

Gender Bias in the Production and Comprehension of Pronominal Referents

Veronica Boyce



Implicit gender bias in linguistic descriptions for expected events: The cases of the 2016 US and 2017 UK election

AUTHORS

Titus von der Malsburg, Till Poppels, Roger Levy


CREATED ON

August 02, 2018

LAST EDITED

December 13, 2018

SUPPLEMENTAL MATERIALS

osf.io/gx5tr/ 

- ▶ Gender stereotype
- ▶ Up-to-date world knowledge
- ▶ Discourse

How do these sources of information jointly shape linguistic expectations?

The US presidential race 2016:

A large-scale natural experiment in which our collective beliefs about the world were being manipulated.



*"This time next year, someone else will be standing here in this very spot. And it's anybody's guess who **she** will be."*

Data collection:

- ▶ June 2016 – February 2017
- ▶ Amazon Mechanical Turk
- ▶ $N > 24,000$, single shot

Three components:

1. Event priors for female president (IV)
2. Production of pronouns (DV)
3. Comprehension of pronouns (DV)

Component 1: Perceived probability of female president

How likely do you think are the following candidates to become the next US president?

0% likely

0

25

50

75

100% likely


100

Donald Trump


Bernie Sanders

Hillary Clinton

Component 1: Perceived probability of female president



Component 1: Perceived probability of female president



Component 2: Pronoun production

*The next **US President** will be sworn into office in January 2017.*

Component 2: Pronoun production

*The next **US President** will be sworn into office in January 2017.*

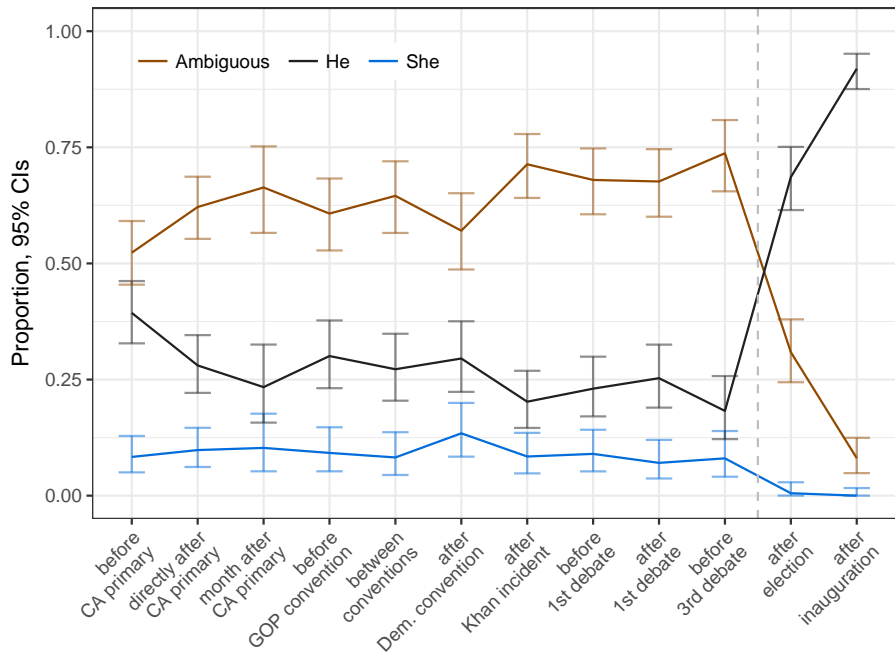
After moving into the Oval Office, one of the first things that ...

