How to take a hint: A case study in linking it all together

http://compcogscisydney.org/psyc2071/

Danielle Navarro

A partial list of important topics

attention

sensation

similarity

perception

reasoning

learning

social cognition

decision making

sensation

perception

attention

similarity

reasoning

learning

social cognition

decision making

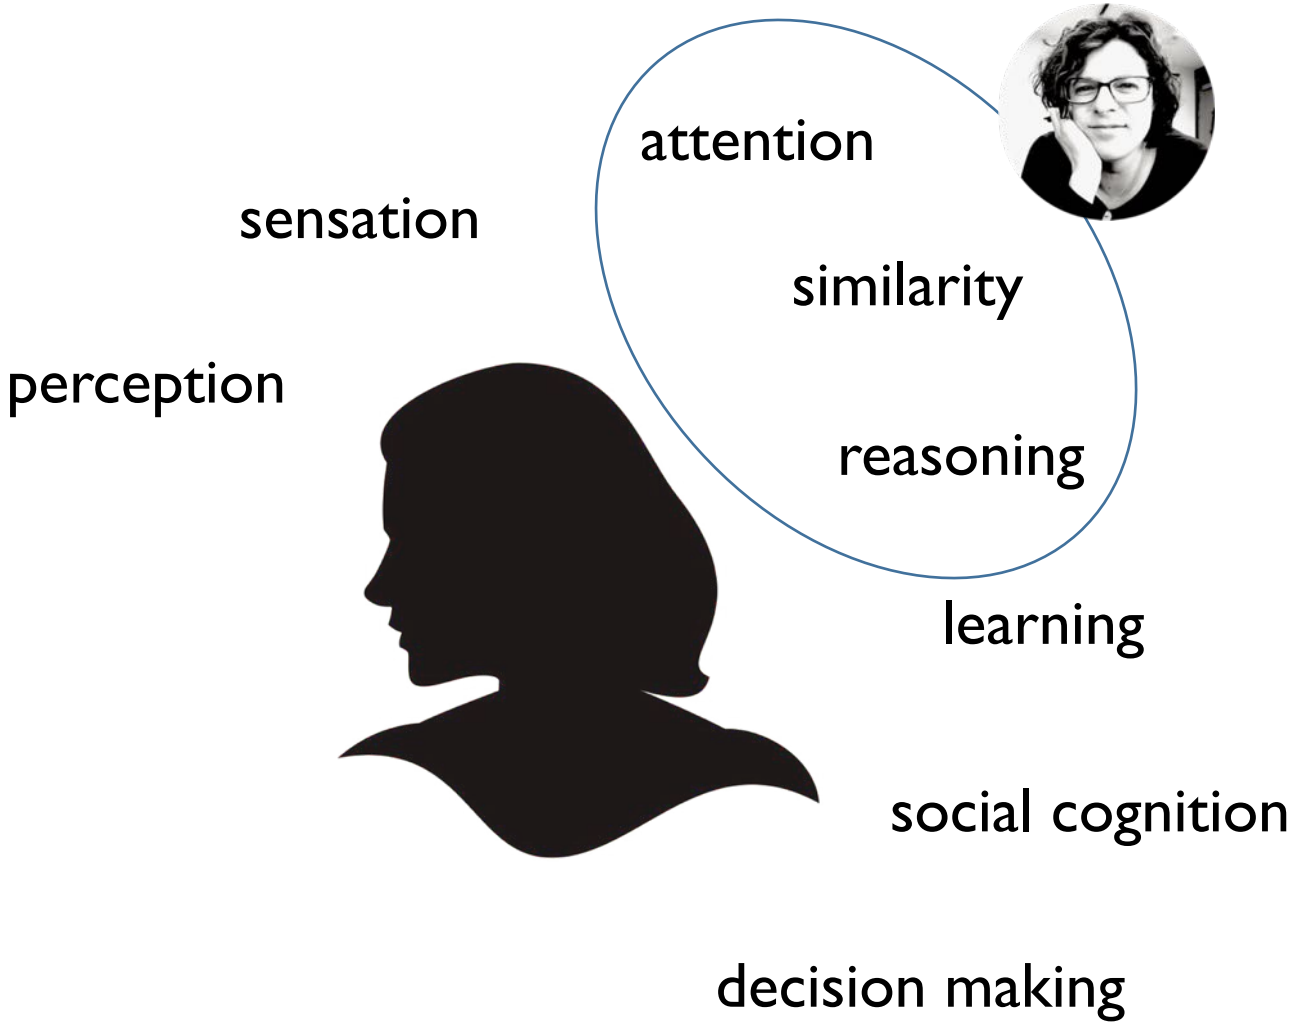
attention
sensation
similarity
perception
reasoning
learning
social cognition
decision making

From topics to tasks…

attention

sensation

similarity

perception

reasoning

- Wason selection task
- Property induction
- Argument evaluation
- Etc…

learning

social cognition

decision making

Today we'll reverse this

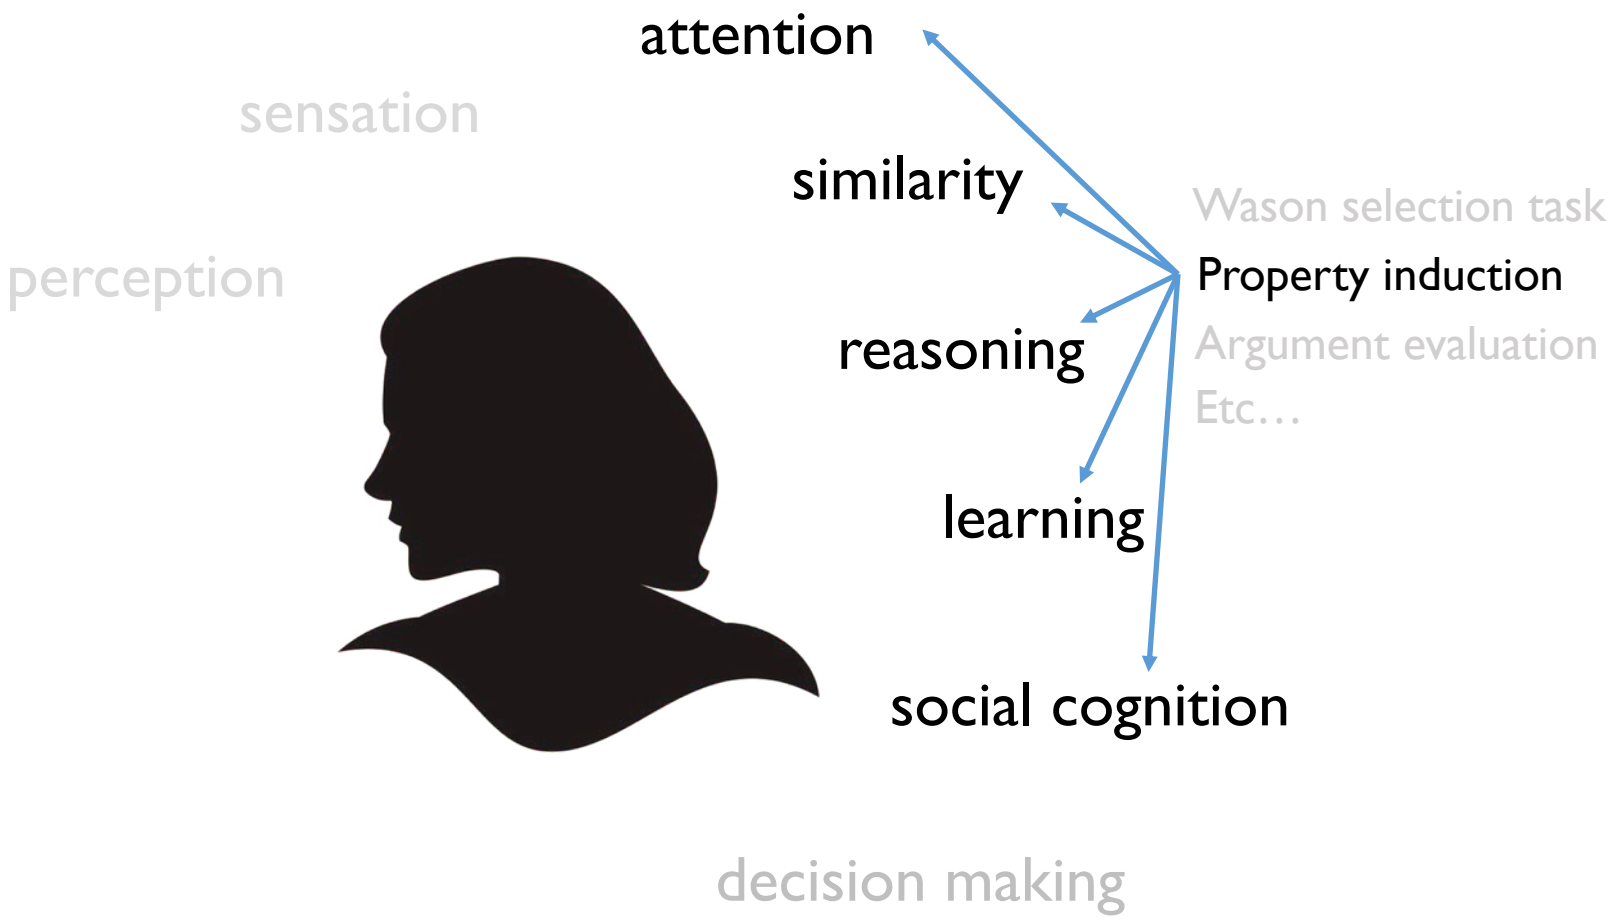

The case study…

Cognitive Science Explore this journal >

Brief Report

Leaping to Conclusions: Why Premise Relevance Affects Argument Strength

Keith J. Ransom, Amy Perfors, Daniel J. Navarro

First published: 16 October 2015 Full publication history

DOI: 10.1111/cogs.12308 View/save citation

Cited by: 0 articles Citation tools

Funding Information

Early View

Browse Early View Articles
Online Version of Record published before inclusion in an issue

Abstract

Everyday reasoning requires more evidence than raw data alone can provide. We explore the idea that people can go beyond this data by reasoning about how the data was sampled. This idea is investigated through an examination of *premise non-monotonicity*, in which adding premises to a category-based argument weakens rather than strengthens it. Relevance theories explain this phenomenon in terms of people's sensitivity to the relationships among premise items. We show that a Bayesian model of category-based induction taking premise sampling assumptions and category similarity into account complements such theories and yields two important predictions: First, that sensitivity to premise relationships can be violated by inducing a weak sampling assumption; and second, that premise monotonicity should be restored as a result. We test these predictions with an experiment that manipulates people's assumptions in this regard, showing that people draw qualitatively different conclusions in each case.

The most exciting phrase to hear in science, the one that heralds new discoveries, is not '*Eureka*!' but '*That's funny...*'

- Isaac Asimov

Premise monotonicity

Dolphin cells contain TH4 hormone
Therefore cow cells contain TH4 hormone?


Dolphin cells contain TH4 hormone
Mouse cells contain TH4 hormone
Bat cells contain TH4 hormone
Therefore cow cells contain TH4 hormone?




Adding evidence usually *strengthens* an inductive argument

Premise *non*-monotonicity

Dolphin cells contain TH4 hormone
Therefore cow cells contain TH4 hormone?

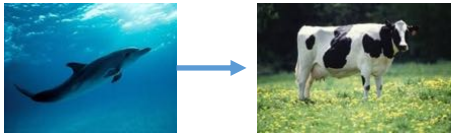

Dolphin cells contain TH4 hormone
Whale cells contain TH4 hormone
Seal cells contain TH4 hormone
Therefore cow cells contain TH4 hormone?

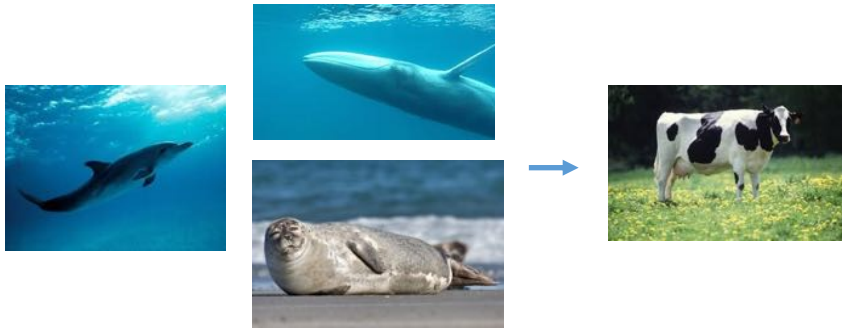

Sometimes adding evidence *weakens* an argument?

"that's funny?"

But *whyyyyyy*?

A tale of similarity, attention and social cognition

Observation #1:
Similarity shapes reasoning

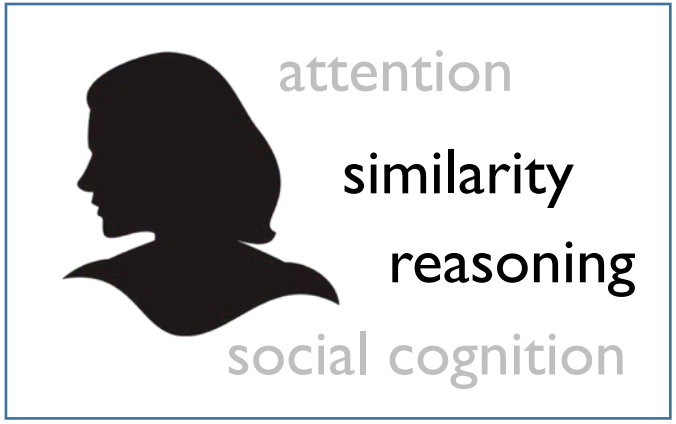

Similarity is relevant to reasoning

Dolphins and cows are dissimilar.

So this feels unreasonable

Similarity is relevant to reasoning




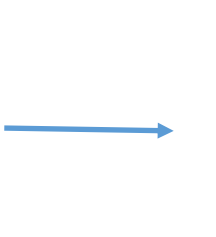


Bats and mice are very to dissimilar cows too, but they're also dissimilar to dolphins.

Suggests the TH4 hormone is common?
… so the argument gets stronger

Similarity is relevant to reasoning

Seals and whales are very dissimilar to cows too, but they are very similar to dolphins and to each other.

Suggests that TH4 is possessed by a narrow range of animals that does not include cows.
... so the argument gets weaker

Similarity is relevant to reasoning

If the conclusion item (dugong) is "sufficiently similar" to the premise items then monotonicity is restored

… this is also a strong argument

Observation #2:
Similarity directs attention* to a particular category

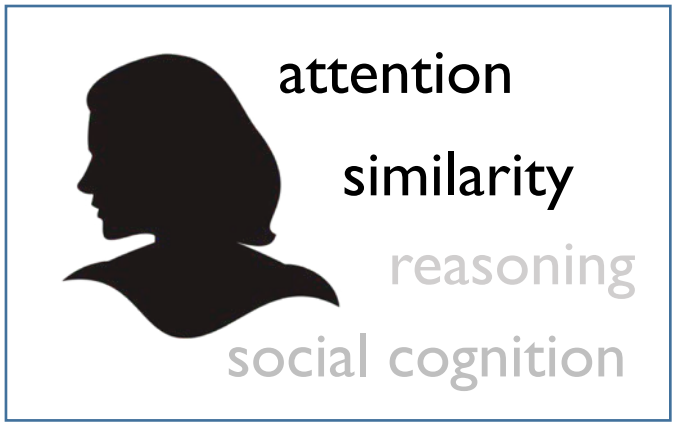

**internal* attention!

Dolphins are …

… marine mammals?

... intelligent animals?

... mammals?

... cute?

Many possible categories that could indicate which animals have TH4 and which don't

… marine mammals?

... intelligent animals?

... mammals?

... cute?

Adding bats and mice calls attention to **mammal**

… marine mammals?

... intelligent animals?

... mammals?

... cute?

