EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Use and interpretation of probability expressions under higher-order uncertainty

Michele Herbstritt, Michael Franke

XPrag.de Workshop
Rationality, Probability, and Pragmatics
Berlin, May 25, 2016

Introduction: WHY probability expressions?

- Example: *The Litvinenko Inquiry*. Delivered by Sir Robert Owen to the Home Secretary of the United Kingdom, January 19, 2016.

- In 329 pages:

  84 occurrences of *probable/probably/likely*

  103 occurrences of *might/possible/possibly*

Introduction: previous literature

- Linguistics: gigantic body of research about modality. More recently about probability expressions too (Swanson, 2006, Yalcin, 2010, Lassiter, 2011, Moss, 2015).

  Probability expressions convey that the probability of a proposition is bigger than a certain value.

- Psychology: studies about probability expressions (Beyth-Marom, 1982, Teigen, 1988, Windschitl and Wells, 1998).

  Main focus: first-order uncertainty.

Introduction: higher-order uncertainty

- Knowing the objective chance is first-order uncertainty.

Introduction: higher-order uncertainty

- Knowing the objective chance is first-order uncertainty.
- What if the objective chance is not known? What if an agent has subjective uncertainty about the objective chance?
- Real-life example: talking about the weather.

Introduction: higher-order uncertainty

- Our setting: 100 red and blue balls in an urn, with partial access.

  Bob draws 8 balls and observes that 5 are red.
  Carol draws 80 balls and observes that 50 are red.

Introduction: higher-order uncertainty

- Our setting: 100 red and blue balls in an urn, with partial access.

  Bob draws 8 balls and observes that 5 are red.
  Carol draws 80 balls and observes that 50 are red.

  (1) A randomly drawn ball might be red.

  (2) A randomly drawn ball will probably be red.

Introduction: higher-order uncertainty

- Our setting: 100 red and blue balls in an urn, with partial access.

  Bob draws 8 balls and observes that 5 are red.
  Carol draws 80 balls and observes that 50 are red.

  (1) A randomly drawn ball might be red.

  (2) A randomly drawn ball will probably be red.

- Bottom line: it seems that best guess about objective chance is not all that matters.

Experiment 1: use of probability expressions

- Goal: explore participants' choices of probability expressions under higher-order uncertainty.

Experiment 1: use of probability expressions

- Goal: explore participants' choices of probability expressions under higher-order uncertainty.
- Expressions: *possibly*, *probably (not)*, *certainly (not)*.

Experiment 1: use of probability expressions

- Goal: explore participants' choices of probability expressions under higher-order uncertainty.
- Expressions: *possibly*, *probably (not)*, *certainly (not)*.
- Scenario: two-player cooperative game "guess the content of the urn"

Expression trial

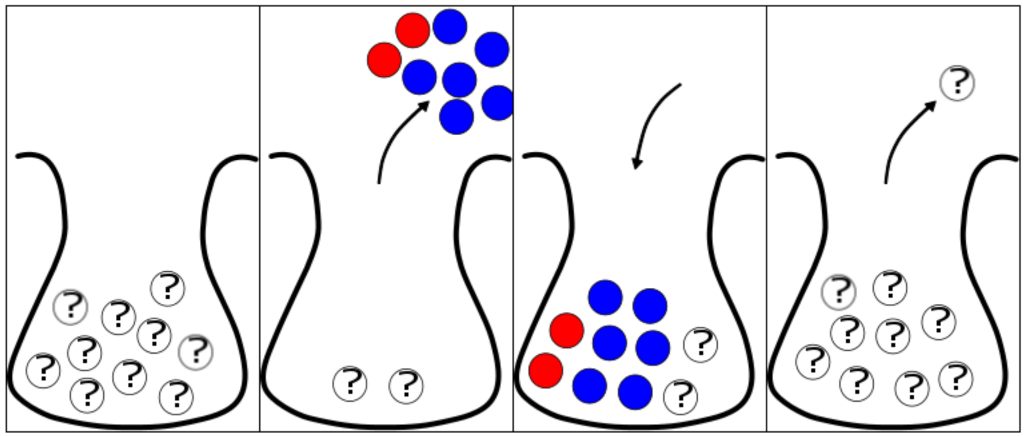

You draw 8 balls and observe that 2 of them are red.

Which message do you send?

The next ball will [▾] be red.

