Slurs in the architecture of grammar

Maria Esipova

University of Genoa

EPITHETS & STAL

May 8, 2025

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?
- How do we formalize these different types of meaning in a way that captures their fundamental architectural properties?

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?
- How do we formalize these different types of meaning in a way that captures their fundamental architectural properties?

Some of the specific topics that are particularly informative when trying to address the broad questions above:

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?
- How do we formalize these different types of meaning in a way that captures their fundamental architectural properties?

Some of the specific topics that are particularly informative when trying to address the broad questions above:

- at-issue vs. not-at-issue meaning

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?
- How do we formalize these different types of meaning in a way that captures their fundamental architectural properties?

Some of the specific topics that are particularly informative when trying to address the broad questions above:

- at-issue vs. not-at-issue meaning
- truth-conditional vs. non-truth-conditional meaning

Big picture questions

Some of the **broad questions** that I think should underlie much of our theorizing in semantics:

- What are the different types of meaning that humans can express?
- How does expression of these different types of meaning get operationalized architecturally, at different levels of representation and their interfaces?
- How do we formalize these different types of meaning in a way that captures their fundamental architectural properties?

Some of the specific topics that are particularly informative when trying to address the broad questions above:

- at-issue vs. not-at-issue meaning
- truth-conditional vs. non-truth-conditional meaning
- and the difference between the two

Truth-conditional vs. non-truth-conditional meaning

For a while now, I've been trying to defend the **truth-conditional vs. non-truth-conditional distinction as a fundamental architectural distinction**, which is not reducible to the at-issue vs. not-at-issue distinction and which thus needs to be reflected in our theories, accordingly

Truth-conditional vs. non-truth-conditional meaning

For a while now, I've been trying to defend the **truth-conditional vs. non-truth-conditional distinction as a fundamental architectural distinction**, which is not reducible to the at-issue vs. not-at-issue distinction and which thus needs to be reflected in our theories, accordingly

One aspect of this architectural distinction is **(non-)compositionality**: not-at-issue (but truth-conditional) meanings still have to interact semantically with their syntactic surroundings while non-truth-conditional meanings don't (see Esipova 2024 for more on this)—suggesting distinct cognitive resources at use in the two cases:

Truth-conditional vs. non-truth-conditional meaning

For a while now, I've been trying to defend the **truth-conditional vs. non-truth-conditional distinction as a fundamental architectural distinction**, which is not reducible to the at-issue vs. not-at-issue distinction and which thus needs to be reflected in our theories, accordingly

One aspect of this architectural distinction is **(non-)compositionality**: not-at-issue (but truth-conditional) meanings still have to interact semantically with their syntactic surroundings while non-truth-conditional meanings don't (see Esipova 2024 for more on this)—suggesting distinct cognitive resources at use in the two cases:

(1) *Context: The speaker sees that someone is about to hurt their dog.*
   a. Step away from my {fucking / (god)damn / bloody} dog!
      → The speaker's affect is towards their dog (or anything else in the sentence).
   b. Step away from my {lovely / awesome / obnoxious / disgusting} dog!
      → The speaker's attitude is towards their dog.

Truth-conditional vs. non-truth-conditional meaning

For a while now, I've been trying to defend the **truth-conditional vs. non-truth-conditional distinction as a fundamental architectural distinction**, which is not reducible to the at-issue vs. not-at-issue distinction and which thus needs to be reflected in our theories, accordingly

One aspect of this architectural distinction is **(non-)compositionality**: not-at-issue (but truth-conditional) meanings still have to interact semantically with their syntactic surroundings while non-truth-conditional meanings don't (see Esipova 2024 for more on this)—suggesting distinct cognitive resources at use in the two cases:

(1) *Context: The speaker sees that someone is about to hurt their dog.*
   a. Step away from my {fucking / (god)damn / bloody} dog!
      → The speaker's affect is towards their dog (or anything else in the sentence).
   b. Step away from my {lovely / awesome / obnoxious / disgusting} dog!
      → The speaker's attitude is towards their dog.

Today, I will focus on another phenomenon that sets apart truth-conditional and non-truth-conditional meanings architecturally, but in a more complex way, that also cross-cuts with the at-issue vs. not-at-issue distinction: **recoverability during ellipsis/anaphora resolution**

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) $\rightarrow$ Pam used to smoke.

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) $\rightarrow$ Pam used to smoke. (ii) $\rightarrow$ Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) → Pam used to smoke. (ii) → Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

b. Lea regrets leaving, {but Mia doesn't / and Mia does, too / and so does Mia}.
(i) → Lea left.

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) → Pam used to smoke. (ii) → Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

b. Lea regrets leaving, {but Mia doesn't / and Mia does, too / and so does Mia}.
(i) → Lea left. (ii) → Mia left.

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) → Pam used to smoke. (ii) → Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

b. Lea regrets leaving, {but Mia doesn't / and Mia does, too / and so does Mia}.
(i) → Lea left. (ii) → Mia left.

c. Zoe knows that she is in danger, {and Ash does, too / and so does Ash}, (#although Ash is not actually in danger, they just believe that they are). (sloppy reading)
(i) → Zoe is in danger.

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) → Pam used to smoke. (ii) → Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

b. Lea regrets leaving, {but Mia doesn't / and Mia does, too / and so does Mia}.
(i) → Lea left. (ii) → Mia left.

c. Zoe knows that she is in danger, {and Ash does, too / and so does Ash}, (#although Ash is not actually in danger, they just believe that they are). (sloppy reading)
(i) → Zoe is in danger. (ii) → Ash is in danger.

Different types of content under ellipsis

Different types of content behave differently during ellipsis/anaphora resolution—hf. "under ellipsis" (see, e.g., Esipova 2019 for an overview and refs therein):

- *Some* **presuppositions** of, a.o., predicates denoting stages of events or factives can never be ignored under ellipsis:

(2) a. Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
(i) → Pam used to smoke. (ii) → Kim used to smoke. (modulo local accommodation: presupposition not ignored, but at-issue)

b. Lea regrets leaving, {but Mia doesn't / and Mia does, too / and so does Mia}.
(i) → Lea left. (ii) → Mia left.

c. Zoe knows that she is in danger, {and Ash does, too / and so does Ash}, (#although Ash is not actually in danger, they just believe that they are). (sloppy reading)
(i) → Zoe is in danger. (ii) → Ash is in danger.
(Note that under the standard story, *know that p* and *believe that p* are truth-conditionally equivalent in their local contexts.)

Different types of content under ellipsis

- **Expressive contributions** of items like *fucking*, *damn*, etc. are always ignored under ellipsis:

Different types of content under ellipsis

- **Expressive contributions** of items like *fucking*, *damn*, etc. are always ignored under ellipsis:

(3) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
(i) $\rightarrow$ A is experiencing strong emotions.
(ii) $\not\rightarrow$ B is experiencing strong emotions.

Focus of today's talk: slurs in elliptical responses

Today, I will focus on **slurs** (denotational component + prejudice component; see, e.g., Orlando & Saab 2020 for a collection of papers on slurs)

In particular, I will be asking whether/to what extent the **prejudice component** of slurs is preserved in different types of elliptical responses and what that tells us about how this prejudice component is operationalized architecturally and how it should therefore be analyzed formally

Focus of today's talk: slurs in elliptical responses

Today, I will focus on **slurs** (denotational component + prejudice component; see, e.g., Orlando & Saab 2020 for a collection of papers on slurs)

In particular, I will be asking whether/to what extent the **prejudice component** of slurs is preserved in different types of elliptical responses and what that tells us about how this prejudice component is operationalized architecturally and how it should therefore be analyzed formally

(4) *Context: In the 'Blade Runner' universe, 'skinjob' is a slur for synthetic humans (neutral term 'replicant').*
A: Did you see a skinjob?
B: No, I didn't. / Yes, I did. / Yes, I saw one.
*Question probing the presence/strength of the **prejudice inference**:*
Is B prejudiced against replicants? (How likely? / How prejudiced?)

