



# Separability and the Effect of Valence

An Experimental Study on Thick Concepts

Pascale Willemsen  
& Kevin Reuter

## Separability and the Effect of Valence An Empirical Study of Thick Concepts

Pascale Willemsen (Pascale.Willemsen@uzh.ch)    Kevin Reuter (Kevin.Reuter@uzh.ch)

Institute of Philosophy, University of Zurich, Zürichbergstrasse 43, 8044 Zurich, Switzerland

### Abstract

Thick concepts like *courage* and *intolerance* are at the heart of a variety of debates in linguistics, philosophy of language, and metaethics. Central to these debates is the question of how the descriptive and evaluative components of thick concepts are related. So far, no empirical data on how thick terms are used in ordinary language has been collected to inform these debates. In this paper, we present the first set of empirical studies, designed to investigate whether the evaluative component of thick concepts is communicated by means of semantics or pragmatics. Our studies not only favor the semantic view, they also reveal an effect of valence, indicating that people reason differently about positive and negative thick terms. Three follow-up studies were conducted to explain this effect. We conclude that the effect of valence is best accounted for by a difference in the social norms guiding evaluative language.

**Keywords:** Thick concepts; moral judgments; experimental metaethics

### Introduction

Philosophers and linguists usually distinguish two types of evaluative concepts: “thin” and “thick” concepts (Eklund, 2011; Väyrynen, 2019). Thin concepts evaluate an object as, for instance, “permissible”, “right”, “wrong”, “good”, “bad”, or “blameworthy”, yet they do not explicate in what way the object is right or wrong. If a speaker evaluates an instance of lying as wrong, they convey no information as to why they think so. The speaker might think that people have a right to be told the truth, that it ruins friendships, etc. Thick concepts do not merely evaluate, they also provide information on why the entity is evaluated. Typical examples are ethical terms, such as “rude”, “cruel”, “courageous”, or “trustworthy”. Calling an agent courageous evaluates them positively for being willing to take risks – “reckless” also ascribes willingness to take risks yet assigns a negative evaluation to it.

While there is widespread consensus that thick concepts form an additional class of concepts, a heated disagreement exists over the way in which the evaluative and the descriptive component of thick terms are connected. According to one group of researchers, the evaluative component of a thick term is part of its semantic meaning; according to another, the evaluation is not part of the meaning but conveyed through pragmatic means. Arguments in favor of either position heavily rely on linguistic intuitions about how thick terms can be used. Such intuitions often circle around the question of whether the evaluation of a thick term

can be cancelled without yielding contradiction. Whether this is possible is not only relevant for the linguistic debate about thick terms. Assumptions about the nature of thick terms provide the argumentative cornerstones in metaethical and normative-ethical debates as well. Therefore, by putting these assumptions to the test, we can provide a more solid basis for theorizing about thick concepts in various disciplines. In this paper, we present the first empirical data of this sort.

### Separating the Evaluative from the Descriptive?

Thick terms, such as “honest”, “friendly”, “cruel” or “rude”, do not only evaluate an entity, they further describe *in virtue of what* this entity is evaluated as positive or negative. This descriptive richness is what distinguishes thick from thin concepts. And it is the evaluative component that distinguishes thick from merely descriptive concepts.

Some philosophers explain this descriptive richness by assuming that thick terms are a combination of thin evaluations plus some description. According to this position, thin terms express the basic moral attitudes a speaker can have. Add some description to this attitude and the result is a thick term. According to such a view, thick terms can be, at least in principle, divided into two distinct components, namely the evaluative and the descriptive Hare, 1952; Blackburn, 1992; Elstein & Hurka 2009). Call this the *Separabilist View* (Väyrynen, 2019). Separabilist Views often assume that the descriptive and the evaluative are connected by pragmatic means, for instance, by conversational implicature. Conversational implicatures are part of the speaker meaning and need to be inferred beyond what is literally said (Grice, 1985). By saying that an agent is rude, one ascribes some descriptive properties, and one further communicates the implicature that the agent is bad in virtue of having these properties. However, as other conversational implicatures, the negative evaluation can be cancelled without creating a contradiction. Therefore, a speaker who utters “What Tom did was rude, but by that I’m not saying something negative about Tom” makes a felicitous statement.

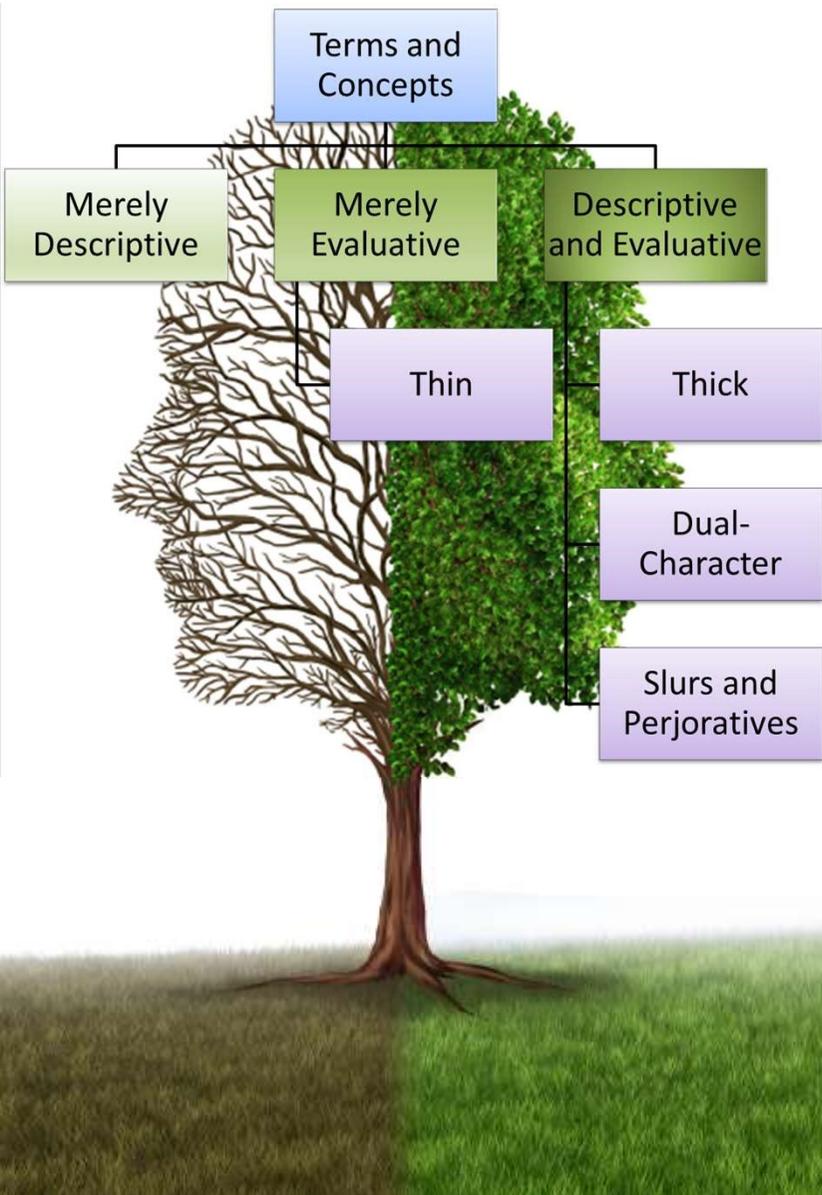
The contrary position, the *Inseparabilist View*, denies that thick terms can be further divided into distinct components and holds that thick terms are basic (Williams, 1985; Putnam, 2002; Kirchin, 2010; Roberts, 2011). Those who argue for Inseparabilism typically argue that thick terms cannot be further reduced because their evaluation is part of the

# Agenda

---

1. What are thick concepts?
2. Why are they philosophically interesting?
3. An empirical study on cancellability
4. Three potential explanations of a surprising result
5. Lots of open questions





# Thick and Thin Terms & Concepts

- **Descriptive:** round, green, short-sighted
- **Thin:** right, wrong, good, bad
- **Thick:** discreet, cautious, industrious, chaste, just, grateful



# Thick Terms & Concepts

- **Epistemology:** open/close-minded, gullible, critical, reflective, rational, dumb
- **Aesthetics:** banal, gracious, creative, elegant, dainty
- **Ethics:** friendly, helpful, generous, cold-hearted, cruel

# Note: Things are not so simple

- If it is true that „ought“ implies „can“, it seems that even ought is thick in some sense

# Thick vs. Thin

- Allan Gibbard: thick terms and concepts praise or condemn an action as having certain properties
- Jonathan Dancy: thick terms and concepts have an evaluative dimension which is inseparable from their descriptive dimension
- Bernard Williams: thick terms and concepts are both 1) world-guided and 2) action-guiding



# What's the Connection?

## Inseparabilists

- Thick terms and concepts are inherently evaluative
- The evaluation is part of the semantic meaning
- It is impossible to separate the evaluative from the descriptive

McDowell 1979, 1981; Williams 1985; Dancy 1996, Putnam 2002, Kirchin 2010

## Separabilists

- Thick terms and concepts are descriptive, yet they presuppose, implicate, etc to convey evaluative content
- The evaluation is not part of the meaning
- It is possible to separate the evaluative from the descriptive

