain informed consent from their patients, but also when they hold a morally erroneous belief about other agents, including their patients who have disabilities.

Thus, the expressivist objection* does not entail that the promotion of prenatal testing for selective abortion is *always* wrong. But it also does not entail that the promotion of prenatal testing for selective abortion can be morally permissible. All it entails is that the promotion of prenatal testing for selective abortion can be wrong. It’s thus worth remembering that not all cases of promoting prenatal testing involve the aim of terminating a pregnancy on grounds of disability; medical professionals may also promote prenatal testing with the end of learning more about the fetus, just in case the child will have a disability, so that expecting parents can be better prepared to welcome their child Medical professionals who wish to avoid the risk of engaging in a moral wrong may thus choose to promote the use of prenatal testing for the sake of learning more about the fetus or for the sake of respecting parental autonomy, rather than for the sake of terminating the pregnancy if there is a positive result for a disability. The promotion of prenatal testing for the sake of learning more about the fetus seems, at least prima facie, less likely to involve the expression or hyper-implicature of an erroneous or misleading attitude. Alternatively, medical professionals may choose to avoid making recommendations about prenatal testing altogether and just *offer* this service in a neutral way. But then it’s important to note that the act of *offering* prenatal testing is a different kind of action and may require an entirely separate analysis.

Chapter 4: Indirect Benefits and Double Jeopardy in the Allocation of Scarce, Lifesaving Resources

Introduction

The outcomes of allocating scarce, lifesaving resources, such as organs for transplantation, often involve varying kinds of benefits, including direct and indirect benefits.[75] In this paper, I will understand direct benefits to be those that benefit the recipient of a resource, e.g., by extending her life or improving her quality of life; and indirect benefits as those that, as a result of being allocated to a direct beneficiary, benefit third parties, e.g., by saving their lives or by improving socioeconomic conditions.[76] For example, imagine we had one dose left of a life-saving vaccine and we had to choose between saving *A* or *B*, both of whom, if saved, would live another forty years. *A* is a thirty-year-old analyst who, if saved, would continue to lead a life of minimal social contribution, and *B* is a fifty-year-old diplomat who, if saved, would be able to attend a meeting at which he would prevent World War III, thereby saving millions of lives.[77] In this case, saving *A* would generate a direct benefit to *A*:

---

[75] Benefits can also be distinguished as health-related or non-health-related, but, for the sake of simplicity, I will not discuss that distinction here. For a discussion of that distinction, see Brock (2003) and Bognar (2008).

[76] Cf. Brock (2003) and Kamm (1993). For a discussion of different conceptions of direct/indirect benefits, see Du Toit and Millum (2016), who argue that there is no morally relevant distinction between direct and indirect benefits.

[77] This example is brought up in Basson (1979).

i.e., forty years of life; but saving *B* would generate not only a direct benefit to *B:* i.e., forty years of life, but also an *indirect* benefit to society or third parties: i.e., the prevention of WWIII, thereby saving millions of lives. As this example suggests, counting indirect benefits may at least sometimes be permissible: saving *B* would lead to much better consequences than not doing so, since it would prevent WWIII and thereby spare millions of lives.

Bioethicists tend to assume that it would be wrong to decide between candidates for a scarce, lifesaving resource on the basis of the indirect benefits that he or she may generate: e.g., on the basis of her social contribution. One kind of worry is practical: it's not clear which factors indicative of social contribution would constitute suitable allocation criteria. Another kind of worry is that such an allocation scheme would disregard the value of the life of the direct beneficiary by giving undue preference to the desires and interests of third parties. However, as I argue in this paper, neither of these kinds of worries is compelling enough against the idea that indirect benefits such as social contribution should count when resources are scarce and not all can be saved. But there is another objection to the view that indirect benefits such as social contribution should count in scarce resource allocation: the double jeopardy objection.[78] Imagine that we had one dose left of a life-saving

---

[78] The 'double jeopardy objection' was first offered by Harris (1987) against the use of the Quality-Adjusted Life-Year (QALY) measure for priority setting in health care, which I discuss in Section 2. For other fairness-based objections, see Du Toit and Millum (2016), Lippert-Rasmussen and Lauridsen (2010), and Broome (1994).

vaccine, but we had to choose between saving *C* or saving *D*, both of whom, if saved, would live another twenty years. Imagine, further, that *C* is employed and *D* is unemployed, and that, if saved, *C* would immediately be able to go back to work and thereby contribute to economic productivity, whereas *D*, if saved, would not (or not immediately) be able to do so. A policy that counted indirect benefits in this case would have a reason to choose *C* over *D*: choosing *C* would, it would be reasonable to believe, lead to better consequences, e.g., in terms of economic productivity. But, the double jeopardy objection would go, choosing *C* would put *D* in a kind of double jeopardy: not only is she *already* worse off than *C*: i.e., she is unemployed, but now *because* she is unemployed and thereby worse off, this allocation scheme would give her lower consideration for a lifesaving drug.[79] The idea, then, is that an allocation scheme that counted indirect benefits such as social contribution would, in this case, put an individual who is already at a disadvantage into a further disadvantage *on the basis of* her prior disadvantage.

In this paper, I argue that a careful examination of this double jeopardy objection might reveal a stronger foundation for the view that, when resources are

---

[79] It’s worth noting that one might here appeal to a prioritarian ideal according to which priority should be given to those who are worse off. For example, according to Parfit, the moral significance of a benefit depends not only on “how great this benefit would be,” but also on “how well off the person is to whom this benefit comes” (1997, 213). In other words, the weight of a benefit is, on a prioritarian view, not uniform; it varies depending on need.

scarce and not all can be saved, counting indirect benefits such as social contribution, at least in certain social conditions, would be wrong.

The paper develops as follows. Section 1 begins by surveying some of the arguments for and against counting indirect benefits such as social contribution in scarce resource allocation. Section 2 then introduces the double jeopardy objection to counting such benefits, arguing that this objection can be understood in a weak sense, according to which a certain rationing scheme might *result* in injustice, and in a strong sense, according to which a rationing scheme might *compound* injustice. The strong sense of double jeopardy, I argue, seems worse than the weak sense of double jeopardy when the agent that compounds injustice is the same agent that brings about the original injustice. Next, in Section 3, I recruit data from the U.S. Bureau of Labor to argue, first, that counting indirect benefits such as social contribution would, in the contemporary U.S., be likely to put certain social groups, e.g., in terms of race, in strong double jeopardy. I then argue that, in the contemporary U.S., such a policy would be worse than one that put such social groups in double jeopardy in the weak sense because of who the agent who compounds injustice is. Section 4 addresses some objections, and Section 5 considers the implications of my view for allocation schemes that make assumptions about the value of life with disability.

1. <u>Why and why not count indirect benefits in scarce resource allocation</u>

The case in which we must save either the diplomat who would prevent World War III and thereby spare millions of lives, or some other person who would not contribute much to society, is a prima facie case for counting indirect benefits such as social contribution when resources are scarce and not all can be saved.[80] The strength of this case is undeniable; it works even if the diplomat's expected direct benefit from the resource is lower than the other person's. For example, even if the life of the diplomat were expected to be extended by only two years, whereas the life of the other person were expected to be extended by twenty, many of us would still find it permissible to save the diplomat. Some might even say that, given the type of indirect benefit that he would bring about, the diplomat is of "greater value" (Basson 1979). This section begins by discussing the main argument for counting indirect benefits such as social contribution when resources are scarce and not all can be saved. I then examine two types of objection to this argument, arguing that neither is compelling enough to support the idea that counting indirect benefits such as social contribution would be morally objectionable.

---

[80] It's worth noting right away, though, that there is also a prima facie case *against* counting indirect benefits such as social contribution or, e.g., social worth. Suppose we had to choose between two individuals who are equal in all respects except that one has slightly greater social worth. Suppose, further, that this slightly greater social worth came from her slightly better manners. In such a case, counting indirect benefits in terms of social worth, i.e., slightly better manners, is clearly not a morally legitimate consideration for favoring one person over another when they are competing for a scarce, lifesaving resource. Kamm (1993) makes this point. While I will not discuss this further, note that one might respond, in favor of counting such indirect benefits, that we are dealing with something very different when the *lives* of third parties depend on the people competing for resources.

The main argument in favor of counting indirect benefits such as social contribution is consequentialist in nature.[81] The idea is that when resources are scarce and not all can be saved, counting indirect benefits, e.g., in terms of social contribution, would produce better consequences than not counting them.[82] Importantly, according to this argument, when resources are scarce and not all can be saved, there is no principled reason for only counting indirect benefits that are lifesaving: e.g., the indirect benefits of saving the diplomat—that is, sparing millions of lives (Basson 1979). In other words, the idea is that if we find it permissible to count the indirect benefits of an allocation scheme when doing so would spare millions of lives (i.e., bring about better consequences than not counting such benefits), it seems like we should, on the same grounds, find it permissible to count the indirect benefits of an allocation scheme when doing so would bring about other kinds of benefits (i.e., better consequences than not counting such benefits). Such indirect benefits may take various forms: for example, they may manifest in terms of economic growth or societal wellbeing. But the point is that not counting them would amount to a kind of waste of resources—or so many would seem to agree.

---

[81] Note that one might also appeal to a kind of fairness consideration in defense of counting indirect benefits: that the possibility of receiving resources has been made possible through social investment, and so society's interests, or the interests of third-parties, should play a role in how these resources are allocated or used. For this view, see Shatin (1966) and Rescher (1969).

[82] Lippert-Rasmussen and Lauridsen (2010) refer to this kind of view as health non-exceptionalism, according to which health- and non-health-related indirect benefits count.

