# Cultural transmission

http://compcogscisydney.org/psyc3211/

A/Prof Danielle Navarro
d.navarro@unsw.edu.au
compcogscisydney.org

# Where are we?

- L1: Connectionism
- L2: Statistical learning
- L3: Semantic networks
- L4: Wisdom of crowds
- L5: Cultural transmission
- L6: Summary

# Structure of the lecture

- Introduction to the topic
- Revealing inductive biases?
- Example: function learning
- Caveat: distortion by individual differences
- Cumulative cultural evolution

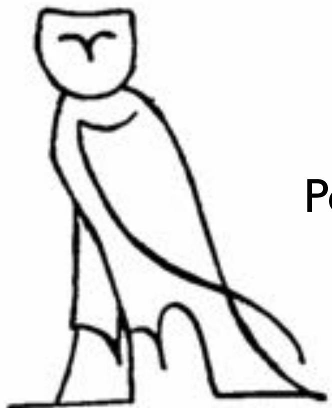

Person #1 draws an owl

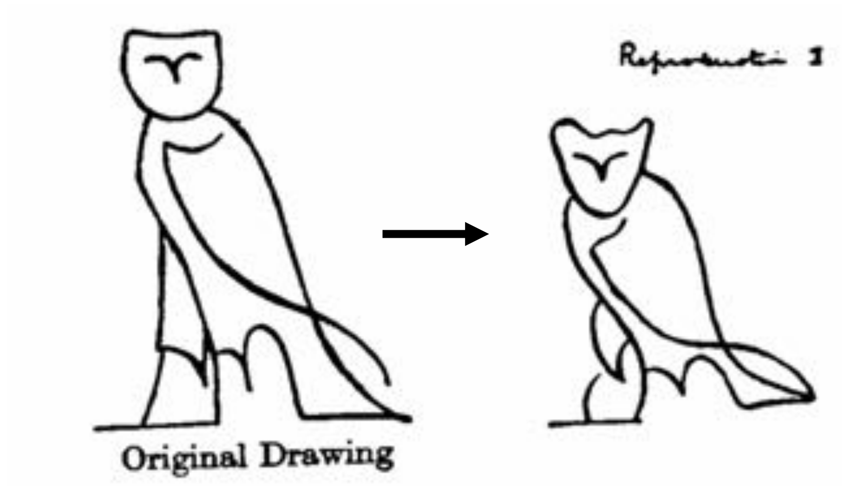


Person #2 attempts to copy it

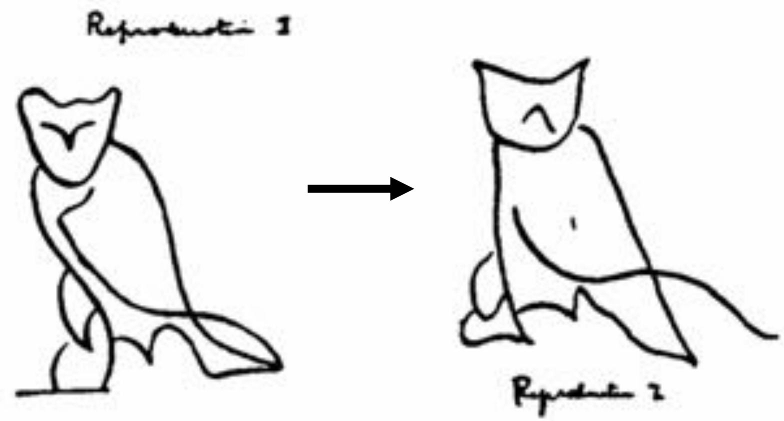

Person #3 makes a copy of the copy

Reproduction 2

Reproduction 3

Person #4 makes a copy of the copy of the copy…

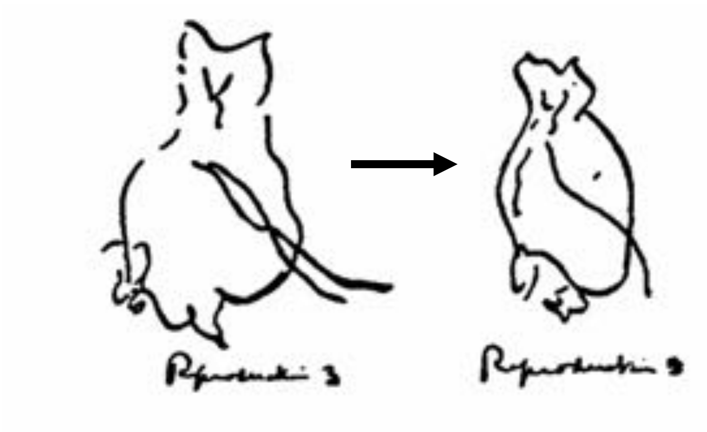

Person #5

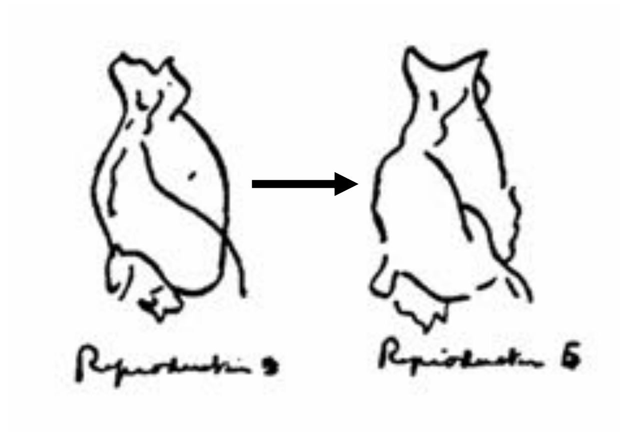

Person #6

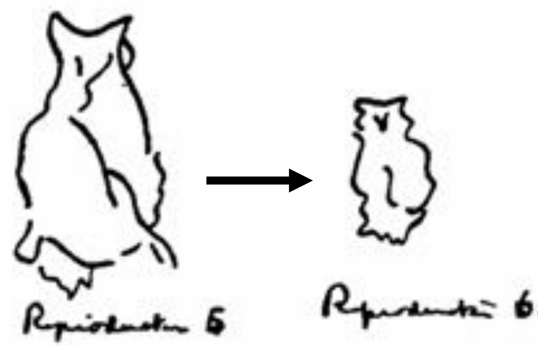


Person #7

Person #8

Person #9

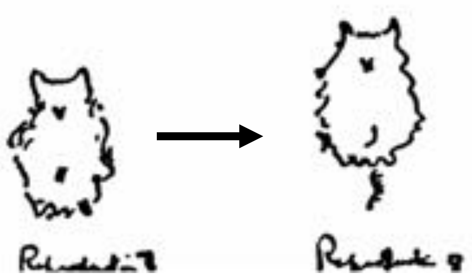

Person #10

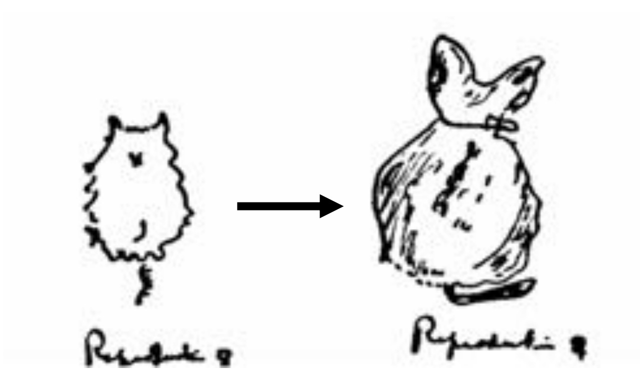

Person #22

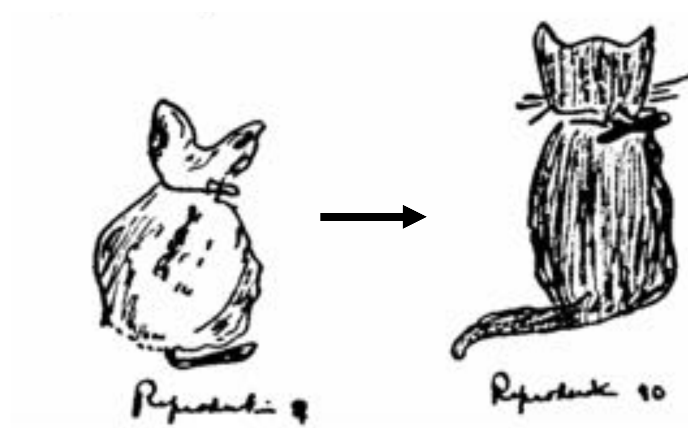

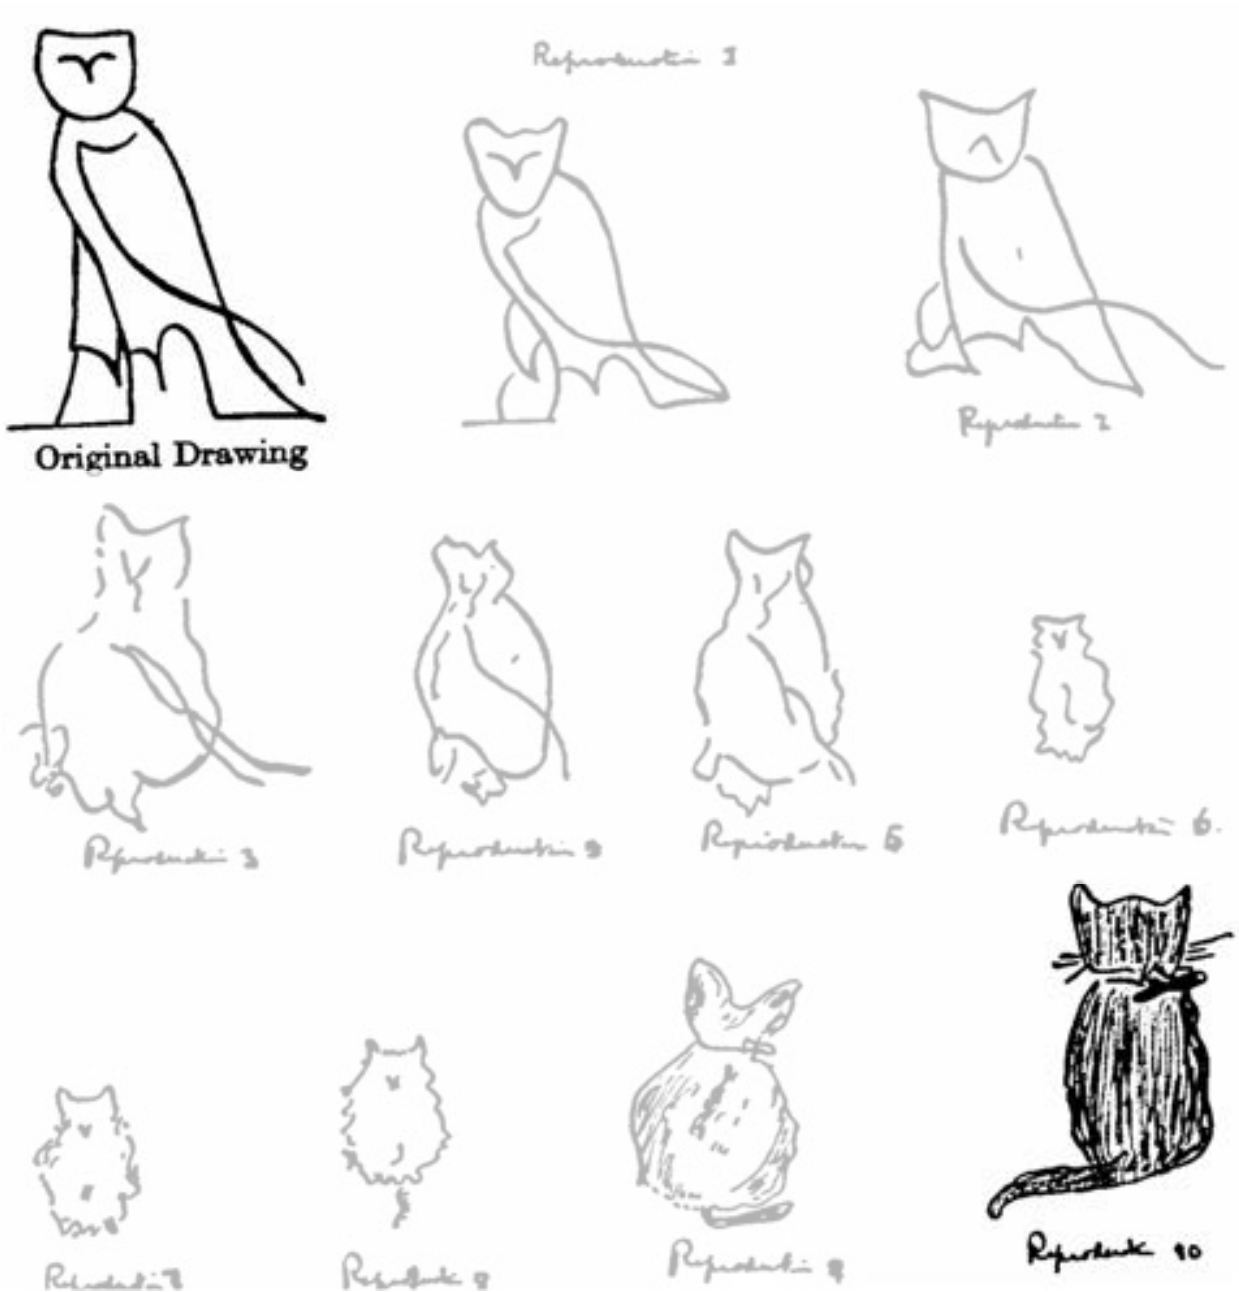


Owl → Cat

Are these changes random?
Are they meaningful?
What processes are going on here?

