

ARTICLE

Models as Dogwhistles

Marina DiMarco

Department of History and Philosophy of Science, University of Pittsburgh, Pittsburgh, PA, USA
Email: marina.dimarco@pitt.edu

(Received 22 January 2022; revised 11 March 2023; accepted 06 August 2023; first published online 08 September 2023)

Abstract

Some scientific models and some claims about model–target relations are fruitfully diagnosed as dogwhistles. *Dogwhistles*, broadly speaking, are speech acts that send different, conflicting, and often differentially inflammatory messages to listeners. I distinguish two ways in which scientific models can be dogwhistles: representational dogwhistling and fit-for-purpose dogwhistling. I illustrate both kinds of dogwhistling using an example from computational social science, the diversity trumps ability theorem. I argue that dogwhistling threatens the objectivity of science, and I propose some ameliorative strategies.

1. Introduction

In this article, I argue that some scientific models and modeling claims are fruitfully diagnosed as a particular kind of speech act, namely, as dogwhistles.¹ *Dogwhistles*, broadly speaking, are communicative acts that send different, conflicting, and often differentially inflammatory messages to different listeners, as when Republican politicians in the United States use “inner city” to evoke racial stereotypes (Saul 2017, 2018). I argue that dogwhistling can arise in the context of scientific modeling in at least two ways. First, models may be dogwhistles when their targets are characterized in ways that allow differentially inflammatory interpretations among listeners, which I call *representational dogwhistling*. Second, claims about the adequacy of particular models for certain purposes may be dogwhistles when these purposes are characterized in ways that allow for conflicting interpretation, which I call *fit-for-purpose dogwhistling*. Importantly, both representational and fit-for-purpose dogwhistling can occur independently of any ill intentions of either speakers or listeners.

If I am right, this diagnosis has important consequences for our understanding of both dogwhistling and scientific communication. Because dogwhistling may

¹ In doing so, I join a cadre of philosophers of science who have begun to analyze the ways in which scientific communication might be fruitfully understood in terms of pragmatic speech acts and, relatedly, how we ought to evaluate scientific claims in the context of the ethics and politics of speech (see, e.g., Grasswick 2010; Franco 2017, 2019; John 2018, 2019; Khalifa and Millson 2020; Dang and Bright 2021).

effectively insulate a speaker from certain kinds of criticism, it may threaten the epistemic as well as the moral goods of science (Longino 1990). Furthermore, Saul (2018) and, to a lesser extent, Santana (2021) argue that diagnosing dogwhistles can play a critical role in allowing us to challenge and potentially neutralize harmful dogwhistles. If this is correct, scientists and philosophers of science might make good use of the dogwhistle diagnosis to ameliorate some harmful models and modeling claims. Furthermore, I propose that the possibility of models as dogwhistles also has important consequences for philosophers of language: it shows that unintentional dogwhistling is worthy of far more philosophical attention than it has received so far.

In section 2, I offer a brief primer on dogwhistles, drawing on recent work in linguistics and philosophy of language, especially Saul (2018). In section 3, I show that models and modeling claims are the sorts of things that could, in principle, be dogwhistles. In section 4, I propose representational and fit-for-purpose dogwhistling as two ways in which this might arise. In section 5, I illustrate both representational and fit-for-purpose dogwhistling using an example from computational social science: the diversity trumps ability theorem (Hong and Page 2004). I conclude by reflecting on possible implications of models as dogwhistles and of dogwhistling more broadly for the social epistemology of science, namely, the threats it may pose to the objectivity of science and the ameliorative strategies that might prevent dogwhistles from escaping criticism in the context of peer review.

2. What are dogwhistles?

In linguistics and philosophy of language, the notion of a dogwhistle has come to refer to a particular kind of speech that only some listeners can hear. For our purposes, it is useful to think of dogwhistles as speech acts that have different, often differentially inflammatory, effects on different listeners. The differences in these effects often correspond to listeners' affiliations with distinct social groups, their possession of some relevant background knowledge, or, as Quaranto (2022) suggests, broader "socio-politically structured linguistic practices" in their communities. For instance, the use of the phrase "inner city" by Republican politicians in the United States is now widely understood to be used in place of "Black" or "African American" in a way that invites a racist response from constituents who hear the dogwhistle, while maintaining deniability of such racism on the part of the speaker. So, when American congressman Paul Ryan remarked in 2014 that "we have got this tailspin of culture, in our inner cities in particular, of men not working and just generations of men not even thinking about working or learning the value and the culture of work," some audience members heard only the vague "inner cities," whereas others, who were sensitized to the dogwhistle (perhaps by racist or antiracist awareness of stereotype), heard Ryan allege that African Americans living in inner cities were lazy.²

To be more precise, following Saul (2018, 377–78), I take it that we ought to think of a dogwhistle as a *perlocutionary speech act*. A speech act, in the tradition indebted to J. L. Austin (1955), is a speaker's action in making an utterance. For instance, if I say "I would like a double espresso" to a barista, I do not merely describe my desire; rather,

² I owe this example to Henderson and McCready (2018, 231–32); the original remarks were made on Bill Bennett's *Morning in America* radio show (Corn 2014).

I make a request. A *perlocutionary act* (or a perlocutionary effect) is the reaction I actually produce in my listeners with my speech act, for instance, convincing my colleagues of a particular belief. Though speakers may intend certain perlocutionary effects, we can think of these as entirely dependent on listeners' reactions (Pagin 2014).³

A dogwhistle is best understood in terms of these perlocutionary effects, specifically, as a type of perlocutionary act that produces different effects in different listeners (Saul 2018). In a canonical instance, like the use of a code word, this might depend on listeners who are "in the know" making an inference that others cannot (Khoo 2017, 2021). In more complicated cases, everyone may "hear" a dogwhistle that violates a social norm, and listeners may respond differently because they differ in their commitment to the norm (Santana 2021). Santana's approach, though not defined in terms of perlocutionary effects, captures a key feature of dogwhistles that distinguishes them from other, more ordinary cases in which a speech act produces multiple effects or interpretations. This is because Santana points out that a critical function of a dogwhistle is to provide plausible deniability to a speaker for an interpretation that might constitute a *norm violation* when heard by the wrong audience (see also Fear 2007; Witten 2023; Quaranto 2022; cf. Saul 2018; Drainville and Saul, *forthcoming*). So, following Saul, I take a dogwhistle to be a speech act that produces at least two distinct perlocutionary effects, and following Santana, I expect that one of these effects amounts to a norm violation and that ambiguity about which effect was intended affords a speaker plausible deniability about violating the norm.

