

# Negation and semantic relatedness in eye-tracking- while-reading

Erika Mayer, Adrian Staub, and Brian Dillon

# Semantic operators

- Semantic operators are words like *not*, *even*, *only*, and many more.
- Like nouns and verbs, they alter the meaning of sentences.
- However, unlike many nouns and verbs, their meanings are abstract and often change or add to the sentence's **logical properties**.
- **How do comprehenders process and understand these abstract, complex operators during reading and listening?**
- **How do these operators interact with the semantics of lexical items like nouns?**

# The semantic operator negation

- What does it mean to change a sentence's logical properties?
- We can look at negation as an example.
- Sentential negation modifies sentences and reverses their truth conditions.

True or false?


A squirrel is a mammal.



A squirrel is a reptile.



True



False


True or false?

A squirrel is not a mammal.



False

A squirrel is not a reptile.



True

# How is negation processed online?

- The fact that sentential negation has such a drastic effect on the proposition it modifies raises an important question: how and when do comprehenders take negation into account in online processing?

# Fischler et al. (1983)

- Event-related potential (ERP) study investigating the online processing of negation
- The N400: an event-related potential associated with semantic anomalies
  - Elicited for sentences like *A giraffe is a large sock.*
  - Reflects semantic processing of a word
  - Occurs 300 – 500ms after word onset

# Fischler et al.: Design

- Factors: Truth value (true or false), Negation (negative or affirmative)
- Participants read sentences word by word. They then judged whether the sentences were true or false.
- ERPs were measured at the **object noun critical region**.

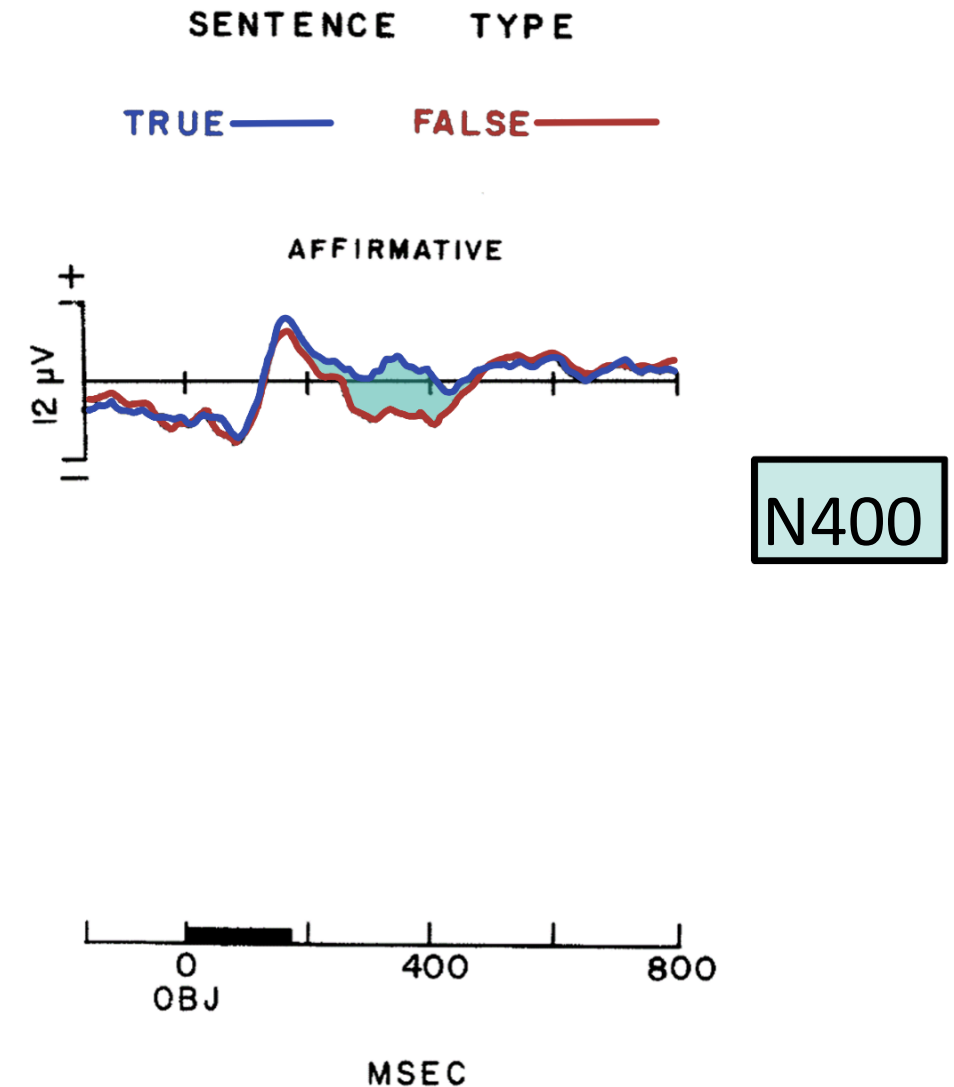
Subject noun	Negative/ affirmative	Object critical region		True or false
A trout	is a	fish		True
A trout	is a	tree		False
A trout	is not a	fish		False
A trout	is not a	tree		True



# ERP Analyses

- For affirmative sentences, the N400 appeared to be related to sentence truth condition: false sentences were associated with a larger N400 (i.e., were more negative) than true sentences.
- For negative sentences, the opposite held: true sentences were associated with a larger N400 than false sentences.

Subject noun	Negative/affirmative	Object critical region		True or false
A trout	is a	fish		True
A trout	is a	tree		False
A trout	is not a	fish		False
A trout	is not a	tree		True



**Figure 3.** Averaged ERPs for true- versus false-sentence trials, all subjects, location C<sub>z</sub>, from 400 msec before to 800 msec following the onset of the sentence object (O frame).

# Fischler et al.: Takeaways

- Fischler et al. suggest that these results are consistent with a multi-step model of sentence processing.
  - In this model, if comprehenders hear a sentence like *A giraffe is not a sock*, they first process the core affirmative proposition *A giraffe is a sock*. Then the negation is incorporated.

# Negation and semantic relatedness

- However, Fischler et al.'s results also make sense if **semantic relatedness** has a greater influence on the N400 than a proposition's **truth value**.
- **Adding sentential negation to a sentence reverses its truth conditions, but does not change how semantically related the words in the proposition are.**
- Semantically related = close in meaning or conceptually connected.

A squirrel is a mammal.

squirrel + mammal = semantically related



A squirrel is a reptile.

squirrel + reptile = semantically unrelated



A squirrel is not a mammal.

squirrel + mammal = semantically related



A squirrel is not a reptile.

squirrel + reptile = semantically unrelated



# Returning to Fischler et al.

- Instead of looking at truth and falsity, we can look at the **semantic relatedness** of the subject and the noun.
- These results can be explained if we assume that semantic relatedness between the subject and object has the same effect regardless of negation.

Subject noun	Negative/ affirmative	Object critical region		True or false
A trout	is a	fish		True
A trout	is a	tree		False
A trout	is not a	fish		False
A trout	is not a	tree		True

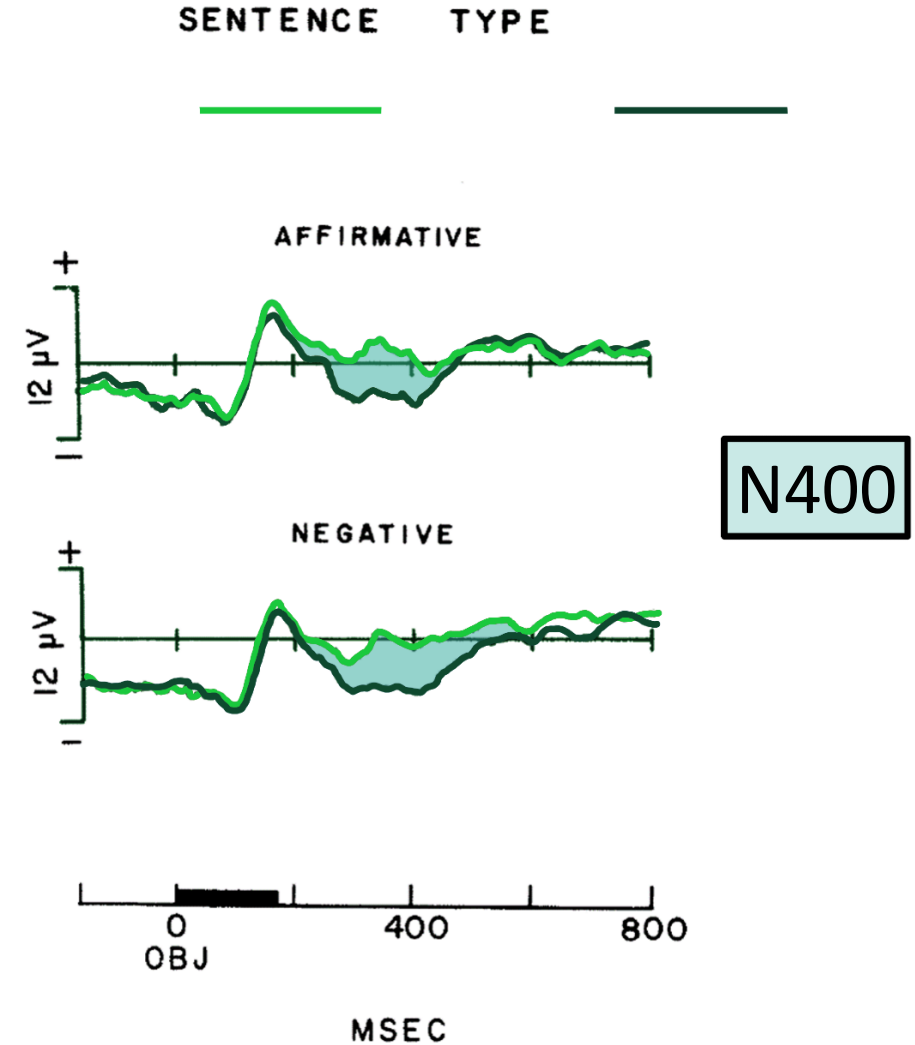


Subject noun	Negative/ affirmative	Object critical region		Related/ unrelated
A trout	is a	fish		Related
A trout	is a	tree		Unrelated
A trout	is not a	fish		Related
A trout	is not a	tree		Unrelated

# Semantic relatedness

- *A trout is a fish* is less negative than *A trout is a tree*
- *A trout is not a fish* is also less negative than *A trout is not a tree*

Subject noun	Negative/affirmative	Object critical region		Related/unrelated
A trout	is a	fish		Related
A trout	is a	tree		Unrelated
A trout	is not a	fish		Related
A trout	is not a	tree		Unrelated



**Figure 3.** Averaged ERPs for true- versus false-sentence trials, all subjects, location C<sub>z</sub>, from 400 msec before to 800 msec following the onset of the sentence object (O frame).