Hare 1952,1981, Blackburn 1992, Elstein and Hurka 2009, Väyrynen 2013

# A Thick Job in Moral Philosophy

- Thick concepts have been used to argue against
  - Non-Cognitivism
  - The (ontological) fact/value distinction
  - Ethical reductionism
- And to argue for
  - Moral Objectivism
  - The nature of reasons for actions (Internalism vs. Externalism)
- Are essential for Virtue Ethics and other character-focused ethics



# A Thick Job in Moral Philosophy

- Thick concepts have been used to argue against
  - Non-Cognitivism
  - The (ontological) fact/value distinction
  - Ethical reductionism
- And to argue for
  - Moral Objectivism
  - The nature of reasons for actions (Internalism vs. Externalism)
- Are essential for Virtue Ethics and other character-focused ethics

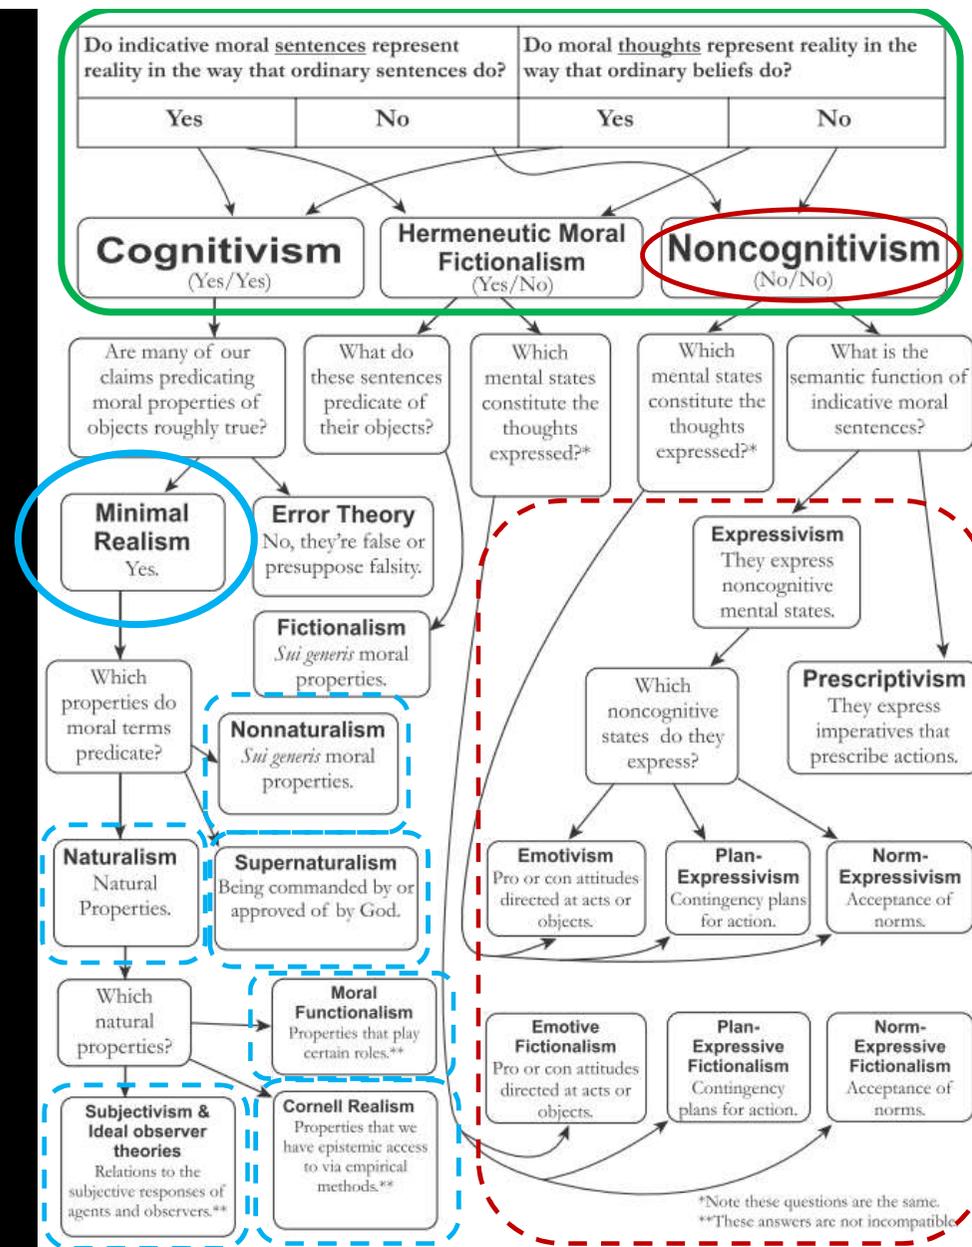


# Cognitivism vs. Non-Cognitivism

## ■ Cognitivism:

- Moral concepts are genuine concepts
- They pick out a real feature of the world
- Values are in the world
- SELFISH picks our selfishness, HONEST picks out honesty, etc.
- Moral sentences can be true or false

