

2022

Est-ce que Vous Compute? Code-Switching, Cultural Identity, and AI

Arianna Falbo
Bentley University
afalbo@bentley.edu

Travis LaCroix
Dalhousie University
tlacroix@dal.ca

Recommended Citation

Falbo, Arianna, and Travis LaCroix. 2022. "Est-ce que Vous Compute? Code-Switching, Cultural Identity, and AI." *Feminist Philosophy Quarterly* 8 (3/4). Article 9.

Est-ce que Vous Compute? Code-Switching, Cultural Identity, and AI

Arianna Falbo, and Travis LaCroix

Abstract

Cultural code-switching concerns how we adjust our overall behaviours, manners of speaking, and appearance in response to a perceived change in our social environment. We defend the need to investigate cultural code-switching capacities in artificial intelligence systems. We explore a series of ethical and epistemic issues that arise when bringing cultural code-switching to bear on artificial intelligence. Building upon Dotson’s (2014) analysis of testimonial smothering, we discuss how emerging technologies in AI can give rise to epistemic oppression, and specifically, a form of self-silencing that we call “cultural smothering.” By leaving the sociodynamic features of cultural code-switching unaddressed, AI systems risk negatively impacting already-marginalised social groups by widening opportunity gaps and further entrenching social inequalities.

Keywords: code-switching, artificial intelligence, machine learning, natural language processing, epistemic oppression, testimonial smothering, cultural smothering

1. Introduction

In linguistics, *code-switching* describes using two (or more) languages (or language varieties) within a single text, conversation, or utterance.¹ This purely linguistic sense of code-switching has been widely discussed in the context of artificial intelligence (AI)—especially for natural language processing (NLP) models.² AI/NLP research on code-switching typically concentrates on the morphosyntactic features of

¹ The term “code-switching” (sometimes referred to as code-mixing, codeshifting, language alternation, language mixture, or language switching) appears in the 1950s; however, observations of these linguistic phenomena in academic writing predate this baptism by several decades; see Benson (2001) for an historical discussion. Code-switching, like many other communicative phenomena, appears to be governed by several linguistic and extralinguistic features; see discussion in, e.g., Gumperz (1977), Pfaff (1979), Poplack (1980), and Benson (2001).

² See Jose et al. (2020) for a recent survey.

code-switched linguistic data at the sentential (or subsentential) level. However, a broader characterisation of code-switching phenomena is primarily *social*—namely, *cultural* code-switching. Cultural code-switching relates to how we adapt our overall behaviour, manner of speaking, and appearance in response to a perceived change in our social environment. Unlike the merely linguistic sense of code-switching, cultural code-switching involves “a much more profound shift than toggling between languages or dialects does” (Morton 2019, 76).³ Cultural code-switching is more intimately connected to conceptions of oneself. It is also a mechanism by which we can conform to dominant cultural norms, and by the same token, it is a means by which we can defect and resist the status quo.

As AI-based language technologies continue to improve and become more pervasive in society, it will be increasingly essential that they fluidly interact with humans—that is, in the way that humans typically interact with one another. However, a significant part of such interactions involves nonlinguistic communication—that is, the additional features of communication that are relevant for cultural code-switching. Engagement in cultural code-switching is especially common among minority communities, and it may be done for a variety of reasons—conforming to majority-group norms, gaining social acceptance, avoiding hostility or discrimination, or achieving social mobility, to name a few. Accordingly, the extant biases of emerging technologies, which typically arise from dominant-group conventions, may be further exacerbated.⁴ To mitigate, and not further contribute to, these biases, it is important for researchers in AI—especially those working in the machine learning paradigm and other emerging methods—to take cultural code-switching into account in models of NLP. This is particularly true when these systems are deployed, especially in high-stakes social environments (e.g., legal settings, job interviews, educational institutions) or when they are otherwise integrated into standard consumer-facing technologies or applications.

The goals of this paper are twofold, in light of these considerations. First, we underscore the importance of cultural code-switching in research on the ethics of AI. Despite its ubiquity and, more importantly, the unique value and costs that cultural code-switching has for those from marginalised backgrounds, surprisingly little attention has been given to these phenomena, either in the technical literature on

³ See also the definitions and related discussions found in Morton (2014), McCluney et al. (2019), and Falbo (2021).

⁴ It is, by now, well documented that many instances of machine bias result from training, where the technological artefact mirrors extant biases in datasets. See discussion in, e.g., Friedman and Nissenbaum (1996), Angwin et al. (2016), O’Neil (2016), Caliskan, Bryson, and Narayanan (2017), Cave and Dihal (2020), and Gabrielle Johnson (2021).

NLP models or in the philosophical literature on AI more broadly. This paper bridges this gap by defending the need to investigate cultural code-switching capacities in AI systems (AIS).

Second, drawing upon the growing literature on epistemic oppression and related work on structural oppression, we canvas the potential moral and epistemic risks involved in implementing—or, more importantly, *failing* to implement—code-switching capacities in AIS. By leaving the sociodynamic features of cultural code-switching unaddressed, AIS risk negatively impacting already marginalised social groups by widening opportunity gaps and further entrenching social inequalities.⁵ More generally, reflecting upon the need for cultural code-switching in AIS allows us to enrich our understanding of these phenomena beyond the context of human-AI interaction. Cultural code-switching influences not only first-personal identity—particularly what Dembroff and Saint-Croix (2019) have called *agential identities*—but also broader systems of social coordination within and between groups.

The paper proceeds as follows. Section 2 provides a brief overview of present-day artificial intelligence and machine learning methods, highlighting extant biases that have surfaced. This section also clarifies that the AI research area that is most relevant to code-switching is *natural language processing*. However, as mentioned, the code-switching phenomena we are primarily concerned with are broader than the linguistic context alone. Accordingly, section 3 clarifies what *cultural* code-switching is and some of the reasons we do it, using specific examples. Section 4 draws upon the literature on epistemic oppression and, in particular, on Dotson’s (2014) analysis of *testimonial smothering*. In so doing, we broaden Dotson’s framework to include a species of self-silencing, which, following Falbo (2021), we call *cultural smothering*. This occurs when cultural code-switching behaviours manifest as a form of self-censoring: one alters aspects of their cultural identity in response to an unwelcoming or hostile social atmosphere. We also discuss how pressures to code-switch in this sense, and thus to engage in or succumb to cultural smothering, often have the structure of a *double bind* choice situation (Frye 1983).

Section 5 draws some tentative conclusions concerning the relationship between cultural code-switching and emerging technologies. We must be sensitive to how practices of silencing are sustained and reinforced in society and how this maintains unjust social arrangements, thereby limiting access to social goods needed to improve the material conditions of marginalised communities (e.g., safety, education, job opportunities). A natural extension of this carries over to the domain

⁵ See recent work by Justin Bruner, Cailin O’Connor, and Travis LaCroix (Bruner and O’Connor 2017; Bruner 2019; O’Connor and Bruner 2019; LaCroix and O’Connor, forthcoming) for discussion on how the constitution of groups and group dynamics alone may give rise to inequitable conventions.

of emerging technologies, particularly AI and NLP. We should be critical of how these same patterns of silencing and epistemic oppression are reproduced and potentially rendered more potent and efficacious in emerging technologies—especially as these technologies become increasingly widespread and seamlessly integrated into everyday life.

