

Remember 'him', forget 'her':  
Gender bias in the comprehension of pronominal  
referents

Veronica Boyce (MIT)  
Titus von der Malsburg (University of Potsdam)  
Till Poppels (UCSD)  
Roger Levy (MIT)

March 29, 2019

von der Malsburg, Poppels, and Levy (2017)

- Presidential election season in 2016

von der Malsburg, Poppels, and Levy (2017)

- Presidential election season in 2016
- Natural experiment: Gender of next president is unknown

von der Malsburg, Poppels, and Levy (2017)

- Presidential election season in 2016
- Natural experiment: Gender of next president is unknown
- How did people refer to the next president?  
He? She? They?


von der Malsburg, Poppels, and Levy (2017)

- Presidential election season in 2016
- Natural experiment: Gender of next president is unknown
- How did people refer to the next president?  
He? She? They?
- How did pronoun choice relate to election outcome expectations?




*How likely is ... to be the next president?*

*How likely is ... to be the next president?*





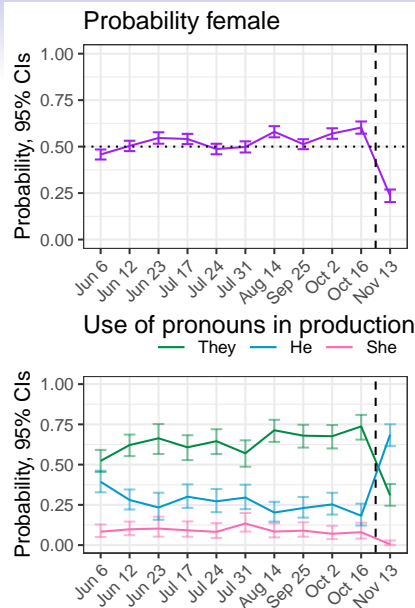
*How likely is ... to be the next president?*



*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ... \_\_\_\_\_*

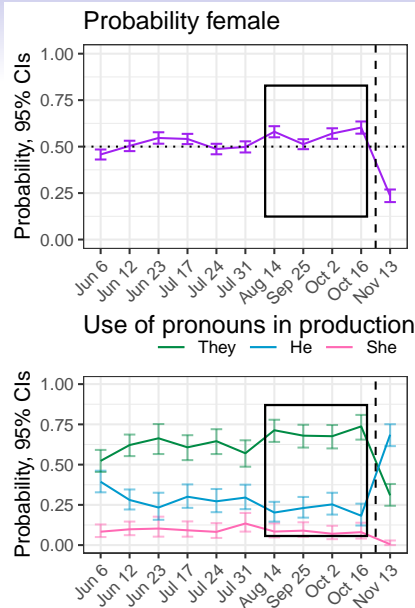
*How likely is ... to be the next president?*

*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ... \_\_\_\_\_*



*How likely is ... to be the next president?*

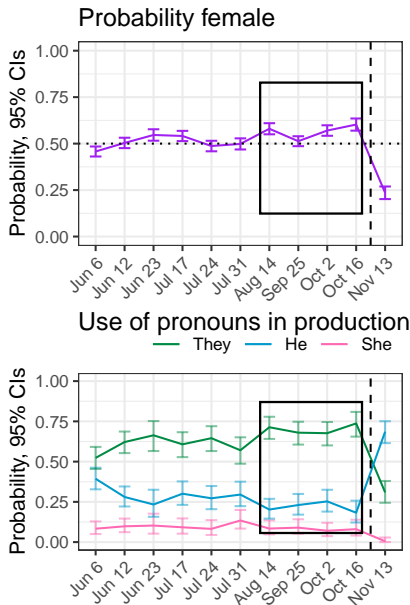
*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ... \_\_\_\_\_*