18.02.2020

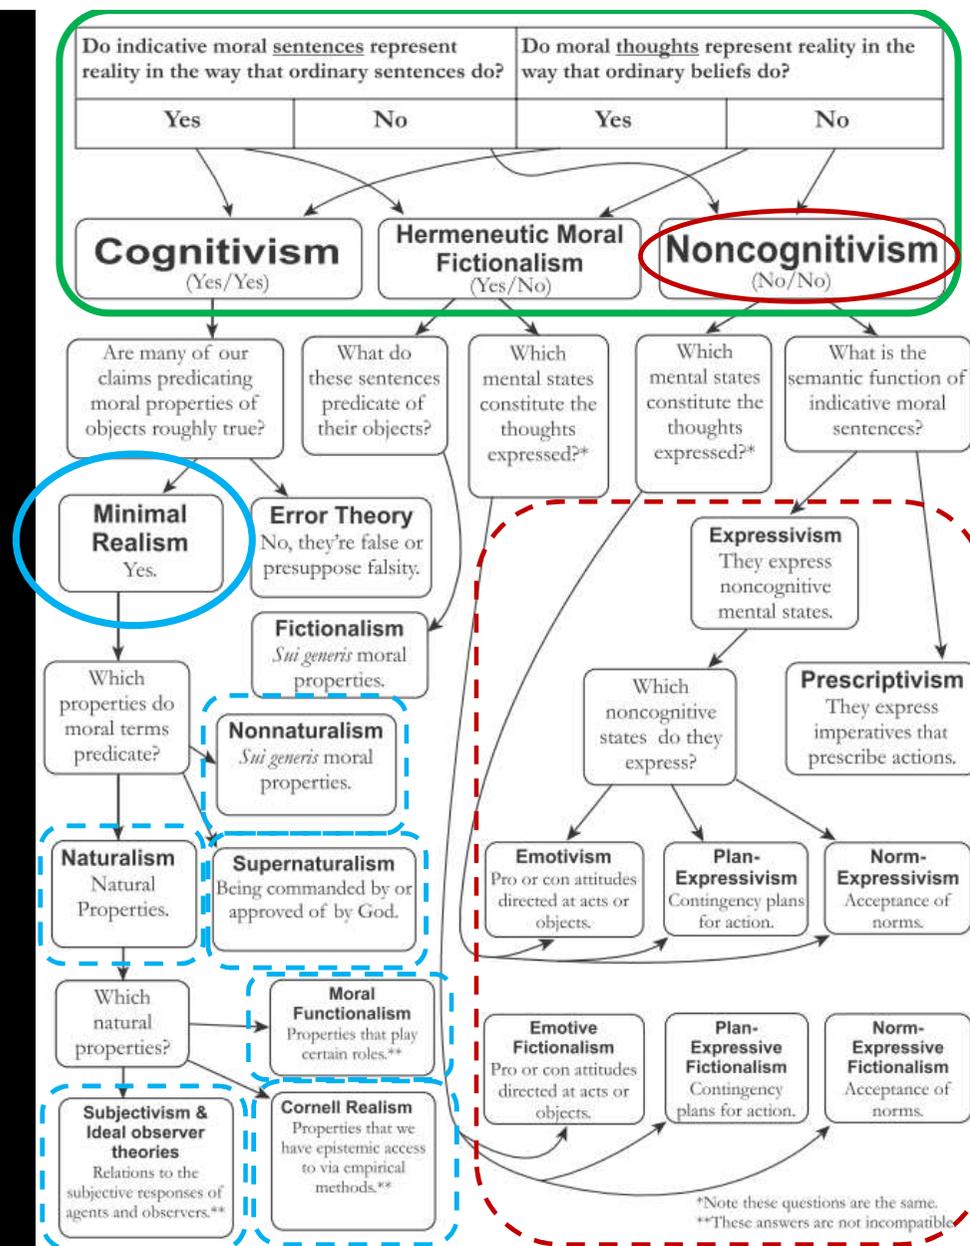


# Cognitivism vs. Non-Cognitivism

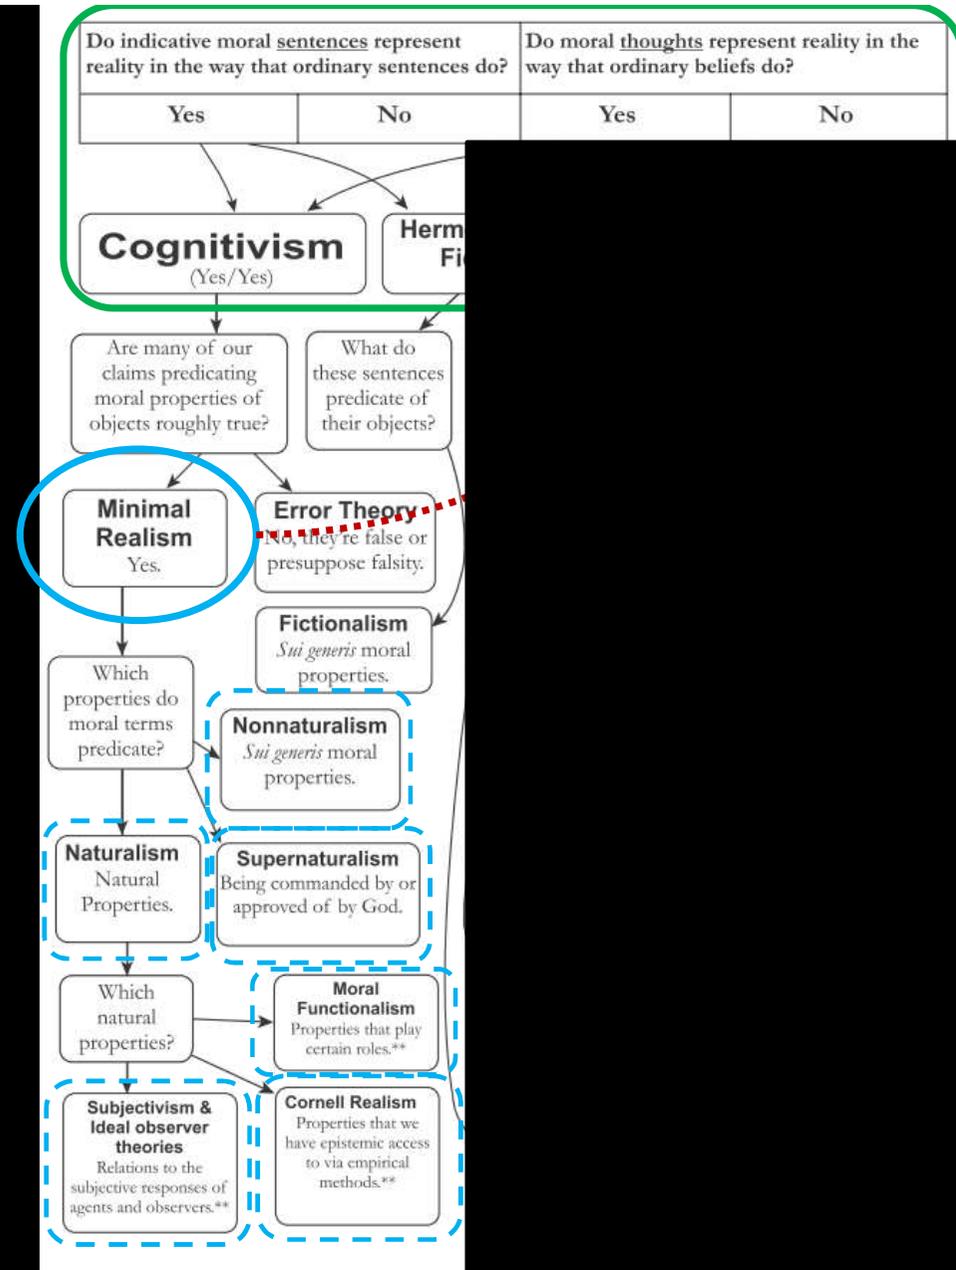
## ■ Non-Cognitivism

- Deny that moral concepts are genuine concepts
- Moral concepts do not represent the world
- Yet, they are a convenient tool for communication
- Values are not in the world
- Moral sentences cannot be true or false

18.02.2020



# Cognitivism vs. Non-Cognitivism



- Theories in metaethics presuppose the correctness of the semantic or the pragmatic understanding of thick concepts
- The use of language seems like an obvious field of empirical research
- Testing whether this is indeed correct will have significant impact on those debates

18.02.2020





## Pre-Registered Experiment

# Methods

- Pre-registration: [osf.io/9pbq2](https://osf.io/9pbq2)
- 200 participants recruited on Mturk, 50 subjects per condition
- 4 × 1 between-subject design with
  - IV: Relation
    - Semantic Entailment
    - Conversational Implicature
    - Thick Positive Concepts
    - Thick Negative Concepts
  - DV: Contradiction

## Thick Positive

honest,  
generous,  
courageous,  
friendly

Please imagine that Sally said the following sentence:

**"Kathy's behavior last week was friendly, but by that I am not saying something positive about Kathy's behavior that day."**

Does Sally contradict herself?

1 = definitely not	2	3	4	5	6	7	8	9 = definitely yes
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Run,  
Widow,  
Couch,  
Lake

Semantic Entailment

## Thick Negative

intolerant,  
rude,  
cruel,  
egoistic

Please imagine that Sally said the following sentence:

**"Amy's behavior last week was egoistic, but by that I am not saying something negative about Amy's behavior that day."**

Does Sally contradict herself?

1 = definitely not	2	3	4	5	6	7	8	9 = definitely yes
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Please imagine that Sally said the following sentence:

**"I am hungry, but by that I am not saying that we should get something to eat."**

Does Sally contradict herself?

1 = definitely not	2	3	4	5	6	7	8	9 = definitely yes
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Hungry,  
Dark,  
Door,  
Chocolate

Conversational Implicature



# Predictions

---

- There is a significant difference between SE and CI, such that Contradiction ratings are significantly higher for SE than for CI.