Component 2: Pronoun production

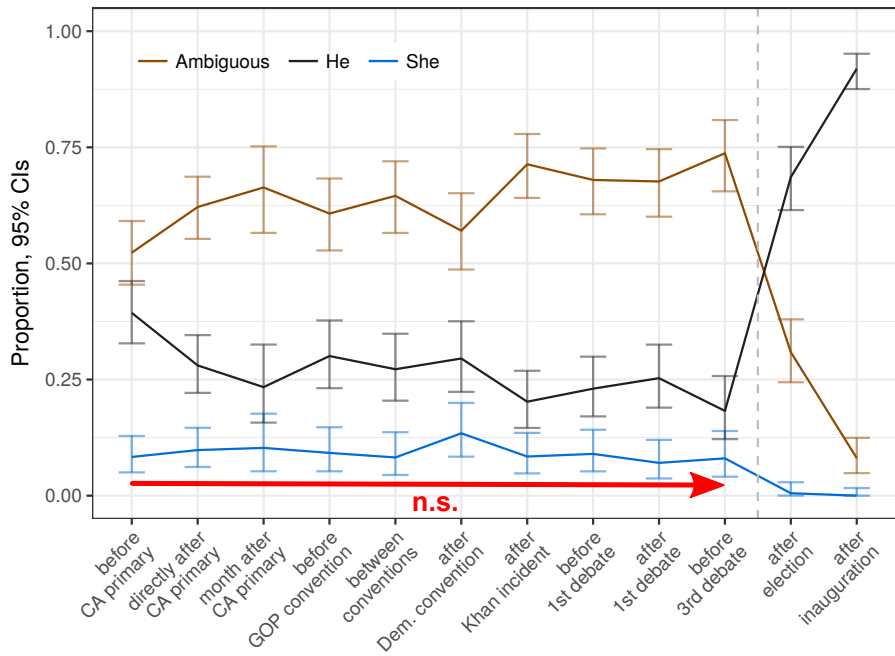
*The next **US President** will be sworn into office in January 2017.*

*After moving into the Oval Office, one of the first things that ... **he/she/they** will do is hold a staff briefing.*

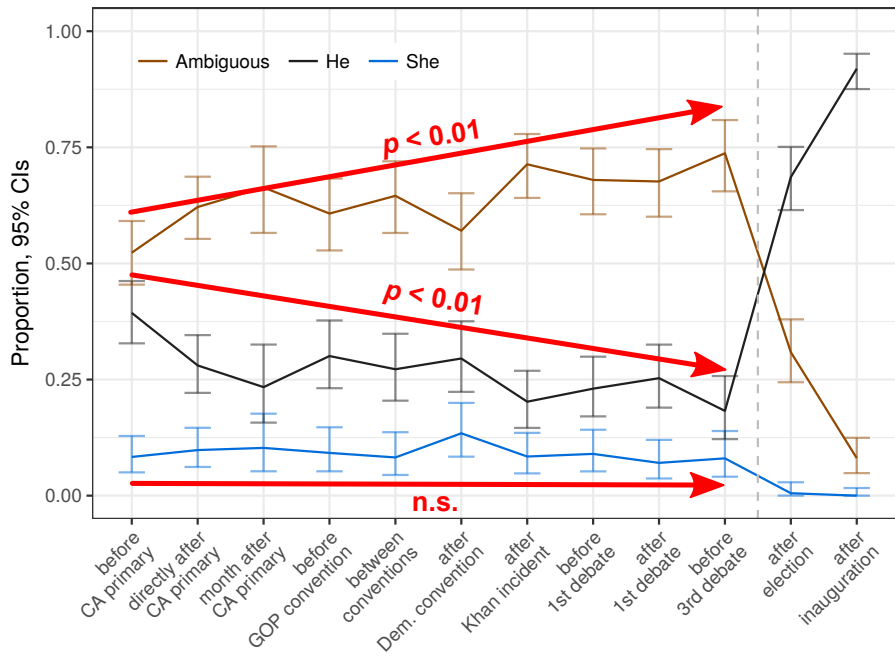
Component 2: Pronoun production



Component 2: Pronoun production



Component 2: Pronoun production



Component 3: Comprehension

*The next **US President** will be sworn into office in January 2017.*

*After moving into the Oval Office, one of the first things that **he/she/they** will do is hold a staff briefing.*

*During the inauguration speech, the president will emphasize **his/her/their** commitment to resolve outstanding issues.*

Component 3: Comprehension

	Sentence 1	Sentence 2	Sentence 3	Discourse
a)	President	he	he	match
b)	President	he	she	mismatch
c)	President	she	she	match
d)	President	she	he	mismatch
e)	President	they	they	match

