Information Choice and Motivated Beliefs

Spencer Y. Kwon, William E. Murdock III & Pierre-Luc Vautrey

February 22, 2019

Today

1. Introduction
2. Experimental Design
3. Preliminary Results
4. Challenges and Questions

Motivation

- Growing body of (mixed) evidence on the formation of overconfident beliefs through asymmetric updating
- While it's hard to tell whether updating biases are purely automatic or not, in real life people definitely have a lot of agency on the information they choose to receive and pay attention to
- Thus it seems crucial to study the role of these choices in forming motivated beliefs

An Example

- Think of someone who has to decide which news channel to listen to
- Fox News? CNN?
- Many people have a sense that these news sources are somewhat biased
- How do they choose?
- How to they subsequently form beliefs?
  - Lab evidence: selection neglect

An Example

- Think of someone who has to decide which news channel to listen to
- Fox News? CNN?
- Many people have a sense that these news sources are somewhat biased
- How do they choose?
- How to they subsequently form beliefs?
  - Lab evidence: selection neglect
- Hypothesis: do people exploit such choices to manufacture preferred beliefs?

An Example

- Think of someone who has to decide which news channel to listen to
- Fox News? CNN?
- Many people have a sense that these news sources are somewhat biased
- How do they choose?
- How to they subsequently form beliefs?
  - Lab evidence: selection neglect
- Hypothesis: do people exploit such choices to manufacture preferred beliefs?
  - Preview: in our experiment, they do!

Today

IQ Tests and Beliefs about Ranked Performance

- Participants answer questions commonly found in IQ tests
  Raven's Matrices
- For each IQ question, they are ranked in a comparison sample
  - Ranked lexicographically by number of incorrect attempts and then time
- After each IQ question, individuals are incentivized to report their probabilistic beliefs about whether they ranked above a fixed threshold

Belief Elicitation

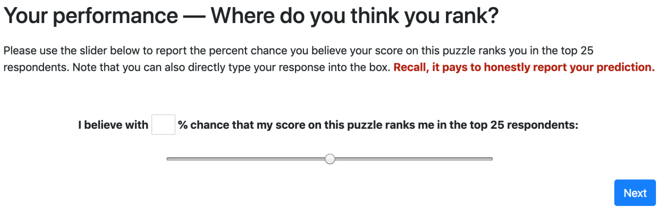

- Participants receive 3 binary signals, drawn from an urn with replacement, and re-report beliefs after each signal
  - Urn: 5 truthful signals and 3 lying signals
- Beliefs incentivized using the Lottery Method (Mobius et al. (2014), Coutts (2018)...)
- Narratives and animations used to help with comprehension

Gremlins | Receiving Signals (Before) | Receiving Signals (After)

Choice vs. No Choice Treatments

After a practice round, the urn is biased in a positive or negative manner for the remainder of the experiment.

- 2 signals are added to the urn that are always positive OR always negative.
- **Choice Treatment**: Participants choose which bias they want
- **No Choice Treatment**: Participants are randomly assigned to a bias

Gremlins | Comprehension Check | Choice Screen | Bayes Choice

Pilot Study

- Recruited 135 participants from MTurk
- Randomly assigned them to Choice and No Choice treatments

Sample

Summary Statistics by Treatment Arm

	Motivated Choice	Motivated No Choice	*p-val*
	(1)	(2)	(3)
Age	38.507	36.476	0.264
Female	0.423	0.419	0.971
White	0.806	0.778	0.694
College Degree	0.625	0.587	0.657
Income Less Than 50k	0.597	0.508	0.301
Democrat	0.451	0.476	0.770
Fox News	0.194	0.190	0.954
NYTimes	0.431	0.444	0.872
Observations	72	63	

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Raven Performance

Today

1. Introduction
2. Experimental Design
3. Preliminary Results
4. Challenges and Questions

Bias Choice

	Proportion
Yay Sayer	0.56
Nay Sayer	0.44
p-val	0.144
Num. obs.	72

P-value calculated using a one-sided test and the binomial CDF.