The method of serial reproduction

Bartlett (1920)

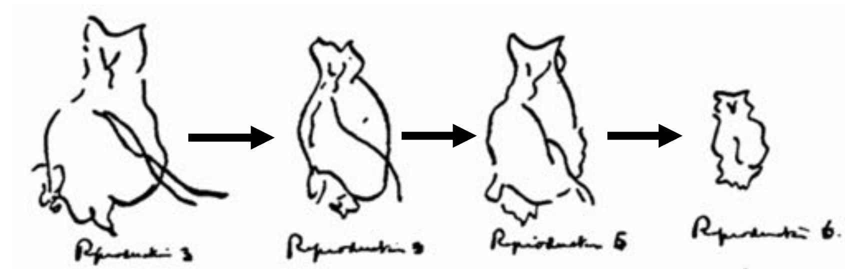

Cultural transmission by iterated learning

Kalish et al (2007)

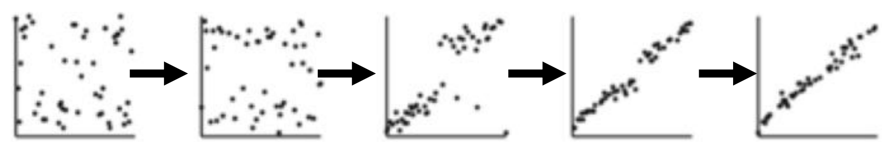

Theoretical claim:

# Iterated learning reveals inductive biases?

# Iterated learning

- Sequential experimental design used to study cultural transmission
- Each person has to learn something… then produce responses
- The responses from one person become the data for the next person

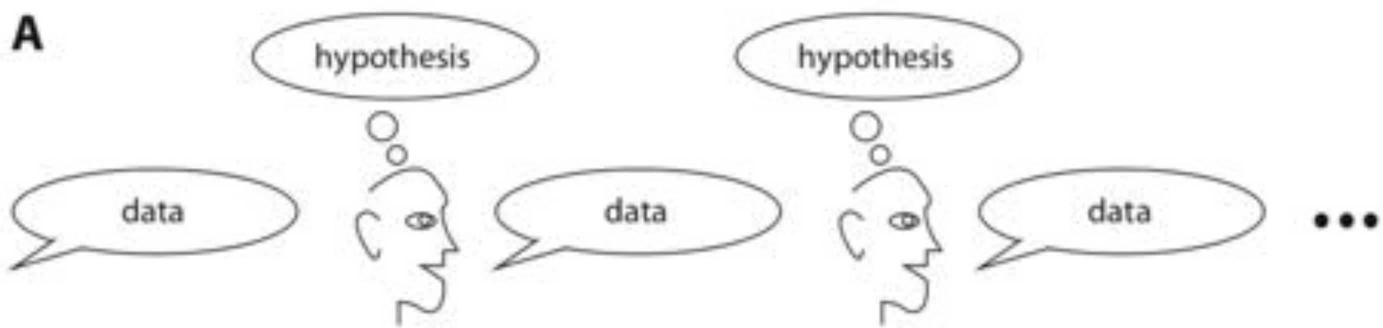

# Bayesian iterated learning

(Kalish et al 2007; Griffiths & Kalish 2007)

- How do these "chains" of learners behave?
  - Suppose each person is a Bayesian reasoner
  - Each person has a prior P(h), then sees data d
  - Responses generated by sampling from P(h|d)

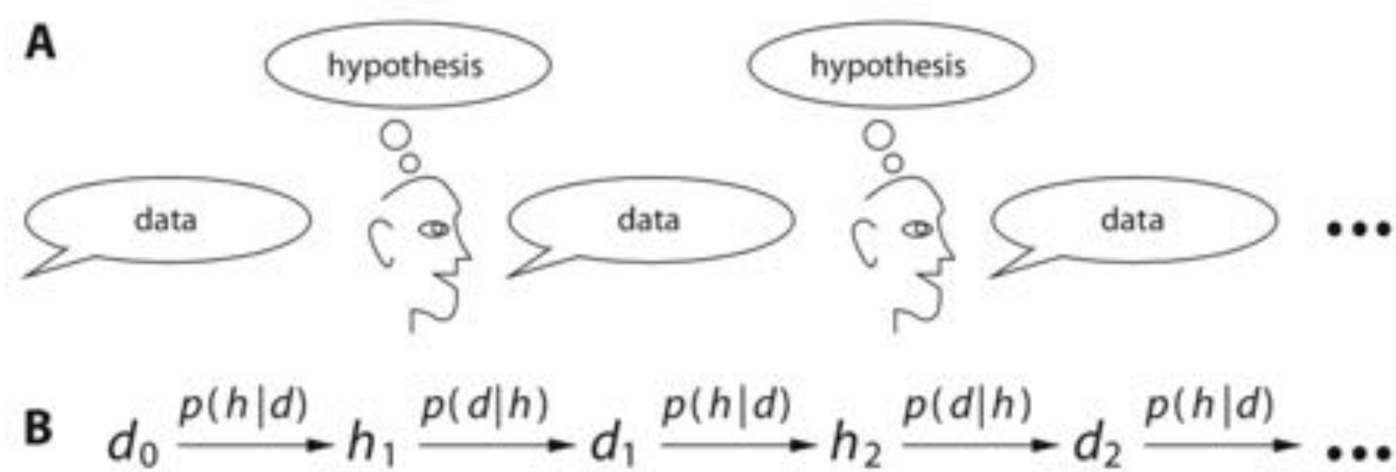

# Bayesian iterated learning

(Kalish et al 2007; Griffiths & Kalish 2007)

They did some maths

$$
\begin{aligned}
P(h_n = i) &= \sum_j P_{\text{samp},PA}(h_n = i \mid h_{n-1} = j) P(h_{n-1} = j) \\
&= \sum_j \sum_{d \in \mathcal{D}} P_{\text{samp}}(h_n = i \mid d) P_{PA}(d \mid h_{n-1} = j) P(h_{n-1} = j) \\
&= \sum_{d \in \mathcal{D}} P_{\text{samp}}(h_n = i \mid d) \sum_j P_{PA}(d \mid h_{n-1} = j) P(h_{n-1} = j) \\
&= \sum_{d \in \mathcal{D}} P_{\text{samp}}(h_n = i \mid d) P_{PA}(d) \\
&= \sum_{d \in \mathcal{D}} \frac{P_{PA}(d \mid h_n = i) P(h_n = i)}{P_{PA}(d)} P_{PA}(d) \\
&= P(h_n = i) \sum_{d \in \mathcal{D}} P_{PA}(d \mid h_n = i),
\end{aligned}
$$

The formal details don't matter for this class, but the take home message is that when Bayesian learners all share the same prior P(h), an iterated learning chain eventually starts to reflect the biases in that prior

* There are other conditions too

Empirical test:

# A function learning task

# Function learning problems

How do people learn to the relationship between different (continuous) quantities?

This function might be tell you the of how many properties you could afford to rent, plotted as a function of your income

Output value

(e.g., properties you can afford)

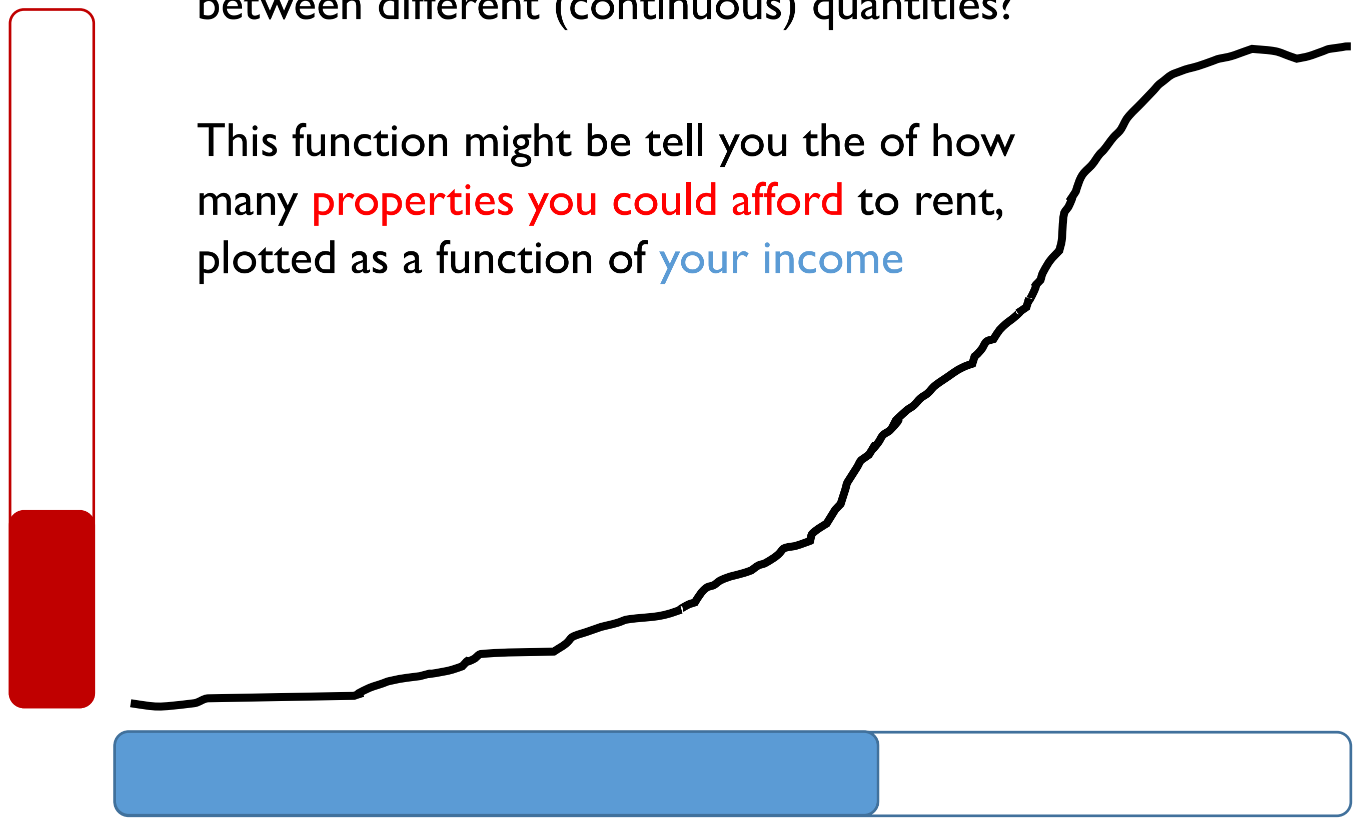

Input value

(e.g., income)

# Function learning problems

- There are many possible types of function
- Consistent finding in the function learning literature is that people find it easiest to learn positive linear functions

Positive linear

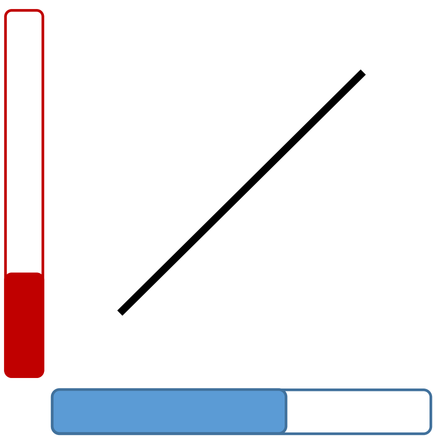

Negative linear

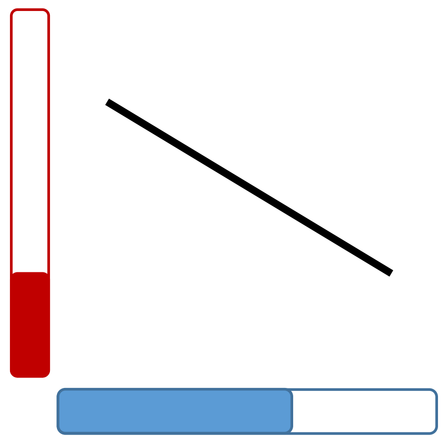

U-shaped

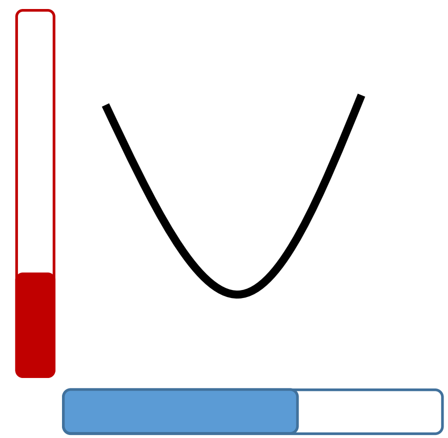

- Prediction: iterated learning chains for a function learning experiment should be biased towards positive linear

# Function learning task

(Kalish et al 2007)