The _____

___ next _____

US

----- president -----

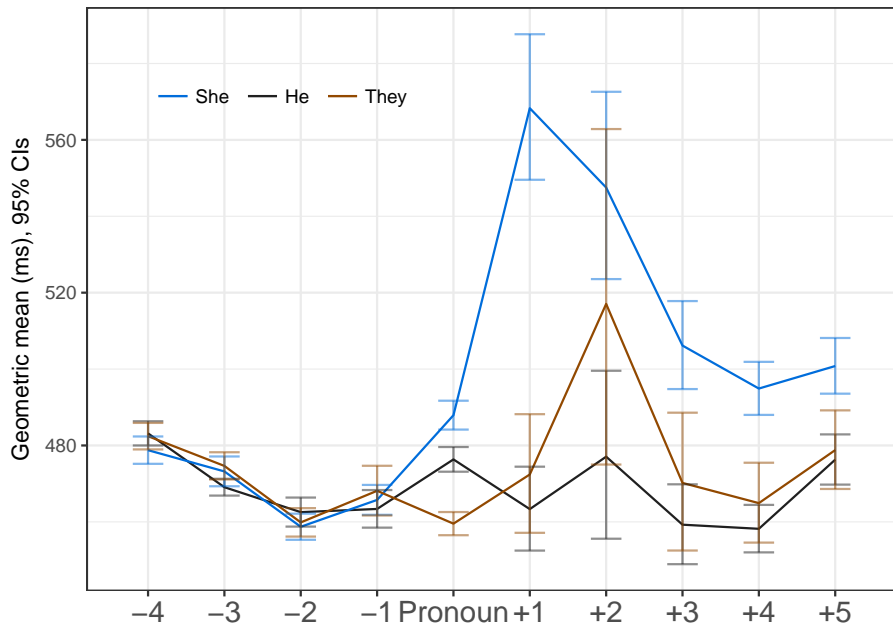
----- will -----

be

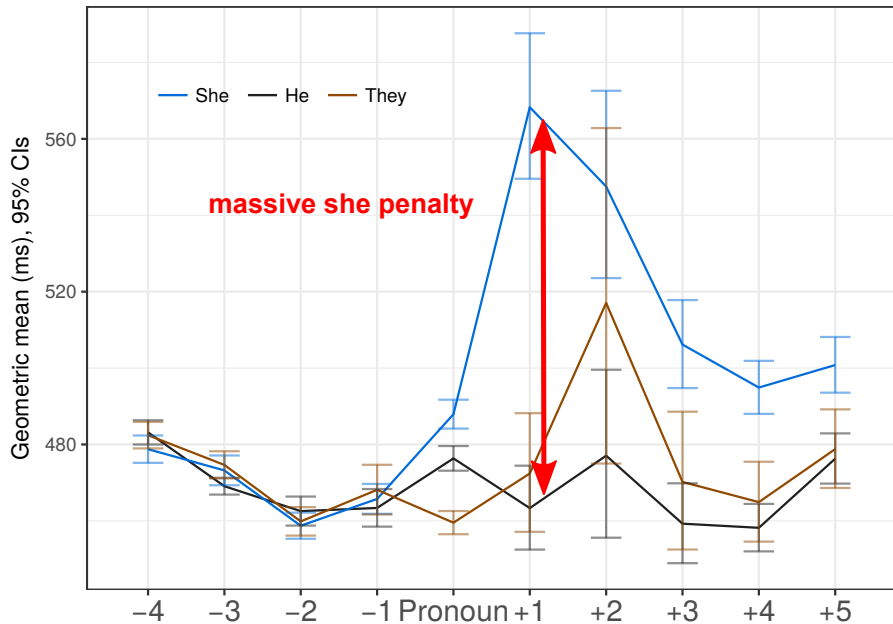
----- Sworn -----

----- into -----

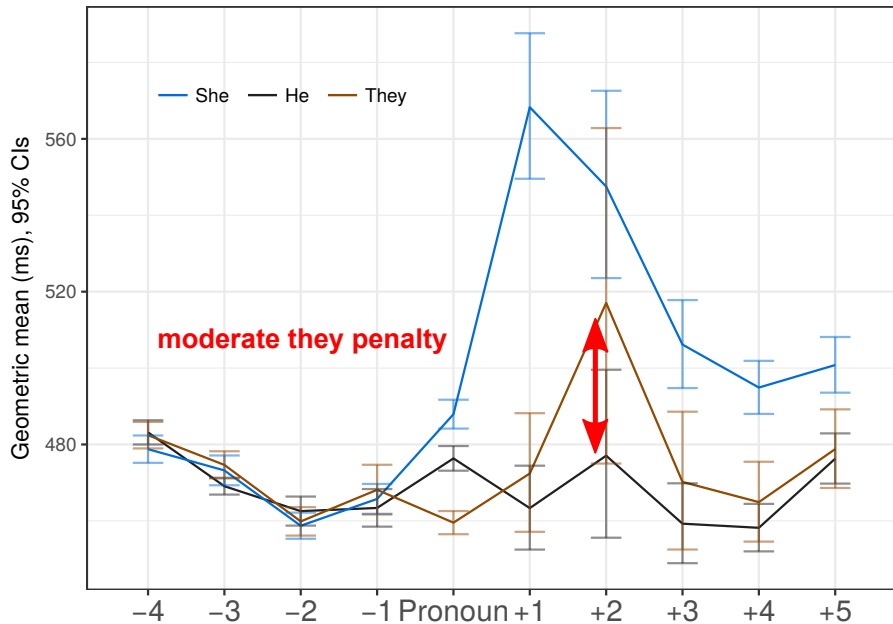
Reading times sentence 2 (first pronoun)