- certainly not
- probably not
- possibly
- probably
- certainly

25 %

Likelihood trial

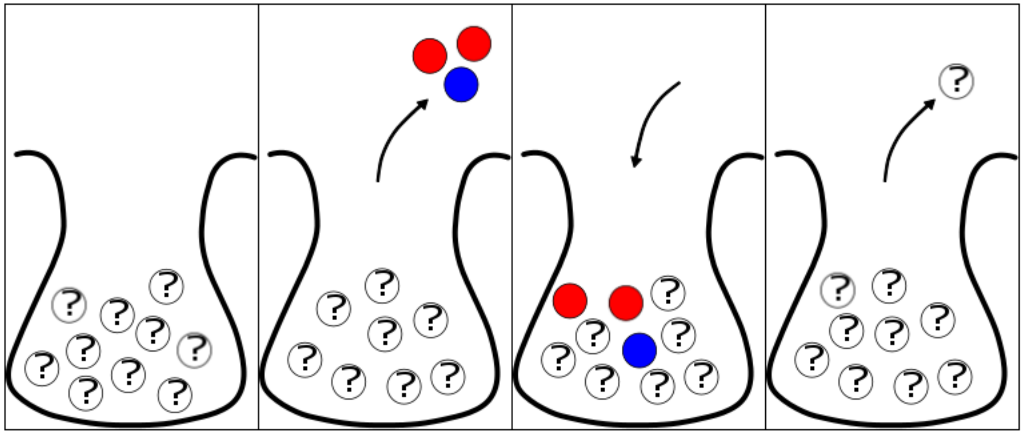

You draw 3 balls and observe that 2 of them are red.

Which message do you send?

It's approximately [▾] % likely that the next ball will be red.

Next

8 % completed.

Experiment 1: design

- 14 conditions: 7 proportions x 2 uncertainty levels

	0	0.25	0.33	0.5	0.67	0.75	1
high	0/2	1/4	1/3	2/4	2/3	3/4	2/2
low	0/10	2/8	3/9	4/8	6/9	6/8	10/10

Experiment 1: design

- 14 conditions: 7 proportions x 2 uncertainty levels

	0	0.25	0.33	0.5	0.67	0.75	1
high	0/2	1/4	1/3	2/4	2/3	3/4	2/2
low	0/10	2/8	3/9	4/8	6/9	6/8	10/10

- 50 self-reported English native speakers on MTurk.
- Two kinds of trials: expression and likelihood.
- Each participant completed 6 trials for each kind, seeing 12 out of 14 conditions, in random order.

