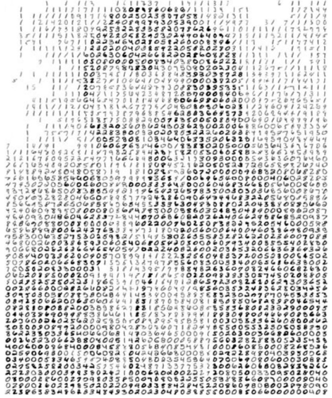

Overweighting extreme events reflects rational use of cognitive resources

Falk Lieder, Ming Hsu, Tom Griffiths

NEUROECONOMICS LABORATORY

Berkeley
UNIVERSITY OF CALIFORNIA

Extreme events influence people as if they were far more likely than they really are.

38% of Americans say they are less likely to travel overseas because of 9/11.

Availability Bias

Expected Utility Theory

utility of
outcome O

Violated!

Take action $\underset{a}{\operatorname{argmax}}\, \mathbf{E}_{p(O|a)}\left[u(O)\right]$

expected
value

$$\int p(o \mid a) \cdot u(o)\, do$$

Intractable!

EU can be approximated by sampling

$$\mathrm{EU} = \int p(o \mid a) \cdot u(o)\, do$$

$$o_1, \cdots, o_s \sim p(o \mid a)$$

$$\hat{a}^* = \underset{a}{\operatorname{argmax}}\ \hat{U}(a)$$

simulated outcomes → EU estimates → decision

$$\hat{U}(a) = \frac{1}{s} \sum_{i=1}^{s} u(o_s)$$

finite time ➔ finitely many simulated outcomes

In small samples variance kills you

vs.

bias

variance

Utility estimation by importance sampling

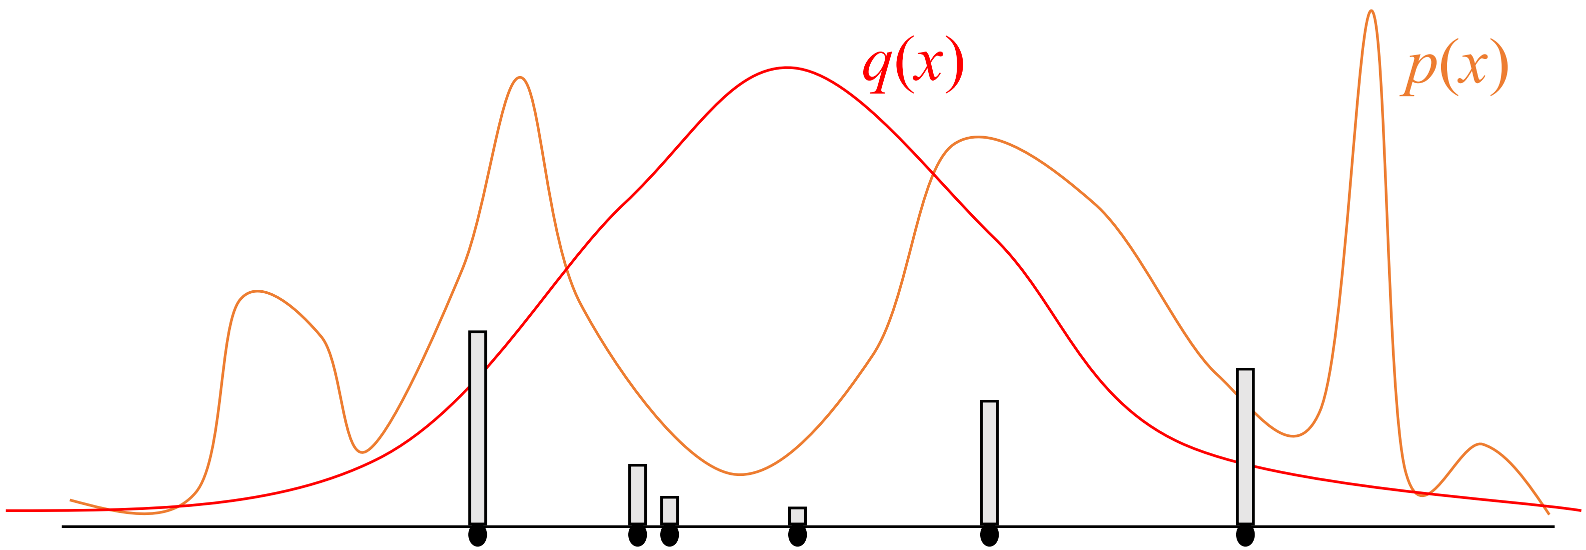

$$w^{(i)} = \frac{p(x^{(i)})}{q(x^{(i)})}$$

$$E[U] \approx \frac{\sum_i w^{(i)} u(x^{(i)})}{\sum_i w^{(i)}}$$

Utility estimation by importance sampling

$$\hat{a}^{\star} = \arg\max_{a} \hat{U}^{\mathrm{IS}}_{q,s}(a)$$

$o_1, o_2, o_3 \sim q$

simulated outcomes → EU estimates → decision

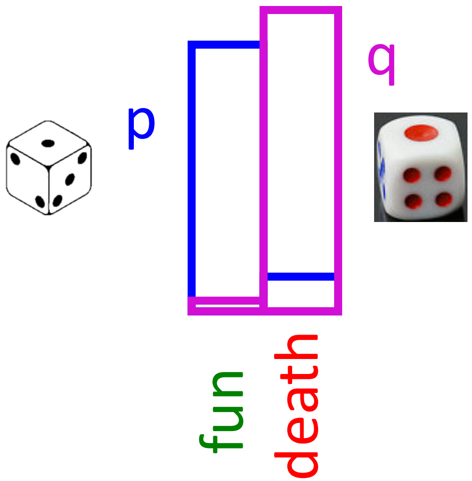

$$\hat{U}^{\mathrm{IS}}_{q,s}(a) = \frac{1}{\sum_{j=1}^{s} w_j} \sum_{j=1}^{s} w_j \cdot u(o_j)$$

$$w_1 = p(o_1) / q(o_1)$$

$$\dots$$

$$w_3 = p(o_3) / q(o_3)$$

Which distribution should the brain sample from?

Answer: Utility-Weighted Sampling (UWS)

probability

$$\tilde{q}(o) \propto p(o) \cdot |\Delta u(o)|$$

simulation frequency

extremity

Lieder, Hsu, and Griffiths (2014)

A simple optimal heuristic

When choosing between two options (e.g., shirt vs. jacket):

1. Imagine a few of possible events (e.g., *rain, sunshine, wind*).
2. For each imagined scenario, evaluate which action would fare better. (*jacket, shirt, jacket*).
3. Count how often the first action fared better than the second one. (*2 out of 3 times*)
4. If the first action fared better more often than the second action, then choose the the first action, else choose the second action. (*Wear a jacket!*)

(Lieder, Hsu, & Griffiths, 2014; Lieder, Griffiths, Hsu, 2015 & under review)

UWS captures biases in decision-making, memory, and frequency estimation

- Overestimation of the frequency of extreme events
- Over-weighting of extreme events in decisions from experience
- Extreme events come to mind first
- Temporal dynamics of risk preferences in decisions from experience
- Inconsistent risk preferences in decisions from description

Biases in frequency estimation

- Prediction:
  - Frequency overestimation increases monotonically with extremity.
- Method:
  - Recruited 100 participants on MTurk
  - 37 life events: 30 stressful, 4 lethal, and 3 mundane
  - Tasks:
    1. How many Americans experienced each of these events in 2015?
    2. Extremity of goodness/badness?
  - DV: relative overestimation = $\frac{\hat{f}-f}{f}$

