

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Dogwhistles: Persona and Ideology

Elin McCready
Aoyama Gakuin University

May 8, 2021

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

This talk presents joint work with Robert Henderson.

George Bush's 2003 State of the Union address contains the following line.

- (1) Yet there's power—wonder-working power—in the goodness and idealism and faith of the American people.

To most people this sounds like, at worst, a civil-religious banality, but to a certain segment of the population the phrase *wonder-working power* is intimately connected to their conception and worship of Jesus. When someone says (1), they hear (2).

- (2) Yet there's power—Christian power—in the goodness and idealism and faith of the American people.

On a 2014 radio program, Representative Paul Ryan said the following.

- (3) 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.

He was criticized shortly after for making a “thinly veiled racial attack”: the phrase *inner-city* is code or euphemism for African American neighborhoods (especially stereotypically racialized views of such neighborhoods).

- (4) We have got this tailspin of culture, in our African American neighborhoods 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.

These examples illustrate the notion of a *dogwhistle*—that is, language that sends one message to an outgroup while at the same time sending a second (often taboo, controversial, or inflammatory) message to an ingroup.

- ▶ There are proposals re DWs coming from philosophy:
- ▶ Stanley (2015) provides a semantic / pragmatic proposal, where dogwhistles are Pottsian CIs, contributing an at-issue component for the outgroup audience and a non-at-issue component that potentially only the ingroup is sensitive to.
- ▶ Khoo (2017) provides a purely pragmatic account, where dogwhistles involve certain default inferences.

Also proposals (e.g., Saul (2018)), which takes dogwhistles to be simple Gricean implicatures.. [redacted: time]

We think none of these proposals is correct, though exploring them is important because they expose certain tensions.

- ▶ We will see that dogwhistles cannot involve conventionalized TC-meaning (either at-issue or not-at-issue)
- ▶ but also that dogwhistles require some kind of conventionalization.

After exploring these previous accounts, we propose our own combining aspects of McCready (2012a), Burnett (2016, 2017) which we think better accounts for their core properties, while resolving this tension about conventionalization.

In broad strokes, we make the novel proposal that dogwhistles come in two types.

- ▶ The first type—*identifying dogwhistles*—concerns covert signals that the speaker has a certain persona, which we model by extending the *Sociolinguistic Signaling Games* of Burnett (2016, 2017).
- ▶ Today's main focus: the second type—*enriching dogwhistles*—involves sending a message with an enriched meaning whose recovery is contingent on recognizing the speaker's covertly signaled persona, which requires a further extension.

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

Plan:

- ▶ Briefly sketch the accounts of Stanley and Khoo, and our own
- ▶ Provide a formal theory of ideology and show how it induces enrichment in DWs
- ▶ Epistemic vigilance and hypervigilance via ideology matching

The conventional implicature account

Stanley (2015) argues that dogwhistle language involves a conventional non-at-issue component along the lines of more familiar expressions like slurs, honorifics, etc.

- ▶ A slur like *kraut* would have AI-component “German” and a NAI-component “I hate Germans”.
- ▶ A dogwhistle like *welfare* would have AI-component “the SNAP program” and a NAI-component “African Americans are lazy”.

We disagree with this characterization.

‘What is said’ by a dogwhistle?

- ▶ The use of dogwhistles is prompted by a desire to ‘veil’ a bit of content, but still to convey it in some manner. Deniability is essential.
- ▶ If a bit of content is conventional, it’s not deniable any longer. This can be seen with pejoratives, which clearly carry conventional NAI content.

(5) A: Angela Merkel is a kraut.

B: What do you have against Germans?

A: # I don’t have anything against Germans. I’m just talking about Merkel’s nationality.

Such dialogues are fine with dogwhistles; in the following, there seems to be no entailment that A has the relevant attitude.

- (6) A: Elin is living high on the hog on welfare again.
B: What do you have against poor people?
A: I don't have anything against poor people. I'm just saying Elin is on welfare and I saw her buying steak at the store.

- ▶ By this test, dogwhistles can be concluded not to be conventional.