That being said, I follow Santana (2021) closely but not blindly. Santana argues that we should restrict the definition of dogwhistling to speech involving norm violations having to do with egalitarianism. I disagree. Although many dogwhistles are political, they need not be: for instance, Witten (2023) gives examples of coded messages for adult audiences within children's television programming, such as sexual innuendos that children do not "hear," and Drainville and Saul (*forthcoming*) argue that artistic practices like iconography can also constitute visual dogwhistles. As I'll argue, dogwhistles that promote multiple interpretations of a model are bad because they insulate scientists from criticism, and they do this even when the interpretations themselves are not anti-egalitarian.

Thinking of dogwhistles as perlocutionary acts is helpful for two reasons. First, Saul (2018) convincingly argues that this description can adequately account for much of what philosophers take to be interesting and important about dogwhistles: that they may be intentional or unintentional and that they may be overt (transparent to the listener) or covert (opaque to the listener, such that the listener is not conscious of their effects). What Saul characterizes as unintentional dogwhistles, or "unwitting use of words and/or images that, used intentionally, constitute an intentional dogwhistle, where this use has the same effect as an intentional dogwhistle" (368), may be just as pernicious as intentional dogwhistles. Because of this, it is important that characterizing dogwhistles in terms of perlocutionary effects also allows us to diagnose dogwhistling without knowledge of speaker intentions. Knowing other people's intentions is often epistemically risky business at best and intractable at worst. Furthermore, we want our account of dogwhistling to at least potentially accommodate a good deal of actual scientific communication, and the widespread practices of collaboration and coauthorship make

³ For a richer taxonomy of dogwhistles, see Taylor Koles's (2023) account.

intentions even more unpalatable for these purposes. Relatedly, this framework allows us to challenge a dogwhistle without accusing a speaker of ill intentions, which they can often simply deny. For instance, if our challenge to Paul Ryan's dogwhistle hangs on our allegation that he intended to say that African Americans are lazy, he can deny (as he did; see Corn 2014) that he intended any such meaning. Although intentions likely matter for other important projects, I want to diagnose dogwhistles independently of speaker intentions; otherwise, it may be impossible in practice.

Though perhaps unintuitive, this independence from intentions is especially important in developing a diagnosis for dogwhistles in scientific practice. This is because listener uptake of public scientific communication is often subject (at least in theory) to a norm of charitability of interpretation.⁴ That is, broadly speaking, listener uptake and criticism of scientific communication are or ought to be governed by an attitude of good faith, which might be directed toward either what meaning we think the speaker intends to convey or what we think the speaker intends to do with their speech act (i.e., what perlocutionary effects they intend). Ideally, we would like our diagnosis of scientific dogwhistling to come apart from the inflammatory listener responses that might result from uncharitable or bad faith interpretation.⁵ To accommodate this, I need my account of dogwhistles in science to be a bit stronger still: I will say that in addition to achieving deniability for a norm violation, a dogwhistle is a perlocutionary speech act that has different and differentially inflammatory effects among different *good faith listeners*. This adaptation of the perlocutionary account still allows us to cover a lot of ground. If good faith is directed toward the meaning of the speech itself, dogwhistles might allow for conflicting good faith interpretations of the same speech act when different listeners have different background knowledge (I'll elaborate on this possibility in section 4). If this good faith is directed toward the effects the speaker might intend to bring about, then we can remain charitable while hearing a dogwhistle because dogwhistling can occur whether the speaker intends it to or not. So, the perlocutionary framework can account for the rather obvious fact that scientific speech can have important consequences the speaker didn't intend, even when speaking to a good faith charitable audience.

3. Models can be speech acts

With this account of dogwhistling more or less in hand, I'll now argue that both claims about models and models themselves can be dogwhistles. In so doing, I do not by any means purport to exhaust the possibilities for dogwhistling in science; rather, it seems to me that models might not strike us as obvious candidates for a dogwhistling diagnosis, relative to other sorts of scientific speech acts. To the extent that this is the case, it seems to me that they make for a more interesting and, potentially, surprising philosophical target.

Many scientific practices go by the name of modeling, and philosophers recognize a rich taxonomy of models (e.g., minimal models [Batterman and Rice 2014; Chirimuuta 2014], toy models [Reutlinger, Hangleiter, and Hartmann 2018], physical models [Contessa 2010]). Despite this heterogeneity, philosophers of science generally

⁴ Thanks to Kevin Zollman for suggesting this and the point about intentions in collaborations.

⁵ Of course, there is always the possibility that this charitability is undeserved or unwarranted in a particular context. It's a defeasible norm.

(though not unanimously) agree that many models have in common a few features that are particularly important for our analysis. First, models depart in some way from complete description or reproduction of their targets, which is some aspect of the world (and where the ontological statuses of both “model” and “target” are compatible with a diverse array of metaphysical commitments or lack thereof) (Morgan and Morrison 1999; Contessa 2010; Giere 2010).⁶ Second, models are good for something, and some models or combinations of models are better than others for a given purpose (Giere 2004, 2010; Weisberg 2012, 2013; Parker 2009, 2015, 2020). This notion is perhaps best developed in Parker’s (2009) well-known adequacy-for-purpose account, where she writes that “a model is a tool in that it is used in the service of particular goals or purposes; typically these purposes involve answering some limited range of questions about the target system” (235).

To show that models and modeling claims might function as dogwhistles in the sense described, I’ll first need to show that at least some models and claims about models can be appropriately understood as the sorts of things that could, in principle, be dogwhistles. Next, I’ll sketch two ways in which we can subject models to a speech act-theoretic framework. We can understand models and modeling claims as two kinds of speech acts: assertion and representation. If this seems obvious, I invite you to skip ahead to section 4.

3.1 Assertion

First, claims about models can be understood as *assertions*. *Assertion* is a term of art in the speech act theory literature, where it refers to a constative speech act, or an action aimed at saying something about matters of fact (Pagin 2014). For instance, I might assert that my dog does not enjoy baths. Franco (2017, 2019) convincingly argues that we can understand many scientific utterances as assertions and further argues that the commitments of assertion play an important role in the social epistemology of science, including publication.⁷ Following Franco, I take it that among the things that scientists assert (and for our purposes, we can concern ourselves only with what they assert publicly, i.e., in print) are claims about models. Because assertions can produce multiple, differentially inflammatory perlocutionary effects, and this ambiguity can afford a speaker plausible deniability about a norm violation, claims about models, as assertions, are the sorts of things that we can properly understand as dogwhistle candidates.

3.2 Representation

Second, modeling itself can be a speech act. Although I suspect that we could gerrymander assertion to accommodate modeling, I will argue that scientific modeling is a particular class of speech act with its own specific commitments: *representation*. How might this work? First, a model itself can count as an utterance. According to Pagin (2014), for Austin and many of his interpreters, an utterance can be so permissive as to include any “physical item, event or state of affairs by means of

⁶ For an interesting exception to this trend, see Watkins (2023) on paleoclimate analogues as full-scale, full-complexity climate models.