Adding whales and seals calls attention to **marine mammal**

… marine mammals?

... intelligent animals?

... mammals?

... cute?

Notice: seals and whales are also intelligent cute mammals. These possibilities aren't ruled out, we just **ignore** them

A scientific question…

Why does this similarity-driven attention influence our reasoning?

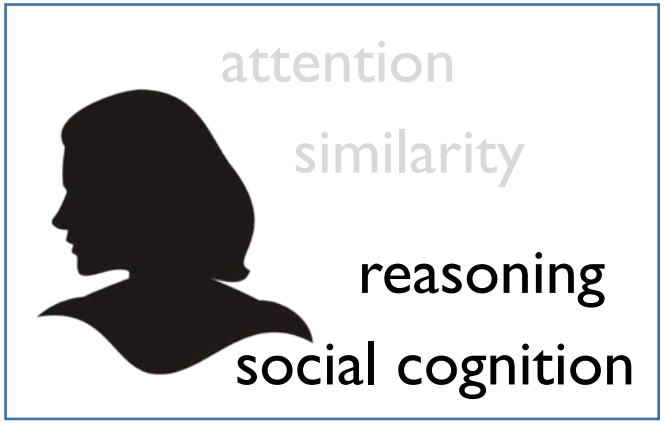

On the origins of data

Information from the world

Information from people

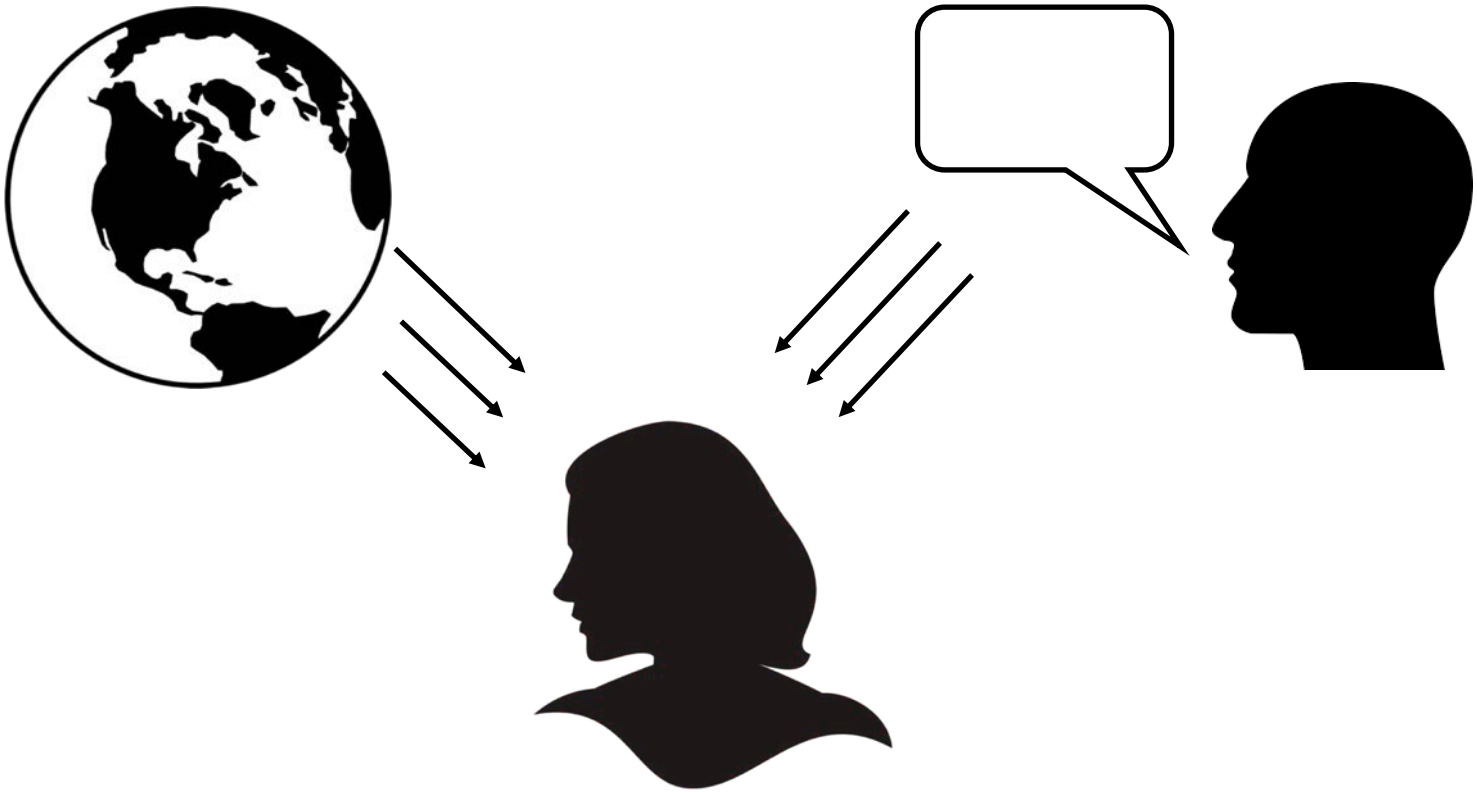

Humans are intelligent agents with complex goals and a rich language. We "transmit" information to each other via a complicated mechanism...

persuasion

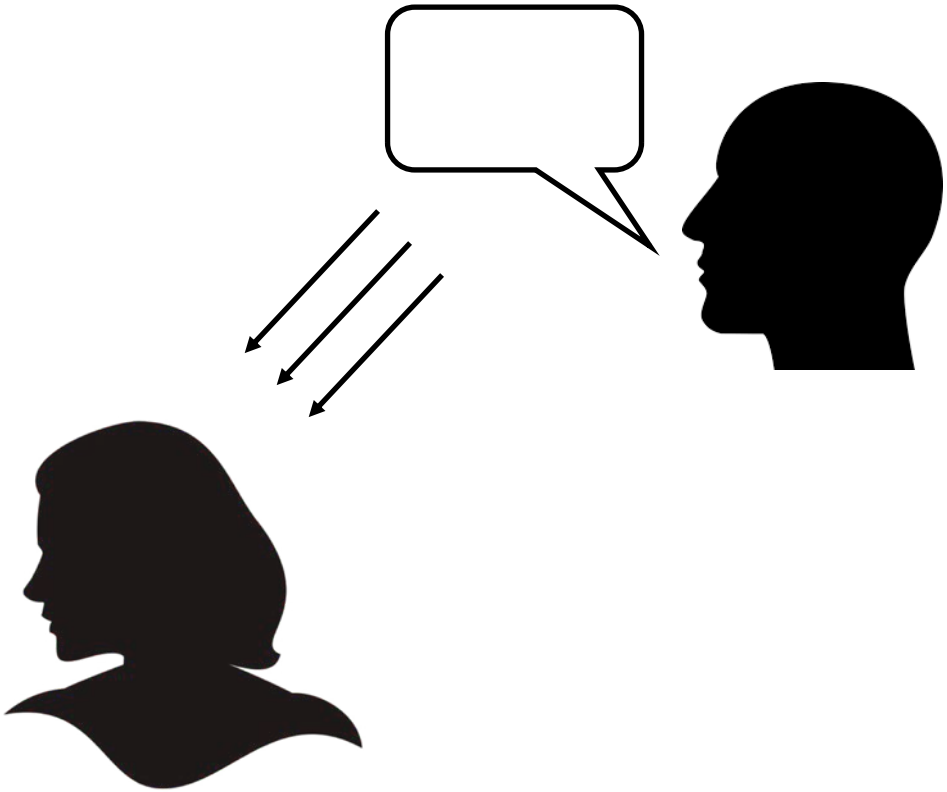

How does this communication work?

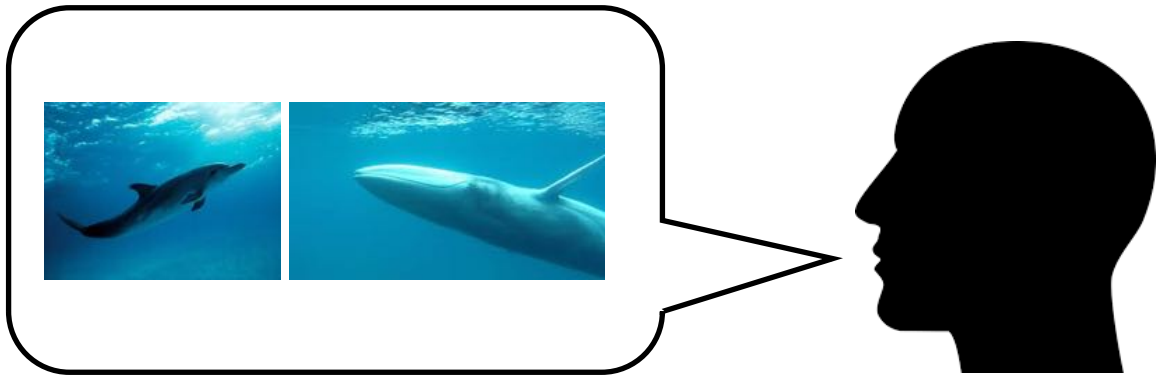

If I *choose* these similar animals…

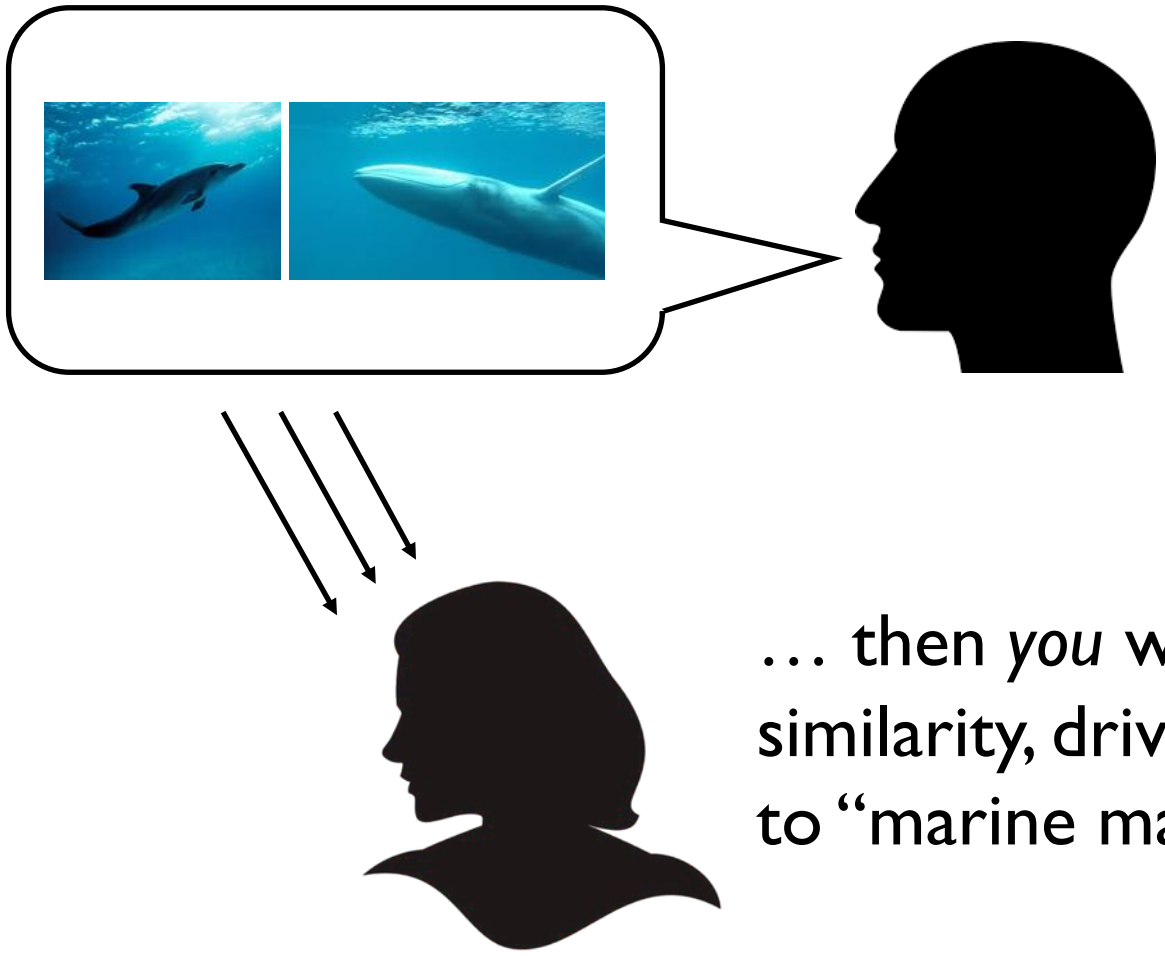

… then *you* will notice the similarity, driving your ***attention*** to "marine mammals"