# The current study: Eye-tracking-while-reading

- Due to the design limitations of ERP analyses, Fischler's analyses were limited to one word – the object critical region.
- An eye-tracking-while-reading study gives us the ability to look at the time-course of reading for pre-critical and post-critical regions as well.
- Eye-tracking-while-reading will also allow us to investigate whether eye-tracking measures are consistent with the N400 for negated sentences.



# The current study: Eye-tracking-while-reading

- 47 participants
- 20 test items
- 65 fillers and other items
- Occasional comprehension questions

# Design

Matrix/carrier sentence	Embedded sentence				Continuation
	Subject noun	Negative/ affirmative	Critical region	Spillover	
The professor told the student that	a squirrel	is a	mammal	on Saturday	during office hours.
The professor told the student that	a squirrel	is a	reptile	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	mammal	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	reptile	on Saturday	during office hours.

# A proposition is embedded within a matrix proposition

Matrix/carrier sentence	Embedded sentence			Continuation of matrix sentence	
	Subject noun	Negative/affirmative	Critical region	Spillover	
The professor told the student that	<b>a squirrel</b>	<b>is a</b>	<b>mammal</b>	on Saturday	during office hours.
The professor told the student that	<b>a squirrel</b>	<b>is a</b>	<b>reptile</b>	on Saturday	during office hours.
The professor told the student that	<b>a squirrel</b>	<b>is not a</b>	<b>mammal</b>	on Saturday	during office hours.
The professor told the student that	<b>a squirrel</b>	<b>is not a</b>	<b>reptile</b>	on Saturday	during office hours.

# Embedded proposition

- The embedding proposition provides multiple benefits:
  - The sentence is less pragmatically odd when it is a report of something that someone else has said.
  - The matrix sentence allows for a longer spillover region.
- One key difference between this study and Fischler et al. was the use of an embedding sentence.
- The overall sentence, then, is not obviously true or false in the real world.
- It is the embedded proposition that is true or false.

# Factor 1: Affirmative vs. negative

Matrix/carrier sentence	Embedded sentence			Continuation of matrix sentence	
	Subject noun	Negative/ affirmative	Critical region	Spillover	
The professor told the student that	a squirrel	is a	mammal	on Saturday	during office hours.
The professor told the student that	a squirrel	is a	reptile	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	mammal	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	reptile	on Saturday	during office hours.

# Factor 2: Semantically related or unrelated object noun

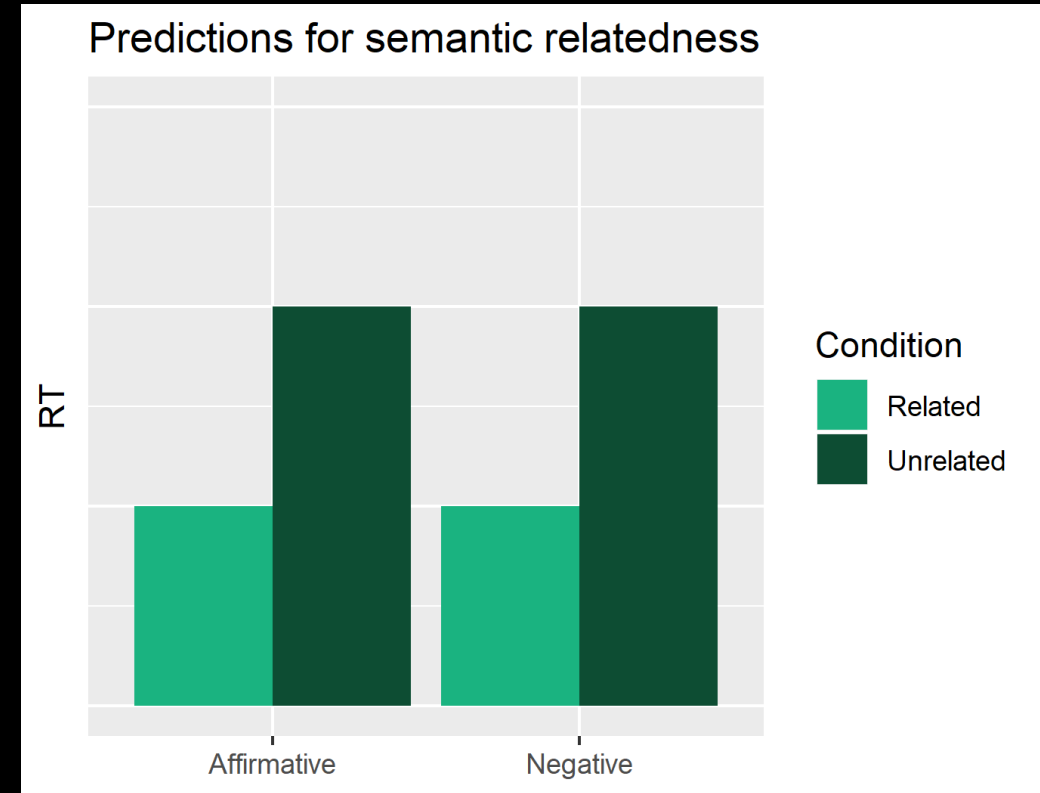
- Roughly matched in length
- Lexical features were controlled

Matrix/carrier sentence	Embedded sentence			Continuation of matrix sentence	
	Subject noun	Negative/ affirmative	<b>Critical region</b>	Spillover	
The professor told the student that	a squirrel	is a	<b>mammal</b>	on Saturday	during office hours.
The professor told the student that	a squirrel	is a	<b>reptile</b>	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	<b>mammal</b>	on Saturday	during office hours.
The professor told the student that	a squirrel	is not a	<b>reptile</b>	on Saturday	during office hours.

# Predictions

# If only semantic relatedness matters

- For a replication of Fischler et al.:
  - We should see main effects of semantic relatedness.
  - Unrelated object nouns should be read more slowly than related object nouns.

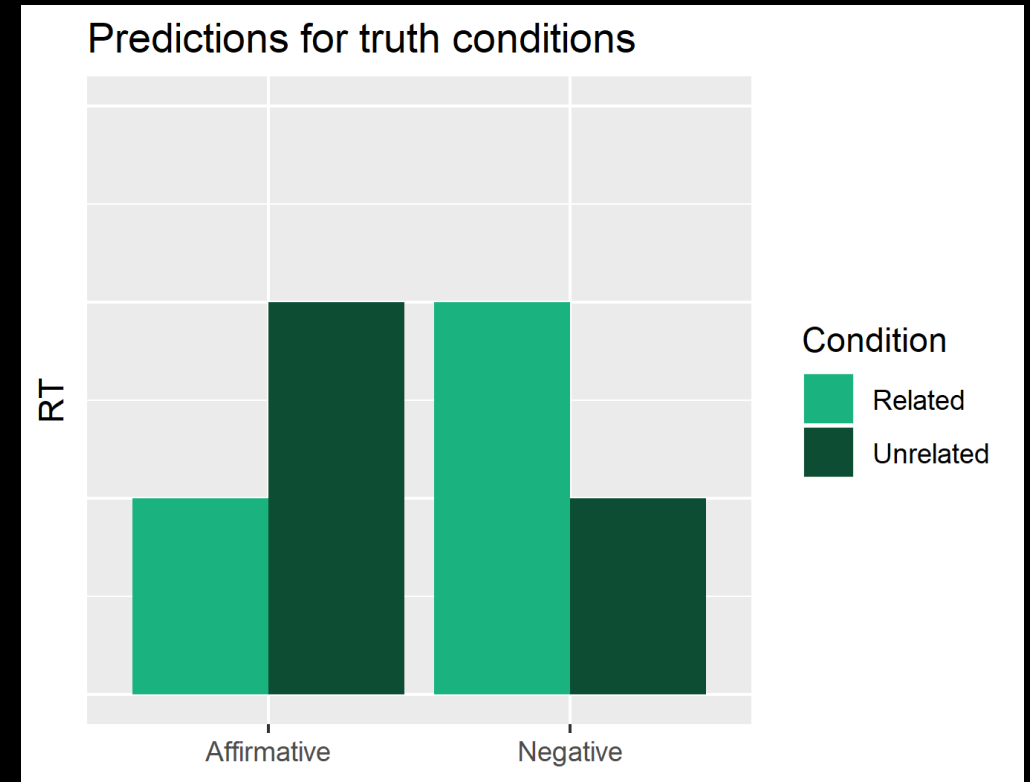


	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.