*How likely is ... to be the next president?*

*$P(\text{female}) \sim 60\%$*

*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ... \_\_\_\_\_*

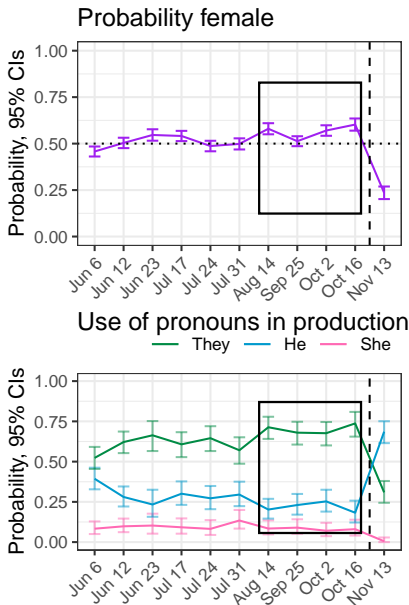


*How likely is ... to be the next president?*

$P(\text{female}) \sim 60\%$

65% They, 25% He, 10% She

*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ...*

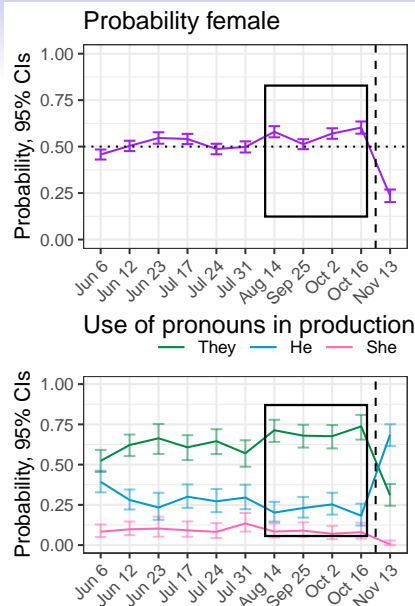


*How likely is ... to be the next president?*

$P(\text{female}) \sim 60\%$

65% They, 25% He, 10% She

*The next US president will be sworn into office in January 2017. After moving into the Oval Office, one of the first things that ...*



**This production asymmetry may be a surprise.  
But is it a concern?**

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, ...



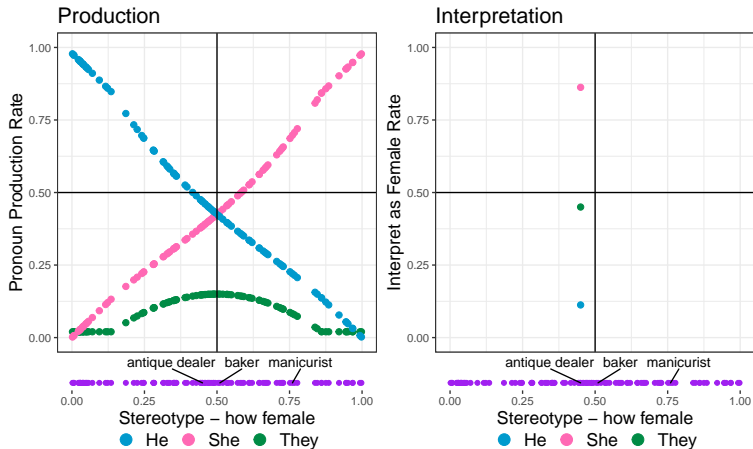
Symmetric

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?






# Are speakers' referent gender expectations faithfully conveyed?


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

Speaker  $\rightarrow$

$\leftarrow$  Listener



Symmetric




Symmetric

Are speakers' referent gender expectations faithfully conveyed?


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

Speaker  $\rightarrow$

$\leftarrow$  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  $\rightarrow$


$\leftarrow$  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  $\rightarrow$


$\leftarrow$  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  $\rightarrow$


$\leftarrow$  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  $\rightarrow$