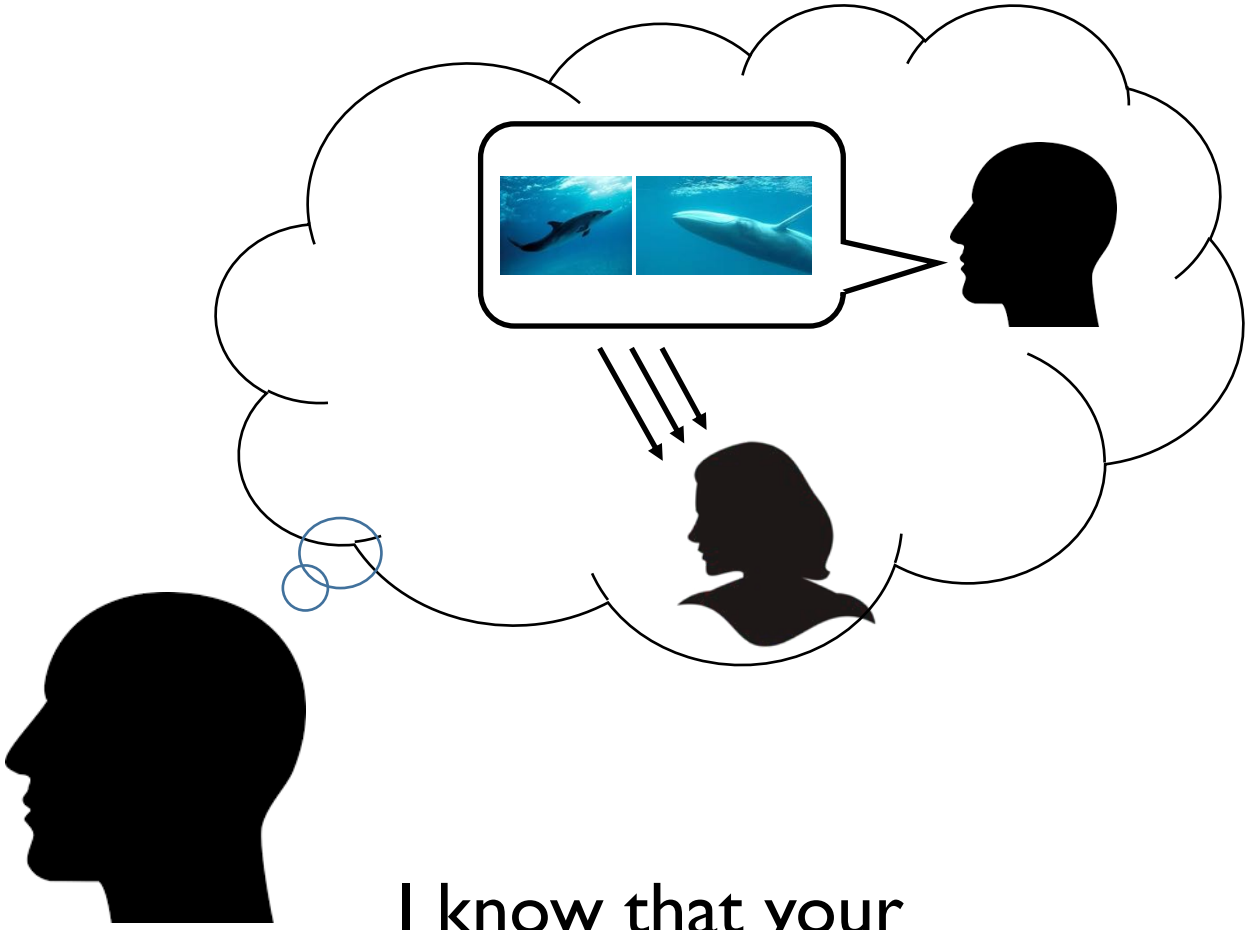

I know that your
mind works like this

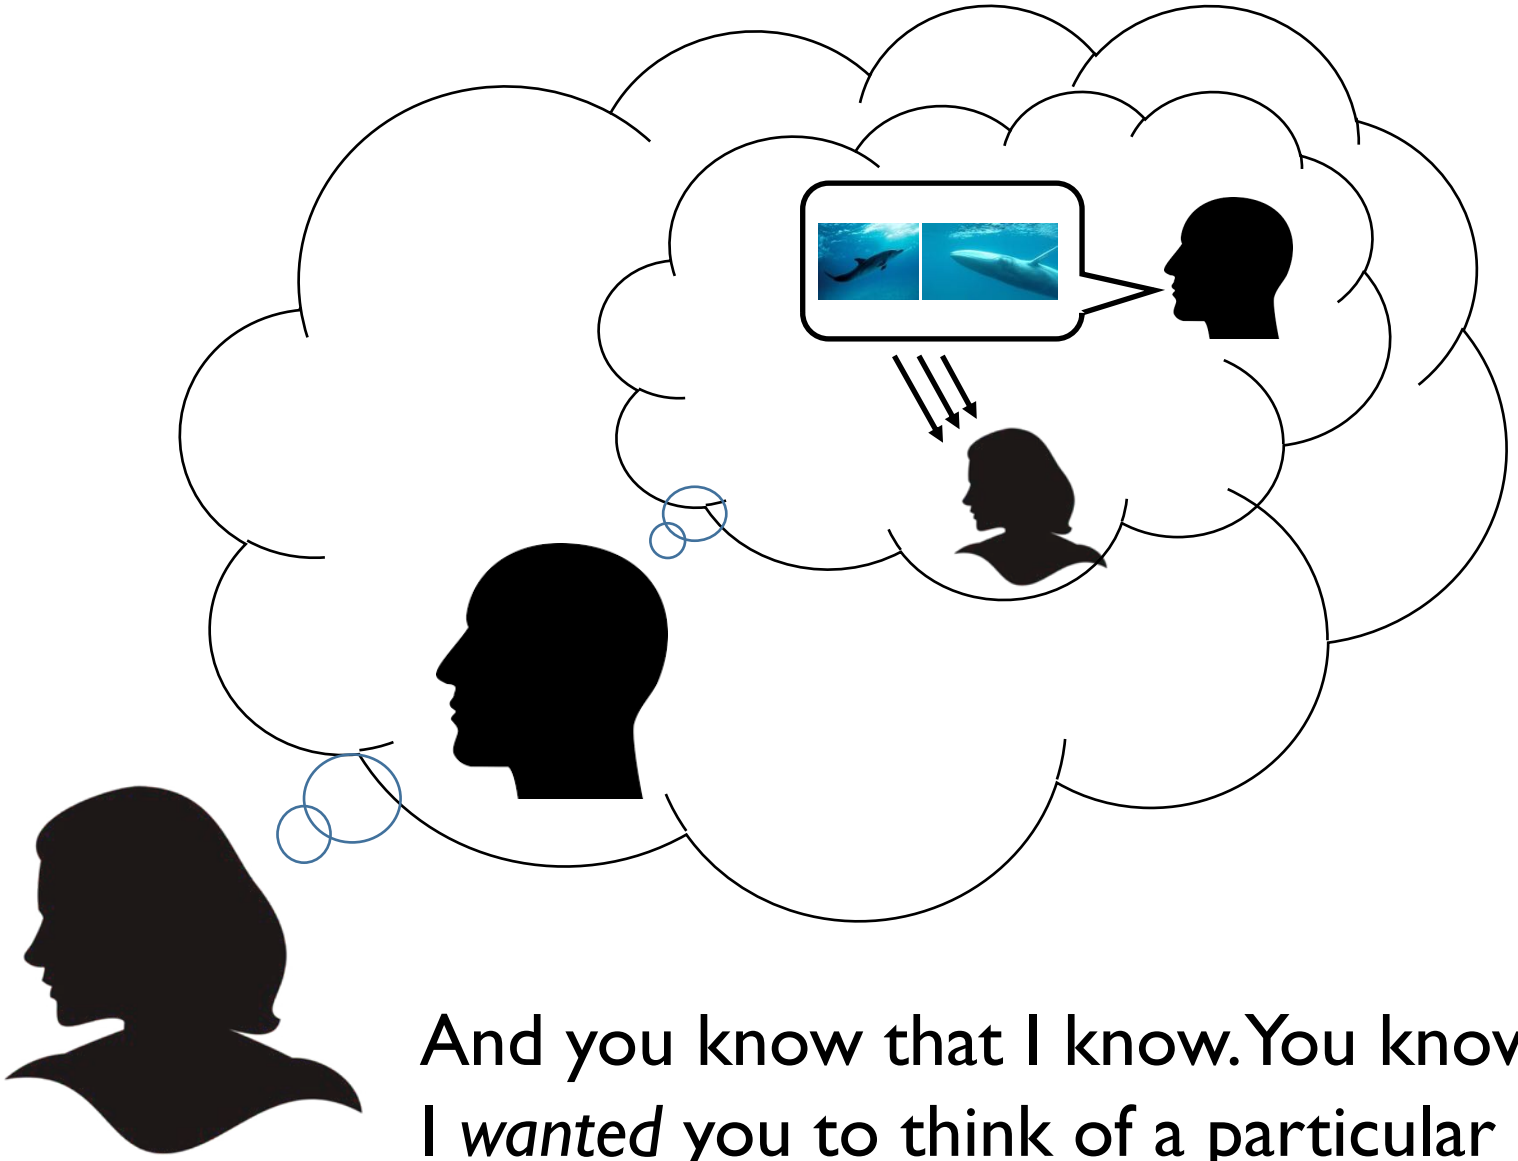

And you know that I know. You know I *wanted* you to think of a particular category … so you can "take a hint"

Theory of mind!

We have intuitive theories about the workings of each other's minds, so we can select **relevant** information that *drives* attention to the right answer

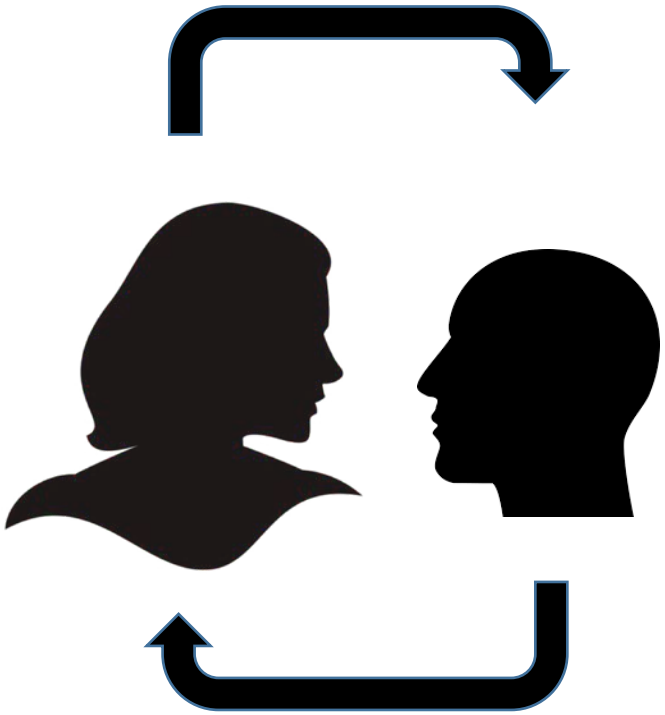

So humans do this…

"I've studied TH4 hormone for many years… and I have discovered it in the cells of whales, seals and dolphins.

I want you to believe that dugongs will produce TH4 hormone"

And we do this…

"I've studied TH4 hormone for many years… and I have discovered it in the cells of mice, bats and dolphins.

I want you to believe that cows will produce TH4 hormone"

We don't do *this:*

"I've studied TH4 hormone for many years… and I have discovered it in the cells of whales, seals and dolphins.

I want you to believe that kittens will produce TH4 hormone"

Okay, so how does this other mechanism work?

The world is dumb. It does not care what you believe. It does not give "hints"

You know that it does not care what you believe. You do not expect the world to behave like an intelligent or helpful agent.

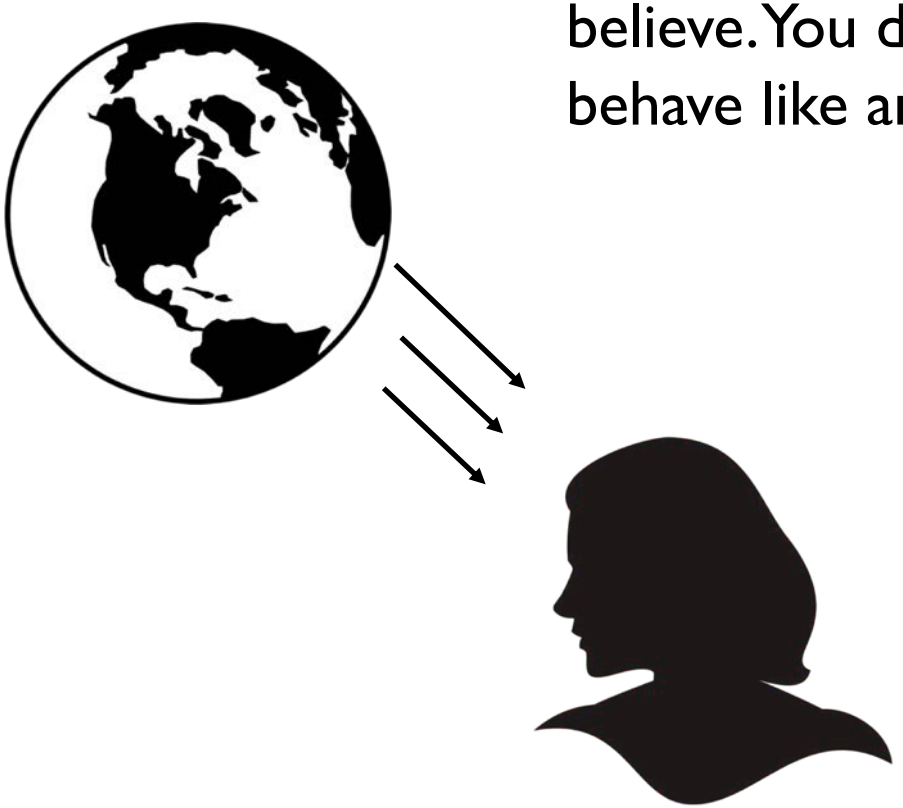

The world generates data randomly

My cat's breath smells like cat food

Unicorns are awesome

Dolphins are cute

The world generates data randomly

My cat's breath smells like cat food

Unicorns are awesome

Dolphins are cute

??????

The world doesn't care what you believe, and it doesn't try to influence you: it's just a big dumb rock

Other humans *do* care what you believe, and they do try to shape your beliefs by choosing the right words

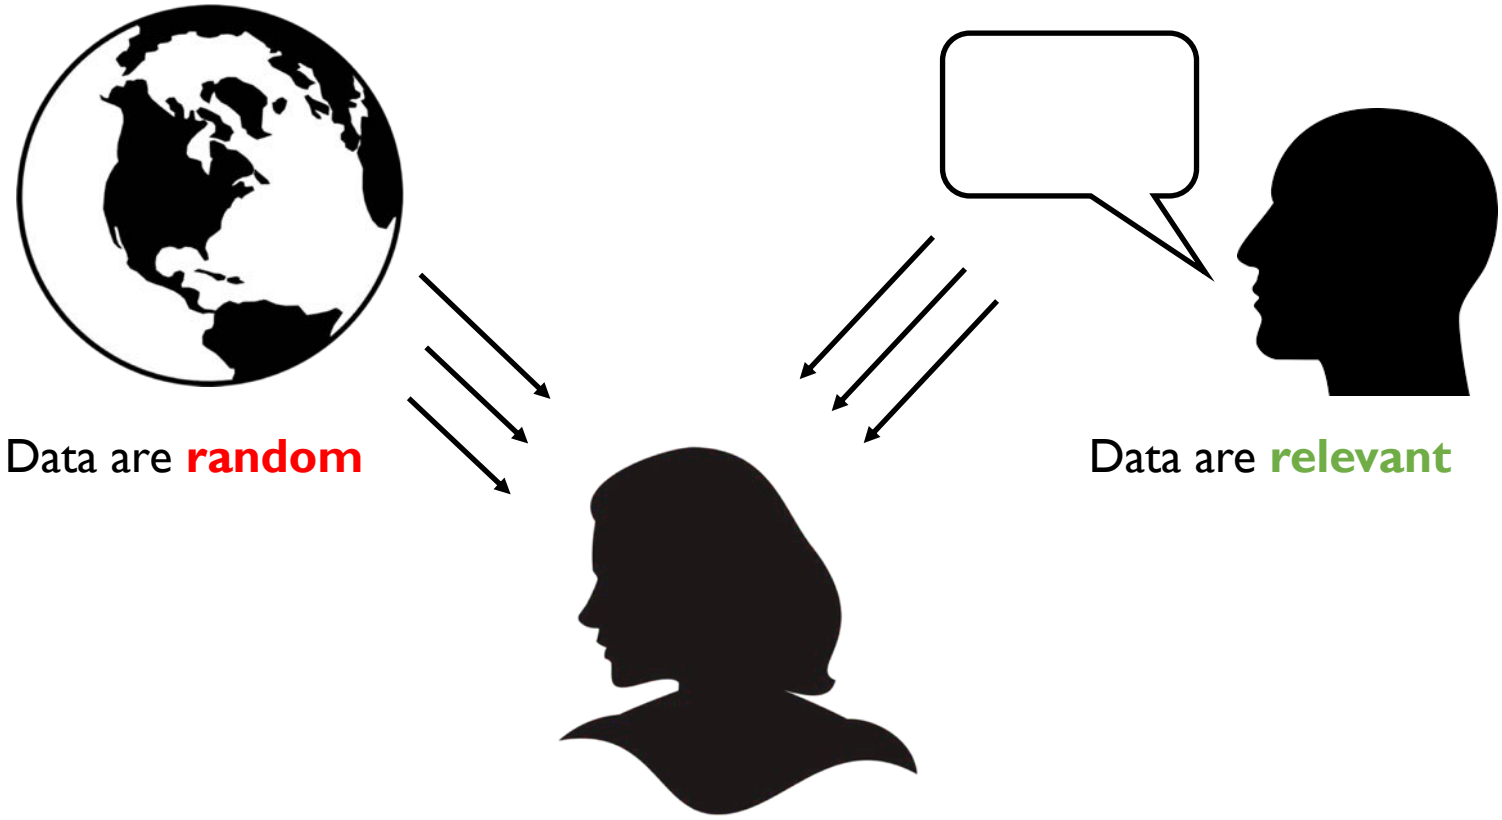

Data are **random**

Data are **relevant**

Do I reason differently in these two situations?

(& can this explain non-monotonic reasoning?)
(& if so, is that the right thing to do?)