# If only truth condition matters

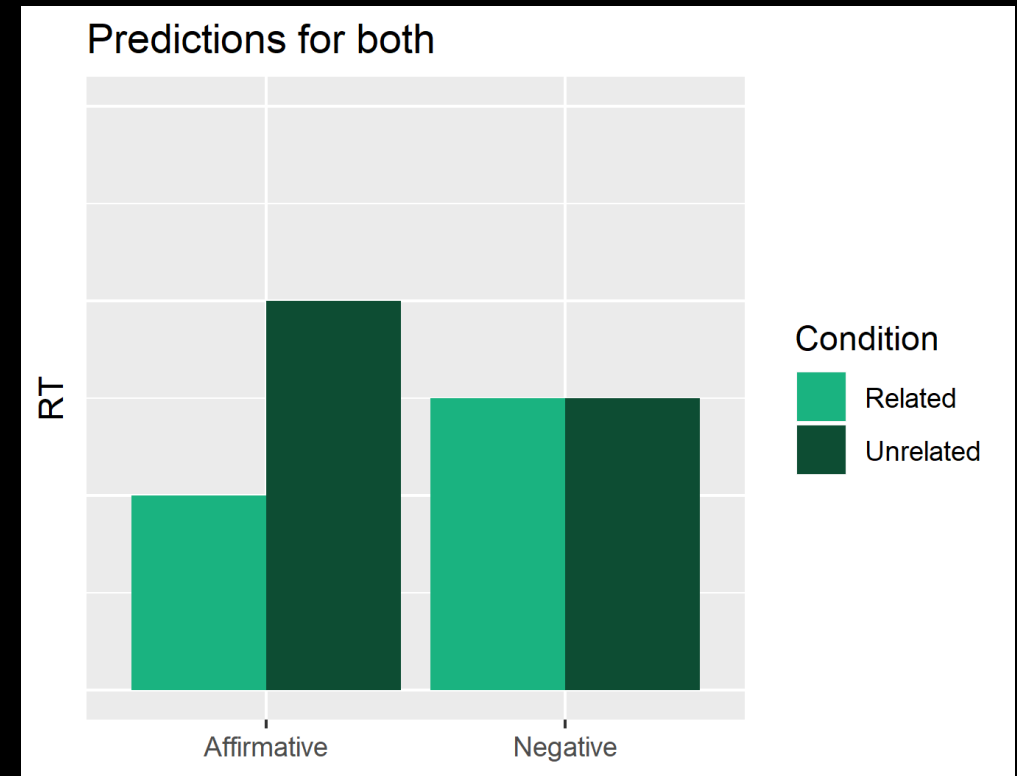
- If truth conditions matter more than semantic relatedness:
  - We should see an interaction between semantic relatedness and negative/affirmative conditions.



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# If both truth condition and semantic relatedness matter

- If both semantic relatedness and truth conditions matter:
  - We should see an interaction between semantic relatedness and negative/affirmative conditions **and** a main effect of semantic relatedness.




	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Results

# Early reading time measures

# First fixation

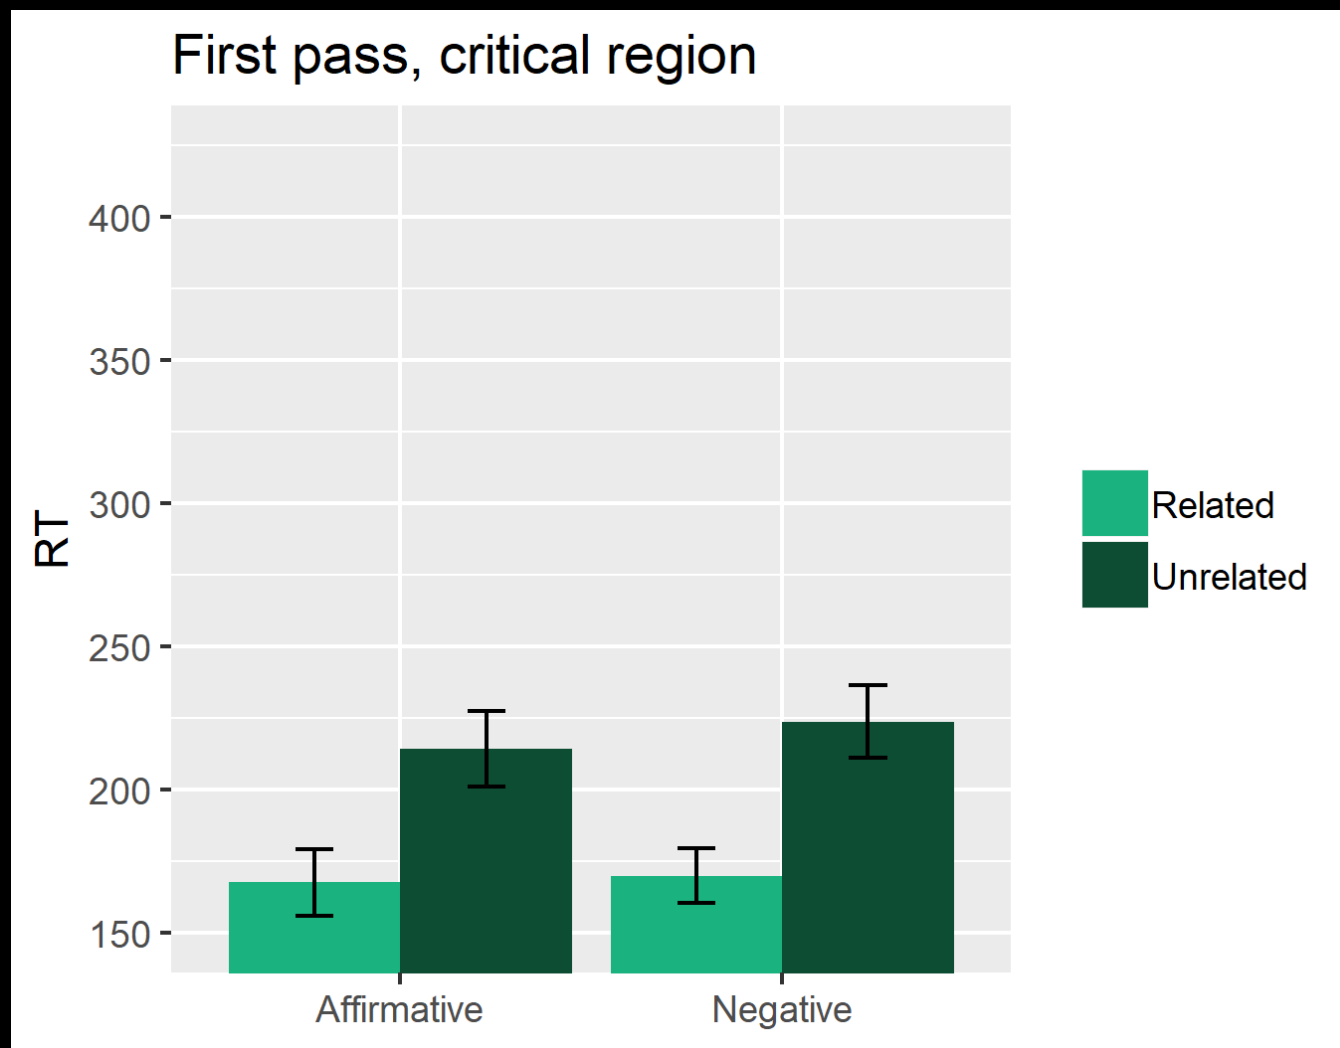
- Main effect of semantic relatedness: semantically related words were read faster than semantically unrelated words.
- There was no main effect of negation.
- Crucially, there was no interaction: in other words, there was no significant penalty for **false** embedded propositions.



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# First pass


- First pass RTs had the same results as first fixation.
- Main effect of semantic relatedness
- No main effect of negation.
- No interaction



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Skipping

- There was only a main effect of semantic relatedness: semantically related words were skipped more often than semantically unrelated words.



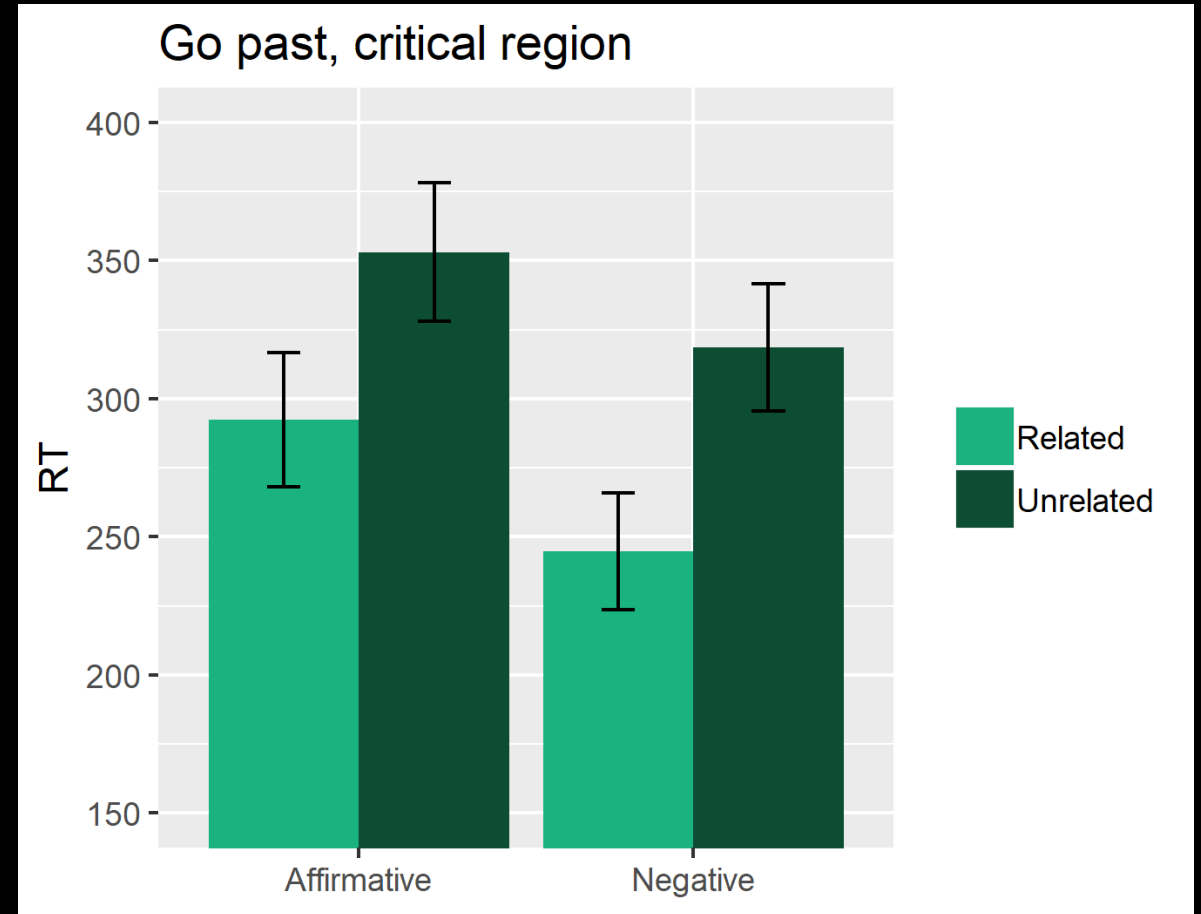
	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