To be clear, it is not our intention to defend any concrete solutions to the problems we identify. We don't offer any optimistic plans for developing or training AI to avoid or ameliorate epistemic oppression. We are not confident that any such solution ultimately exists. Instead, we offer these arguments as a serious caution, as a lesson concerning the limitations and harms of AI technologies when they are adopted and implemented against the backdrop of a nonideal, hierarchically structured social world, filled with unjust divisions on the bases of race, gender, sexual orientation, ability, religion, and class, among others. Unless and until the social world changes for the better, we should expect these problems to persist.

2. Artificial Intelligence, Machine Learning, and Natural Language Processing

Artificial intelligence describes both a property of computer systems—displaying intelligent behaviour⁶—as well as a set of techniques that researchers use to achieve this property (Gabriel 2020). Machine learning (ML) is a branch of artificial intelligence wherein algorithms use data to learn gradually. The three main approaches to machine learning are supervised learning, unsupervised learning, and reinforcement learning.

In supervised learning, the algorithm is trained on labelled examples; this is the *training data*. The model is evaluated on how well it generalises what it learns from the training data to previously unseen examples (called test data). For example, one might train an image-recognition model on millions of labelled images; if it performs well, it should correctly label novel images not in the training set. In unsupervised learning, an algorithm learns underlying patterns or correlations from unlabelled data. For example, recommendation systems categorise users based on their viewing patterns in order to recommend similar content. Reinforcement learning depends upon sparse rewards for actions. For example, a chess game has a reward of +1 if a player wins, -1 if a player loses, and 0 if a player draws. A reinforcement learning model can learn to play chess merely by playing hundreds of thousands of games and receiving a reward at the end of each game.

Machine learning techniques stand in contrast to *symbolic systems* (or “good old-fashioned AI”), where explicit, handcrafted rules had to be hard-coded into the machine—typically in the form of complex, nested if-then statements. The current

⁶ Where “intelligence” is understood as “an agent’s ability to achieve goals in a wide range of environments” (Legg and Hutter, 2007).

driver of AI research is so-called *deep learning*, an approach to machine learning that utilises deep neural networks modelled (roughly) after neurons in the human brain. Deep learning uses layers of algorithms to process data—information is passed through each subsequent layer in a neural network, with the previous layer’s output providing input for the subsequent layer. One of the key advantages of deep learning techniques is that they do not require the heavily handcrafted features used by traditional machine learning methods. Deep convolutional neural networks are differentiated from their shallow counterparts primarily in terms of their depth, heterogeneity, and sparse connectivity (Buckner 2019).

In the early days of machine learning, researchers could focus on the fundamental aspects of their work without much concern for the social or ethical consequences of these systems because they were relatively encapsulated in the confines of the research lab. However, Luccioni and Bengio (2019) highlight that these algorithms are increasingly being deployed in society due in part to the promise of the unprecedented economic impacts of ML applications (Bughin et al. 2018; Szczepański 2019; Russell 2019). Consequently, these systems’ ethical and social impacts have started to be examined by researchers from a wide range of disciplines.

Since 2016, emerging technologies’ ethical and social consequences have just begun to have a more prominent role in AI research more generally (Christian 2020). One worry is that algorithms may latch on to spurious correlations (overfitting the training data), which can lead to behaviour that we would call, for example, sexist or racist (even if implicitly so) when performed by a human. Another slightly different worry is that the data used to train an algorithm is inherently biased.⁷ Hence an optimally trained model will perform exactly how it should, given the data it has received, again leading to potentially problematic patterns of “behaviour.” A final worry is that the data is simply incomplete or unrepresentative, meaning that the trained system may be “oblivious” to certain (potentially huge) swaths of society.

For example, a health algorithm that used health costs as a proxy for health *needs* learned to be biased against Black patients in the United States (Obermeyer et al. 2019). Similarly, a recidivism-prediction algorithm is twice as likely to have false positives (labelling individuals as high risk for re-offence when they are in fact low risk) for Black individuals and twice as likely to have false negatives (labelling individuals as low risk for re-offence when they are in fact high risk) for white individuals (Angwin et al. 2016).

To put it another way, at *best*, a sophisticated ML system can be an efficient tool that reflects exactly what the data tell it about the real world, meaning that if the data are biased—which they typically are—then the ML system will perpetuate the existing biases or inequalities in the social system that gave rise to those data.

⁷ See discussion in Green (2019) and Gabrielle Johnson (2021).

As mentioned in section 1, cultural code-switching is especially common among minority communities, implying differential values and risks for individuals in minority groups compared with those in majority groups. Thus, biases in AI technologies can come to bear in significant ways on cultural code-switching (which we discuss in more detail in the next section). Owing to the communicative nature of code-switching, the field of artificial intelligence in which an analysis of code-switching will be most relevant is *natural language processing*, which studies natural language interactions between computers and humans.

Some of the key areas of study in NLP research include speech recognition, which may involve recognising spoken language and translating it into text format; natural language understanding or interpretation, which is necessary for, for example, question-and-answer interactions or machine translation; and natural language generation, which outputs text. NLP has many important applications, including virtual assistants, spam detection, speech recognition, named-entity recognition, sentiment analysis, question answering, automatic text summary, autocomplete, predictive typing, relationship extraction, and machine translation, to name a few.

For many years, researchers have utilised shallow models (such as support vector machines or logistic regression) or explicit symbolic representations (such as first-order logic) to solve particular NLP tasks. However, the advent of *deep learning* techniques has allowed for significant progress in NLP research in the last decade. You are probably familiar with NLP, even if you have not heard of it—NLP models are necessary for consumer-facing applications like Google Assistant, Siri, or Alexa to work.

Consider the last of these. Amazon’s Alexa is a virtual assistant AI system that makes significant use of NLP models. It constantly monitors signals in the environment and processes them to minimise ambient noise (such as the conversation on the television) that is not relevant to its task. This preprocessing is always going on in the background. For the software to turn on, it requires a “wake word”—for example, *hey Alexa!* Thus, before Alexa even does anything, it must be capable of differentiating speech from background noise, detecting very specific signals, and responding to them. Once Alexa wakes, it converts the recorded audio to text by analysing certain features like frequency and pitch—this is *speech recognition*. Once the speech has been converted to text, Alexa must interpret the meaning of (at least parts of) the text to respond appropriately—this is natural language *understanding*. Once the request is processed and understood, Alexa may need to produce a “voiced” response.

However, NLP research generally starts with the assumption that language is solely about the transfer of information—that is, the *content* of the message. As mentioned in the introduction, there are many extralinguistic factors involved in cultural code-switching, so this emphasis on content is myopic. Even granting the view

that language is primarily a way of communicating or transferring information, language is but one way of doing so.