Reading times sentence 2 (first pronoun)



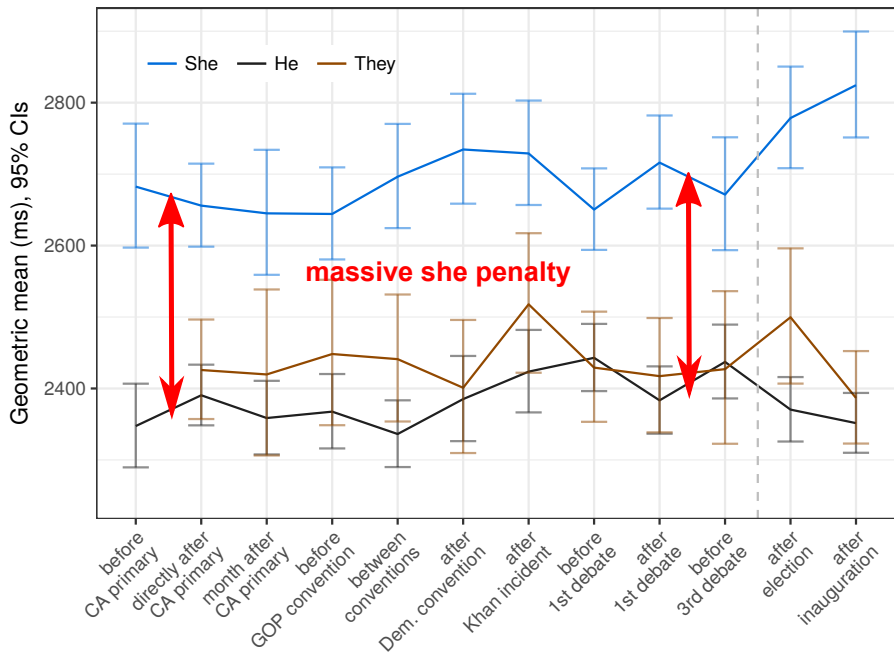
Reading times sentence 2 (first pronoun)




RTs sentence 2 over time (first pronoun + 4 words)



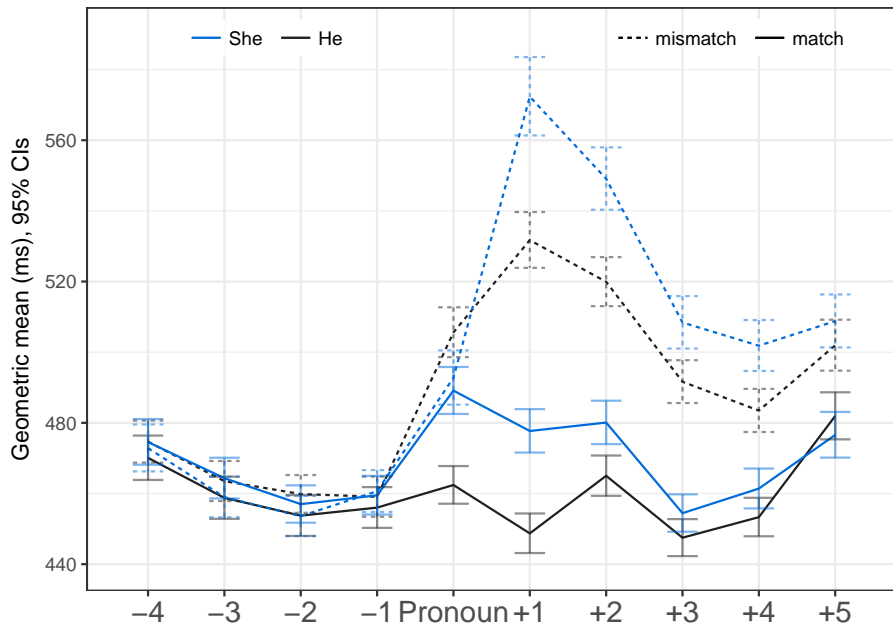
RTs sentence 2 over time (first pronoun + 4 words)