This talk

1. Introduction
2. Background: more on different types of content under ellipsis
3. Experiment 1: fictional slurs
4. Experiment 2: real-life slurs
5. Discussion: the prejudice component of slurs has both a truth-conditional and a non-truth-conditional component, and is partially, but not fully preserved if the slur itself is not uttered, but is recovered during ellipsis resolution—need for hybrid analysis
6. Conclusion: limitations of the present experiments and moving forward

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
→ Kim used to smoke.

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
$\rightarrow$ Kim used to smoke.

(6) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
$\not\rightarrow$ B is experiencing strong emotions.

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
→ Kim used to smoke.

(6) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
$\not\rightarrow$ B is experiencing strong emotions.

1. Unlike *stop* in (5), *fucking* in (6) is an **adjunct** (?) inside an nP that, in turn, is targeted by *one* (?) or is inside the vP targeted by VPE/*do so*.

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
→ Kim used to smoke.

(6) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
↛ B is experiencing strong emotions.

1. Unlike *stop* in (5), *fucking* in (6) is an **adjunct** (?) inside an nP that, in turn, is targeted by *one* (?) or is inside the vP targeted by VPE/*do so*. Cf. truth-conditional, but not-at-issue modifiers (see Esipova 2019, 2021; Sailor & Colasanti 2020; also Esipova 2024 on how truth-conditional evaluative modifiers differ from expressives like *fucking*):

   (7) *Context: A loves all scotch; they are in a bar with B.*
   A: Could you get me another delicious shot of Laphroaig?
   B: OK, I {will / will do so / will get you one}. (I don't know how you drink this stuff, it's disgusting.)
   ↛ B considers all (or any) shots of Laphroaig delicious.

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
$\rightarrow$ Kim used to smoke.

(6) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
$\not\rightarrow$ B is experiencing strong emotions.

2. The presupposition of *stop* is an **"ontological precondition"** for its at-issue content (see Roberts & Simons 2023 on some types of presuppositions as ontological preconditions), but that's not the case for the expressive contribution of *fucking*. Note, however, that this is not necessarily the case for *know* either. But not being such a precondition is at least a prerequisite for ignorability

More on different types of content under ellipsis

Let's zoom in on the relevant differences between presuppositions of items like *stop* and the expressive component of items like *fucking* in (6):

(5) Pam stopped smoking, {but Kim didn't / and Kim did, too / and so did Kim}.
$\rightarrow$ Kim used to smoke.

(6) A: Did you bring a fucking gun to my house?
B: No, I didn't. / Yes, I did. / Yes, I did so. / Yes, I brought one.
$\not\rightarrow$ B is experiencing strong emotions.

3. Acts of producing expressives like *fucking* in (6) are **performative**: the speaker achieves their expressive goal (e.g., to let out their frustration) by producing a given form (use via mention) and can't do so w/o performing this act (no use w/o mention). But anaphoric phenomena like ellipsis hinge on not producing the form (they are instances of use w/o mention)

Another note of caution: anaphoric reference to forms/acts

NB: Anaphoric reference to forms—or to acts of producing them—is, possible; i.e., we can refer to the performative effects associated with the act of producing a certain form without obtaining them

Another note of caution: anaphoric reference to forms/acts

NB: Anaphoric reference to forms—or to acts of producing them—is, possible; i.e., we can refer to the performative effects associated with the act of producing a certain form without obtaining them

E.g., in (8), B refers to A's expressive act to assert that they have the same mental state as A, but they do not themselves obtain the tension-relieving benefits of the expressive act they are referring to:

(8) A: Fuck! / *punches a wall*
B: Yep, {same / that}.

So, what about slurs?

Slurs are a more complex case, however:

So, what about slurs?

Slurs are a more complex case, however:

1. **A slur doesn't have to be an adjunct**, and, in fact, can be the "lexical noun" of the constituent targeted by different types of ellipsis, as in noun slurs in nPs targeted by *one*-replacement (like *stop*, unlike *fucking* in the examples above)—less ignorable?

So, what about slurs?

Slurs are a more complex case, however:

1. **A slur doesn't have to be an adjunct**, and, in fact, can be the "lexical noun" of the constituent targeted by different types of ellipsis, as in noun slurs in nPs targeted by *one*-replacement (like *stop*, unlike *fucking* in the examples above)—less ignorable?
2. Despite that, **the prejudice component of a slur is not an ontological precondition for the at-issue component** (unlike the presupposition of *stop*, like the expressive component of *fucking*)—so, in principle ignorable

So, what about slurs?

Slurs are a more complex case, however:

1. **A slur doesn't have to be an adjunct**, and, in fact, can be the "lexical noun" of the constituent targeted by different types of ellipsis, as in noun slurs in nPs targeted by *one*-replacement (like *stop*, unlike *fucking* in the examples above)—less ignorable?
2. Despite that, **the prejudice component of a slur is not an ontological precondition for the at-issue component** (unlike the presupposition of *stop*, like the expressive component of *fucking*)—so, in principle ignorable
3. **Slurs definitely have a performative component**: they can be used performatively to intentionally cause offense (use via mention) and can even have a performative effect of offense by virtue of being uttered in the absence of such intent on the speaker's part (mention w/o use), but **it is unclear if the prejudice component can be preserved if a slur is recovered but not uttered (use w/o mention), i.e., if it is purely performative**

Prior literature

Saab (2020) looks at Spanish data involving dialogues with elliptical responses to antecedents that contain ethnic slurs, as well as stylistic pairs like the informal *morfar* vs. the stylistically neutral *comer* ('eat'):

Prior literature

Saab (2020) looks at Spanish data involving dialogues with elliptical responses to antecedents that contain ethnic slurs, as well as stylistic pairs like the informal *morfar* vs. the stylistically neutral *comer* ('eat'):

(9) A: ¿A cuántos sudacas viste en la fiesta?
to how-many South Americans$_{\text{PEJ}}$ saw.2SG in the party
"How many South Americans$_{\text{PEJ}}$ did you see at the party?"

B: Vi a tres <sudacas>, pero podrías evitar ese
saw.1SG to three <South American$_{\text{PEJ}}$> but could.2SG avoid that
modo de hablar de los sudamericanos. Yo nunca hablo así de ellos.
way of speaking of the South American I never speak so of them
"I saw three, but you could avoid that way of speaking about South Americans. I never talk that way." (Saab 2020, (43))

Prior literature

Saab (2020) looks at Spanish data involving dialogues with elliptical responses to antecedents that contain ethnic slurs, as well as stylistic pairs like the informal *morfar* vs. the stylistically neutral *comer* ('eat'):

(9) A: ¿A cuántos sudacas viste en la fiesta?
to how-many South Americans$_{\text{PEJ}}$ saw.2SG in the party
"How many South Americans$_{\text{PEJ}}$ did you see at the party?"

B: Vi a tres <sudacas>, pero podrías evitar ese
saw.1SG to three <South American$_{\text{PEJ}}$> but could.2SG avoid that
modo de hablar de los sudamericanos. Yo nunca hablo así de ellos.
way of speaking of the South American I never speak so of them
"I saw three, but you could avoid that way of speaking about South Americans. I never talk that way." (Saab 2020, (43))

(10) A: Qué morfaste?
what ate.2SG.INFORMAL
'What did you eat$_{\text{INFORMAL}}$?'