Two of the leading language models in machine learning are *Bidirectional Encoder Representations from Transformers*—also known as BERT (Devlin et al. 2019)—and *Generative Pre-trained Transformer 3*—also known as GPT-3 (Brown et al. 2020)—which is the third iteration of an autoregressive language model, produced by OpenAI. Let’s consider the latter of these.

The full version of GPT-3 makes use of around 175 billion parameters—two orders of magnitude larger than its predecessor, GPT-2, and one order of magnitude larger than the next-largest NLP model, Microsoft’s *Turing NLG*.⁸ The majority of training data for this model (around 60 percent) comes from a dataset called *Common Crawl*—a corpus generated from crawling the Internet—which, as of April 2021, contains around 320 *tebibytes* (TiB) of data taken from 3.1 billion web pages. GPT3 used a filtered version of the Common Crawl dataset, resulting in approximately fourteen billion tokens (though they do not provide details of how the Common Crawl dataset was curated).

To give a sense of this dataset’s scale, GPT-3 was also trained, in part, on the *entire* English-language Wikipedia, which consists of three billion tokens and accounted for only around 3 percent of the training data for the model. Additional datasets used to train GPT-3 include WebText2—consisting of the text of web pages from all outbound Reddit links from posts with three or more “upvotes”—and Books1 and Books2—which are two Internet-based books corpora.

Note, however, that the Internet—and therefore, the data used to train GPT3—is predominantly *English*: around 45 percent of HTML pages in the Common Crawl dataset have English as their primary language. It is estimated that English is used by over 61 percent of all the websites whose content language we know, with the next most predominant language being Russian, at 5.5 percent (Web Technology Surveys 2022).

Thus, there is likely an Anglocentric bias in the data used to train large language models. However, it is perhaps relevant to note that it is very difficult to make claims about the training of these models with any degree of certainty. For example, OpenAI does not describe the filtering process for the Common Crawl dataset, and the WebText2 dataset and the books corpora are not publicly available and lack official documentation (Bandy and Vincent 2021).

⁸ Since the release of GPT-3, this model has been surpassed in size by Google’s 1.7 trillion-parameter language model. This show of one-upmanship is criticised by Bender et al. (2021), who highlight, among other things, the environmental costs of training large language models.

Furthermore, as others have pointed out, the Common Crawl dataset will overrepresent those populations with higher Internet usage rates (Luccioni and Viviano 2021). And those countries with the highest Internet penetration rates are European (Northern, Western, and Southern) and North American, with the lowest rates coming from Africa and South Asia (J. Johnson 2021). This discrepancy leads to a “digital language divide” on the Internet—despite there being around seven thousand languages spoken worldwide, only a few hundred of these are represented on the Internet (Young 2015). Thus, the datasets being used to train these language models are inherently biased, and there has (to date) been relatively little analysis of these biases.

In some ways, the simple case is the one where an algorithm is trained on a specific language; yet the technologies are highly biased in favour of Anglophones. More difficult cases involve particular accents or dialects, and there is already overwhelming evidence that NLP systems are biased—this time in favour of a particular *type* of Anglophone: one whose idiolect resembles something like “broadcast English” (Harwell 2018). Thus, the ubiquity of code-switching highlights a further difficulty that these already inherently flawed systems will face.

The biases that we have been discussing in NLP models have been primarily linguistic in nature. However, as mentioned in the introduction, there is a broader sense of code-switching that we take to be pertinent here—namely, cultural code-switching. In the next section, we discuss how the pressures to code-switch—to adapt, minimise, suppress, mask, hide, or conform—are especially pressing among already disadvantaged populations. Often, code-switching is a skill that is required to gain access to safety and opportunities for improving one’s material conditions. But at what cost?

3. Cultural Code-Switching

Cultural code-switching is ubiquitous in human interactions: it is used to switch between formal and informal contexts; to wield power and to signal asymmetric authority relations (Popa-Wyatt and Wyatt 2018; Tirrell 2012); and can express a sense of agency over one’s social identity and serve to build and reinforce a sense of community and solidarity by signalling in-group affiliation (Dembroff and Saint-Croix 2019). Since cultural code-switching takes account of the socio-pragmatic aspects of communication, it includes extralinguistic cultural artefacts (in addition to linguistic ones). For example, consider how a student may modulate their tone of voice or change the register of their speech by using honorifics like “ma’am,” “sir,” “doctor,” or “professor” when interacting with faculty. Thus, instead of the narrow sense of altering some linguistic features of communication, we take code-switching more broadly to include nonlinguistic communicative signals which convey information—for example, about oneself and one’s relation to others in a given social context.

Why do we code-switch? The reasons for code-switching are vast and complex. One important use of code-switching is to signal in-group membership and affiliation. Consider how teenagers (in a broadly North American context) will often use various slang terms and expressions when talking with each other. A quick search through Urban Dictionary, a crowd-sourced online lexicon used to archive slang words and phrases, reveals a host of in-group expressions of this sort—for example, “lit,” “salty,” “bae,” “simp,” “noob,” “ghosting,” “dank,” “taking an L,” and many more. If you don’t know what these expressions mean, that is part of the point. When young people use these expressions among friends, they are code-switching (whether they realise it or not). They are signalling to one’s audience that one is “in the know,” “one of them,” or as they might colloquially put it, “cool.”

Consider another example from the LGBTQ2+ community, particularly the subculture of drag queens (*Them* 2018). If you go to a drag show, there are a range of unspoken cultural norms and expectations that are operative. If it is your first time, you might be taken aback by the playful, mock-impolite, bantering style that drag performers often use. In some cases, drag queens might be “throwing shade”: a kind of indirect insult toward other drag performers or audience members. This phenomenon is discussed in Jennie Livingston’s (1990) *Paris Is Burning*, a documentary on drag culture in New York City. Dorian Corey, a drag queen interviewed in the film, gives the following example of throwing shade.

If I were to say in a terribly condescending voice, “Oh honey, I’m so glad you saved up to buy those glasses,” that’s blatant shade. I didn’t insult the glasses, or you, directly. It’s implied by my voice and the context of what I said. You know they’re ugly.

It is important to note the central role that tone of voice and context play in this example. This is a key issue which we will return to in discussing the relationship between the phenomena of cultural code-switching and human-to-AI interactions. More broadly, this style of communicative engagement has been identified as a kind of pro-social in-group communication, which builds community and solidarity within LGBTQ2+ communities. Additionally, mock impoliteness is also a rhetorical strategy that can be deployed as a form of self-protection; it is a way of preemptively developing a thick skin against bigotry and discrimination (McKinnon 2017; Oliva, Antonialli, and Gomes 2021).

As another example, cultural code-switching in the context of race plays a major role in the plot of Boots Riley’s (2018) black-comedy film, *Sorry to Bother You*. The film centres on a young, Black man named Cassius Green (played by LaKeith Stanfield) who gets a job as a telemarketer and finds major success after he is coached by an older (also Black) coworker (played by Danny Glover) to use his “white voice.”