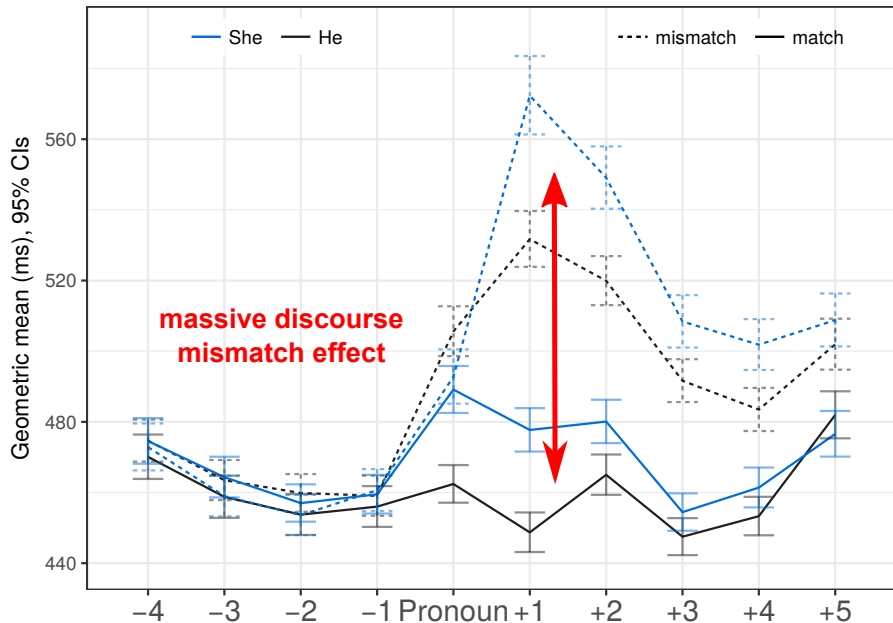
RTs sentence 2 over time (first pronoun + 4 words)



Reading times sentence 3 (second pronoun)



Reading times sentence 3 (second pronoun)



Summary of President study:

Even as female gender is judged more likely than male...

- In comprehension, slower reading times for she vs he or they
- In production, he > she, but most pronouns hedged (they)

What do comprehenders do?

What do comprehenders do?

Two possibilities:

What do comprehenders do?


Two possibilities:


1. Comprehenders “reverse” the production asymmetry in interpretation
 - ‘He’ is a weaker signal of referent gender than ‘She’
 - Production pattern not necessarily a concern

What do comprehenders do?


Two possibilities:

1. Comprehenders “reverse” the production asymmetry in interpretation
 - ‘He’ is a weaker signal of referent gender than ‘She’
 - Production pattern not necessarily a concern
2. Comprehenders do **not** “reverse” the production asymmetry in interpretation
 - ‘He’ is not a sufficiently weaker signal of referent gender than ‘She’
 - Female gender is systematically under-conveyed







Before the baker took out the trash, ...




Before the baker took out the trash, ...




Before the baker took out the trash, ...



Before the baker took out the trash, ...




Before the baker took out the trash, ...





Symmetric

Before the baker took out the trash, ... he swept the floor and wiped
down the counter.





Symmetric

Before the baker took out the trash, ... he swept the floor and wiped down the counter.





Symmetric

Before the baker took out the trash, ... he swept the floor and wiped down the counter.




Symmetric

Before the baker took out the trash, ... he swept the floor and wiped down the counter.




Symmetric

Before the baker took out the trash, ... he swept the floor and wiped down the counter.



Symmetric



Symmetric

Are speakers' referent gender expectations faithfully conveyed?

Are speakers' referent gender expectations faithfully conveyed?

$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$


Speaker \longrightarrow \longleftarrow Listener

Are speakers' referent gender expectations faithfully conveyed?


$$P_S(\text{ref. gender}) \stackrel{?}{=} P_L(\text{ref. gender})$$

Speaker →

← Listener



Symmetric




Symmetric

Are speakers' referent gender expectations faithfully conveyed?


$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$

Speaker →

← Listener



Symmetric



Symmetric




$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Are speakers' referent gender expectations faithfully conveyed?

$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$

Speaker

Listener




$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Are speakers' referent gender expectations faithfully conveyed?


$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$

Speaker →


← Listener



Symmetric



Symmetric




$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Are speakers' referent gender expectations faithfully conveyed?