B: Una pizza <?>, pero no tolero cuando hablás tan
a pizza but not tolerate.1SG when speak.2SG so
informalmente. Yo nunca lo hago.
informally I never it do
"A pizza. But I don't tolerate when you speak informally. I never do it." (Saab 2020, (21))

Prior literature

Saab compares the exchanges above to versions thereof where the target items are repeated in the response, rendering the response contradictory:

Prior literature

Saab compares the exchanges above to versions thereof where the target items are repeated in the response, rendering the response contradictory:

(11) A: ¿A cuántos sudacas viste en la fiesta?
to how-many South Americans$_{\text{PEJ}}$ saw.2SG in the party
B:#Vi a tres sudacas, pero podrías evitar ese
saw.1SG to three South American$_{\text{PEJ}}$ but could.2SG avoid that
modo de hablar de los sudamericanos. Yo nunca hablo así de
way of speaking of the South American I never speak so of
ellos.
them (Saab 2020, (44))

Prior literature

Saab compares the exchanges above to versions thereof where the target items are repeated in the response, rendering the response contradictory:

(11) A: ¿A cuántos sudacas viste en la fiesta?
to how-many South Americans$_{\text{PEJ}}$ saw.2SG in the party
B:#Vi a tres sudacas, pero podrías evitar ese
saw.1SG to three South American$_{\text{PEJ}}$ but could.2SG avoid that
modo de hablar de los sudamericanos. Yo nunca hablo así de
way of speaking of the South American I never speak so of
ellos.
them (Saab 2020, (44))

(12) A: Qué morfaste?
what ate.2SG.INFORMAL
B:#Una pizza morfé, pero no tolero cuando
a pizza ate.1SG.INFORMAL but not tolerate.1SG when
hablás tan informalmente. Yo nunca lo hago.
speak.2SG so informally I never it do (Saab 2020, (22))

Prior literature

Saab's conclusions:

Prior literature

Saab's conclusions:

- The prejudice component of slurs and the stylistic effect of words like *morfar* are ignored under ellipsis ("ellipsis is an apt strategy to nullify the bias encoded in some lexical items")

Prior literature

Saab's conclusions:

- The prejudice component of slurs and the stylistic effect of words like *morfar* are ignored under ellipsis ("ellipsis is an apt strategy to nullify the bias encoded in some lexical items")
- Slurs and their neutral counterparts, as well as pairs like *morfar* and *comer*, are thus different phonological realizations of the same root, with the meaning differences between the two holding only at PF

Prior literature

Saab's conclusions:

- The prejudice component of slurs and the stylistic effect of words like *morfar* are ignored under ellipsis ("ellipsis is an apt strategy to nullify the bias encoded in some lexical items")
- Slurs and their neutral counterparts, as well as pairs like *morfar* and *comer*, are thus different phonological realizations of the same root, with the meaning differences between the two holding only at PF
- In our terms: this amounts to saying the prejudice component of slurs is purely performative (but his claim is even stronger, as one can still maintain that a slur and its neutral counterpart are distinct roots, but their meaning differences lie exclusively in the performative dimension)

Prior literature

Issues with Saab's claim:

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors
- The way we probe the extent to which said bias is preserved matters! In Saab's examples, B simply denies talking in a certain way

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors
- The way we probe the extent to which said bias is preserved matters! In Saab's examples, B simply denies talking in a certain way
  - But this can be taken to mean that they simply avoid producing certain forms, in which case the contrast between (9)/(10) and (11)/(12) is trivial and tells us nothing about whether there is no non-performative "bias" component in the target items above that is still preserved when said item is recovered during ellipsis resolution

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors
- The way we probe the extent to which said bias is preserved matters! In Saab's examples, B simply denies talking in a certain way
  - But this can be taken to mean that they simply avoid producing certain forms, in which case the contrast between (9)/(10) and (11)/(12) is trivial and tells us nothing about whether there is no non-performative "bias" component in the target items above that is still preserved when said item is recovered during ellipsis resolution
  - Also, we can can situationally do things we are in general opposed to; cf.:

    (13) A: How is Pam doing?
    B: She's doing fine, as far as I know, but can you please stop asking me about her, you know I don't want to talk about her.

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors
- The way we probe the extent to which said bias is preserved matters! In Saab's examples, B simply denies talking in a certain way
  - But this can be taken to mean that they simply avoid producing certain forms, in which case the contrast between (9)/(10) and (11)/(12) is trivial and tells us nothing about whether there is no non-performative "bias" component in the target items above that is still preserved when said item is recovered during ellipsis resolution
  - Also, we can can situationally do things we are in general opposed to; cf.:

    (13) A: How is Pam doing?
    B: She's doing fine, as far as I know, but can you please stop asking me about her, you know I don't want to talk about her.

  - In a similar vein, you can also do something and attempt to mitigate the effect by adding a direct objection to doing that

Prior literature

Issues with Saab's claim:

- It's categorical, but the extent to which the "bias" of a given lexical item is preserved in a response to an antecedent utterance containing said item might very well be gradient and affected by a variety of factors
- The way we probe the extent to which said bias is preserved matters! In Saab's examples, B simply denies talking in a certain way
  - But this can be taken to mean that they simply avoid producing certain forms, in which case the contrast between (9)/(10) and (11)/(12) is trivial and tells us nothing about whether there is no non-performative "bias" component in the target items above that is still preserved when said item is recovered during ellipsis resolution
  - Also, we can can situationally do things we are in general opposed to; cf.:

    (13) A: How is Pam doing?
    B: She's doing fine, as far as I know, but can you please stop asking me about her, you know I don't want to talk about her.

  - In a similar vein, you can also do something and attempt to mitigate the effect by adding a direct objection to doing that

The bottom line: we need a more fine-grained empirical picture, which I sought to provide in my experiments

Paradigms of interest

Set-up (simplified): The exchanges are set in a fictional universe where humans co-exist with centaurs, dwarves, elves, orcs, etc. and happen in the context of a criminal investigation. A detective (human) is questioning different witnesses (all human).

Paradigms of interest

Set-up (simplified): The exchanges are set in a fictional universe where humans co-exist with centaurs, dwarves, elves, orcs, etc. and happen in the context of a criminal investigation. A detective (human) is questioning different witnesses (all human).

(14) *Context: 'Tusky' is a slur for orcs.*
Detective: Did you see a tusky?
Witness: Yes. ('Bare') / Yes, I did. ('VPE') / Yes, I saw one. ('One') / Yes, I saw a tusky. ('Slur') / Yes, I saw an orc. ('Nonslur')

Paradigms of interest

Set-up (simplified): The exchanges are set in a fictional universe where humans co-exist with centaurs, dwarves, elves, orcs, etc. and happen in the context of a criminal investigation. A detective (human) is questioning different witnesses (all human).

(14) *Context: 'Tusky' is a slur for orcs.*
Detective: Did you see a tusky?
Witness: Yes. ('Bare') / Yes, I did. ('VPE') / Yes, I saw one. ('One') / Yes, I saw a tusky. ('Slur') / Yes, I saw an orc. ('Nonslur')

(15) *Context: 'Tusky' is a slur for orcs. This slur can also be used as a verb meaning 'to crawl' (for any race), because orcs are stereotyped as living in caves and, thus, having to crawl through narrow spaces all the time. The detective is asking a question about a human.*
Detective: What happened next? Did he tusky under the table?
Witness: Yes. ('Bare') / Yes, he did. ('VPE') / Yes, he did so. ('So') / Yes, he tuskied under the table. ('Slur') / Yes, he crawled under the table. ('Nonslur')

Paradigms of interest

Set-up (simplified): The exchanges are set in a fictional universe where humans co-exist with centaurs, dwarves, elves, orcs, etc. and happen in the context of a criminal investigation. A detective (human) is questioning different witnesses (all human).