The inferentialist account

Khoo 2017 argues that dogwhistles involve default inferences:

- ▶ Speaker claims that x is C and the interpreter believes that C 's are R 's, then the interpreter will conclude that x is R ; it's this kind of inference that Khoo thinks that dogwhistles license.
- ▶ If the interpreter believes that *inner-city* neighborhoods are African American neighborhoods. Then the speaker saying that people who live in inner-city neighborhoods lack a culture of work licenses the inference that people who live in African American neighborhoods lack a culture of work.

Non-substitutability argument

Problem: Khoo's inference follows from the expression TCs. Thus, any expression with the same TCs should dogwhistle.

- ▶ This is not true. A phrase like *downtown neighborhoods* doesn't dogwhistle like *inner city* does. The same for *welfare* and paraphrases like *assistance to the poor*

This suggests that while dogwhistles must not bear conventionalized content, some expressions are singled out as something like “dogwhistle expressions”, and so there is some kind of conventionalization.

Core properties to be accounted for:

- ▶ Dogwhistles are not part of conventional content, so speakers are able to avoid (complete) responsibility for what they convey.
- ▶ Dogwhistles can be identified as such, even if not bearing conventional content.

And further:

- ▶ Dogwhistles are semi-cooperative—that is, they are meant to be under-informative to one segment of the audience, while communicating a particular message to another.
- ▶ While deniable, dogwhistles are risky. Being detected using a dogwhistle by the wrong party should be costly.

SSGs

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

In recent work, Burnett (2016, 2017) pioneers the use of Bayesian signaling games to model identity construction through sociolinguistic variation.

- ▶ We take identifying dogwhistles to be only slightly more complex versions of sociolinguistic identity construction through variation of the kind Burnett (2016, 2017) discuss;
- ▶ enriching DWs are a special sort of identifying DW which interact with ideological background.

Burnett's Social Meaning Games which have the following simplified architecture (which we modify / elaborate further below):

- ▶ Players: a speaker S , a listener L
- ▶ Actions for players
 - ▶ The speaker chooses a persona p from the space of personas P
 - ▶ Based on their persona, the speaker chooses a message $m \in M$ to send to the listener.
 - ▶ Based on the message, the listener chooses a response $r \in R$, which in the simplest case we can identify with selecting an element of P —i.e., identifying the speaker's persona.

Utility functions for players: U_S/U_R —functions from $P \times M \times R$ to \mathbb{R} , which represents payoffs for every possible combination of actions.

- ▶ The speaker's utility is maximized by picking a message that sends the most information to the listener about the persona they want them to assign to them.
- ▶ The listener's utility is maximized if they extract the most information they can about a speaker's persona given their message.

We want the dogwhistle effect to arise from listeners being unaware (or uncertain) about the close connection between some bit of language and a persona.

⇒ We want listeners to have beliefs about a speaker's persona. . .

- ▶ . . . but also beliefs about how personas and messages are connected.

That is, listeners have prior over P , but also beliefs about $P(m|p)$ —namely how closely messages are linked to particular personas.

We can now update a listener's belief about the speaker's persona given their message by doing bayesian inference.

$$(7) \quad P(p|m) \propto P(p)P(m|p)$$

'The probability of a persona given a message is proportional to prior probability of the persona and the likelihood of sending that message given that persona'

- ▶ This is an extension of Burnett (2016, 2017), who takes social meanings to be fully lexicalized, i.e., the likelihood $P(m|p) = 1$ when p and m are consistent.

The final ingredient we need to provide utility functions.

- ▶ **listener** is straightforward—utility is maximized by extracting as much information from a message as possible about a speaker's persona—that is, by doing doing bayesian inference as just described.
- ▶ **speakers**: utility is more complex because unlike in many signaling games, the speaker doesn't just pick messages based on some type assigned by nature—i.e., they don't just *report* their personas.
- ▶ Instead, speakers have preferences for different personas, some of which may be dependent on how the listener would react to that persona.

Thus, we must allow for speakers to “construct” a persona in concert with their listeners.

- ▶ Speakers want to present themselves in a certain way.
- ▶ Speakers will also be sensitive to whether listeners will approve of that persona or not.
- ▶ In adversarial contexts, a speaker might have to juggle presenting a safe persona with a persona they might prefer to present (or prefer to present to another audience that might be listening)—this is when dogwhistle language become useful.