Nevertheless, this consequentialist argument for counting indirect benefits in scarce resource allocation raises several worries. One of them is practical: it is unclear which factors indicative of indirect societal benefit would constitute suitable allocation criteria. For example, would number of dependents or ability to contribute to the economy be relevant factors? One response to this worry is that deciphering suitable allocation criteria is at least prima facie not an impossible task, as we might form committees of laypeople in order to come up with allocation criteria that are in accordance with societal values (Basson 1979, Rescher 1969, Shatin 1966). Of course, the worry then is that such committees might generate criteria that are biased or unreflective, such as assuming that societal benefit is just a matter of wealth or power.[83] Nevertheless, while it may be difficult or impossible to eliminate such biases altogether, it does not seem unreasonable to think that efforts could be undertaken to reduce such biases.[84] For example, as I argue in Section 3, while it may seem reasonable to think that economic contribution would be an acceptable way to gauge indirect benefit to society, a careful analysis of what this would bring about in the

---

[83] See, for example, Alexander (1977) for a discussion of many questionable factors that may have played a role in the allocation of dialysis machines in an Oregon hospital in the 1960s, including ability to pay and moral character.

[84] Suggestions to this end include conducting opinion surveys (Shatin 1966) and forming committees of experts (Basson 1979). But note that there is reason to think that conducting opinion surveys alone would not circumvent the worry that forming committees of laypeople could lead to criteria that are immoral. A committee of laypeople formed in Nazi Germany might reason that doctors should treat an Aryan over a Jew (Basson 1979).

contemporary U.S., for instance, might reveal some important moral concerns. I will thus set aside this practical worry with counting indirect benefits such as social contribution.

Another kind of objection to the consequentialist argument for counting indirect benefits such as social contribution involves a set of worries regarding the idea that social contribution might be assessed by counting number of dependents (Harris 1985). Imagine we had to choose between giving a lifesaving drug to Jane, who has five children who are dependent on her, or to some other person who has no dependents. A policy that counted indirect benefits such as social contribution in terms of number of dependents would seem to generate a reason to save Jane: saving Jane, unlike saving the other person, would bring about indirect benefits to her five children.[85] Nevertheless, as Harris (1985) writes,

> If the fact that Jane has children who want her to live, and are dependent on her, is a good reason for saving her rather than me, why shouldn't the fact that she has children who want her dead be a good reason for saving me rather than her?

Harris' point is that if we do not find it appropriate to count the desires and interests of third parties as justifications *against* someone (i.e., choosing not to save her), then

---

[85] Harris uses the term 'social usefulness' instead of 'social contribution,' but I take these to have roughly the same meaning.

we also should not find it appropriate to count the desires and interests of third parties in *favor* of someone (i.e., choosing to save her over someone else).[86]

But Harris' worry is misguided. While it does seem inappropriate to count Jane's children's *wants* as justifications against her, the point seems less strong when it comes to counting her children's *interests* against her—that is, their dependence on her. To clarify, suppose Jane abuses her children. In this case, regardless of whether or not the children *want* Jane dead, many of us might find it appropriate to discount her claim on the drug on grounds that saving her would not be in the interest of her children.[87] In other words, even if Harris is right that counting the desires of third

---

[86] Consider another example, adapted from Lippert-Rasmussen and Lauridsen (2010). Imagine a former popstar who is in need of a scarce, lifesaving resource and who, if saved, would go on to make a new album. Now, suppose that one of the things we know about this former popstar is that last month a representative poll revealed that many people in our society have a strong preference *against* her making a new album, and that very few people have a slight preference *for* her making a new album. Should the fact that many people have a strong preference against the former popstar making a new album be counted against her? Harris' point is that if we do not find it appropriate to count the interests and desires of third parties (i.e., people in our society) as justifications *against* someone (i.e., choosing not to save the popstar), then we should not find it appropriate to count other indirect benefits *in favor* of someone (i.e., choosing to save him over someone else).

[87] In the case of the popstar, the idea would be that the reason we find it inappropriate to count the interests and desires of society is not that society doesn't *want* the popstar to make a

parties against someone would be inappropriate, there may be good reasons for counting their *interests*. Harris' third-party desires and interests objection to counting number of dependents is thus not conclusive against the consequentialist argument for counting indirect benefits, either.

But Harris thinks that there is an additional problem with counting indirect benefits such as number of dependents when resources are scarce and not all can be saved. He thinks that doing so might, at the policy level, amount to a "systematic family preference." This idea takes two forms. First, Harris thinks that counting number of dependents would lead to two kinds of citizens: those with families and those without, where those without families would carry the burden of never being chosen when competing for resources against those with families. Note that to agree with Harris one must assume that this would happen if all other relevant factors—e.g., prognosis—were equal; otherwise, it's not clear how this family preference would arise.[88] In any case, Harris' point is that while it is true that dependents would be caused great misery and loss if their loved one is not saved, this does not entail that these dependents have *a right to*, or ought to, be "protected against such misery and

---

new album, but rather that it would not be really bad for society, i.e., in terms of dependence, if the popstar were saved and went on to make a new album.

[88] It's generally agreed upon that prognosis is a relevant factor when allocating a lifesaving drug. For instance, if we had to save X or Y, both of whom are 20 years old, but X had a prognosis of 2 years if saved, while Y had a prognosis of 40 years if saved, many people would agree that Y should be saved.

loss, at the expense of the lives of others who also would be miserable and suffer a great loss if they were to die prematurely" (Harris 1985).

Note, though, that this point depends on the kind of misery and loss in question. One might think that some kinds of misery and loss may be acceptable considerations for choosing to save one person over another: e.g., if such misery and loss included life-threatening conditions for the family members or other third parties. Thus, even if Harris were right that counting number of dependents would lead to individuals without families carrying the burden of never being chosen for a lifesaving treatment, such a preference may, at least in some cases, be justified.

Harris' second formulation of the worry that counting number of dependents would lead to a systematic family preference is that counting number of dependents might serve as a reason for people to see the acquisition of dependents as a kind of insurance for a higher priority rating. Harris' worry is that although this may lead to positive consequences—e.g., more durable marriages, more intense custody battles, or higher rates of adoption—these consequences are unlikely to be preferable, morally, to giving individuals lower priority based on a low number of dependents.

But this point seems unwarranted, too. Accumulating a steady number of dependents would quickly become a quite expensive burden, thereby disqualifying this as an option for many people. Moreover, it seems unlikely that individuals would consider seriously such expensive, long-term commitments *just in case* they happen to need some scarce lifesaving resource in the future. Harris' worry that counting number of dependents may lead people to see the acquisition of dependents as a kind of insurance for a higher priority rating is thus also short of compelling.

To sum up: the underlying argument for counting indirect benefits when resources are scarce and not all can be saved is consequentialist in nature: the claim is that counting indirect benefits such as social contribution would bring about better consequences than not counting them. Two types of objections have been raised against this argument: a practical one about how to gauge indirect benefits, and one about counting social contribution in the form of number of dependents; but neither is conclusive against the claim that indirect benefits such as social contribution should count. Might counting indirect benefits such as social contribution be morally objectionable for other reasons? In the next section, I recruit Harris' (1987) double jeopardy objection to the use of the Quality-Adjusted Life-Year measure for priority setting in health care in order to offer a stronger case for the idea that it would be wrong for a scarce resource allocation scheme to count indirect benefits such as social contribution.

2. *Two concepts of 'double jeopardy'*

My goal in this section is to lay the groundwork for the idea that certain kinds of allocation schemes, e.g., those that count indirect benefits such as social contribution, are morally objectionable because they risk putting certain people in double jeopardy. To this end, I make a distinction between a 'weak' sense of double jeopardy, according to which an allocation scheme brings about injustice, and a 'strong' sense of double jeopardy, according to which an allocation scheme brings about *additional* injustice—or, equally, *compounds* injustice. As I argue, strong double jeopardy seems worse if the allocating agent or entity is the same agent or entity that brought about an original injustice against the same group.

The double jeopardy objection originates in a paper by Harris (1987) against the use of the Quality-Adjusted Life-Year (QALY) measure for priority setting in healthcare. The QALY measure seeks to offer allocation recommendations that are based on the number of ‘QALYs’ that a certain allocation scheme would generate. QALYs are calculated by taking into account the life expectancy as well as the expected quality of life of an individual or group of individuals who would benefit from a health treatment.[89] On this system, a high priority healthcare activity is one that generates a low cost per QALY; a low priority healthcare activity is one that generates a high cost per QALY. Harris’ worry is that because the QALY system measures expected quality of life, and because people who have disabilities are more likely to have a lower expected quality of life than people who lack disabilities, using

---

[89] Harris (1987) quotes Alan Williams, “the architect of QALYs”:

> The essence of a QALY is that it takes a year of healthy life expectancy to be worth one, but regards a year of unhealthy life expectancy as worth less than 1. Its precise value is lower the worse the quality of life of the unhealthy person (which is what the "quality adjusted" bit is all about). If being dead is worth zero, it is, in principle, possible for a QALY to be negative, i.e., for the quality of someone's life to be judged worse than being dead.
>
> The general idea is that a beneficial health care activity is one that generates a positive amount of QALYs, and that an efficient health care activity is one where the cost per QALY is as low as it can be. A high priority health care activity is one where the cost-per-QALY is low, and a low priority activity is one where cost-per-QALY is high. (119)

the QALY system would effectively lead to a systematic preference against people who have disabilities—both in terms of making allocation decisions between candidates, and in terms of making decisions about which treatments to fund. For example, imagine we had to save either a person who just suffered an accident and who, if saved, would live with paraplegia, or some other person who, if saved, would make a perfect recovery. Furthermore, imagine that each candidate would live another twenty years if saved. As Harris puts it,

> QALYs dictate that because an individual is unfortunate, because she has once become a victim of disaster, we are required to visit upon her a second and perhaps graver misfortune. The first disaster leaves her with a poor quality of life and QALYs then require that in virtue of this she be ruled out as a candidate for lifesaving treatment, or at best, that she be given little or no chance of benefiting from what little amelioration her condition admits of. Her first disaster leaves her with a poor quality of life and when she presents herself for help, along come QALYs and finish her off! (Harris 1987, 120)

Harris' point is that the QALY measure puts the candidate who became paraplegic as a result of misfortune in double jeopardy. Not only does she suffer a misfortune: i.e., is left paraplegic and thereby with a lower expected quality of life, but *because* of this misfortune, the QALY measure presents her with an injustice: it gives her unfair or unjust lower consideration for a scarce lifesaving treatment.[90] Put

---

[90] Harris (1995) offers two arguments for thinking that the QALY system is unjust. One is based on his idea that the value of the life of one person (who wants to go on living) is equal to the value of the life of any other person (who also wants to go on living), and that such

differently, not only is the patient who was a victim of an accident *already* worse off than the patient who would make a full recovery, but *because* she is worse off, the QALY measure makes her even worse off: she now gets a lower priority score for a lifesaving treatment. In still other words, the QALY measure puts the person who is already at a disadvantage into a further disadvantage *on the basis o*f her prior disadvantage (Singer, et al. 1995).