⁷ Similarly, Contessa (2010, 215) explicitly notes that scientists make assertions about their models, though Contessa does not claim assertion in a speech-act-theoretic sense of the term per se.

which a speaker communicates.”⁸ For instance, Dixon (2022) argues that art can be understood in a speech act theoretic framework; see also Novitz (1977) and Drainville and Saul (forthcoming). Second, in putting forth a model, scientists are (at least sometimes) doing so with the aim of representing. Giere’s philosophy of modeling is helpful here. Giere (2004) argues that as a matter of scientific practice, models are used to represent and that *representing* is a fundamental activity of science. On his (2010) account, he cashes this out as follows: “Agents (1) intend; (2) to use model, *M*; (3) to represent a part of the world, *W*; (4) for some purpose, *P*” (Giere 2010, 269). Giere rather handily ties this account of modeling to linguistic practice for us.⁹

The point of accounting for modeling as representation is not to proliferate distinctions but rather to suggest that representation, as its own kind of speech act, might come with a specific set of commitments.¹⁰ I propose that we can think of representing as committing the speaker(s) (or modelers) to taking responsibility for defending or providing justification for the model in terms of its adequacy for some particular purpose à la Parker (2020). I suspect that this is an important departure from assertion insofar as it more explicitly allows us to proceed without expecting models to be true, approximately true, or even the sort of thing that could, in principle, be true or false (see, e.g., Elgin 2007; Contessa 2010; Frigg and Nguyen 2016). Because various sorts of departures from the whole truth are often taken to be a defining feature of models, it is important that representation not commit modelers to justifying models as true or false in a sense that we might want to leave open for assertion. Instead, following Parker (2009, 2020), I propose that representing invites listeners to proceed as if the model were adequate for the purposes described. Understood as an act of representation, a model (or, perhaps more accurately, modeling) is the sort of thing that we can properly understand as a dogwhistle candidate. As with assertion, this is not simply because the purposes of models may be vague but more specifically because when this ambiguity produces multiple interpretations, one purpose may be a norm violation (say, an illegitimate goal), while another may not; this kind of ambiguity gets a modeler off the hook for the norm violation.

4. Two kinds of dogwhistles

Now that I have established that both models and modeling claims are dogwhistle candidates, let me elaborate a couple of ways in which I think they may be particularly likely to actually become dogwhistles. Recall that, for our purposes, a dogwhistle is a perlocutionary speech act that has different effects among different good faith listeners, that it provides plausible deniability about norm-violating effects, and that these effects might very well be unintentional. I think that there are at least two means by which models can produce different (and differentially

⁸ Drainville and Saul (forthcoming) develop an account of visual dogwhistles that does not depend on the use of language; this might be another avenue for models.

⁹ We might worry about the role of intentions here, given our goal of agnosticism about the intentions of listeners and dogwhistlers. But merely having a role for intentions in an account of representation is perfectly compatible with agnosticism about what those intentions are and how they might relate to perlocutionary effects.

¹⁰ Commitments here in Brandom’s (1983) sense. Franco (2017) likewise discusses these sorts of commitments in relation to assertion. See Pagan (2014) for an overview of commitments in speech act theory.

inflammatory) effects in distinct groups of charitable listeners, both of which concern the adequacy-for-purpose of models in relation to their targets. The first, which I'll call *representational dogwhistling*, concerns the characterization of targets. The second, which I'll call *fit-for-purpose dogwhistling*, concerns the characterization of the purposes for which the model is supposed to be adequate. Characterizing representational targets and purposes are obviously closely related in scientific practice, and my concern is not to impose an arbitrary distinction between the two; rather, I aim to suggest that they are worth considering separately for philosophical purposes, even if they are often two sides of the same coin.

4.1 *Representational dogwhistling*

Recall that, generally speaking, models differ in some way from whatever it is they might be thought to relate to and, furthermore, that this model-target relation is often (if not always) something other than complete description of this target. Whatever we might take a target to be (e.g., “part of the world” [Giere 2010], “possible concrete objects” [Contessa 2010]), an important aspect of representation is specifying what the model is, or could be, a model of. In Giere’s parlance, the representation or claim that M is a model of W that is good for some purpose P entails that M is a model of W , and representing involves characterizing what part of the world W in which we are interested. Models and modeling claims can produce different and conflicting understandings of a model’s target in different listeners. Furthermore, depending on the general nature of the target, these interpretations may be differentially inflammatory or norm violating. I’ll call these *representational dogwhistles*.

Perhaps the most intuitive sense in which this might be the case is when modelers characterize their targets in terms that are particularly vague. This would seem to invite different and perhaps conflicting interpretations from listeners with different background knowledge. For instance, suppose an economist put forth a model for which they characterized their target as “employment in inner-city populations in the United States” (which, as far as I know, nobody has actually done). Those of us who are familiar with the historical use of *inner city* to mean “African American” might interpret the model as being a model of African American employment, and, depending on the context, this might produce an inflammatory response, perhaps by inciting our racism or (we can hope) by arousing our moral indignation at a perceived instance of racism. By contrast, those who are not aware of this more specific meaning might interpret the model in the more general, innocuous sense.

However, this might be an unsatisfying example because it seems parasitic on a preexisting dogwhistle, *inner city*. So, consider a different case. A modeler wants to represent some target W that they do not take to exist, say, a possible future scenario, entertained for the sake of argument. Call this target *possible-W*. If the modeler presents M as a model of W without specifying that they take W to be a possible, not an actual, case, then some listeners may take M to represent W as a feature of the world, whereas others may take M to represent *possible-W*. At least in some discursive contexts (namely, those that are given to frequently modeling both W s and *possible-W*s), I submit that these conflicting interpretations may both be charitable ones. If this is right, then in the case that positing *possible-W* as an actual feature of the world is inflammatory (say, the scenario is one in which *possible-W* is the condition that

women are essentially worse at analytic philosophy), we can expect these interpretations to be differentially inflammatory among different listeners. Furthermore, although this possible–actual distinction concerns the target, it could also concern the sort of explanation a model might provide (see Bokulich 2014). Note that this needn't be an intention of the modeler: it may just be some intellectual sloppiness on their part, or it may be that they intend their model for an audience who would have the additional background knowledge that *Ws* are never to be taken as actual in this particular context, but some nonexpert listeners sneak in or eavesdrop. It seems that our nonexperts are not properly thought of as uncharitable in this context but rather as less well informed. If this is right, then we could have a nonparasitic case of unintentional representational dogwhistling on our hands.¹¹

4.2 *Fit-for-purpose dogwhistling*

The second sense in which models can function as dogwhistles depends on our characterization of models as adequate for some purpose and, more specifically, on what we take this purpose to be. If we stick with Giere, model-purpose characterization is an important component of representation, as representing involves not just the claim that *M* is a model of *W* but the claim that it is good for *P*. Here we are interested in characterizing *P*.