Results

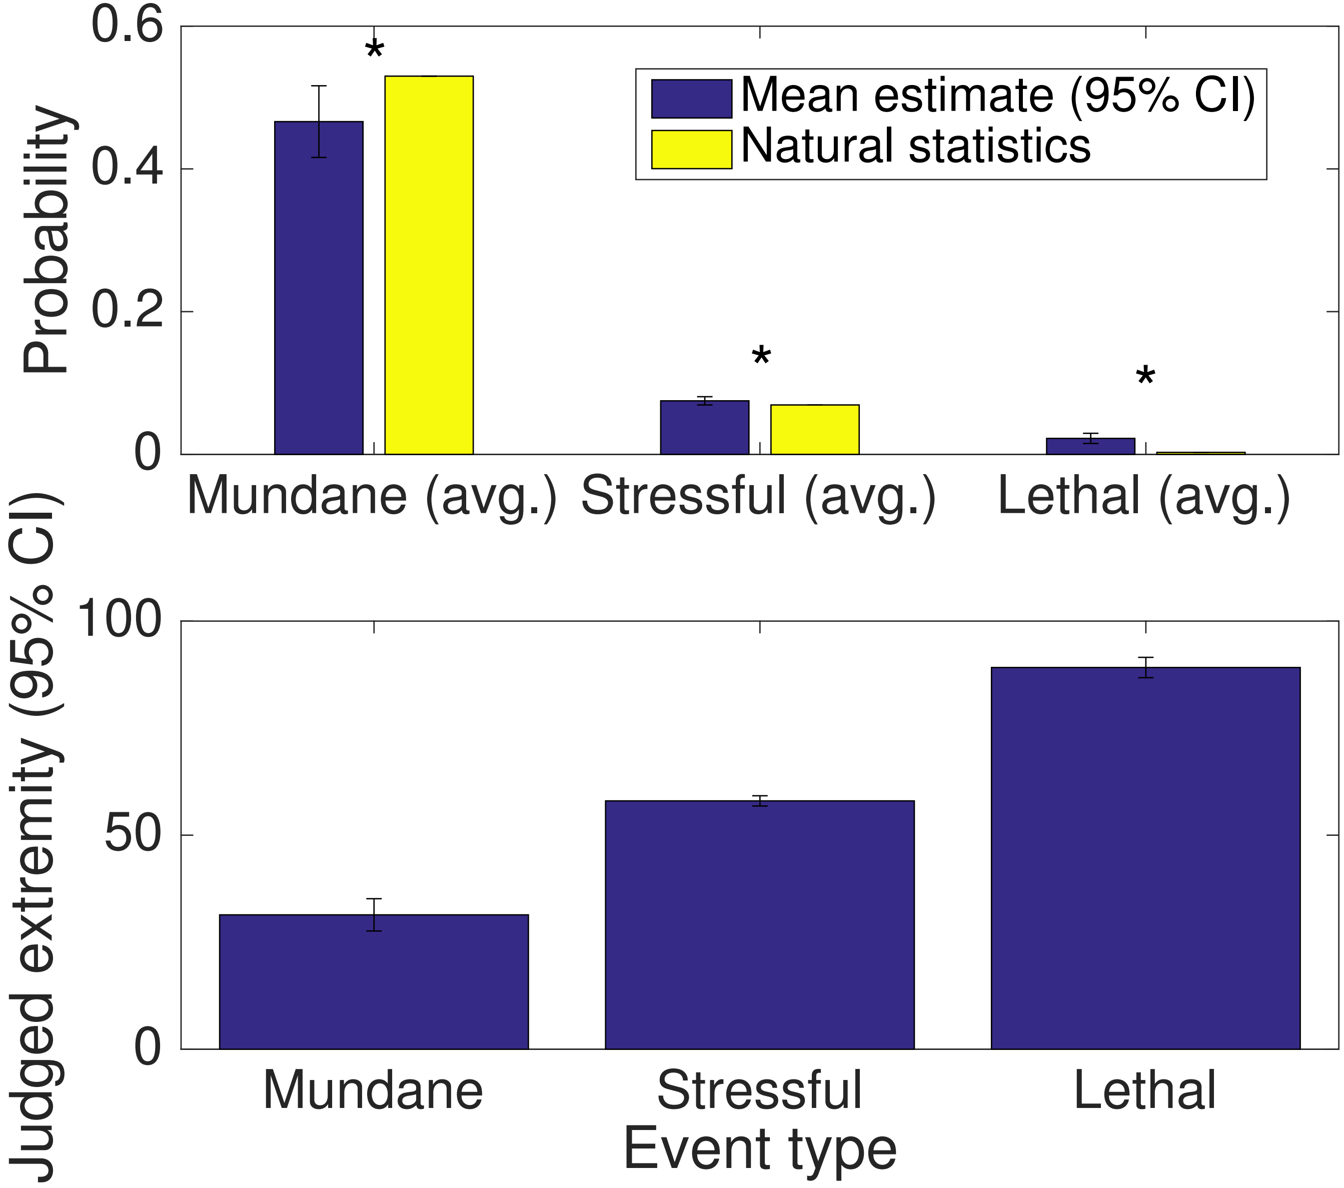

0.6
0.4
0.2
0
Probability
*
*
*
Mean estimate (95% CI)
Natural statistics
Mundane (avg.)
Stressful (avg.)
Lethal (avg.)
100
50
0
Judged extremity (95% CI)
Mundane
Stressful
Lethal
Event type