Consider the double jeopardy objection against an allocation scheme that counted indirect benefits such as social contribution. Recall the case of *C* or *D* from above. *C* is employed, *D* is unemployed, and, if saved, each would live an additional twenty years. A rationing scheme that counted indirect benefits, e.g., in terms of economic productivity, would in the case of *C* or *D* generate a reason for saving *C*: it would be reasonable to believe that saving *C*, because he is employed, would bring about better consequences for society, i.e., in terms of economic contribution, than saving *D*. However, the double jeopardy objection would go, saving *C* on these grounds would amount to putting *D* in a kind of double jeopardy: not only would *D*

---

value does not vary depending on, e.g., how many more years an individual will live or what her expected quality of life will be. He thus seems to think that a system (or any system) that gives higher priority to one individual over another is unjust, for it gives inappropriate weight to the life of one person over the life of another. The other argument that Harris offers for thinking that the QALY system is unjust is that if a community "values individuals differentially according to the success of their lives and its quality and predicted length of un-elapsed time, this is highly likely to have a disastrous effect on [those individuals'] sense of personal worth and their sense of security" (1995, 156).

be *already* worse off than *C*, i.e., she is unemployed; but *because* she worse off, she would now receive a lower priority score for a lifesaving treatment. In other words, choosing *C* on grounds of the indirect benefits that saving her would generate would put *D*, who is already at a disadvantage, into a further disadvantage *on the basis of* her prior disadvantage.

A supporter of counting indirect benefits such as social contribution might respond that in order for this double jeopardy objection to work, we would need a reason to think that *D*'s being unemployed actually makes her already worse off than *C* (Du Toit and Millum 2016, Singer, et al. 1995). If *D* is poor and unemployed, for instance, then giving higher priority to *C* may well make *D* even worse off. But suppose *D* happened to be the recipient of a very large inheritance. In this case, even though *D* is unemployed, she is not already worse off than *C*. Giving higher priority to *C* would thus not make *D* even worse off—in other words, it would not put *D* in double jeopardy. Thus, this reply to the double jeopardy objection would go, the double jeopardy objection is unsuccessful because it does not always work; counting indirect benefits such as economic contribution may or may not further disadvantage an individual who is already at a disadvantage. In what follows, I will refer to this reply to the double jeopardy kind of objection as the *does-not-always-apply reply*.

I do not think that we should accept the does-not-always-apply reply for two reasons. First, the relationship between counting indirect benefits and making someone worse off than she already is, on the basis that she is already worse off, does not have to be necessary to be worth examining. But, second, and more importantly, I think that there are two concepts of 'double jeopardy' at play in the notion of double jeopardy. Understanding each of these concepts might highlight that some cases of double jeopardy may be worse than others.

Let's begin by noting that Harris' notion of double jeopardy involves a combination of factors that results in an injustice.[91] On the one hand, there is the misfortune that the candidate for the scarce lifesaving treatment suffers: she is in an accident that leaves her paraplegic and thereby with a lower expected quality of life. On the other hand, there is the kind of allocation scheme in place: in this case, the QALY measure, which generates an allocation recommendation on the basis of the candidate's expected quality of life. When combined, these factors result in an injustice: the candidate who is paraplegic is now denied equal consideration for a lifesaving treatment *on the basis of* her expected quality of life—an expected quality of life that was the result of misfortune. In other words, not only is this candidate in an accident that leaves her with a lower expected quality of life, but now *because* she is in an accident that leaves her with a lower expected quality of life, she receives lower consideration for the scarce lifesaving treatment. According to this understanding of double jeopardy, then, one injustice occurs: the accident victim is denied equal consideration for a lifesaving treatment.

---

[91] It's worth keeping in mind that Harris thinks that such a policy would be unjust because it would fail to give equal weight to the interests if each person, failing to account for "the root of democratic theory and of most conceptions of justice: that each person is as morally important as any other" (1987, 118). For the purposes of this paper, I will not question this assumption.

But there is a second, stronger sense of double jeopardy—or so I believe. Suppose we had to save either a person whose back was broken by an attacker and who, if saved, would live with paraplegia, or some other person who, if saved, would make a perfect recovery. Suppose, further, that each candidate would live another forty years if saved. In this case, the QALY system would seem to generate a reason for giving the attack victim lower priority, as she would generate fewer QALYs (since she would live the same number of years as the other candidate, but with a lower quality of life). This case, like the accident case above, also involves a combination of factors that results in an injustice. One factor is that as a result of being attacked, this candidate is left with paraplegia and thereby with a lower expected quality of life. The other factor is that there is a certain kind of allocation scheme in place: the QALY measure, which generates an allocation recommendation on the basis of the candidate's lower expected quality of life. When combined, these factors result in an injustice just like in the case above: the attack victim is denied equal consideration for a lifesaving treatment *on the grounds that* she was attacked and left with a lower expected quality of life.

But note that while the attack case and the accident case are similar in that both bring about an injustice (i.e., both candidates are denied equal consideration for a lifesaving drug), there is an important difference between these candidates: whereas it is *unfortunate* that the accident victim is left paraplegic and thereby with a lower expected quality of life, it is not merely unfortunate that the attack victim is left paraplegic and thereby with a lower expected quality of life. Accidents are unfortunate; but attacks are not just unfortunate, they are unjust. Thus, while the accident victim receives lower consideration for the lifesaving drug as a result of a prior misfortune, the attack victim receives lower consideration for the lifesaving drug as a result of a prior *injustice*. Thus, insofar as the QALY system makes priority recommendations based on an individual's expected quality of life, it puts the accident victim in double jeopardy in a *weak* sense, and it puts the attack victim in double jeopardy in a *strong* sense. In other words, the QALY system brings about an injustice for the accident victim but brings about *additional* injustice—or, equally, *compounds* injustice—for the attack victim.

Now, does this mean that disfavoring the attack victim (i.e., putting her in strong double jeopardy) is worse than disfavoring the accident victim (i.e., putting her in weak double jeopardy) if resources are scarce and not all can be saved? One might not think so. After all, in these examples the allocating agent or entity does not seem to have anything to do with the misfortune that the accident victim previously suffered or with the injustice that the attack victim previously suffered. In fact, it is precisely because no wronging occurs in the accident case that that is a case of weak double jeopardy. Thus, one might think that even if one grants that the QALY system brings about an injustice for both the accident victim and the attack victim (insofar as both get a lower priority score based on their expected quality of life), one can still hold that it is not the place of an allocating entity to try to remedy past misfortune or injustice—especially if the allocating entity is not the same entity that brought about such misfortune or injustice.

So far so good. But imagine a case in which the allocating entity *is* the same entity that brought about an original injustice on which the allocation injustice is based. To be more precise, imagine, for example, that in the case of the attack victim, the entity that makes an allocation decision based on the attack victim’s expected quality of life is the same entity that attacked the victim in the first place. In other words, in the case we are imagining now, the allocating entity makes an allocation decision based on the attack victim’s expected quality of life, but it is the allocating entity itself that brought about the attack victim’s lower expected quality of life (i.e., by attacking the victim).

Now, would it be worse for the allocating entity to disfavor the attack victim in *this* scenario than it would be for it to disfavor the accident victim, whose lower expected quality of life was the result of mere misfortune? I believe so. While the accident victim is wronged once by the allocating entity—insofar as she does not receive equal consideration for the lifesaving drug—the attack victim is wronged *twice* by the allocating entity: first, she is attacked by this entity and left with paraplegia and thereby with a poor expected quality of life, and then *because* of this attack, she does not receive equal consideration for the lifesaving drug. Put differently, the accident victim is put in weak double jeopardy: she suffers an injustice based on a prior misfortune; but the attack victim is put in strong double jeopardy: she suffers an injustice on the basis of a prior injustice. In other words, the allocating entity brings about an injustice for the accident victim, but it *compounds* the injustices that the attack victim suffers; it wrongs the attack victim a second time *on the basis* of having already wronged her. Consequently, strong double jeopardy would seem to be worse than weak double jeopardy.

What does this suggest for the topic of this paper—namely, the idea that indirect benefits such as social contribution should count in scarce resource allocation? I believe that it suggests that an allocation scheme that counted indirect benefits such as economic contribution and thereby put certain individuals in double jeopardy in the weak sense would be morally objectionable, but that it would be worse if such a policy put certain individuals in double jeopardy in the strong sense *when the agent who brings about the second injustice is the same agent that brought about the first injustice*, which is the basis for the second injustice. In the next section, I offer a concrete, real-world case in support of this claim.

To sum up this section, then, one way in which one might object to the consequentialist argument that indirect benefits such as social contribution should count when allocating scarce, lifesaving resources is by appealing to Harris' double jeopardy objection to the use of the QALY measure. But there are two concepts of 'double jeopardy' that are relevant when determining how bad a certain allocation scheme would be. There is a weak sense of double jeopardy, according to which an allocation policy may bring about injustice, and a strong sense of double jeopardy, according to which an allocation policy may compound injustice. The latter would seem to be worse if the entity that brings about the second injustice is the same entity that brings about the first injustice.

3. The contemporary U.S.: When counting indirect benefits compounds injustice

There are various ways in which indirect benefits such as social contribution might figure in an allocation scheme: they may be taken as the sole factor, as one of several factors, or as a kind of tiebreaker when deciding between candidates. My aim in this section is threefold. First, I offer a concrete case in support of the idea that an allocation scheme that counted indirect benefits such as social contribution—whether as the sole factor, as one of several factors, or as a kind of tiebreaker—would, in the contemporary U.S., be morally objectionable because it would put certain social groups, e.g., in terms of race, in weak double jeopardy.[92] Next, I argue that there is reason to believe that such an allocation scheme would, in the contemporary U.S., put blacks in *strong* double jeopardy. Ultimately, the claim I seek to defend is that, in the contemporary U.S., an allocation scheme that put black Americans in strong double jeopardy would be worse than one that put them in weak double jeopardy because the allocating entity would seem to be the same entity that brought about an original injustice on which it bases its allocation decision.