Later reading time measures



# Go past

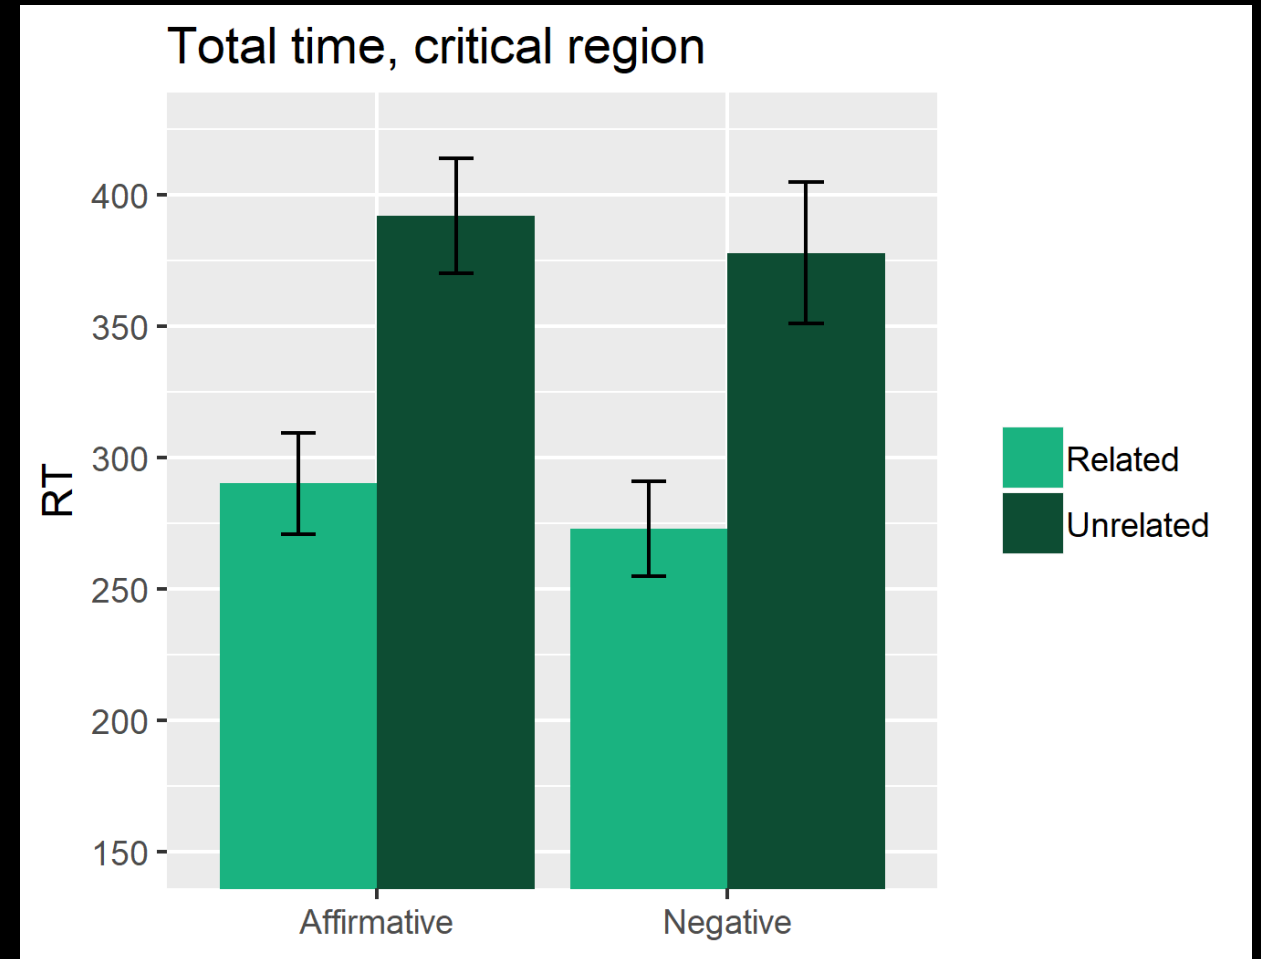
- Again, a main effect of semantic relatedness in the expected direction.
- Again, no interaction – there was no penalty for **false** sentences in the negation condition.
- There was a main effect of negation: the critical regions of negative sentences were read **faster** than those of affirmative sentences.



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Total time

- Again, a main effect of semantic relatedness in the expected direction.
- No main effect of negation
- No interaction

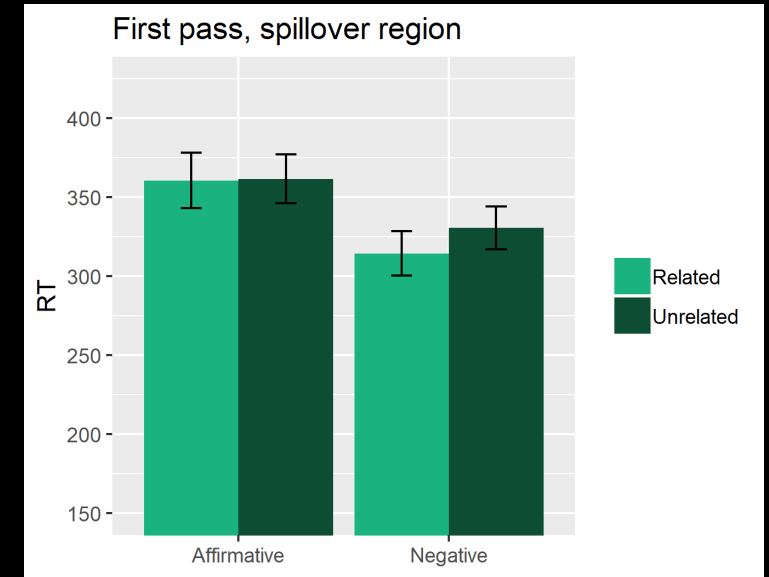
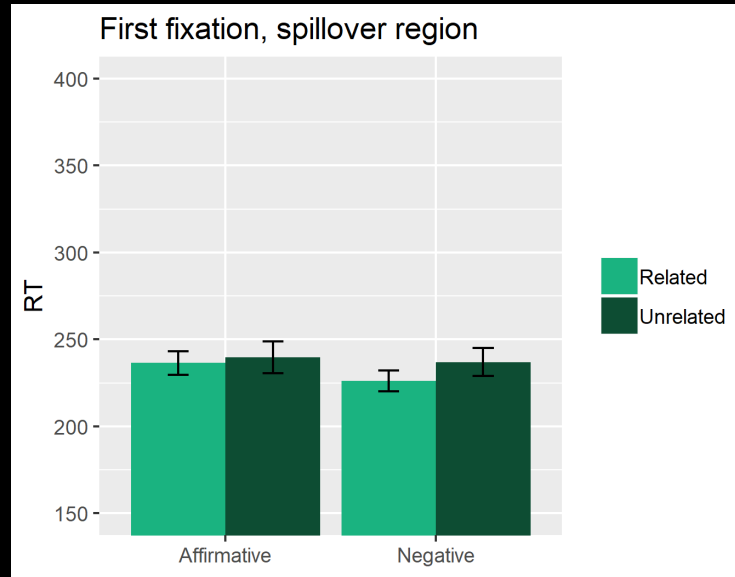


	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

Spillover region

# First fixation and first pass

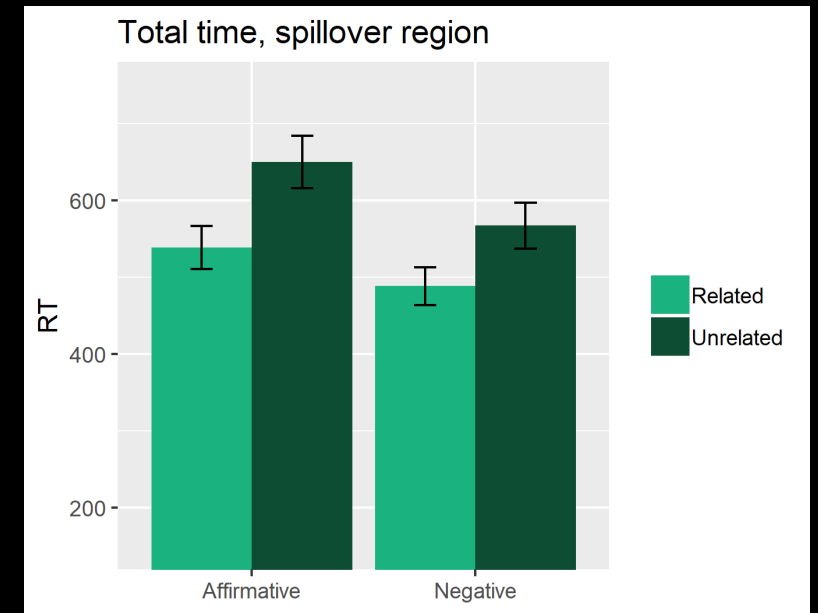
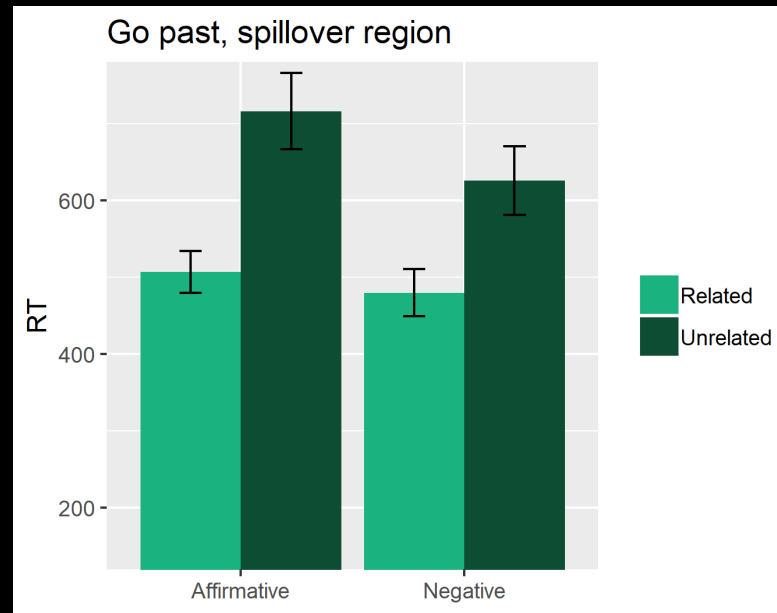
- First fixation: no significant effects
- First pass: significant effect of negation