Results

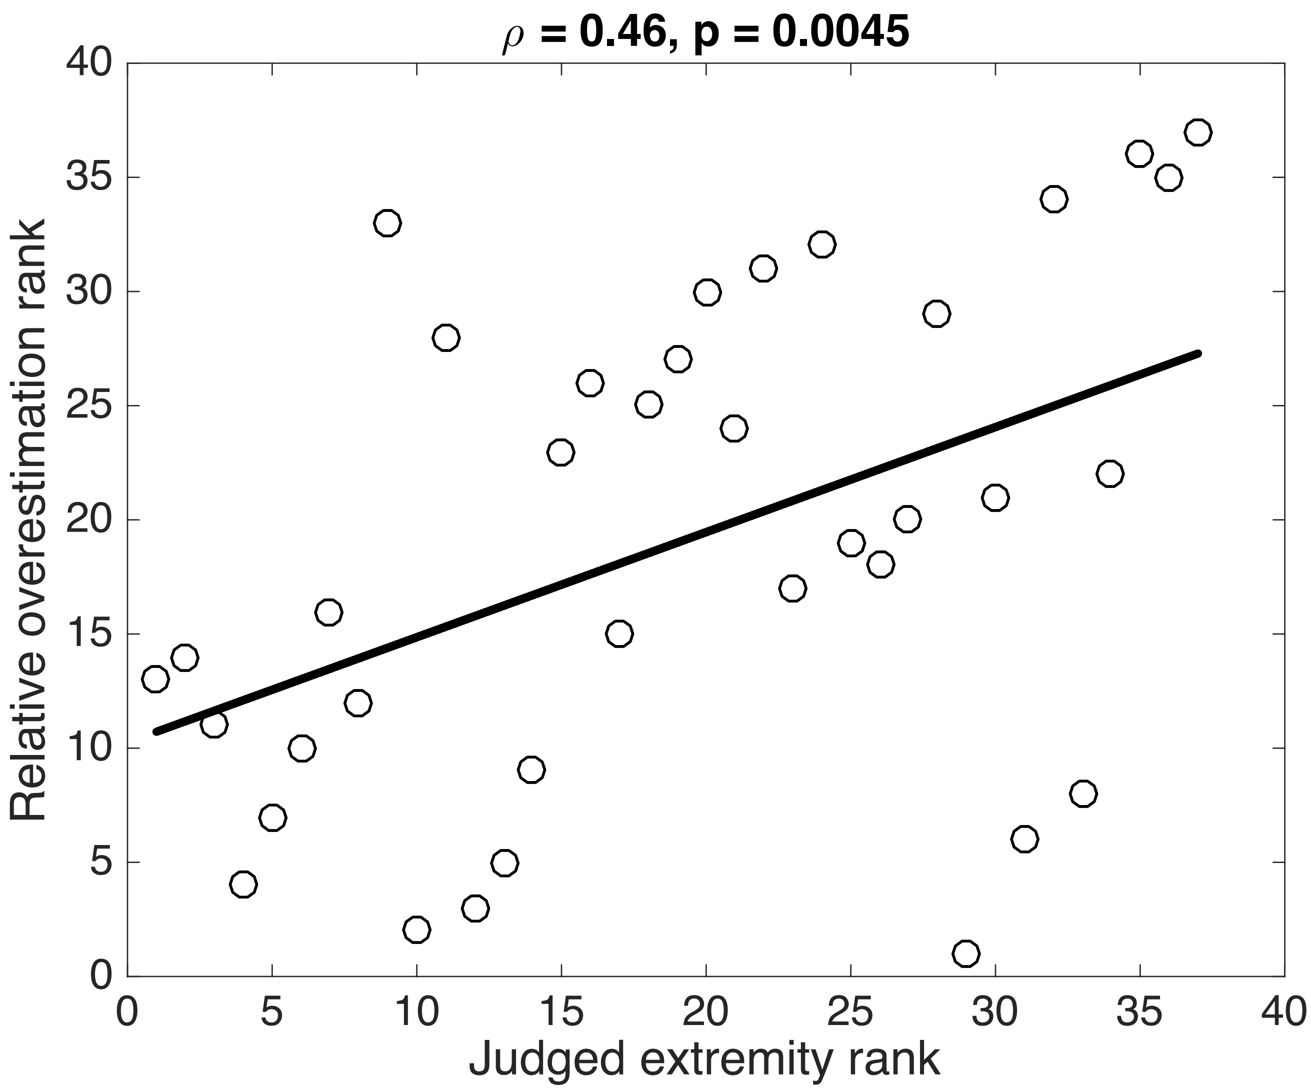

ρ = 0.46, p = 0.0045
Relative overestimation rank
Judged extremity rank
0
5
10
15
20
25
30
35
40

Decisions from Experience (Ludvig, et al., 2014)

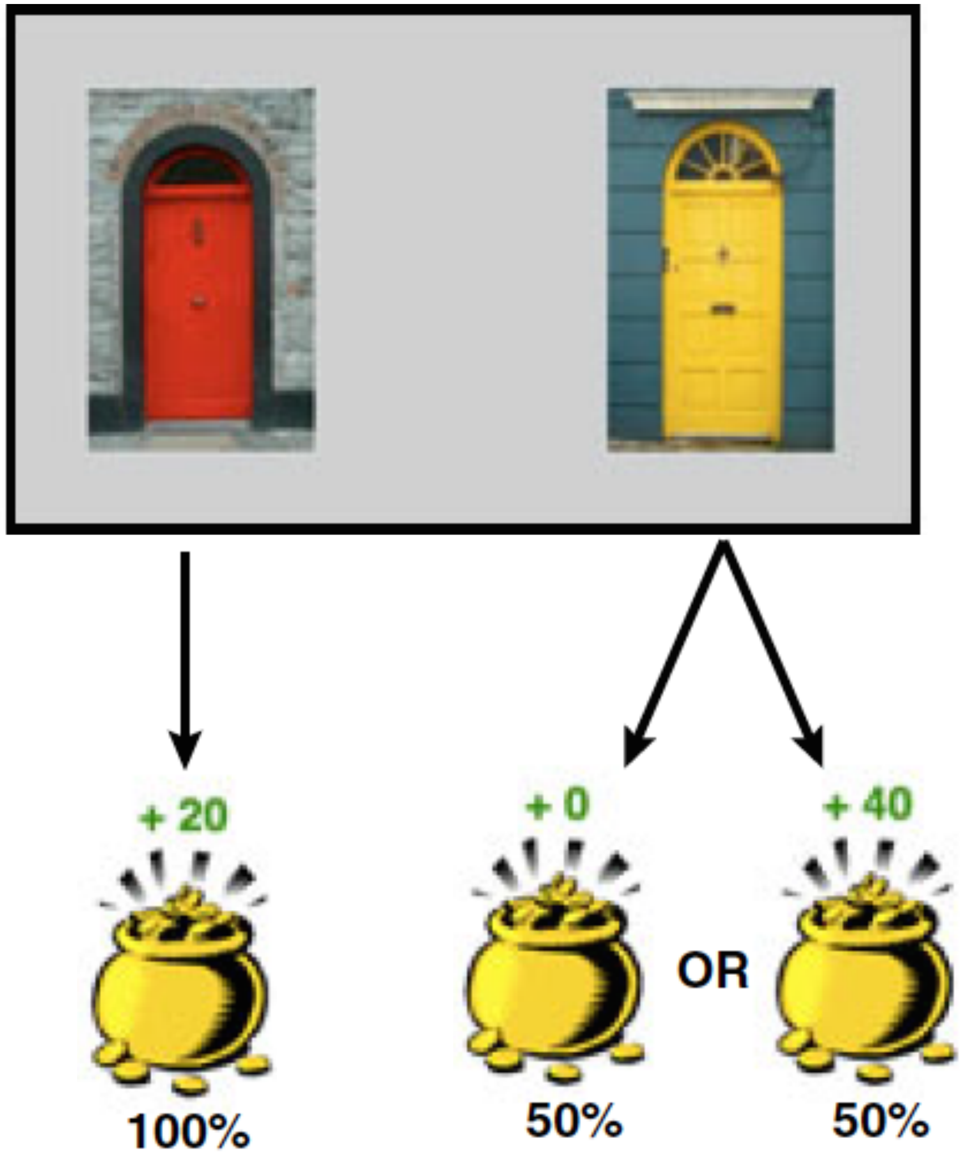

Decisions from Experience (Ludvig, et al., 2014)