Along these lines, we follow Burnett (2017); Yoon et al. (2016) in assuming that the utility calculation takes into account the message's social value, which is given by two functions:

- ▶ The speaker has a function v_S that assigns a positive real number to each persona representing their preferences.
- ▶ The listener has a function v_L that assigns a real number (positive or negative) to each persona representing their (dis)approval.

We can now calculate the speaker's utility.

The utility is dependent on the affective values of the range of personas consistent with the message and the likelihood that the particular persona is recovered given the message, as follows:

$$(8) \quad U_S^{Soc}(m, L) = \sum_{p \in [m]} P(p|m) + v_S(p)P(p|m) + v_L(p)P(p|m)$$

When only one listener is addressed, dogwhistles reduce to ordinary social meaning; the speaker should choose a signal which maximizes U_S^{Soc} .

- ▶ Dogwhistles come into their own when speakers address groups of individuals with mixed preference over personas, different priors for the speaker's persona, and different experiences about the likelihood of a persona given a message.
- ▶ The simplest way to assign utilities to the group case is to sum over all listeners; we will assume this metric in the following.

$$(9) \quad U_S^{Soc}(m, G) = \sum_{L \in G} U_S^{Soc}(m, L)$$

With this utility function, the basic prediction is:

- ▶ Speakers will use language that maximizes their social utility wrt a group of listeners.
- ▶ For the dogwhistle case, this happens when using the dogwhistle allows gain of higher social utility than otherwise wrt the entire group,
- ▶ i.e., when the dogwhistle gives benefit for some 'savvy' listeners while avoiding deficits that would come from speakers disliking the persona but oblivious to the dogwhistle.

Detailed formal example redacted for time reasons.

Enriching dogwhistles

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

To deal with enriching dogwhistles, we believe that we must achieve a new kind of synthesis.

- ▶ We think the Khoo-style account we critiqued earlier, which focuses on enriching dogwhistles, fails short for not social meanings, and for not building off a social meaning account of identifying dogwhistles.
- ▶ At the same time, our previous account of enriching dogwhistles in terms of pragmatic enrichments falls short as it failed to make the precise connection between enrichment and personas clear.

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

We take enriching dogwhistles to be *identifying dogwhistle*⁺ in the following sense:

- ▶ On use, a savvy listener detects the dogwhistle and assigns the speaker a relevant persona.
- ▶ Those personas are associated with ideologies, which come with background assumptions.
- ▶ The listener, consciously or not, learns what ideological grounds the speaker is speaking on.
- ▶ The savvy listener can then draw inferences about the speaker's intended content.

Our task is to:

- ▶ make clear, and work into our formal model, what ideologies are and how they can be entangled with personas.
- ▶ "bridge the gap" between speaker and hearer, that is, understand how recognizing the ideological grounds on which the speaker is speaking can cause the listener to make inferences about the speaker's communicative intent, and, sometimes, influence listener behavior.

Ideologies: formal treatment

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

What do ideologies do? And what happens when we recognize a persona associated with an ideology?

- ▶ Ideologies indicate affect and (dis)approval of various actions or people, but also bring in more global assumptions about the world.
- ▶ Thus, to understand what effects assigning personas to discourse agents has, we need at minimum:
 1. a way of valuating actions and individuals and
 2. a way of introducing beliefs and world knowledge to our models.

Assigning affective values

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

We need a function that can assign affective values to objects relevant to ideologies and personas.

- ▶ We will use ρ ('rate') for our new function.
- ▶ This function takes individuals as input and yields real number as value: we allow both positive and negative real numbers here, as with the listener valuation function v_L on personas.

Values to ρ : attitudes toward particular individuals (Trump), but also attitudes toward behaviors, groups of people, and properties.

- ▶ We can treat these as individuals by making use of the kind-mapping function ‘ \cap ’ (Chierchia and Turner (1988); Chierchia (1998)).
 - ▶ standard: nominalizations (self-predication), bare nominals (Chinese/Japanese)

(10) Being nice is nice.
nice(\cap *nice*)

We use \cap to produce kinds for behaviors/groups/properties in general.