White actors voice the “white voices” for these characters. One of the major themes of this film is how being successful in a predominantly white culture requires conforming to the expectations of the dominant group—in this case, a marginalised Black man “sounding” white to achieve success at work. According to Touré (2018), “[putting] on a white voice” means to embrace “the ease that white privilege brings. It means sounding as if you’re entitled to the good life. It means feeling calm way down in your soul. It means never having to be afraid someone will call the police on you just because you’re breathing.” However, the film also explores the social consequences of code-switching with respect to one’s in-group: Green ends up alienating (or being alienated from) his friends, family, and social in-group.

This resonates with a broader experience of *double alienation* that individuals from marginalised groups often confront as they find themselves existing between disparate social worlds with distinct and often incompatible cultural norms, values, and expectations. Morton (2019) describes how first-generation scholars often confront ethical compromises in the process of trying to gain upward social mobility through pursuing a degree in higher education. In conforming to the dominant norms operative in institutes of higher education (especially within elite institutions with predominately upper-class and white student and faculty populations), students from first-generation or low-income backgrounds often feel isolated and less able to connect with family and loved ones. This is perhaps especially true within fields like philosophy, where first-generation and low-income academics may be unable to fully share their academic life with family and loved ones, who may not understand the point or value in pursuing a career as an academic philosopher. In a recent paper, Morton, who is a first-generation scholar, says of her family: “We love each other, but I am now part of a world whose logic is mysterious to them” (2021, 10). Thus, the process of code-switching can give rise to serious ethical trade-offs and compromises.⁹

More generally, one key use of code-switching is cultivating a sense of belonging and community with those in one’s social group. One code-switches to express a shared social identity with others; to make others feel comfortable, at ease, or “at home” in a given social environment. This is just one—distinctively *positive* and *self-affirming*—use of code-switching. Other uses of code-switching are starkly different.

In some cases, one might code-switch not as a means to increase a sense of community with others but rather to navigate an unfamiliar and potentially unwelcoming (or even hostile) social environment. In this sense, the “masking” behaviours that autistic individuals may perform—faking eye contact, mirroring

⁹ For further discussions of related issues, see Morton (2019); and for discussion on the case of academic philosophy specifically, see the recent special issue of the *APA Newsletter on Feminism and Philosophy* (Falbo and Stewart 2021).

gestures and expressions, scripting conversations, disguising “stimming” behaviour, enduring sensory discomfort, and so on—are a type of cultural code-switching (Hull et al. 2017), which may be employed in a range of social situations, including feeling safe in a culture that has, historically, highly stigmatised autistic individuals (Silberman 2015; Donvan and Zucker 2016).

Importantly, for our purposes, each of the behaviours just described is *nonlinguistic*. In this case, regular practice of cultural code-switching can have inherently negative side effects, including increased stress, anxiety, depression, and exhaustion or “autistic burnout” (Bargiela, Steward, and Mandy 2016; Cage, Di Monaco, and Newell 2018; Cage and Troxell-Whitman, 2019; Raymaker et al. 2020), or loss of identity and increased risk of suicidal thoughts (Cassidy et al. 2020). In this case, too, certain intersectional groups may feel differential social pressure to code-switch—several studies have suggested that people who identify as women are more prone to code-switching than those who identify as men, leading women to be misdiagnosed and partially causing the gender gap in autism diagnoses (Gould and Ashton-Smith 2011; Dworzynski et al. 2012; Lai et al. 2015; Hull et al. 2019).

Thus, code-switching is a crucially important skill for negotiating the social dynamics of high-stakes environments—for example, interactions with the police or in courtrooms, educational institutions, job interviews, and more. Moreover, the pressures to code-switch are not felt uniformly but are disproportionately experienced by members of historically marginalised groups, who often must assimilate to dominant cultural norms to gain social acceptance and access certain resources or benefits or avoid potential harms and mistreatment.

For example, empirical studies and first-person testimonies have shown that Black women with natural hairstyles (e.g., Afros, braids, twists) are less likely to get job interviews than Black women with straightened hair, especially compared to white women (Koval and Rosette 2021). Black women with natural hair received lower scores on assessments of “professionalism” and “competence,” and they were less likely to be invited for job interviews as a result. Most notably, this occurred when the norms of the job required a more conservative and formal dress code. One striking example is the 2017 case of Destiny Tompkins, a young Black woman who was told by her manager (a white woman) at Banana Republic that her braided hair was “unkempt” and “too urban” for the store’s image. The manager told her to take out her braids or else she would stop scheduling her for shifts (Samotin 2017). In this context, Black women might decide to code-switch—for example, by straightening their hair for job interviews—in order to assimilate to norms of “professionalism” in white-dominated spaces. This is often needed to ensure job security and social acceptance.

Of course, code-switching in high-stakes settings, especially where particular opportunities for social mobility and benefits reside, involves more than changes in

one's physical appearance. Dembroff and Saint-Croix (2019) emphasise the relationship between cultural code-switching and what they call *agential identities*. They define agential identities as “the self-identities we make available to others—they bridge what we take ourselves to be with what others take us to be” (Dembroff and Saint-Croix 2019, 572). They discuss how code-switching is an effective way to negotiate and switch between *entire social identities*, in some cases moving between genuine identities and merely apparent or superficial social identities.

To make an explicit connection between the concepts discussed in this section and the claims of section 2, consider the following concrete example of how code-switching is implicated in NLP. We highlighted how NLP models typically encode the conventions of *dominant* groups. As a result, NLP is severely Anglocentric.¹⁰ Furthermore, we emphasise how researchers typically understand language primarily in terms of its *syntactic* features. But, even within the Anglophone community, different *manners* of communication abound, and these differences are not necessarily exclusive to syntax. Thus, suppose an NLP model is trained on Common Crawl and other Anglocentric datasets, as with GPT-3. The model will, of course, be biased in favour of dominant-group conventions—in this case, communication conventions, which include morphosyntactic features of particular languages but also, as we have highlighted, other features relevant to communication within and between groups. Then, any application that uses this model will encode those biases, leading to the risk of *double alienation* for those individuals who are not members of the dominant group but who have to interact with the application.¹¹

This feature of NLP models reflects exactly the insight that code-switching pressures are felt primarily by those who are historically marginalised. Thus, these models, and the applications that utilise them, give rise to negative feedback loops which perpetuate and reproduce harms faced by these groups, thus further entrenching societal structural inequalities. For example, consider the application of AI tools that make use of NLP to detect toxic speech online on social media platforms like Twitter. Recent studies have shown that AI tools are far more likely to generate false positives—that is, misidentifying tweets as containing offensive or toxic content—when the tweets are written in African American English (Sap et al. 2019). In some cases, when such content is flagged, this can ultimately serve to censor and, thus, silence those who are speaking from nondominant social positions—specifically,

¹⁰ NLP models, in particular, are often significantly biased in several other ways as well, including gender and cultural biases in the datasets on which these models are trained. See discussion in LaCroix and Luccioni (2022).