Like target characterizations, purpose characterizations may produce different and conflicting understandings of a model's purpose in different listeners. Depending on the nature of these interpretations, they may be differentially inflammatory. I'll call these *fit-for-purpose dogwhistles*. Like representational dogwhistles, fit-for-purpose dogwhistles can be parasitic on preexisting dogwhistles, but they need not be so. Fit-for-purpose dogwhistles might also arise in cases in which the purpose of a model is sufficiently unclear so as to allow conflicting listener understandings of what that purpose might be. This might include cases in which the purpose of a model is not explicitly specified at all. This is because, if I am right about representing, modeling entails a commitment that the model is good for *some* purpose, even if that purpose is not specified, and it therefore authorizes listeners to infer that the modeler has some purpose in mind.

To see how model-purpose characterizations might produce differentially inflammatory interpretations of a model's purpose, consider a model for which the modelers take its purpose to be figuring as a premise in a conditional argument with other inflammatory or norm-violating premises. If a listener believes an inflammatory premise *I*, then they ought to be convinced by the model *M* to behave in a particular way. The modelers take the purpose of *M* to be convincing a subset of listeners who believe *I*, and the modelers themselves are neutral or skeptical with regard to the goodness of *I*. If the conditional status of the purpose of *M* is unclear, a listener may instead infer that the modelers actually believe *I* and therefore that the purpose of *M* is to convince everyone, taking *I* for granted. A listener who infers that the modelers take for granted that the listener believes *I* may (or may not) be offended by the perceived taking-for-granted of an inflammatory premise.

¹¹ Again, note that it could be unintentional. I suspect that modelers are at least sometimes intentionally vague about their targets to preserve this sort of uncertainty and to invite multiple interpretations. As O'Connor (2017) points out, there are also contexts in which the possible versus the actual status of the model doesn't really matter, given its role in a particular argument.

We can intuitively extend this possibility to sufficiently vague claims about the further purposes to which a model might be put, especially when it may be unclear in virtue of what the model is fit for some further purpose. For instance, if a paper is framed in terms of a purpose of general social interest, such as intervening on the possible causal relationship between socioeconomic status (SES) and obesity, and the paper offers a model of a much more specific phenomenon, such as the correlation between SES and obesity for a number of Pittsburgh zip codes, for the purpose of understanding the distribution of both SES and obesity, it may be unclear in virtue of what the more specific model is adequate for the more general purpose.¹² Extending a model to additional purposes might involve asserting or assuming that two purposes (or two targets) are similar or analogous in some relevant way. Because claims about the similarity or difference of potential targets may be inflammatory, we might reasonably expect extrapolation from one purpose to another to be another potential source of fit-for-purpose dogwhistling.

These two mechanisms almost certainly do not exhaust the possibilities of model-related dogwhistling in science, but they seem as good a place to start as any other. In section 5, I'll illustrate how representational and fit-for-purpose dogwhistling might arise in a case of mathematical modeling: Hong and Page's (2004) diversity trumps ability theorem.

5. The case of the diversity trumps ability theorem

In their 2004 article "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers," Lu Hong and Scott Page present what they call "a general framework for modeling functionally diverse problem-solving agents" (16385). Using a mathematical model and computer simulations, Hong and Page claim to show that a group of diverse agents performs better than a group of highest-performing agents on a hard problem. This is commonly known as the *diversity trumps ability theorem* (Anderson 2006). We can use the diversity trumps ability theorem to illustrate the possibility of both representational and fit-for-purpose dogwhistling in science. Although this result, and, in particular, the mathematical theorem behind it have been the target of both celebration and criticism (e.g., Thompson 2014; cf. Kuehn 2017), my concern is to demonstrate how both the target of their model and the purpose for which the model is supposed to be adequate may invite different and differentially inflammatory interpretations. Whether the model is actually a good representation of any target for any particular purpose is beside the point (for evaluation, see Grim et al. 2019; Reijula and Kuorikoski 2021; Kassam, forthcoming).¹³

5.1 Diversity, ability, and representational dogwhistling

Let's start with the target or targets of the model. At least one target of the model in Hong and Page (2004) is described in the title as "diverse problem solvers." The

¹² For one account of the ethics and epistemology of framing in scientific communication, see McKaughan and Elliott (2013).

¹³ I offer this case as a proof of possibility, but I suspect that dogwhistling also occurs elsewhere in science, especially given debates about the interpretation and purposes of models of complex, socially salient phenomena in other domains (e.g., sociogenomics).

characterization of “diverse problem solvers” has resulted in different and conflicting interpretations of both “diverse” and “solvers.” Specifically, it has been variously interpreted to refer to functionally diverse algorithms (where *functionally diverse* refers to differences in search strategies for a specific sequential search problem), functionally diverse people, and demographically diverse people (at least to the extent that demographic or “identity” diversity causes relevant cognitive diversity; see Page 2007a). For instance, Anderson (2006, 12) characterizes the paper as “initially constructed to model problem solving within firms.” By contrast, mathematician Abigail Thompson (2014, 1026) writes, “A ‘problem-solver’ strongly suggests an individual person. However as Hong and Page are using the word, a ‘problem-solver’ is an algorithm.” Thompson thinks the move from algorithms to people in this paper is a norm violation: “an example of the misuse of mathematics” (1024), inappropriately applying the idealized model to the context of diversity initiatives in universities. “What the paper emphatically does not contain is information that can be applied to any real-world situation involving actual people,” she writes (1029). In response, Kuehn (2017) argues that Thompson’s (2014) interpretation of the model is incorrect because she does not understand the norms of mathematical modeling in social science. Specifically, Kuehn claims that models in social science “typically illustrate a non-exhaustive set of mechanisms at work in society that have too many moving parts to be fruitfully expressed or convincingly demonstrated using prose,” as opposed to models in the physical sciences, which “reflect actual physical phenomena quite closely” (73). Some (if not all) of these different and conflicting interpretations seem to be charitable: Anderson (2006) cites Hong and Page to extoll the value of their argument for the epistemology of democracy; Kuehn (2017) argues for this interpretation to defend the paper against Thompson, and although Thompson has been accused of being uncharitable in her critique (which also concerns more technical features of the 2004 paper), interpretations of the target as either persons or algorithms are consistent with the language of the paper.