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Go past and total time

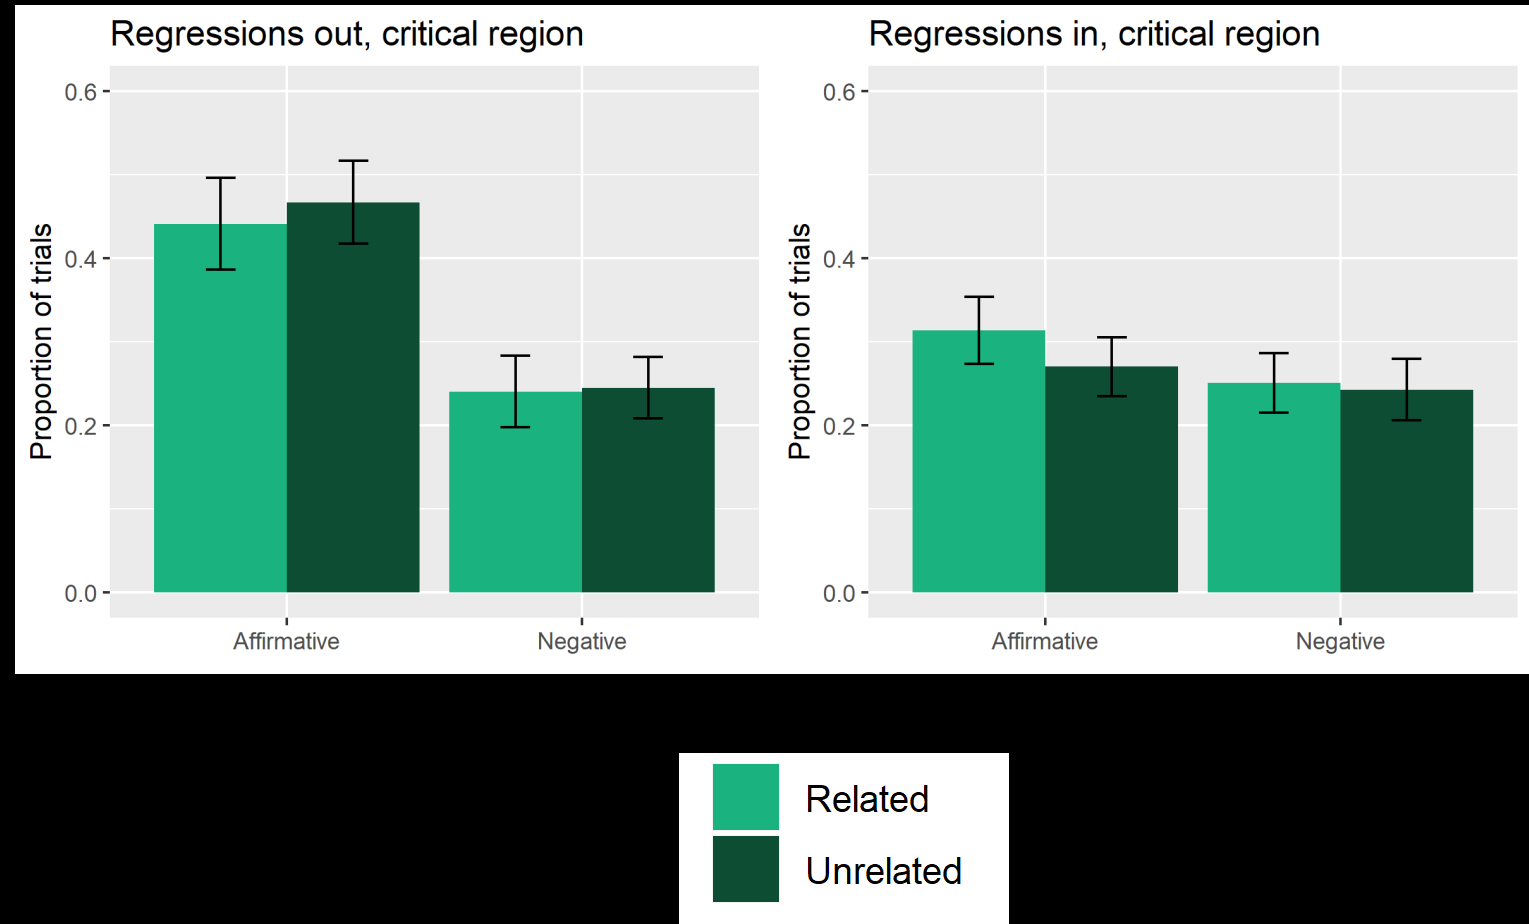
- Main effects of semantic relatedness and negation



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Regressions: Critical region

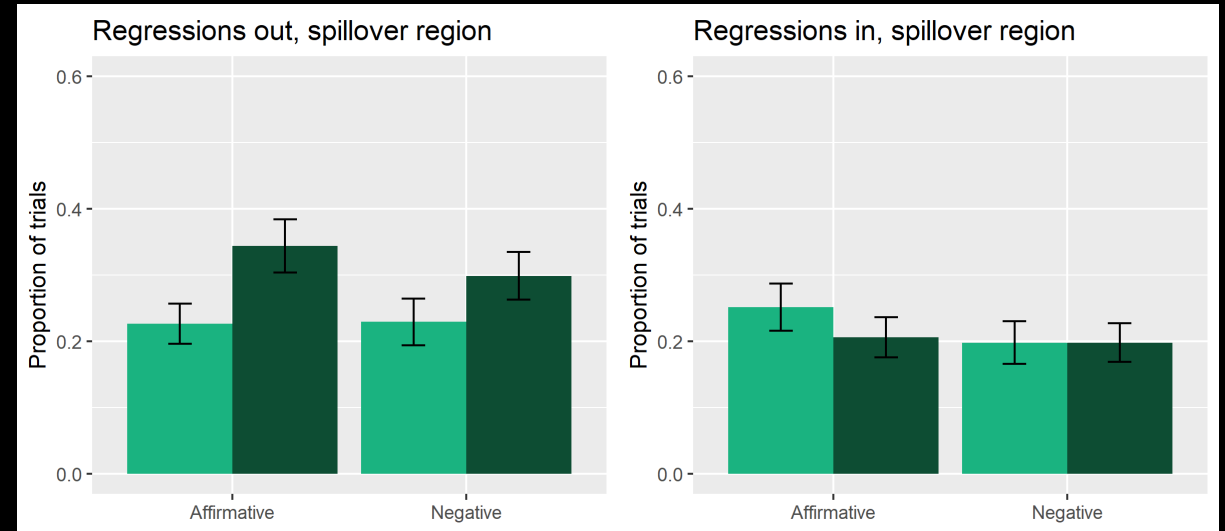
- Regressions out of the critical region
  - Main effect of negation
  - More regressions for affirmative condition
- Regressions into the critical region
  - All effects n.s.



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Regressions: Spillover region

- Regressions out of the spillover region
  - Main effect of semantic relatedness
- Regressions into the spillover region
  - All effects n.s.