$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$

Speaker →


← Listener



Symmetric



Symmetric






$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Are speakers' referent gender expectations faithfully conveyed?

$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$

Speaker

Listener




$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Are speakers' referent gender expectations faithfully conveyed?

$$P_S(\text{ref.gender}) \stackrel{?}{=} P_L(\text{ref.gender})$$


Speaker

Listener




Symmetric

Calibrated



Symmetric




Faithfully Conveyed


$$\sum_{pro} P(\text{pro} | P_S(\text{ref.gender})) \cdot P(\text{ref.gender} | \text{pro}) = P_L(\text{ref.gender})$$

Interpretation


Symmetric



Symmetric




Production





Interpretation

Symmetric




Symmetric



Calibrated


Production

Asymmetric




Interpretation


Symmetric



Symmetric





Production




Calibrated

Asymmetric




Interpretation


Symmetric





Symmetric



Production




Asymmetric




Interpretation



Symmetric



Asymmetric





Symmetric




Production

Asymmetric




Interpretation




Symmetric



Asymmetric





Symmetric




Production

Asymmetric




Interpretation




Symmetric



Asymmetric





Symmetric




Production

Asymmetric




Interpretation




Symmetric



Asymmetric






Symmetric




Production

Asymmetric




Interpretation




Symmetric



Asymmetric






Symmetric




Production

Asymmetric




Interpretation




Symmetric






Asymmetric



Symmetric




Production




Asymmetric

Interpretation

Symmetric






Asymmetric



?




Symmetric



Production



Asymmetric



Experiments

Experiments

Test with 80 role nouns (ex. *diplomat*, *butler*, *nanny*, *reporter*)

Experiments

Test with 80 role nouns (ex. *diplomat*, *butler*, *nanny*, *reporter*)

a. Explicit Norm

If you encounter a baker, how likely are they to be female?

Experiments

Test with 80 role nouns (ex. *diplomat*, *butler*, *nanny*, *reporter*)

a. Explicit Norm

If you encounter a baker, how likely are they to be female?

b. Production

Before the baker took out the trash, ... _____

Experiments

Test with 80 role nouns (ex. *diplomat, butler, nanny, reporter*)

a. Explicit Norm

If you encounter a baker, how likely are they to be female?

b. Production

Before the baker took out the trash, ... _____

c. Gender recall

After the baker ... he ...

1, 4, or 8 stories

$5 * 7 + 12 - 3 = \underline{\hspace{1cm}}$

working memory task

In the story you read,
what was the gender of
the baker?

One question
per story

Female / Male / Other

Experiments

a. Explicit Norm

If you encounter a **baker**, how likely are they to be female?

- 51 participants each rated all 80 nouns
- Half asked “how male”, half asked “how female”
- We averaged results into a 0-1 ‘how female’ norm

Experiments

b. Production

After the shop on High Street closed for the night, a **baker** stayed to tidy up.
Before the **baker** took out the trash, ...

- 149 participants each completed 20 items
- We coded pronoun completions

Experiments

c. Gender recall

After the shop on High Street closed for the night, a **baker** stayed to tidy up. Before the **baker** took out the trash, **he** swept the floor and wiped down the counter.

1,4, or 8 stories

Experiments

c. Gender recall

$$(70 - 4) \div 3 + 9 =$$

Experiments

c. Gender recall

In the story you read, what was the gender of the **baker**?

Female Male Other

One question per story

Experiments

c. Gender recall


In the story you read, what was the gender of the **baker**?

Female Male Other

One question per story


- 712 participants
- We only consider male, female responses

Production




Pronoun ● He ● She ● They

Production




Pronoun ● He ● She ● They

Interpretation




Interpretation



Asymmetric

Asymmetric




$$\sum_{pro} P(pro|P_S(ref.gender)) \cdot P(ref.gender|pro) = P_L(ref.gender)$$

Asymmetric

Asymmetric

Mis-calibrated




$$\sum_{pro} P(pro|P_S(ref.gender)) \cdot P(ref.gender|pro) = P_L(ref.gender)$$

Asymmetric

Asymmetric

Mis-calibrated




Gender of referent not faithfully conveyed!

$$\sum_{pro} P(pro|P_S(ref.gender)) \cdot P(ref.gender|pro) = P_L(ref.gender)$$

Asymmetric

Asymmetric


Mis-calibrated



Gender of referent not faithfully conveyed!
Female gender expectations are under-recovered!

$$\sum_{pro} P(pro|P_S(ref.gender)) \cdot P(ref.gender|pro) = P_L(ref.gender)$$


Production



Pronoun




Interpretation



Explicit Norm – How Female

Weighted Interpretation



1
trial

4
trial

8
trial

Conclusion

Gender of referent not faithfully conveyed!