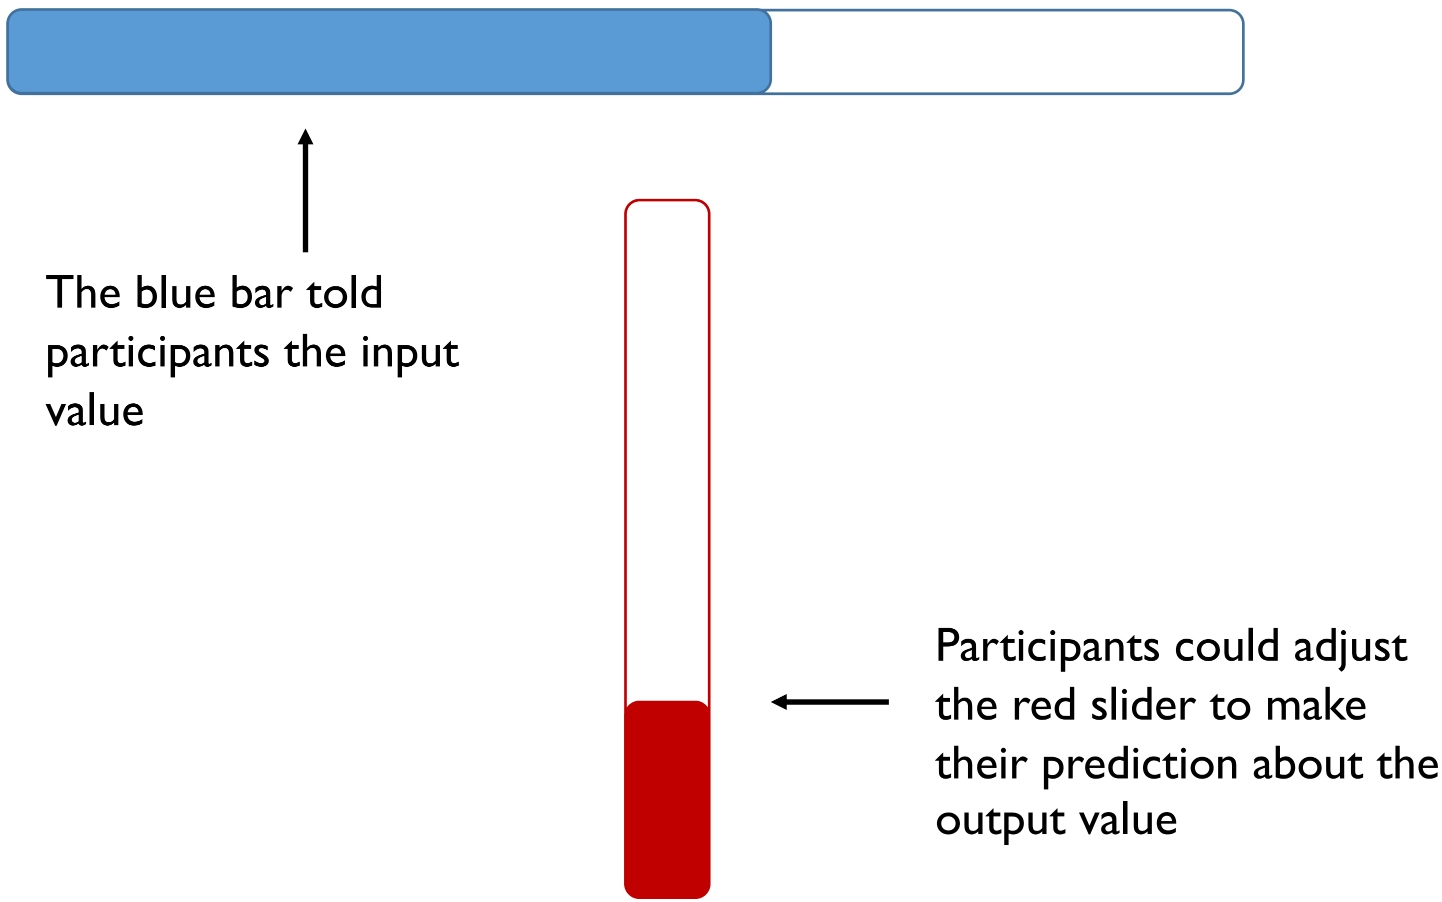

# Function learning task

(Kalish et al 2007)

Show stimulus and get response

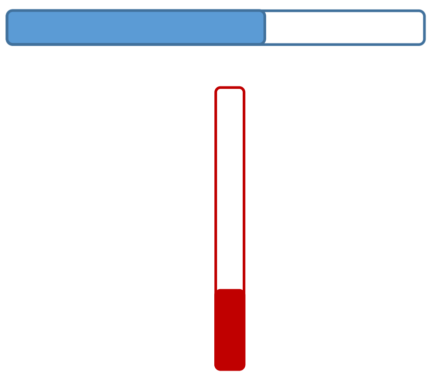

Provide feedback indicating what the correct response should have been

Unbeknownst to participants, this feedback was taken from the responses of the previous participant

- Repeat for many trials
- Followed by a test phase where no feedback is given
- Test phase responses become feedback for next person

# Bayesian iterated learning

(Kalish et al 2007)

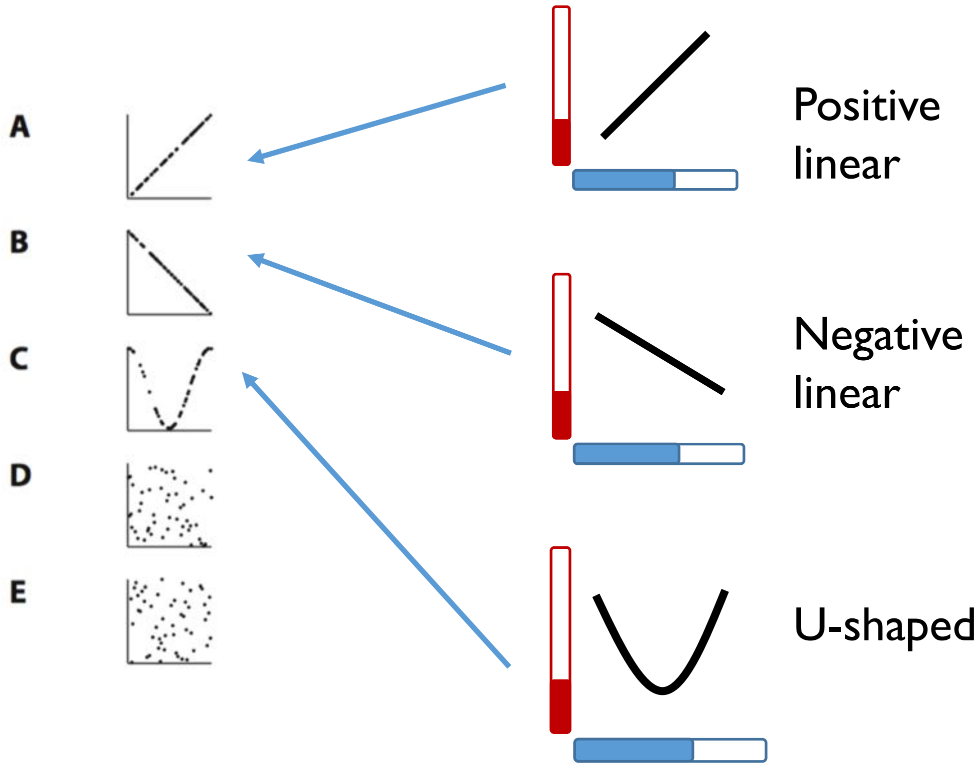


Several iterated learning chains were "initialized" with different functions (i.e., used as feedback for the first person)

# Bayesian iterated learning

(Kalish et al 2007)

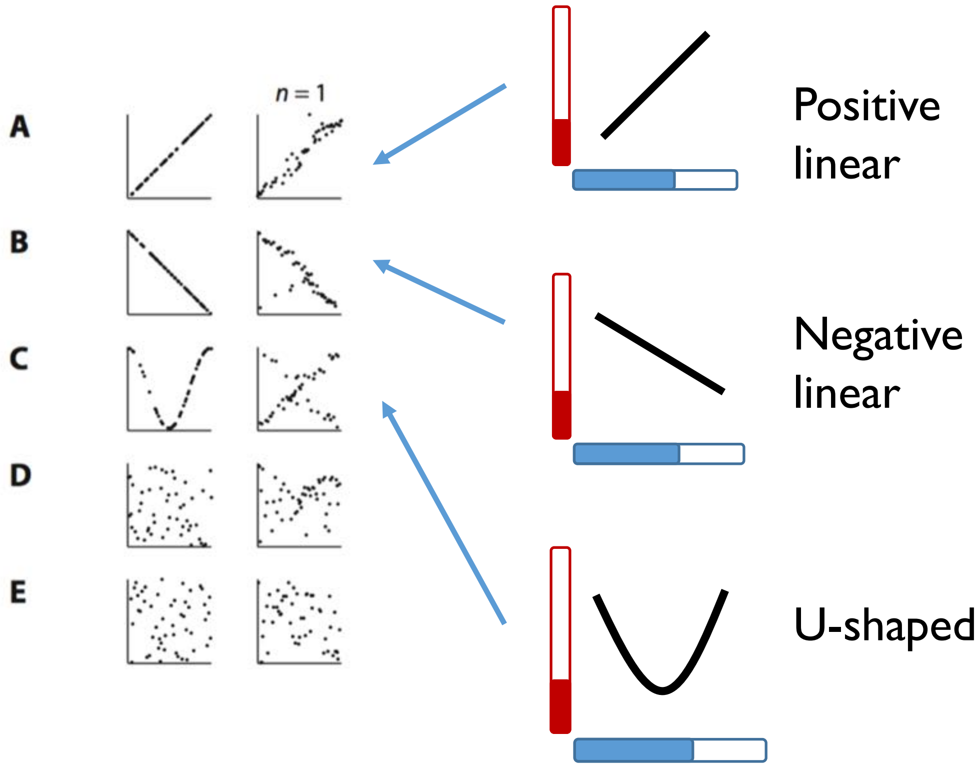


Responses from the first person distort the functions in a systematic way…

These now become data for the next person

# Bayesian iterated learning

(Kalish et al 2007)

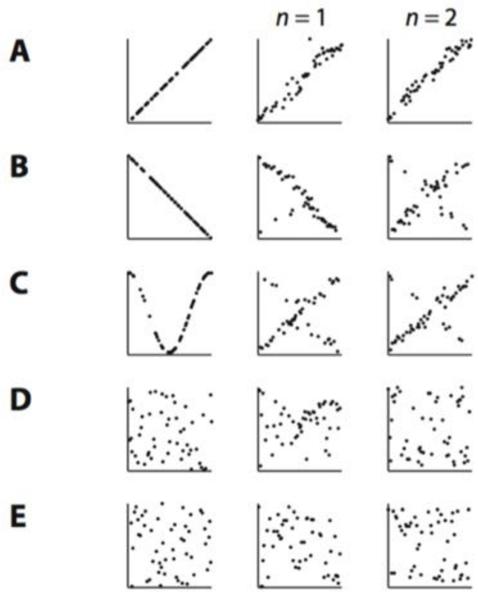


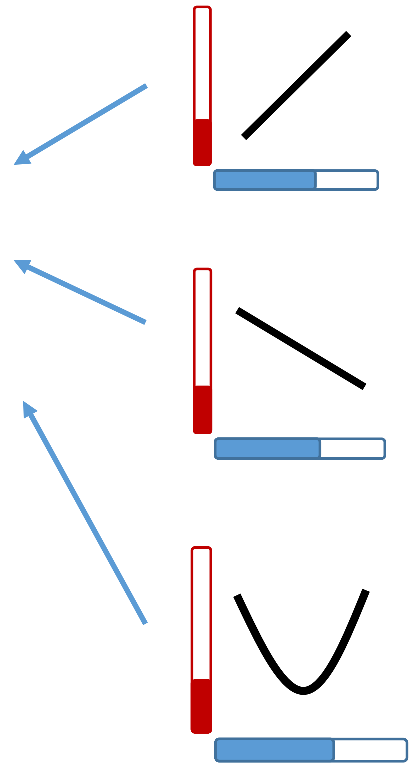

Positive linear

Negative linear

U-shaped

The second person's responses distort the function a bit more…

# Bayesian iterated learning

(Kalish et al 2007)

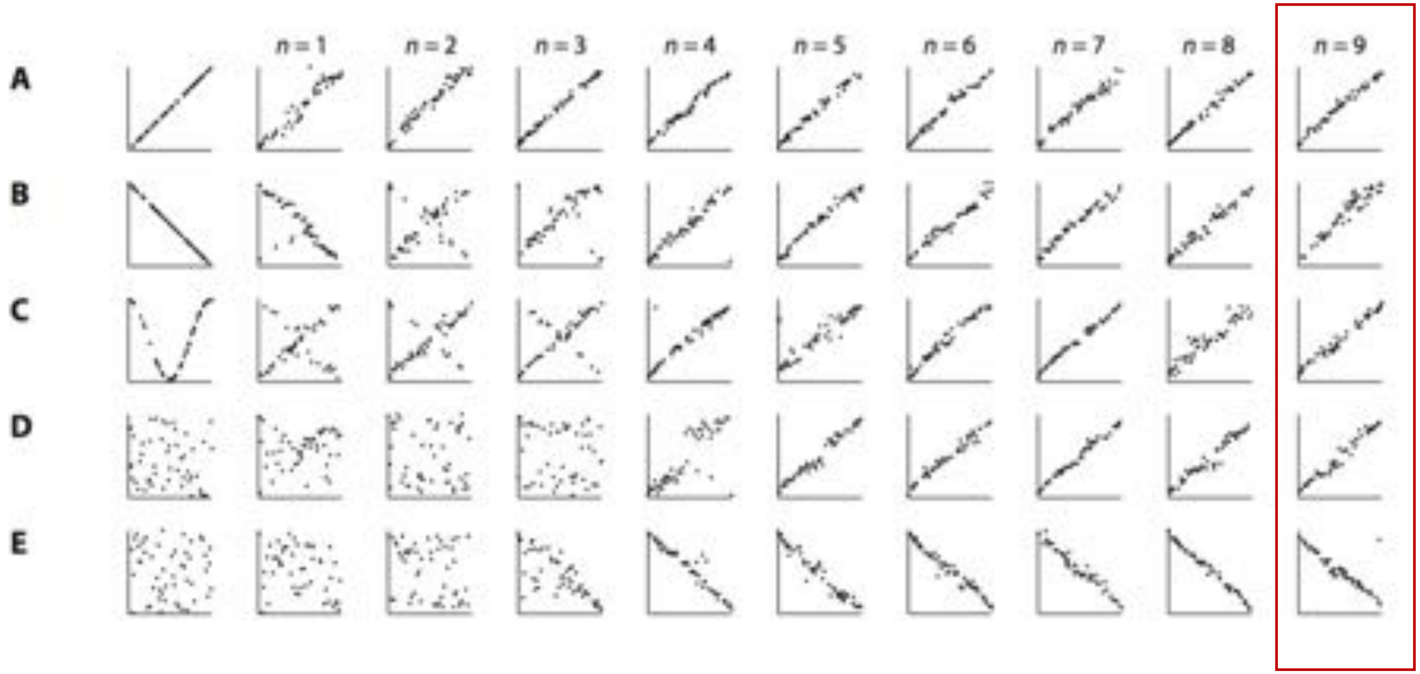


… by the ninth person, the biases of the participants have overwhelmed the input

# What happens when people have different biases?

# Individual differences matter

(Navarro et al 2018)