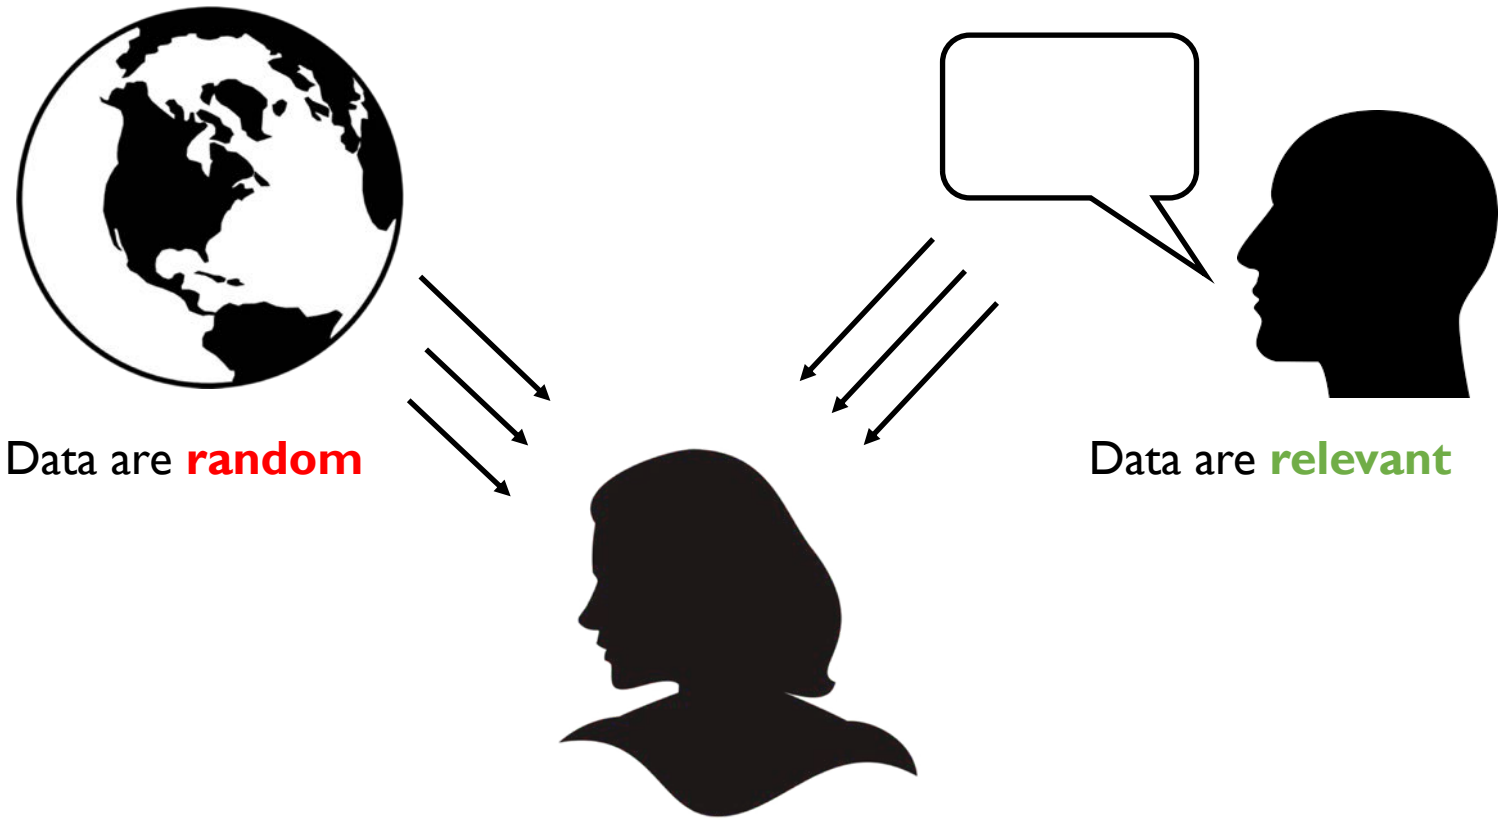

Data are **random**

Data are **relevant**

To the laboratory!

The reasoning task

Item A has property P

Item X has property P

How likely is it that item X has property P?

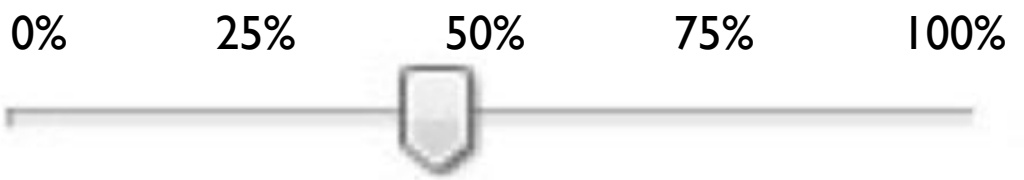

Participants are first asked to rate an inductive argument with a single premise

The reasoning task

Item A has property P
Item B has property P

Item X has property P

How likely is it that item X has property P?

0% 25% 50% 75% 100%

A second item is added and they are asked to revise their estimate

Argument 1b



Item A has property P
Item B has property P

Item X has property P

Argument 1a


Item A has property P

Item X has property P

DV is the *difference* score

Argument 1b



Item A has property P
Item B has property P

Item X has property P

Argument 1a


Item A has property P

Item X has property P

Response b minus response a

>0 : premise monotonicity

0 : no change

<0 : premise non-monotonicity

Our hypothesis

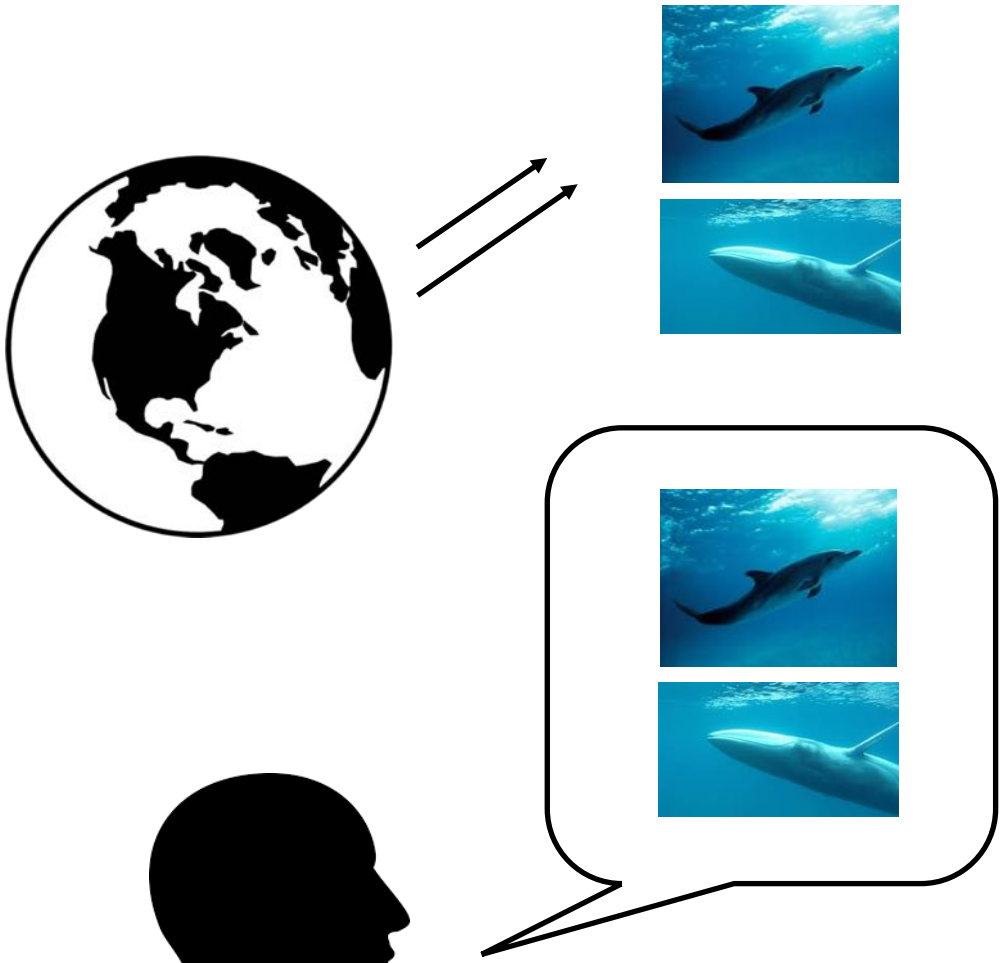

When an indifferent world generates random data, the *similarity* between premise items will be deemed *irrelevant*, and people will revert to premise monotonicity

When a helpful human makes an argument, the *similarity* between premise items will be deemed relevant, and the premise non-monotonicity effect will appear

Manipulated variable

Participant believes data are random

Participant believes data are relevant

DV

>0 : premise monotonicity

0 : no change

<0 : premise non-monotonicity

Okay… but how do we manipulate people's beliefs???

Problem: an *experiment* is designed by an experimenter....

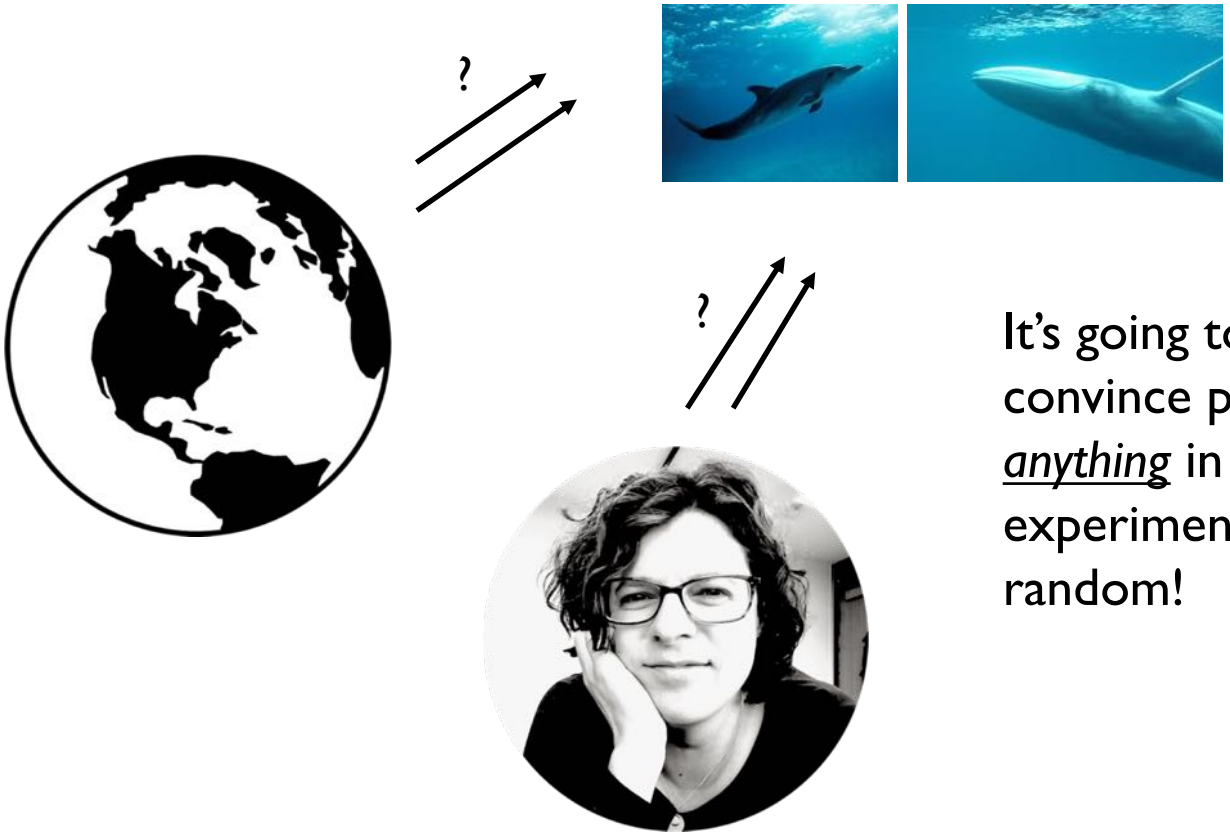

It's going to be hard to convince people that *anything* in a psychology experiment is truly random!

"Cover story" manipulation

(independent variable #1)

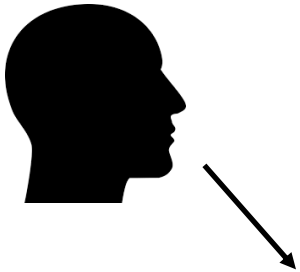

What do we ***tell*** people about the origin of the second premise?

- **Relevant** story: It is a *hint* from a previous participant
- **Neutral** story: Don't tell them where it comes from
- **Random** story: It is chosen at random from a database

"Experience" manipulation

(independent variable #2)

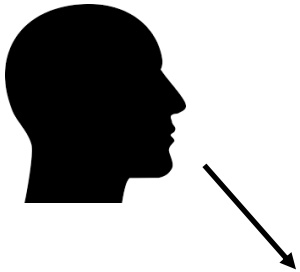

What do we ***show*** people about the nature of the second premise?

- **Relevant** data: Previous examples have been helpful
- **Random** data: Previous examples have been stupid

“Flow” of the experiment

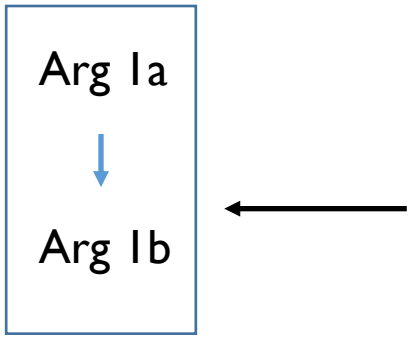

Cover story manipulation appears here:

The “a” arguments have one premise; when the second premise is added to create the “b” argument, we “remind” the participants that the data [is a hint / ??? / is random]

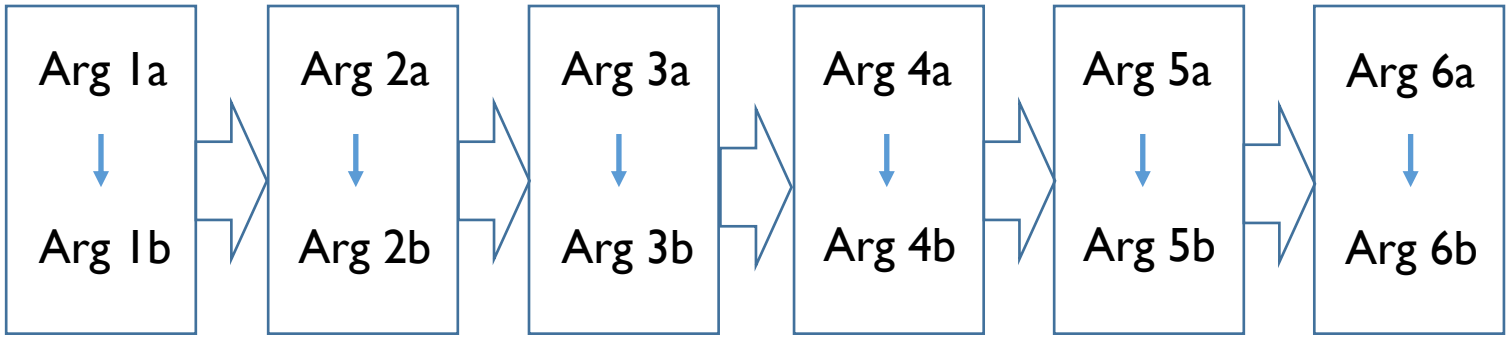

Each participant is shown 6
of these argument pairs

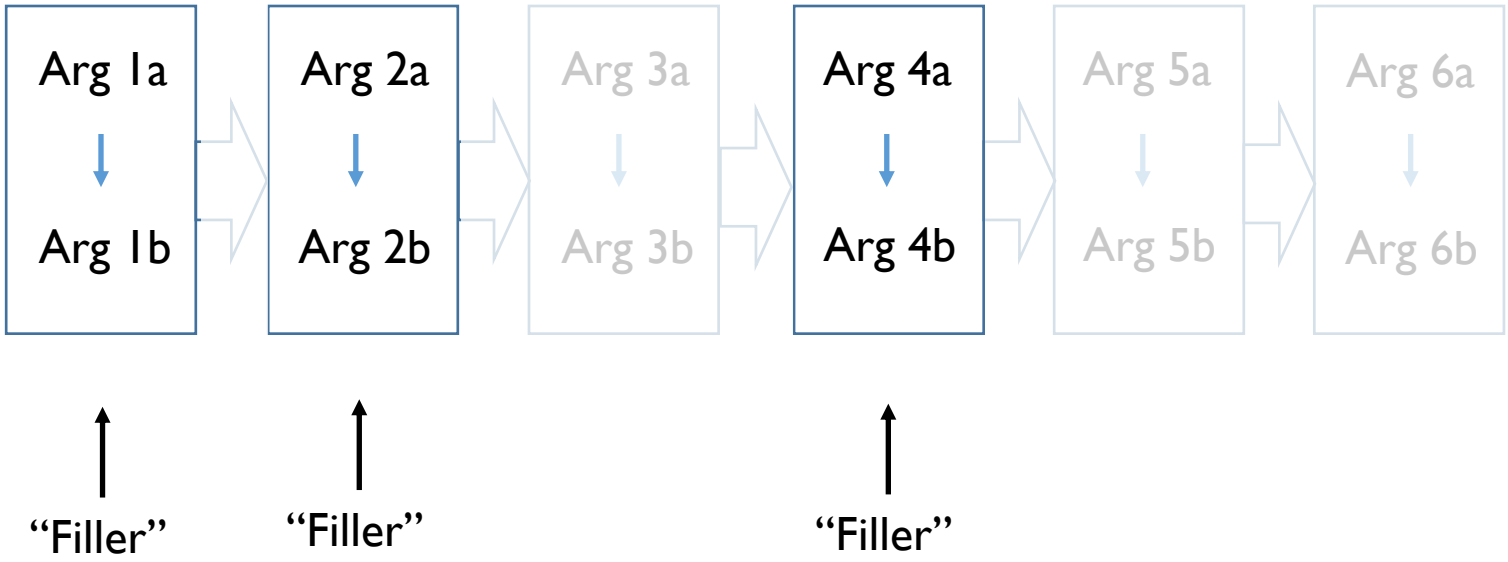

The experience manipulation appears here: the 1st, 2nd and 4th arguments were "filler" items designed to highlight the [relevance / randomness] of the second premise