To these ends, I will assume that whether we are concerned with societal wellbeing or with economic growth, social contribution can be gauged in terms of

---

[92] Notice that there is a question of whether or not compounding injustice might be justified if overall health benefit is maximized. For instance, one might think that compounding injustice would be justified if the overall health benefit in terms of number of lives saved increases. Properly addressing this suggestion is a project for another time, but for a discussion of this idea with regard to the use of kidney vouchers, see Kerstein (2017). Note, though, that while Kerstein considers the permissibility of potentially *exacerbating inequalities* by using kidney vouchers, the relevant point for this paper is whether a policy that counted indirect benefits such as social contribution and thereby *compounded injustice* might be justified.

economic contribution.[93] This assumption is not uncontroversial; it circumvents the difficult question of what constitutes an appropriate measure of a society's wellbeing.[94] Nevertheless, I think it is at least plausible that a society's wellbeing goes hand in hand with its higher standard of living, and that a higher standard of living goes hand in hand with economic growth, such that counting economic contribution—whether in terms of past contribution, future contribution, or both—would be a plausible way to gauge social contribution, or indirect societal benefit.[95]

---

[93] For one proposal in which social usefulness might make a difference, see Persad, Wertheimer and Emanuel (2009).

[94] For instance, is an appropriate measure of a society's wellbeing its economic power? The Human Development Index? The Quality of Life Index? Moreover, it is not at all obvious that the higher a person's salary or purchasing power, the more socially useful she is. One might think, for instance, that someone who is more likely to increase her society's wellbeing is more socially useful. For example, a preschool teacher might have a lot more social worth than an investment banker; but if all we consider in order to determine social usefulness are economic factors, we might deem the banker to be more socially useful than the teacher.

[95] Further evidence of how counting economic benefit might make a difference in the real world, see Du Toit and Millum (2016), who write: "For example, in a joint news release from the World Health Organization (WHO), 'Stop TB Partnership, and the World Bank,' Dr. Margaret Chan, Director-General of the WHO, is quoted saying: '[t]here were already compelling reasons to fight TB, which causes massive human suffering. Now, as a further incentive, there are strong indications that investment in meeting the Millennium Development Goal related to TB carries important economic benefits.'"

Several factors may be indicative of economic contribution. The U.S. Bureau of Labor's report on the characteristics of the labor force offers some relevant information with respect to such factors (U.S. Bureau of Labor Statistics 2017). For example, according to the 2016 report, over 90 percent of whites, blacks, and Asians age 25 and older who were part of the labor force were reported to have at least a high school diploma, compared to only 72 percent of Hispanics.[96] These statistics are useful for our purposes because, according to the same report, individuals who have higher levels of education are generally more likely to be employed in higher-paying jobs than are individuals who have lower levels of education.[97] Moreover, level of education for all of these groups is also directly related to employability, such that the higher the level of education, the greater the likelihood of employment and the lower the likelihood of unemployment.[98] The report thus offers enough information for at least four possible ways of measuring economic contribution: (1) educational attainment (insofar as it correlates with higher salaries), (2) occupation (insofar as the report sorts occupations into higher- and lower-paying categories), (3) wages (as indicative of purchasing power), and (4) employment status (also as indicative of

---

[96] Moreover, 60 percent of Asians but only 38 percent of whites, 27 percent of blacks, and 19 percent of Hispanics in the labor force had a bachelor's degree or higher.

[97] Even so, blacks and Hispanics generally had lower earning when compared with whites and Asians at nearly all educational levels.

[98] Although according to this report the relation between employability and level of education was true of every group, blacks and Hispanics were still more likely to be unemployed than were whites and Asians.

purchasing power). As we will soon see, however, whites and Asians fare better across all of these characteristics than blacks and Hispanics. This suggests that a scarce, lifesaving resource allocation scheme that counted economic contribution and took these factors as indicative of it would likely have the effect of giving higher priority to whites and Asians than to blacks and Hispanics. Moreover, note that these characteristics, insofar as they represent patterns across educational attainment, occupation, wage, and employment status, could be taken to be indicative of both past and future or expected economic contribution.

Consider educational attainment first. Because Hispanics are less likely than whites, blacks, and Asians to have a high school diploma, an allocation scheme that measured (past or future) economic contribution in terms of, e.g., 'educational attainment above a high school diploma,' would give lower priority to Hispanics.

Consider 'occupation' next. According to the report, the majority of Asians were employed in the highest paying category, and only one-third of blacks and one-fifth of Hispanics were employed in the same category.[99] Moreover, 2 out of 5 blacks and Hispanics were employed in the lowest-paying category, but only, approximately,

---

[99] To be more precise, while 51 percent of Asians were employed in the highest-paying major occupational category, 'management, professional, and related occupations,' 39 percent of whites, 30 percent of blacks, and only 21 percent of Hispanics were employed in the same category. A total median of $1,137, although lowest for blacks and Hispanics, at $903 and $937, respectively.

1 out of 4 whites and Asians were employed in the same category.[100] Thus, if we were to allocate scarce, lifesaving resources by measuring economic contribution in terms of ‘occupation,’ giving priority to the highest-paying category, the report suggests, we would likely give lower priority to blacks and Hispanics, and higher priority to whites and Asians.

A similar trend follows if we take ‘wages’ as indicative of economic contribution. The same report indicates that blacks and Hispanics generally had lower earnings when compared with whites and Asians at nearly all education levels, and at nearly all the different occupation categories. Only three—of twenty-two—occupations listed blacks or Hispanics as having higher wages than either whites or Asians.[101] This suggests that on an allocation scheme that took ‘wages’ as an

---

[100] To be more precise, 41 percent of blacks and 42 percent of Hispanics were employed in the categories of ‘service’ or ‘production, transportation, and material moving’—the two categories reported to have the lowest weekly earnings of full-time wages and salaries—compared to 27 percent of whites and 27 percent of Asians. Service occupations had the lowest weekly earnings of full-time wages and salaries (total median of $505, although lower for blacks and Hispanics, at $483 and $454 respectively); and ‘production, transportation, and material-moving’ occupations had the second-lowest weekly earnings of full-time wages and salaries (total median of $642, although lower for blacks and Hispanics, at $597 and $540 respectively).

[101] In ‘protective service occupations,’ average wages for Hispanics were higher than they were for Asians, but lower than they were for whites; in ‘farming, fishing, and forestry occupations,’ average wages for blacks were higher than they were for whites and Hispanics,

indicator of economic contribution, blacks and Hispanics would again be likely to receive lower priority than whites and Asians.

Finally, consider ‘employment status.’ Although, according to the same report, the total unemployment rate was 6.2, the unemployment rates for whites and Asians were lower than this average (5.3 and 5.0, respectively), but higher for blacks (11.3—that is, almost twice the total!) and Hispanics (7.4). In fact, the unemployment rate for whites has been lower than that of Hispanics, which has been lower than that of blacks, since 1973. In short, according to the report, blacks and Hispanics were more likely to be unemployed than were whites and Asians. This suggests that an allocation scheme that measured economic contribution in terms of ‘employment status’ would again be likely to give lower priority to blacks and Hispanics than to whites and Asians.

In short, any one of these ways of measuring economic contribution—that is, educational attainment, occupation, wages, and employment status—would likely have the effect of giving lower consideration to blacks and Hispanics, and higher consideration to whites and Asians. In other words, since blacks and Hispanics are already at a disadvantage, a policy that counted economic productivity would seem to put them in a kind of double jeopardy, at least in the weak sense. To clarify, since blacks and Hispanics are *already* at a disadvantage across all of the different work force characteristics that might be relevant for gauging economic contribution, and

---

but lower than they were for Asians; and in ‘production occupations,’ wages for blacks were slightly higher than they were for Asians but much lower than they were for whites.

since such a policy would generate an allocation recommendation based on economic contribution, what would result is that blacks and Hispanics would likely be presented with an injustice: they would likely be denied equal consideration for a lifesaving drug.

Now, I believe that blacks and Hispanics are not individually responsible for having lower scores across all of these workforce characteristics and assume so in what follows. Thus, the question I wish to consider is whether or not their lower scores across all of these work force characteristics are the result of injustice. If they are, then a policy that counted economic contribution in these ways would put blacks and Hispanics in double jeopardy in the strong sense, since it would deny them equal consideration for a scarce, lifesaving resource on the basis of a prior *injustice*. Moreover, as I will argue, putting blacks in strong double jeopardy would seem to be worse than putting them in weak double jeopardy because the allocating entity would seem to be the same entity that brought about the injustices on which the allocation injustice is based.

To see why one might think that an allocation policy that counted economic contribution would bring about an injustice to blacks that is based on a prior injustice—not on a prior misfortune—it's important to consider the history of the concept of 'race' in the U.S. In what follows, I will focus only on the two main racial groups in the U.S.: whites (the racial majority in the U.S.) and blacks (the largest racial minority in the U.S.).[102] Although a similar case can be made about other nonwhite groups in the U.S.—e.g., about Hispanics and Native Americans—an overview of that history is beyond the scope of this paper.

The concept of 'race' as it is understood today arguably emerged with the Atlantic slave trade and the financial need to produce cotton, tobacco, rice, and sugar (Fields 1990). 'Race' was at that time understood as a biological characteristic that was essential, inherited, and unchanging, thus serving to classify human beings into subgroups—that is, racial groups (Blum 2002, Shelby 2002). As a biological characteristic, race was also taken to determine the physical, mental, and moral capabilities of the individual members of each racial group, such that the groups could be ranked in a hierarchy, with those with light skin at the top, and those with dark skin at the bottom (Blum 2002, James 2017, Shelby 2002). Understood this way, 'race' arguably served as a justification for the subordination and exploitation of nonwhites, which, in turn, made possible the production of cotton, tobacco, rice, and sugar (Blum 2002, James 2017, Shelby 2002).