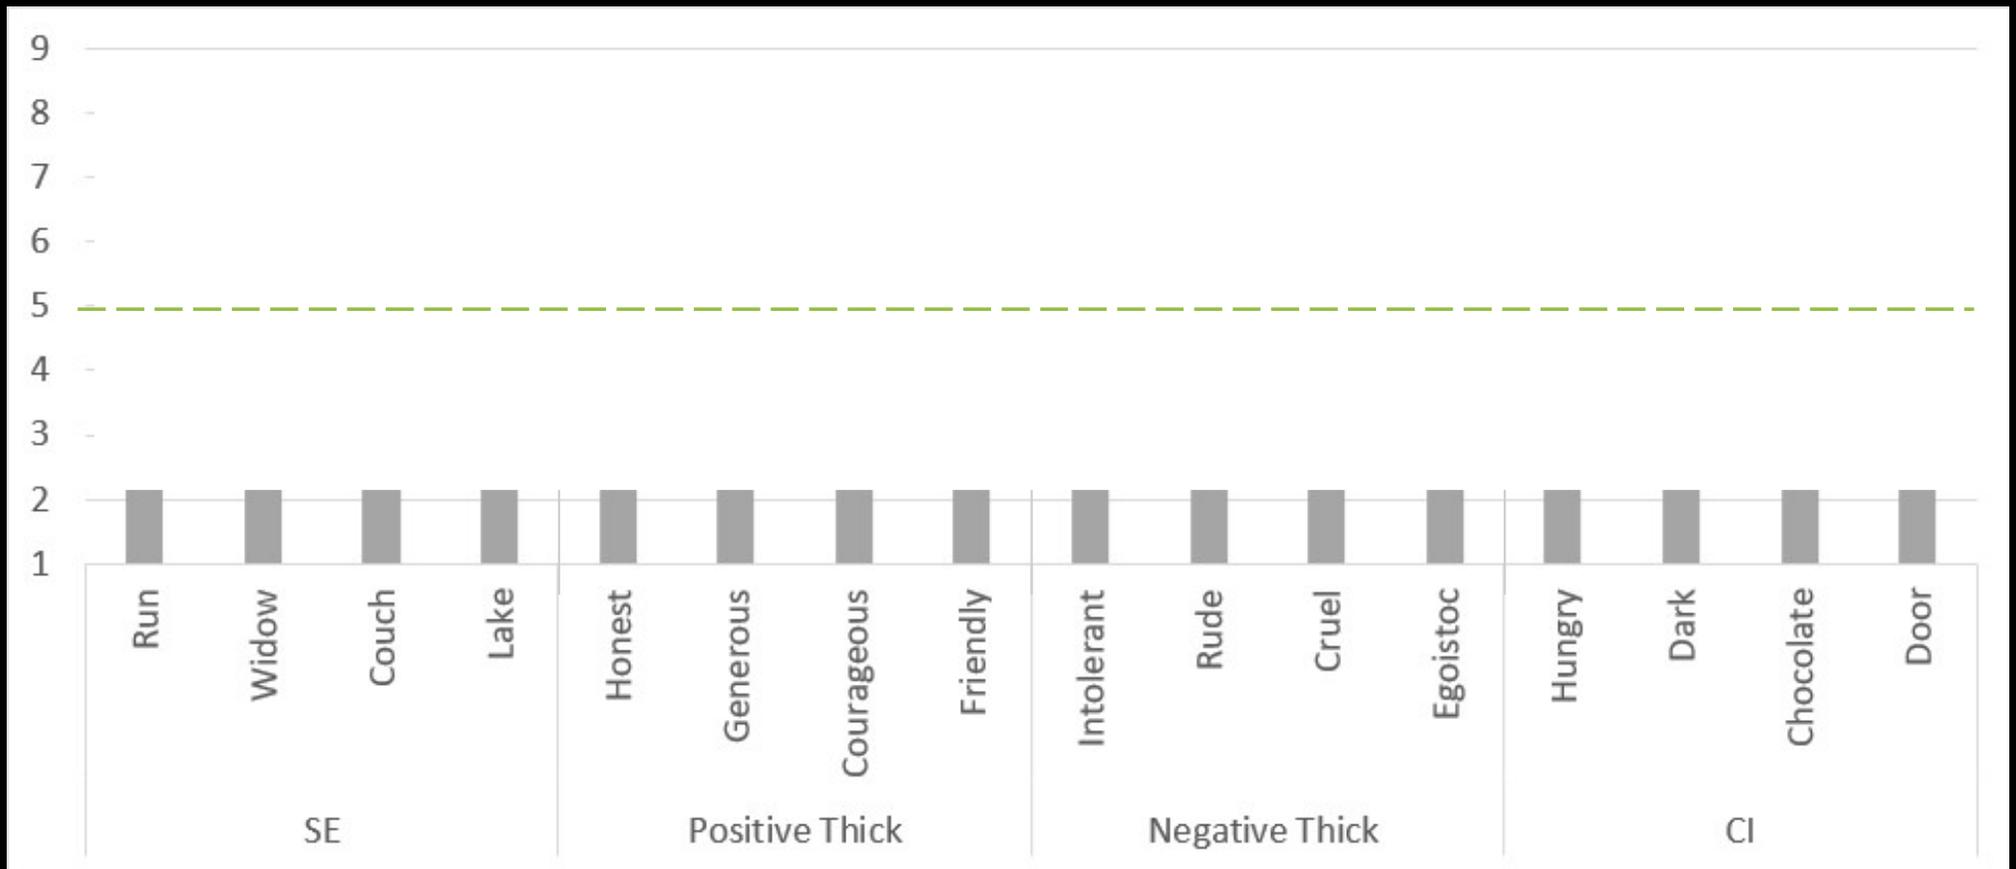
If the Semantic View is correct,

- There is no significant difference between SE and TC

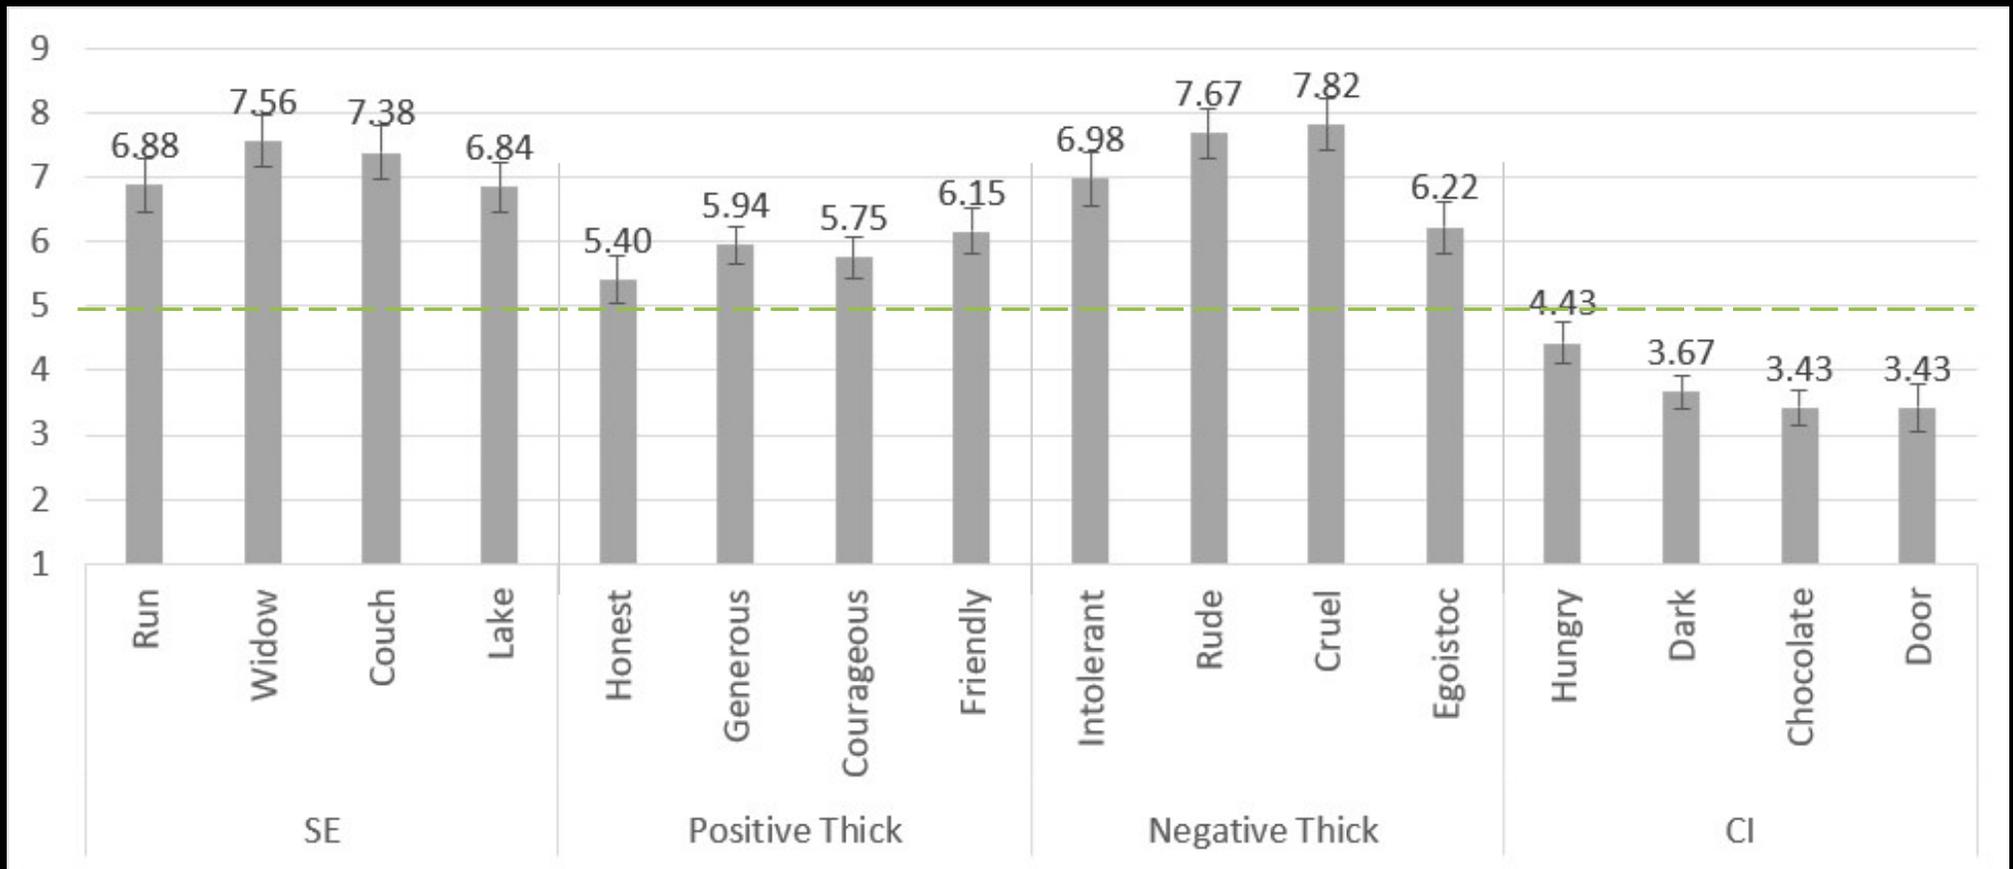
If the Pragmatic View is correct,

- There is no significant difference between CI and TC

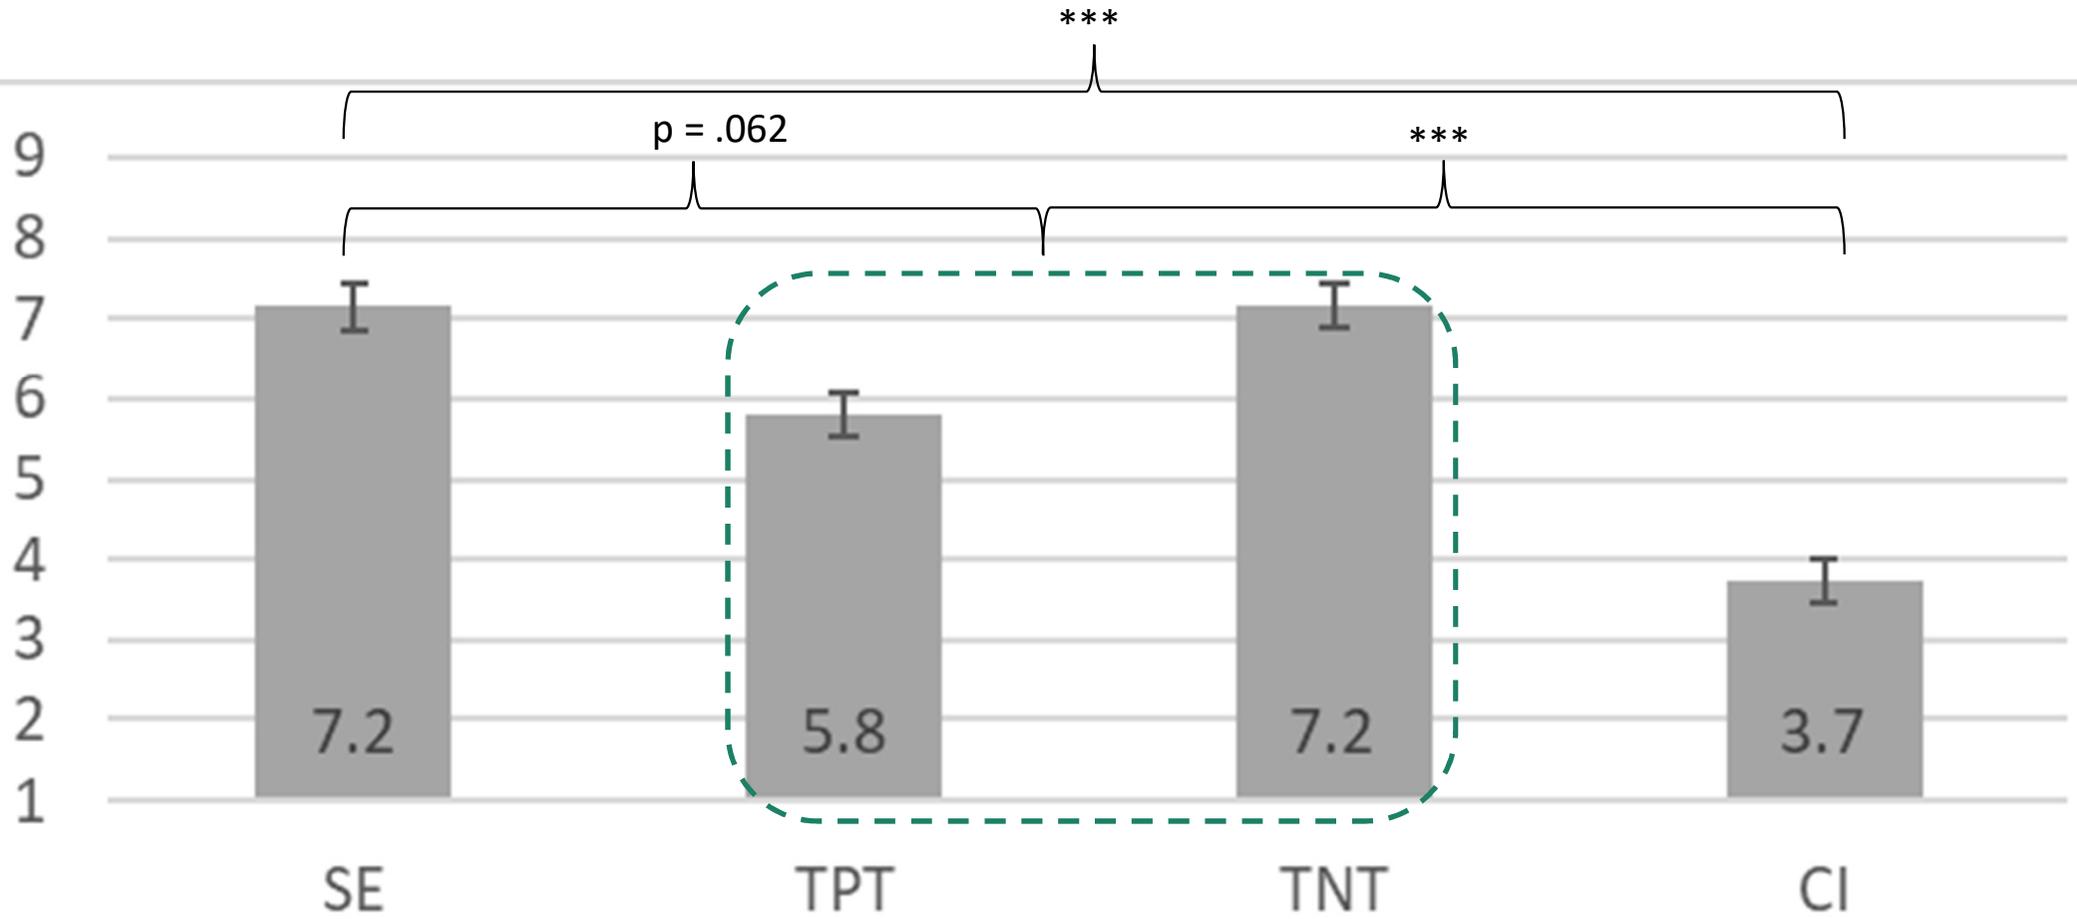
# Results



# Results



# Results





## Predictions

- There is a significant difference between SE and CI, such that Contradiction ratings are significantly higher for SE than for CI.

If the Semantic View is correct,