This discussion reflects two possible ambiguities about the model’s target. First, it is unclear whether demographic diversity specifically counts among the legitimate targets of the model. Consider that Hong and Page (2004) frame their paper as an answer to a question about people: “Can a functionally diverse group whose members have less ability outperform a group of people with high ability who may themselves be diverse? The main result of our paper addresses exactly this question” (16385). They characterize their result in terms of agents—“we find that a random collection of agents drawn from a large set of limited-ability agents typically outperforms a collection of the very best agents from that same set” (16386)—and operationalize agents in their mathematical model and computer simulation just in terms of their strategies for a particular difficult search task. While Hong and Page claim to exclusively model functional cognitive diversity and merely *frame* this in terms of discussions of identity diversity and affirmative action cases, Page (2007a, 2007b) further elaborates conditions under which identity diversity might yield functional diversity. In his book *The Difference*, Page (2007a) explicitly discusses admissions as an application of his argument, which draws on the 2004 paper as well as other models. “The logic in this book can be used to support an instrumental argument for affirmative action The extent to which this logic supports race-based affirmative

action depends on empirical facts: either identity diversity correlates with cognitive diversity or it does not,” he writes (49).

Second, it is ambiguous whether the model should be understood as figuring in a possible or actual explanation of its target(s). Page (2007a, 394) presents empirical data with the caveat that “the models that we covered do not provide the only candidate explanations for the empirical phenomena we discuss. Other models might better explain some or all of these empirical regularities.” This aligns with Page’s (2015) argument that the model provides sufficient conditions for a phenomenon, which would suggest that it provides a how-possibly explanation (by analogy to, e.g., Wu 2023). Page (2015, 10) responds to Thompson along these lines:

The note accuses me of misusing mathematics, claiming that I imply that the mathematical results are somehow fact in the world of people. The accusation is baseless. In my book, I caution readers to apply mathematical models carefully, highlighting the subtleties of moving from the starkness of mathematical logic to the richness of human interactions. Not everyone understands the role of mathematical claims in the social sciences.

Nonetheless, the primary motivation for the book seems to be the actual salience of these models to cognitive diversity in admissions, firms, and teams, and Page devotes careful (even critical) attention to empirical support for their relevance in these contexts. Given that the book was also written for a broad audience, who seem no more likely than mathematicians to be familiar with the norms of modeling in social science, it is plausible that charitable readers could form different interpretations as to whether the model could figure in an actual explanation of anything about its target(s). Kuehn and Page both interpret Thompson’s complaint as stemming from her understanding of the role of the mathematical model; Page uses this to deny the misuse of mathematics. If they are right, and her charge depends on either of these conflicting interpretations of the model’s target or its relation to it, the model is a representational dogwhistle.¹⁴

5.2 Diversity, ability, and fit-for-purpose dogwhistling

Let’s turn to the purpose of the model. In section 4, I argued that one way fit-for-purpose dogwhistling can occur is when the modelers take a model’s purpose to be figuring as a premise in a conditional argument with another inflammatory premise, *I*. If the conditional status of the purpose of the model *M* is unclear, a listener may

¹⁴ Some listeners might also notice a technical feature of the model in Hong and Page’s (2004, 13689) presentation, recently highlighted by Kassam (forthcoming), namely, a “tension” between the diversity of groups of problem solvers and the problem solvers’ individual abilities in the mathematical model, in which, in the limit as the pool of problem solvers becomes increasingly large, a group of the highest-ability problem solvers will just be identical copies of the same highest-performing search strategy. Given that the paper is framed in terms of affirmative action case law, those who are aware of Justice Scalia’s suggestion that there might be a trade-off between the diversity and the ability of the University of Michigan Law School and of applicants to the University of Texas (see Liptak 2015) may find this inflammatory. However, Hong and Page (2004) explicitly disclaim that they model only functional diversity, and Page (2007b, 2017) has explicitly criticized “trade-off” framings of affirmative action, so it is less obvious that this meets the charity condition for a dogwhistle.

instead infer that the modelers actually believe *I* and therefore that the purpose of *M* is to convince everyone, taking *I* for granted. If some listeners are offended by the perceived taking-for-granted of an inflammatory premise, this can produce a fit-for-purpose dogwhistle. Now I'll argue that we can use the diversity trumps ability paper as an example of this sort.

Many philosophers and social scientists have interpreted the purpose of the model in Hong and Page (2004) to be to provide justification for preferring epistemic and sometimes demographic diversity in epistemic communities. (For a concise review, see Grim et al. 2019, 98–99.) For instance, Thompson (2014, 1024) characterizes the paper as an “[attempt] to find a mathematical justification for ‘diversity’ as practiced in universities”; Anderson (2006, 12) writes that it “supports the claim that democracy, which allows everyone to have a hand in collective problem solving, is epistemically superior to technocracy, or rule by experts”; and Grim et al. (2019, 98) argue that the result “suggests that an organization is epistemically better off by recruiting a diverse set of candidates instead of just selecting the best individual performers.”¹⁵ We might worry that these examples reflect these authors’ making use of the model for their own purposes or extending it beyond what Hong and Page (2004) originally authorized. But, at a minimum, Hong and Page assert that their model has possible consequences for organizations and for educational policy. For instance, they write that the result of their model “has implications for organizational forms and management styles, especially for problem-solving firms and organizations” (16389). In addition to the firm example, they write that “even if we were to accept the claim that IQ tests, Scholastic Aptitude Test scores, and college grades predict individual problem-solving ability, they may not be as important in determining a person’s potential contribution as a problem solver as would be measures of how differently that person thinks” (16389). They conclude, “We should do more than just exploit our existing diversity. We may want to encourage even greater functional diversity, *given its advantages*” (16389, emphasis added).

Thus it is at least understandable that some listeners take Hong and Page to be arguing that their result justifies increasing diversity and that this justification rests on the epistemic advantage conferred by the diversity of the group of problem solvers. We can interpret this in at least two ways. On one interpretation, Hong and Page assert that this justification for increasing diversity plays a role in a conditional argument like the following:

- (1) *If you think exploitability is a good reason to increase the diversity of an epistemic community, then this model provides justification for doing so.*

On another interpretation, Hong and Page assert that this justification for increasing diversity plays a role in an argument like the following:

- (2) *Because exploitability is a good reason to increase the diversity of an epistemic community, this model provides justification for doing so.*

¹⁵ Grim and colleagues (2019) criticize Anderson’s (2006) extension of the model to expertise, but this seems to be primarily because they want a more demanding notion of expertise, not because they find her interpretation particularly inflammatory.

For listeners who are committed to the relative priority of various moral or exploitability-independent reasons for increasing the diversity (epistemic or otherwise) of epistemic communities, (2) may be inflammatory in a way that (1) is not. These conflicting interpretations of Hong and Page's characterization of the purpose of their model stand in very different relationships to the affirmative action cases against which the paper is framed; whereas (1) is agnostic about the appropriateness of exploitability (as opposed to, say, considerations of justice or the intrinsic value of diversity) as a reason to support affirmative action, (2) instrumentalizes diversity.