¹¹ This is not unrelated to the Cultural Red King effect, discussed by Justin Bruner and Cailin O'Connor (Bruner and O'Connor 2017; O'Connor 2017; Bruner 2019; O'Connor and Bruner 2019).

users who make use of linguistic practices which are unfamiliar to these AI tools and which fall outside the status quo.¹²

At the most general level, cultural code-switching influences one's overall mode of existence in social space—how one chooses (or is forced) to exist in the social world. And technology is increasingly implicated in these social spaces and the interactions within them. Particular social environments and cultural contexts that incentivise cultural code-switching may thus be sites of epistemic oppression where silencing thrives. As long as NLP seeks to furnish AI systems with human-level linguistic capacities, it will be important to attend to code-switching behaviour, since cultural code-switching often depends upon the communicative signals that NLP seeks to analyse and “understand.”

4. Cultural Smothering and Double Binds

When one is forced to code-switch in the ways discussed in the previous section—to mask, blend in, or otherwise express a superficial identity or persona to be accepted, make a living, or avoid harms—it results in a form of self-censorship or self-silencing. This phenomenon is akin to what Dotson (2011) has analysed as *testimonial smothering*. On Dotson's analysis, testimonial smothering is a form of self-silencing that occurs when a speaker truncates or self-censors their testimony because their audience is taken to lack the required competence needed to properly understand what one is saying. This amounts to a kind of *epistemic violence* when the testimonial recipient exhibits *pernicious ignorance*. Pernicious ignorance, according to Dotson, is a form of reliable ignorance, or reliable insensitivity to the truth, that results in a harm.¹³

Dotson's analysis of testimonial smothering highlights the importance of testimonial competence—the ability of an interlocutor to *get what you mean* (in the way you mean it)—for successful communicative exchanges. Furthermore, this analysis describes how particular speakers, especially those who are testifying from disadvantaged positions, are prone to vulnerabilities and harms when their audience lacks such a competence. Dotson (2011, 238) writes,

Speakers are vulnerable in linguistic exchanges because an audience may or may not meet the linguistic needs of a given speaker in a given exchange. . . . To communicate we all need an audience willing and capable of hearing us.

¹² Also, see Oliva, Antonialli, and Gomes (2021) for a related discussion concerning the linguistic practice of “mock impoliteness” in queer communities.

¹³ Also see, for example, Dotson (2014), Medina (2013), Pohlhaus (2012), and Fricker (2007) for related discussion.

Moreover, engaging in communicative exchanges where one's audience lacks testimonial competence (or where it is unclear if one's interlocutor is competent) can be very risky and potentially unsafe (Dotson, 2011, 244–45). For example, drawing upon the work of Crenshaw (1991), Dotson explains how Black women's testimony about sexual violence perpetrated by Black men is potentially unsafe or risky in social contexts where it may be interpreted as reinforcing racist stereotypes.¹⁴

Dotson's analysis focuses on *testimonial* exchanges—how testifiers from oppressed groups are often forced to capitulate to the testimonial incompetency of their audiences by altering their testimony to include only that which is likely to be given proper uptake. However, with cultural code-switching in view, we can broaden Dotson's analysis to apply not only to the truncating or smothering of testimony but also to the smothering of aspects of one's cultural identity more broadly. This is salient in cases of forced code-switching, where one masks aspects of their culture in response to dominant norms and pressures within the social environment, especially where code-switching is required to gain some benefit or avoid harms. In such cases, one may not only smother their testimony but also engage in a broader form of what we will refer to as *cultural smothering* (Falbo 2021). When cultural code-switching behaviours manifest as cultural smothering, they take on a form of self-silencing.

The pressures to engage in cultural smothering through code-switching arise in nonideal-choice situations. Such situations often have the structure of a *double bind*, which is described as a "situation in which options are reduced to a very few and all of them expose one to penalty, censure or deprivation" (Frye 1983, 2). For example, when deciding whether to have (heterosexual) sex, Frye discusses how women are often susceptible to social criticism and scrutiny no matter what they do. If they have sex, they are prone to be called "easy" or a "slut," but if they do not, they are seen as "prudish," "uptight," or "frigid." No matter what a woman chooses, she loses. The disadvantage that double binds give rise to is a function of their structural properties; therefore, they are sites of systemic oppression. They result from, reinforce, and sustain oppressive systems and unjust social arrangements.

Recently, Hirji (2021, 645) has argued that what makes double binds unique and effective vehicles of oppression, compared to other difficult choice situations, is that "whatever an agent does necessarily undermines their own objective interests." Hirji develops the notion of an "imperfect choice," which is a kind of choice situation where it's impossible to advance one's interests and where one will inevitably be worse off as a result. Imperfect choices constrain one's agency while nonetheless leaving aspects of one's autonomy intact. One is still free to choose, but this freedom is hollow and illusory because all available options undermine one's interests.

¹⁴ For related discussion of rape myths in relation to epistemic injustice, also see, for example, Jenkins (2017), Yap (2017), and Falbo (2022).

Moreover, in so choosing, one is forced to be complicit in their own oppression. Double binds present a series of options, all of which are nonideal and further contribute to systems of oppression that function to make one, and members of one's social group, worse off.

How does this relate to cultural code-switching? In cases where one must code-switch to gain some advantage or avoid some harm in the social environment, one typically faces an imperfect choice. They are in a double bind. On the one hand, one might choose to code-switch. If so, one engages in a form of self-silencing or self-censorship: they culturally smother in response to (likely or actual) pernicious ignorance. They conceal, truncate, mask, or alter aspects of their social identity, presenting an edited version of themselves that accords with dominant cultural standards. On the other hand, if one decides *not* to code-switch, they resist dominant cultural norms and expectations within the social environment. But by not code-switching, one risks social exclusion, unemployment, forms of material deprivation, and even safety (as was discussed in section 3).

No matter the choice, one's interests are undermined. By code-switching, one is forced to be complicit in reinforcing dominant cultural norms, further entrenching broader patterns of inequality. But by not code-switching, one risks losing important material goods and opportunities for social advancement.

The pressure to code-switch highlights how members of marginalised groups must navigate unfamiliar and potentially hostile social environments to pursue upward mobility. It is also important to remember that code-switching can be *exhausting*, especially when it manifests as cultural smothering. It involves a great deal of psychologically taxing work in which one's comparatively privileged counterparts simply need not engage.

5. Tentative Lessons and Concluding Remarks

AI systems are increasingly ubiquitous and more complex; so too are the social issues to which they give rise. This includes not only more direct human-to-AI interactions via smartphones, chatbots, or digital voice assistants but also the ways these technologies are implicated "behind the scenes" through social media platforms, streaming services, predictive search algorithms, video games, apps for rideshare programs and online dating, email communications, banking and finance, e-commerce, and more. These advancements have undoubtedly led to many positive changes, making previously arduous tasks much more streamlined and efficient. But at the same time, these technologies can hinder social progress.