It's approximately . . . % likely that the next ball will be red

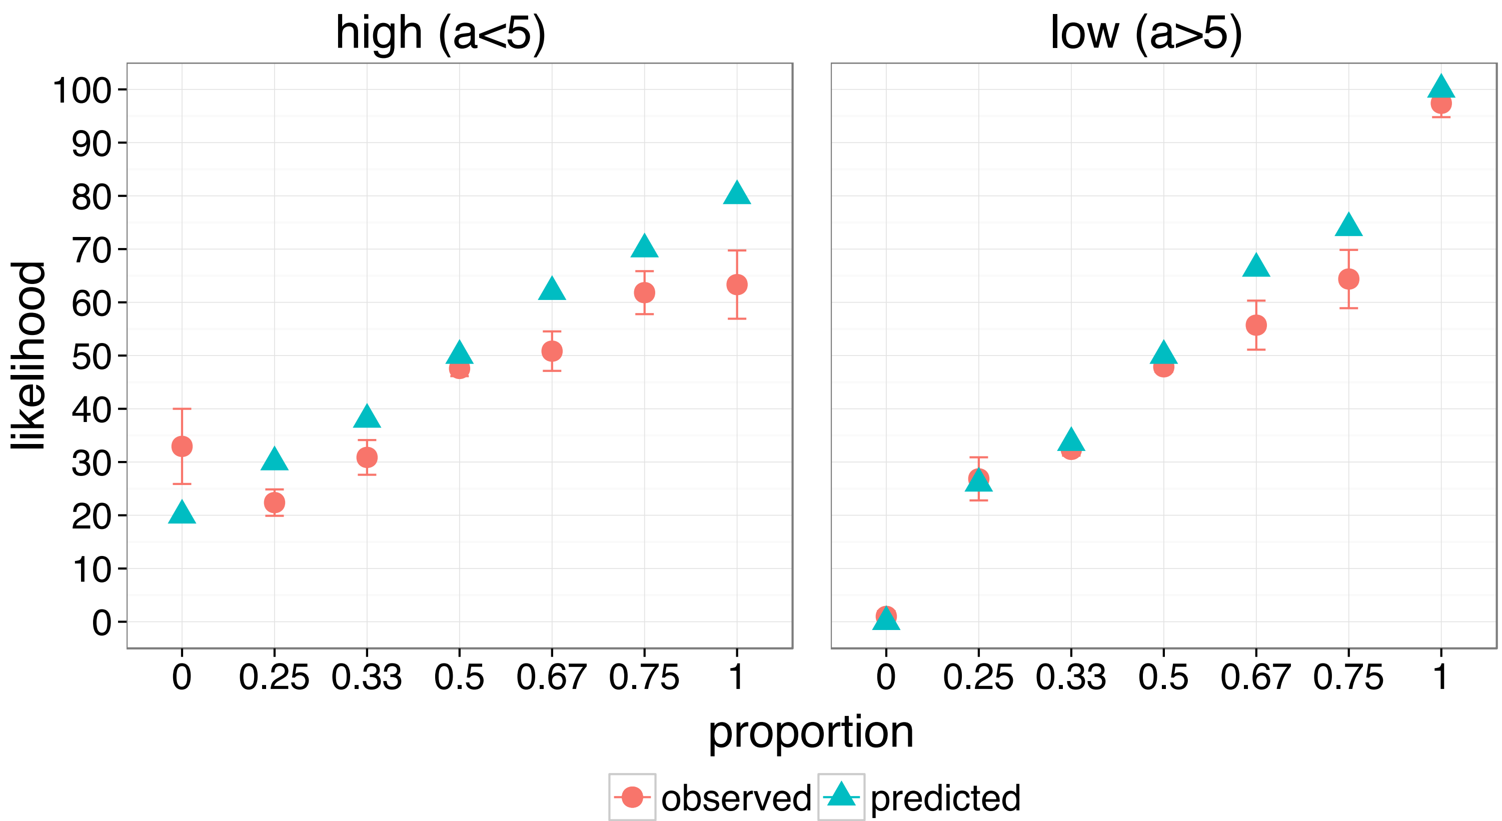

(predictions are expected values under ideal belief given observation)