(14) *Context: 'Tusky' is a slur for orcs.*
Detective: Did you see a tusky?
Witness: Yes. ('Bare') / Yes, I did. ('VPE') / Yes, I saw one. ('One') / Yes, I saw a tusky. ('Slur') / Yes, I saw an orc. ('Nonslur')

(15) *Context: 'Tusky' is a slur for orcs. This slur can also be used as a verb meaning 'to crawl' (for any race), because orcs are stereotyped as living in caves and, thus, having to crawl through narrow spaces all the time. The detective is asking a question about a human.*
Detective: What happened next? Did he tusky under the table?
Witness: Yes. ('Bare') / Yes, he did. ('VPE') / Yes, he did so. ('So') / Yes, he tuskied under the table. ('Slur') / Yes, he crawled under the table. ('Nonslur')

Question: How likely do you think that this witness is prejudiced against orcs?

Hypothesis and predictions

1. **Slurs do have performative effects**, so the prejudice likelihood is highest when the witness utters the slur themselves ('Slur')

Hypothesis and predictions

1. **Slurs do have performative effects**, so the prejudice likelihood is highest when the witness utters the slur themselves ('Slur')
2. **The prejudice likelihood is lowest when the witness indirectly challenges the detective by using the neutral term** instead ('Nonslur'), in an attempt to minimize complicity (see, e.g., Cepollaro 2020 and refs therein)

Hypothesis and predictions

1. **Slurs do have performative effects**, so the prejudice likelihood is highest when the witness utters the slur themselves ('Slur')
2. **The prejudice likelihood is lowest when the witness indirectly challenges the detective by using the neutral term** instead ('Nonslur'), in an attempt to minimize complicity (see, e.g., Cepollaro 2020 and refs therein)
3. **The prejudice component of slurs is not exclusively performative** (allows for use w/o mention). So, when the slur itself is obligatorily recovered during ellipsis resolution, the prejudice likelihood is higher than when it isn't

Hypothesis and predictions

1. **Slurs do have performative effects**, so the prejudice likelihood is highest when the witness utters the slur themselves ('Slur')
2. **The prejudice likelihood is lowest when the witness indirectly challenges the detective by using the neutral term** instead ('Nonslur'), in an attempt to minimize complicity (see, e.g., Cepollaro 2020 and refs therein)
3. **The prejudice component of slurs is not exclusively performative** (allows for use w/o mention). So, when the slur itself is obligatorily recovered during ellipsis resolution, the prejudice likelihood is higher than when it isn't
    - Further assumption: abstract lexical identity (in the sense of Harley 2014) is only required for the "main stem" of the constituent targeted by *one*-replacement/VPE/*do so*-replacement; only this structure needs to be lexically recovered for resolving the relevant anaphor (*one* or *do*), while some other form of identity (e.g., properly constrained form of truth-conditional identity) is sufficient for the entire constituent.
        - For noun slurs, we only require stem identity in *one*-replacement, but not in VPE or bare particle responses
        - For verb slurs, we only require stem identity in VPE/*do so*-replacement, but not in bare particle responses

Hypothesis and predictions

1. **Slurs do have performative effects**, so the prejudice likelihood is highest when the witness utters the slur themselves ('Slur')
2. **The prejudice likelihood is lowest when the witness indirectly challenges the detective by using the neutral term** instead ('Nonslur'), in an attempt to minimize complicity (see, e.g., Cepollaro 2020 and refs therein)
3. **The prejudice component of slurs is not exclusively performative** (allows for use w/o mention). So, when the slur itself is obligatorily recovered during ellipsis resolution, the prejudice likelihood is higher than when it isn't
   - Further assumption: abstract lexical identity (in the sense of Harley 2014) is only required for the "main stem" of the constituent targeted by *one*-replacement/VPE/*do so*-replacement; only this structure needs to be lexically recovered for resolving the relevant anaphor (*one* or *do*), while some other form of identity (e.g., properly constrained form of truth-conditional identity) is sufficient for the entire constituent.
     - For noun slurs, we only require stem identity in *one*-replacement, but not in VPE or bare particle responses
     - For verb slurs, we only require stem identity in VPE/*do so*-replacement, but not in bare particle responses

(16) *Predicted prejudice likelihood ratings (from lowest to highest)*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur'

b. Verbs: 'Nonslur' < 'Bare' < 'VPE'/'So' < 'Slur'

Additional considerations

- I also expected verb slurs to be harder to judge than noun slurs (as there are few, if any counterparts of such slurs in English—but Elin McCready in a p.c. pointed out the existence of the verb *to n*****-lip* to me) and possibly less "offensive" (due to the less direct link b/n their meaning and the targeted group), so no direct comparison of nouns and verbs was planned

Additional considerations

- I also expected verb slurs to be harder to judge than noun slurs (as there are few, if any counterparts of such slurs in English—but Elin McCready in a p.c. pointed out the existence of the verb *to n*****-lip* to me) and possibly less "offensive" (due to the less direct link b/n their meaning and the targeted group), so no direct comparison of nouns and verbs was planned
- I also conjectured that for some people shorter responses might be another strategy to minimize complicity, which might introduce further gradient distinctions across 'Bare' vs. 'VPE' vs. 'One/So' and potentially obscure some of the contrasts predicted by the hypothesis

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)
- Each participant saw 2 trials per condition + 2 attention checks (22 trials total)

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)
- Each participant saw 2 trials per condition + 2 attention checks (22 trials total)
- Instructions described the universe, the general context of the exchanges, and explicitly said that, regardless of their views, the witnesses did not feel comfortable openly challenging the detective on their use of the slur

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)
- Each participant saw 2 trials per condition + 2 attention checks (22 trials total)
- Instructions described the universe, the general context of the exchanges, and explicitly said that, regardless of their views, the witnesses did not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)
- Each participant saw 2 trials per condition + 2 attention checks (22 trials total)
- Instructions described the universe, the general context of the exchanges, and explicitly said that, regardless of their views, the witnesses did not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'
- Participants were recruited on Prolific (final $N = 128$) and paid £1.25 for completing the task

Methods

- 10 conditions (or rather 2 sub-experiments for nouns and verbs, w/5 response type conditions for each)
- Items similar to (14) and (15) (4 nonslur–slur pairs: centaur–hoofy, dwarf–stunty, elf–leafy, orc–tusky)
- Each participant saw 2 trials per condition + 2 attention checks (22 trials total)
- Instructions described the universe, the general context of the exchanges, and explicitly said that, regardless of their views, the witnesses did not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'
- Participants were recruited on Prolific (final $N = 128$) and paid £1.25 for completing the task
- Basic sociodemographic info about participants was collected, as well (age, gender, ethnicity)

Typical trial

progress

Context: 'Stunty' is a slur for dwarves. This slur can also be used as a verb meaning 'to fall' (for any race), because dwarves are stereotyped as clumsy and, thus, prone to falling. The detective is asking a question about a human.

Detective: What happened next? Did he stunty down the stairs?
Witness: Yes, he stuntied down the stairs.

Question: How likely do you think that this witness is prejudiced against dwarves?