6. Dogwhistles threaten the objectivity of science

In sections 4 and 5, I hope to have shown that we can (at least sometimes) diagnose models as dogwhistles. Now I'll briefly argue that we should. One reason to diagnose dogwhistling seems obvious, especially because we began with the example of Ryan's racist dogwhistle: dogwhistling can cause serious harm, and it might be the sort of thing for which we'd like to assign responsibility, sanctions, or blame, or at least take care to avoid. Given the tremendous epistemic authority and practical import of scientific communication, we might be particularly worried about the harms of dogwhistling in a scientific context when the stakes are high; in this sense, possible dogwhistling might be a source of epistemic risk or of more broadly pragmatic, as well as ethical, considerations (Biddle and Kukla 2017). Inflammatory characterizations of particular targets (i.e., populations) or the perception of untrustworthy purposes may contribute to public distrust of science. Our desire to avoid the harms of dogwhistling is an important potential avenue for so-called nonepistemic social and political values in science, and perhaps also for our thinking about the epistemic and testimonial responsibilities of scientific experts (Daukas 2011; Almassi 2012; Franco 2017, 2019; John 2018, 2019; Kelsall 2020; Khalifa and Millson 2020; Dang and Bright 2021). Dogwhistling may jeopardize the trustworthiness or increase the epistemic risk of scientific communication, and it may purposefully or inadvertently perpetuate harmful speech. We can think of these as harms in a traditionally ethical sense of the word.

But diagnosing dogwhistles also has important consequences in a more restricted epistemic sense. This is because at least some dogwhistles, in some contexts, may be immune to criticism. As many feminist philosophers of science and social epistemologists have argued, criticism and dissent play important roles in the well-functioning social epistemology of science (e.g., Longino 1990; Solomon 2001; Intemann 2011; and [small world] Anderson 2006). Perhaps most famously, Longino (1990) argues that the possibility of transformative criticism and the social structures that guarantee the possibility of sincere uptake of scientific criticism are so important that a scientific community cannot be said to be objective in any meaningful sense without them. If some dogwhistles escape criticism by avoiding detection, or, as Santana (2021) convincingly argues, if they are sometimes designed to be immune to criticism even when they *are* detected, then they pose a serious threat to the epistemic as well as the moral goods of science.

One way dogwhistling might escape criticism is, as we have seen, by allowing scientists plausible deniability about both their targets and the intended purposes of

their work. This threatens the objectivity of science because it insulates dogwhistlers from criticism by allowing them to deny that norm-violating interpretations of their work were intended: if a critic points out that a model has, say, antiegalitarian representations or purposes, the modelers can point to the alternative interpretation to claim that they are being misunderstood. The intentions framework also helps dogwhistlers to make challenging dogwhistling a costly activity for would-be critics, as dogwhistlers can argue that critics are being uncharitable, inappropriately political, or dogmatically motivated in their assessment of speaker intentions when they “hear” a dogwhistle (Khoo 2021). Furthermore, successful dogwhistles also insulate dogwhistled content from scrutiny via what Howdle (2023) calls “information balkanization.” In other words, if successful dogwhistles are not “heard” by would-be critics, they not only evade criticism but may not even be subject to it.

A second way that dogwhistles can escape criticism is when they are what Saul (2018, 361) calls “covert dogwhistles.” On her account, whereas overt dogwhistles are those coded messages that the listener recognizes as such, covert dogwhistles are more slippery. In Saul’s historical example, a political advertisement that used the phrase “violent crime” while showing an image of a Black person functioned as a covert dogwhistle because, for some viewers, it appealed to racist stereotypes without doing so in a way that was consciously accessible to those viewers. Saul argues that undiagnosed covert dogwhistles escape certain sorts of criticism precisely because we do not see them at work. For Saul, a covert dogwhistle might “[allow] for people to be manipulated in ways that they would resist if the manipulation was carried out more openly—often drawing on racist attitudes that are consciously rejected” (361). This seems especially likely when listeners consciously hold, for example, egalitarian values but still maintain something like implicit bias. Importantly, for Saul, covert dogwhistles do not survive transparency: they do not have the same effect when they are made overt. When the more inflammatory potential of a speech act is made clear, dogwhistles can at least sometimes be challenged and even ameliorated. This suggests an important epistemological role for the dogwhistle diagnosis in science: it may enable certain kinds of resistance to, and criticism of, some models and claims about models that might otherwise persist in the shadowy realm of the covert. In this regard, the dogwhistle diagnosis may thus support the objectivity of scientific inquiry in the same way as so many other mechanisms that protect and encourage criticism.

In light of this, I’ll briefly propose three specific practical recommendations by which we might hope to avoid or ameliorate dogwhistling in science. First, it should be obvious by now that we ought to be on the lookout for dogwhistles in scientific communication. Specifically, because I’ve argued that dogwhistling may arise due to differences in listener background knowledge, we ought to be especially wary of unintentional dogwhistling in cases in which scientific communication might be expected to reach a broad audience who probably have such differences. For instance, Hong and Page’s 2004 paper appeared in the *Proceedings of the National Academy of Sciences of the United States of America*, a generalist journal with a broad readership that spans many academic disciplines, and Page’s (2007a) book further popularized the model. Authors, editors, and reviewers for venues where such a broad audience is not only likely but expected might take special care to notice and ameliorate potentially inflammatory ambiguities before such papers go to press. Similarly, scientific communication that is intended or expected to reach a nonexpert or public audience

might warrant special care. Second, journals and editors can allow longer articles, which give authors more space to clarify the scope of their models and to head off potentially inflammatory interpretations.¹⁶ Third, in cases in which it is particularly important to avoid dogwhistling, editors might select additional reviewers precisely in virtue of their demographic differences or differences in expertise, because such differences might make it more likely that a reviewer will “hear” a dogwhistle before a paper is published (Wylie 2012). In light of our discussion of the diversity trumps ability theorem, it strikes me as particularly interesting to note that, if one is in the business of justifying the need for diversity of particular epistemic communities, this might be another (instrumental) reason to do so.

7. Conclusion

In this article, I have argued that we can and sometimes ought to diagnose models as dogwhistles. I’ve characterized two possible forms of dogwhistling with an eye toward modeling in scientific practice: representational and fit-for-purpose dogwhistling. We saw how these might work in the example of a mathematical model of diverse problem solvers. I strongly doubt that this is the only case of dogwhistling in the modeling literature. Because dogwhistling poses both a moral and an epistemic threat to science, I have argued that we ought to be on the lookout for potential dogwhistles and suggested a couple of places we might look. Because this account of dogwhistles is compatible with unintentional dogwhistling and charitable listeners, it is more amenable to challenging dogwhistle speech than it otherwise might be.

The dogwhistle diagnosis ought to be a fruitful one for scientists and for philosophers of science—but not *only* for philosophers of science: by extending the notion of dogwhistles to scientific practice in this particular way, we have learned some things about dogwhistling itself. I have tried to show how much ground we can cover with the notion of an unintentional dogwhistle and with a presumption of charitability. If this is right, then unintentional dogwhistling seems to be much more varied, interesting, and worthy of philosophical attention than it has so far been considered. Though I stop far short of a complete account of dogwhistling in science, these possibilities lead me to hope that future work might continue to be productive for scientists and for philosophers of language as well as for philosophers of science.