Monetary Gains Optimal Bayes

Updating in each Treatment

	Choice Yay	Choice Nay	Forced Yay	Forced Nay
Prior			0.90	0.90
			(0.04)	(0.04)
Yay-Yes			1.02	
			(0.16)	
Yay-No			0.64	
			(0.13)	
Nay-Yes				0.74
				(0.16)
Nay-No				1.09
				(0.13)
R^2	0.86	0.89	0.83	0.86
Num. obs.	588	456	568	316

Regression Specification Asymmetric Updating Bias Neglect

Updating in each Treatment

	Choice Yay	Choice Nay	Forced Yay	Forced Nay
Prior	0.89	0.94	0.90	0.90
	(0.02)	(0.02)	(0.04)	(0.04)
Yay-Yes	1.05		1.02	
	(0.14)		(0.16)	
Yay-No	0.69		0.64	
	(0.12)		(0.13)	
Nay-Yes		0.79		0.74
		(0.12)		(0.16)
Nay-No		0.69		1.09
		(0.17)		(0.13)
R^2	0.86	0.89	0.83	0.86
Num. obs.	588	456	568	316

Regression Specification Asymmetric Updating Bias Neglect

Biased Beliefs

	Mean Bias		
	Choice	Forced	Pr(Choice=Forced)
Yay	0.22	0.21	0.49
Nay	0.14	−0.15	0.01
Pooled	0.18	0.05	0.12

Today

1. Introduction
2. Experimental Design
3. Preliminary Results
4. **Challenges and Questions**

Challenges and Questions

- Non-motivated Control Arm
- Meta-cognition
- Strategy method for belief?
- Real-world extensions

Raven's Matrices

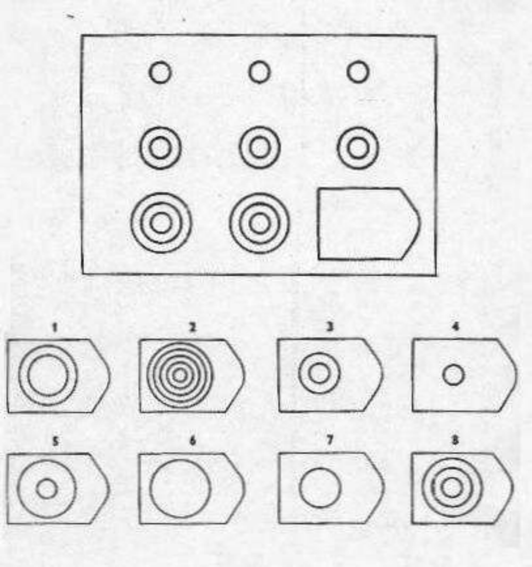

Back

Peformance on Raven's Matrices

Matrix	# Incorrect	Time (Sec.)	Payment ($)
Prac. Easy	0.246	13.093	0.446
Prac. Hard	0.875	27.206	0.338
34	1.153	21.534	0.331
45	0.912	20.245	0.376
47	2.277	34.975	0.215
50	0.956	20.519	0.367
55	1.781	31.216	0.267
59	4.307	47.458	0.102

PE PH M34 M45 M47 M50 M55 M59

Back to Sample

Practice Matrix (Easy)

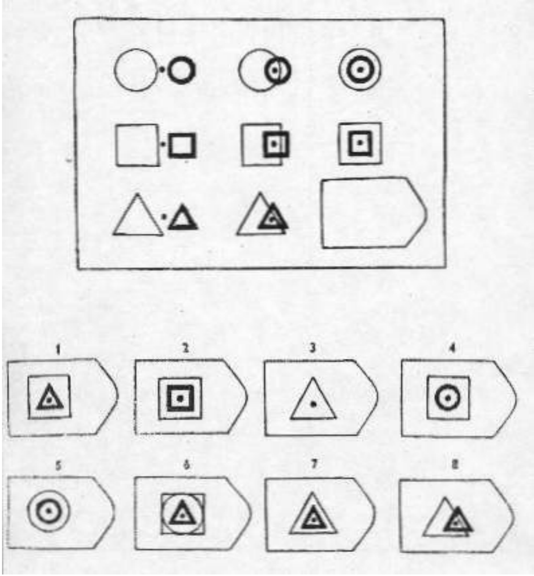

Back to table

Practice Matrix (Hard)