The next ball will . . . be red

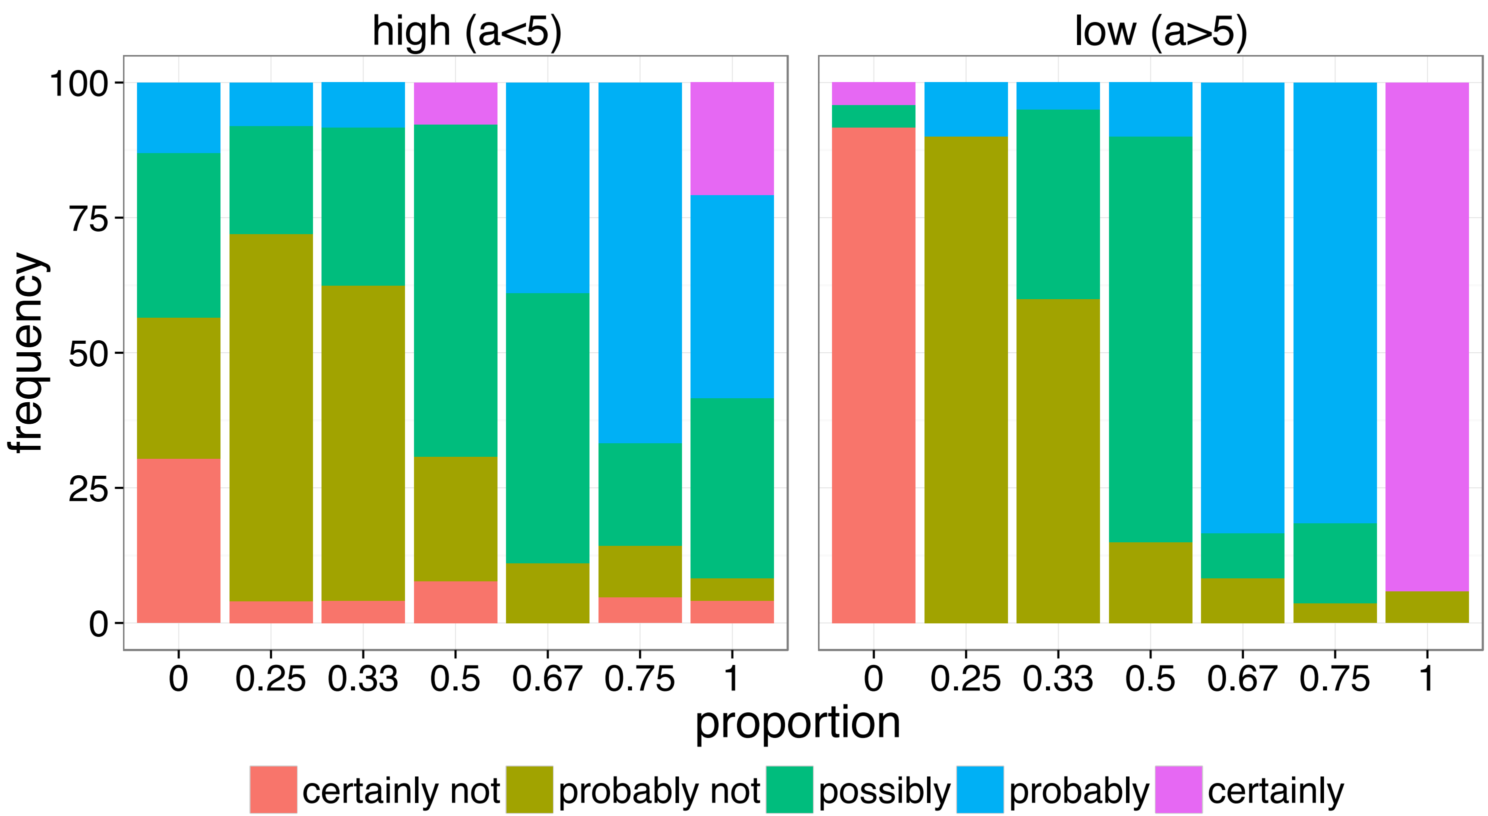

Experiment 1: analysis

- Question: are participants' choices governed only by their guess about the likelihood or does higher-order uncertainty play a role?

Experiment 1: analysis

- Question: are participants' choices governed only by their guess about the likelihood or does higher-order uncertainty play a role?

- We compared three multinomial logistic regression models:

  simple: `answer~belief`
  complex: `answer~access+observation`
  interaction: `answer~access*observation`

Experiment 1: analysis

- Question: are participants' choices governed only by their guess about the likelihood or does higher-order uncertainty play a role?

- We compared three multinomial logistic regression models:

  simple: `answer~belief`
  complex: `answer~access+observation`
  interaction: `answer~access*observation`

- AIC scores:

	simple	*complex*	*interaction*
AIC	662.78	646.69	635.93

Experiment 1: analysis

- Question: are participants' choices governed only by their guess about the likelihood or does higher-order uncertainty play a role?
- We compared three multinomial logistic regression models:

  simple: `answer~belief`
  complex: `answer~access+observation`
  interaction: `answer~access*observation`

- AIC scores:

	simple	*complex*	*interaction*
AIC	662.78	646.69	635.93

- Question: what exactly is the role of access and observation?

Model: basics

- Based on RSA by Goodman and Stuhlmüller (2013).
- State space $S = \{0, 1, \ldots, 10\}$, # of red balls in the urn.
- For any $s \in S$, $s/10$ is objective chance of a ball to be red.
- The speaker doesn't know s, draws a balls (access) and observes that o are red (observation).

Model: uncertain belief

- Based on the observation, the speaker forms uncertain rational belief about the content of the urn, ie distribution over S given o, a:

$$\text{speaker.bel}(s|o, a) \propto P(o|a, s) * \text{prior}(s) \tag{1}$$

Model: literal meaning and literal listener

- Simple threshold semantics of messages, θ is free:

$$[\![\text{certainly}(p)]\!]_s = 1 \text{ iff } P(p) = 1 \text{ in } s$$

$$[\![\text{probably}(p)]\!]_s = 1 \text{ iff } P(p) > \theta \text{ in } s$$

$$[\![\text{possibly}(p)]\!]_s = 1 \text{ iff } P(p) > 0 \text{ in } s$$

Model: literal meaning and literal listener

- Simple threshold semantics of messages, θ is free:

$$[\![\text{certainly}(p)]\!]_s = 1 \text{ iff } P(p) = 1 \text{ in } s$$

$$[\![\text{probably}(p)]\!]_s = 1 \text{ iff } P(p) > \theta \text{ in } s$$

$$[\![\text{possibly}(p)]\!]_s = 1 \text{ iff } P(p) > 0 \text{ in } s$$

- Literal listener's belief:

$$\text{literal.bel}(s|m) = \mathcal{U}(s|[\![m]\!]_s = 1) * \text{prior}(s) \tag{2}$$