---

[102] According to the U.S. Census Bureau, whites (not including Hispanics) make up 60.7% of the U.S. population and blacks make up 13.4%. Racial minorities in the U.S. include the following categories: Asian (5.8%), American Indian and Alaskan Native (1.3%), Native Hawaiian and Other Pacific Islander (0.2%). According to this report, Hispanics, who may be of any race, make up 18.1% of the U.S. population. https://www.census.gov/quickfacts/fact/table/US/PST045218

But this ideological understanding of race was arguably sustained over the centuries in part as a result of slavery, and later Jim Crow—laws that existed until just fifty years ago (Shelby 2002).[103] For example, in the U.S., the military wasn't integrated until 1948; segregation in public schools wasn't deemed unconstitutional until *Brown v. Board of Education* in 1954; voter suppression wasn't deemed illegal until the Voting Rights Act of 1965; and discrimination in housing wasn't deemed illegal until the Fair Housing Act of 1968. Even today, when scientists and scholars tend to agree that race does not have a biological basis, there is reason to believe that the ideological understanding of race continues to influence the politics, race relations, and economic conditions of the U.S.[104] The disparities reported by the U.S. Bureau of Labor, discussed above, are arguably just some of the ways in which this ideology continues to influence the social conditions of the U.S. Evidence also

---

[103] I here adopt Shelby's use of 'ideology,' which, as he explains, is a Marxist understanding of 'ideology.' As I understand it, 'ideology' in the Marxist sense is a system of beliefs that shape how we perceive or make sense of the world, and which is produced, stabilized, maintained, and reproduced by the historical, economic and material conditions of society (which seem to include a ruling class that benefits from the ideology). If this is right, then to understand 'race' as an ideology is to understand it as a set of beliefs that emerges from and is maintained by the historical, economic and material conditions of society (e.g., the production of cotton, tobacco, etc.), and which serves to make sense of the world. See (Shelby 2002) and Fields (1990).

[104] See, for example, Kaplan and Winther (2014), Haslanger (2008), Glasgow (2006), Appiah (1990), Taylor (2011). Cf. Sesardic (2010).

suggests that black Americans have higher rates of poverty and incarceration, as well as lower rates of access to healthcare, than white Americans.[105]

Thus, while African Americans historically have been and currently are generally worse off than whites in the contemporary U.S., these inequalities are not inevitable. In other words, they are not the result of some innate or biological factor, as the original concept of race would have us believe; they are historical and socially

---

[105] In terms of health, consider that according to the Center for Disease Control and Prevention HIV/AIDS has been a leading cause of death among young, black, adult women for more than a decade. See Shrage (2017). Shrage theorizes that there is a relationship between the very high HIV/AIDS rates among African Americans and their very high rates of incarceration. Specifically, she argues that that there is a kind of vicious cycle that occurs between the low amount of health resources and the high rates of legal punishment in African American communities, and the poor living conditions in U.S. prisons, which include very high rates of HIV infection, sexual and physical violence, and a mentally unhealthy environment. As Shrage puts it, "When a particular community or ethnic group is subject to both legal and extrajudicial punishment at a higher rate, because of a criminal justice system that racially profiles this group, imposes 'zero tolerance' sentencing … for nonviolent drug offenses not imposed on other groups, which then makes a considerable percentage of this group ineligible to vote and serve on juries… this amounts to political and social persecution" (at p. 455). For the purposes of this paper, we need not agree with Shrage that this amounts to "political and social persecution"; all I am trying to suggest here is that these disparities exist today in part as a result of the role that racist ideology has played in the social conditions of the U.S. since it was founded.

controllable. Moreover, given the history of the U.S. and the role that racist ideology has played in it, it is reasonable to believe that the disparities that we find between blacks and whites today are the result of *unjust* social, legal, and economic patterns that were ratified by governmental institutions.[106]

If this is right, then counting economic contribution in scarce resource allocation would not only put blacks in double jeopardy in the weak sense—that is, by giving them lower consideration for a scarce, lifesaving resource; it would also likely put them in double jeopardy in the *strong* sense, as it would likely *compound*

---

[106] These continued disparities have been theorized to be at least partly the result of bias and stereotyping. For interesting discussions of how bias and stereotyping may play a role in employment status and the like, see Young (2001) and Anderson (2010). It's worth noting that research suggests that social and economic conditions can have an effect even on people's *risk* of illness, for instance, by causing long-term stress (Wilkinson and Marmot 2003). According to this research, psychosocial factors affect physical health by "raising the heart rate, mobilizing stored energy, diverting blood to muscles and increasing alertness," which, when maintained for extended periods of times, make people "more vulnerable to a wide range of conditions including infections, diabetes, high blood pressure, heart attack, stroke, depression and aggression" (Wilkinson and Marmot 2003). See also Marmot (2015) and The World Health Organization (2017). According to the latter, higher income has been linked to good health, and lower education and lack of employment have been linked to "poor health, more stress and lower self-confidence." It would not be unreasonable to suppose that blacks and Hispanics or other individuals who belong to disadvantaged groups may generally be at higher risk of illness than other groups.

the injustices that they already face.[107] In other words, blacks *already* have lower scores than whites across all of the characteristics that one might reasonably appeal to in order to gauge economic contribution. And these lower scores are arguably the result of injustice insofar as they are the result of an unjust social and legal order. On an allocation scheme that counted such factors, then, blacks would be wronged a second time on the basis of a prior injustice they would have suffered: they would be denied equal consideration for a scarce, lifesaving resource on the basis of unjustly faring poorly across the various ways one might reasonably take to be indicative of economic contribution.

I take it that many of us would agree that a concern with justice suggests that, all things considered, past wrongs should be ameliorated, or at the very least not compounded, so a scarce resource allocation scheme that put certain groups of people in double jeopardy in this strong sense would be gravely unjust. But the key question is this: Would it be worse if a policy that counted economic contribution put blacks in double jeopardy in the strong sense than it would be if it put blacks in double jeopardy only in the weak sense?

---

[107] Now, an opponent might say that while it may be unfortunate that blacks fare worse across all of the characteristics one might reasonable take to be indicative of economic contribution, it is not *unjust*. But I think that many of us would agree that, given the U.S.'s history of (social legal, and economic) mistreatment of blacks, that blacks continue to perform poorly across all of these work force characteristics is more than just unfortunate; it is unjust.

I think it would be if the agent who brings about the second injustice is the same agent who brings about the first injustice, on which this second injustice is based. And, in the case of blacks in the contemporary U.S., the allocating entity—that is, a governmental institution—would seem to be the same entity that ratified the prior injustices that blacks already face (i.e., an unjust social order that gave rise to blacks' lower scores across all the various work force characteristics).

To clarify, imagine that the reason that blacks performed poorly today across all of the various measures one might take to be indicative of economic contribution was not an unjust social order, but rather sheer misfortune. An allocation scheme that counted economic contribution in such a case would, it would be reasonable to believe, give higher priority to whites than to blacks, thereby putting blacks in a kind of double jeopardy. But such a policy would not, in this imaginary case, put blacks in *strong* double jeopardy. The disadvantage that blacks face is not, in this imaginary case, the result of an injustice; rather, it is the result of misfortune. In the real case, however, the disadvantage that blacks face *is* the result of an injustice—an injustice that is brought about by the same entity that brings about the second injustice. Thus, in the imaginary case, the reason that blacks are at a prior disadvantage across all of the relevant economic contribution factors is not an injustice; it is just misfortune. But, in the real case, the reason that blacks are at a prior disadvantage is an unjust social order that the allocating agency itself brought about. A policy that counted economic contribution and put blacks in strong double jeopardy would thus seem to be worse; it would bring about not one, but *two* injustices to blacks: the second based on the first.

Let me sum up this section. Examining the social conditions of the contemporary U.S. suggests that an allocation scheme that counted economic contribution in terms of educational attainment, wages, occupation, and employment status, would, in the contemporary U.S., likely give lower consideration to blacks and Hispanics than to whites and Asians, insofar as blacks and Hispanics fare poorly across all of these characteristics. Such an allocation scheme would thus seem to put blacks and Hispanics in double jeopardy in the weak sense. But, given that the U.S. has a long history of unjust mistreatment of blacks, there is reason to think that such an allocation scheme would also put blacks in double jeopardy in the strong sense—that is, it would compound the injustices that blacks face. Since the allocating entity would, in this case, seem to be the same entity that brought about the original injustice on which the second injustice is based, a policy that put blacks in double jeopardy in a strong sense seems worse than a policy that put blacks in double jeopardy in a weak sense.

4. Objections

A number of objections might be raised to the idea that a policy that counted indirect benefits in the form of economic contribution would, in the contemporary U.S., put blacks in double jeopardy in the strong sense. An obvious objection is that such a policy would not be doubly unjust to blacks because blacks are not worse off than whites across economic contribution factors as a result of injustice. For instance, such critics might assert that blacks are morally responsible for not informing themselves about better ways to find employment or about the benefits of education. However, in part because of the impact of the concept of race discussed in the previous section, I believe that these criticisms are unsuccessful, and assume so in this article.

Nevertheless, the opponent might ask: even if blacks' lower scores across potential economic contribution measures *were* the result of unjust social, legal, and economic patterns, if our resources are limited and we cannot save everyone, is it really *unfair* or *unjust* to give lower priority to saving the lives of those who would be less likely to contribute to economic growth?[108] After all, no matter what, someone will die, and saving those who would be less likely to contribute to economic growth would lead to worse consequences for society than saving those who would be more likely to contribute to economic growth. Put differently, the objection would seem to be that while blacks may have been subject to a prior injustice, it is not the place of an allocation scheme to rectify that injustice. It *is* the place of an allocation scheme to be forward-looking and bring about the best outcomes. Blacks may have suffered quite a bit as a result of having lower scores on all measures of economic contribution under discussion, but "[n]othing can be done about past suffering, whereas (often) something can be done about present and future suffering" (Singer, et al. 1995). Choosing to save those who would contribute to society would maximize the amount of future suffering we can prevent. Indirect benefits—e.g., in the form of economic contribution—should count.