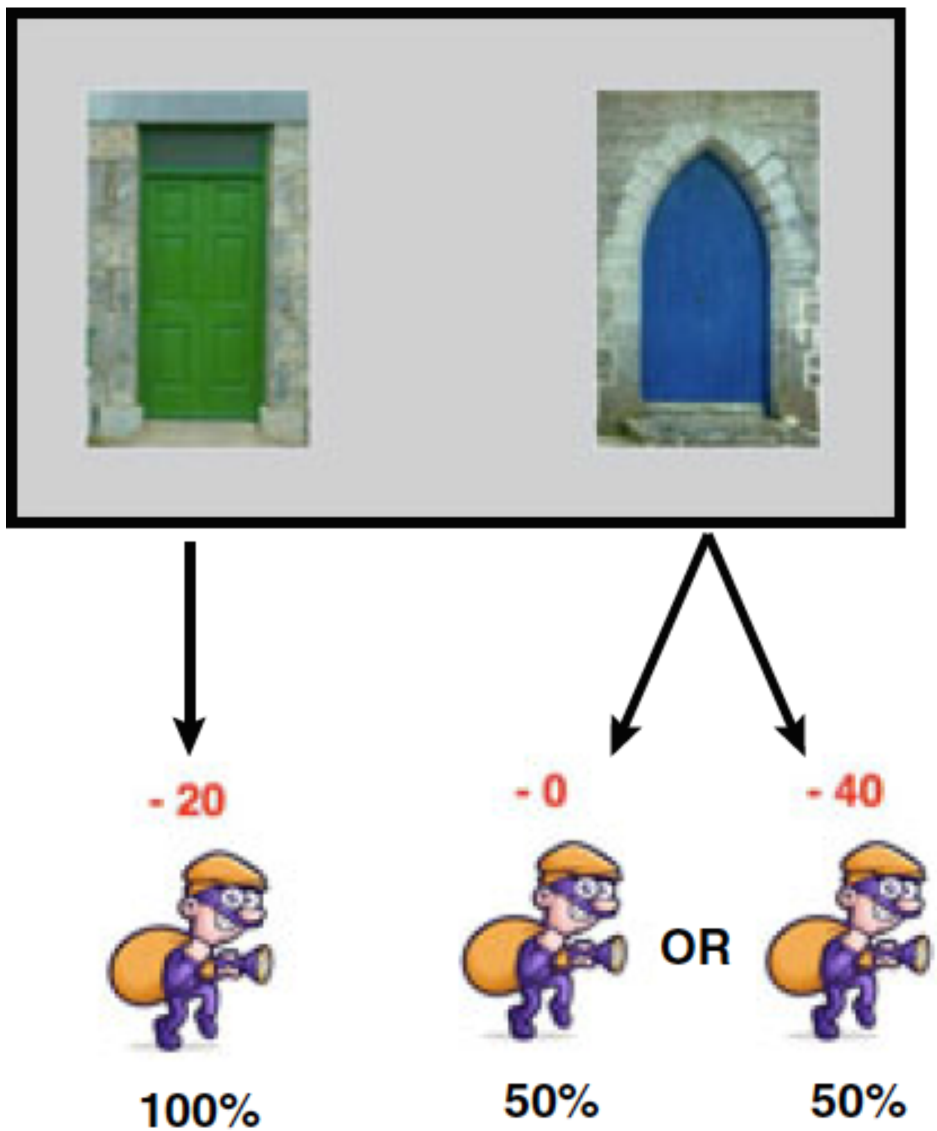

Learning → Inconsistent Risk Preferences

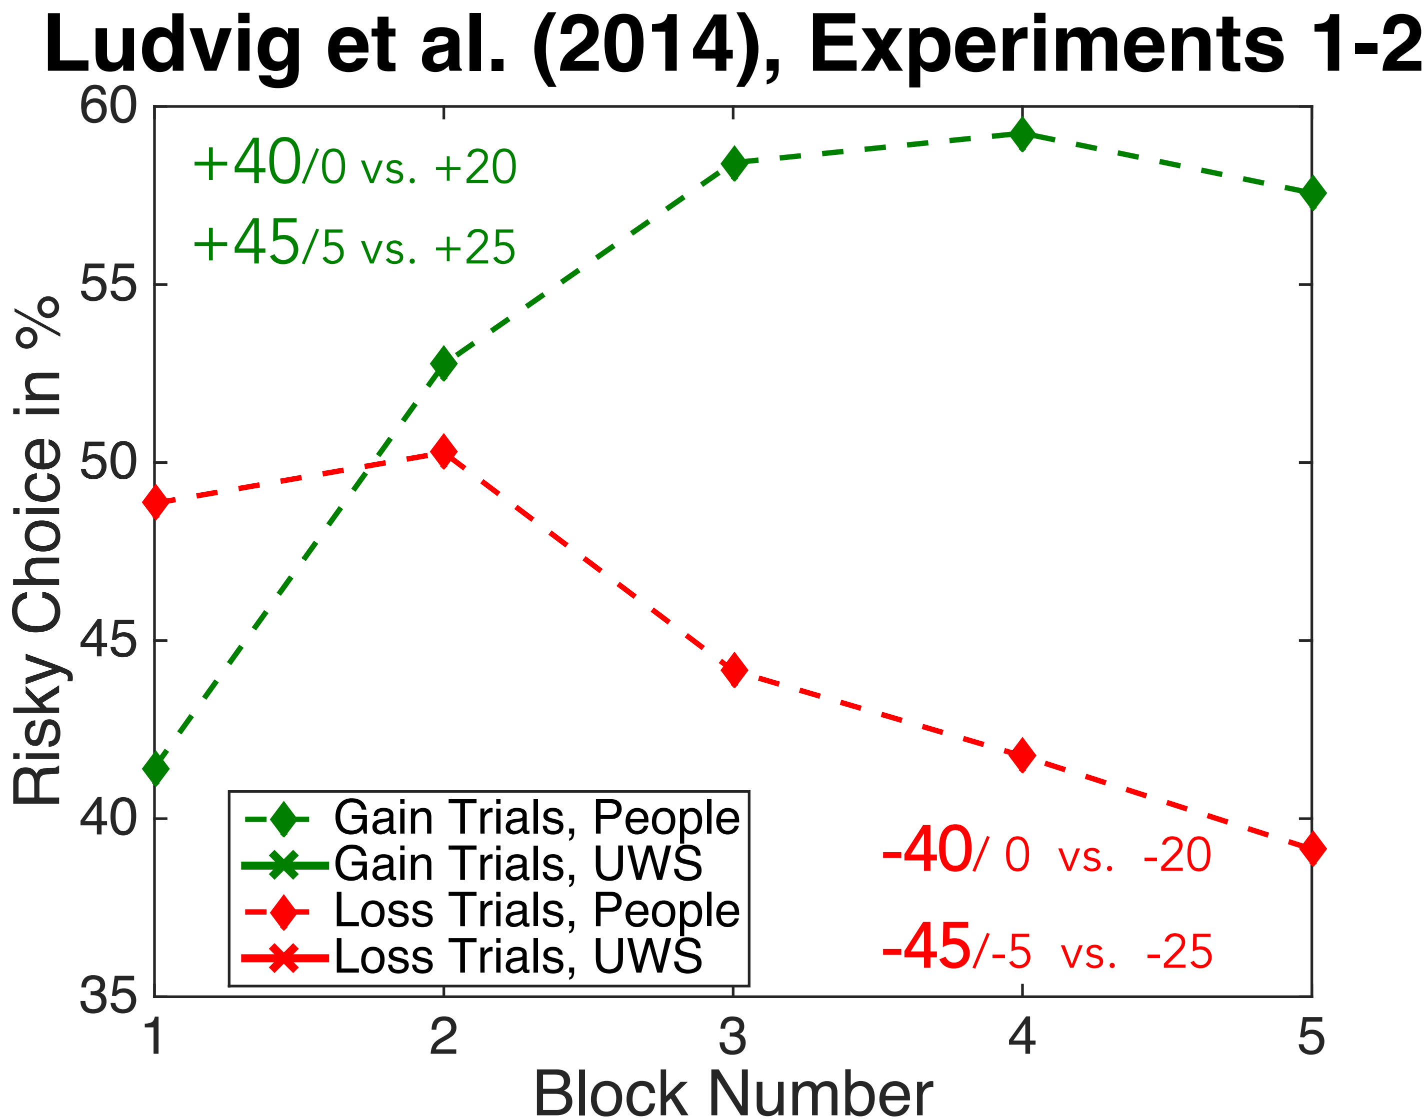

UWS Can Emerge from Reward-Modulated Associative Learning

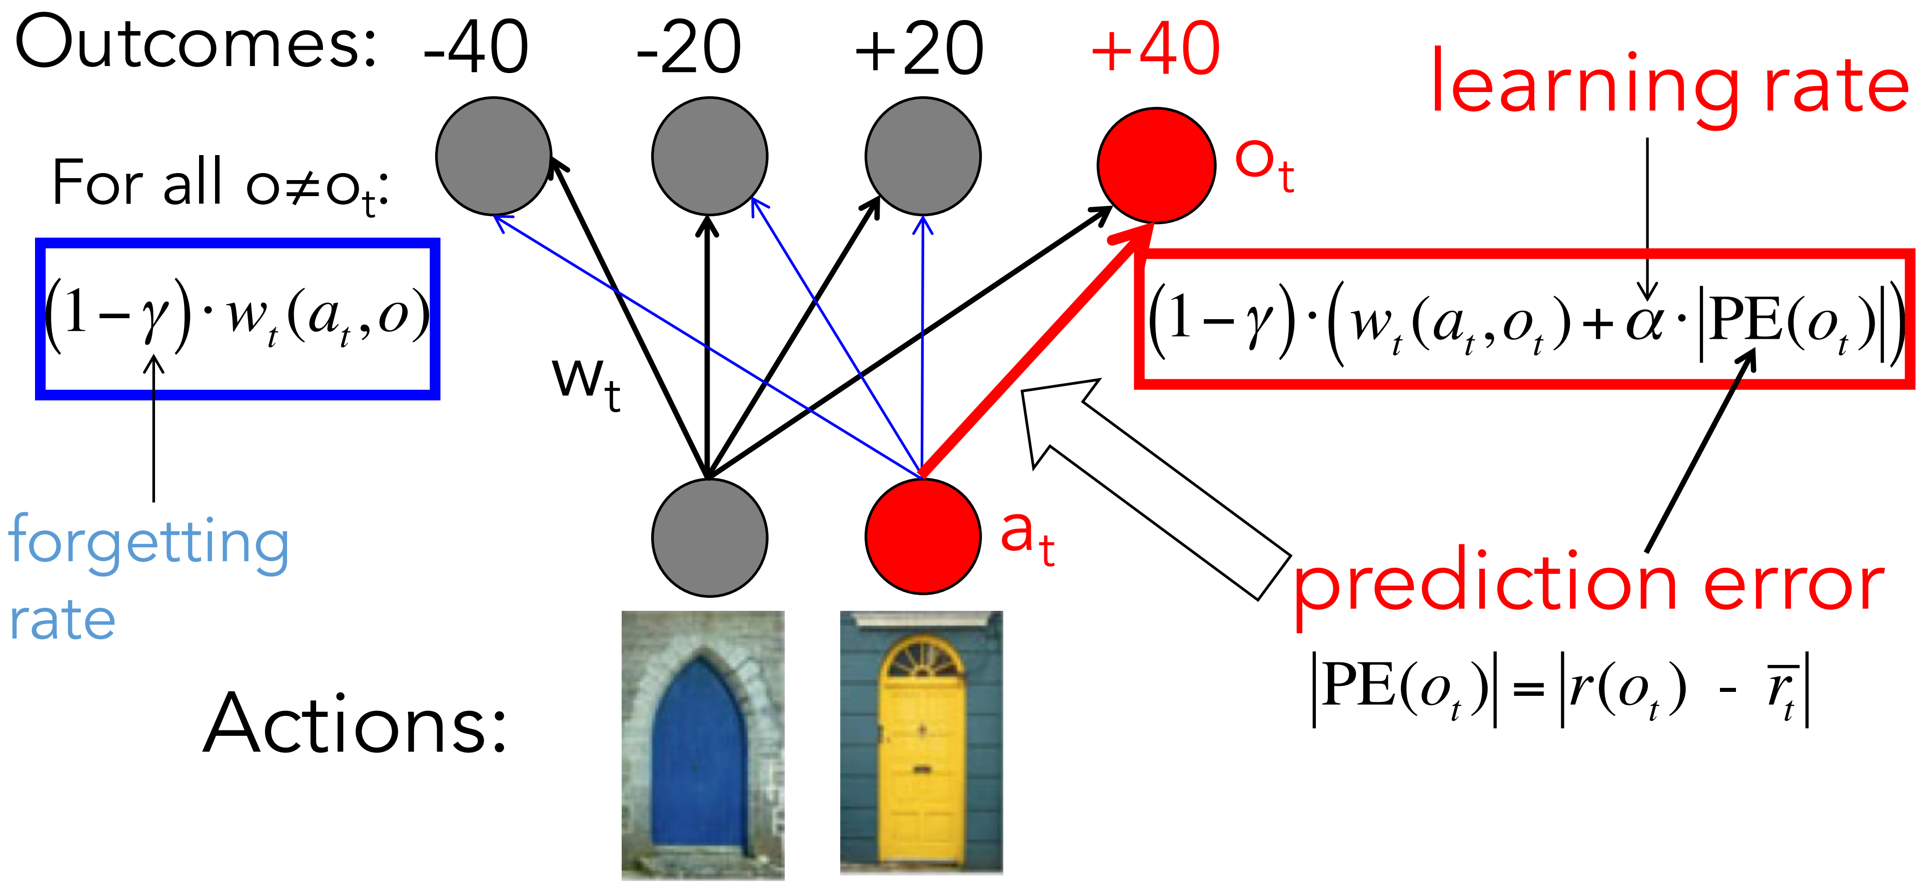

Learning Rule Convergences to Utility-Weighted Sampling

Utility-weighted learning converges to

$$w_{a,o} \propto p(o \mid a) \cdot |u(o)| \quad \text{with} \quad u(o) = \mathrm{PE}(o)$$

$$\Downarrow$$

with activation function $P(Y=1) \propto \mathbf{w}^t \cdot \mathbf{x}$ the network learns to perform utility-weighted sampling.

Efficient coding (Summerfield & Tsetsos, 2015)

$$\left|\mathrm{PE}(o_t)\right| = \left|r(o_t)\ \text{-}\ \bar{r}_t\right|$$

$$r(o) = \frac{o}{o_t^{\max} - o_t^{\min}} + \varepsilon$$

$$\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$$

$$\bar{r}_t \ = \ \bar{r}_{t\text{-}1} + \eta \cdot \left(r_t \text{-} \bar{r}_{t\text{-}1}\right)$$

Model fitting

Maximum-Likelihood-Estimation of $s, \alpha, \gamma, \lambda,$ and σ_ε^2 from block-by-block choice frequencies in Experiments 1-4 by Ludvig et al. (2014).

A single set of parameters fits all experiments.

People learn to overweight extreme events

High Extreme (HX):

+80/+40 vs. +60

0/-40 vs. -20

$$u(o) = \frac{o - \bar{o}}{o_t^{\max} - o_t^{\min}}$$

Low Extreme (LX):