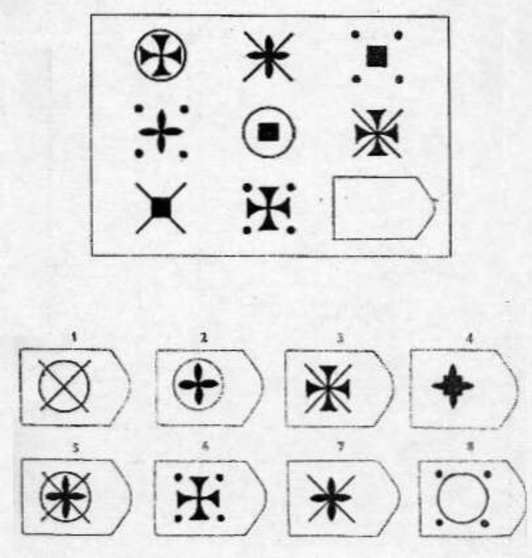

Back to table

Gremlins

	**Truth Teller:** This type of gremlin will look at all 100 responses and will always honestly report to you whether you are ranked in the top 25 respondents or not.
	**Liar:** This type of gremlin will look at all 100 responses and will always report to you the opposite of what they observe. For example, if you are truly ranked in the top 25 respondents, this type of gremlin would tell you that you are not ranked in the top 25.
	**Yay Sayer:** This type of gremlin will always answer "yes", regardless of whether it is true or not. They are lazy and positive — they won't even look at the data and just tell you that you are ranked in the top 25 respondents.
	**Nay Sayer:** This type of gremlin will always answer "no", regardless of whether it is true or not. They are lazy and negative — they won't even look at the data and just tell you that you are not ranked in the top 25 respondents.

Back (Beliefs) Back (Choice/Forced)

Receiving Signals (Before)

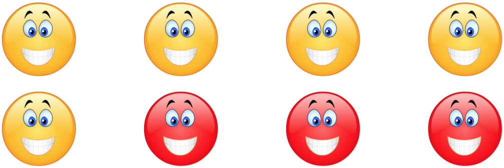

Talk with a Gremlin

Gremlin 1 Gremlin 2 Gremlin 3

I now believe with ☐ % chance that my score on this puzzle ranks me in the top 25 respondents:

Back

Receiving Signals (After)

Gremlin 1 Gremlin 2 Gremlin 3

I now believe with ☐ % chance that my score on this puzzle ranks me in the top 25 respondents:

Back

Raven Matrix 34

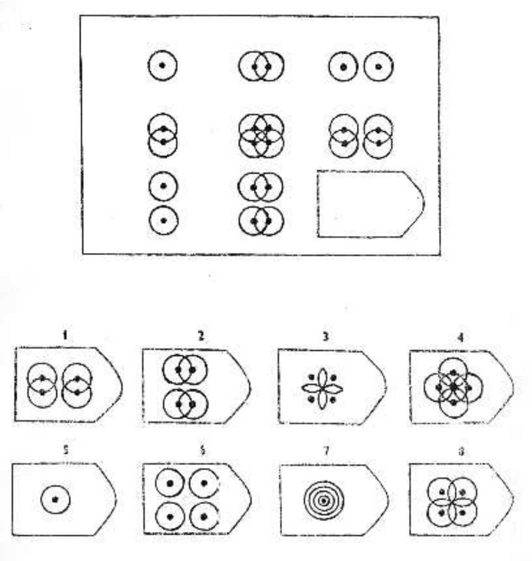

Back to table

Raven Matrix 45

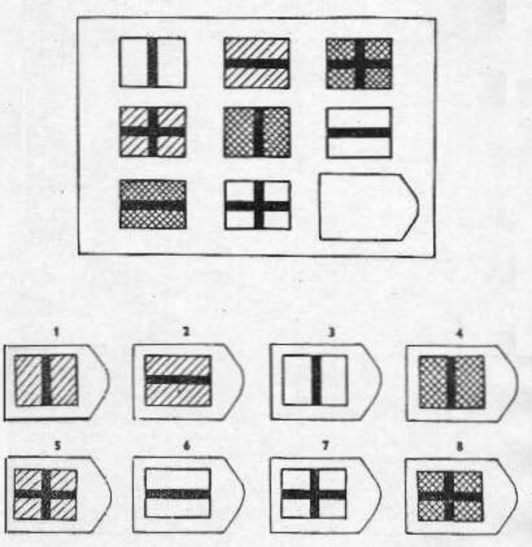

Back to table

Raven Matrix 47

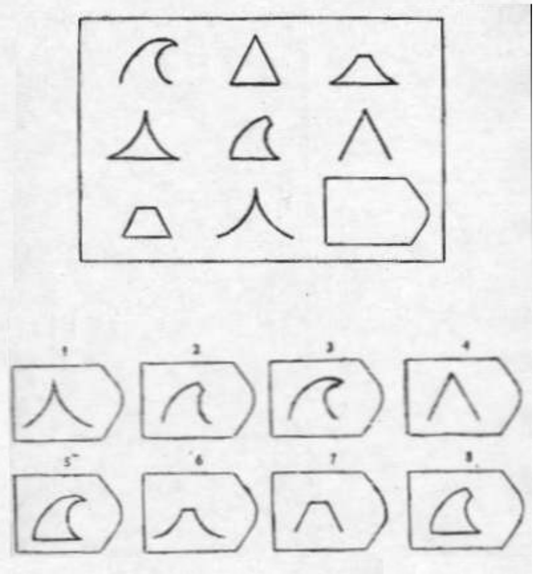

Back to table

Raven Matrix 50

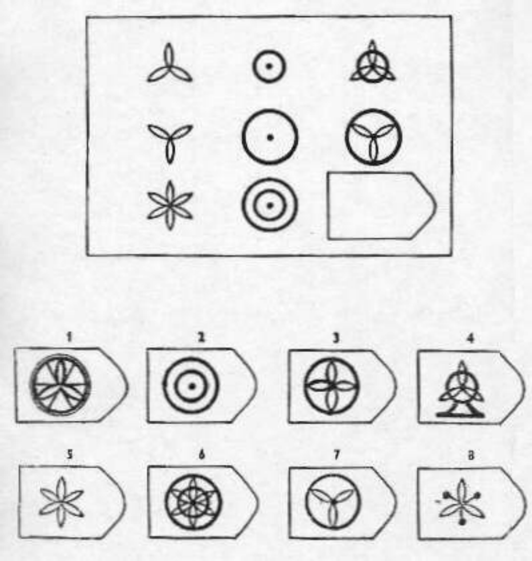

Back to table

Raven Matrix 55

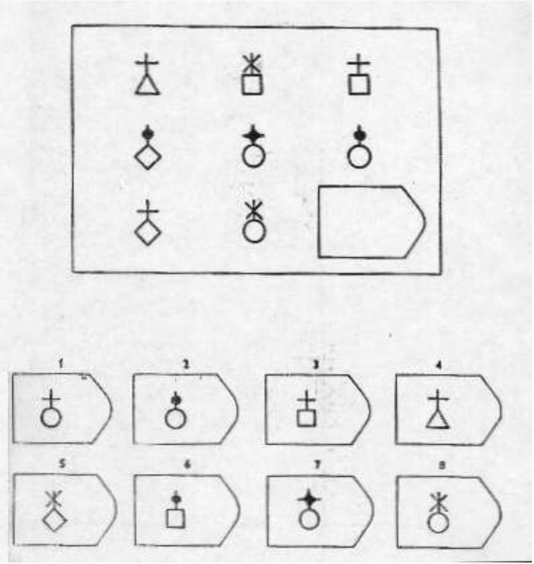

Back to table

Raven Matrix 59

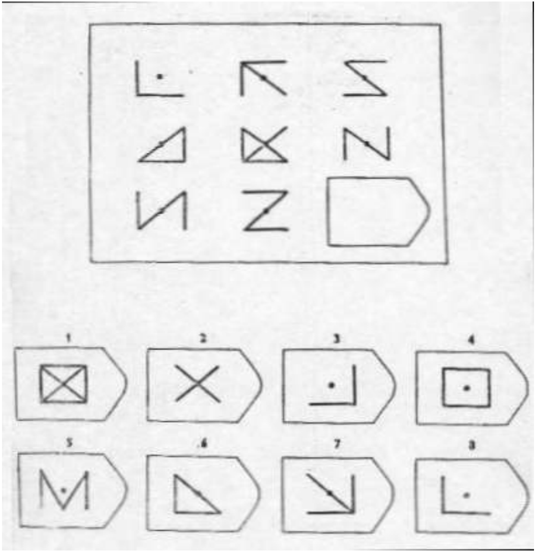

Back to table

Gremlin Comprehension Check

Gremlin Example

Suppose your performance on the puzzle ranks you 80 out of 100. What would each type gremlin report to you after looking at the 100 responses? Remember, they are answering the question "are you ranked in the top 25 respondents?"