$\leftarrow$  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




Calibrated





# Interpretation

Symmetric




Symmetric



Calibrated


# Production

Asymmetric





# Interpretation

Symmetric





Symmetric



Calibrated


Production

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

$5 * 7 + 12 - 3 = \underline{\quad}$

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

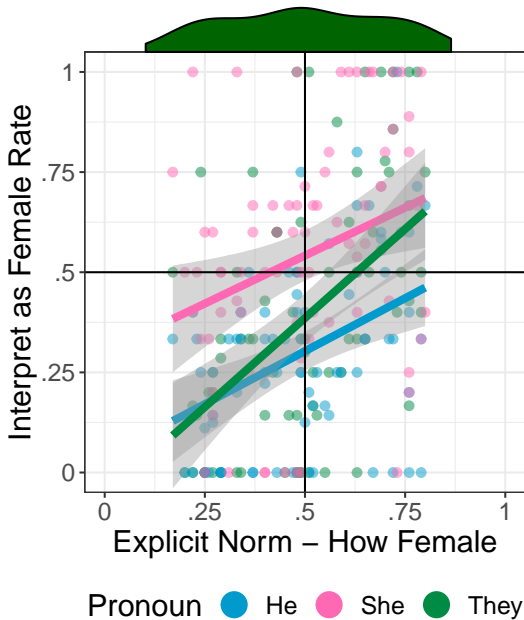


Pronoun ● He ● She ● They






# Interpretation



# Interpretation






## Asymmetric

## Asymmetric

## Weighted Interpretation




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

## 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**



Pronoun ● He ● She ● They




**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




Pronoun ● He ● She ● They

**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

- He
- She
- They

Interpretation




Interpret as Female Rate



Explicit Norm – How Female

Weighted Interpretation

Interpret as Female Rate



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

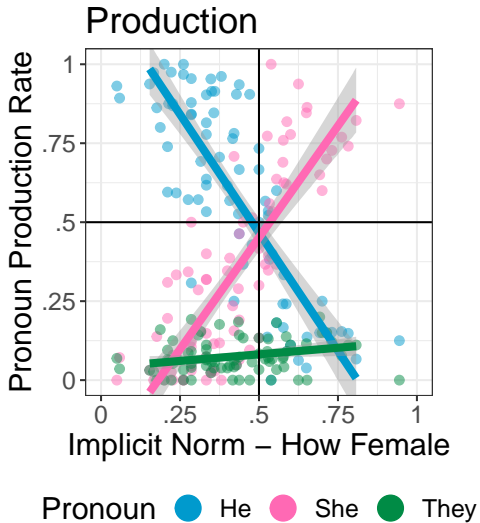
In the story you read, what was the gender of the **baker**?