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated
 - A challenge for rational theories that expect listeners and speakers to calibrate for effective communication

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated
 - A challenge for rational theories that expect listeners and speakers to calibrate for effective communication
- What might be the source of the miscalibration?

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated
 - A challenge for rational theories that expect listeners and speakers to calibrate for effective communication
- What might be the source of the miscalibration?
 - Grammatical gender is often not central to communicative goals, so may not be under much pragmatic pressure?

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated
 - A challenge for rational theories that expect listeners and speakers to calibrate for effective communication
- What might be the source of the miscalibration?
 - Grammatical gender is often not central to communicative goals, so may not be under much pragmatic pressure?
 - Speakers and listeners are rational, but they consider misgendering a male as female as costlier than vice versa?

Conclusion

Gender of referent not faithfully conveyed!

- Pronoun production and interpretation are not mutually calibrated
 - A challenge for rational theories that expect listeners and speakers to calibrate for effective communication
- What might be the source of the miscalibration?
 - Grammatical gender is often not central to communicative goals, so may not be under much pragmatic pressure?
 - Speakers and listeners are rational, but they consider misgendering a male as female as costlier than vice versa?
- Socially concerning – what are consequences of this miscalibration?

Implicit Norm

d. Implicit Norm

After the shop on High Street closed for the night, a **baker** stayed to tidy up. Before the **baker** took out the trash, the **baker** swept the floor and wiped down the counter.

1,4, or 8 stories

Implicit Norm

d. Implicit Norm

$$(70 - 4) \div 3 + 9 =$$

Implicit Norm

d. Implicit Norm

In the story you read, what was the gender of the **baker**?

Female Male Other

One question per story

Implicit Norm

d. Implicit Norm

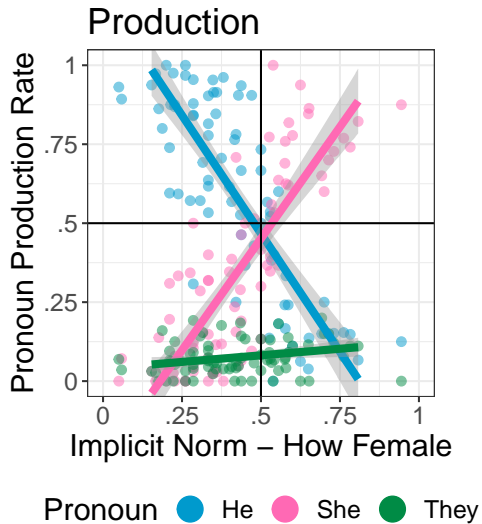
In the story you read, what was the gender of the **baker**?