---

[108] See Singer et al. (1995). Although Singer et al. do not ask this specific question, they suggest it, in the context on the QALY measure, when they ask if it is really *unfair* to prioritize people who will have a better quality of life when we can only save some.

But, while maximizing the amount of future suffering we can prevent is surely important, the point here is not that allocation schemes must be retrospective and compensate people for past suffering. The point is simply that a concern with justice suggests that disadvantage should be ameliorated, or at the very least not worsened—especially if such disadvantage is unjust, and especially if such disadvantage was brought about by the same entity that seeks to allocate resources in a fair way.

Another rejoinder to the idea that an allocation policy that counted economic contribution would, in the contemporary U.S., put blacks in double jeopardy in the strong sense is a form of the does-not-always-apply-reply, discussed in Section 2. The idea would be that while it may be true that counting economic contribution in the contemporary U.S. may *sometimes* lead to compounding injustice, this wouldn't *always* be the case. Suppose we had to choose between two individuals to give our last lifesaving drug: one who had a very high income, the other who had a very low income. If the former were white and the latter were black, taking 'wages' into account as an indicator of economic contribution may well compound injustice. But if it turns out that the person with the higher income is black and the person with the lower income is white, then giving priority to the one with the higher income would *not* compound injustice. Thus, counting indirect benefits—here in the form of 'wages' as an indicator of economic contribution—would not, even in the contemporary U.S., always lead to compounding injustice.

In response to this objection, let me highlight again that this point amounts to the idea that there is no reason not to count economic contribution because doing so does not *necessarily* compound injustice. But, again, that the relation between counting economic contribution and compounding injustice is not necessary does not mean that it's not worth considering. Even if prioritizing by economic contribution would benefit blacks in *some* cases, the data show that, typically, it would not do so. In other words, even if an allocation scheme that counted economic contribution would not compound injustice in *all* cases, there is reason to believe that in the contemporary U.S. it would be likely to do so.

An opponent may further note that just because there are practical worries with a method for counting economic contribution as indicative of indirect benefit does not mean that counting economic contribution is wrong *in itself*.[109] That is, in an ideal society in which there were no correlations between race and various factors indicative of economic contribution, we would not, in virtue of counting economic contribution, face the problem of compounding injustice. Indirect benefits such as economic contribution should count.

But, while it may be true that in an ideal society we would not face the problem of putting some people in double jeopardy in the strong sense, it may also be true that in an ideal society I could leave my laptop unattended at the airport while I go grab some coffee, and that my laptop would be there every time when I returned. Alas, we don't live in such a society, and given the way people and societies work I have no reason to think that my laptop will always be there when I return; so, I err on the side of caution and either take my laptop with me or just skip the coffee. Similarly, we don't live in an ideal society. In our society, as a result of prior social

---

[109] See, for example, Du Toit and Millum (2016).

and legal patterns, benefits and burdens are not distributed equally, and greater economic contribution accrues with benefits, putting those who are generally less able to contribute economically as a result of injustice at an even further disadvantage. There is thus no reason to think that we should count economic contribution in the ways I've described it here just because doing so is not *in principle* wrong.

Another objection to the idea that a policy that counted economic contribution would put blacks in strong double jeopardy may be that having a rule according to which we ought to take economic contribution into account only among people who have not already suffered a relevant injustice would prevent or diminish the likelihood of compounding injustice. Suppose we had to choose between two white male graduates of a prestigious university, neither of whom belongs to a group that has historically faced prior economic injustice—that is, in terms of educational attainment, wages, occupation, or employment status. Suppose, further, that the only difference between these candidates is their likelihood to contribute to the economy: while both have jobs that are predicted to create a constant rate of new jobs for the next twenty-five years, the predicted job-creation rate of one of these individuals happens to be slightly higher than the other's. In this case, choosing the individual who is more likely to contribute more to the economy, the objection would go, would not compound injustice (because neither candidate has already suffered a relevant injustice). But it would lead to better consequences (given the higher predicted job-creation rate of one of the candidates). The objection that counting economic contribution would be likely to compound injustice would thus not apply to this case; we should count indirect benefits in this case.

This point is not really an objection. It amounts to the idea that it would be permissible to count economic contribution so long as it did not compound injustice—a point that seems to implicitly accept the badness of counting economic contribution if it compounds injustice. Nonetheless, it's worth noting a practical worry with the suggestion of having a rule according to which economic contribution would not play a role if it were likely to compound injustice. And that is: How, exactly, would prior injustice be measured? For instance, would it suffice that a candidate belongs to a group that has a history of unjust disadvantage? If so, how would we proceed in cases in which an individual belongs to a historically disadvantaged group but has never suffered a disadvantage? Giving that individual priority would seem to constitute a case of unfair or unjust *advantage*. Similarly, how would we proceed in cases where an individual who belongs to historically *advantaged* group has faced various unjust disadvantages in his life? Giving this individual priority would constitute a case of unfair disadvantage. Of course, these worries aren't decisive against this proposal, but they do question the promise of its implementation.

Finally, an opponent may note that while it may be true that counting economic contribution when allocating scarce lifesaving resources would amount to compounding injustice against *some* blacks, it does not follow that it would amount to compounding injustice against this group as a whole. Instead, allocating resources based on economic contribution might actually *benefit* this group as a whole. Imagine we could save either a patient who is not very likely to contribute to the economy and is black or a patient who is likely to contribute to the economy and is white. Imagine, further, that the reason the white patient is more likely to contribute to the economy is that he is a very dedicated counselor at a predominantly black and Hispanic, inner-city school, where he has helped a large percentage of high school seniors get into good universities, receive scholarships, and eventually obtain high-paying jobs. A policy that counted economic contribution in terms of, for example, 'ability to help inner city children who belong to historically disadvantaged groups,' would actually make blacks (and Hispanics) *better* off as a whole. In other words, although such a policy may initially seem to compound injustice (i.e., the black patient does not obtain the resource), it would *not* compound injustice in terms of blacks as a group; rather, it would seem to ameliorate the prior injustice that blacks are likely to have faced, as there would presumably be higher rates of blacks and Hispanics attending college and obtaining higher-paying jobs—factors that, according to most measures of wellbeing, go hand in hand with a better quality of life. In short, the idea would seem to be that in this case prioritizing the white counselor does not at the group level compound injustice; rather, it ameliorates it.

But this point, too, is short of an objection. The policy it describes measures economic productivity in terms of 'ability to help inner city children who belong to historically disadvantaged groups'—not in terms of any of the factors that I appealed to as indicative of economic contribution, across which blacks (and Hispanics) fare poorly. This policy is thus by definition aimed at avoiding compounding injustice; it thus implicitly acknowledges the moral significance of putting someone in double jeopardy in the strong sense. Nevertheless, this case also raises various worries. First, it does not at all seem likely that an allocation scheme that counted economic contribution would measure it in terms of ‘ability to help inner city children who belong to historically disadvantaged groups.’ And even if it did, practical worries would arise. For instance, how, exactly, would such an ability be measured? Would it suffice for an individual to be, e.g., a counselor, a teacher, or a nurse at an inner-city school? Would awards for community service at minority-group events do the trick? These questions aren’t decisive against this possibility, but they do highlight that implementing it does not seem a promising task.

In short, even if counting economic contribution is not *necessarily* or *in principle* wrong, there are reasons for being wary of counting it: in the contemporary U.S., counting it would likely compound the injustices that blacks already face. Moreover, although there may be ways to count indirect benefits, including economic contribution, that would avoid compounding injustice, implementing such policies does not seem promising.

5. Implications

Recall the accident victim and the attack victim from Section 2. The idea was that on the QALY system, both would receive lower consideration for a lifesaving drug on the basis of their prior disadvantage: i.e., on the basis that they are paraplegic and thereby have a lower expected quality of life than the individuals who would make a full recovery. What I argued in that section is that the QALY system would seem to put the accident victim in weak double jeopardy, since it would deny her equal consideration for a lifesaving drug on the basis of a prior *unfortunate* disadvantage. But in the case of the attack victim, I argued, the QALY system would seem to put the victim in *strong* double jeopardy, since it would deny her equal consideration for a lifesaving drug on the basis of a prior *unjust* disadvantage. One question that I was after in that section was whether it would really be *worse* for an allocation scheme to disfavor an attack victim than it would be for it to disfavor an accident victim. And I suggested that, setting aside issues about quantifying health states, it did seem worse for an allocation scheme to disfavor the attack victim than it did for it to disfavor the accident victim if the allocating entity was the same entity that brought about the original injustice on which the second injustice was based. In support of this idea, we imagined a case in which it was the allocating entity that attacked the attack victim, thereby leaving her with a poor quality of life. The idea was that this would be worse than a case in which the allocating entity brought about only once injustice to the victim. I now want to examine this idea in more detail, for I think that at least in the case of people with disabilities as a group, there is reason to think that the pre-existing disadvantages that they face are the result of an unjust legal, social, political, and economic patterns, and not of misfortune. Specifically, the question I want to examine is: Does an allocating entity that uses the QALY measure really bring about only one injustice for an individual such as the accident victim, or does it compound the injustices that such an individual faces?

In answering this question, it’s important to note that today, there is a near consensus among disability rights theorists that how we perceive an individual’s ‘disability’ will depend to a great extent not only the disability’s impact on her health, but also on physical, social and attitudinal factors (Bickenbach 2016). On this “social model” of disability, the idea is that ‘disability’ is conceptually linked to an individual’s health state, such that “without an underlying health problem or impairment there is no disability"; but it is not conceptually linked to the individual's quality of life (Bickenbach 2016). An individual's quality of life will depend not just on the impact of the disability on her health state, but also, at least in part, on physical, social, and attitudinal factors.