+40/0 vs. +20

-40/-80 vs. -60

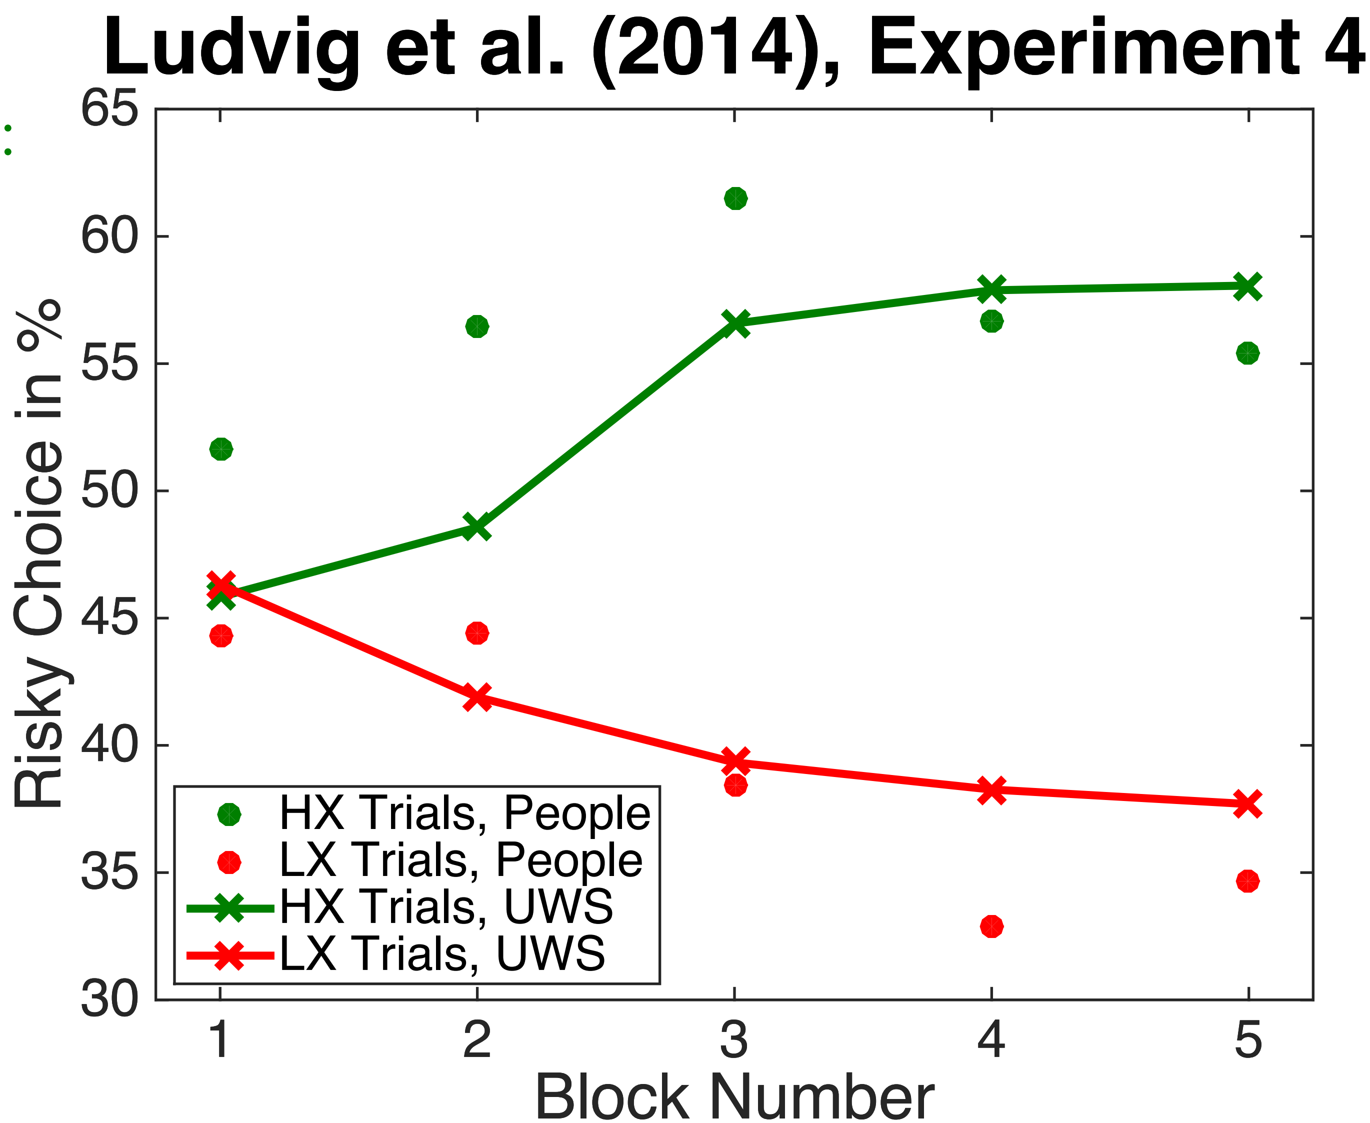

UWS captures that people learn to overweight extreme outcomes

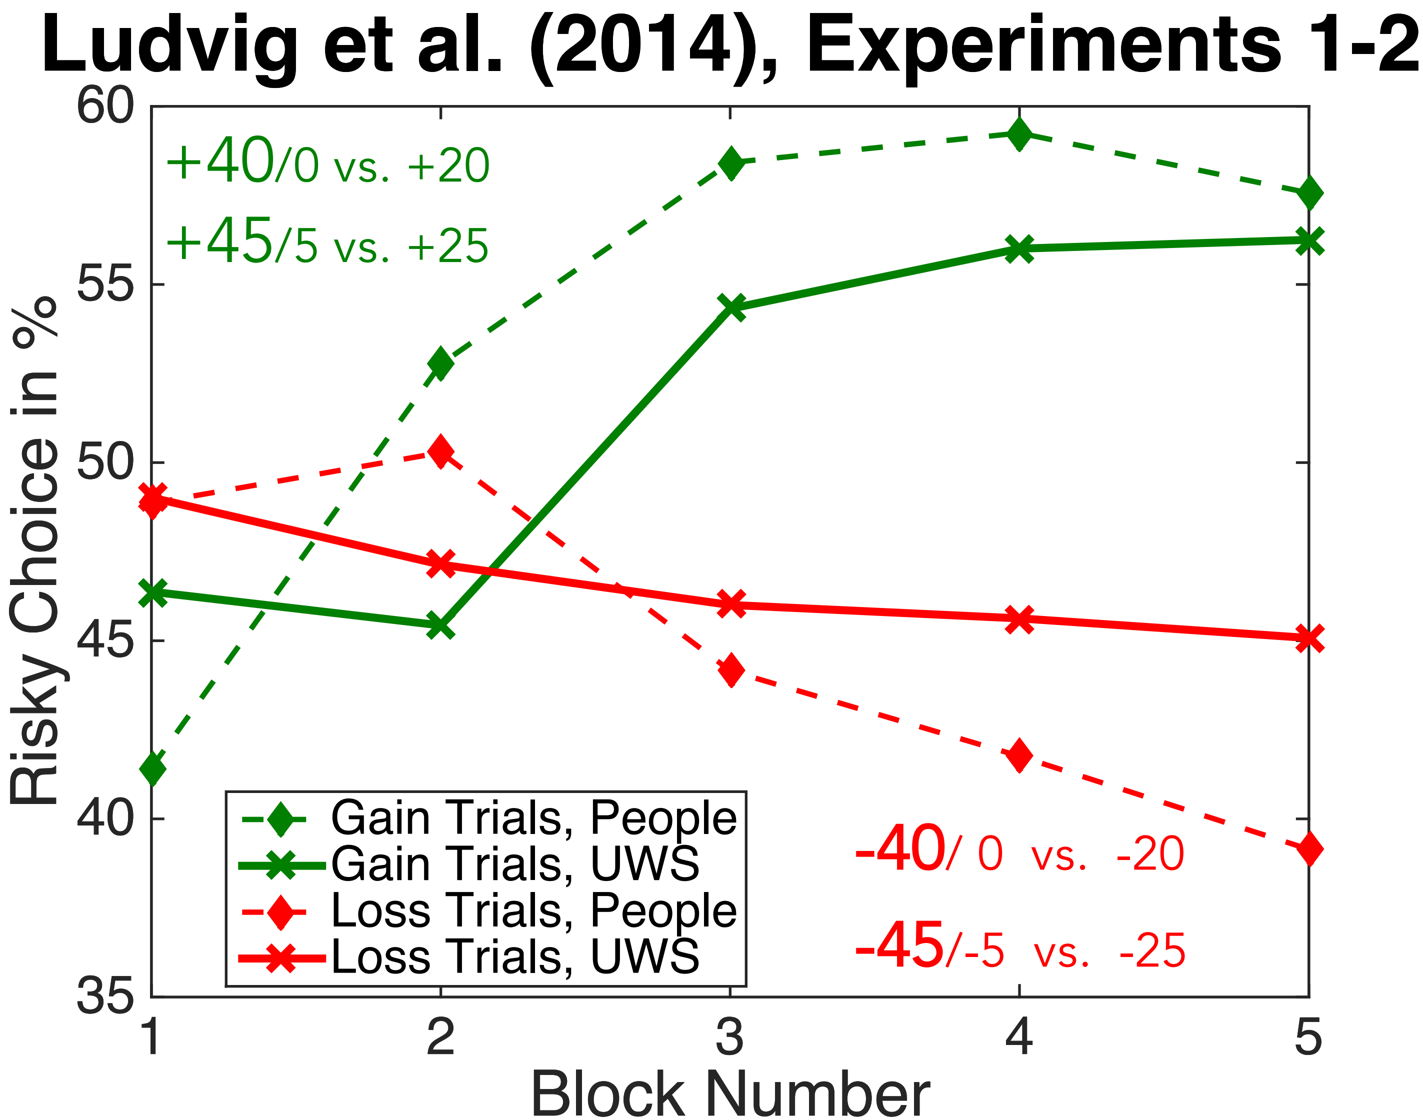

Extremity is relative

Extreme:

+80/0 vs. +40

-80/0 vs. -40

Nonextreme:

+40/0 vs. +20

-40/0 vs. -20

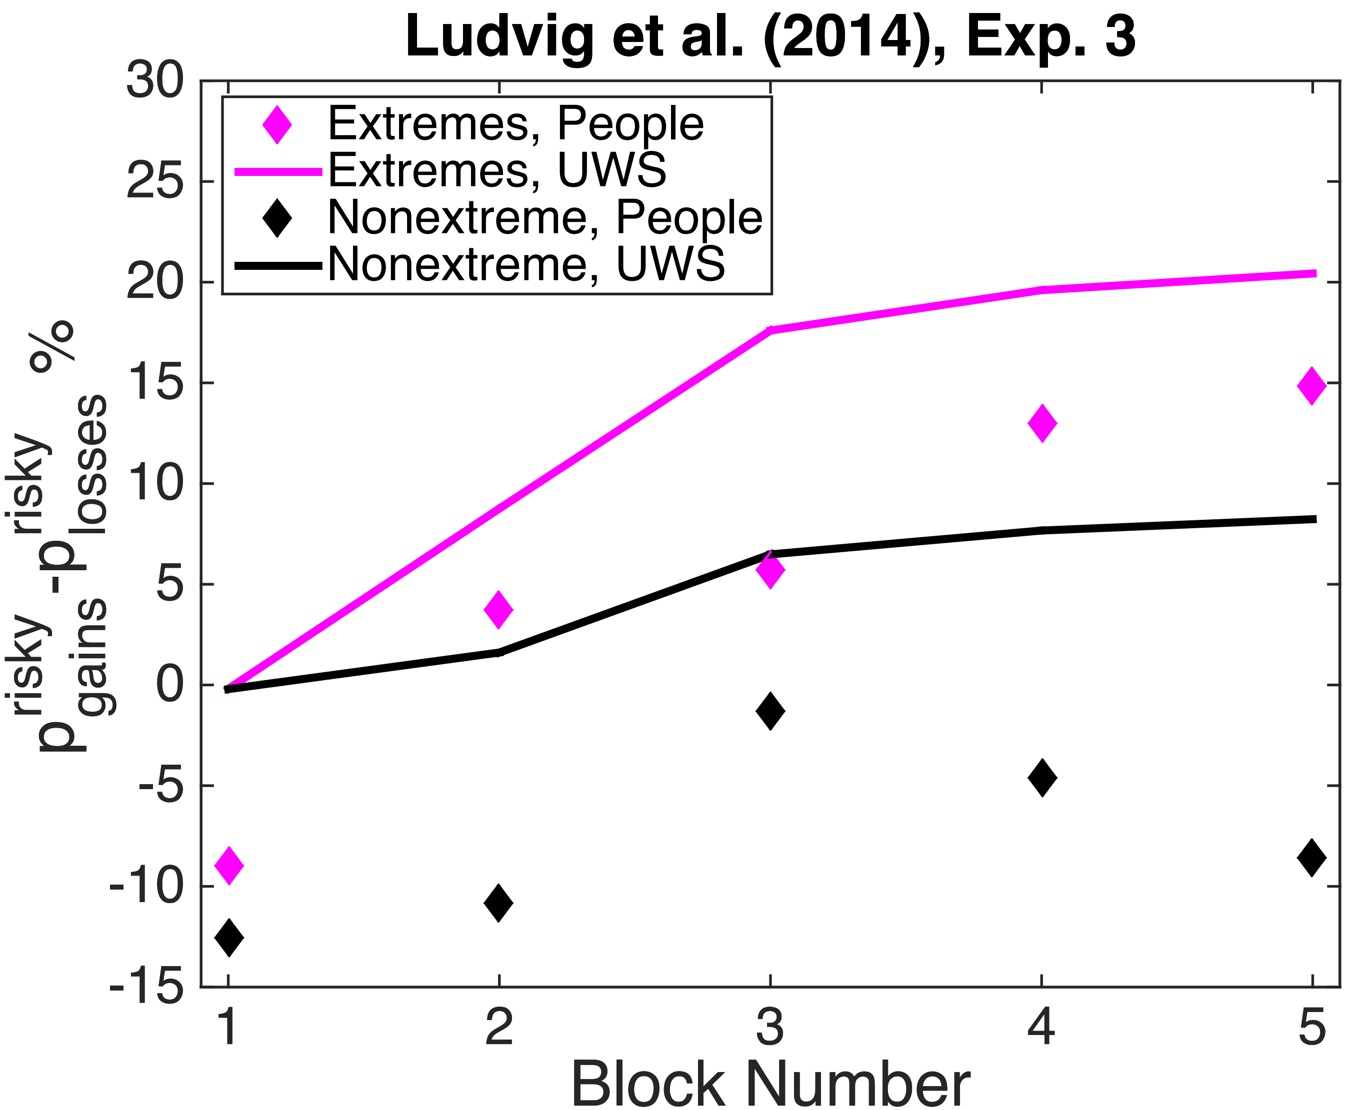

Memory Biases (Madan et al. 2014)

Which outcome comes to mind first?

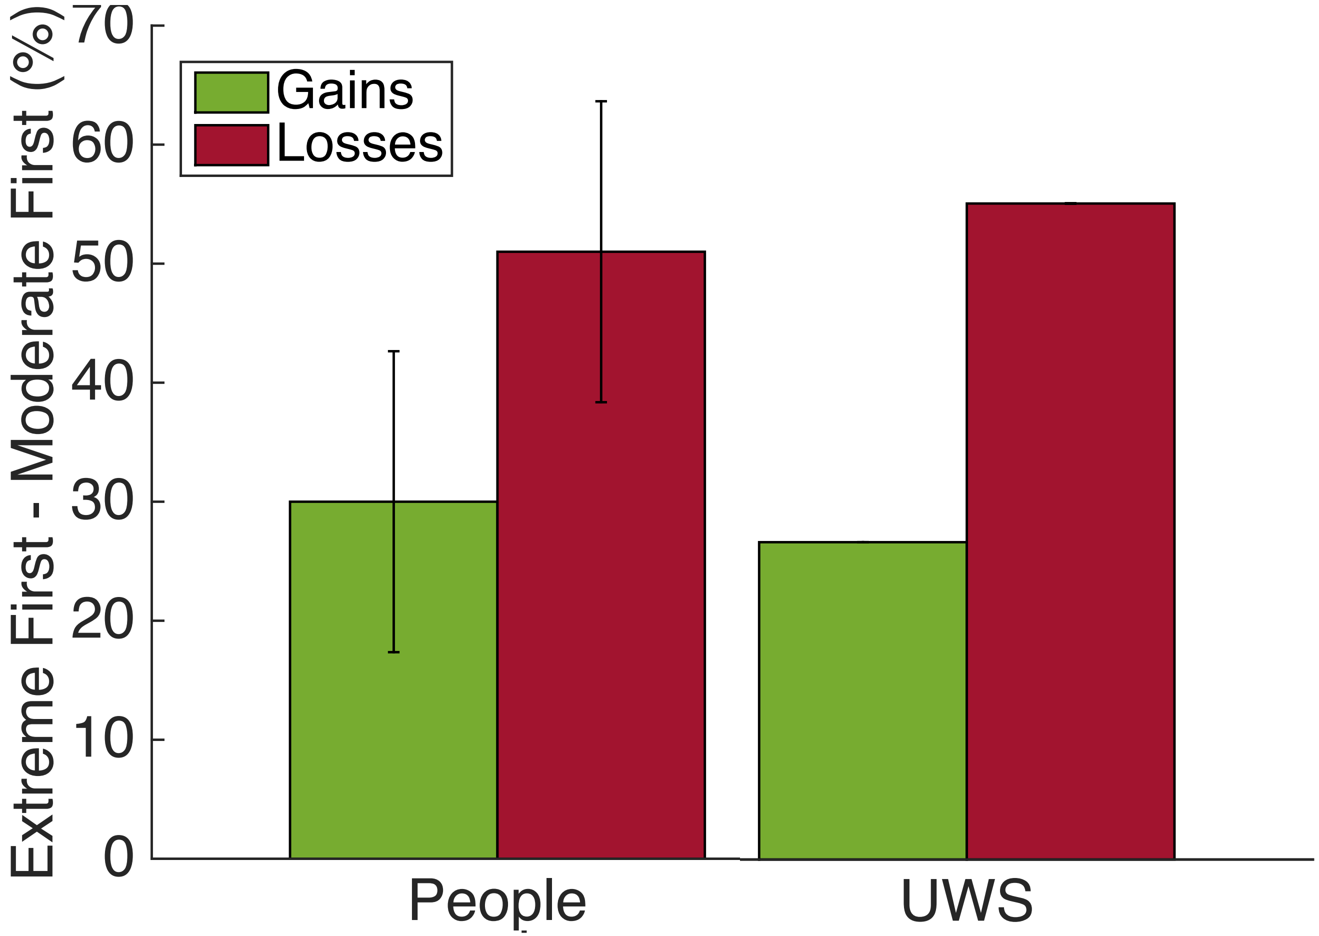

Frequency Estimation Bias (Madan et al. 2014)

How often did this door lead to each outcome?