Ideologies: epistemic bases

Ideologies assign value; they also comprise sets of beliefs about how the world is:

- ▶ the kinds of things that make it up, the properties of kinds of people, systems, and objects, and the causal elements that induce and condition change.

The truth-evaluable elements which make up an ideology are modelable as sets of such propositions.

- ▶ We call each set of this kind the *basis* of an ideology.

What sorts of propositions form the basis of ideologies? The answers to this question are as various as ideologies themselves.

- ▶ QAnon ideology takes the existence of a conspiracy with bizarre goals as a given.
- ▶ Racist ideologies involve beliefs about the relative value and superiority of ethnic groups, and so on.
- ▶ We use the notation \mathcal{B} for ideological bases.

All these beliefs can function to bridge gaps in reasoning and connect things that without the ideology would be nonobvious.

- ▶ We will argue that it is these sorts of beliefs, and ideological bases in general, that trigger enriching dogwhistles.

Upshot: the ideologies related to personas have the form $\iota = \langle \rho, \mathcal{B} \rangle$ and so consist of pairs of affect-assigning functions and ideological bases.

- ▶ The propositions comprising the basis of an ideology can be somewhat indeterminate and vary from individual to individual depending on where they have acquired their ideological beliefs.
- ▶ So we must think in terms of related but possibly non-identical ideologies, which we can view as ideological equivalence classes.
- ▶ We thus define the basis of an ideology as the set of beliefs common to all its variants (here Π is a projection function).

$$(11) \quad \Pi_2(\iota) =_{df} \bigcap \Pi_2(\iota'), \text{ where } \iota' \sim \iota.$$

Social sincerity

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

We need to make one assumption about the relation between persona and belief.

- ▶ What kind of personas are available for an individual? That is, in a linguistic context, what kinds of personas can a speaker assume or signal?
- ▶ We assume here that speaker personas are required to be sincerely assumed, ie that the basis of that persona correlates with the speaker's actual beliefs.
- ▶ This is an analogue of Gricean Quality for the domain of social meaning, which we will call *Social Sincerity*.

Formally speaking, this amounts to requiring the personas compatible with the speaker's utterance, $\mathbf{emf}(u)$, to associate with bases which have some relationship to the speaker's beliefs.

(12) *Social Sincerity*

$$\forall s, u, \pi [\text{utter}(s)(u) \wedge \pi \in \mathbf{emf}(u) \wedge \iota_\pi \rightarrow \text{MOST}(p \in \Pi_2(\iota_\pi))(\text{Bel}(s, p))]$$

'If a speaker utters a sentence compatible with persona π , they believe a significant number of the propositions comprising the basis for π .'

Two comments on this principle.

- ▶ It is relatively weak in the sense that it simply requires the speaker to hold most of the beliefs associated with the ideology. Possible modifications:
 - ▶ use a different quantifier
 - ▶ use different underlying theory, eg use a contextually determined parameter for sincerity in the manner of Kennedy (2007) on vague predicates or McCready (2015) for reliability of information source.
- ▶ It also treats all beliefs in $\Pi_2(1)$ identically, but likely some of these beliefs are more 'core' to the ideology than others. which could be modeled by weighting them as in e.g. the belief revision literature on entrenchment (Gärdenfors, 1988)

With these formal elements, we can return to enrichment.

- ▶ We argue that enrichment is a multistep process, the first part of which is shared with identifying dogwhistles.
 1. Listener identifies speaker's persona on the basis of their utterance [IDs]
 2. Listener calls up basis of ideology associated with that persona
 3. If the basis, plus the utterance content, allows inferences to be drawn: enrichment.

Case.

inner city:

- ▶ Suppose, for the (quasi)racist persona and corresponding ideology \mathfrak{t} communicated by this DW to a savvy listener, $live_inner_city(x) > black(x) \in \Pi_2(\mathfrak{t})$.
- ▶ This extra premise licenses an inference from ‘inner city people don’t work’ to ‘Black people living in cities don’t work’, which is the enriched meaning.