The "experience" manipulation

A "filler" trial with a relevant second premise:

A "filler" trial with a random second premise:

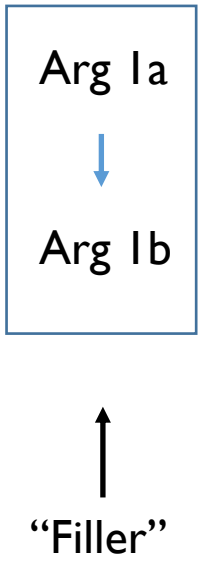

The "experience" manipulation

A "filler" trial with a relevant second premise:

A "filler" trial with a random second premise:

Arg 1a

Arg 1b

"Filler"

Eagles have multiple foveas

Doves have multiple foveas

Eagles have multiple foveas
Hawks have multiple foveas

Doves have multiple foveas

The "experience" manipulation

A "filler" trial with a relevant second premise:

A "filler" trial with a random second premise:

Arg 1a

Arg 1b

"Filler"

Eagles have multiple foveas

Doves have multiple foveas

Eagles have multiple foveas
Hawks have multiple foveas

Doves have multiple foveas

This suggests the involvement of a helpful human because *hawks* seem relevant to the context

The "experience" manipulation

A "filler" trial with a relevant second premise:

A "filler" trial with a random second premise:

Eagles have multiple foveas

Doves have multiple foveas

↓

Eagles have multiple foveas
Tortoises do *not* have multiple foveas

Doves have multiple foveas

Arg 1a

↓

Arg 1b

↑

"Filler"

The "experience" manipulation

A "filler" trial with a relevant second premise:

A "filler" trial with a random second premise:

Arg 1a

↓

Arg 1b

↑

"Filler"

Eagles have multiple foveas

Doves have multiple foveas

↓

Eagles have multiple foveas
Tortoises do *not* have multiple foveas

Doves have multiple foveas

This is the <u>*worst*</u> hint ever. Why would anyone think <u>*tortoises*</u> are relevant? Clearly, the second premise is randomly chosen

Incomplete 2x3 design

Story \ Experience		
	Relevant story, Relevant items	
	Neutral story, Relevant items	Neutral story, Random items
		Random story, Random items

Incomplete 2x3 design

Story \ Experience		
	Relevant story, Relevant items	
	Neutral story, Relevant items	Neutral story, Random items
		Random story, Random items

Q1: Why not 3x3? Where are the "neutral items" conditions?
Q2: Why is it incomplete? Why did we leave two empty cells here?

Now, how should we measure the effect?

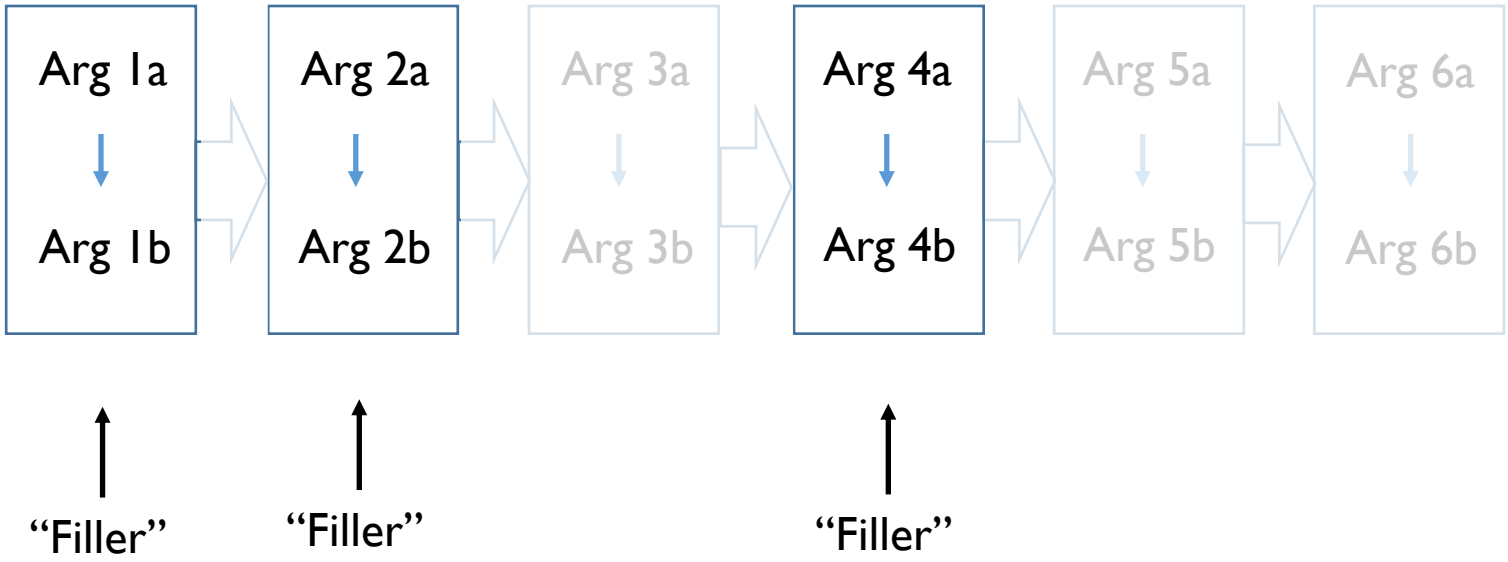

We used these three arguments for the IV

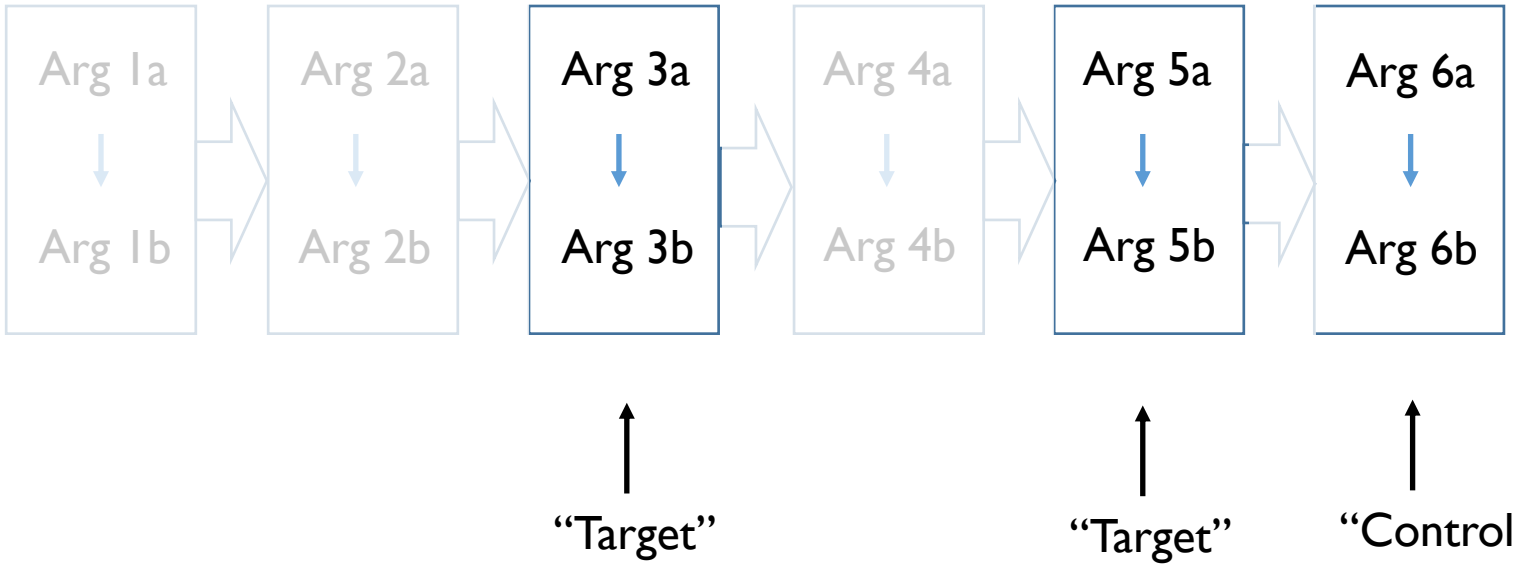

We used these three argument for the DV

For **target** items the second premise was similar to the first premise item but *not* the conclusion

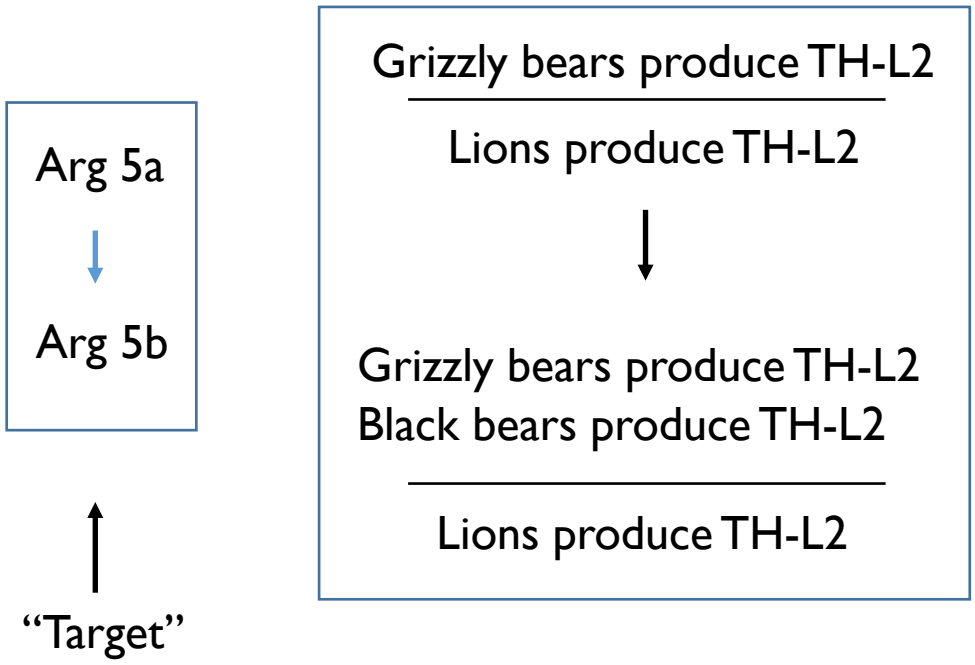

For **target** items the second premise was similar to the first premise item but *not* the conclusion

Arg 5a

↓

Arg 5b

↑

"Target"

Grizzly bears produce TH-L2

Lions produce TH-L2

↓

Grizzly bears produce TH-L2
Black bears produce TH-L2

Lions produce TH-L2

For **target** items the second premise was similar to the first premise item but *not* the conclusion

If this similarity is deemed relevant… it strongly suggests TH-L2 is a property of bears so the extra evidence weakens the conclusion... non monotonicity

For **target** items the second premise was similar to the first premise item but *not* the conclusion

If this similarity is deemed irrelevant because the data are random … the extra "evidence" is almost useless

So we expect no difference or a weak monotonicity effect because at least there's some extra evidence that TH-L2 is not rare

For the **control** item the second premise was similar to the first premise *and* the conclusion

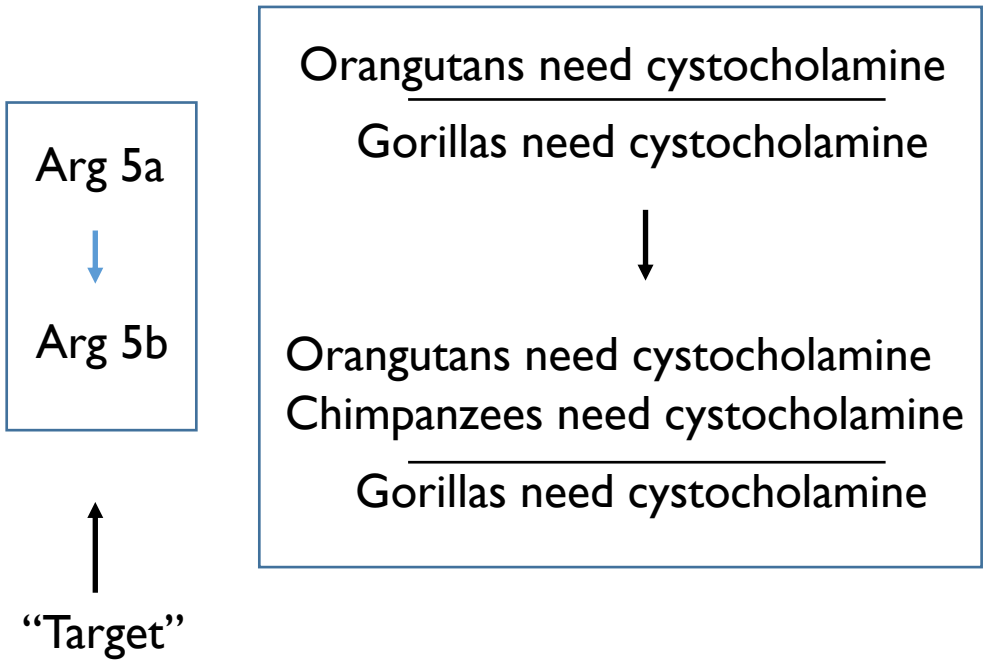

For the **control** item the second premise was similar to the first premise *and* the conclusion

Arg 5a

↓

Arg 5b

↑

"Target"

Orangutans need cystocholamine

Gorillas need cystocholamine

↓

Orangutans need cystocholamine
Chimpanzees need cystocholamine

Gorillas need cystocholamine

For the **control** item the second premise was similar to the first premise *and* the conclusion

If similarity is considered relevant, it calls attention to "primates", but this is the intuitively obvious answer anyway. So it really doesn't make a difference:

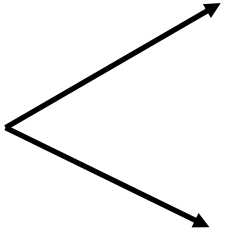

... So we expect premise monotonicity in **both** conditions because either way it reinforces the thing you already believed!