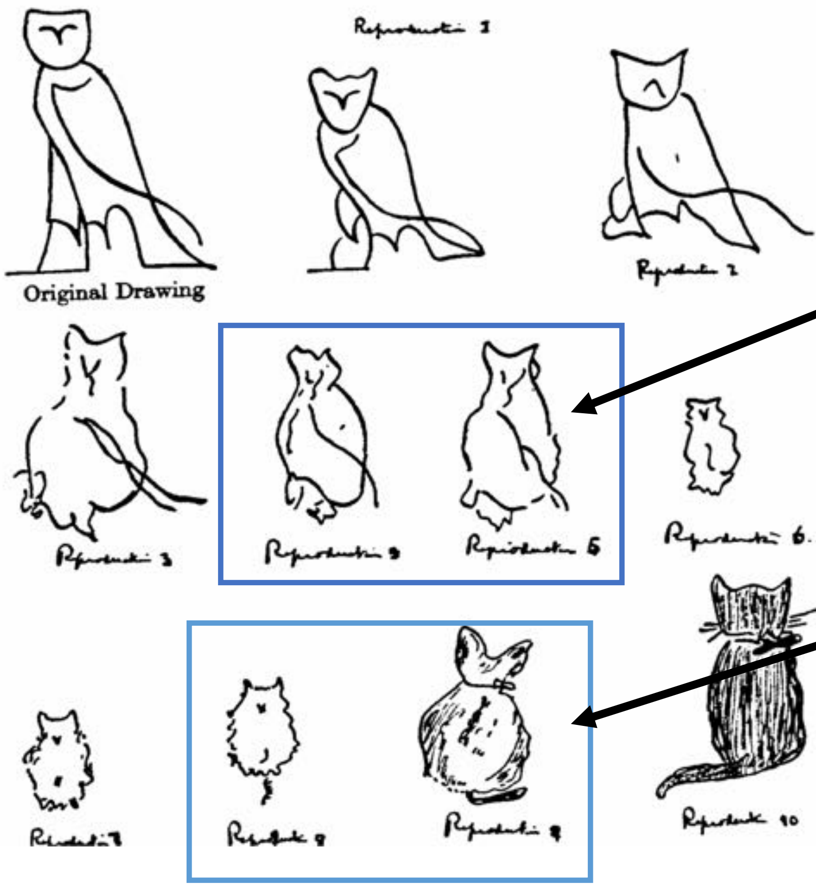


Some people reproduce faithfully

Others do not

# Individual differences matter

(Navarro et al 2018)

Who will win the
2016 election?

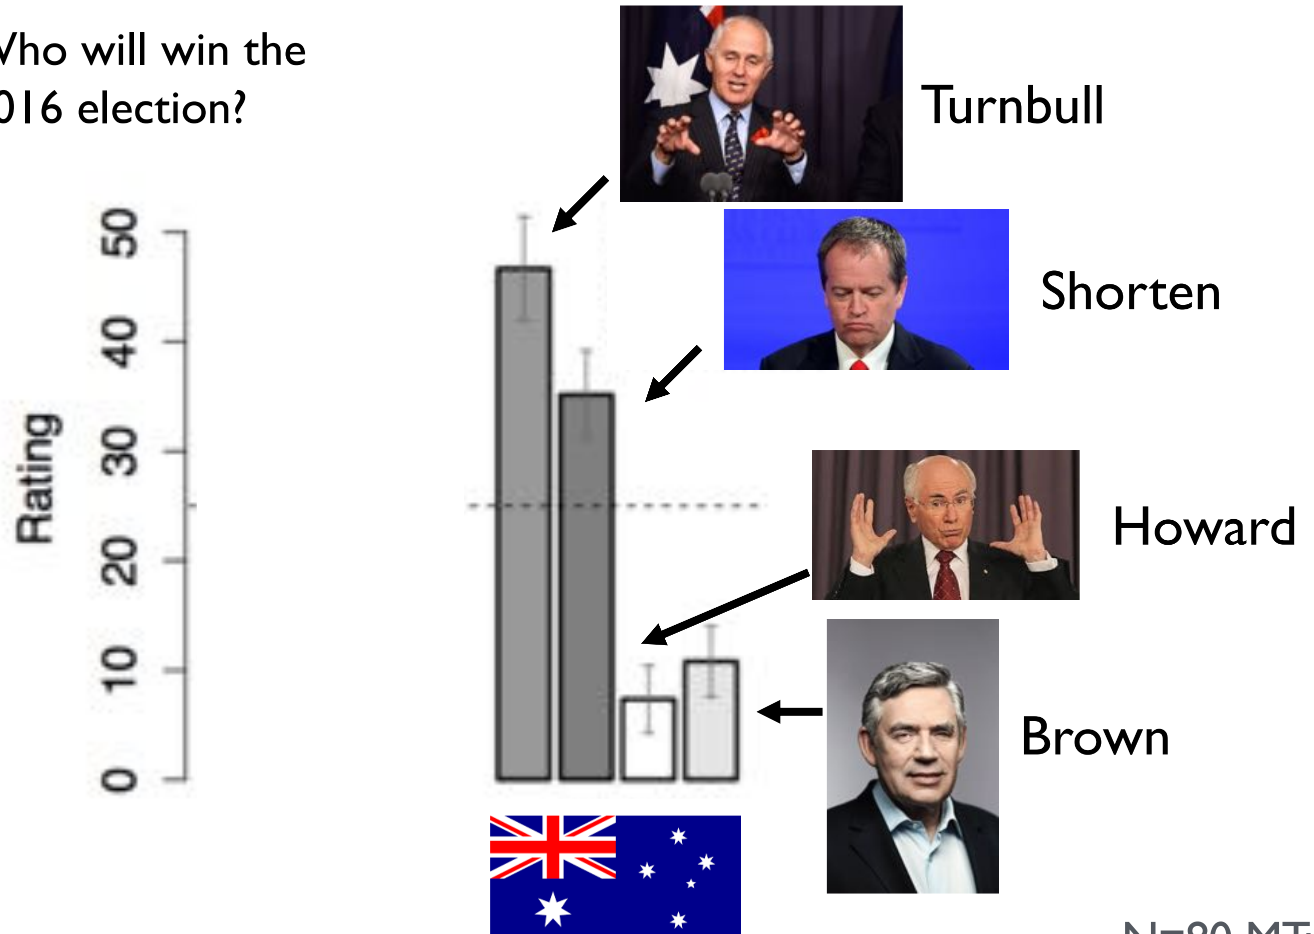


N=80 MTurk workers
and UNSW students

# Individual differences matter

(Navarro et al 2018)

US participants have no knowledge of Australian politics

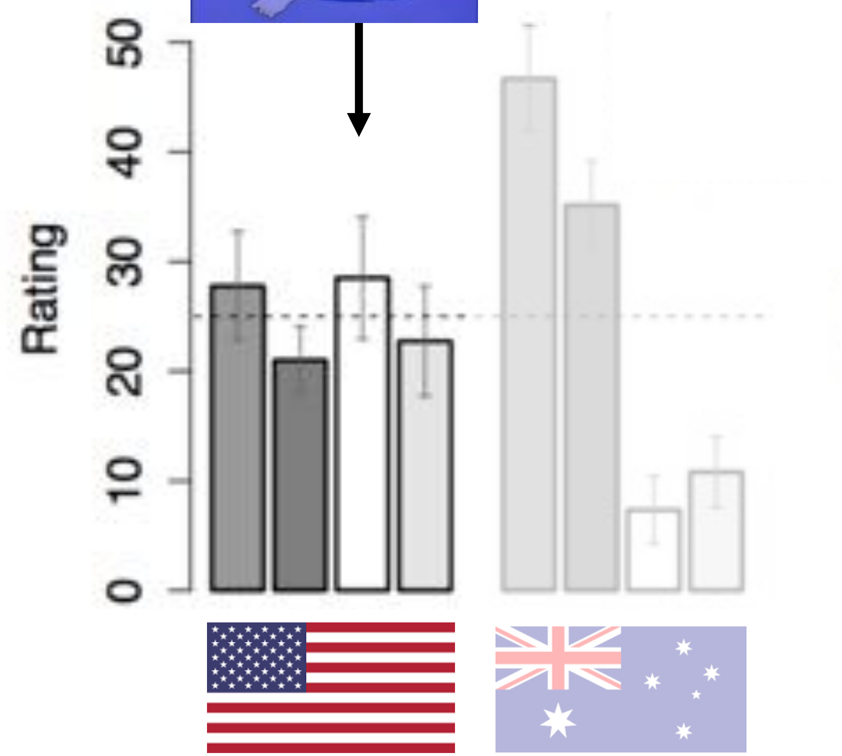


N=80 MTurk workers and UNSW students

# Individual differences matter

(Navarro et al 2018)