Acknowledgments. Thanks to Kevin Zollman, Carlos Santana, Mike Dietrich, Dasha Pruss, Alnica Visser, Taylor Koles, Will Conner, Greg Kehne, and audiences at the PSA 2020/2021 poster sessions for their generous suggestions, ideas, and expertise. Thanks also to the reviewers, who were very kind, constructive, and helpful in playing along with this article.

References

- Almassi, Ben. 2012. “Climate Change, Epistemic Trust, and Expert Trustworthiness.” *Ethics and the Environment* 17 (2):29–49. <https://doi.org/10.2979/ethicsenviro.17.2.29>.
- Anderson, Elizabeth. 2006. “The Epistemology of Democracy.” *Episteme* 3 (1–2):8–22. <https://doi.org/10.3366/epi.2006.3.1-2.8>.
- Austin, J. L. 1955. *How to Do Things with Words*. Cambridge, MA: Harvard University Press.
- Batterman, Robert W., and Collin C. Rice. 2014. “Minimal Model Explanations.” *Philosophy of Science* 81 (3):349–76. <https://doi.org/10.1086/676677>.

¹⁶ Thank you to Kevin Zollman for suggesting this.

- Biddle, Justin B., and Rebecca Kukla. 2017. "The Geography of Epistemic Risk." In *Exploring Inductive Risk: Case Studies of Values in Science*, vol. 1, edited by Kevin C. Elliott and Ted Richards, 215–38. Oxford: Oxford University Press.
- Bokulich, Alisa. 2014. "How the Tiger Bush Got Its Stripes: 'How Possibly' vs. 'How Actually' Model Explanations." *The Monist* 97 (3):321–38. <https://doi.org/10.5840/monist201497321>.
- Brandom, Robert. 1983. "Asserting." *Notis* 17 (4): 637–50. <https://doi.org/10.2307/2215086>.
- Chirimuuta, M. 2014. "Minimal Models and Canonical Neural Computations: The Distinctness of Computational Explanation in Neuroscience." *Synthese* 191 (2):127–53. <https://doi.org/10.1007/s11229-013-0369-y>.
- Contessa, Gabriele. 2010. "Scientific Models and Fictional Objects." *Synthese* 172 (2):215–29. <https://doi.org/10.1007/s11229-009-9503-2>.
- Corn, David. 2014. "More Evidence of Paul Ryan's 'Inner Cities' Problem." *Mother Jones*, March 27. <https://www.motherjones.com/politics/2014/03/paul-ryan-inner-cities-problem/>.
- Dang, Haixin, and Liam Kofi Bright. 2021. "Scientific Conclusions Need Not Be Accurate, Justified, or Believed by Their Authors." *Synthese* 199 (3):8187–8203. <https://doi.org/10.1007/s11229-021-03158-9>.
- Daukas, Nancy. 2011. "Altogether Now: A Virtue-Theoretic Approach to Pluralism in Feminist Epistemology." In *Feminist Epistemology and Philosophy of Science*, edited by Heidi E. Grasswick, 45–67. Dordrecht, Netherlands: Springer. https://doi.org/10.1007/978-1-4020-6835-5_3.
- Dixon, Daisy. 2022. "Artistic (Counter) Speech." *The Journal of Aesthetics and Art Criticism* 80 (4):409–419.
- Drainville, Ray, and Jennifer M. Saul. Forthcoming. "Visual and Linguistic Dogwhistles." In *The Oxford Handbook of Applied Philosophy of Language*, edited by Luvel Anderson and Ernie Lepore. Oxford: Oxford University Press.
- Elgin, Catherine. 2007. *True Enough*. Cambridge, MA: MIT Press.
- Fear, Josh. 2007. "Under the Radar: Dog-Whistle Politics in Australia." Discussion Paper 96, Australia Institute, Manuka, ACT. <https://australiainstitute.org.au/report/under-the-radar-dog-whistle-politics-in-australia/>.
- Franco, Paul L. 2017. "Assertion, Nonepistemic Values, and Scientific Practice." *Philosophy of Science* 84 (1):160–80. <https://doi.org/10.1086/688939>.
- Franco, Paul L. 2019. "Speech Act Theory and the Multiple Aims of Science." *Philosophy of Science* 86 (5):1005–15. <https://doi.org/10.1086/705452>.
- Frigg, Roman, and James Nguyen. 2016. "The Fiction View of Models Reloaded." *The Monist* 99 (3):225–42. <https://doi.org/10.1093/monist/onw002>.
- Giere, Ronald N. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71 (5):742–52. <https://doi.org/10.1086/425063>.
- Giere, Ronald N. 2010. "An Agent-Based Conception of Models and Scientific Representation." *Synthese* 172 (2):269–81. <https://doi.org/10.1007/s11229-009-9506-z>.
- Grasswick, Heidi E. 2010. "Scientific and Lay Communities: Earning Epistemic Trust through Knowledge Sharing." *Synthese* 177 (3):387–409. <https://doi.org/10.1007/s11229-010-9789-0>.
- Grim, Patrick, Daniel J. Singer, Aaron Bramson, Bennett Holman, Sean McGeehan, and William J. Berger. 2019. "Diversity, Ability, and Expertise in Epistemic Communities." *Philosophy of Science* 86 (1):98–123. <https://doi.org/10.1086/701070>.
- Henderson, R., and Elin McCready. 2018. "How Dogwhistles Work." In *New Frontiers in Artificial Intelligence*, edited by Sachiyo Arai, Kazuhiro Kojima, Koji Mineshima, Daisuke Bekki, Ken Satoh, and Yuiko Ohta, 231–40. Lecture Notes in Computer Science. Cham, Switzerland: Springer.
- Hong, L., and S. E. Page. 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences of the United States of America* 101 (46):16385–89. <https://doi.org/10.1073/pnas.0403723101>.
- Howdle, Giles. 2023. "Microtargeting, Dogwhistles, and Deliberative Democracy." *Topoi*, February. <https://doi.org/10.1007/s11245-023-09889-3>.
- Intemann, Kristen. 2011. "Diversity and Dissent in Science: Does Democracy Always Serve Feminist Aims?" In *Feminist Epistemology and Philosophy of Science*, edited by Heidi E. Grasswick, 111–32. Dordrecht, Netherlands: Springer. https://doi.org/10.1007/978-1-4020-6835-5_6.
- John, Stephen. 2018. "Scientific Deceit." *Synthese*, November. <https://doi.org/10.1007/s11229-018-02017-4>.
- John, Stephen. 2019. "Science, Truth and Dictatorship: Wishful Thinking or Wishful Speaking?" *Studies in History and Philosophy of Science, Part A* 78 (December):64–72. <https://doi.org/10.1016/j.shpsa.2018.12.003>.