	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

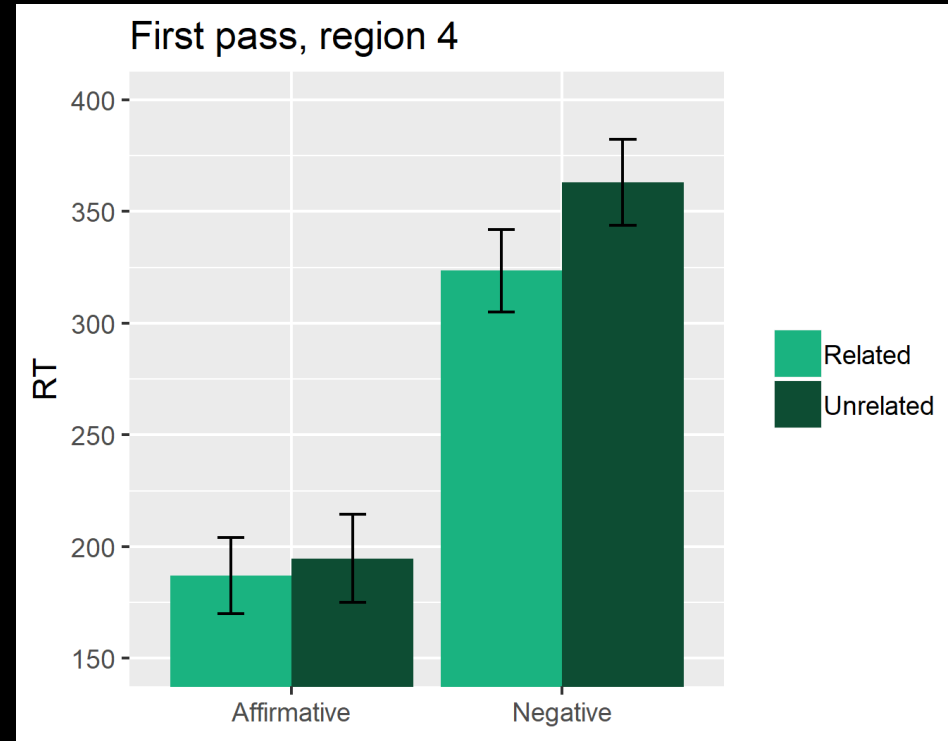
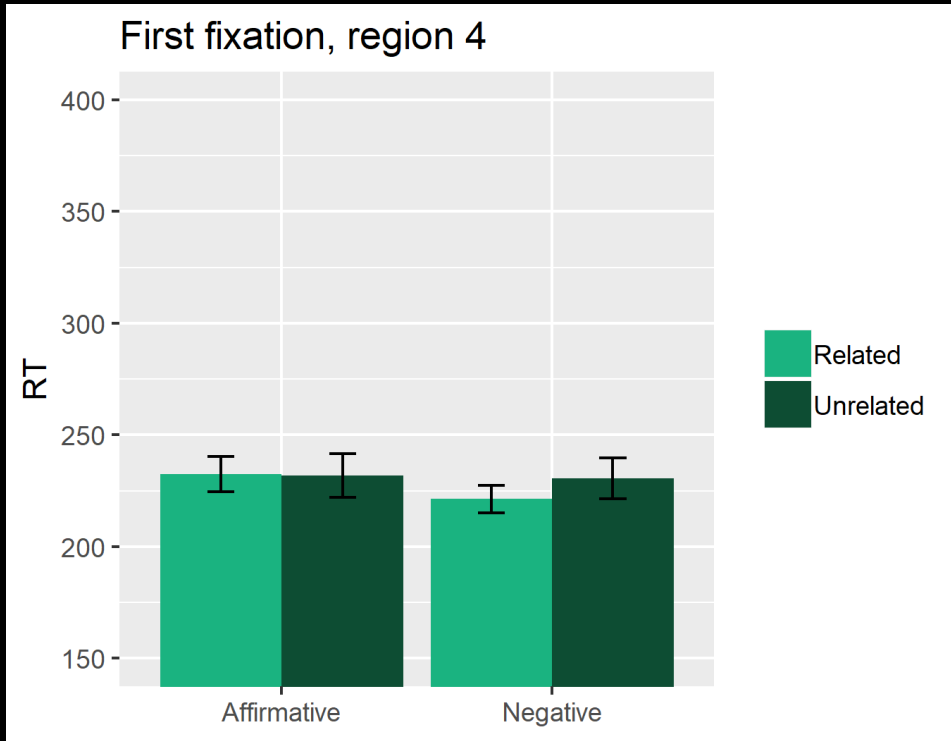
# Discussion: why was there a main effect of negation?

- Participants spent **much** longer on the previous region when it included negation.
- Perhaps this slowdown led them to speed up later on.
  - “Catching up”
  - Parafoveal preview
- Perhaps negation was more natural in an embedded context.

	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.



# Results: Early RT measures for negative or affirmative region (region 4)



- Region 4:
  - is a
  - is not a

	Critical region	Spillover	
The professor told the student that a squirrel is {not} a	{mammal/reptile}	on Saturday	during office hours.

# Discussion

- There was **no evidence** for an interaction between semantic relatedness and negation in any reading time measure.
- The truth or falsity of the embedded proposition did not significantly affect reading times.
- This pattern replicates the ERP findings of Fischler et al. in eye-tracking-while-reading.

# Discussion

- There are many components of semantic processing:
  - Truth and falsity
  - Pragmatic acceptability
  - Predictability
  - Semantic relatedness
- In this study, reliable effects of **semantic relatedness**, but not truth or **falsity**, were found.

# Offline cloze norms

# Negation and offline cloze norms

- After asking whether negation influences **online** processing, it is important to check whether (and to what extent) it influences **offline** judgments in an experimental setting.
- To do this, we can use an offline **cloze norm task**.
- In this task, 40 participants were presented with 20 sentence fragments and asked to fill in a word.
- Some of the sentences were affirmative, while some were negative.

The professor told the student that a squirrel is a \_\_\_\_\_

## Responses

mammal      mammal      rodent  
mammal      thief      rodent  
mammal      animal  
rodent      rodent      mammal  
mammal      animal      animal  
mammal      mammal      rodent  
animal      mammal      rodent

## Count

mammal (9)  
rodent (6)  
animal (4)  
thief (1)

## Cloze probability Modal response

mammal – 45%  
rodent – 30%  
animal – 20%  
thief – 5%

The professor told the student that a squirrel is not a \_\_\_\_\_

## Responses

amphibian      mammal      mammal  
rodent              mammal  
rodent      animal              toy  
carnivore              animal      pet  
excuse              mammal      pet  
animal      bird              dog      large  
rodent      nut      predator

## Count

mammal (4)  
rodent (3)  
animal (2)  
pet (2)  
amphibian (1)  
bird (1)  
carnivore (1)  
dog (1)  
excuse (1)  
large animal (1)  
nut (1)  
predator (1)  
toy (1)

## Cloze probability Modal response

Mammal – 20%  
Rodent – 15%  
animal – 10%  
pet – 10%  
amphibian – 5%  
bird – 5%  
Carnivore – 5%  
dog – 5%  
excuse – 5%  
large animal – 5%  
nut – 5%  
predator – 5%  
toy – 5%

# Cloze norm analysis

- How can we tell whether readers are sensitive to negation in a cloze norm task?
- We explore three ways:
  - Entropy
  - Modal responses
  - Truth values of the response

The professor told the student that a squirrel is {not} a \_\_\_\_\_



By truth condition

# By truth condition


- Responses were coded based on whether they made the embedded proposition true or false.
- Only the affirmative version of the sentence preamble was considered during coding.
  - Sample response: The supervisor said that a penny is not a coin.
  - Affirmative version: *A penny is a coin*
  - Coded as **true**
- Items were randomized and conditions were hidden to avoid bias.
- Responses could be coded as *True, False, or Other*.

The professor told the student that a squirrel is {not} a \_\_\_\_\_

# By truth condition: Predictions

- We predict that there should be more false responses for negative embedded sentences than for affirmative ones (when judging based on a core affirmative proposition).

# Truth condition analysis: results



# Shannon entropy

# Shannon entropy

- Shannon entropy provides a way to quantify the amount of uncertainty in the probability distribution of a variable.

$$H(X) = - \sum_{i=0}^{N-1} p_i \log_2 p_i$$

Where  $p_i$  is the cloze probability of response  $i$ .