The advisor task…

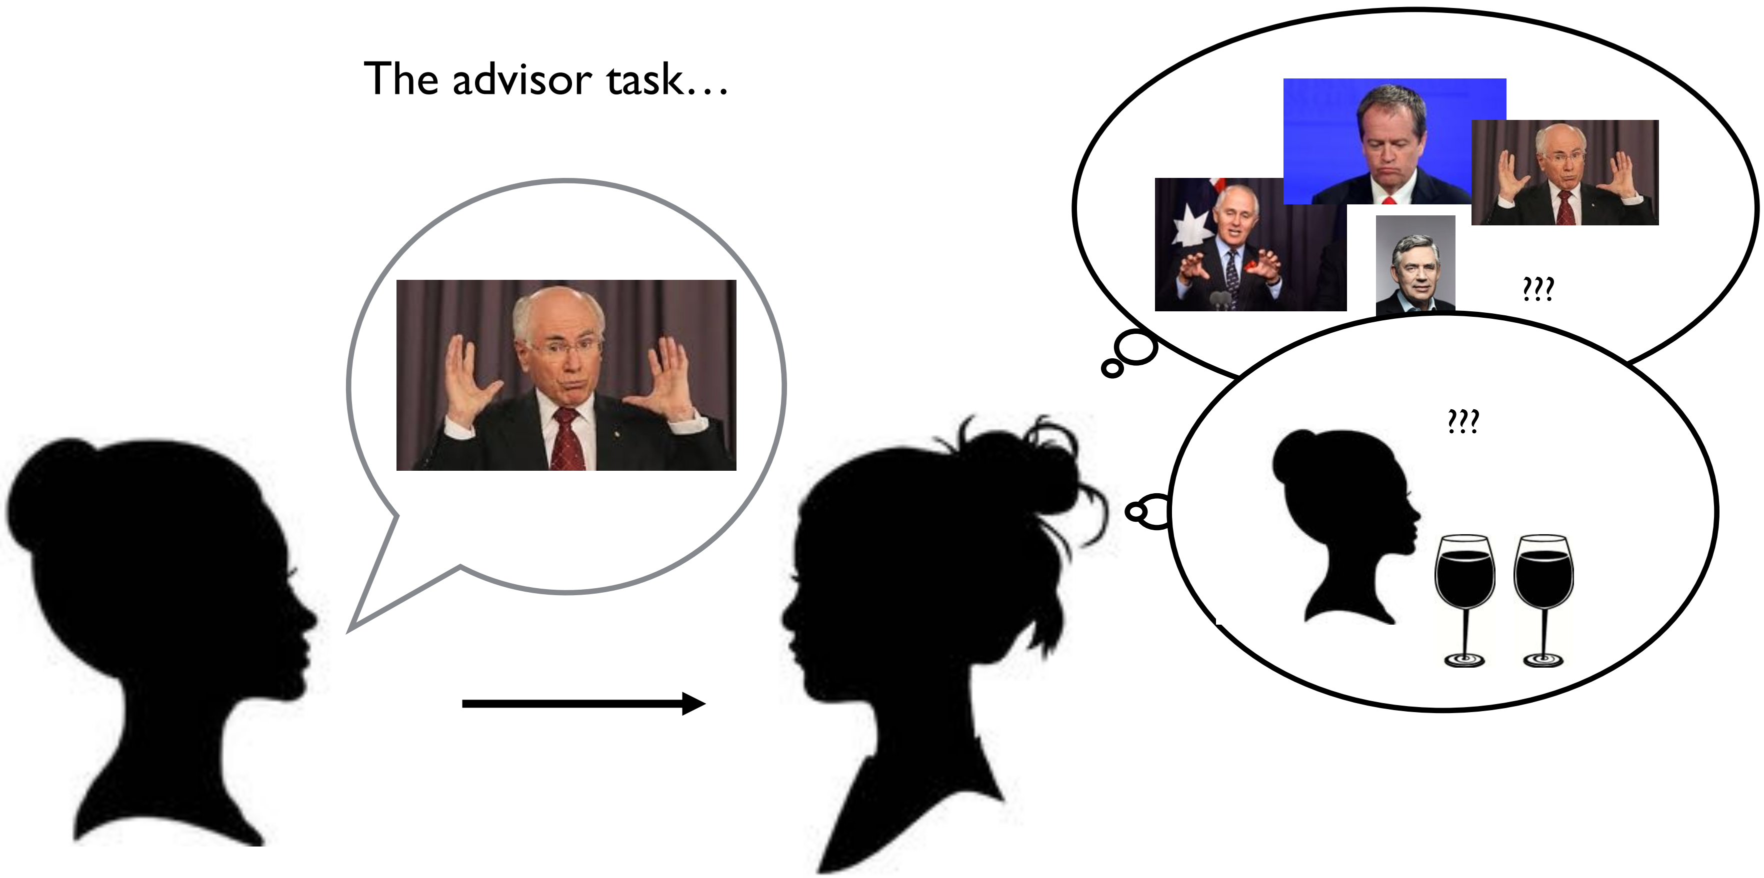

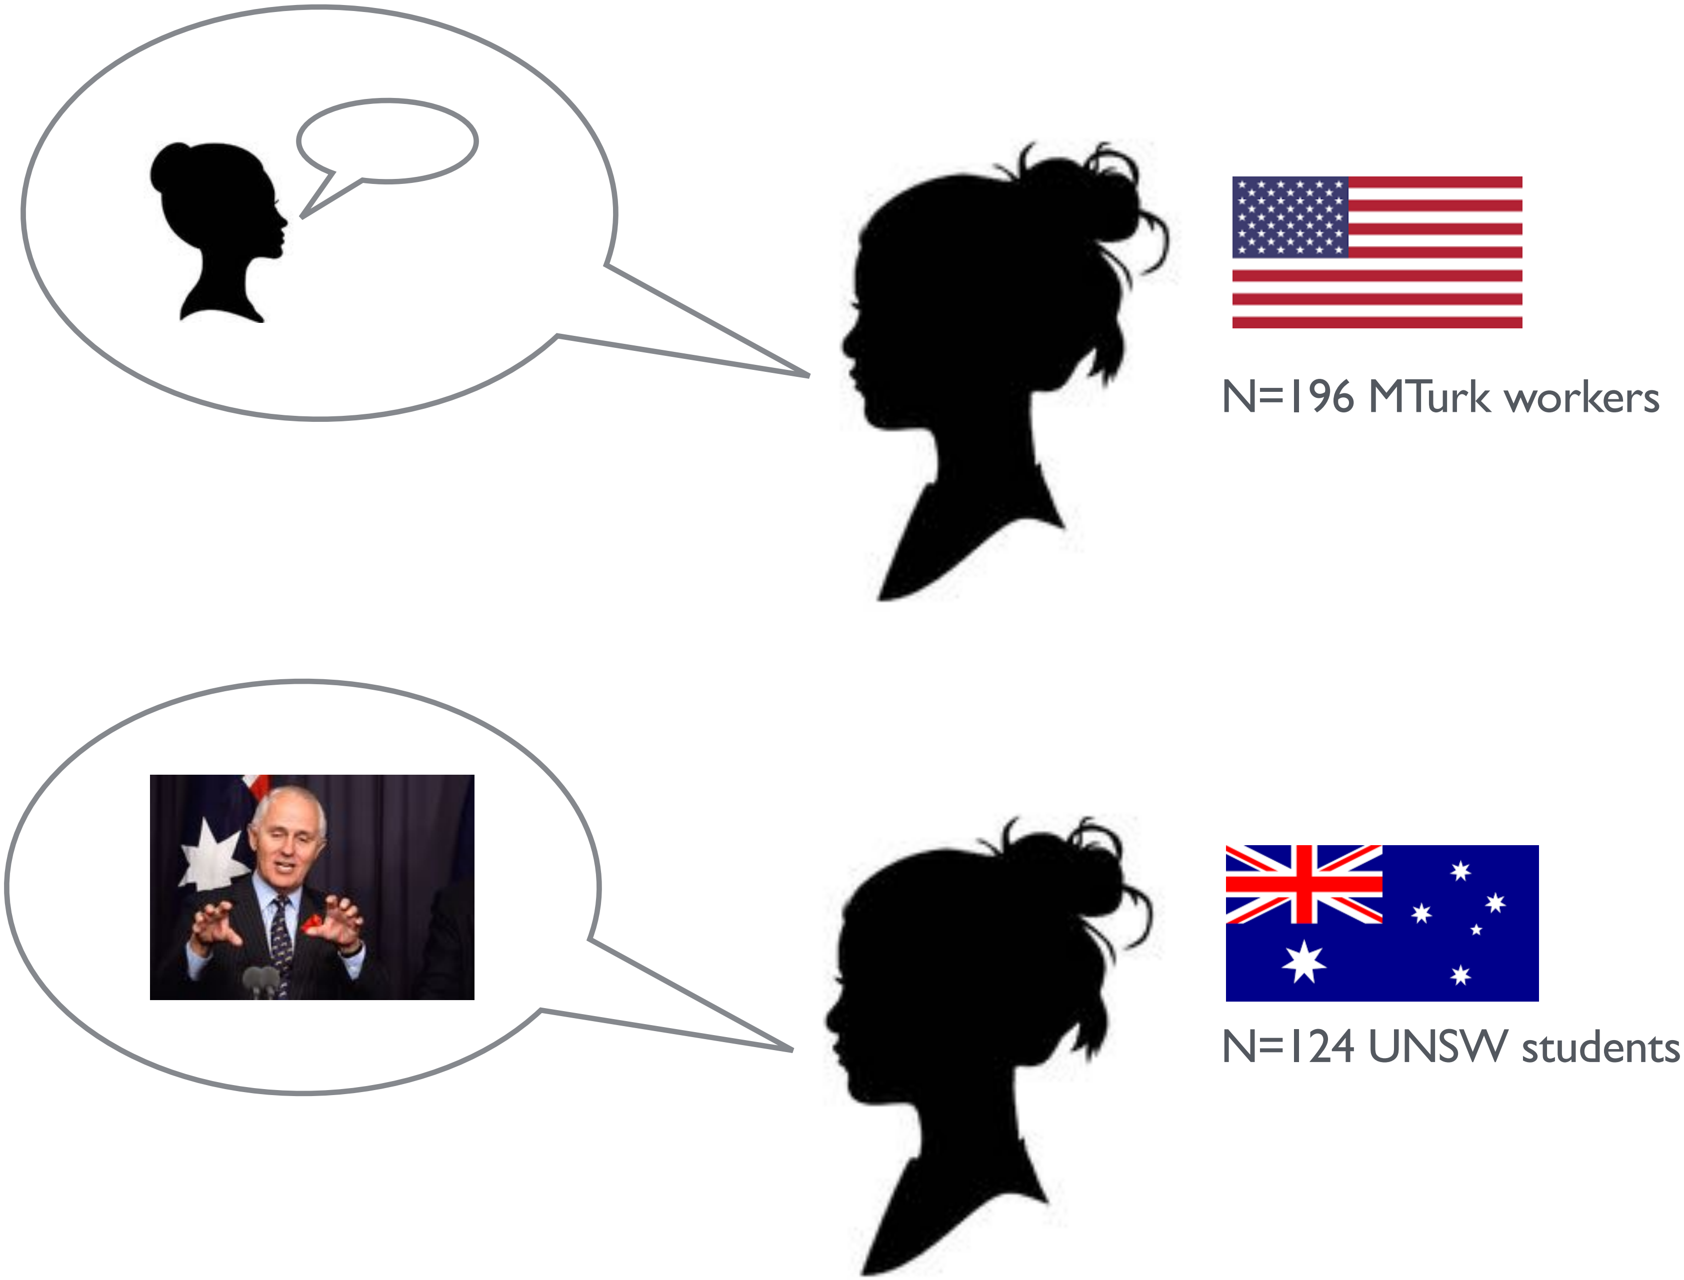
N=196 MTurk workers
N=124 UNSW students

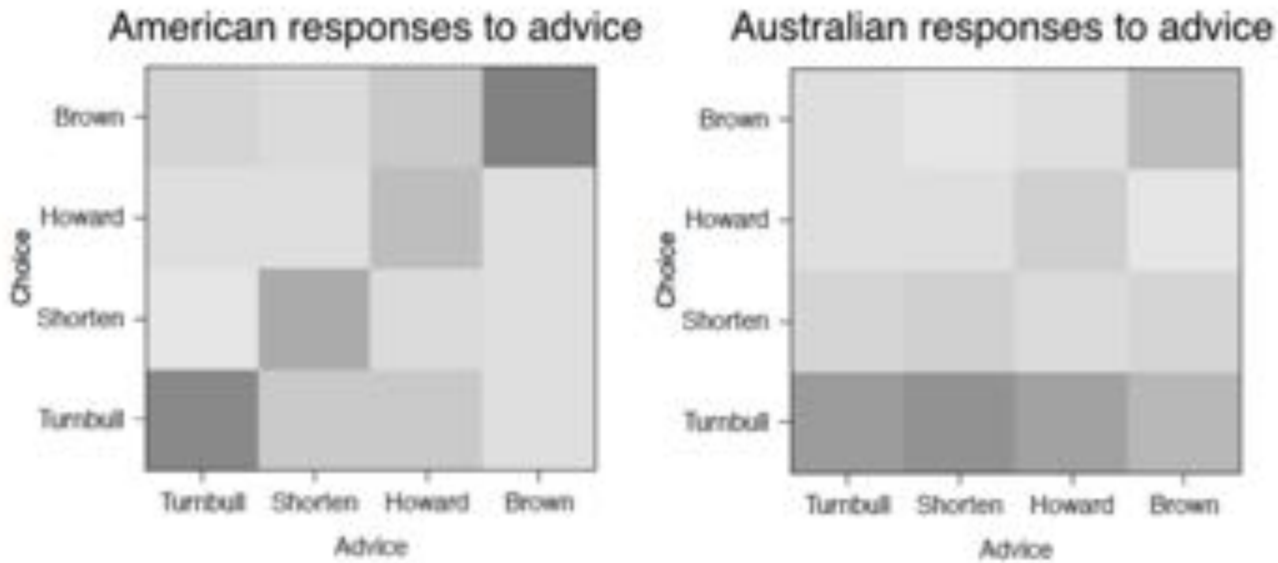


*Figure 15*. Empirical transition matrices estimated from the ADVICE TAKING task. The advice is shown on the $x$ axis, while the choice of the participant is on the $y$ axis. Americans (left panel) tended to follow the advice they were given about the Australian election, as reflected in the darker colors on the diagonal. By contrast, Australians (right panel) were likely to pick Malcolm Turnbull regardless of what they were advised to do. This reflects differences in confidence about their prior beliefs, which affects the degree to which they are influenced by their input.

We can “remix” the responses in different proportions to see what happens when we mix learners with different biases together