Notice, though, that people with certain disabilities, much like blacks in the U.S., have historically faced unjust mistreatment. Not long ago, appropriate public accommodations for people with disabilities were not legally required in the U.S., which made it very difficult for people with certain disabilities to obtain a public education and get to work, among other things. It thus does not seem unreasonable to think that the more general lack of environmental resources that (certain) people with disabilities encounter today and which reduce their quality of life is, at least in part, the result of prior *unjust* social, political, legal and attitudinal patterns that have been ratified by governmental institutions. What I mean to suggest is that it may be that much like the ideological understanding of 'race' as a biological concept, which arguably continues to influence the politics, race relations, and economic conditions of the U.S., an ideological understanding of 'disability' as a characteristic that is entirely responsible for an individual's quality of life may have shaped past and present environmental conditions for people with disabilities.[110]

---

[110] If this is right, then to understand 'disability' as an ideology would be to understand it as a set of beliefs that emerges from and is maintained by the historical, economic and material conditions of society, and which serves to make sense of the world.

The idea is, then, that while people with certain disabilities have historically been at a disadvantage, these disadvantages are not *entirely* inevitable, since they are, at least in part, environmentally controllable. But if this is right, then in the accident case above, it's not clear that the QALY measure brings about only *one* injustice for the accident victim. Rather, it seems reasonable to believe that the accident's victim's expected quality of life is lower than the individual who would make a full recovery's as a result of unjust environmental conditions: i.e., a lack of resources that would have generated a higher quality of life score for the accident victim (on a system like the QALY measure). And if *this* is right, then it would not be clear that the attack victim's prior (unjust) disadvantage is any worse than the accident victim's prior (unfortunate) disadvantage, as in both cases that prior disadvantage is in part the result of unjust environmental conditions which, had they been just, could have increased those individuals' expected quality of life. Thus, when the QALY system gives lower priority to people with certain disabilities because they have a lower expected quality of life—a quality of life that is lower because there is a lack of environmental resources that could have increased it—it may, after all, be putting them in double jeopardy in the strong sense, as it would deny them equal consideration for a lifesaving drug on the basis that they were already subjected to a prior *injustice*. In short, it may be that an allocating entity that uses the QALY system puts people with certain disabilities in strong double jeopardy regardless of whether the cause of their disability was misfortune or injustice.

6. *Concluding remarks*

To sum up, while bioethicists tend to agree that counting indirect benefits such as social contribution when resources are scarce and not all can be saved is morally objectionable, the arguments for this view are surprisingly weak. As I argued, however, a policy that counted indirect benefits such as social contribution might put certain individuals in a kind of double jeopardy. But ‘double jeopardy’ can be understood in a weak sense, according to which an allocation scheme may bring about injustice, and in a strong sense, according to which an allocation scheme may compound injustice. In the contemporary U.S., for example, a policy that counted indirect benefits, e.g., in terms of economic contribution, would likely put black Americans in strong double jeopardy. This, I argued, would be worse than if such an allocation policy were to put black Americans in weak double jeopardy. This is because the allocating entity would seem to be the same entity that brought about the original injustice on which the second injustice is based. This does not mean that putting someone in double jeopardy in the weak sense would not be bad. All it means is that putting someone in double jeopardy in the strong sense seems worse if the agent that compounds injustice is the same agent that brought about the original injustice. A lot more needs to be said about how to understand prior injustice, but I hope to have shown at the very least that an allocation policy that puts people in double jeopardy may be wrong in two distinct ways.

Bibliography

Alexander, Shana. 1977. "They Decide Who Lives, Who Dies." In *Ethical Issues in Medicine*, by Robert Hunt and John Arras. Palo Alto, CA: Mayfield Publishing Co.

Anderson, Elizabeth. 2010. *The Imperative of Integration.* Princeton, N.J: Princeton University Press.

Appiah, Anthony. 1990. ""But Would That Still Be Me?" Notes on Gender, "Race." Ethnicity, as Sources of "Identity"." *The Journal of Philosophy* 87 (10): 493-499. https://www.jstor.org/stable/2026866.

Asch, Adrienne. 2003. "Disability equality and prenatal testing: contradictory or compatible." *Florida State University Law Review* 30 (2): 315-42.

Asch, Adrienne. 2000. "Why I haven't changed my mind about prenatal diagnosis: reflections and refinements." In *Prenatal Testing and Disability Rights*, by Erik Parens and Adrienne Asch, 234-258. Washington, DC: Georgetown University Press.

Asch, Adrienne, and David Wasserman. 2015. "Reproductive Testing for Disability." In *The Routledge Companion to Bioethics*, by J D Assas, E Fenton and R Kukla, 417-432. New York: Routledge.

Austin, John Langshaw. 1962. *How to Do Things with Words.* Cambridge: Harvard University Press.

Bach, Kent. 1998. "Speech acts." *Routledge Encyclopedia of Philosophy* 385-390.

Baily, Mary Ann. 2000. "Why I Had Amniocentesis." In *Prenatal Testing and Disability Rights*, by Erik Parens and Adrienne Asch, 64–71. Washington, DC: Georgetown University Press.

Basson, Marc D. 1979. "Choosing among Candidates for Scarce Medical Resources." *Journal of Medicine and Philosophy* 4 (3): 313-333.

Basu, Rima. 2019. "What we epistemically owe to each other." *Philosophical Studies* 1-17.

Basu, Rima, and Mark Schroeder. 2019. "Doxastic wronging." In *Pragmatic Encroachment in Epistemology*, by Brian Kim and Matthew McGrath, 181-205. Routledge.

Bickenbach, Jerome. 2016. "Disability and Health Care Rationing." Edited by Edward N Zalta. *The Stanford Encyclopedia of Philosophy.* https://plato.stanford.edu/archives/spr2016/entries/disability-care-rationing/.

Blum, Lawrence. 2002. *"I'm not a racist but…": The moral quandary of race.* Ithaca: Cornell University Press.

Bognar, Greg. 2008. "Well-being and health." *Health Care Analysis* 16: 97–113.

Brock, Dan. 2003. "Separate Spheres and Indirect Benefits." *Cost Effectiveness and Resource Allocation* 1 (4).

Broome, John. 1994. "Fairness versus doing the most good." *The Hastings Center Report* 24: 36–39.

Buchanan, Allen. 1996. "Choosing who will be disabled: Genetic intervention and the morality of inclusion." *Social Philosophy and Policy* 13: 18–45.

Burns, Kelly. 2014. "Minimizing and Managing Microaggressions in the Philosophy Classroom." *Teaching Philosophy* 37 (3).

Campbell, Bradley, and Jason Manning. 2014. "Microaggression and Moral Cultures." *Comparative Sociology* 13 (6): 692-726.

Carruthers, Peter. 2011. *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford: Oxford University Press.

Darwall, Stephen. 1977. "Two kinds of respect." *Ethics* 88: 36–49.

DeGrazia, David. 2012. *Creation Ethics: Reproduction, Genetics, and Quality of Life.* Oxford: Oxford University Press.

Du Toit, Jessica, and Joseph Millum. 2016. "Are Indirect Benefits Relevant to Health Care Allocation Decisions?" *Journal of Medicine and Philosophy* 41 (5): 540–557. doi:10.1093/jmp/jhw018.

Edwards, S D. 2004. "Disability, identity and the 'expressivist objection'." *Journal of Medical Ethics* 30 (4): 418-420.

Eggleton, Ian, Sue Roberston, Ryan Justine, and Ralph Kober. 1999. "The impact of employment on the quality of life of people with an intellectual disability." *Journal of Vocational Rehabilitation* 13: 95–107.

Fields, Barbara Jeanne. 1990. "Slavery, race, and ideology in the United States of America." *New Left Review* 95-118.

Fisher, Anne. 2015. "How Microaggressions Can Wreck Your Business." *Fortune*, Novermber 19. http://fortune.com/2015/11/19/microaggressions-talent-business/.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Glasgow, Joshua. 2006. "A Third Way in the Race Debate." *The Journal of Political Philosophy* 14 (2): 163-185.

Glover, Jonathan. 2008. *Choosing Children: Genes, Disability, and Design.* Oxford: Oxford University Press.

Green, Emma. 2016. "When doctors refuse to treat LGBT patients." *The Atlantic*, April 19. https://www.theatlantic.com/health/archive/2016/04/medical-religious-exemptions-doctors-therapists-mississippi-tennessee/478797/.

Grice, Herbert Paul . 1989. "Logic and Conversation." In *Studies in the Way of Words*, by H P Grice, 22–40. Cambridge: Harvard University Press.

Grice, Herbert Paul . 1957. "Meaning." *The Philosophical Review* 66: 377-388.

Haidt, Jonathan. 2017. "The Unwisest Idea on Campus: Commentary on Lilienfeld." *Perspectives on Psychological Science* 12 (1): 176-177.

Haidt, Jonathan, and Lee Jussim. 2016. "Hard Truths About Race on Campus." *The Wall Street Journal* .

Hare, Richard Mervyn. 1952. *The Language of Morals.* Oxford: Clarendon Press.

Harris, John. 1987. "QALYfying the value of life." *Journal of Medical Ethics* 13: 117-123.

—. 1985. *The Value of Life: An Introduction to Medical Ethics.* London: Routledge.

Haslanger, Sally. 2008. "A Social Constructionist Analysis of Race." *Revisitng Race in the Genomic Age* 56-69.

Heil, Emily. 2018. "ABC cancels 'Roseanne' after its star, Roseanne Barr, went on a vitriolic and racist Twitter rant." *The Washington Post*, May 29. https://www.washingtonpost.com/news/reliable-source/wp/2018/05/29/roseanne-went-on-a-vitriolic-and-racially-charged-twitter-rampage-now-shes-apologizing/?noredirect=on&utm_term=.bdaa4c0d894e.

Hellman, Deborah. 2008. *When is discrimination wrong?* Cambridge: Harvard University Press.

Hofmann, Bjørn . 2017. "'You are inferior!' Revisiting the expressivist argument." *Bioethics* 31: 505-514.

Holm, Soren. 2008. "The expressivist objection to prenatal diagnosis: Can it be laid to rest?" *Journal of Medical Ethics* 234 (1): 24-25.

James, Michael. 2017. "Race." Edited by Edward N Zalta. *The Stanford Encyclopedia of Philosophy.* https://plato.stanford.edu/archives/spr2017/entries/race/.

Kamm, Frances M. 1993. *Morality, Mortality. Volume 1: Death and Whom to Save From It?* . Oxford: Oxford University Press.

