Clarity, Ambiguity, and Hindsight Bias

Kevin Dorst
kmdorst@mit.edu

23.223 Rationality

I. Clarity and Ambiguity

Some judgments under uncertainty are *clear*; others are *ambiguous*:

Clear	**Ambiguous**
Clear chances. I'm about to flip a fair coin. What's your probability it'll land heads?	*Ambiguous chances.* What's your probability that it'll rain in Berlin next Monday?
Clear ignorance. What's your probability that the 27th digit of π is between 1–6?	*Ambiguous ignorance.* What's your probability that I used a spoon on the 27th of last month?

Suppose, in each case, there's a fact about what our subjective probability is—e.g. sampling propensity in generative model.

For each q, $P(q)$ is your probability.

Idea: Clear judgments are *higher-order certain* (HOC): we know what our (rational) probabilities are.
Ambiguous judgments are *higher-order **un**certain* (HOU).

HOC about q: $\exists a$: $P\Big(P(q) = a\Big) = 1$
HOU about q: $\forall a : P\Big(P(q) = a\Big) < 1$

Hypothesis: If Hedden is right, uncertainty about your(rational) prior drives hindsight bias. So people should exhibit hindsight bias under ambiguity, but *not* under clarity.

II. Primer on Higher-Order Probability

If we're to model a quantity X as uncertain, then it must be a *variable* that takes different values in different 'worlds' ('states').

In most models, your prior is treated as a *constant*, not a variable.

There's a strategy for making it a variable, and taming the infinite hierarchy (probabilities of probabilities of...) that ensues. It's rarely used in the *single*-agent case since, under clarity, the hierarchy is trivial.

What you need to know:

1) Under ambiguity, the hierarchy matters: uncertainty about your probability in q affects your probability in q.
2) 'P' is the variable for your prior—you're unsure of its value; '$P_@$' is a constant for what is (unbeknownst to you) your *actual* prior.
3) Assume that **elicitations are noisy**—so we can distinguish your *true* probability, $P(q)$, from your *elicited* probability $\epsilon(q)$.

'H' is a variable for my height
'72' is a constant for my *actual* height

$P(q)$ is what you really think;
$\epsilon(q)$ is what you write down.

III. The Finding

Hindsight Bias (HB): Learning that q leads people to increase their

The "I knew it all along" effect.

estimate for their *prior* probability in q.

Focus on pre-test / post-test designs.

Formalized: Your prior estimate for your prior in q is less than your posterior (given q) estimate for your prior in q: $\mathbb{E}_{P_@}\Big(P(q)\Big) < \mathbb{E}_{P_@}\Big(P(q)\Big|q\Big)$

$\mathbb{E}_{P_@}\Big(P(q)\Big)$ = prior estimate of $P(q)$ $= \sum_w P_@(w) \cdot P_w(q)$.

$\mathbb{E}_{P_@}\Big(P(q)\Big|q\Big)$ = *posterior* estimate, given q, of $P(q)$ $= \sum_w P_@(w|q) \cdot P_w(q)$.

As we saw, something *like* hindsight bias is clearly rational in (1) third-person cases, and (2) forgetting.

IV. Prediction: Hindsight rational $\Leftrightarrow$ prior is ambiguous

Hindsight bias can happen while *knowing* what information you had.

No forgetting.

Under clarity, Bayesians won't do this: if $P_@$ certain that $\langle P(q) = a \rangle$, then $P_@(\cdot|q)$ is too.

So $\mathbb{E}_{P_@}\Big(P(q)|q\Big) = a = \mathbb{E}_{P_@}\Big(P(q)\Big)$.

Under ambiguity?

Fact. A Bayesian commits hindsight bias on q iff their prior thinks that their prior in q is correlated with q's truth-value.

How much hindsight bias they'll do is bounded by how uncertain they are of $P(q)$.

Formally:
The *covariance* of X and Y, is $Cov_{P_@}(X,Y) = \mathbb{E}_{P_@}(XY) - \mathbb{E}_{P_@}(X) \cdot \mathbb{E}_{P_@}(Y)$.
The truth-value of q, $\mathbb{1}_q$, is a variable. So is your prior, $P(q)$.
Fact: $\mathbb{E}_{P_@}\Big(P(q)\Big) < \mathbb{E}_{P_@}\Big(P(q)\Big|q\Big) \Leftrightarrow Cov_{P_@}\Big(P(q), \mathbb{1}_q\Big) > 0$.

Empirical Predictions:

1) When people's opinion in q is *clear*, HB will be minimal.
2) When people's opinion is *ambiguous*, HB will be common.
3) The *extent* of HB will be modulated by degree of ambiguity.