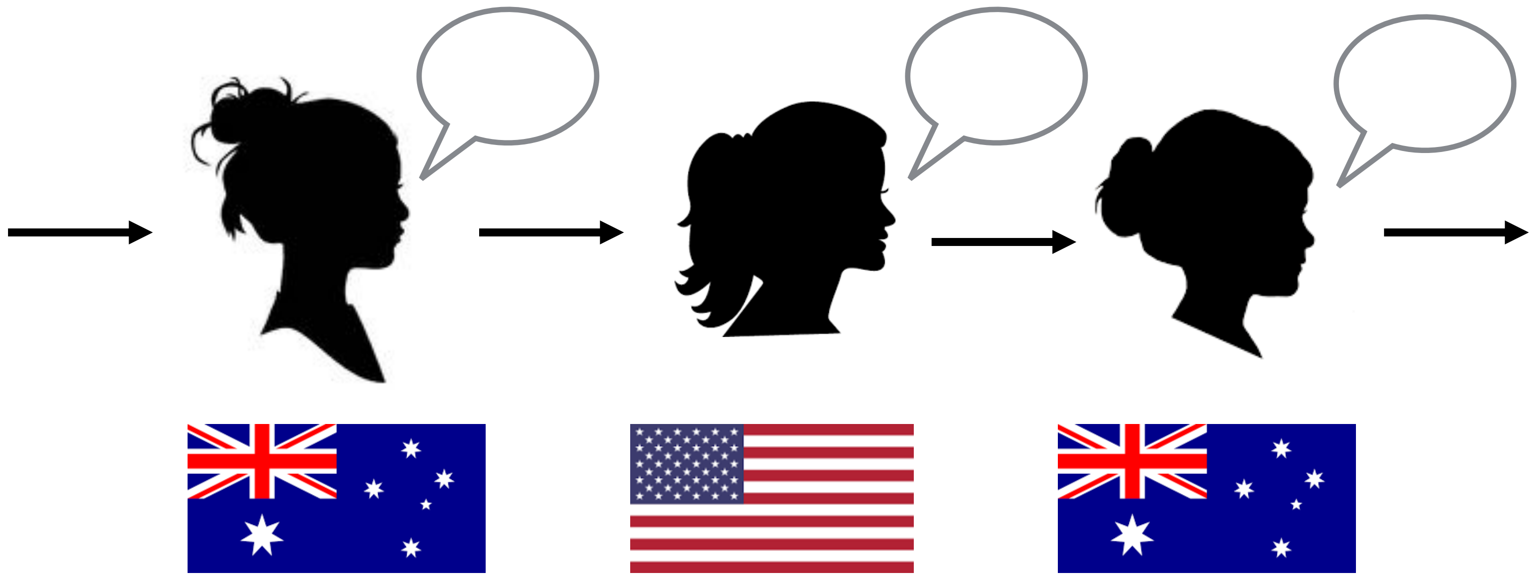

Americans *claim* to be totally ignorant about Australian politics…

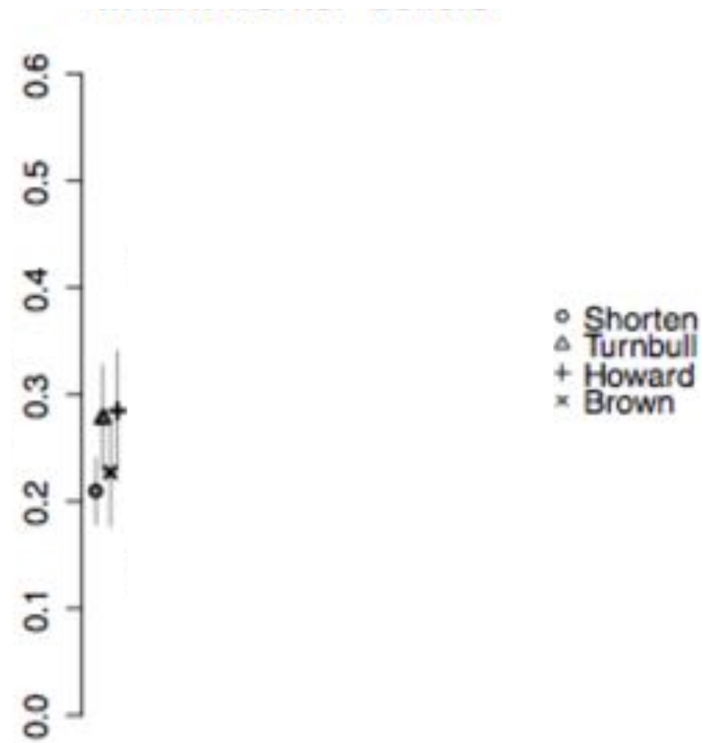

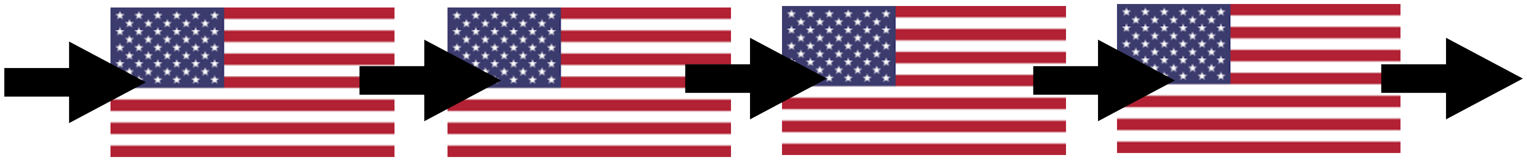

… and an all American iterated learning chain "reveals" a "preference" for *Gordon Brown* …

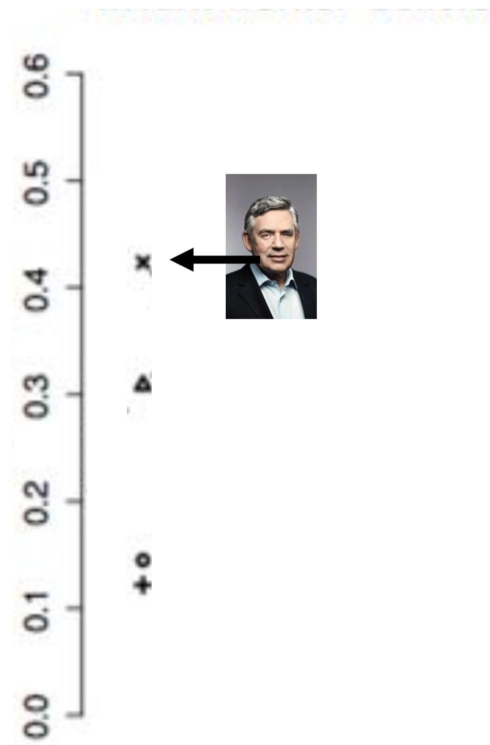

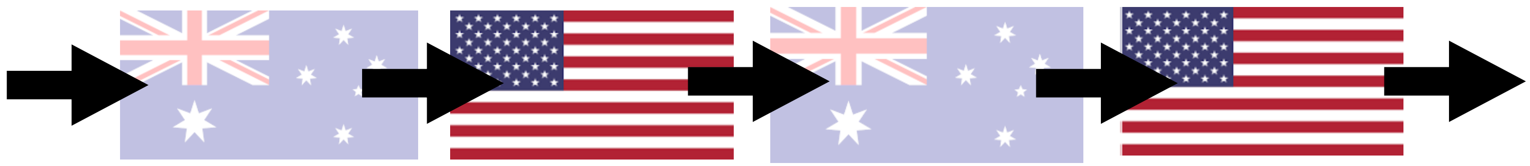

If we mix some Australians into the chain the Americans endorse *Malcolm Trunbull*