You must submit the correct answer to proceed.

	**Truth Teller:** No
	**Liar:** Yes
	**Nay Sayer:** No
	**Yay Sayer:** Yes

Check Answer

That's correct!

Back

Bias Choice Screen

Would you prefer the other 2 gremlins to be always negative Nay Sayers or always positive Yay Sayers?

Nay Sayers

OR

Yay Sayers

Please click on the type of gremlin you would like for the other 2 gremlins and then click the "Next" button.

Back

What a Bayesian Would Do

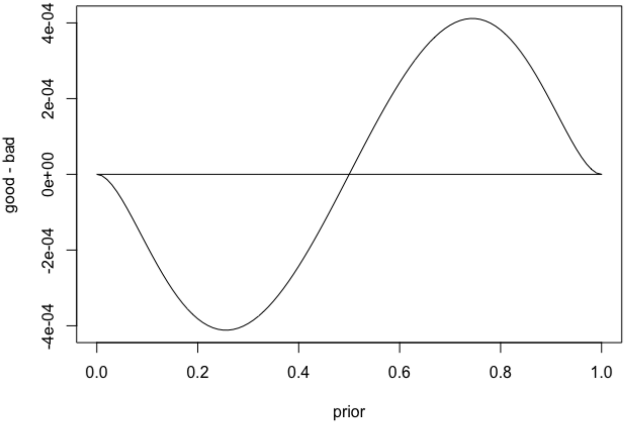

Back

Choose to Maximize Expected Monetary Gains?

	Chose Yay
Intercept	0.13
	(0.15)
Monetary Edge	467.24
	(503.02)
Num. obs.	72

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Back

Optimal Bayesian Choice?

	Optimal Bayes	
	Yay-Sayer	Nay-Sayer
Yay Sayer	23	14
Nay Sayer	15	12

Notes: Participants who reported a prior belief of 50 were dropped.

Back

Updating Specification

We use similar specifications as in Mobius et al. (2014) and Coutts (2018):

- In sequential updating problems, Bayes' rule can be written as

$$\frac{\mu_t}{1-\mu_t} = \frac{\mu_{t-1}}{1-\mu_{t-1}} \cdot LR_k$$

  where μ_t is the posterior, μ_{t-1} is the prior, and LR_k is the likelihood ratio of observing signal $s_t = k \in \{0, 1\}$. In our case, a signal of 0 corresponds to a gremlin saying "No" and a signal of 1 corresponds to a gremlin saying "Yes".

- Taking logs, this motivates the following regression:

$$\ln\left(\frac{\mu_{it}}{1-\mu_{it}}\right) = \delta \ln\left(\frac{\mu_{i,t-1}}{1-\mu_{i,t-1}}\right) + \beta_1 I(s_{it}=1)\ln(LR_1) + \beta_0 I(s_{it}=0)\ln(LR_0) + \epsilon_{it}$$

- Where, for a Bayesian, $\delta = \beta_1 = \beta_0 = 1$.

Back

Asymmetric Updating

	Pooled	Choice	Forced (Balanced)
Prior	0.91	0.92	0.91
	(0.01)	(0.01)	(0.02)
Good Signal	0.90	0.91	0.84
	(0.05)	(0.06)	(0.09)
Bad Signal	0.72	0.69	0.85
	(0.05)	(0.07)	(0.10)
R^2	0.86	0.87	0.85
Num. obs.	1928	1044	632
Pr(Good = Bad)	0.011	0.014	0.941

Back

Bias Neglect

	Pooled
Prior	0.91
	(0.01)
Congruent Signal	0.97
	(0.06)
Incongruent Signal	0.71
	(0.05)
R^2	0.86
Num. obs.	1928
Pr(Con. = Incon.)	0.000

Back