As is now well known, AIS encode biases. This is unsurprising. These systems are not getting their training data in a vacuum. Instead, they are sourcing it from places like the Internet, which reflects society's myriad biases and prejudices, and produces selection effects favouring certain privileged populations (e.g., English

speakers, those who have Internet access, etc.) over others. As these technologies become increasingly advanced and indispensable to everyday tasks, researchers must be cognizant of how social structures that maintain unjust social arrangements are reproduced, reflected, or potentially made more potent and harmful within these technologies and their implementation.¹⁵

These insights underscore the need to attend to broader structural inequalities and mechanisms of oppression that reach well beyond AI technologies. It suggests that the solution to these issues does not rest within some manipulation or reconfiguration of training data or some alternative algorithm. The problems we identify are not fundamentally problems resulting from AI technologies and, as a result, won't give way to technological solutions. As O'Neil (2016, 204) has argued, "Big Data processes codify the past. They do not invent the future." Importantly, more data and more accurate or efficient algorithms will not solve any of the problems that we have described. Even if an algorithm had a perfect model of the world, the structures that disadvantage certain groups over others would be reflected in the data. This is because the problems that arise when training and implementing AI technologies in a nonideal and unjust social world are due to the very structure of that world. As a result, the issue is not whether the data are incomplete or unreflective of the real world. Even (perhaps especially) if they are complete and perfectly reflective, these problems will persist.

Moreover, these technologies feed back into society, creating more biased data as input, thus giving rise to a negative feedback loop of the type described by O'Neil (2016). Accordingly, the solution cannot be merely technological because the problem is not merely technological. A responsible first step is to be aware of structural injustice in the world. The point of intervention is not in changing our data to be better—our data will be better when the world is better—but when inequalities persist, so too will inequalities in data emerge to reflect them. All technological progress does is instantiate these inequalities more efficiently.

We have highlighted one important way these technologies can potentially be sites for epistemic oppression: by imposing imperfect choice situations or double binds onto nondominant social groups. Just as in the case of human-to-human interactions, where patterns of silencing emerge in response to the pernicious ignorance of one's audience, so too can pernicious ignorance be encoded into AI technologies, further deepening patterns of silencing on a potentially grander scale and at a more rapid pace. If technologies involving NLP and human-to-AI-interactions

¹⁵ This discussion also raises the question of whether we should want AI technologies—for instance, voice assistants like Alexa and Siri—to be very human-like. One interesting development is the creation of Q, a genderless voice assistant. You can try it out yourself: www.genderlessvoice.com.

require speakers to conform to dominant expectations and norms—or, in other words, if such technologies encode pernicious ignorance—then nondominant groups will need to engage in cultural smothering to interact with these technologies.

Acknowledgements

We are thankful to the participants at the Feminism, Social Justice, and Artificial Intelligence Workshop (July 2021), as well as two anonymous reviewers, and the editor of this special issue, Carla Fehr, for incredibly helpful feedback, questions, and comments. Thanks also to the Schwartz Reisman Institute at the University of Toronto for funding this research, and to Mila - Québec Artificial Intelligence Institute for providing generous resources and funding.

References

- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. “Machine Bias: There’s Software Used across the Country to Predict Future Criminals. And It’s Biased against Blacks.” *Pro Publica*, May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
- Bandy, Jack, and Nicholas Vincent. 2021. “Addressing ‘Documentation Debt’ in Machine Learning Research: A Retrospective Datasheet for BookCorpus.” arXiv preprint, arXiv:2105.05241 [cs.CL], submitted on May 11, 2021. <https://doi.org/10.48550/arXiv.2105.05241>.
- Bargiela, Sarah, Robyn Steward, and William Mandy. 2016. “The Experiences of Late-Diagnosed Women with Autism Spectrum Conditions: An Investigation of the Female Autism Phenotype.” *Journal of Autism and Developmental Disorders* 46 (10): 3281–94. <https://doi.org/10.1007/s10803-016-2872-8>.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *FAccT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Association for Computing Machinery, 610–23. New York: ACM. <https://doi.org/10.1145/3442188.3445922>.
- Benson, Erica J. 2001. “The Neglected Early History of Codeswitching Research in the United States.” *Language & Communication* 21 (1): 23–36. [https://doi.org/10.1016/S0271-5309\(00\)00012-4](https://doi.org/10.1016/S0271-5309(00)00012-4).
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. “Language Models Are Few-Shot Learners.” arXiv preprint, arXiv:2005.14165v4 [cs.CL], submitted on 22 Jul 2020. <https://doi.org/10.48550/arXiv.2005.14165>.

- Bruner, Justin P. 2019. "Minority (Dis)advantage in Population Games." *Synthese* 196 (1): 413–27. <https://doi.org/10.1007/s11229-017-1487-8>.
- Bruner, Justin, and Cailin O'Connor. 2017. "Power, Bargaining, and Collaboration." In *Scientific Collaboration and Collective Knowledge: New Essays*, edited by Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg, 135–57. Oxford: Oxford University Press.
- Buckner, Cameron. 2019. "Deep Learning: A Philosophical Introduction." *Philosophy Compass* 14 (10): e12625. <https://doi.org/10.1111/phc3.12625>.
- Bughin, Jacques, Jeongmin Seong, James Manyika, Michael Chui, and Raoul Joshi. 2018. "Notes from the AI Frontier: Modeling the Impact of AI on the World Economy." Discussion paper. McKinsey Global Institute. <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>.
- Cage, Eilidh, Jessica Di Monaco, and Victoria Newell. 2018. "Experiences of Autism Acceptance and Mental Health in Autistic Adults." *Journal of Autism and Developmental Disorders* 48 (2): 473–84. <https://doi.org/10.1007/s10803-017-3342-7>.
- Cage, Eilidh, and Zoe Troxell-Whitman. 2019. "Understanding the Reasons, Contexts and Costs of Camouflaging for Autistic Adults." *Journal of Autism and Developmental Disorders* 49 (5): 1899–911. <https://doi.org/10.1007/s10803-018-03878-x>.
- Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-Like Biases." *Science* 356 (6334): 183–86.
- Cassidy, S. A., K. Gould, E. Townsend, M. Pelton, A. E. Robertson, and J. Rodgers. 2020. "Is Camouflaging Autistic Traits Associated with Suicidal Thoughts and Behaviours? Expanding the Interpersonal Psychological Theory of Suicide in an Undergraduate Student Sample." *Journal of Autism and Developmental Disorders* 50 (10): 3638–48. <https://doi.org/10.1007/s10803-019-04323-3>.
- Cave, Stephen, and Kanta Dihal. 2020. "The Whiteness of AI." *Philosophy and Technology* 33 (4): 685–703. <https://doi.org/10.1007/s13347-020-00415-6>.
- Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values*. New York: W. W. Norton.
- Crenshaw, Kimberle. 1991. "Mapping the Margins: Intersectionality, Identity, Politics, and Violence against Women of Color." *Stanford Law Review* 43 (6): 1241–99.
- Dembroff, Robin, and Catharine Saint-Croix. 2019. "'Yep, I'm Gay': Understanding Agential Identity." *Ergo: An Open Access Journal of Philosophy* 6 (20): 571–99. <https://doi.org/10.3998/ergo.12405314.0006.020>.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding."