This statement of course can then be used to understand the speaker’s political views, draw conclusions about their policy decisions, etc.

- ▶ Note: *social sincerity* assumption is necessary here.

We can already see how this view improves on Khoo and our previous work.

- ▶ Khoo: the inference depends purely on semantic content.
 - ▶ No way to explain how semantically coextensive terms trigger/don't trigger these inferences;
 - ▶ for us, the mediation through persona, which is only enabled by DWs for savvy listeners, makes coextensive phrases act differently in the inferences they trigger via ideologies.
- ▶ Old us: mediated by DWs, but how?
 - ▶ Now us: principled explanation available of how dogwhistles trigger conclusions about personas, and thus ideologies, and thus inferences.

Taking the relation between persona and ideology seriously gives the necessary ingredients for an explanation that satisfies our desiderata.

Crucial pieces of evidence that we need something like this from Hurwitz and Peffley (2005)

- ▶ White people in studies on racial dogwhistles are more likely to unconsciously assign stereotypes to racial minorities on a post-test, but that you can correct this by telling them that the dogwhistle is racist before they hear it.
- ▶ African American in these same studies show no effect when hearing the dogwhistle in assigning racial stereotypes, yet they are aware that the word in questions is, in fact, a dogwhistle.

Clearly, we must separate "hearing the dogwhistle" from "making inferences based on the dogwhistle". Linking personas with these background ideologies can help us do this.

- ▶ When a white person hears a racial dogwhistle, they learn that the speaker, in virtue of their persona, is willing to have a conversation from a "white chauvinist" standpoint.
- ▶ Without objection, the default effect might be that this ideology becomes the ground for the conversation, dragging along all the inferences associated with that ideology—hence, the effect we see on implicit bias post-tests.

In contrast,

- ▶ when white participants are explicitly warned about the dogwhistle, resisting the associated ideology becomes salient. Listeners still "hear" the dogwhistle, but they are primed to not slide onto the ideological grounds the persona associated with that whistle invites.
- ▶ We see a similar effect with African American participants. They "hear" the dogwhistle, but by default resist grounding the conversation in a racist ideology by. This makes sense for obvious reasons.

Epistemic vigilance and hypervigilance.

When observing a speaker using an expression we know to be a dogwhistle, what is the proper reaction?

- ▶ It is possible that the expression is being used innocently: speakers often hear dogwhistles without recognizing them, and may pick up expressions from political discourse without knowing their dogwhistley quality.
- ▶ Some uses of dogwhistles, then, are innocent.
- ▶ How can these be distinguished from cases where dogwhistles are genuinely used to deceive, as covert signals of identity?

For a dogwhistle to be useful, the speaker must judge the speech situation to be one in which it is to their advantage to covertly signal.

- ▶ ie. a context which is, at least potentially, one where revealing their true identity could have negative effects.

A (partial) zero-sum situation is required for DW to be useful.

- ▶ A consequence of this observation is that, for a listener to judge a particular use of an expression as a dogwhistle, she must think that the speaker is intending to deceive.
- ▶ She must take the speaker to believe the interests of some discourse participants not to be aligned.

The result is that judging an expression a dogwhistle already imputes hostile intent to the speaker.

- ▶ If this is wrong, it can both create arguments and disagreement where there might have been none.
- ▶ It can also dispose the hearer to systematically misconstrue other utterances of the speaker due to shifts in the probabilities she assigns to speaker personas and consequently to what she guesses the speaker's intentions to be (as in discussion of *fucking* in McCready (2012b)).

Thus: while it is important to be epistemically vigilant about dogwhistles, it is also important not to be hypervigilant.

- ▶ Under what conditions does hypervigilance arise?
- ▶ Experience: the buildup of bad interactions which get associated with a particular term (Twitter).
 - ▶ This boils down to shifts in priors, which is already expressed in our model.
- ▶ Value: sometimes, it is beneficial to the speaker to search out dogwhistles.

We want to look at the latter case in a bit more detail now.