(* This is analogous to the dugongs example earlier)

What we expect to see if our theory is *right*

Dependent variable is the difference score…

>0 : premise monotonicity

0 : no change

<0 : premise non-monotonicity

Target arguments should show a systematic change when we manipulate the perceived origin of the data…

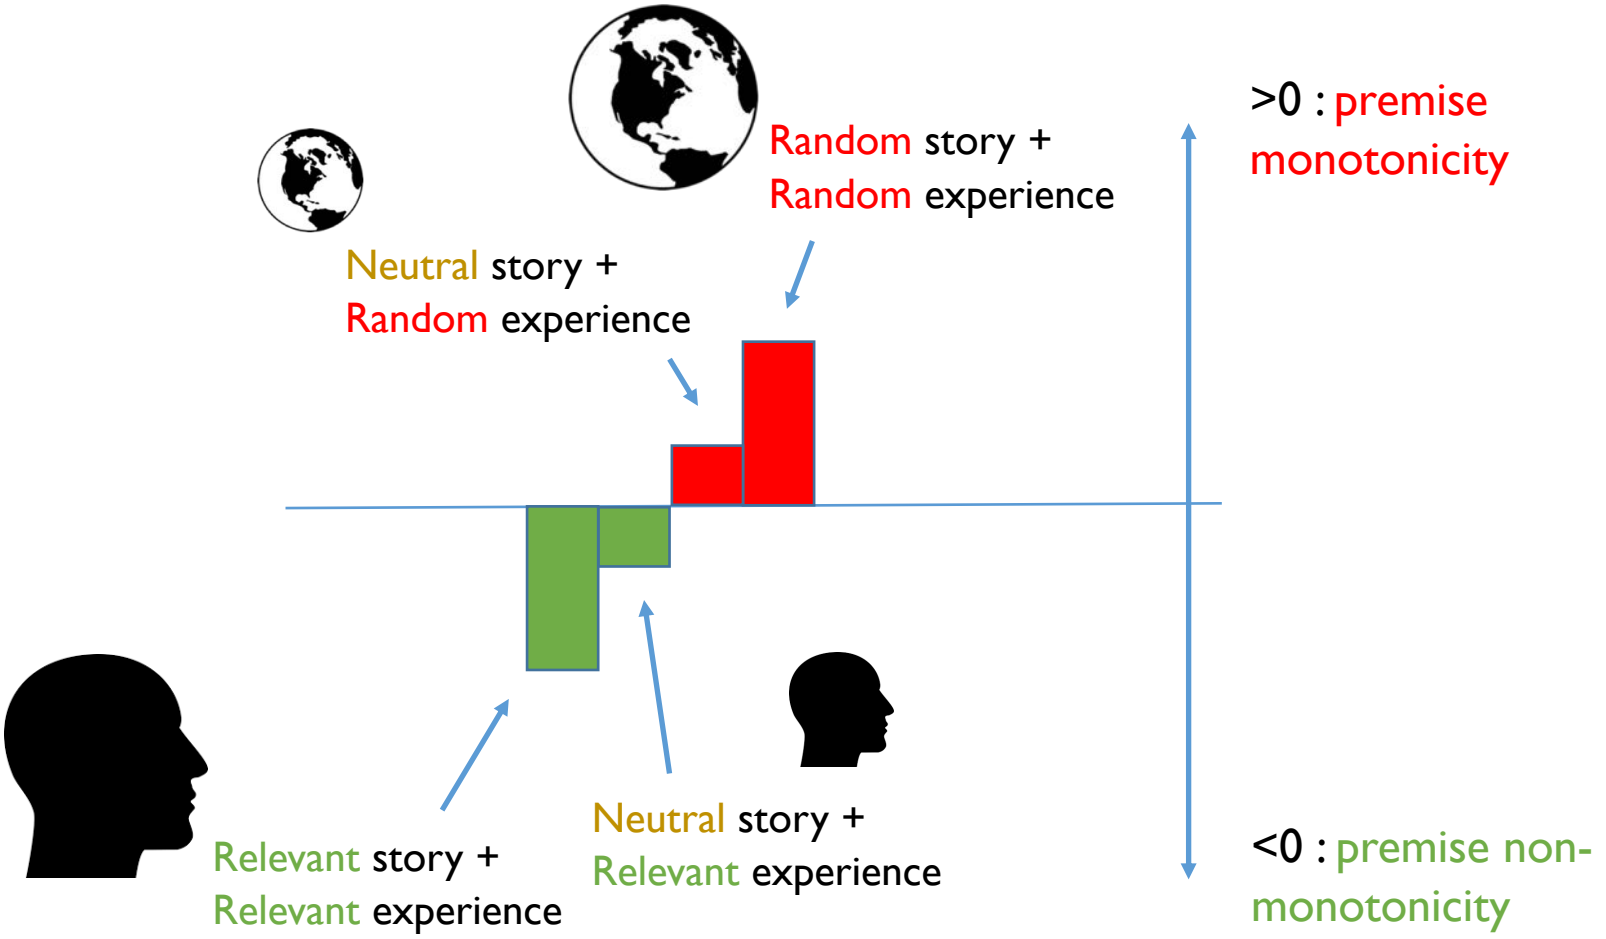

Control argument should produce premise monotonicity under all conditions

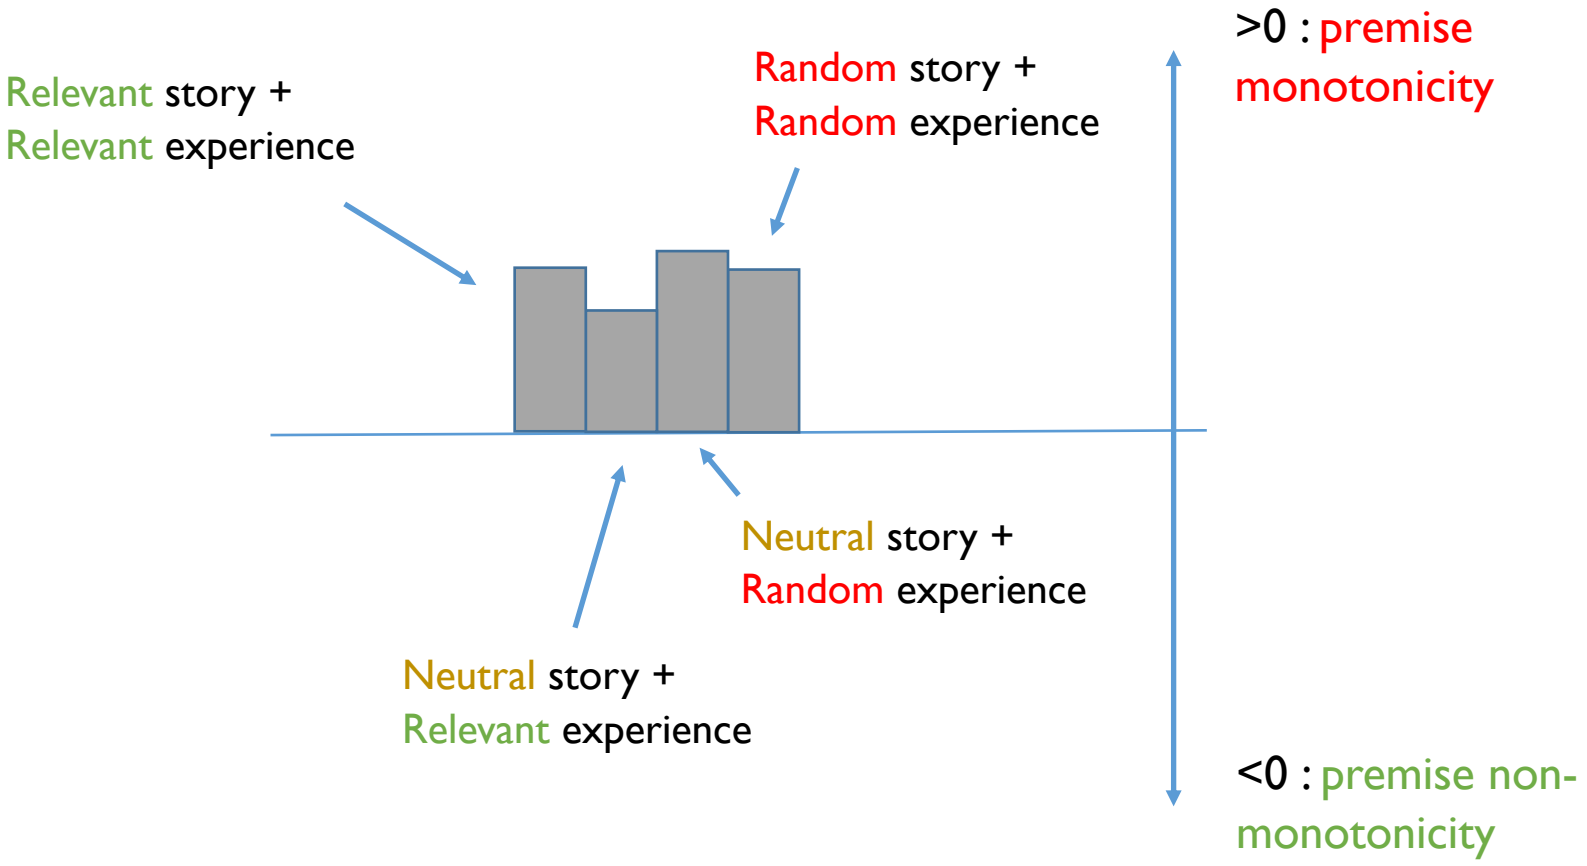

What we expect to see if our theory is *wrong*

If similarity always drives attention and reasoning in the same way, the IVs should make no difference…

>0 : premise monotonicity

<0 : premise non-monotonicity

If similarity always drives attention and reasoning in the same way, the IVs should make no difference…

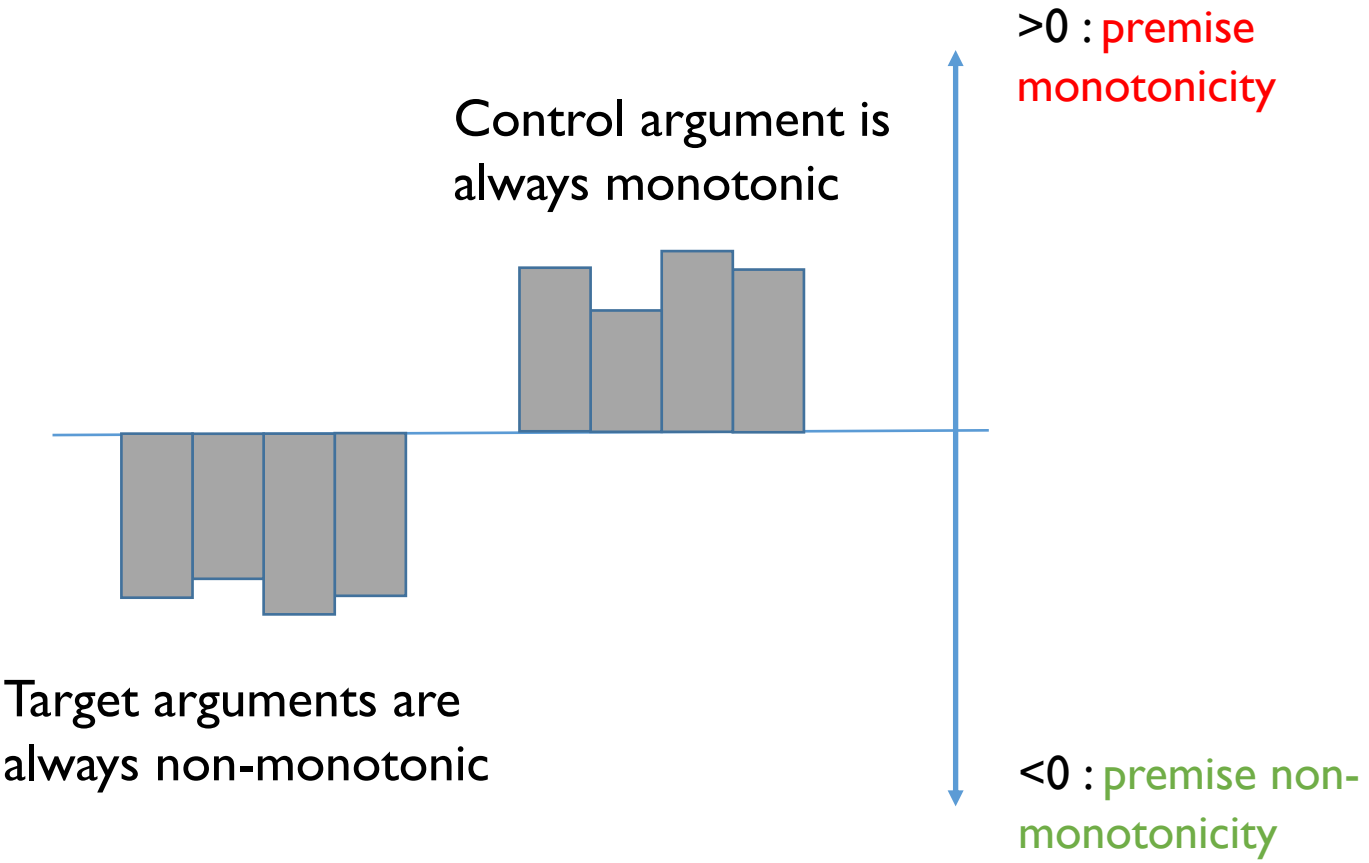

What we expect to see if our experiment is total garbage

If we made some weird mistake, the data should show no meaningful patterns

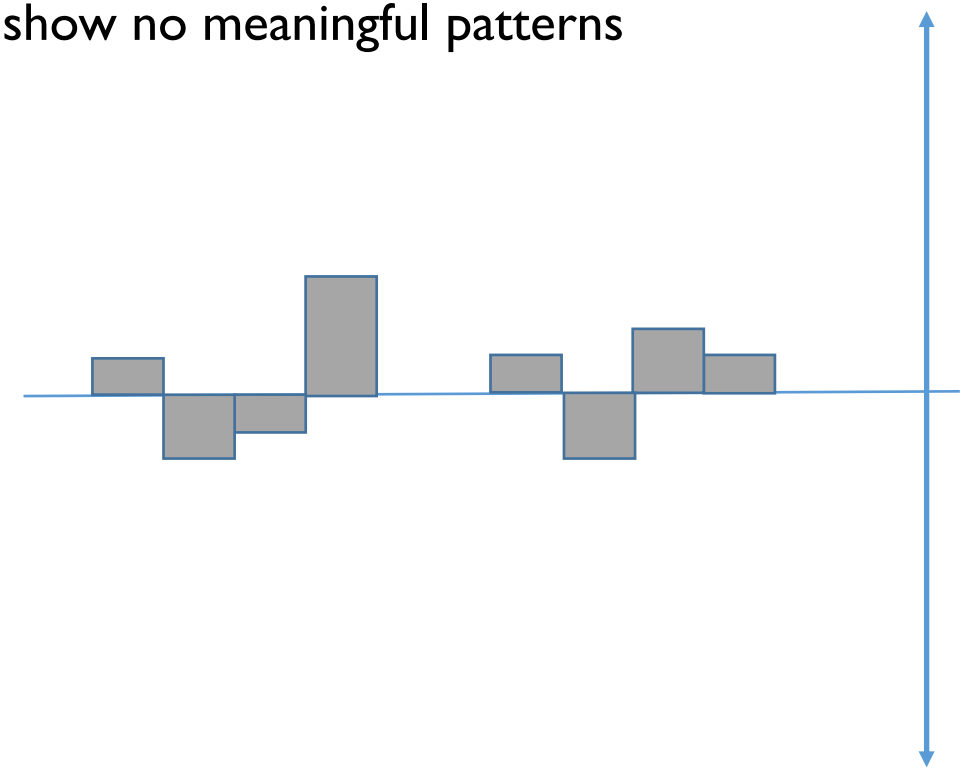

>0 : premise monotonicity

<0 : premise non-monotonicity

What were our results?