- arXiv preprint, arXiv:1810.04805v2 [cs.CL], submitted on 24 May 2019. <https://doi.org/10.48550/arXiv.1810.04805>.
- Donvan, John, and Caren Zucker. 2016. *In a Different Key: The Story of Autism*. New York: Crown.
- Dotson, Kristie. 2011. "Tracking Epistemic Violence, Tracking Practices of Silencing." *Hypatia* 26 (2): 236–57.
- . 2014. "Conceptualizing Epistemic Oppression." *Social Epistemology* 28 (2): 115–38.
- Dworzynski, Katharina, Angelica Ronald, Patrick Bolton, and Francesca Happé. 2012. "How Different Are Girls and Boys above and below the Diagnostic Threshold for Autism Spectrum Disorders" *Journal of the American Academy of Child & Adolescent Psychiatry* 51 (8): 788–97.
- Falbo, Arianna. 2021. "Fitting In Is Human: Forcing Someone to Fit In Is Oppression." *Psyche*, September 27, 2021. <https://psyche.co/ideas/fitting-in-is-human-forcing-someone-to-fit-in-is-oppression>.
- . 2022. "Hermeneutical Injustice: Distortion and Conceptual Aptness." *Hypatia* 37, no. 2 (Spring): 343–63. <https://doi.org/10.1017/hyp.2022.4>.
- Falbo, Arianna, and Heather Stewart, eds. 2021. "Outsiders Within: Reflections on Being a First-Generation and/or Low-Income Philosopher." Special issue, *APA Newsletter on Feminism and Philosophy* 20 (3).
- Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Friedman, Batya, and Helen Nissenbaum. 1996. "Bias in Computer Systems." *ACM Transactions on Information Systems* 14 (3): 330–47. <https://doi.org/10.1145/230538.230561>.
- Frye, Marilyn. 1983. *The Politics of Reality: Essays in Feminist Theory*. Trumansburg, NY: Crossing Press.
- Gabriel, Iason. 2020. "Artificial Intelligence, Values, and Alignment." *Minds and Machines* 30 (3): 411–37. <https://doi.org/10.1007/s11023-020-09539-2>.
- Gould, Judith, and Jacqui Ashton-Smith. 2011. "Missed Diagnosis or Mis-diagnosis? Girls and Women on the Autism Spectrum." *Good Autism Practice* 12 (1): 34–41.
- Green, Ben. 2019. "'Good' Isn't Good Enough." Paper presented at NeurIPS 2019 Joint Workshop on AI for Social Good, Vancouver, BC.
- Gumperz, John J. 1977. "The Sociolinguistic Significance of Conversational Code-Switching." *RELC Journal* 8 (2): 1–34.
- Harwell, Drew. 2018. The Accent Gap. *Washington Post*, July 19, 2018. <https://www.washingtonpost.com/graphics/2018/business/alexa-does-not-understand-your-accent/>.

- Hirji, Sukaina. 2021. "Oppressive Double Binds." *Ethics* 131 (4): 643–69. <https://doi.org/10.1086/713943>.
- Hull, Laura, Meng-Chuan Lai, Simon Baron-Cohen, Carrie Allison, Paula Smith, K. V. Petrides, and William Mandy. 2019. "Gender Differences in Self-Reported Camouflaging in Autistic and Non-autistic Adults." *Autism* 24 (2): 352–63. <https://doi.org/10.1177/1362361319864804>.
- Hull, Laura, K. V. Petrides, Carrie Allison, Paula Smith, Simon Baron-Cohen, Meng-Chuan Lai, and William Mandy. 2017. "'Putting on My Best Normal': Social Camouflaging in Adults with Autism Spectrum Conditions." *Journal of Autism and Developmental Disorders* 47 (8): 2519–34. <https://doi.org/10.1007/s10803-017-3166-5>.
- Johnson, Gabrielle M. 2021. "Algorithmic Bias: On the Implicit Biases of Social Technology." *Synthese* 198 (10): 9941–61. <https://doi.org/10.1007/s11229-020-02696-y>.
- Johnson, Joseph. 2021. "Internet Penetration Rate Worldwide 2021, by Region." Statista. <https://www.statista.com/statistics/269329/penetration-rate-of-the-internet-by-region/>.
- Jose, Navya, Bharathi Raja Chakravarthi, Shardul Suryawanshi, Elizabeth Sherly, and John P. McCrae. 2020. "A Survey of Current Datasets for Code-Switching Research." In *2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS)*, 136–41. Piscataway, NJ: Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/ICACCS48705.2020.9074205>.
- Koval, Christy Zhou, and Ashleigh Shelby Rosette. 2021. "The Natural Hair Bias in Job Recruitment." *Social Psychological and Personality Science* 12 (5): 741–50. <https://doi.org/10.1177/1948550620937937>.
- Jenkins, Katharine. 2017. "Rape Myths and Domestic Abuse Myths as Hermeneutical Injustices." *Journal of Applied Philosophy* 34, no. 2 (February): 191–205. <https://doi.org/10.1111/japp.12174>.
- LaCroix, Travis, and Alexandra Sasha Luccioni. 2022. "Metaethical Perspectives on 'Benchmarking' AI Ethics." arXiv preprint, arXiv:2204.05151 [cs.CY], submitted on 11 April 2022. <https://doi.org/10.48550/arXiv.2204.05151>.
- LaCroix, Travis, and Cailin O'Connor. Forthcoming. "Power by Association." *Ergo: An Open Access Journal of Philosophy*. <http://philsci-archive.pitt.edu/id/eprint/14318>.
- Lai, Meng-Chuan, Michael V. Lombardo, Bonnie Auyeung, Bhismadev Chakrabarti, and Simon Baron-Cohen. 2015. "Sex/Gender Differences and Autism: Setting the Scene for Future Research." *Journal of the American Academy of Child & Adolescent Psychiatry* 54 (1): 11–24. <https://doi.org/10.1016/j.jaac.2014.10.003>.