So long as they trust their judgment.

I.e. by $1 - P_@\Big(P(q) = a\Big)$.

III. The Study

Assume that people have underlying probabilities $P_@$, but when probed we get a noisy *elicitation* ϵ of them.

When asked for estimate of X, we get $\mathbb{E}_\epsilon(X)$—a noisy indicator of $\mathbb{E}_{P_@}(X)$.

Design:

Step 1: Training.

- Distinguish their **true probability**[1] ('what they really think') from their **guess** about their true probability.[2]
- Point out that they can be more or less *confident* that their guess matches their true proability.
- Instruct that later in the survey, they'll have a chance to *revise their guess* about what their initial probability *was*.
- Give comprehension checks—exclude participants who fail.

[1] $P(q)$

[2] $\mathbb{E}_\epsilon(P(q))$, elicited estimate for $P(q)$. Estimator for $\mathbb{E}_{P_@}(P(q))$.

In the theory: if $\mathbb{E}_\epsilon(P(q)) = a$, their confidence is $P_@(P(q) = a)$.

$\mathbb{E}_\epsilon(P(q)|q)$; estimator for $\mathbb{E}_{P_@}(P(q)|q)$.

Step 2: Present people with *clear* or *ambiguous* evidence about q, and elicit their guesses about their probabilities. Eg:

$\mathbb{E}_\epsilon(P(q))$.

[Clear:] Emily works on the 7th floor of the TechFrontier office building, where the programmers and accountants work. They employ **40 programmers** and **60 accountants**. Emily was **selected randomly** from this group.

[Ambiguous:] Emily works on the 7th floor of the TechFrontier office building, where the **programmers** and **accountants** work. Please read the following vignette about her:

> Emily sat at her workstation in the center of the bustling office, bathed in the soft glow of dual monitors. Her desk displayed neatly organized post-it notes and a large mug with mathematical jokes. Her fingers moved swiftly across a mechanical keyboard, the clicks melding into the office's ambient sounds.
>
> She occasionally adjusted her black-rimmed glasses and scrutinized the spreadsheet before her, and rubbed her temples as she worked. A worn comic book peeked out from under some papers, and a picture of her Yorkie Poodle smiled up at her from her desk.

Step 3: Elicit their **clarity** in this estimate: 'How confident are you that your stated guess captures what you really think?' [0–1 slider]
Step 4: Repeat total of 5 times with different scenarios.

This is the **clarity** of their judgment. **Ambiguity** = 1 − clarity.

Step 5: Tell them what happened in each scenario, remind them of the evidence they had previously, and ask if they'd like to revise their guess for what they *initially* thought.

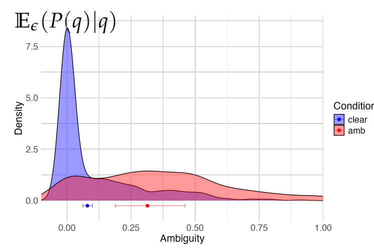

Predictions and results:

P1 Manipulation check: Ambiguous condition had higher ambiguity than Clear [»].

P2 People should be more likely to exhibit a *positive hindsight shift* ($\mathbb{E}_\epsilon(P(q)) < \mathbb{E}_\epsilon(P(q)|q)$) in ambiguous than clear scenarios [»].

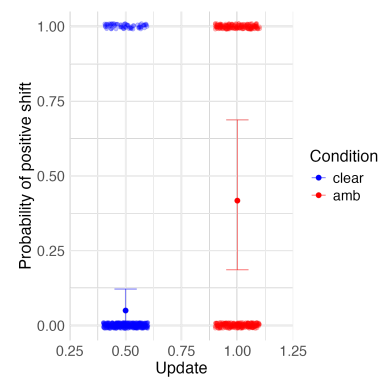

P3 Regressing on (self-reported) ambiguity should show that it's associated with greater probability of a positive hindsight shift:

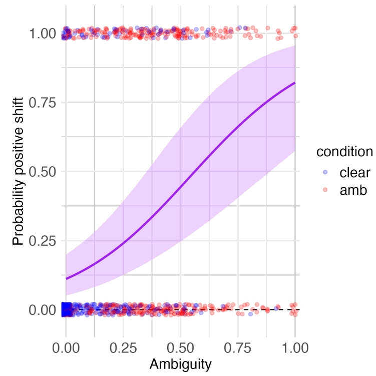

P4 Linear regression between ambiguity and *degree* of hindsight shift ($\mathbb{E}_\epsilon(P(q)|q) - \mathbb{E}_\epsilon(P(q))$) should be positive [»]