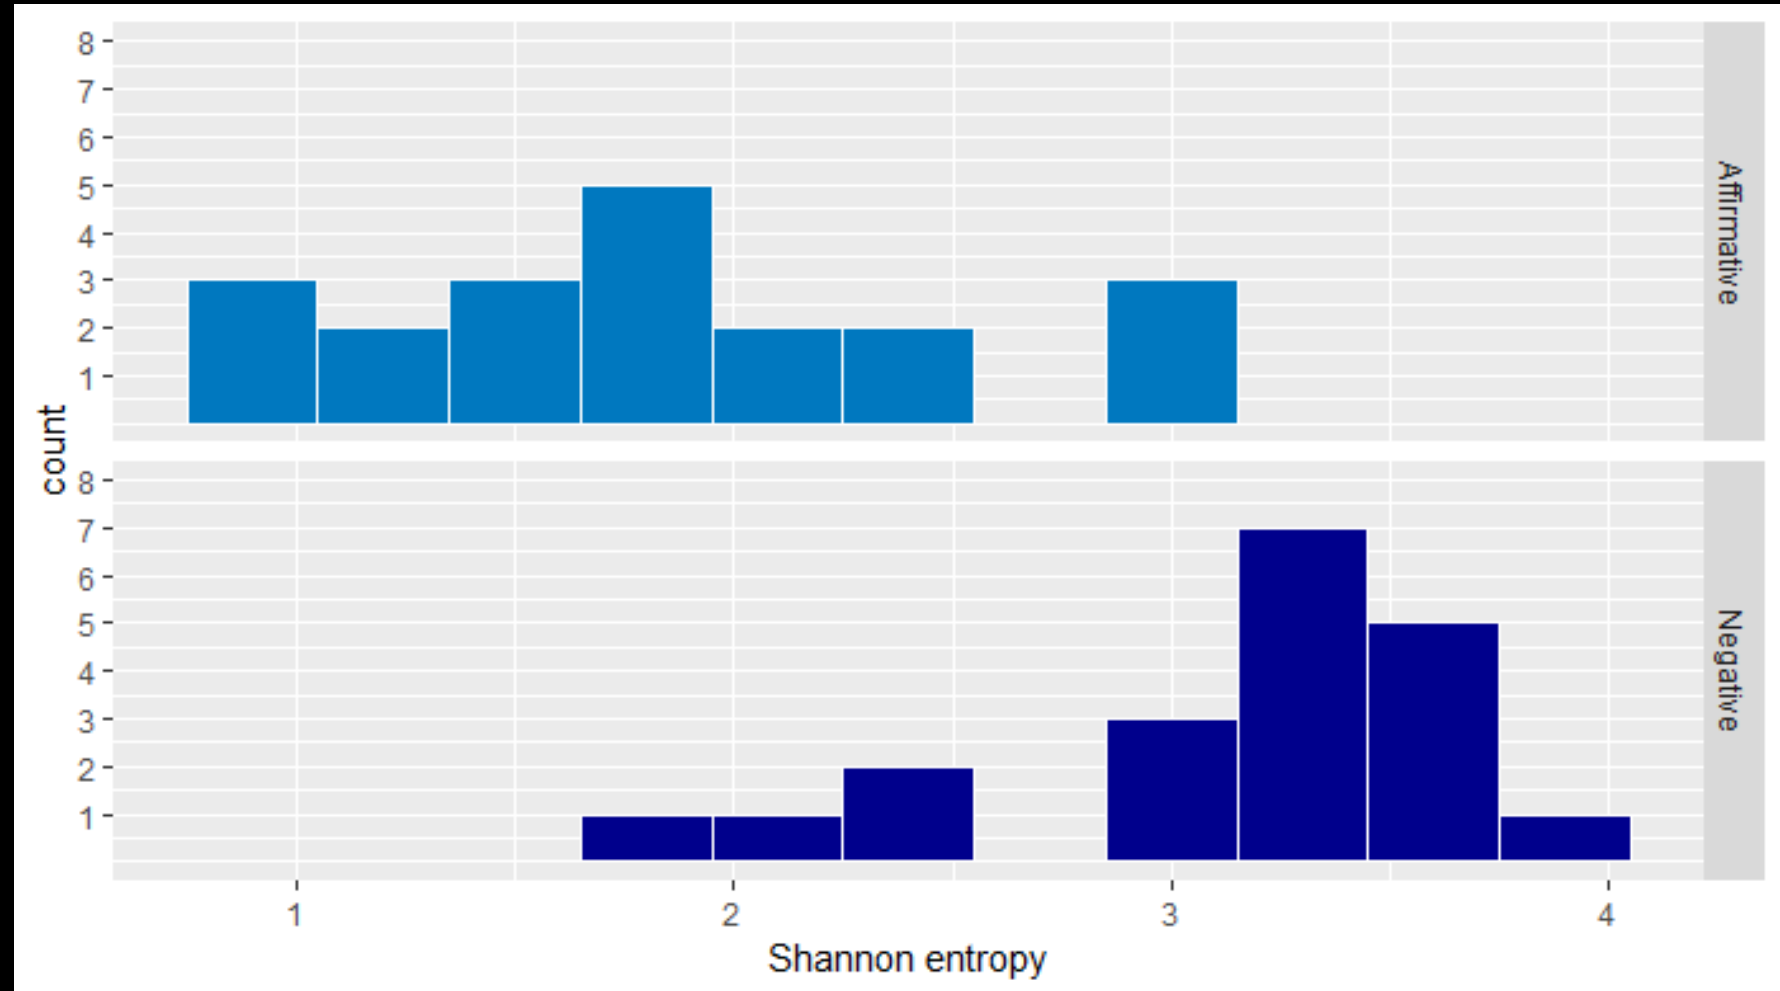
- Entropy is calculated for each item/condition pairing.
- The more variable the responses, the higher the entropy.

# Entropy quantifies variation

- Prediction: Responses to the negative condition should have a higher entropy than responses to the affirmative condition.
  - A sparrow is a \_\_\_\_ (bird, small bird, animal)
  - A sparrow is not a \_\_\_\_ (reptile, mammal, large bird, squirrel, lemonade, planet, dog, snail, armchair, etc.)

# Entropy results

A paired t-test revealed that, as predicted, the negative condition had significantly more entropy in its responses than the affirmative condition.