Model: speaker's utility and behavior

- Communicative goal: maximize the information transferred to the listener.
- Expected utility of a message is negative Hellinger Distance between speaker's belief and literal listener's belief:

$$\text{EU}(m; o, a) = -\text{HD}[\text{speaker.bel}(s|o, a),\ \text{literal.bel}(s|m)] \quad (3)$$

Model: speaker's utility and behavior

- Communicative goal: maximize the information transferred to the listener.
- Expected utility of a message is negative Hellinger Distance between speaker's belief and literal listener's belief:

$$\text{EU}(m; o, a) = -\text{HD}[\text{speaker.bel}(s|o, a),\ \text{literal.bel}(s|m)] \quad (3)$$

- Speaker's behavior is soft-max of EU:

$$\text{speaker.prob}(m|o, a) \propto exp(\lambda * \text{EU}(m; o, a)) \quad (4)$$

Model & Data: parameter estimation

- Mean values of *M*'s parameters, with HDIs:

	θ	α	λ
mean	0.55	2.84	4.32
HDI	0.50-0.60	1.25-4.88	3.66-4.99

- Why are these interesting?

Model & Data: parameter estimation

- Mean values of *M*'s parameters, with HDIs:

	θ	α	λ
mean	0.55	2.84	4.32
HDI	0.50-0.60	1.25-4.88	3.66-4.99

- $\theta = 0.55$ is intuitive, usually assumed in literature!

Model & Data: how good is the model?

- In terms of AIC, the model is much better than the best regression model:

	interaction	*M*
AIC	635.93	305.57

- Pearson's product-moment correlation between data and predictions of M fitted with inferred parameters:

$$r = 0.89\text{; 95\% ci:0.83-0.93; } p < 0.001$$

Model: listener's inference

- Joint inference of access, observation, state:

$$\text{listener.prob}(s, o, a|m) \propto \text{speaker.prob}(m|o, a) * \text{priors} \qquad (5)$$

Experiment 2: interpretation of probability expressions

- Goal: explore participants' interpretation of probability expressions along two dimensions: infer the state of the world, infer the uncertainty state of the speaker.

Experiment 2: interpretation of probability expressions

- Goal: explore participants' interpretation of probability expressions along two dimensions: infer the state of the world, infer the uncertainty state of the speaker.
- 5 conditions, one for each message *certainly not*, *probably not*, *possibly*, *probably*, *certainly*.
- Two kinds of trials: guess-the-state and guess-the-observation.

Experiment 2: interpretation of probability expressions

- Goal: explore participants' interpretation of probability expressions along two dimensions: infer the state of the world, infer the uncertainty state of the speaker.
- 5 conditions, one for each message *certainly not*, *probably not*, *possibly*, *probably*, *certainly*.
- Two kinds of trials: guess-the-state and guess-the-observation.
- 109 self-reported English native speakers on MTurk.
- Each participant completed 5 trials for each kind, seeing all messages twice, in random order.

EBERHARD KARLS UNIVERSITÄT TÜBINGEN

Guess-the-state trial

Another player sent the following message:

The next ball will probably be red.

How many **red** balls do you think there are **in the urn**?

0 7 10

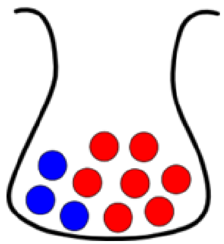

(please adjust the slider in the way that best corresponds to your intuition)

Next

10 % completed.

Guess-the-observation trial

Another player sent the following message:

The next ball will possibly be red.

How many balls do you think **the sender has drawn**?

0 8 10

And how many **of them** do you think were **red**?

0 2 10

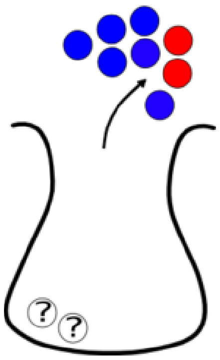

(please adjust the sliders in the way that best corresponds to your intuition)

Next

0 % completed.