Not at all likely  Very likely

Confirm my answer

Results

(17) *Predicted prejudice likelihood ratings*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur'

b. Verbs: 'Nonslur' < 'Bare' < 'VPE'/'So' < 'Slur'

Results

(17) *Predicted prejudice likelihood ratings*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur'

b. Verbs: 'Nonslur' < 'Bare' < 'VPE'/'So' < 'Slur'

(18) *Statistically significant contrasts in prejudice likelihood ratings*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur' — full match

b. Verbs: 'Nonslur' < 'Bare'/'VPE'/'So' < 'Slur' — partial match

Results

(17) *Predicted prejudice likelihood ratings*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur'

b. Verbs: 'Nonslur' < 'Bare' < 'VPE'/'So' < 'Slur'

(18) *Statistically significant contrasts in prejudice likelihood ratings*

a. Nouns: 'Nonslur' < 'Bare'/'VPE' < 'One' < 'Slur' — full match

b. Verbs: 'Nonslur' < 'Bare'/'VPE'/'So' < 'Slur' — partial match

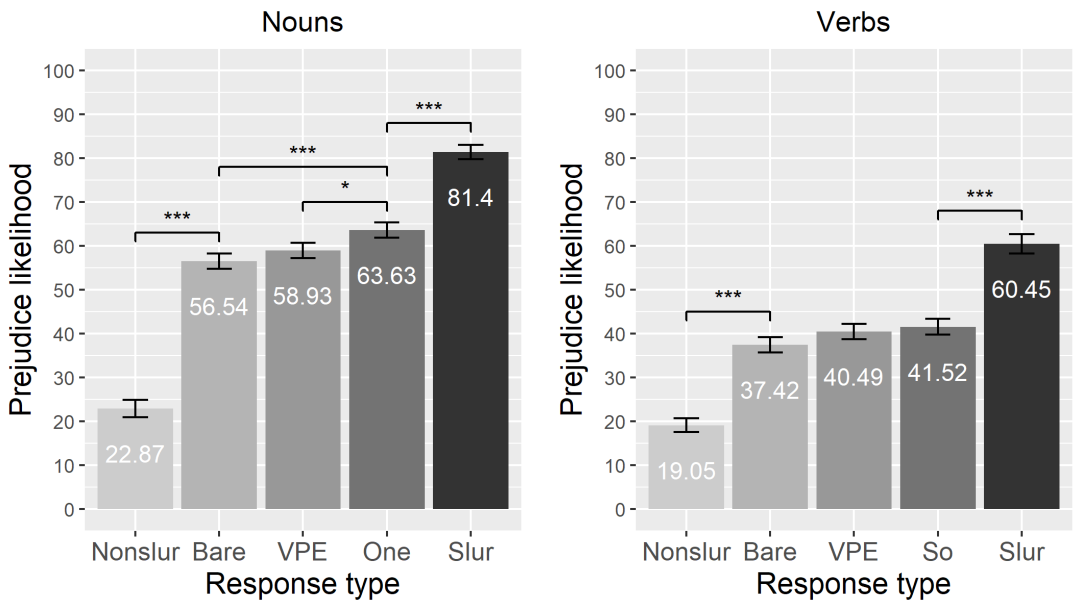

Fig. 1: Mean prejudice likelihood ratings of different types of responses to antecedent utterances with fictional noun and verb slurs. Crucial significant contrasts are indicated. Error bars show standard error.

Replicating the results of Experiment 1 for real-life slurs

One of the concerns you might be having: is using fictional slurs optimal for learning about real-life slurs?

Experiment 2: attempt to replicate Experiment 1 (for nouns only) with real-life slurs

Methods

General set-up:

The exchanges happen in the context of a criminal investigation. The detective (White non-Jewish straight cis man) is questioning different witnesses (all White non-Jewish straight cis men).

Methods

General set-up:

The exchanges happen in the context of a criminal investigation. The detective (White non-Jewish straight cis man) is questioning different witnesses (all White non-Jewish straight cis men).

Sample paradigm:

(19) Detective: Was there a SLUR at the bar?
Witness: Yes. ('Bare') / Yes, there was. ('VPE') / Yes, there was one. ('One') / Yes, there was a SLUR at the bar. ('Slur') / Yes, there was a NONSLUR at the bar. ('Nonslur')

Methods

General set-up:

The exchanges happen in the context of a criminal investigation. The detective (White non-Jewish straight cis man) is questioning different witnesses (all White non-Jewish straight cis men).

Sample paradigm:

(19) Detective: Was there a SLUR at the bar?
Witness: Yes. ('Bare') / Yes, there was. ('VPE') / Yes, there was one. ('One') / Yes, there was a SLUR at the bar. ('Slur') / Yes, there was a NONSLUR at the bar. ('Nonslur')

+ **Controls:**

(20) Detective: Who brought the SLUR?
Witness: Alex. ('ControlShort') / Anna did. ('ControlLong')

Methods

- 5 conditions for the different response types: ‘Bare’, ‘VPE’, ‘One’, ‘Slur’, ‘Nonslur’

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items
- Each participant saw 2 trials per condition (10 experimental items + 4 controls) + 2 attention checks (16 trials total)

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items
- Each participant saw 2 trials per condition (10 experimental items + 4 controls) + 2 attention checks (16 trials total)
- Instructions described the general context of the exchanges and explicitly said that, regardless of their views, the witnesses might not feel comfortable openly challenging the detective on their use of the slur

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items
- Each participant saw 2 trials per condition (10 experimental items + 4 controls) + 2 attention checks (16 trials total)
- Instructions described the general context of the exchanges and explicitly said that, regardless of their views, the witnesses might not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items
- Each participant saw 2 trials per condition (10 experimental items + 4 controls) + 2 attention checks (16 trials total)
- Instructions described the general context of the exchanges and explicitly said that, regardless of their views, the witnesses might not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'
- After completing the 16 trials, participants were also asked to rate the 5 slurs above on their offensiveness (0–100) and answer sociodemographic questions

Methods

- 5 conditions for the different response types: 'Bare', 'VPE', 'One', 'Slur', 'Nonslur'
- Items:
  - 5 slurs: *c***** (Chinese person), *f****** (gay man), *k**** (Jewish person), *n****** (Black person), *t****** (transgender woman)
  - 5 question types
  - 25 experimental items + 10 control items
- Each participant saw 2 trials per condition (10 experimental items + 4 controls) + 2 attention checks (16 trials total)
- Instructions described the general context of the exchanges and explicitly said that, regardless of their views, the witnesses might not feel comfortable openly challenging the detective on their use of the slur
- The prejudice likelihood was assessed by dragging a slider on a pseudo-continuous scale (mapped to 0–100) from 'Not at all likely' to 'Very likely'
- After completing the 16 trials, participants were also asked to rate the 5 slurs above on their offensiveness (0–100) and answer sociodemographic questions
- Participants were recruited on Prolific (final $N = 85$) and paid £1.5 for completing the task