Female Male Other

One question per story


- 569 participants
- Only consider male, female responses

Implicit Norm

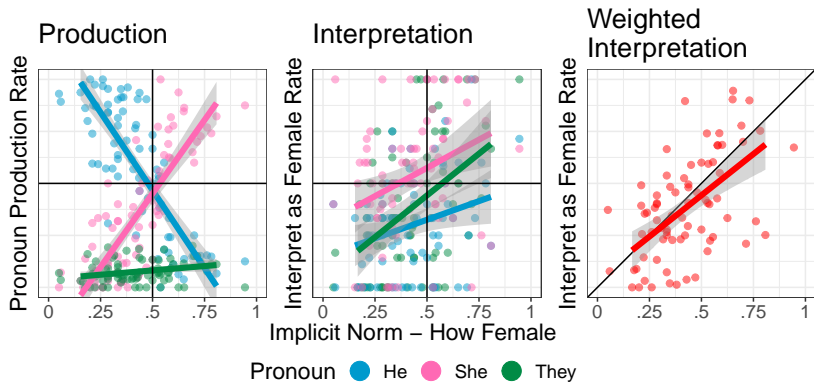


Implicit Norm

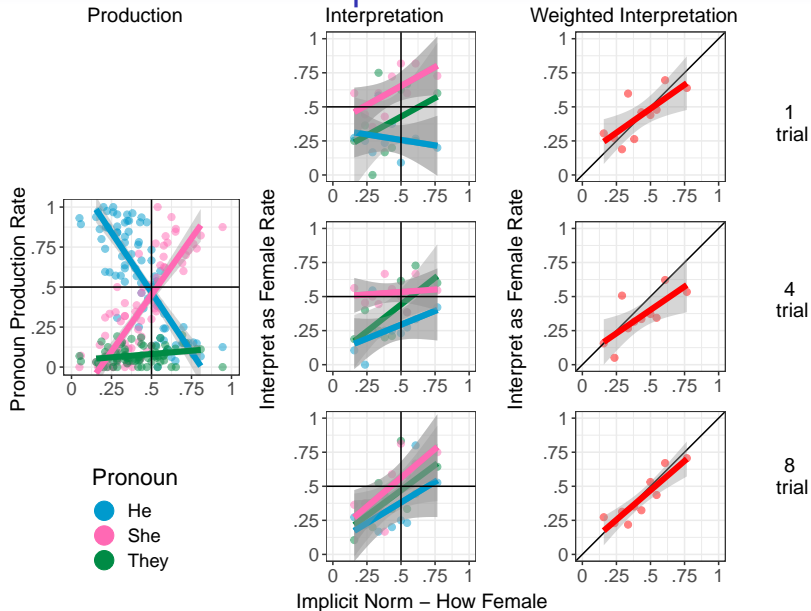
Interpretation



Implicit Norm

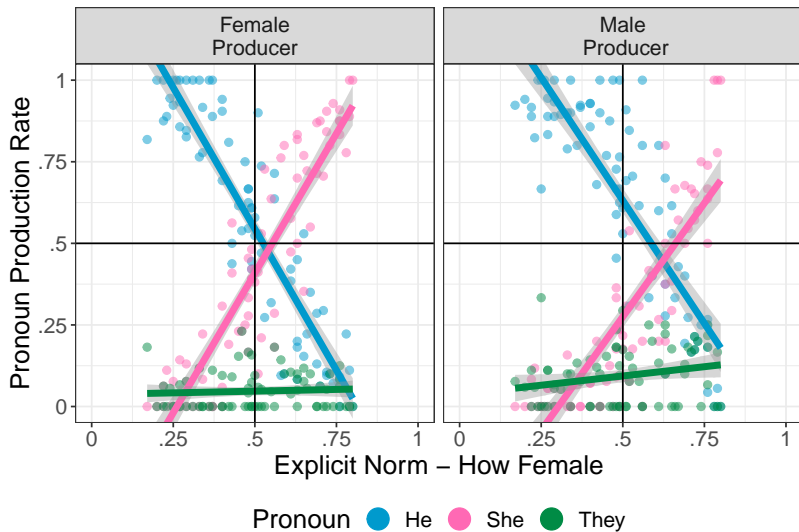


Implicit Norm

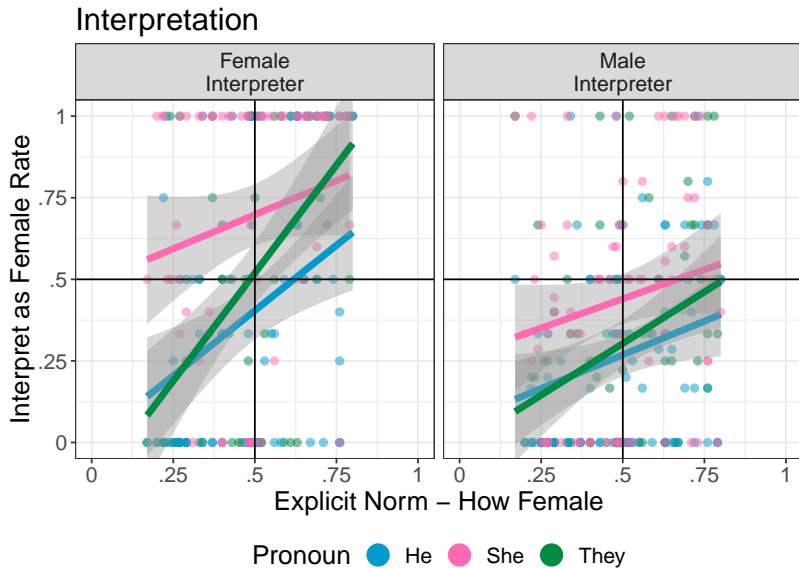


By gender

Production




By gender




By gender

Weighted Interpretation




Participant Demographics

Participants in Norming Study



Participant Demographics

Gender recall study



Participant Demographics

Production study

Participants in Production Study