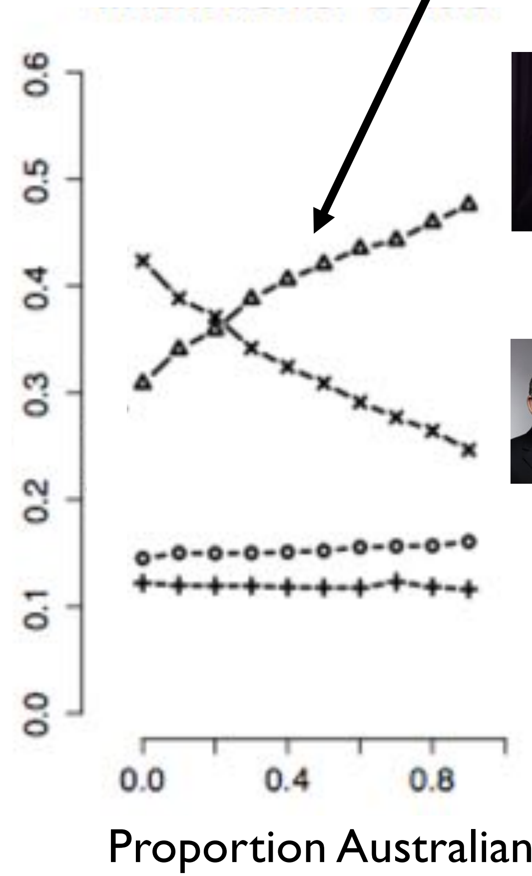


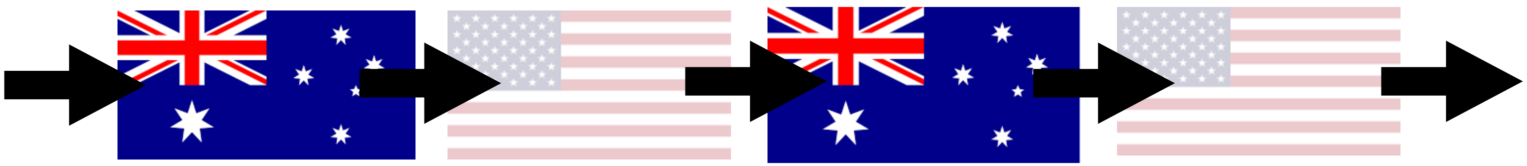

Australians choose Turnbull no matter how many Americans are included

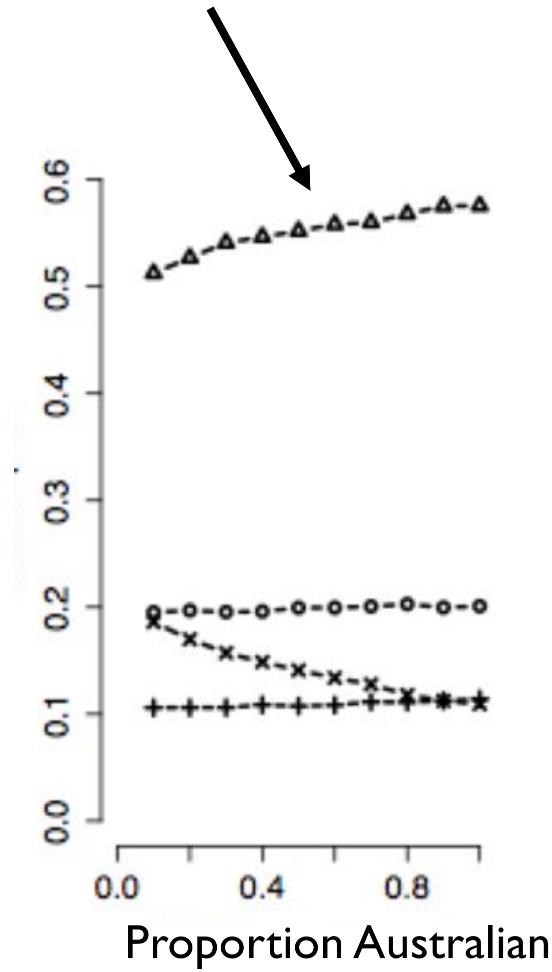

# Stronger biases/beliefs win?

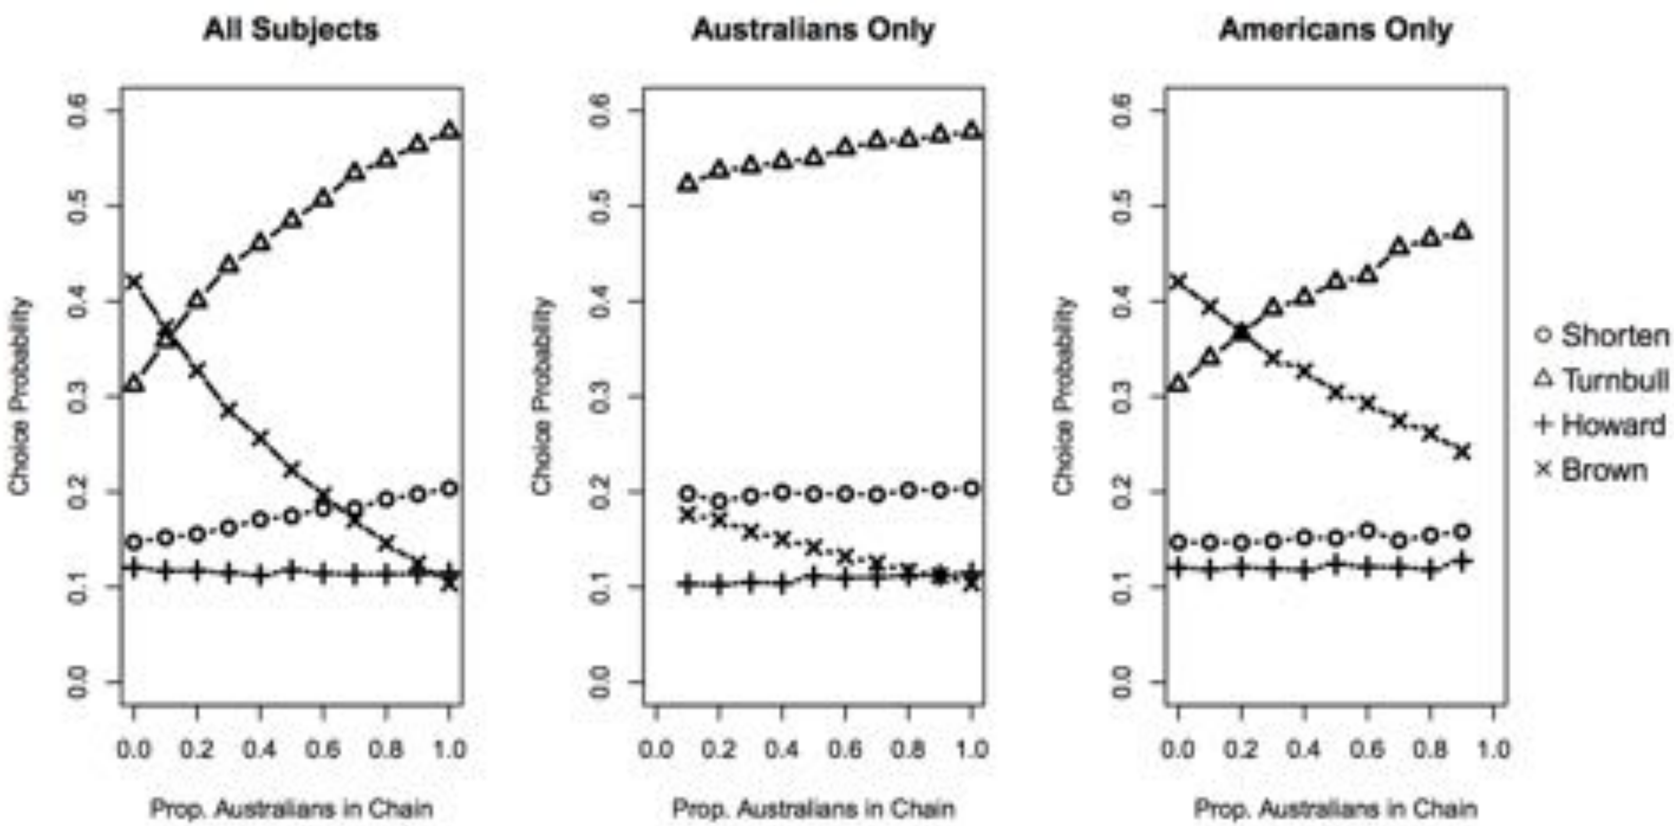

# Stronger biases/beliefs win

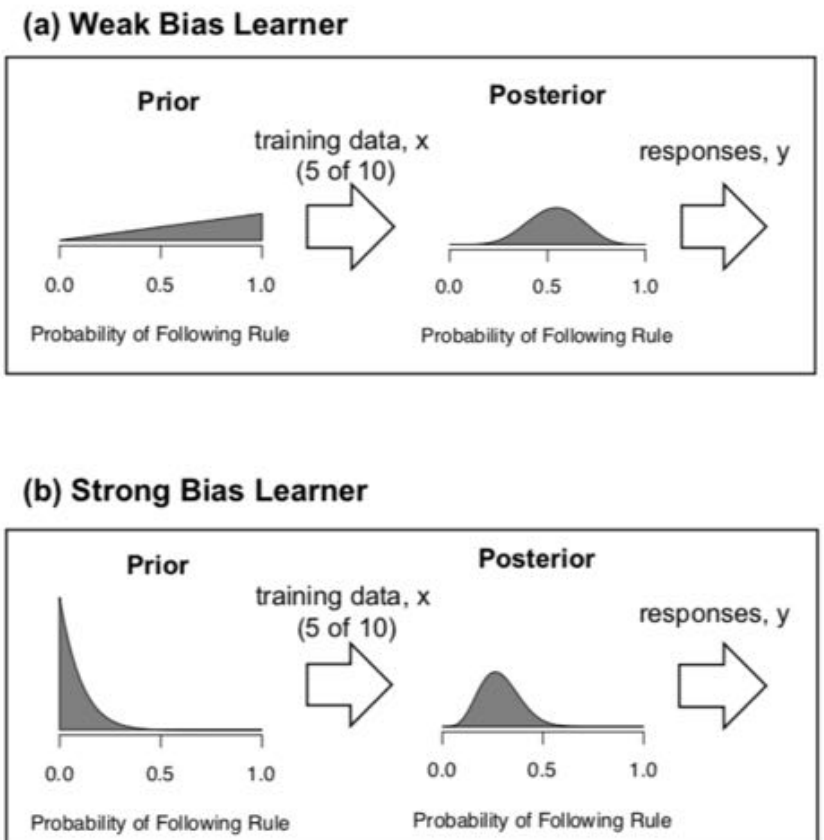

# Stronger biases/beliefs win

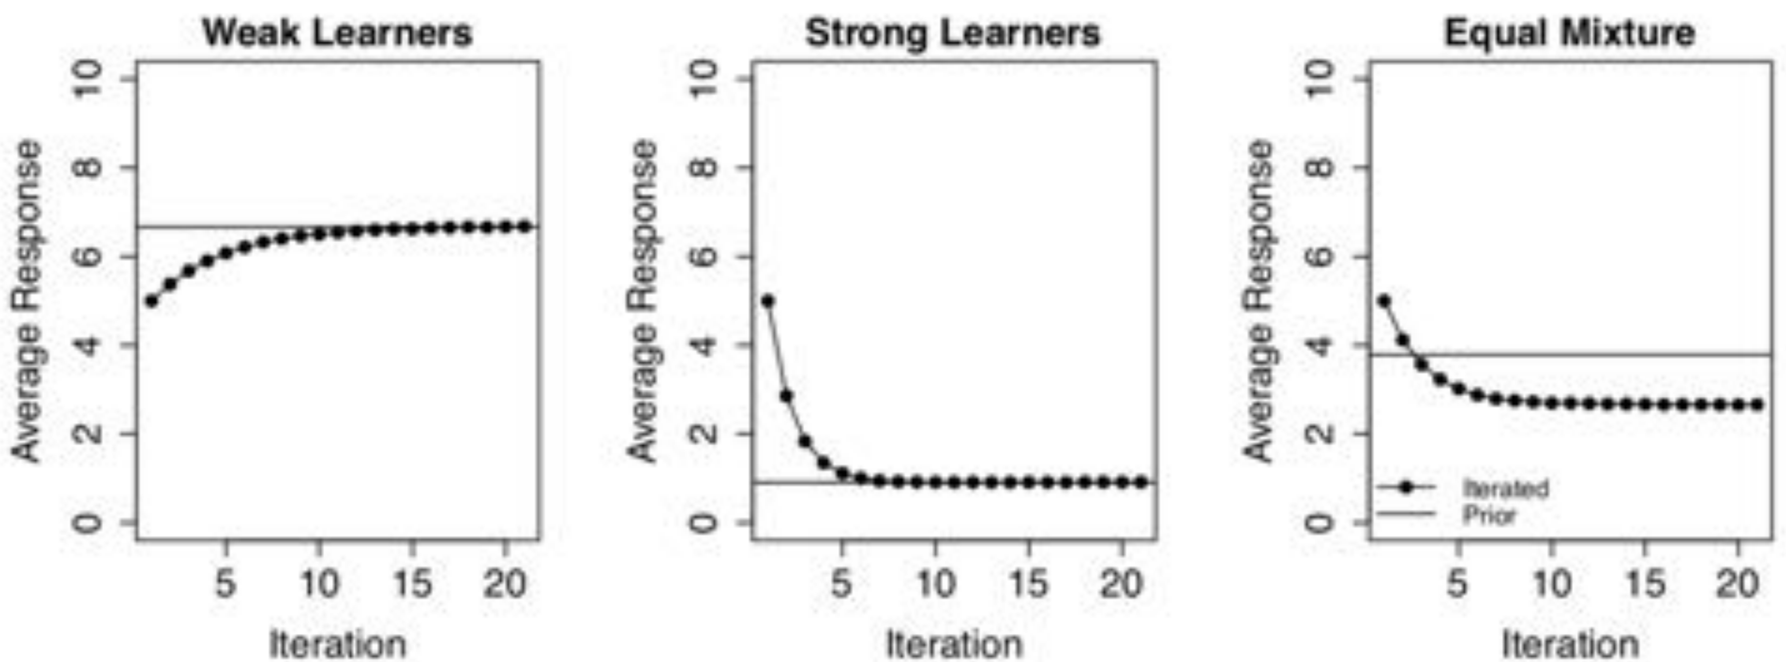

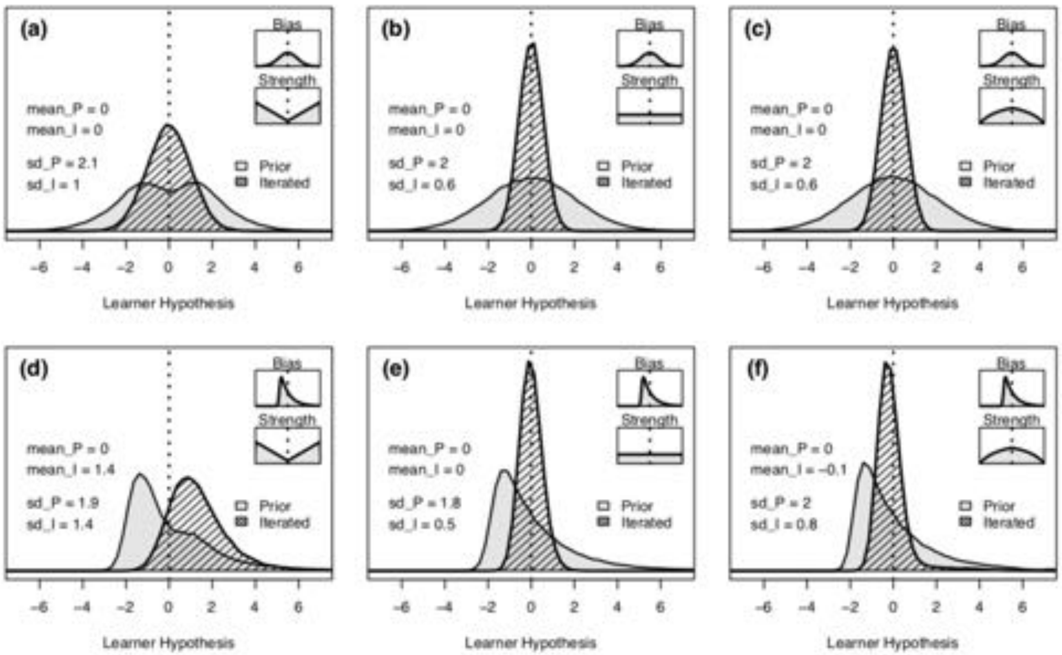


*Figure 12*. Systematic exploration of the behavior of heterogeneous iterated learning chains when individual differences are continuous. In the top row, individual differences are normally distributed (*symmetric* condition), whereas in the bottom row there is a positive skew (*asymmetric* condition). In the middle column, all learners have the same strength of bias no matter where their belief falls on the scale (*equal strength*), whereas on the left it is the learners on the tails of the distribution that have the strongest biases (*confident extremes*), and on the right the strongest biases occur in the middle of the distribution (*confident center*). In each panel, the solid plots depict the marginal *prior* distribution over the latent belief $b$ taken across the entire population, whereas the shaded plots show the corresponding marginal distribution over the beliefs $b$ elicited by an iterated learning procedure. The inset panels indicate which distribution condition (symmetric or asymmetric) and which bias condition (confident extremes, equal strength or confident center) is depicted, and the text annotations indicate the mean and standard deviation for each of the marginal distributions shown. It is evident that in all cases iterated learning distorts the underlying shape of population variation, and when the population is asymmetric and the extremes are more confident (panel d) it also converges to a different mean.

- The distortion depends on lots of factors
- One consistent pattern… if there are people with extreme beliefs and high confidence on “one side” but no corresponding group on the other side, iterated learning chains will favour those with extreme views
- Pretty hard to say how well this generalises to real life though

# Cumulative cultural evolution

# This is kind of depressing???

Does social transmission mean we just live in an echo chamber and all we get out are the biases we put in????

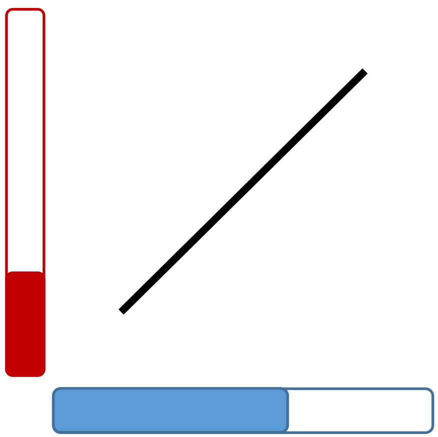

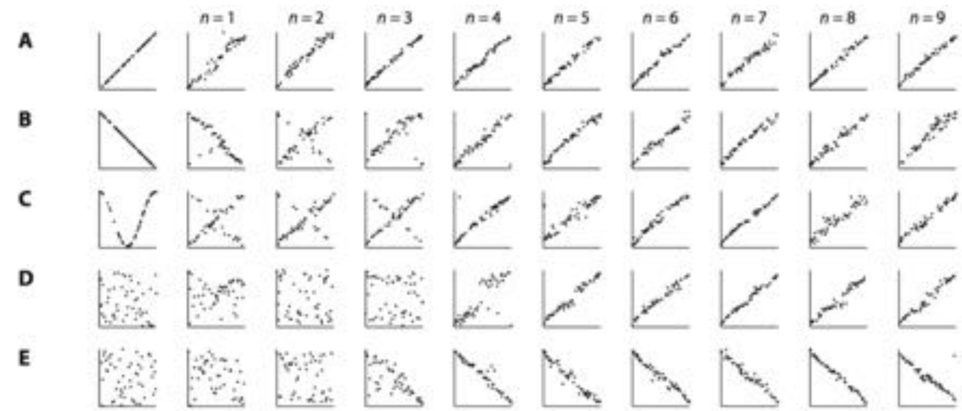


Or worse… one that amplifies the most extreme voices????

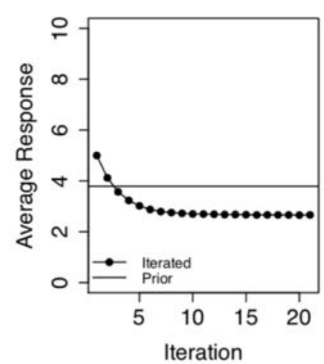

# Not necessarily!

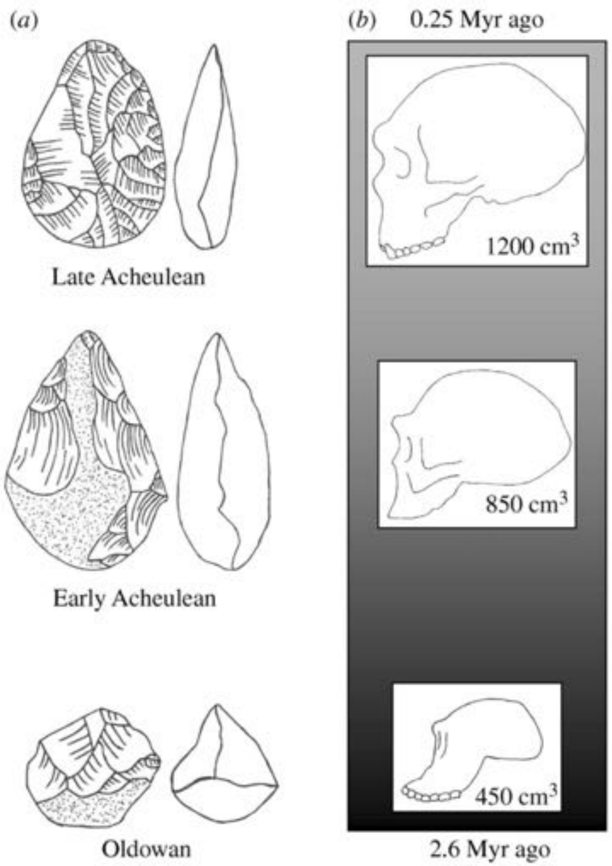


(Stout et al 2008)

- Cumulative cultural evolution *obviously* happens, the only question is how…
- Improvements in stone making technology *might* accompany biological evolution???
- But of course that hardly explains lasers…

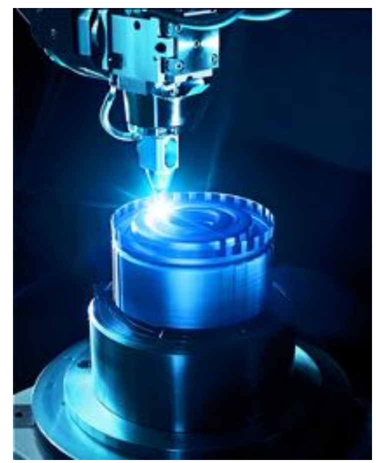

# Cumulative cultural evolution for social artefacts

(Fay et al 2018)