- Legg, Shane, and Marcus Hutter. 2007. "Universal Intelligence: A Definition of Machine Intelligence." *Minds and Machines* 17 (4): 391–444. <https://doi.org/10.1007/s11023-007-9079-x>.
- Livingston, Jennie, dir. 1990. *Paris Is Burning*. Off White Productions. 78 min. DVD: Lionsgate, 2012.
- Luccioni, Alexandra, and Yoshua Bengio. 2019. "On the Morality of Artificial Intelligence." arXiv preprint, arXiv:1912.11945 [cs.CY], submitted December 26, 2019. <https://doi.org/10.48550/arXiv.1912.11945>.
- Luccioni, Alexandra Sasha, and Joseph D. Viviano. 2021. "What's in the Box? An Analysis of Undesirable Content in the Common Crawl Corpus." arXiv preprint, arXiv:2105.02732v3 [cs.CL], submitted on 31 May 2022. <https://doi.org/10.48550/arXiv.2105.02732>.
- Them*. 2018. "InQueery: Trixie Mattel Breaks Down the History of 'Drag.'" *Them*, September 20, 2018. <https://www.them.us/story/inqueery-drag>.
- McCluney, Courtney L., Kathrina Robotham, Serenity Lee, Richard Smith, and Myles Durkee. 2019. "The Costs of Code-Switching." *Harvard Business Review*, November 15, 2019. <https://hbr.org/2019/11/the-costs-of-codeswitching>.
- McKinnon, Sean. 2017. "'Building a Thick Skin for Each Other': The Use of 'Reading' as an Interactional Practice of Mock Impoliteness in Drag Queen Backstage Talk." *Journal of Language and Sexuality* 6 (1): 90–127. <https://doi.org/10.1075/jls.6.1.04mck>.
- Medina, José. 2013. *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. New York: Oxford University Press.
- Morton, Jennifer M. 2014. Cultural Code-Switching: Straddling the Achievement Gap. *Journal of Political Philosophy* 22 (3): 259–81. <https://doi.org/10.1111/jopp.12019>.
- . 2019. *Moving Up without Losing Your Way: The Ethical Costs of Upward Mobility*. Princeton, NJ: Princeton University Press.
- . 2021. "Flourishing in the Academy: Complicity and Compromise." In "Outsiders Within: Reflections on Being a First-Generation and/or Low-Income Philosopher," edited by Arianna Falbo and Heather Stewart, special issue, *APA Newsletter on Feminism and Philosophy* 20 (3): 7–11.
- Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." *Science* 366 (6464): 447–53. <https://doi.org/10.1126/science.aax2342>.
- O'Connor, Cailin. 2017. "The Cultural Red King Effect." *Journal of Mathematical Sociology* 41 (3): 155–71. <https://doi.org/10.1080/0022250X.2017.1335723>.

- O'Connor, Cailin, and Justin Bruner. 2019. "Dynamics and Diversity in Epistemic Communities." *Erkenntnis* 84: 101–19. <https://doi.org/10.1007/s10670-017-9950-y>.
- Oliva, Thiago Dias, Dennys Marcelo Antonialli, and Alessandra Gomes. 2021. "Fighting Hate Speech, Silencing Drag Queens? Artificial Intelligence in Content Moderation and Risks to LGBTQ Voices Online." *Sexuality and Culture* 25 (2): 700–732. <https://doi.org/10.1007/s12119-020-09790-w>.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- Pfaff, Carol W. 1979. "Constraints on Language Mixing: Intrasentential Code-Switching and Borrowing in Spanish/English." *Language* 55 (2): 291–318.
- Pohlhaus, Gaile, Jr. 2012. "Relational Knowing and Epistemic Injustice: Toward a Theory of Willful Hermeneutical Ignorance." *Hypatia* 27 (4): 715–35. <https://doi.org/10.1111/j.1527-2001.2011.01222.x>.
- Popa-Wyatt, Mihaela, and Jeremy L. Wyatt. 2018. "Slurs, Roles and Power." *Philosophical Studies* 175 (11): 2879–906. <https://doi.org/10.1007/s11098-017-0986-2>.
- Poplack, Shana. 1980. "Sometimes I'll Start a Sentence in Spanish y termino en Español: Toward a Typology of Code-Switching." *Linguistics* 18 (7–8): 581–618.
- Raymaker, Dora M., Alan R. Teo, Nicole A. Steckler, Brandy Lentz, Mirah Scharer, Austin Delos Santos, Steven K. Kapp, Morrigan Hunter, Andee Joyce, and Christina Nicolaidis. 2020. "'Having All of Your Internal Resources Exhausted Beyond Measure and Being Left with No Clean-Up Crew': Defining Autistic Burnout." *Autism in Adulthood* 2 (2): 132–43. <https://doi.org/10.1089/aut.2019.0079>.
- Riley, Boots, dir. 2018. *Sorry to Bother You*. Significant Productions. 112 mins. DVD: 20th Century Fox.
- Russell, Stuart. 2019. *Human Compatible: Artificial Intelligence and the Problem of Control*. New York: Viking.
- Samotin, Perrie. 2017. "A Banana Republic Employee Says She Was Told Her Box Braids Looked Too 'Urban.'" *Glamour*, October 7, 2017. <https://www.glamour.com/story/banana-republic-employee-destiny-tompkins-says-she-was-told-box-braids-looked-too-urban>.
- Sap, Maarten, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. "The Risk of Racial Bias in Hate Speech Detection." In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 1668–78, Florence, Italy: Association for Computing Machinery. <https://doi.org/10.18653/v1/P19-1163>.
- Silberman, Steve. 2015. *NeuroTribes: The Legacy of Autism and the Future of Neurodiversity*. New York: Avery.

- Szczepański, Marcin. 2019. *Economic Impacts of Artificial Intelligence (AI)*. European Parliamentary Research Service, Briefing PE 637.967. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2019\)637967](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2019)637967).
- Tirrell, Lynne. 2012. "Genocidal Language Games." In *Speech and Harm: Controversies over Free Speech*, edited by Ishani Maitra and Mary Kate McGowan, 174–221. Oxford: Oxford University Press.
- Touré. 2018. "Sorry to Bother You: Is This the Most Shocking Anti-capitalist Film Ever?" *Guardian*, August 19, 2018. <https://www.theguardian.com/film/2018/aug/19/sorry-to-bother-you-is-this-the-most-anti-capitalist-film-ever>.
- Web Technology Surveys. 2022. "Usage Statistics of Content Languages for Websites." W3Techs. https://w3techs.com/technologies/overview/content_language.
- Yap, Audrey. 2017. "Credibility Excess and the Social Imaginary in Cases of Sexual Assault." *Feminist Philosophy Quarterly* 3 (4): article 1. <https://doi.org/10.5206/fpq/2017.4.1>.
- Young, Holly. 2015. "The Digital Language Divide: How Does the Language You Speak Shape Your Experience of the Internet?" <http://labs.theguardian.com/digital-language-divide/>.

ARIANNA FALBO is an assistant professor at Bentley University in Waltham, Massachusetts. Her research is primarily in epistemology, as well as feminist and social philosophy, with a focus upon the norms governing rational inquiry, as well as the relationship between inquiry and social injustice.

TRAVIS LACROIX: is an assistant professor at Dalhousie University in Halifax, Nova Scotia, where he teaches in the Department of Philosophy and the Faculty of Computer Science. His research centres on philosophical questions arising from sociodynamic systems. This research involves interrelated projects which address the social, philosophical, legal, and ethical implications of artificial intelligence, machine learning, and other data-driven technologies; the evolutionary origins of linguistic communication; and, the dynamics of other (nonlinguistic) social phenomena.