Typical trial

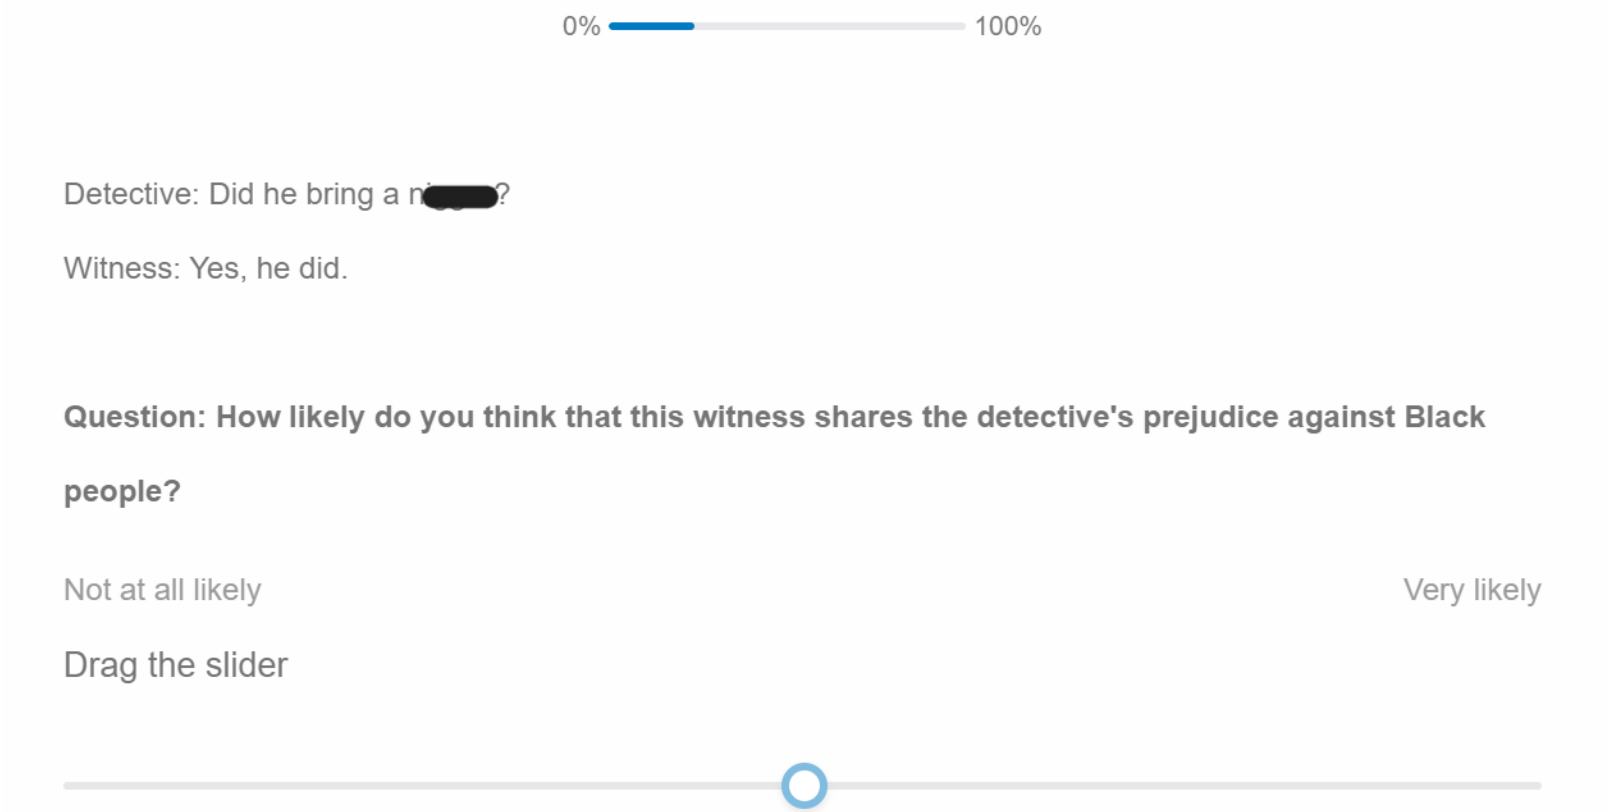

Results

The results were very similar to what we saw for fictional noun slurs + no significant difference between short and long controls:

Results

The results were very similar to what we saw for fictional noun slurs + no significant difference between short and long controls:

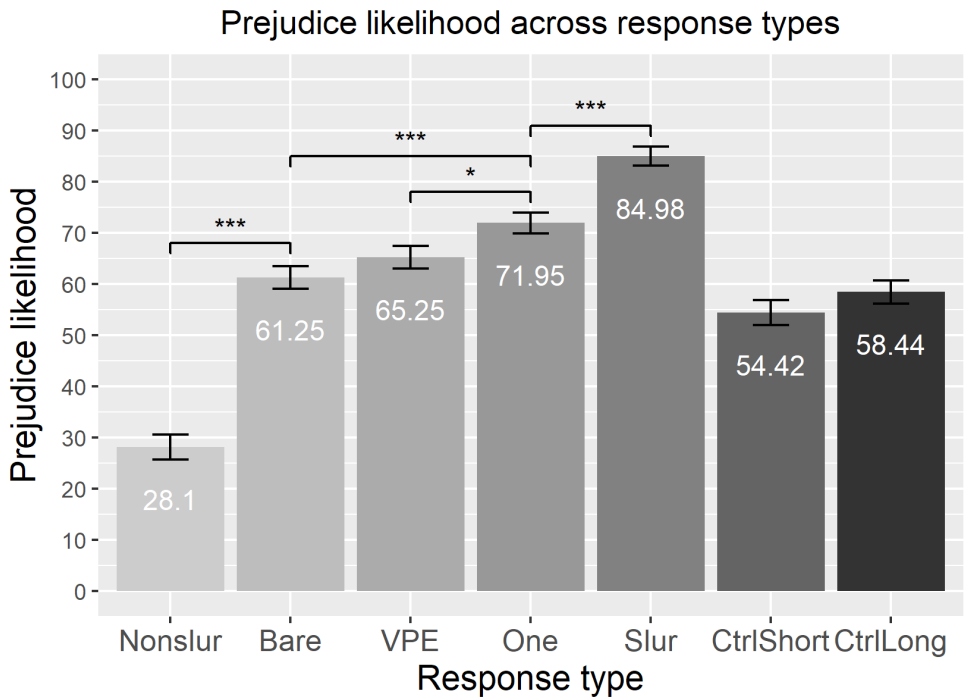

Fig. 2: Mean prejudice likelihood ratings of different types of responses to antecedent utterances with real-life noun slurs. Crucial significant contrasts are indicated. Error bars show standard error.

Results

The results were very similar to what we saw for fictional noun slurs + no significant difference between short and long controls:

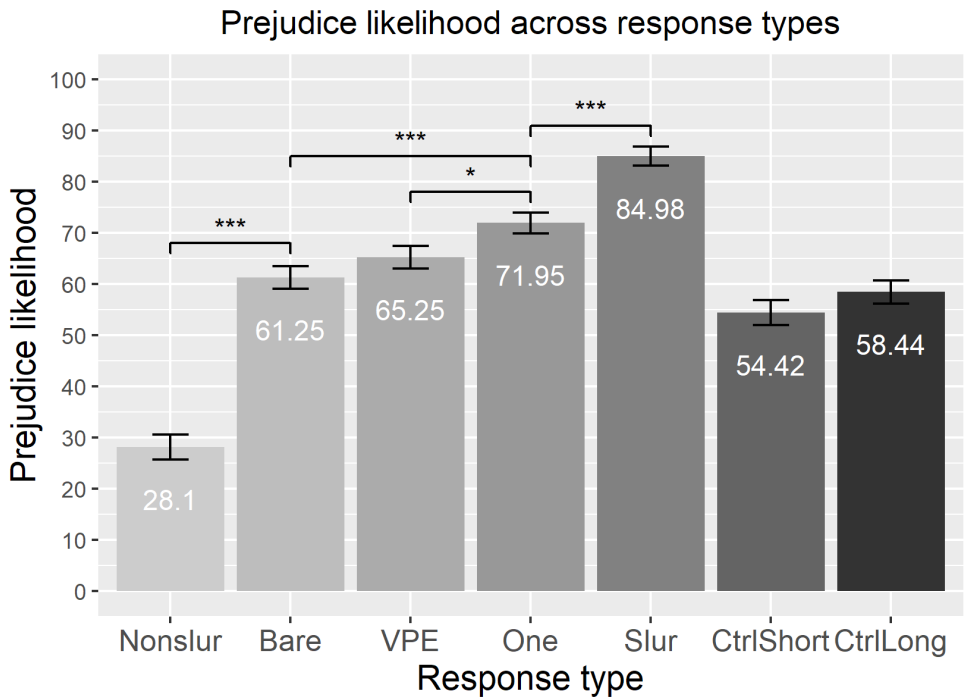

Fig. 2: Mean prejudice likelihood ratings of different types of responses to antecedent utterances with real-life noun slurs. Crucial significant contrasts are indicated. Error bars show standard error.