- There is no significant difference between SE and TC

If the Pragmatic View is correct,

- There is no significant difference between CI and TC

# Discussion

---

## We must not overestimate the data!

- We did only test for Conversational Implicatures, not Conventional Implicatures
- We only looked at a small collection of terms
- Do laypeople's contradiction ratings track semantic contradiction?
- Did laypeople make additional assumptions about the speaker's intentions?



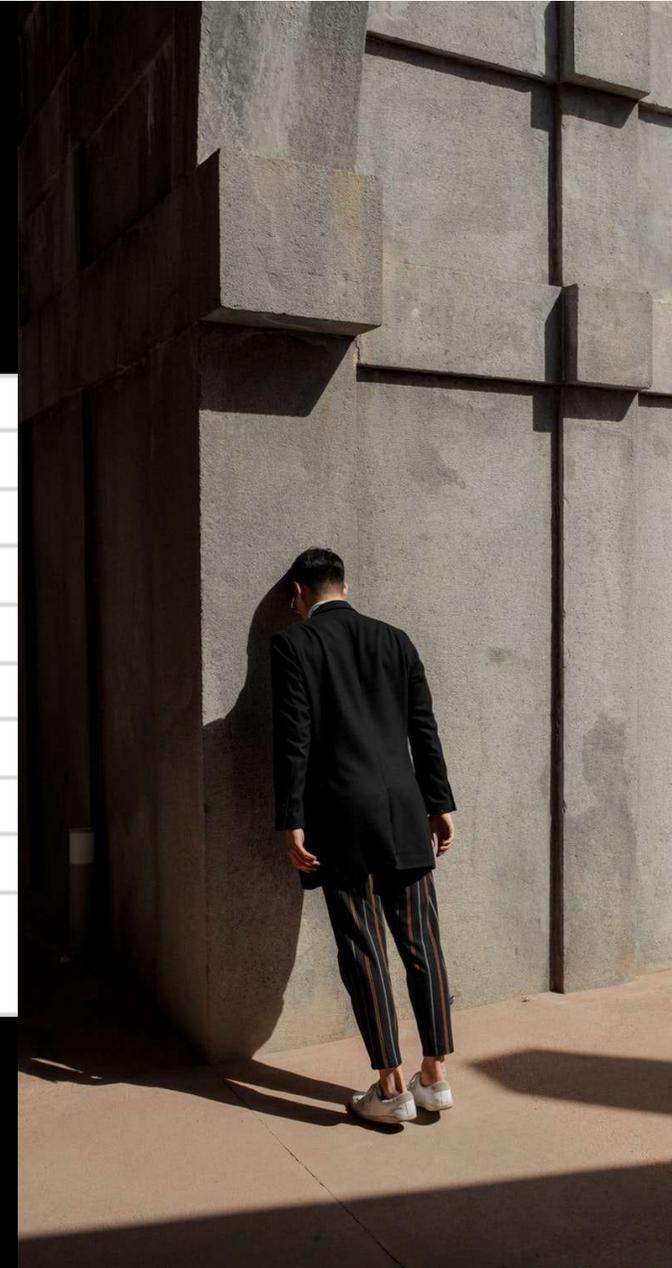
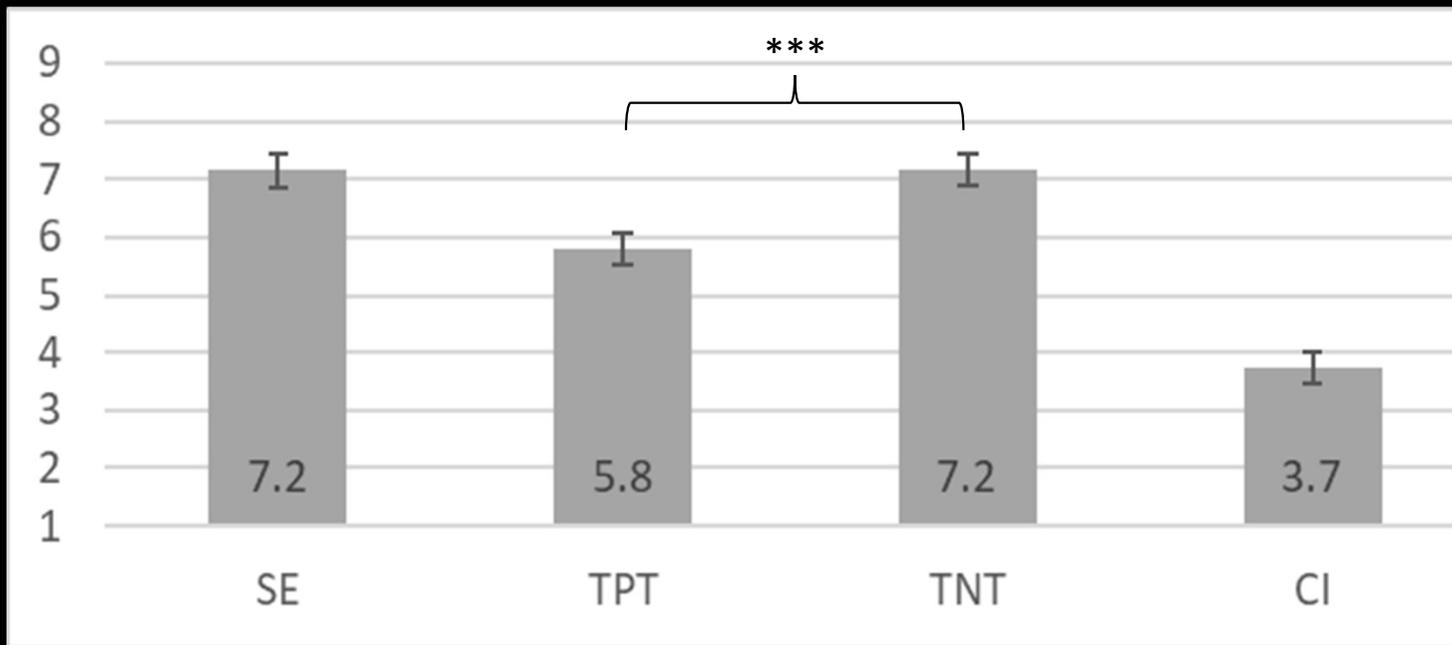
# Discussion