- ▶ Recall that utilities are assigned, in our theory, via a combination of informative content and value assigned to particular social meanings. On what basis?
- ▶ Many possibilities, but many are ideological in a broad sense (tradition/radicalness, political views, social groupings).
- ▶ One metric is similarity: 'I like people who are like me.'
- ▶ If so: we can assign affective values on the basis of similarity metrics between speaker and hearer personas (proposed in Henderson and McCready (2019)).

(of course, only one aspect of value assignment)

Idea: people are sensitive to dogwhistles to precisely the degree they have strong feelings about the ideology.

- ▶ If highly positive, listener can learn that the speaker is sympathetic;
- ▶ if highly negative, that the speaker is to be avoided, or even combatted.
- ▶ This follows pretty directly from our theory: once affective value is incorporated into utilities, learning someone holds a high pos/neg utility ideology is useful to know.

So to the extent that speaker persona is similar/dissimilar to listener persona, utility is high.

- ▶ The result can be hypervigilance.
- ▶ We think there is a mathematical result to be found here, somewhat akin to the credibility result for cheap talk games of Farrell (1993), where a signal is credible to the degree that the interests of the sender align with that of the receiver:
 - ▶ here, perhaps, listeners interpret possible dogwhistles as dogwhistles to the degree that doing so reflects a utility change, ie to the extent that personas are (dis)similar along some dimension.
- ▶ This is ongoing work.

A quick summary. This talk has:

- ▶ Argued against several existing accounts of dogwhistles
- ▶ Distinguished two types of dogwhistle, both of which convey social personas but only one of which has at-issue content which is influenced by the persona recovered
- ▶ Modeled the two types using an extension and variant of Burnett's social meaning games
- ▶ Briefly discussed some implications for unconscious bias and vigilance about dogwhistles.

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

[Intro](#)

[Previous work](#)

[Our account](#)

[Enrichment](#)

[\(Hyper\)vigilance](#)

[References](#)

THANK YOU!!!!!!!!!!!!!!

Dogwhistles:
Persona and
Ideology

Elin McCready
Aoyama Gakuin
University

Intro

Previous work

Our account

Enrichment

(Hyper)vigilance

References

- Heather Burnett. Signalling games, sociolinguistic variation and the construction of style. In *the 40th Penn Linguistics Colloquium, University of Pennsylvania*, 2016.
- Heather Burnett. Sociolinguistic interaction and identity construction: The view from game-theoretic pragmatics. *Linguistics and Philosophy*, 2017.
- Gennaro Chierchia. Reference to kinds across language. 6:339–405, 1998. doi: 10.1023/A:1008324218506.
- Gennaro Chierchia and Raymond Turner. Semantics and property theory. *Linguistics and Philosophy*, 11:261–302, 1988.
- Joseph Farrell. Meaning and credibility in cheap-talk games. *Games and Economic Behavior*, 5(4):514–31, 1993. doi: 10.1006/game.1993.1029.
- Peter Gärdenfors. *Knowledge in Flux*. MIT Press, 1988.
- R. Henderson and E. McCready. Dogwhistles, trust, and ideology. In *Proceedings of the 22nd Amsterdam Colloquium*, 2019.
- Jon Hurwitz and Mark Peffley. Playing the race card in the post-willie horton era the impact of racialized code words on support for punitive crime policy. *Public Opinion Quarterly*, 69(1):99–112, 2005.
- Chris Kennedy. Vagueness and gradability: The semantics of relative and absolute gradable predicates. 30(1):1–45, 2007.
- Justin Khoo. Code words in political discourse. *Philosophical Topics*, 2017.
- E. McCready. *Reliability in Pragmatics*. Oxford University Press, 2015.
- Elin McCready. Emotive equilibria. *Linguistics and Philosophy*, 35(3):243–283, 2012a.
- Elin McCready. Emotive equilibria. 35:243–283, 2012b.
- Jennifer Saul. Dogwhistles, political manipulation, and philosophy of language. *New Works on Speech Acts*, pages 360–383, 2018.
- Jason Stanley. *How propaganda works*. Princeton University Press, 2015.
- Elina J Yoon, Michael Henry Tessler, Noah D Goodman, and Michael C Frank. Talking with tact: Polite language as a balance between kindness and informativity. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Cognitive Science Society, 2016.