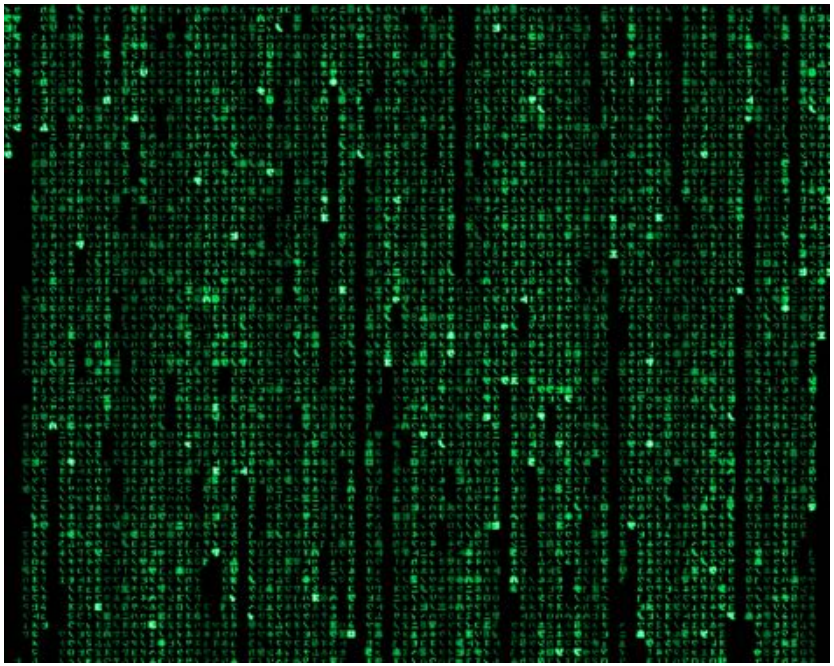

Premise monotonicity

Premise non-monotonicity

Premise monotonicity

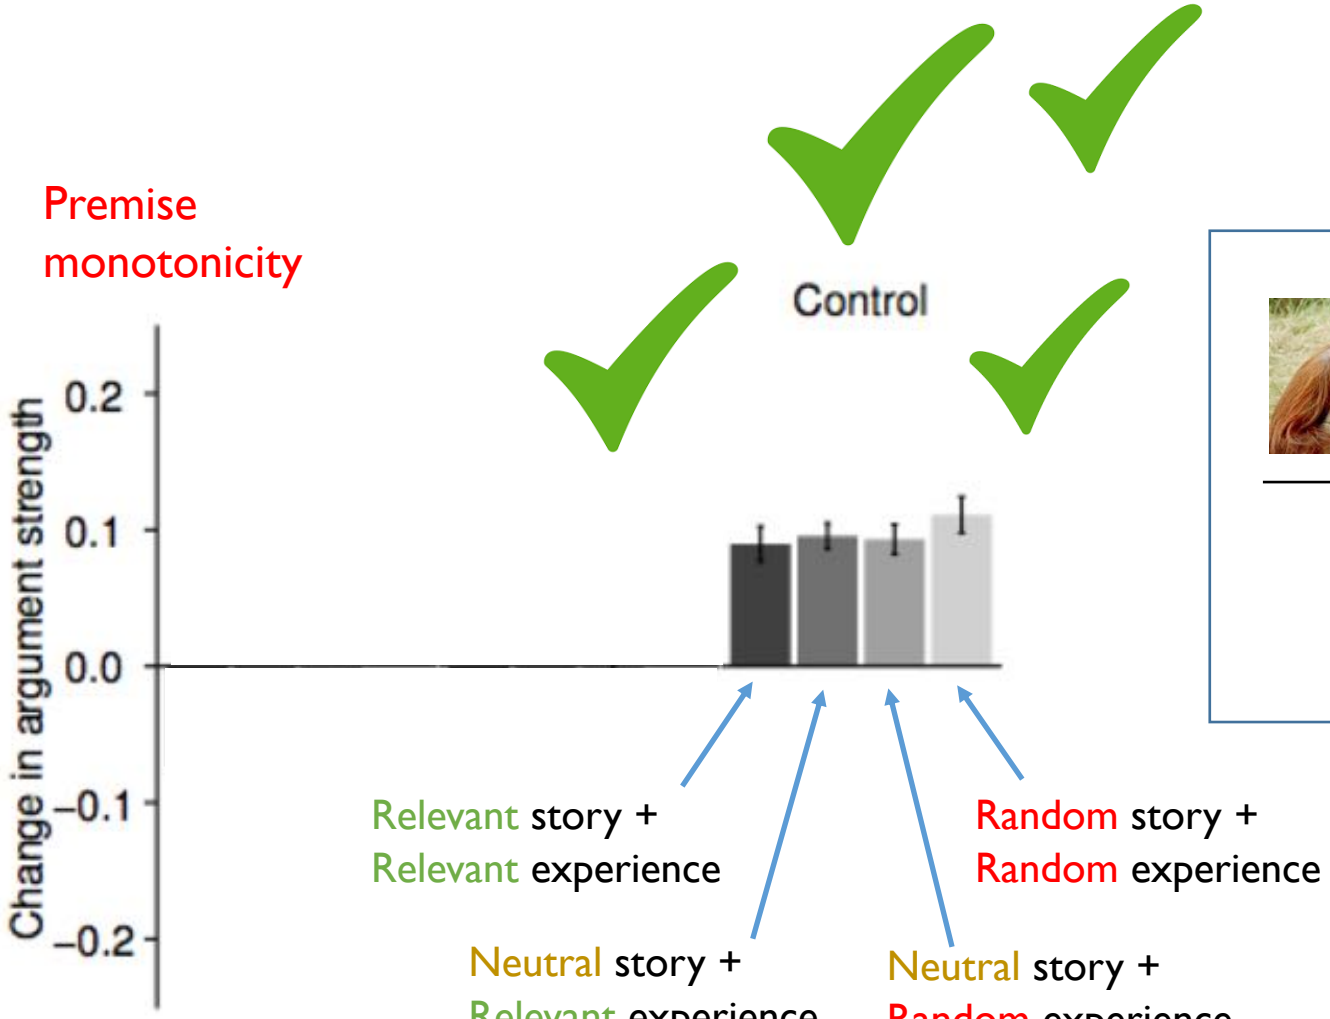

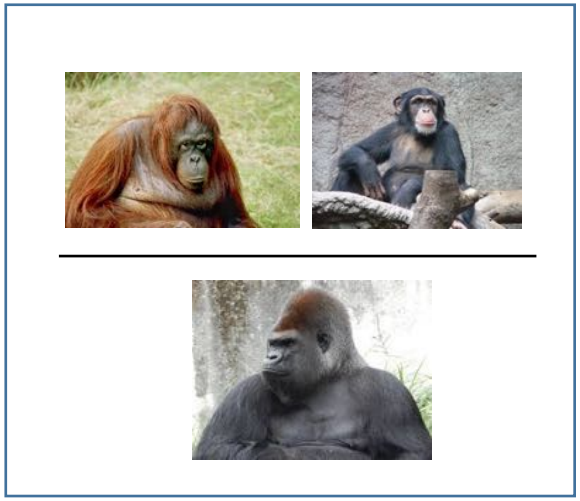

Premise non-monotonicity

The control argument produces monotonic reasoning in all four conditions

Premise monotonicity

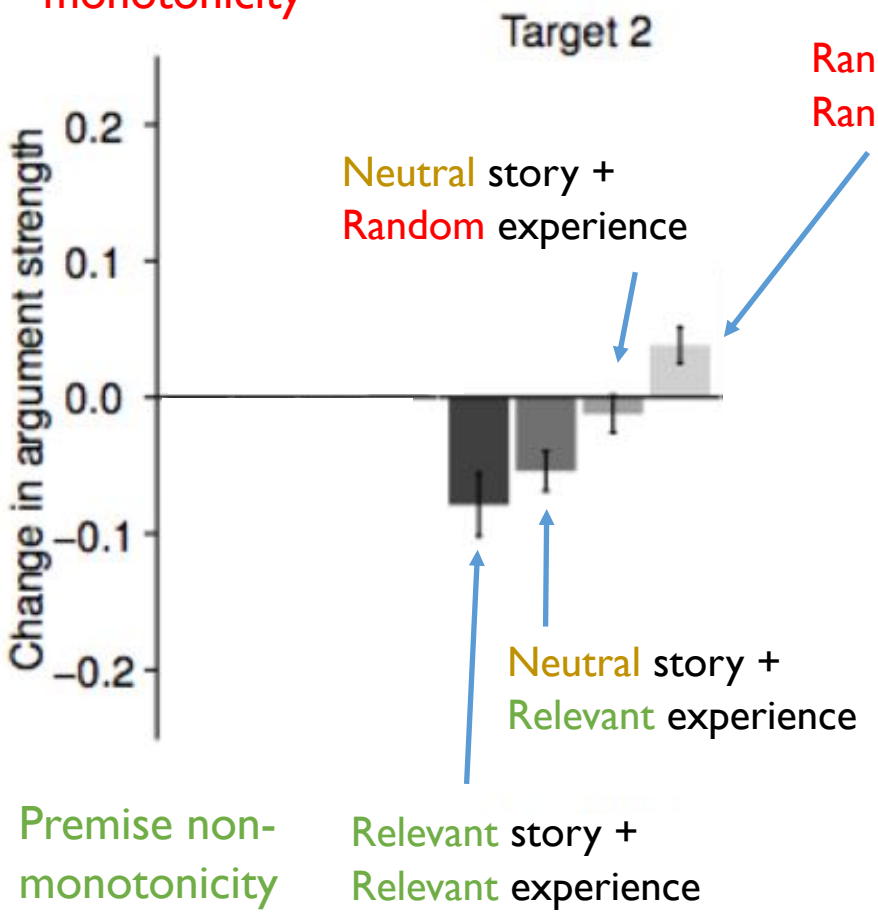

Premise non-monotonicity

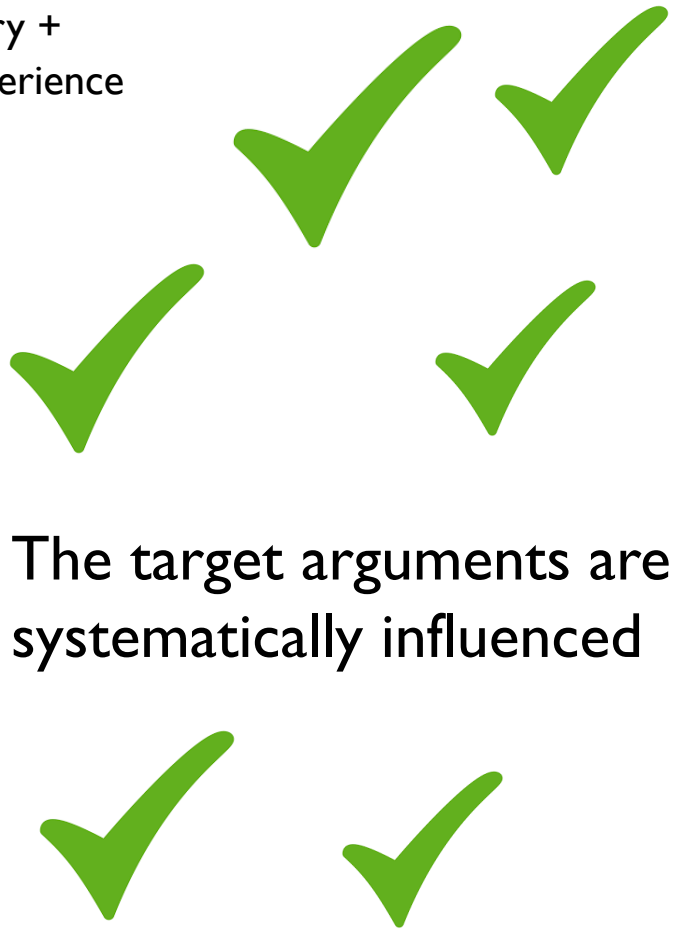

The target arguments are systematically influenced

Premise monotonicity

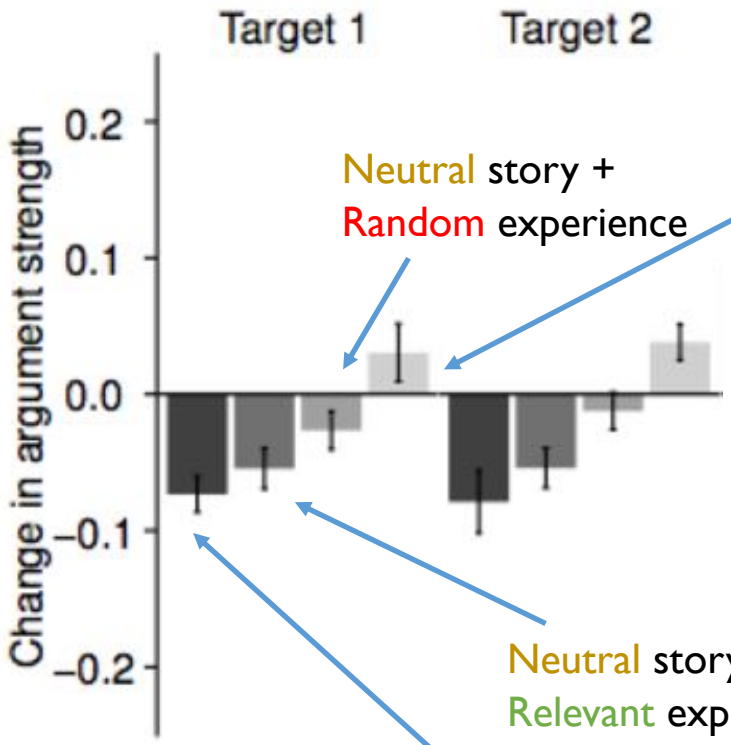

Random story +
Random experience

Neutral story +
Random experience

Neutral story +
Relevant experience

Premise non-monotonicity

Relevant story +
Relevant experience

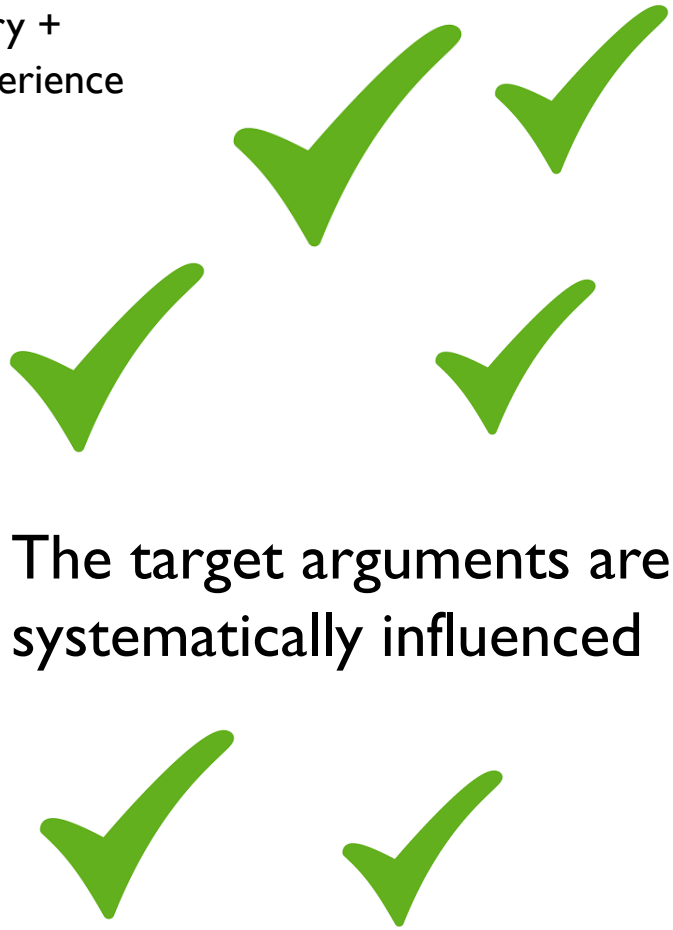

The target arguments are systematically influenced

What does it mean?