---

Why do Positive and Negative Thick Terms behave differently?



# Results



# Discussion

---

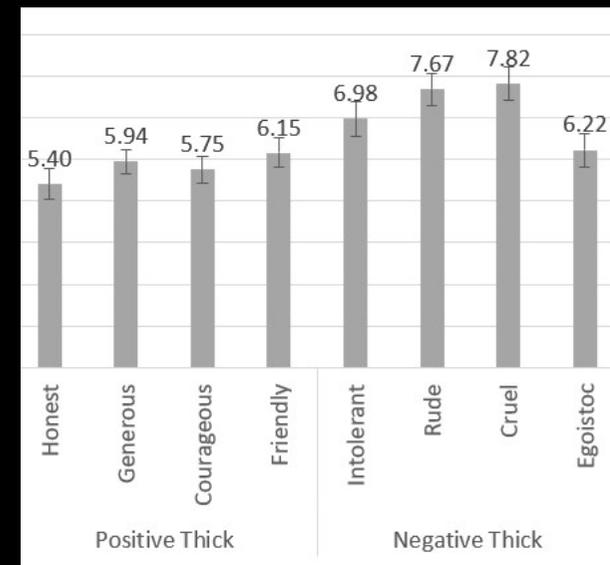
Why do Positive and Negative Thick Terms behave differently?

1. Differences in moral evaluation
2. Availability of counterexamples
3. Differences in the social norms guiding evaluative language



# Explanation 1: Differences in moral evaluation

- The terms that we tested do not only differ in the valence (positive or negative)
- They might also differ in how strongly they evaluate
- The stronger the evaluation, the more contradictory does it seem to cancel it
- Prediction:  
*Negative terms should be rated more negatively than positive terms are rated positively.*



# Explanation 1: Results

- 400 participants recruited on Mturk, 25 subjects per condition
- 2 × 2 between-subject design
  - Question (Behavior, Sentence)
  - Valence (Positive, Negative)
- All native speakers of English

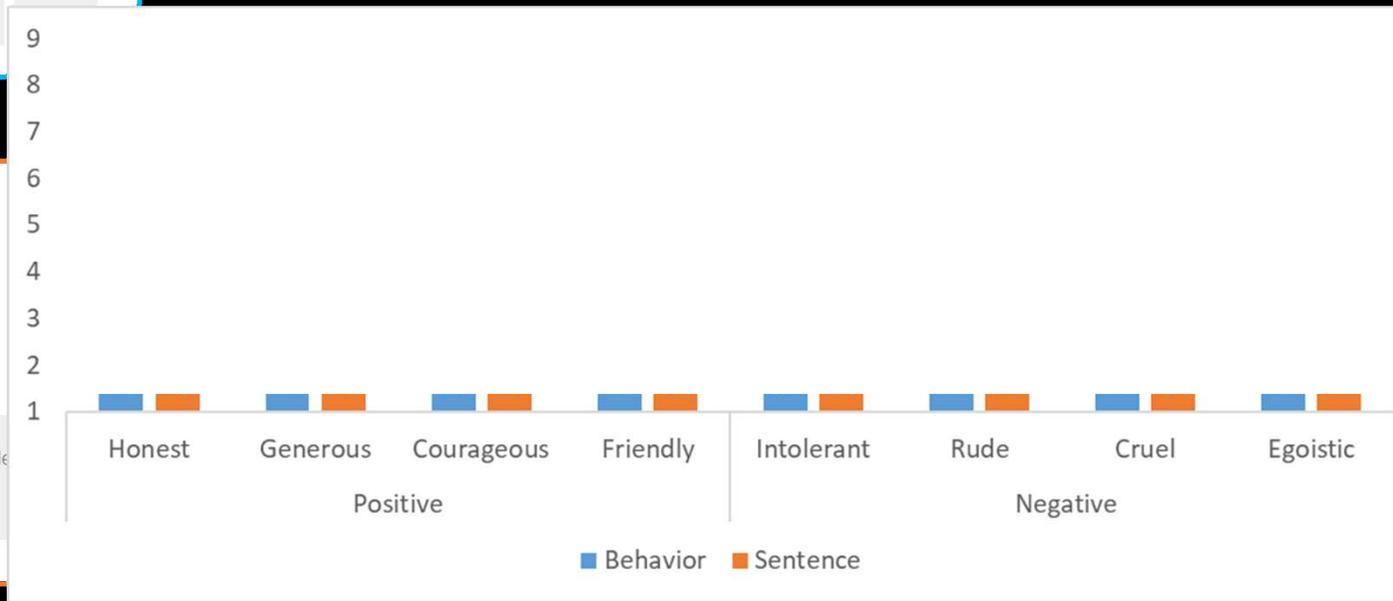
How bad is it if a person's behavior is rude?



Suppose that Sally said the following thing about Tom:

**"What Tom did was cruel."**

Is this a negative statement about Tom?



# Explanation 1: Results

- 400 participants recruited on Mturk, 25 subjects per condition
- 2 × 2 between-subject design
  - Question (Behavior, Sentence)
  - Valence (Positive, Negative)
- All native speakers of English

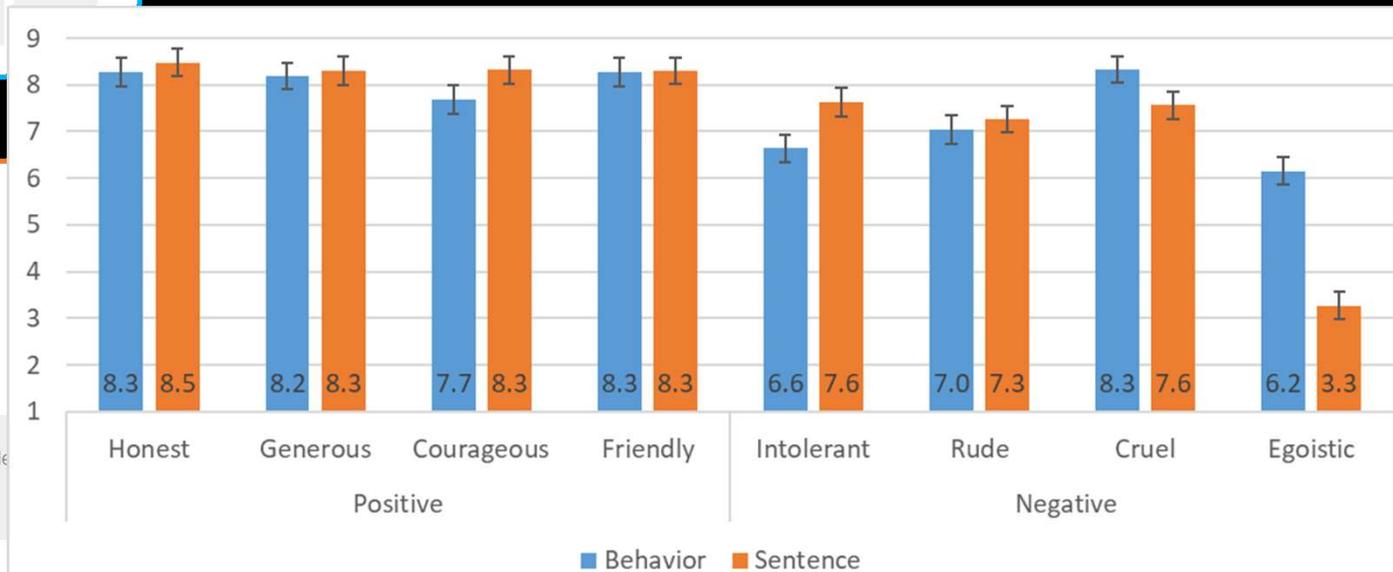
How bad is it if a person's behavior is rude?



Suppose that Sally said the following thing about Tom:

**"What Tom did was cruel."**

Is this a negative statement about Tom?