Kaplan, Jonathan Michael, and Rasmus Grønfeldt Winther. 2014. "Realism, Antirealism, and Conventionalism about Race." *Philosophy of Science* (Philosophy of Science , Vol. 81, No. 5 (December 2014), pp. 1039-1052) 81 (5): 1039-1052. https://www.jstor.org/stable/10.1086/678314.

Keller, Johannes. 2005. "In genes we trust: The biological component of psychological essentialism and its relationship to mechanisms of motivated social cognition." *Journal of Personality and Social Psychology* 88: 686–702.

Kerstein, Samuel J. 2017. "Kidney Vouchers and Inequity in Transplantation." *Journal of Medicine and Philosophy* 42 (5): 559-574. doi:10.1093/jmp/jhx020.

Kneski, Laura. 2013. *University of South Florida News: Choose A Better Word.* April 2. Accessed September 10, 2018. http://news.usf.edu/article/templates/?z=38&a=5326.

Kukla, Rebecca, and Katherine Wayne. 2018. "Pregnancy, Birth, and Medicine." *The Stanford Encyclopedia of Philosophy.* Edited by Edward N Zalta. https://plato.stanford.edu/archives/spr2018/entries/ ethics-pregnancy/.

Lassiter, C. 2012. "Implicating without intending on the Gricean account of implicature." *Empedocles: European Journal for the Philosophy of Communication* 4 (2): 199-215.

Leonhardt , David, and Ian Prasad Philbrick. 2018. "Donald Trump's Racism: The Definitive List." *New York Times*, January 15. https://www.nytimes.com/interactive/2018/01/15/opinion/leonhardt-trump-racist.html.

Leslie, Sarah Jane. 2014. "Carving up the social world with generics." *Oxford Studies in Experimental Philosophy* 1: 208-232.

Lilienfeld, Scott O. 2017. "Microaggressions: Strong Claims, Inadequate Evidence." *Perspectives on Psychological Science* 12 (4): 138-169.

Lippert-Rasmussen, Kasper, and Sigurd Lauridsen. 2010. "Justice and the allocation of healthcare resources: should indirect, non-health effects count?" *Medicine, Health Care and Philosophy* 13 (3): 237–246.

Lukianoff, Greg, and Jonathan Haidt. 2015. "The Coddling of the American Mind." *The Atlantic.*

Marmot, Michael. 2015. *The Health Gap: The Challenge of an Unequal World.* London: Bloomsbury.

Matthews, T J, D C Allain, A L Matthews, A Mitchell, S L Santoro, and L Cohen. 2018. "An assessment of health, social, communication, and daily living skills of adults with Down syndrome." *American Journal of Medical Genetics* 1389–1397.

Mattoni, Alejandra. 2017. "The hardest part of getting my tubes tied was making doctors take me seriously." *Broadly*, July 31. https://broadly.vice.com/en_us/article/xwzexq/the-hardest-part-of-getting-my-tubes-tied-was-making-doctors-take-me-seriously.

McMahan, Jeff. 2005. "Preventing the existence of people with disabilities." In *Quality of Life and Human Difference: Genetic Testing, Heath Care, and Disability*, by David Wasserman, Jerome Bickenbach and Robert Wachbroit, 142–171. Cambridge: Cambridge University Press.

McTernan, Emily. 2017. "Microaggressions, Equality and Social Practices." *The Journal of Political Philosophy.* doi:10.1111/jopp.12150.

Mill, John Stuart. 1861. *Utilitarianism.* London: Fraser's Magazine.

Moller, Dan. 2011. "Abortion and moral risk." *Philosophy* 86: 425-443.

Moller, Dan. 2016. "Dilemmas of Political Correctness." *Journal of Practical Ethics* 4: 87-107.

Nelson, James L. 1998. "The meaning of the act: Reflections on the expressive force of reproductive decision making and policies." *Kennedy Institute of Ethics Journal* 8 (2): 165-182.

Parens, Erik, and Adrienne Asch. 2000. "The Disability Critique of Prenatal Genetic Testing: Reflection and Recommendations." In *Prenatal Testing and Disability Rights*, by Erik Parens and Adrienne Asch, 3-43. Washington, DC: Georgetown University Press.

Parfit, Derek. 1997. "Equality and Priority." *Ratio.*

Persad, Govind, Alan Wertheimer, and Ezekiel Emanuel. 2009. "Principles for allocation of scarce medical interventions." *Lancet* 373: 423-31.

Pierce, Chester. 1970. "Black Psychiatry One Year After Miami." *Journal of the National Medical Association* 62: 471-473.

Reinders, Hans. 2000. *The future of the disabled in liberal society: An ethical analysis.* University of Notre Dame Press.

Rescher, Nicholas. 1969. "The Allocation of Exotic Medical Lifesaving Therapy." *Ethics* 79: 173-86.

Rini, Regina. 2015. "Microaggression; macro harm." *LA Times*, October 12. https://www.latimes.com/opinion/op-ed/la-oe-1012-rini-microaggression-solidarity-20151012-story.html.

Robinson, Nathan J. 2018. "The 'Microaggression' Concept." *Current Affairs*, February 17. https://www.currentaffairs.org/2018/02/the-microaggression-concept.

Rogers, Stephanie E, Angela D Thrasher, Yinghui Miao, W John Boscardin, and Alexander K Smith. 2015. "Discrimination in healthcare settings is associated with disability in older adults: Health and retirement study, 2008–2012." *Journal of General Internal Medicine* 30: 1413–20.

Saxton, Marsha. 1997. "Disability rights and selective abortion." In *Abortion Wars*, by R (ed.) Solinger, 374-395. Berkeley: University of California Press.

Scheiber, Noam, and Rachel Abrams. 2018. "Can Training Eliminate Biases? Starbucks Will Test the Thesis." *New York Times*, April 18. https://www.nytimes.com/2018/04/18/business/starbucks-racial-bias-training.html.

Schroeder, Mark. 2018. "When beliefs wrong." *Philosophical Topics* 46 (1): 115-127.

Scott, Rosamund. 2005. "Prenatal testing, reproductive autonomy, and disability interests." *Cambridge Quarterly of Healthcare Ethics* 14: 65–82.

Searle, John. 1969. *Speech Acts: An Essay in the Philosophy of Language.* Cambridge: Cambridge University Press.

Sesardic, Neven. 2010. "Race: A Social Destruction of a Biological Concept." *Biology & Philosophy* 25: 143-162. doi:10.1007/s10539-009-9193-7.

Shakespeare, Tom. 2006. *Disability Rights and Wrongs.* London: Routledge.

Shatin, Leo. 1966. "Medical Care and the Social Worth of a Man." *American Journal of Orthopsychiatry* 36: 96-101.

Shelby, Tommie. 2002. "Is racism in the 'heart'?" *Journal of Social Philosophy* 33 (3): 4110420.

Shrage, Laurie. 2017. "Race, Health Disparities, Incarceration, and Structural Inequality." In *The Oxford Handbook of Philosophy and Race*, by Naomi Zack.

Skenazy, Lenore, and Jonathan Haidt. 2017. "The Fragile Generation." *Reason: Free Minds and Free Markets*, December. https://reason.com/2017/10/26/the-fragile-generation.

Sloper, P, and S Turner. 1996. "Progress in social-independent functioning of young people with Down's syndrome." *Journal of Intellectual Disability Research* 40: 39-48.

Sue, Derald Wing. 2017. "Microaggressions and ‘Evidence’: Empirical or Experiential Reality." *Perspective on Psychological Science* 12 (1): 170-172.

Sue, Derald Wing, Christina M Capodilupo, Gina C Torino, Jennifer M Bucceri, Aisha M.B. Holder, Kevin L Nadal, and Marta Esquilin. 2007. "Racial Microaggressions in Everyday Life: Implications for Clinical Practice." *American Psychologist* 62 (4): 271-286.

Sue, Derald Wing, Christina M Capodilupo, Gina C Torino, Jennifer M Bucceri, Aisha M.B. Holder, Kevin L Nadal, and Marta Esquilin. 2007. "Racial Microaggressions in Everyday Life: Implications for Clinical Practice." *American Psychologist* 62 (4): 271-286.

Taylor, Peter. 2011. "Rehabilitating a Biological Notion of Race? A Response to Sesardic." *Biology & Philosophy* 26: 469-473. doi:10.1007/s10539-011-9249-3.

Tello, Monique. 2017. "Racism and discrimination in health care: Providers and patients." *Harvard Health Blog*, January 16. https://www.health.harvard.edu/blog/racism-discrimination-health-care-providers-patients-2017011611015.

The World Health Organization. 2017. "Health Impact Assessment (HIA): The determinants of health."

Tschaepe, Mark. 2016. "Addressing Microaggressions and Epistemic Injustice: Flourishing from the Work of Audre Lorde." *Essays in the Philosophy of Humanism* 24 (1): 87–101.

University of Maryland, Division of Student Affairs. 2014. *University of Maryland, Division of Student Affairs.* Accessed October 15, 2018. https://dsa.umd.edu/events/itooammaryland-responses-to-racial-microaggressions.

U.S. Bureau of Labor Statistics. 2017. "Labor force characteristics by race and ethnicity, 2016." Report 1070. https://www.bls.gov/opub/reports/race-and-ethnicity/2016/home.htm.

U.S. Department of Health and Human Services, Centers for Disease Control and Prevention. 2013. "CDC Health Disparities and Inequalities Report — United States, 2013." https://www.cdc.gov/mmwr/pdf/other/su6203.pdf.

Wilkinson, Richard, and Michael Marmot. 2003. *Social Determinants of Health: The Solid Facts, 2nd Edition.* . WHO Library Cataloguing in Publication Data.

Williams, Thomas D. 2018. "Calling America ‘Land of Opportunity’ Is a ‘Microaggression’ Says University of California." *Breitbart*, March 17.

https://www.breitbart.com/politics/2018/03/17/calling-america-land-opportunity-microaggression-says-university-california/.

Wolf, Jonathan, and Avner de-Shalit. 2007. *Disadvantage.* New York: Oxford Unviersity Press.

Young, Iris Marion. 2001. "Equality of Whom? Social Groups and Judgments of Injustice." *Journal of Political Philosophy* 9: 1-18.