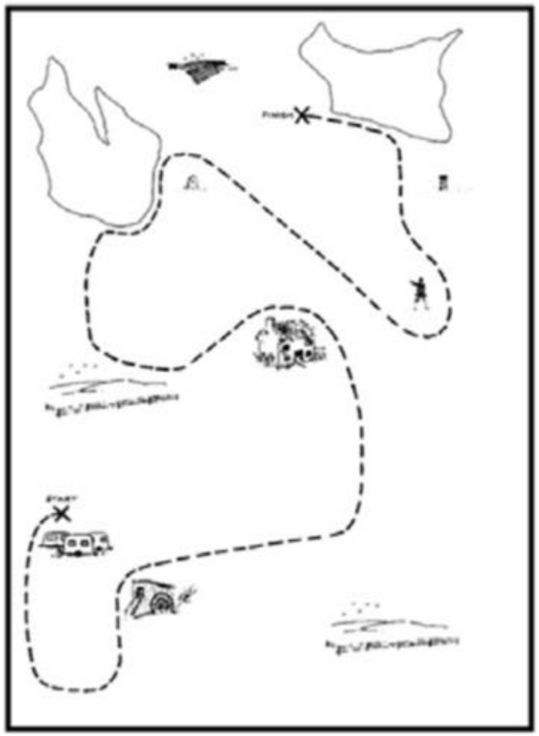

The "instruction giver" has a map with a route marked out on it

The "instruction follower" has the same map without the route

The instruction giver has to describe the path to be drawn via text messaging (in 10 mins)

(observation) The follower cannot send messages

(coordination) They can text back and forth

The follower becomes the new instruction giver for the next iteration using a new map and new route

# Cumulative cultural evolution for social artefacts

(Fay et al 2018)

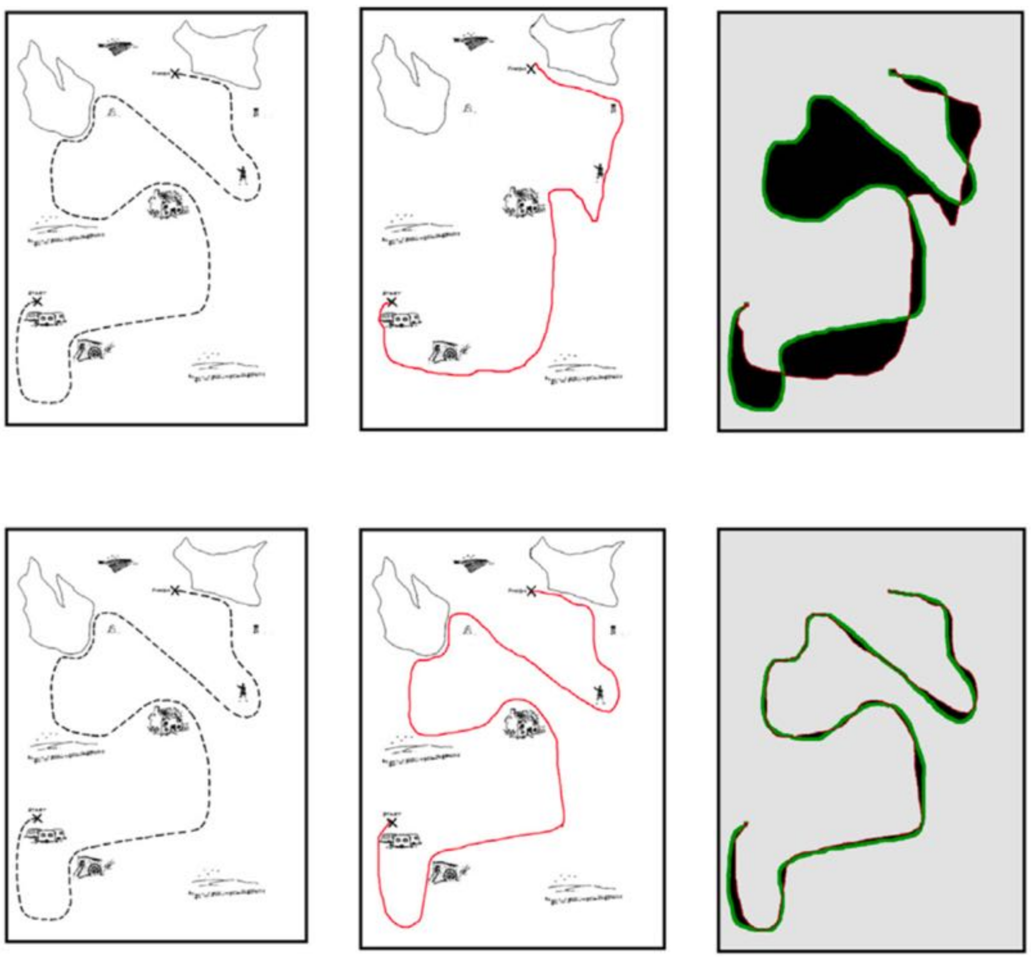

Reproduction accuracy is given by the size of the grey area

Low accuracy map means there are bigger deviations (black) between the two curves

High accuracy map means there are very few black pixels and many grey pixels

# Cumulative cultural evolution for social artefacts

(Fay et al 2018)

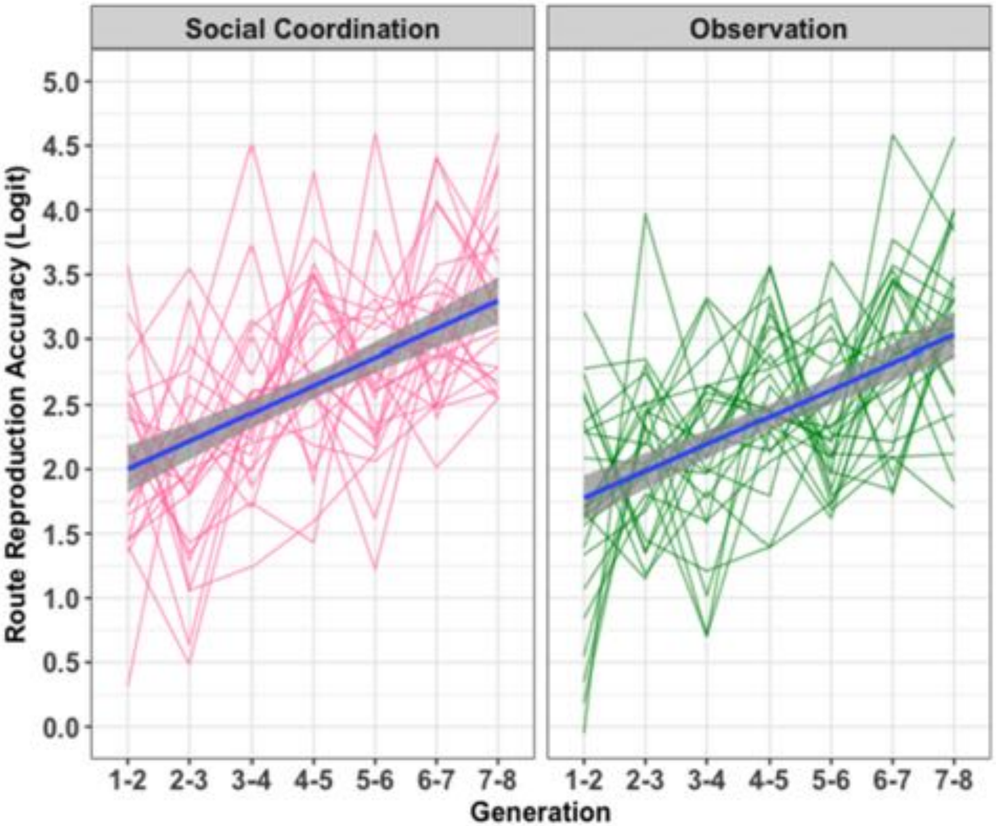


- 408 participants
- 51 x 8-person chains
- 25 coordination chains
- 26 observation chains

- Both conditions show improvement across generations
- Slightly higher fidelity in the social coordination condition
- These two factors account for about 30% of all variability
- But *<u>how</u>*????

# Cumulative cultural evolution for social artefacts

(Fay et al 2018)

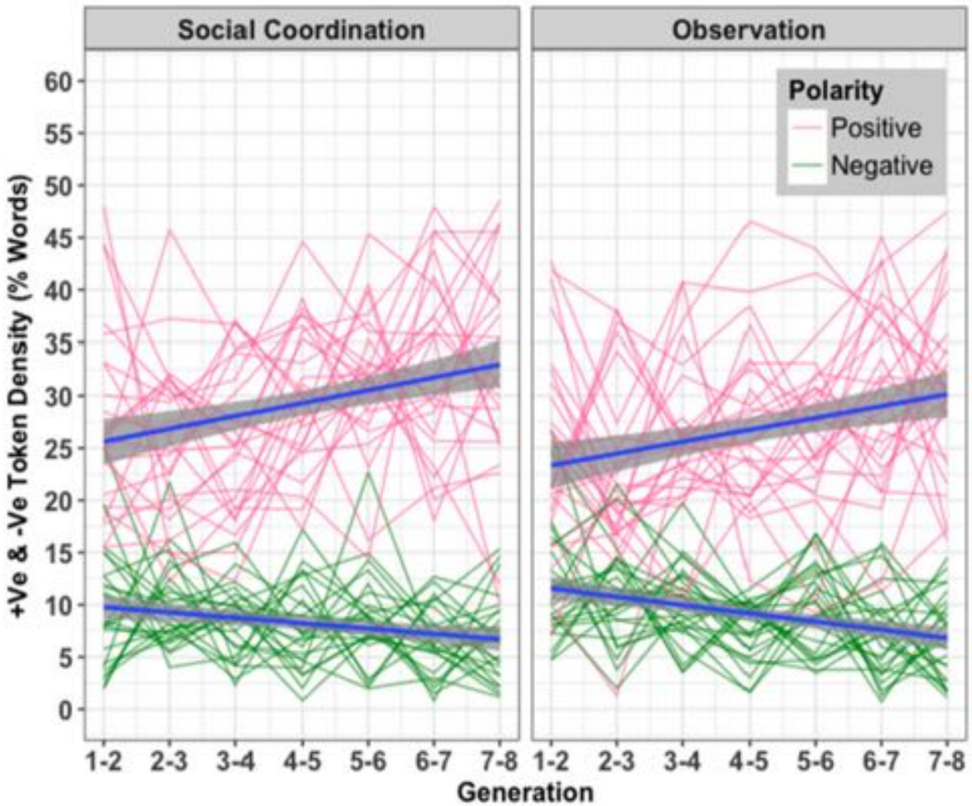


- Use more good words!

- Positive words are those that correlate with better solutions.
- The proportion of positive words increased across generations in both conditions

- Negative words are those that correlate with bad solutions
- The negative words decline across generations but not much

# Cumulative cultural evolution for social artefacts

(Fay et al 2018)

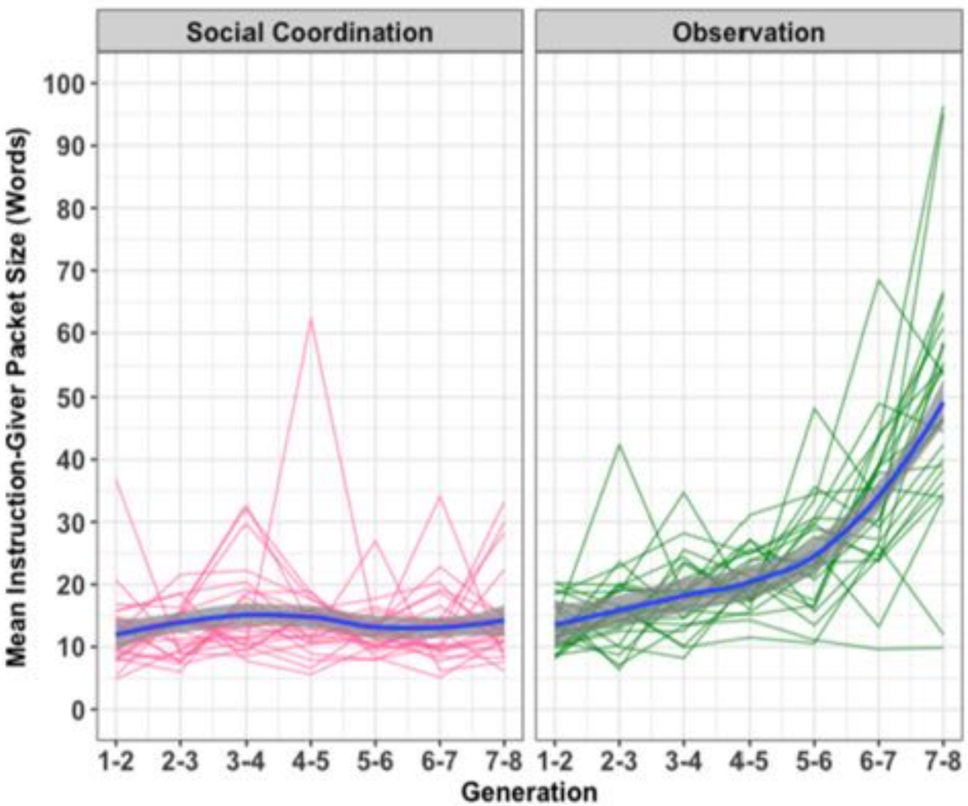


Exploratory analyses

- Curiously, they found that in the observation condition people learned to give "large packets"
- Likely explanation… when there's no actual social interaction, you learn it's more effective to draft the whole instruction beforehand?

Thanks