- Kassam, Alysha. Forthcoming. "Uncertainty and Risk Surrounding the Application of Social Models." *Synthese*.
- Kelsall, Joshua. 2020. "The Trust-Based Communicative Obligations of Expert Authorities." *Journal of Applied Philosophy* 38 (2):288–305. <https://doi.org/10.1111/japp.12465>.
- Khalifa, Kareem, and Jared Millson. 2020. "Explanatory Obligations." *Episteme* 17 (3):384–401. <https://doi.org/10.1017/epi.2019.55>.
- Khoo, Justin. 2017. "Code Words in Political Discourse." *Philosophical Topics* 45 (2):33–64. <https://doi.org/10.5840/philtopics201745213>.
- Khoo, Justin. 2021. "Code Words." In *The Routledge Handbook of Social and Political Philosophy of Language*, 1st ed., edited by Justin Khoo and Rachel Katharine Sterken, 147–60. New York: Routledge. <https://doi.org/10.4324/9781003164869-12>.
- Koles, Taylor. 2023. "On the Varied Species of Dogwhistles." Unpublished manuscript.
- Kuehn, Daniel. 2017. "Diversity, Ability, and Democracy: A Note on Thompson's Challenge to Hong and Page." *Critical Review* 29 (1):72–87. <https://doi.org/10.1080/08913811.2017.1288455>.
- Liptak, Adam. 2015. "Supreme Court Justices' Comments Don't Bode Well for Affirmative Action." *New York Times*, December 9, sec. U.S. <https://www.nytimes.com/2015/12/10/us/politics/supreme-court-to-revisit-case-that-may-alter-affirmative-action.html>.
- Longino, Helen E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, NJ: Princeton University Press. <https://doi.org/10.1515/9780691209753>.
- McKaughan, Daniel J., and Kevin C. Elliott. 2013. "Backtracking and the Ethics of Framing: Lessons from Voles and Vasopressin." *Accountability in Research* 20 (3):206–26. <https://doi.org/10.1080/08989621.2013.788384>.
- Morgan, Mary S., and Margaret Morrison, eds. 1999. *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511660108>.
- Novitz, David. 1977. *Pictures and Their Use in Communication: A Philosophical Essay*. The Hague, Netherlands: Martinus Nijhoff.
- O'Connor, Cailin. 2017. "Modeling Minimal Conditions for Inequity." Unpublished manuscript. <http://philsci-archive.pitt.edu/13473/>.
- Page, Scott E. 2007a. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press. <https://doi.org/10.1515/9781400830282>.
- Page, Scott E. 2007b. "Making the Difference: Applying a Logic of Diversity." *Academy of Management Perspectives* 21 (4):6–20. <https://doi.org/10.5465/amp.2007.27895335>.
- Page, Scott E. 2015. "Diversity Trumps Ability and the Proper Use of Mathematics." *Notices of the American Mathematical Society* 62 (1):9–10.
- Page, Scott E. 2017. *The Diversity Bonus*. Princeton, NJ: Princeton University Press.
- Pagin, Peter. 2014. "Assertion." In *The Stanford Encyclopedia of Philosophy*, Winter 2016 ed., edited by Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/assertion/>.
- Parker, Wendy S. 2009. "Confirmation and Adequacy-for-Purpose in Climate Modelling." *Aristotelian Society* 83 (1):233–49. <https://doi.org/10.1111/j.1467-8349.2009.00180.x>.
- Parker, Wendy S. 2015. "Getting (Even More) Serious about Similarity." *Biology and Philosophy* 30 (2):267–76. <https://doi.org/10.1007/s10539-013-9406-y>.
- Parker, Wendy S. 2020. "Model Evaluation: An Adequacy-for-Purpose View." *Philosophy of Science* 87 (3):457–77. <https://doi.org/10.1086/708691>.
- Quaranto, Anne. 2022. "Dog Whistles, Covertly Coded Speech, and the Practices That Enable Them." *Synthese* 200 (4):330. <https://doi.org/10.1007/s11229-022-03791-y>.
- Reijula, Samuli, and Jaakko Kuorikoski. 2021. "The Diversity-Ability Trade-off in Scientific Problem Solving." *Philosophy of Science* 88 (5):894–905. <https://doi.org/10.1086/714938>.
- Reutlinger, Alexander, Dominik Hangleiter, and Stephan Hartmann. 2018. "Understanding (with) Toy Models." *British Journal for the Philosophy of Science* 69 (4):1069–99. <https://doi.org/10.1093/bjps/axx005>.
- Santana, Carlos. 2021. "What's Wrong with Dogwhistles." *Journal of Social Philosophy* 53 (3):387–403. <https://doi.org/10.1111/josp.12409>.
- Saul, Jennifer M. 2017. "Racial Figleaves, the Shifting Boundaries of the Permissible, and the Rise of Donald Trump." *Philosophical Topics* 45 (2):97–116. <https://doi.org/10.5840/philtopics201745215>.

- Saul, Jennifer M. 2018. "Dogwhistles, Political Manipulation, and Philosophy of Language." In *New Work on Speech Acts*, 1st ed., edited by Daniel Fogal, Daniel W. Harris, and Matt Moss, 360–83. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198738831.003.0013>.
- Solomon, Miriam. 2001. *Social Empiricism*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/6296.001.0001>.
- Thompson, Abigail. 2014. "Does Diversity Trump Ability?" *Notices of the American Mathematical Society* 61 (9):1024. <https://doi.org/10.1090/noti1163>.
- Watkins, Aja. 2023. "Using Paleoclimate Analogues to Inform Climate Projections." *Perspectives on Science*, November, 1–58. https://doi.org/10.1162/posc_a_00622.
- Weisberg, Michael. 2012. "Getting Serious about Similarity." *Philosophy of Science* 79 (5):785–94. <https://doi.org/10.1086/667845>.
- Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford Studies in Philosophy of Science. Oxford: Oxford University Press.
- Witten, Kimberly. 2023. "The Definition and Typological Model of a Dogwhistle." *Manuscript*, July. <https://doi.org/10.1590/0100-6045.2023.v46n3.kw>.
- Wu, Jingyi. 2023. "Epistemic Advantage on the Margin: A Network Standpoint Epistemology." *Philosophy and Phenomenological Research* 106 (3):755–77. <https://doi.org/10.1111/phpr.12895>.
- Wylie, Alison. 2012. "Feminist Philosophy of Science: Standpoint Matters." *Proceedings and Addresses of the American Philosophical Association* 86 (2):47–76.

Cite this article: DiMarco, Marina. 2024. "Models as Dogwhistles." *Philosophy of Science* 91 (1):72–89. <https://doi.org/10.1017/psa.2023.100>