# Modal responses

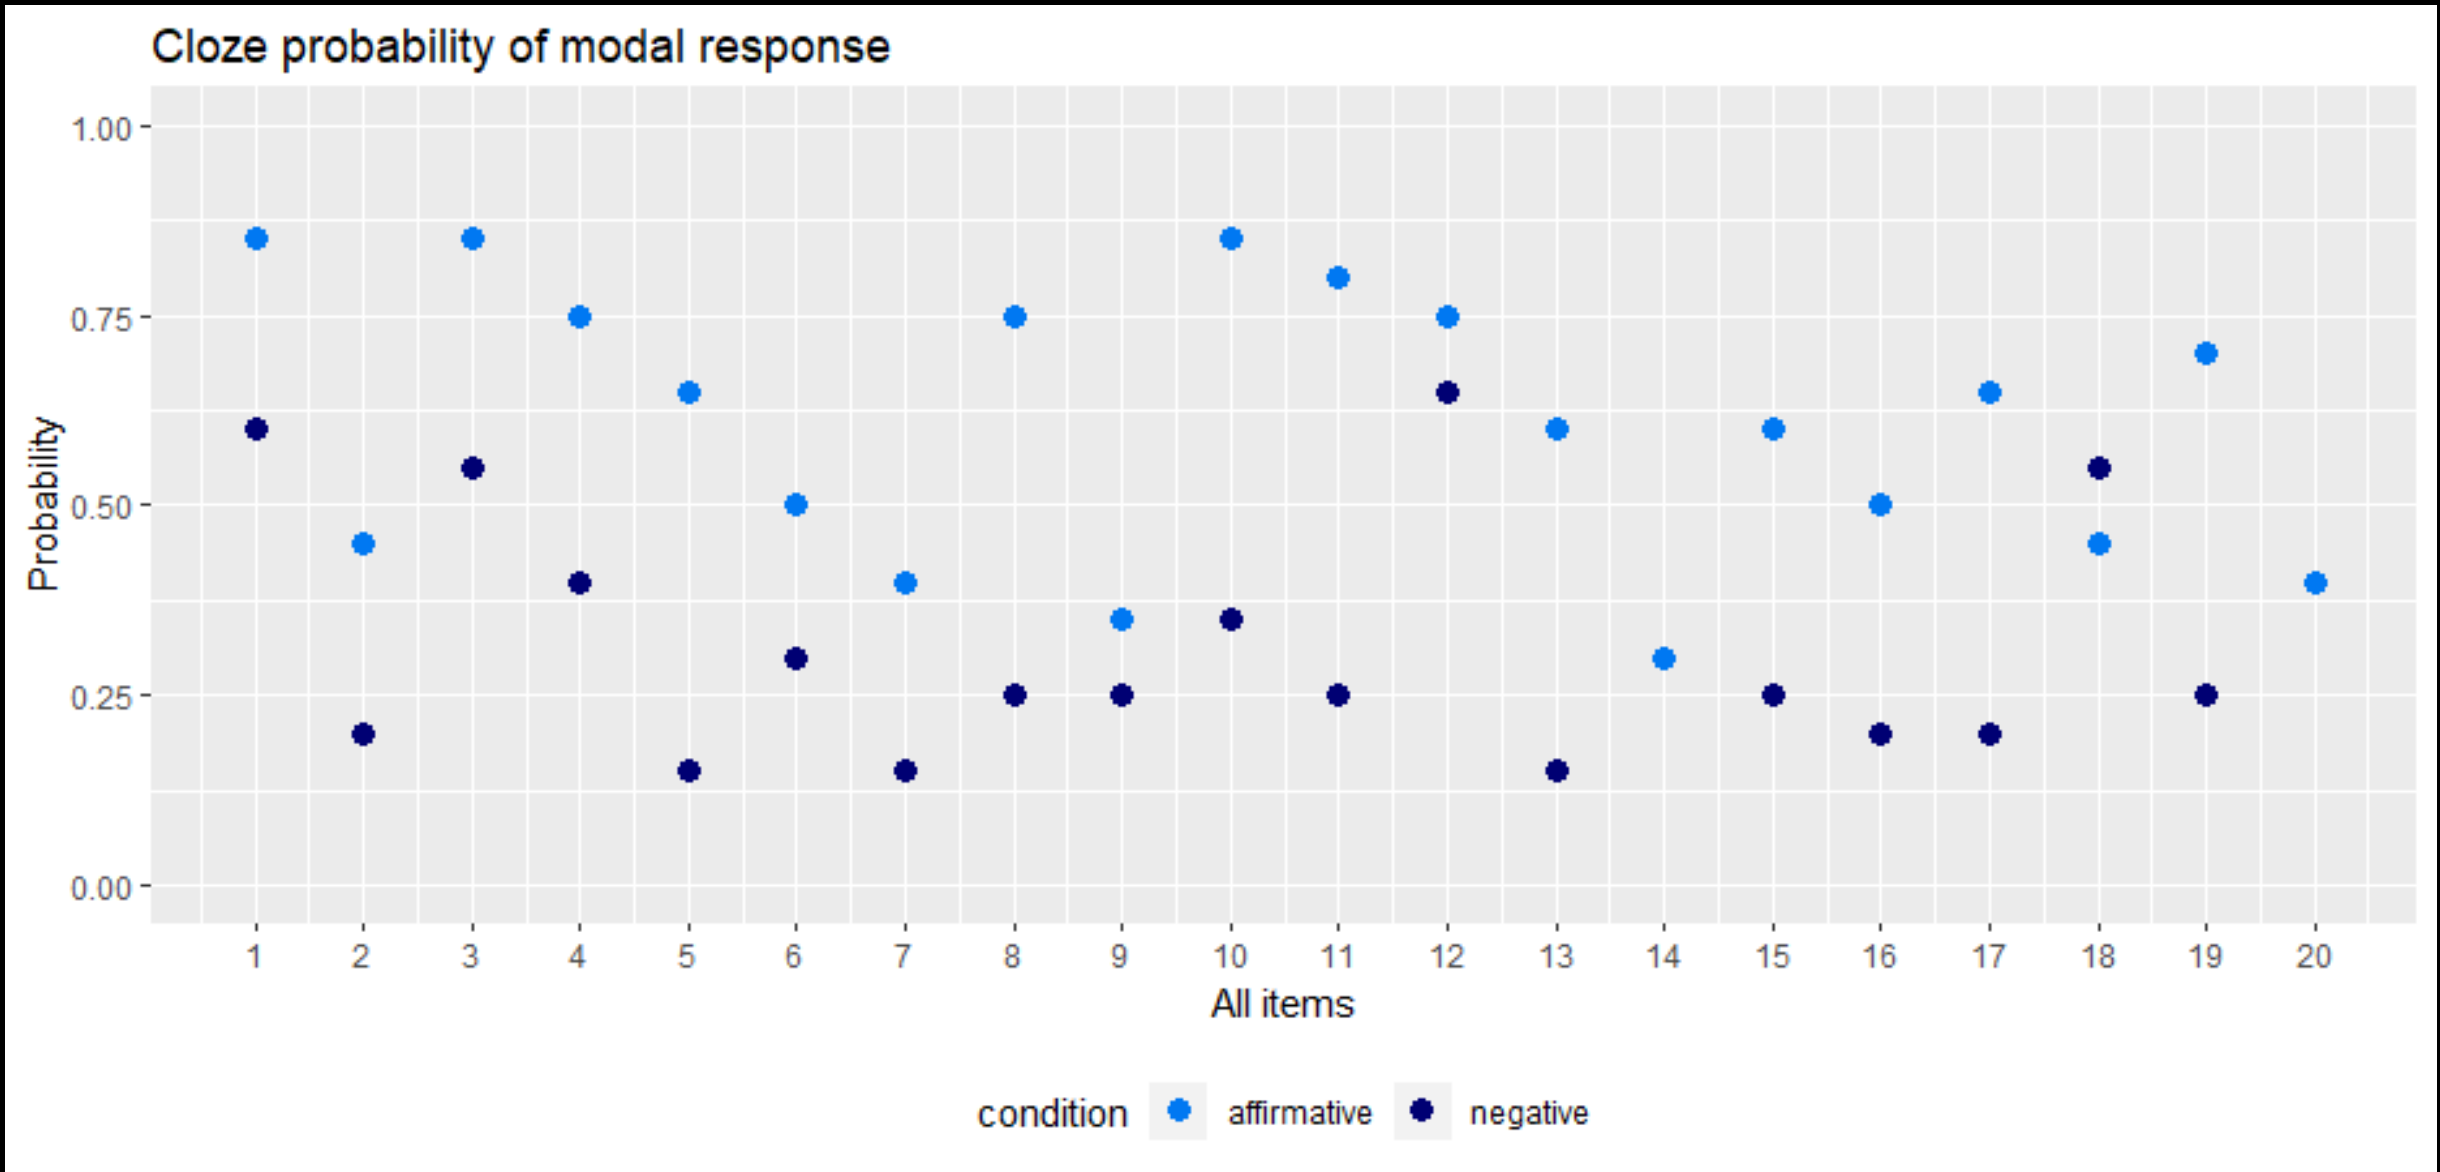
# Modal responses

- A given item's **modal response** is the response given by the most number of participants.
- If participants are sensitive to negation, they should give different modal responses for negative conditions than for affirmative ones.
- Furthermore, the **cloze probability** of a given modal response should be lower for negative conditions.

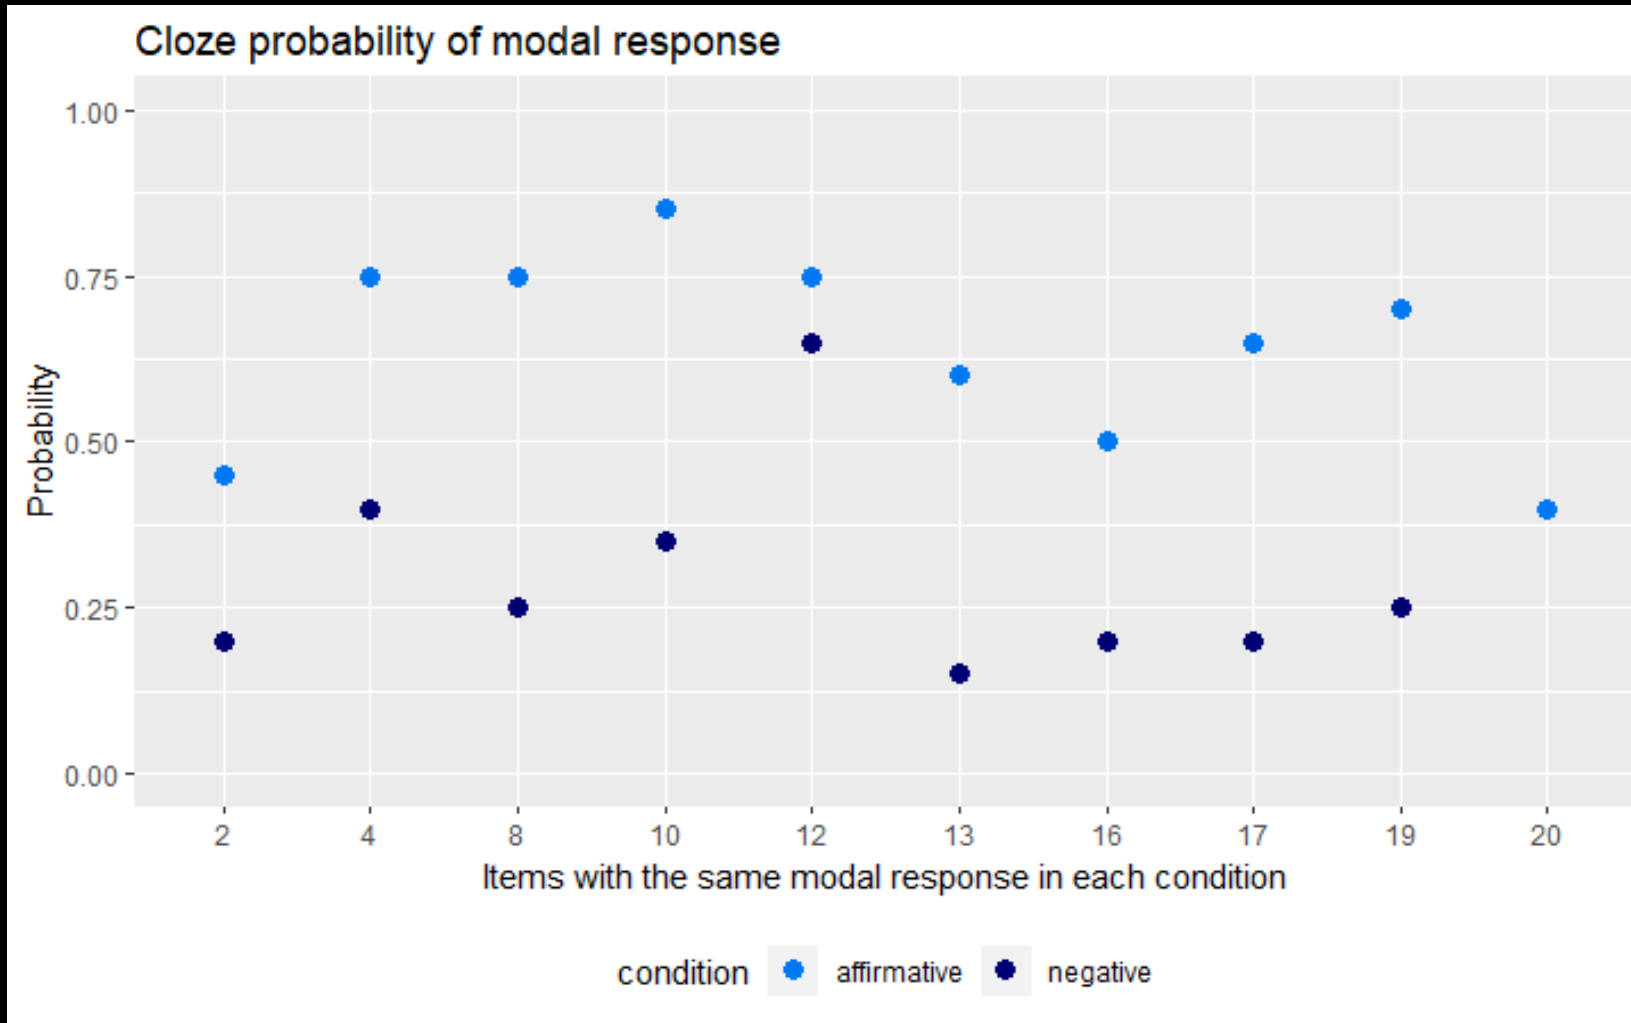
# Modal responses

- 10 of 20 items had the same modal response for affirmative & negative conditions (or, for multi-modal items, at least one response matched)
- In general, the cloze probability of modal responses was higher for affirmative preambles than negative ones.

# Cloze probability of modal response



# Cloze probability of modal response for items with the same modal response in each condition



Takeaway: Readers are sensitive to negation in an offline cloze norm task.

# Discussion

- The online results are a bit of a puzzle.
- Readers **are** sensitive to negation in offline responses. In other words, they notice negation and take it into account.
- Yet there is no evidence that they consider the semantic contribution of negation online.

Possible reasons for this pattern



# Semantic relatedness

- A large role of semantic relatedness could drown out any hint of an interaction.

# The role of context

- The role of the matrix/carrier sentence
  - More pragmatically neutral environment
  - The overall truth conditions of the sentence are unknown
  - However, this matrix sentence did **not** prevent negation from having an effect in offline measures, so it isn't the whole story
- Nieuwland & Kuperberg (2008): When negation is particularly informative in a given context, false critical words elicit a greater N400 than true critical words **for both affirmative and negative sentences.**
- In short, context is extremely important in semantics and for negation, and diverse experimental designs are important.

# A delay in processing and “operator blindness”

- There may simply be a delay in the integration of negation with other material in the sentence.
- Negation is not the only semantic operator that appears to not significantly influence early reading time measures (Mayer, Dillon, and Staub, 2019)
- There may be key differences in the effects or timing of more abstract semantic operators in online processing measures such as eye-tracking-while-reading or ERPs.

Thank you!

**This research was supported by NSF BCS 1732008 to Adrian Staub.**