We used a psychological theory to help us design an experiment

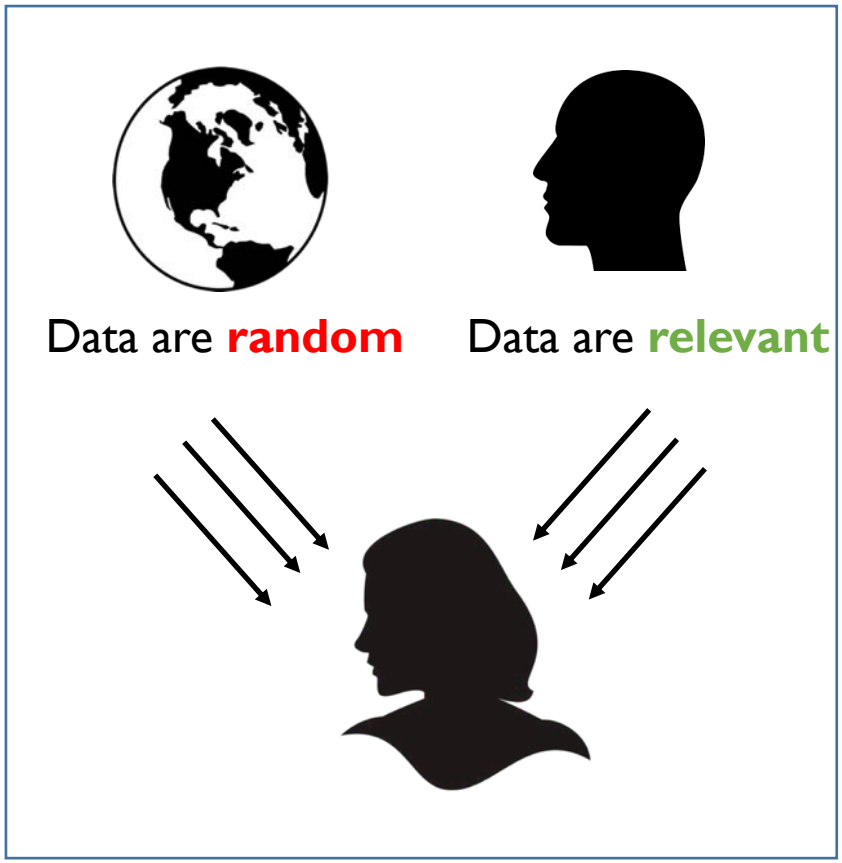

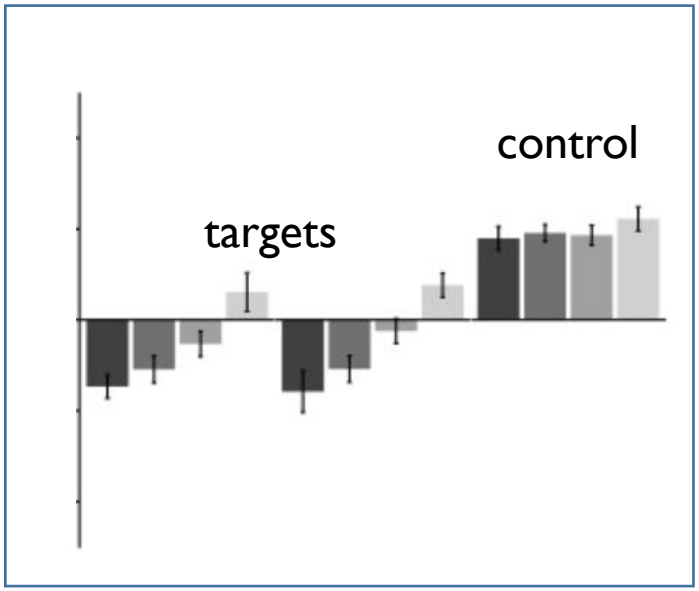

The results agree with our theory and not with competing theories. Yay!

Call back to lecture 1… we can get a more precise test using a *computational model*

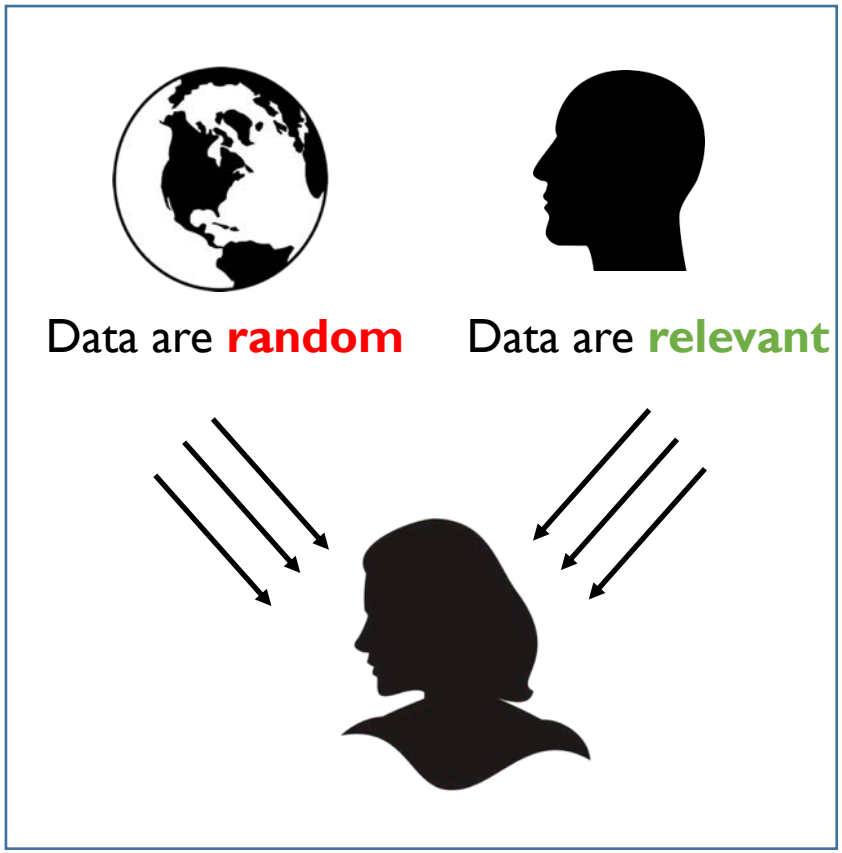

Humans:

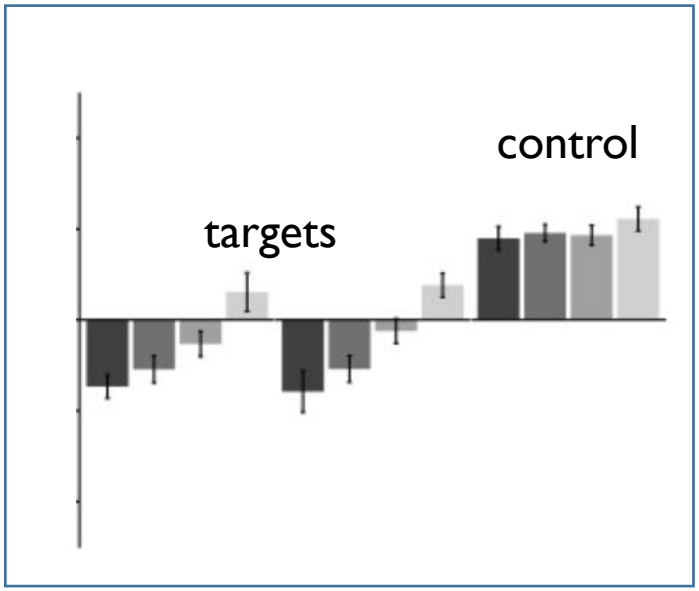

Model:

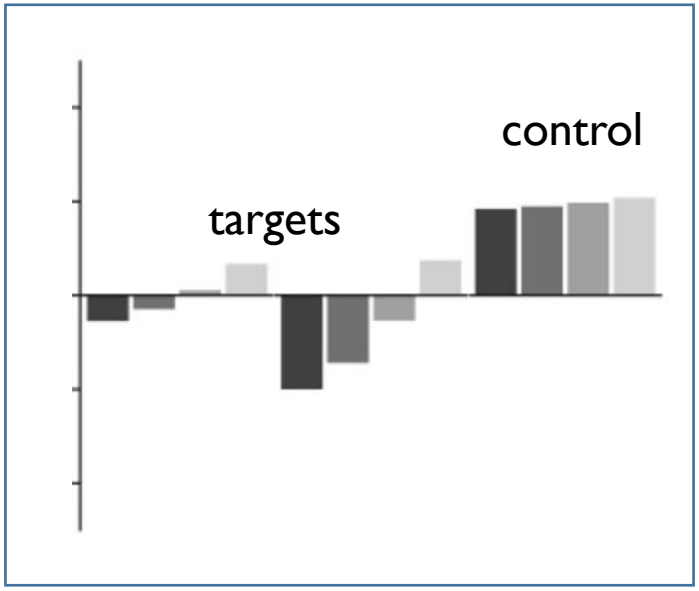

People are smart (& psychology is hard)

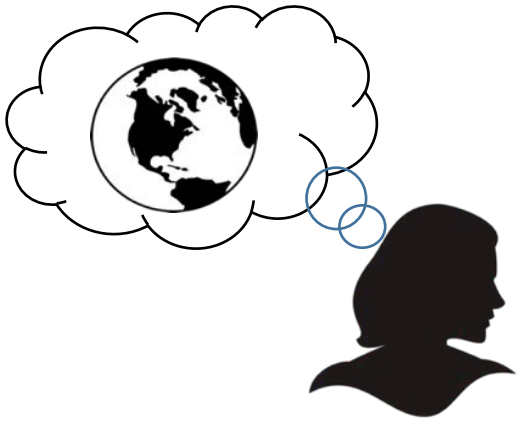

Sometimes inductive reasoning is driven by intuitive theories of the world

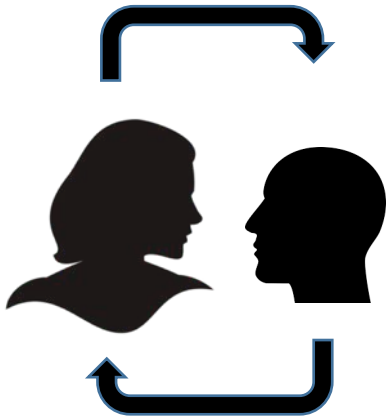

Sometimes inductive reasoning is driven by "theory of mind"

We switch flexibly between these in a sensible way – and the inductive power of an argument changes when we do so.

What are the limitations?

Possible limitations?

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Probably not. We collected data from 538 participants

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Maybe? Our participants were recruited online: diverse in age and gender, but narrow in nationality (USA) and probably above average in education

Important question: is there a plausible reason to think this might matter?

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Probably not.

The non random ordering (i.e., fillers mostly first) was intentional, and was central to the experimental manipulation

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Absolutely not. It would have been absurd to include a "relevant story + random experience" condition… this would introduce a confound colloquially referred to as "lying to participants"

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Definitely a limitation. We used a fixed set of six arguments, all of which were about animals.

Important question: why might this matter? (hint… people have different knowledge)

- Sample size too small?
- Sample not representative?
- Stimulus order not randomised?
- Factors not fully crossed?
- Limited range of arguments?
- Limited range of phenomena?

Definitely a limitation. We only looked at the premise (non) monotonicity effect.

There are good reasons to think the same manipulations should influence other inductive phenomena

Final thoughts?

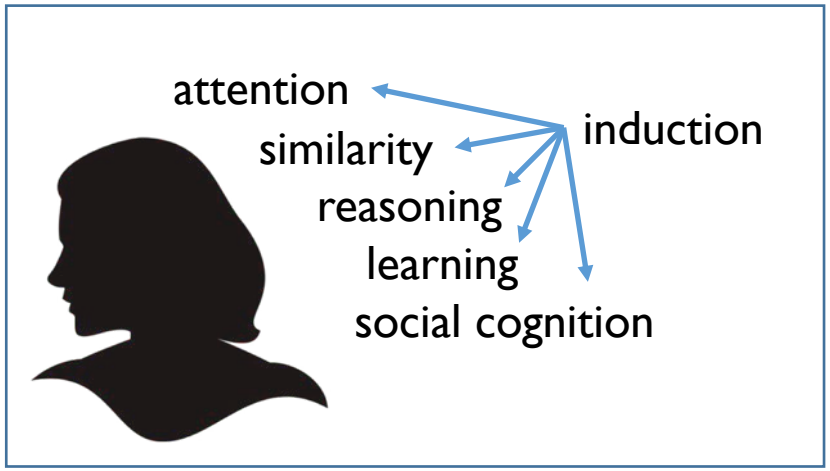

A "single" task often requires psychologists to think about several different aspects of human cognition

A "single" task often requires psychologists to think about several different aspects of human cognition

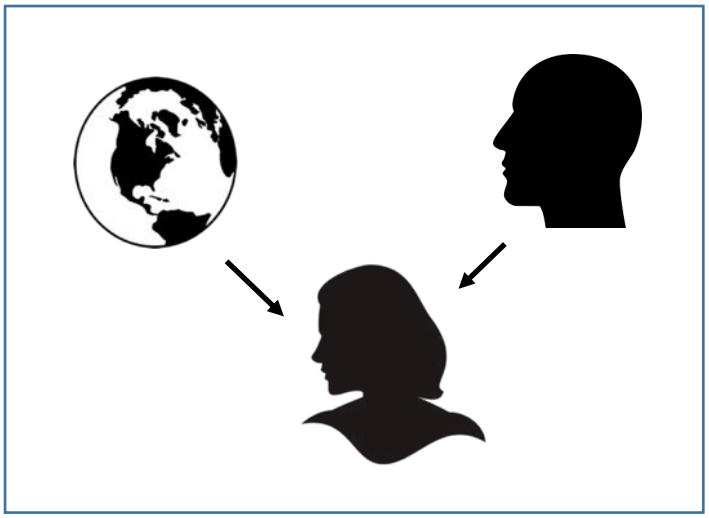

Using psychological theories to guide experimental design yields insight into how cognition works

A "single" task often requires psychologists to think about several different aspects of human cognition

Using psychological theories to guide experimental design yields insight into how cognition works

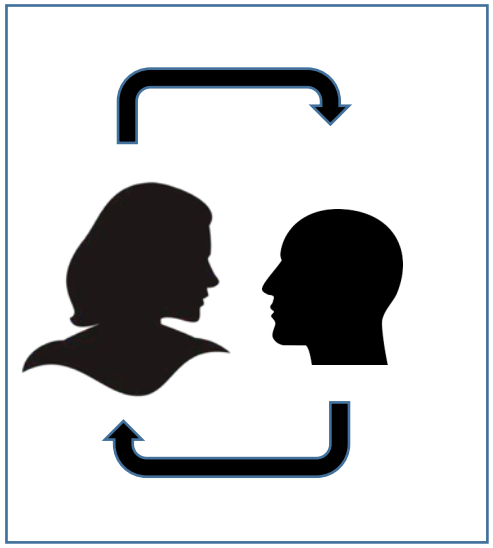

Human reasoning (and cognition generally) can be remarkably complicated

A "single" task often requires psychologists to think about several different aspects of human cognition

Using psychological theories to guide experimental design yields insight into how cognition works

Human reasoning (and cognition generally) can be remarkably complicated

… which is one of the reasons we build models