No main effect of slur type, but here are the offensiveness ratings (0–100): *c***** (Chinese person): 73.7, *f****** (gay man): 82.1, *k**** (Jewish person): 77.2, *n****** (Black person): 85.8, *t****** (transgender woman): 72.2

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Implications:

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Implications:

- **The prejudice component of slurs is partially performative** (like the expressive component of *fucking*, unlike the presupposition of *stop*), **but it is not exclusively performative** (unlike the expressive component of *fucking*, like the presupposition of *stop*)

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Implications:

- **The prejudice component of slurs is partially performative** (like the expressive component of *fucking*, unlike the presupposition of *stop*), **but it is not exclusively performative** (unlike the expressive component of *fucking*, like the presupposition of *stop*)
  - Saab 2020 (about slurs, a.o.): "ellipsis is an apt strategy to nullify the bias encoded in some lexical items"—My results suggest that ellipsis attenuates said bias (by avoiding the performative effects of saying the slur), but doesn't completely nullify it (when the slur itself is recovered)

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Implications:

- **The prejudice component of slurs is partially performative** (like the expressive component of *fucking*, unlike the presupposition of *stop*), **but it is not exclusively performative** (unlike the expressive component of *fucking*, like the presupposition of *stop*)
  - Saab 2020 (about slurs, a.o.): "ellipsis is an apt strategy to nullify the bias encoded in some lexical items"—My results suggest that ellipsis attenuates said bias (by avoiding the performative effects of saying the slur), but doesn't completely nullify it (when the slur itself is recovered)
- This calls for a **hybrid analysis for the prejudice component of slurs** that doesn't reduce it to just some kind of truth-conditional, but not-at-issue content (conventional implicature in Potts 2005; presupposition in Schlenker 2007)—or just a performative effect of uttering a certain form on the context (as in Potts 2007 or Saab 2020)

Discussion: noun slurs

The results for noun slurs corroborate all parts of the original hypothesis

Implications:

- **The prejudice component of slurs is partially performative** (like the expressive component of *fucking*, unlike the presupposition of *stop*), **but it is not exclusively performative** (unlike the expressive component of *fucking*, like the presupposition of *stop*)
  - Saab 2020 (about slurs, a.o.): "ellipsis is an apt strategy to nullify the bias encoded in some lexical items"—My results suggest that ellipsis attenuates said bias (by avoiding the performative effects of saying the slur), but doesn't completely nullify it (when the slur itself is recovered)
- This calls for a **hybrid analysis for the prejudice component of slurs** that doesn't reduce it to just some kind of truth-conditional, but not-at-issue content (conventional implicature in Potts 2005; presupposition in Schlenker 2007)—or just a performative effect of uttering a certain form on the context (as in Potts 2007 or Saab 2020)
  - Note: I am not talking about separating the denotational meaning of slurs from their attitudinal meaning, but about the attitudinal component itself being of a mixed nature

Potts 2005 vs. Potts 2007

Potts 2005 and Potts 2007 are not the same!

Potts 2005 vs. Potts 2007

Potts 2005 and Potts 2007 are not the same!

Potts 2005:

Potts 2005 vs. Potts 2007

Potts 2005 and Potts 2007 are not the same!

Potts 2005:

- **Conventional implicatures (CIs)** are truth-based (two types of truth values, at-issue and CI)

Potts 2005 vs. Potts 2007

Potts 2005 and Potts 2007 are not the same!

Potts 2005:

- **Conventional implicatures (CIs)** are truth-based (two types of truth values, at-issue and CI)
- Expressives as in (22) are analyzed as CIs; same as, e.g., appositives as in (21):

(21) I didn't see any orcs, whom, by the way, I find despicable.
At-issue: I didn't see any orcs.
CI: I find orcs despicable.

(22) I didn't see any fucking orcs.
At-issue: I didn't see any orcs.
CI: I have a negative attitude towards orcs.
(NB: Once again, expressives don't actually have to convey affect towards the denotation of the thing they apparently compose with—or, in fact, anything else in the sentence)

Potts 2005 vs. Potts 2007

Potts 2007:

- Acknowledges that expressives are very different from supplements (appositives, sentence-level adverbs, parentheticals)
- One difference he discusses: "ineffability" (inability to be rephrased using truth-conditional content; *Fuck!* $\neq$ *I am angry*)—in our terms, this is the consequence of their performativity
- Formal insight: expressives directly alter a specific parameter of the context of interpretation
  - Contexts are tuples that store various info, like the speaker, the addressee, the time/place of the utterance, etc.
  - The expressive parameter c_ε is essentially a ledger that tracks the emotional states of the conversation participants of the input context of interpretation c
  - Confusingly, Potts 2007 uses the same notation, namely the bullet operator, as Potts 2005, but it does different things in the two cases!

(23) Expressive semantics in Potts 2007 (simplified)

a. $[\![\textit{expr}_{\langle\sigma,\varepsilon\rangle}]\!]^{c'}([\![\alpha_\sigma]\!]^c)(c) = c''$, where c'' is just like c, except $\text{feels}(c''_s, c''_\varepsilon)$; σ is a truth-conditional type; and ε is an expressive type

b. $[\![\alpha_{\langle\sigma,\varepsilon\rangle}]\!]^{c'} \bullet_\varepsilon [\![\beta_\sigma]\!]^c = [\![\beta]\!]^{[\![\alpha]\!]^{c'}([\![\beta]\!]^c)(c)}$

c. $[\![\textit{fucking}_{\langle et,\varepsilon\rangle}]\!]^{c'} \bullet_\varepsilon [\![\textit{orcs}_{et}]\!]^c = [\![\textit{orcs}]\!]^{c''}$, where c'' is just like c, except $\text{feels}(c''_s, c''_\varepsilon)$

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

I remain neutral on the details, but:

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

I remain neutral on the details, but:

- Such an analysis can only be instantiated in **a system that properly separates performative contributions from truth-conditional ones** (e.g., Potts 2007; Saab 2020; not Potts 2005; Schlenker 2007)

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

I remain neutral on the details, but:

- Such an analysis can only be instantiated in **a system that properly separates performative contributions from truth-conditional ones** (e.g., Potts 2007; Saab 2020; not Potts 2005; Schlenker 2007)
- E.g.:

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

I remain neutral on the details, but:

- Such an analysis can only be instantiated in **a system that properly separates performative contributions from truth-conditional ones** (e.g., Potts 2007; Saab 2020; not Potts 2005; Schlenker 2007)
- E.g.:
  - Performative aspect of the prejudice component: Potts 2007-style context altering

Hybrid analysis of slurs

A hybrid account of the prejudice component of slurs:

- Part of it is **performative** (tied to the act of saying the slur; non-truth-based; disappears under ellipsis)
- Part of it is **truth-conditional, but not-at-issue** (preserved under ellipsis when we require lexical identity for the slur itself)

I remain neutral on the details, but:

- Such an analysis can only be instantiated in **a system that properly separates performative contributions from truth-conditional ones** (e.g., Potts 2007; Saab 2020; not Potts 2005; Schlenker 2007)
- E.g.:
  - Performative aspect of the prejudice component: Potts 2007-style context altering
  - Truth-conditional, but not-at-issue aspect of the prejudice component: e.g., presupposition or Potts 2005-style CI

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

- Perhaps the identity requirements for VPE and *do so*-replacement are not the same as for *one*-replacement?