Female    Male    Other

One question per story


- 569 participants
- Only consider male, female responses

# Implicit Norm




# Implicit Norm

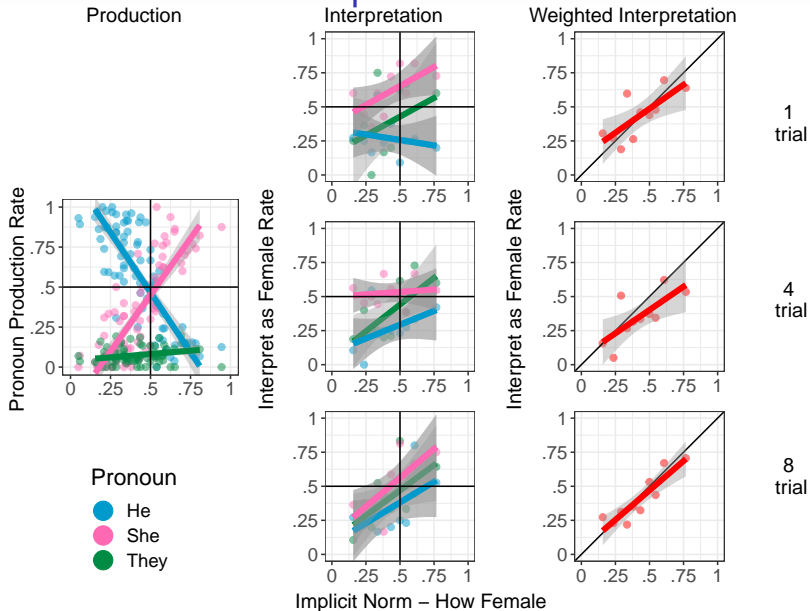
## Interpretation



# Implicit Norm



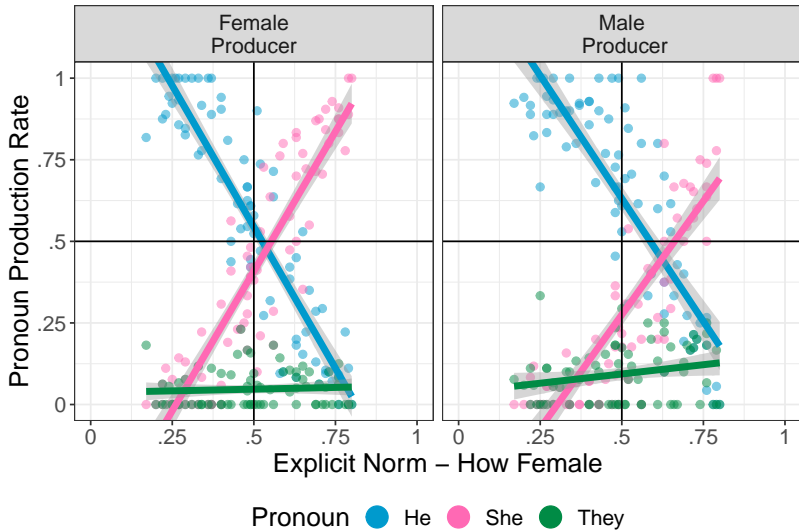
# Implicit Norm





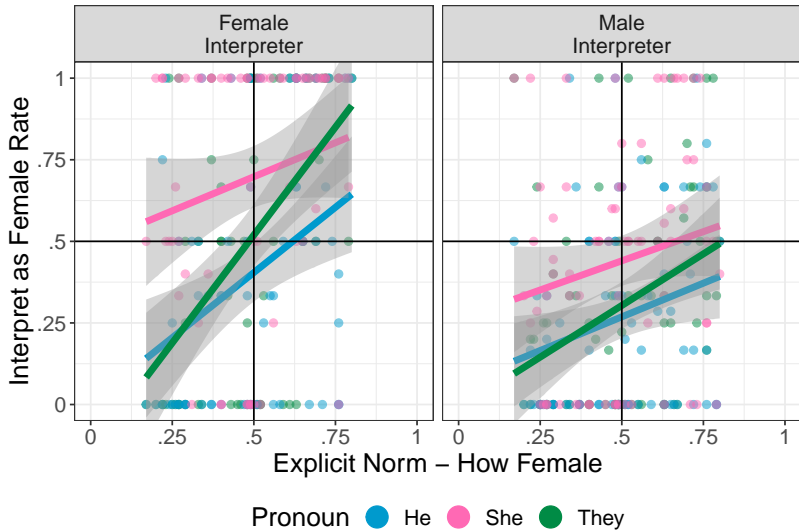
# By gender

## Production



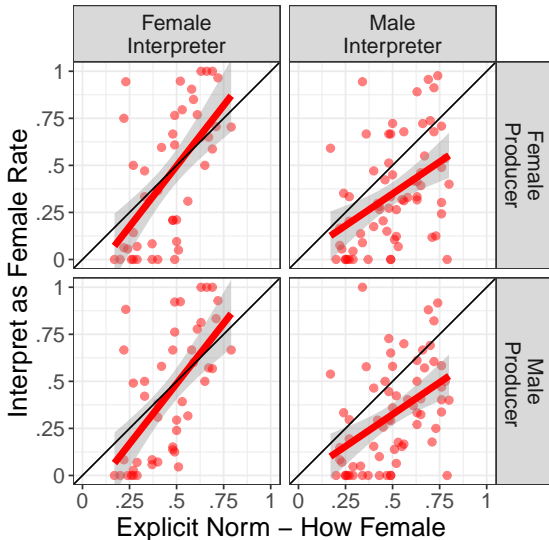
# By gender

## Interpretation




# By gender

## Weighted Interpretation




# Participant Demographics

## Participants in Norming Study



# Participant Demographics

Gender recall study



# Participant Demographics

Production study

Participants in Production Study