There are . . . red balls in the urn

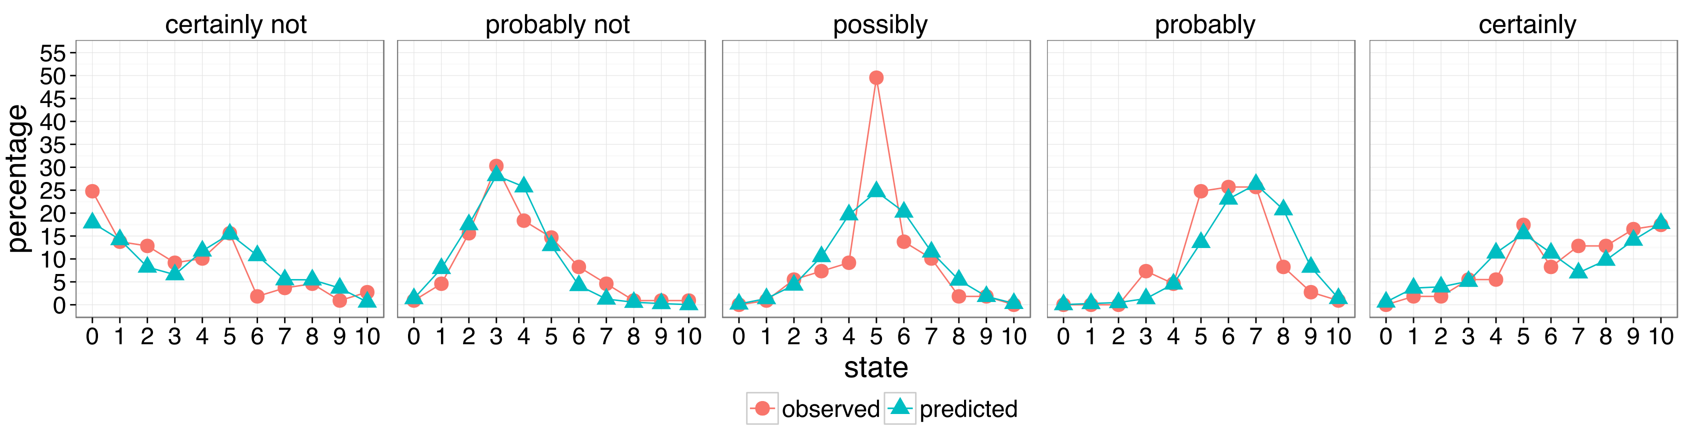

$$r = 0.84;\ 95\%\ \text{ci:}0.73\text{-}0.90;\ p < 0.001$$

The speaker has drawn . . . balls, and . . . of them were red

- For each message, participants had $11 + 10 + 9 + 8 + ... = 66$ possible choices!
- The data are much noisier and harder to visualize
- We focused on one aspect, ie the inference made by the participants about the higher-order uncertainty state of the speaker

The speaker has drawn . . . balls

- Higher-order uncertainty represented by access value, 11 possible choices for each message: $0, 1, 2, \ldots, 10$
- Group access values into four levels of uncertainty:

none: $a = 10$
low: $a > 5$
mid: $a = 5$
high: $a < 5$

Levels of uncertainty

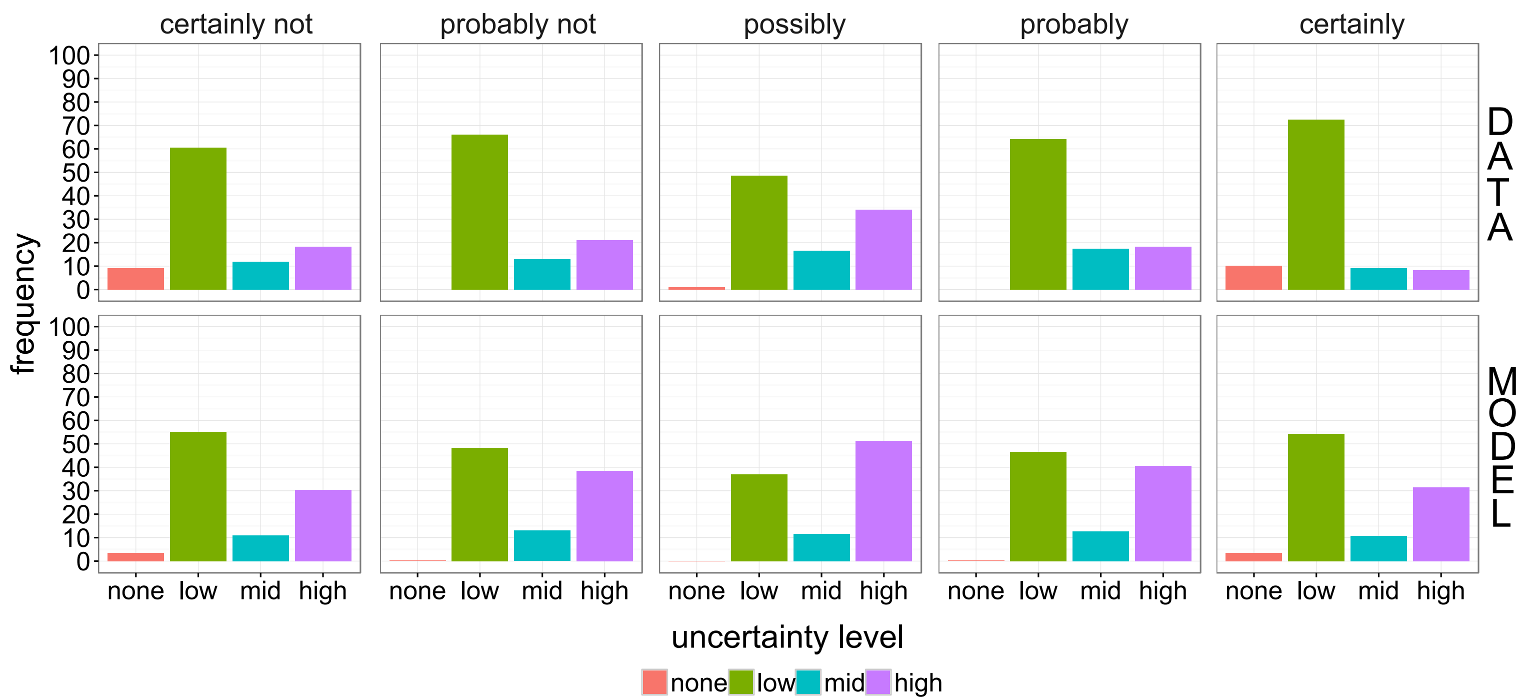

$r = 0.85$; 95% ci:0.65-0.94; $p < 0.001$

Wrapping up: use of probability expressions

- Pre-theoretical intuitions and experimental data support the idea that higher-order uncertainty plays a role in the use of probability expressions
- The same proportions of red balls with different levels of uncertainty give rise to different use patterns of probability expressions
- Our pragmatic model captures this observation in terms of cooperative communication of uncertain beliefs about the world, in the idealized urn setting

Wrapping up: interpretation of probability expressions

- Intuitively, if the speakers communicate beliefs about the world, we can expect listeners to be able to infer some information about the world
- Moreover, if speakers communicate higher-order uncertainty, we can expect listeners to be able to infer it
- Both expectations are borne out in the model.

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Thank you!

Appendix: instructions

This experiment is an interactive two player game of chance. The players cooperate to guess the content of an urn. Both players know that the urn ***always contains 10 balls of different colors (for example, red and blue)****. But only one player (the sender) is allowed to draw a certain number of balls from the urn and look at them. The sender puts the balls back into the urn and gives it a nice shake, then the sender draws a new ball from it. Before looking at it, the sender sends a message to the other player (the receiver). The receiver reads the message and tries to guess the exact content of the urn.*

Appendix: RSA model

(1) $P(o|a, s) = \text{hypergeometric}(o; a, s, 10)$

(2) $\text{speaker.bel}(s|o, a) \propto P(o|a, s) * \text{prior}(s)$

(3) $\text{literal.bel}(s|m) = \mathcal{U}(s|[\![m]\!]_s = 1) * \text{prior}(s)$

(4) $\text{EU}(m; o, a) = -\text{HD}[\text{speaker.bel}(s|o, a),\ \text{literal.bel}(s|m)]$

(5) $\text{speaker.prob}(m|o, a) \propto exp(\lambda * \text{EU}(m; o, a))$

(6) $\text{listener.prob}(s, o, a|m) \propto \text{speaker.prob}(m|o, a) * \text{priors}$

Appendix: HD between discrete distributions

$$\mathrm{HD}(P, Q) = \frac{1}{\sqrt{2}} \sqrt{\sum_i (\sqrt{P_i} - \sqrt{Q_i})^2}$$