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

- Perhaps the identity requirements for VPE and *do so*-replacement are not the same as for *one*-replacement?
- Perhaps *tusky* as a verb is not assumed to be truly co-extensive with *crawl* (cf. *tusky* as a noun and *orc*)

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

- Perhaps the identity requirements for VPE and *do so*-replacement are not the same as for *one*-replacement?
- Perhaps *tusky* as a verb is not assumed to be truly co-extensive with *crawl* (cf. *tusky* as a noun and *orc*)
- The contrasts were overall less pronounced for verb slurs:

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

- Perhaps the identity requirements for VPE and *do so*-replacement are not the same as for *one*-replacement?
- Perhaps *tusky* as a verb is not assumed to be truly co-extensive with *crawl* (cf. *tusky* as a noun and *orc*)
- The contrasts were overall less pronounced for verb slurs:
  - In the absence of perfect English counterparts (and w/a very bare-bones explanation of use), the fictional verb slurs were harder to intuit about

Discussion: verb slurs

The results for verb slurs corroborate parts 1 and 2, but not 3 of the hypothesis. A few relevant considerations:

- Perhaps the identity requirements for VPE and *do so*-replacement are not the same as for *one*-replacement?
- Perhaps *tusky* as a verb is not assumed to be truly co-extensive with *crawl* (cf. *tusky* as a noun and *orc*)
- The contrasts were overall less pronounced for verb slurs:
  - In the absence of perfect English counterparts (and w/a very bare-bones explanation of use), the fictional verb slurs were harder to intuit about
  - Less direct link between the meaning of the slur and the targeted group, hence a lower upper bound for the ratings

Some limitations of the present experiments

- Pretty charged context (by choice), seemingly making the contrasts b/n challenged vs. unchallenged slurs more salient than any other contrasts

Some limitations of the present experiments

- Pretty charged context (by choice), seemingly making the contrasts b/n challenged vs. unchallenged slurs more salient than any other contrasts
- The question was about prejudice likelihood rather than intensity (by choice), although participants could be assessing both

Some limitations of the present experiments

- Pretty charged context (by choice), seemingly making the contrasts b/n challenged vs. unchallenged slurs more salient than any other contrasts
- The question was about prejudice likelihood rather than intensity (by choice), although participants could be assessing both
- Is asking about prejudice inference the best way to probe preservation of the prejudice component of slurs? Not at all, but neither are continuations denying being prejudiced (one can use slurs and deny being prejudiced) or expressing one's general opposition to uttering a slur (recall our objections to Saab's examples)

Some limitations of the present experiments

- I still think response length or perhaps size of the elided constituent containing the slur plays a role; we're seeing a relevant trend, across the board, it just hasn't reached significance. My Experiment 2 controls don't fully address the issue and are too similar to the already existing 'Bare' vs. 'VPE' conditions; I considered using the following controls, but was worried that exposing participants to numerals will make the number interpretation of *one*-replacement more salient:

Some limitations of the present experiments

- I still think response length or perhaps size of the elided constituent containing the slur plays a role; we're seeing a relevant trend, across the board, it just hasn't reached significance. My Experiment 2 controls don't fully address the issue and are too similar to the already existing 'Bare' vs. 'VPE' conditions; I considered using the following controls, but was worried that exposing participants to numerals will make the number interpretation of *one*-replacement more salient:

  (24) Detective: How many SLURS were there?
  Witness: Three. (short) / There were three. (long)

Some limitations of the present experiments

- I still think response length or perhaps size of the elided constituent containing the slur plays a role; we're seeing a relevant trend, across the board, it just hasn't reached significance. My Experiment 2 controls don't fully address the issue and are too similar to the already existing 'Bare' vs. 'VPE' conditions; I considered using the following controls, but was worried that exposing participants to numerals will make the number interpretation of *one*-replacement more salient:

  (24) Detective: How many SLURS were there?
  Witness: Three. (short) / There were three. (long)

- Quality of data from Prolific workers

(Other) potential follow-ups

- Other syntactic configurations (e.g., *–Did you see a picture of a tusky? –Yes, I saw one. / Yes, I saw a picture of one.*; *–Is she a tusky? –Yes, she is. / Yes, she is one.*)

(Other) potential follow-ups

- Other syntactic configurations (e.g., *–Did you see a picture of a tusky? –Yes, I saw one. / Yes, I saw a picture of one.*; *–Is she a tusky? –Yes, she is. / Yes, she is one.*)
- Other types of content, e.g., stylistic variants (see, e.g., Saab 2020 for Spanish *morfar* vs. *comer*), but again, the concerns about true co-extensiveness apply

(Other) potential follow-ups

- Other syntactic configurations (e.g., *–Did you see a picture of a tusky? –Yes, I saw one. / Yes, I saw a picture of one.*; *–Is she a tusky? –Yes, she is. / Yes, she is one.*)
- Other types of content, e.g., stylistic variants (see, e.g., Saab 2020 for Spanish *morfar* vs. *comer*), but again, the concerns about true co-extensiveness apply
- Other languages, in particular, those with other types of fragment responses and ellipsis in general (e.g., Russian)

Some final words

- Inferences about prejudice are an empirically messy phenomenon, with many factors affecting them, so we need to be careful about making categorical empirical claims about them, let alone drawing theoretical conclusions about the semantics of slurs from said claims

Some final words

- Inferences about prejudice are an empirically messy phenomenon, with many factors affecting them, so we need to be careful about making categorical empirical claims about them, let alone drawing theoretical conclusions about the semantics of slurs from said claims
- When looking at slurs under ellipsis, we are essentially trying to use two phenomena we don't understand well to explain one another, which is a reason to be extra careful

Thanks!

References I

Cepollaro, Bianca. 2020. *Slurs and thick terms: When language encodes values*. Lexington Books.

Esipova, Maria. 2019. *Composition and projection in speech and gesture*: New York University dissertation. https://ling.auf.net/lingbuzz/004676.

Esipova, Maria. 2021. On not-at-issueness in pictures. *Glossa: A Journal of General Linguistics* 6(1). 83. doi:10.5334/gjgl.1314.

Esipova, Maria. 2024. Composure and composition. Ms., under revision. https://ling.auf.net/lingbuzz/005003.

Harley, Heidi. 2014. On the identity of roots. *Theoretical Linguistics* 40(3-4). 225–276. doi:doi:10.1515/tl-2014-0010.

Orlando, Eleonora & Andrés Saab (eds.). 2020. *Slurs and expressivity*. Lexington Books.

Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780199273829.001.0001.

Potts, Christopher. 2007. The expressive dimension. *Theoretical linguistics* 33(2). 165–198. doi:10.1515/TL.2007.011.

Roberts, Craige & Mandy Simons. 2023. Preconditions and projection: Explaining non-anaphoric presupposition. Ms. https://ling.auf.net/lingbuzz/007406.

Saab, Andrés. 2020. On the locus of expressivity. Deriving parallel meaning dimensions from architectural considerations. In Eleonora Orlando & Andrés Saab (eds.), *Slurs and expressivity*, 17–44. Lexington Books.

Sailor, Craig & Valentina Colasanti. 2020. Co-speech gestures under ellipsis: a first look. Talk given at *The 94th Annual Meeting of the LSA*, New Orleans.

Schlenker, Philippe. 2007. Expressive presuppositions. *Theoretical Linguistics* 33(2). 237–245. doi:10.1515/TL.2007.017.