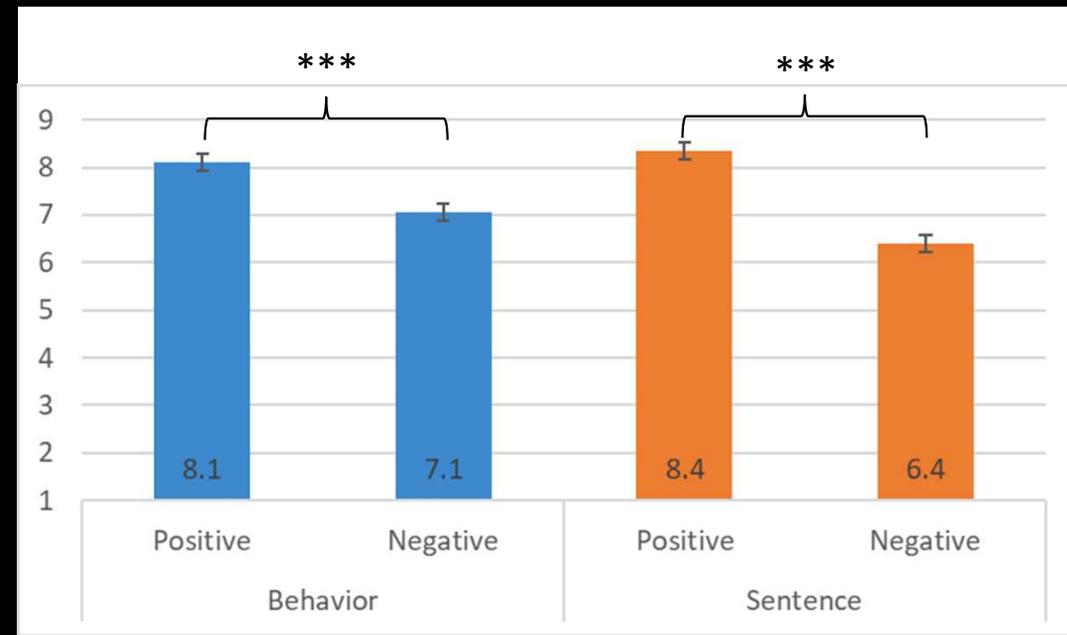
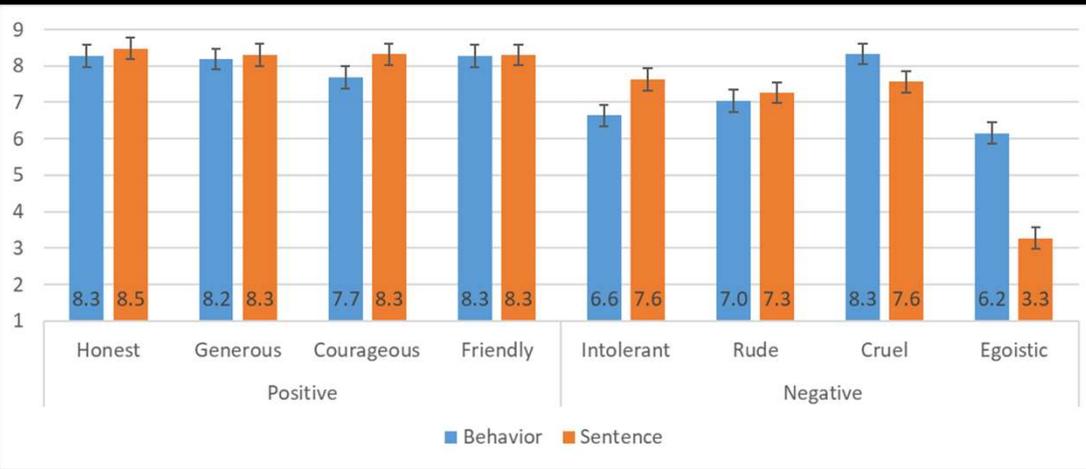


# Explanation 1: Results

- 400 participants recruited on Mturk, 25 subjects per condition
- 2 × 2 between-subject design
  - Question (Behavior, Sentence)
  - Valence (Positive, Negative)
- All native speakers of English

## Prediction:

*Negative terms should be rated more negatively than positive terms are rated positively.*



## Explanation 2: Availability of Counterexamples

- Maybe the results are an artefact of the experimental design
- In one condition, reasoning about counterexamples is triggered

Please imagine that Sally said the following sentence:

**"Kathy's behavior last week was friendly, but by that I am not saying something positive about Kathy's behavior that day."**

Does Sally contradict herself?

1 = definitely not	2	3	4	5	6	7	8	9 = definitely yes
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

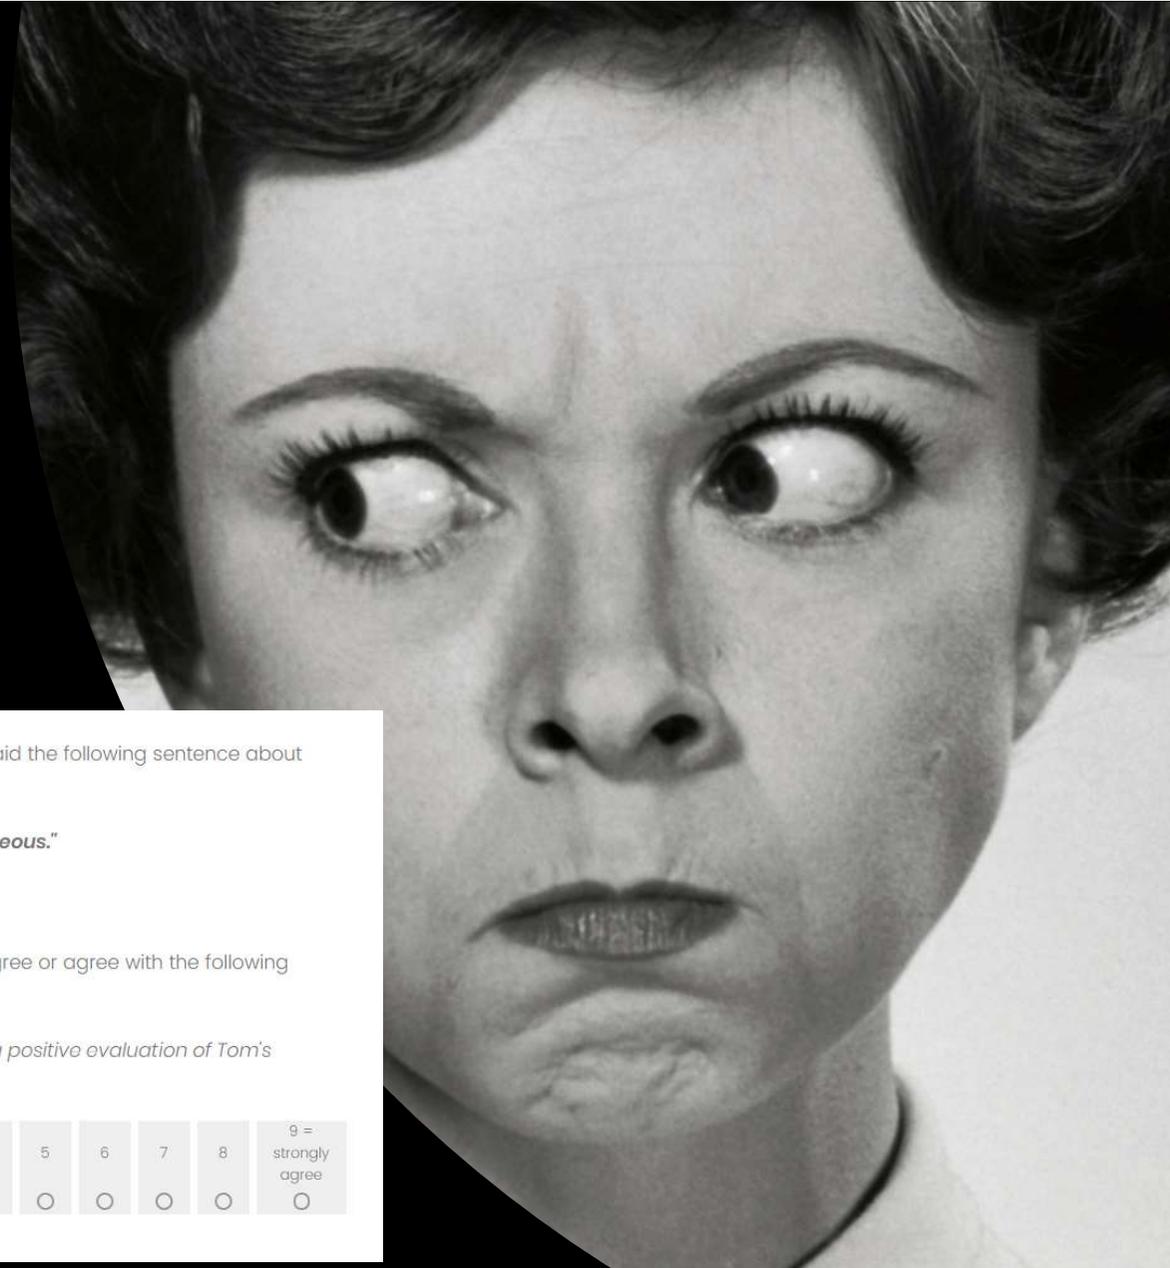
Please suppose that Sally said the following sentence about Tom's behavior:

**"What Tom did was courageous."**

To what extent do you disagree or agree with the following claim:

*Sally committed herself to a positive evaluation of Tom's behavior.*

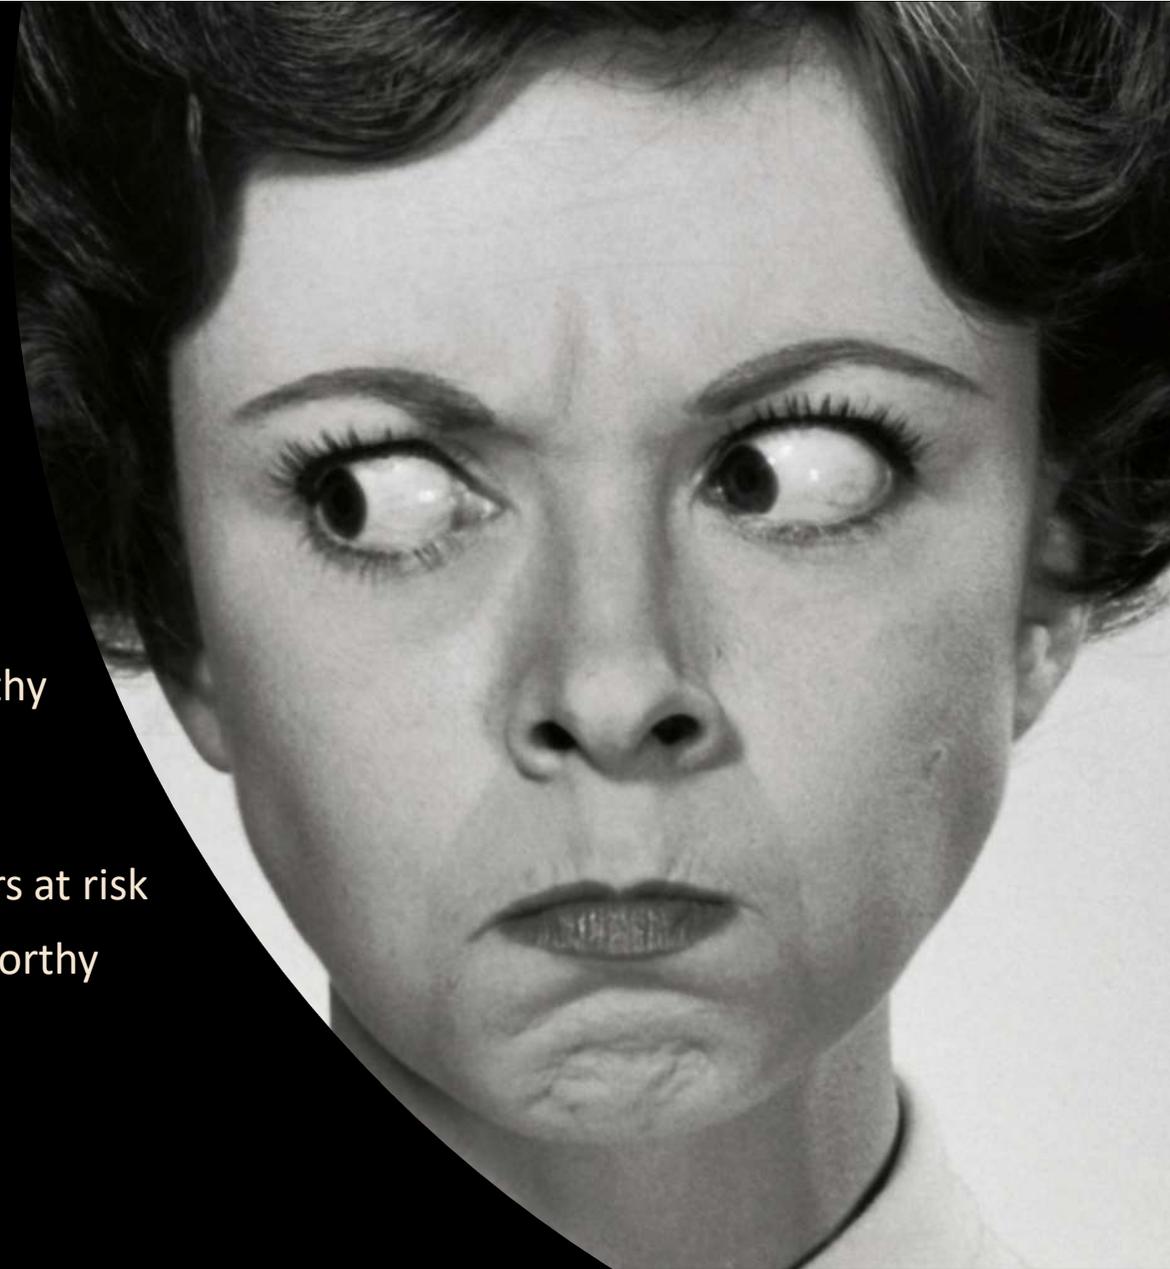
1 = strongly disagree	2	3	4	5	6	7	8	9 = strongly agree
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>



## Explanation 2: Availability of Counterexamples

---

- You are good, but for the wrong reasons
  - Courageous: show-off, want to make others look bad
  - Friendly: manipulative, devious, not trustworthy
- You are overdoing it
  - Honest: mean, offensive, indiscreet
  - Courageous: reckless, foolhardy, putting others at risk
- The the consequences render your action blameworthy



## Explanation 2 – Positive Thick Concepts

- While courageous acts are normally associated with positive things, it's possible to do something negative that is also courageous
- The person could have faked it (friendly)
- Mostly, 'friendly' means nice and good things. But if you were friendly to, say, a Nazi...is that a good thing? Mostly not, probably.
- Context may present the courage displayed as unwise
- Friendliness is not necessarily positive if the person exhibiting it has nefarious intent.
- Honesty isn't always good. He could have been honest and saying something like, "I find your face to be really ugly." It could be honest, but it wouldn't be a good thing to say, and people wouldn't think he was a good person for his honesty.
- Courageous behavior might mean they are being a bit more forward when they should really hold back more.
- Because you can be honest in a negative way. (sometimes the truth hurts, etc..)
- Jim could have been honest about something that should not have been said to be polite. In this case, it is not really positive that he was honest.

## Explanation 2 – Negative Thick Concepts

- You should not be cruel to anyone in any circumstances. That is inherently negative.
- Rudeness is always negative. There's no positive way to be rude.
- being cruel at any time is a negative thing
- being rude is generally considered negative so this would be a contradiction
- In my world, being rude is not a positive behavior and I could think of no world in which it would be. I may lack sufficient imagination.
- Intolerant is a negative adjective.
- i'm not sure how it's possible to be intolerant and have that not be a negative thing
- Well rude and negative is basically the same thing.
- Intolerance is a negative thing, so you can't say that someone is intolerant and then claim you weren't saying something negative about them.

## Explanation 3: Availability of Counterexamples

- If the effect arises because of a specific experimental design, using a different design should make the effect disappear
- To test this, we used an even simpler design
- We contrasted a thick term with its evaluation, using „but“
- Prediction:  
*Acceptability ratings for thick terms should be equally low for thick positive and negative terms*





# Explanation 3: Social Norms and Moral Language

---



- The norms as to when we may say something evaluative differ
  - There is no harm in communicating something positive one doesn't want to commit to
  - There is a lot of harm in communicating something negative one doesn't want to commit to

## Two Social Rules:

1. **Positive thick terms may be used** even if the speaker does not mean to evaluate by doing so
2. **Negative thick terms may not be used** if the speaker does not mean to evaluate by doing so

# Explanation 3: Social Norms and Moral Language

---



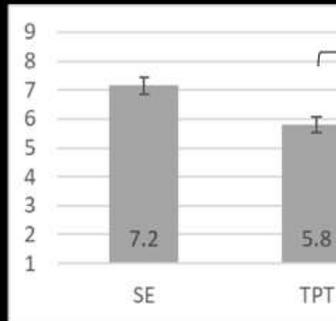
- If we use a positive thick term, we can cancel the evaluation without contradicting ourselves
- If we use a negative thick term, we cannot cancel the evaluation without contradicting ourselves
- Prediction:  
*People give higher „you should not“ ratings when negative thick concepts are used non-evaluatively, compared to positive thick concepts*





## Predictions

## Results



## Discussion

Why do Positive and Negative Thick Terms behave differently?

1. Differences in moral evaluation
2. Availability of counterexamples
3. Differences in the social norms guiding evaluative language

## Discussion

Why do Positive and Negative Thick Terms behave differently?

1. Differences in moral evaluation
2. Availability of counterexamples
3. *Differences in the social norms guiding evaluative language*



*Let's do more of this!*

---

- This results we got are surprising, given that philosophers discuss TC as a unique class of concepts
- We should determine whether research on TC needs to treat positive and negative TC separately
- We should determine what features can make the evaluative component change to neutral





## Get in touch!



[PascaleWillemsen.com](http://PascaleWillemsen.com)



[Pascale.Willemsen@philos.uzh.ch](mailto:Pascale.Willemsen@philos.uzh.ch)

## Special thanks to:

Bianca Cepollaro

Matti Eklund

Markus Kneer

Iza Skoczen

Alex Wiegmann

Pekka Värynen



FONDS NATIONAL SUISSE  
SCHWEIZERISCHER NATIONALFONDS  
FONDO NAZIONALE SVIZZERO  
SWISS NATIONAL SCIENCE FOUNDATION