-40: ___ %
0: ___ %
-20: ___ %

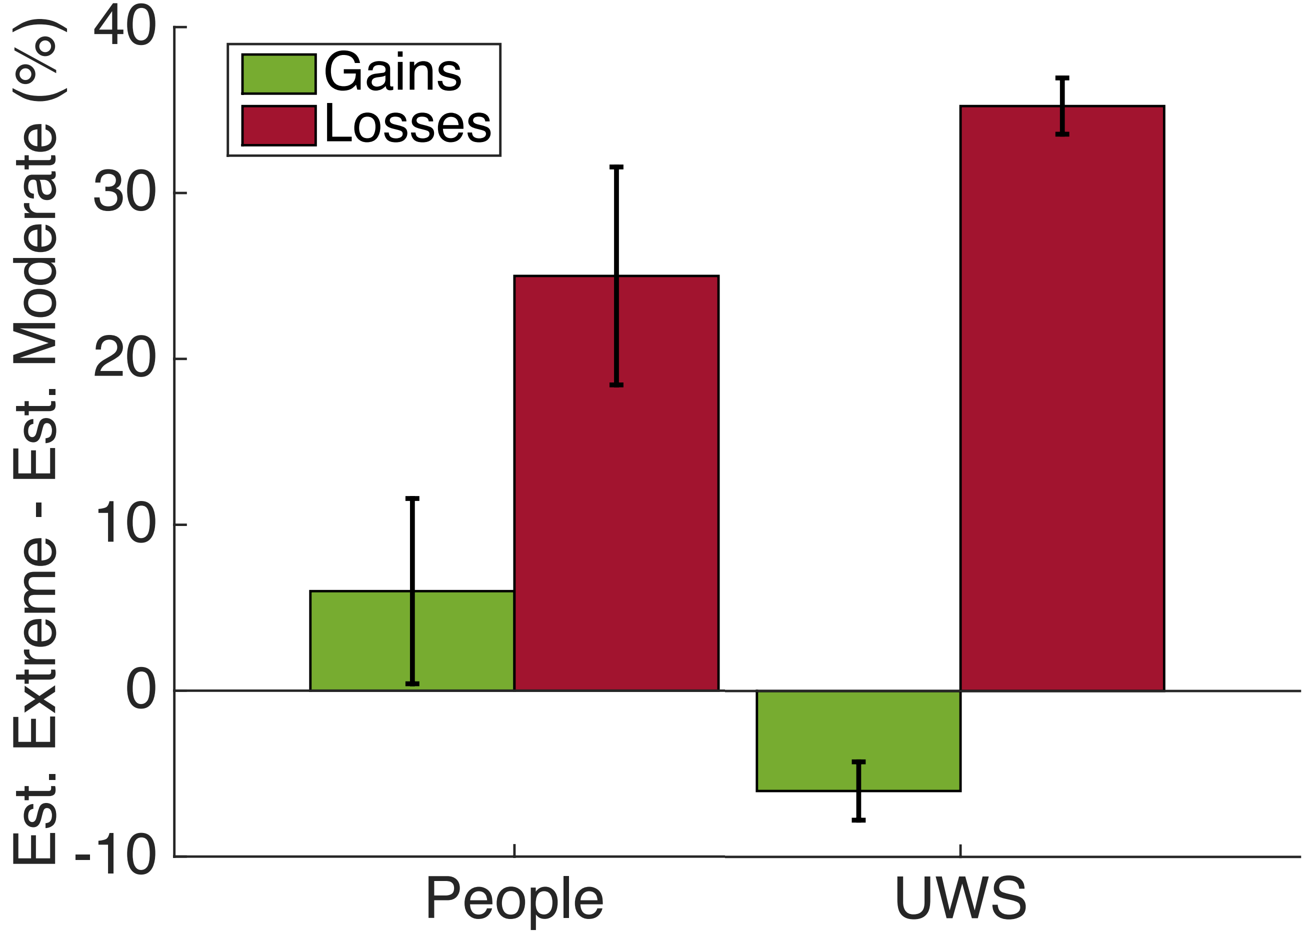

Biased Beliefs Predict Risk Seeking

r_{UWS} = +0.23
r_{people}= +0.16; $p<0.05$

r_{UWS} = -0.44
r_{people} = -0.48; $p< 0.05$

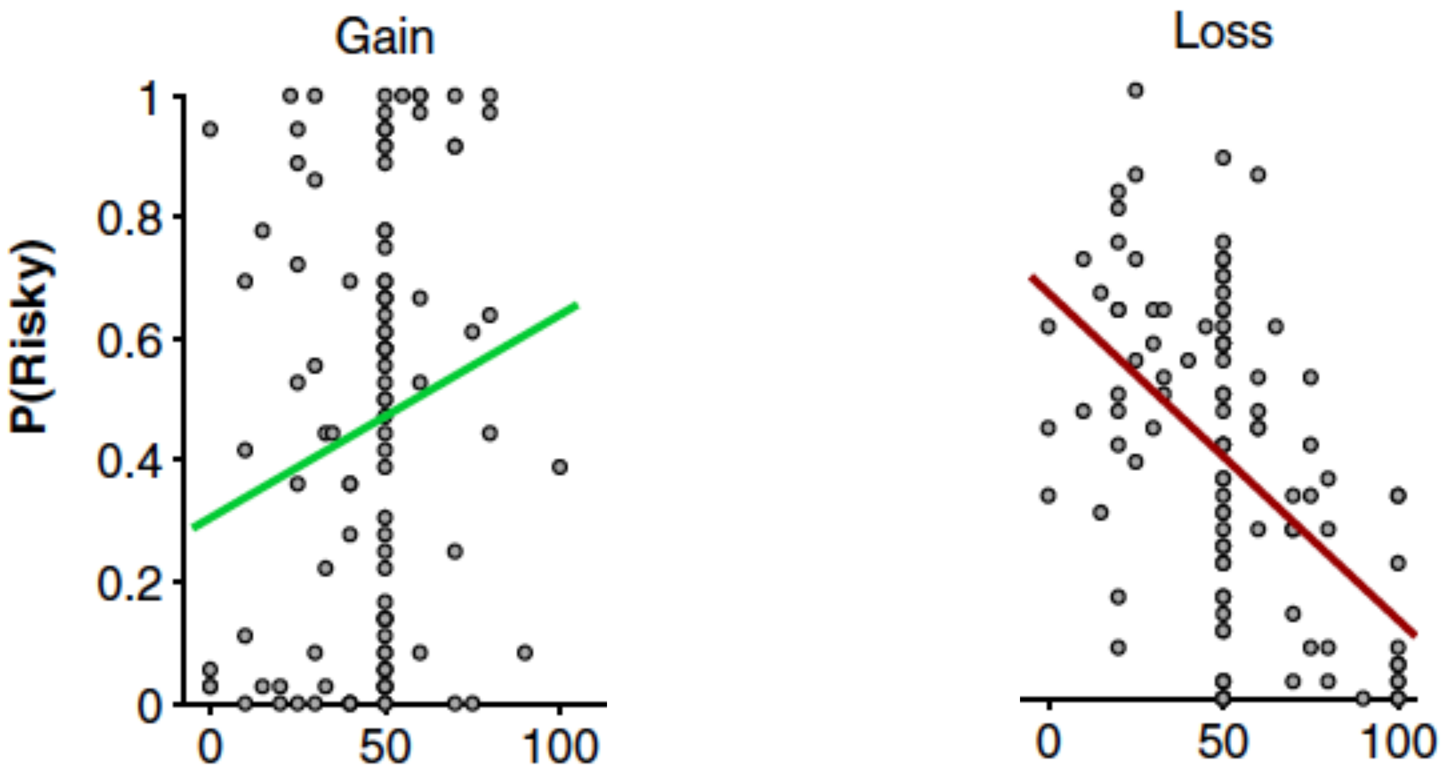

Judged Freq. of High Gain (%) Judged Freq. of Large Loss (%)

Inconsistent risk preferences in decisions from description

- Fourfold pattern of risk preferences (Tversky & Kahneman, 1992)
- Allais paradox (Allais, 1953)
- Preference reversals from pricing to choice
- Outperforms cumulative prospect theory in the Technion prediction tournament (Erev et al., 2010)
- Real-life decisions of contestants in the game show Deal-No-Deal (Post et al., 2008)

(Lieder, Hsu, Griffiths, 2014; Lieder, Griffiths, Hsu, under review)

UWS captures fourfold pattern of risk preferences

\$$o$ with prob. p vs. \$$(p \cdot o)$ for sure

$$\tilde{q}(o) \propto p(o) \cdot |u(o) - u(p \cdot o)|$$
$$\tilde{q}(0) \propto p(o) \cdot |u(p \cdot o)|$$

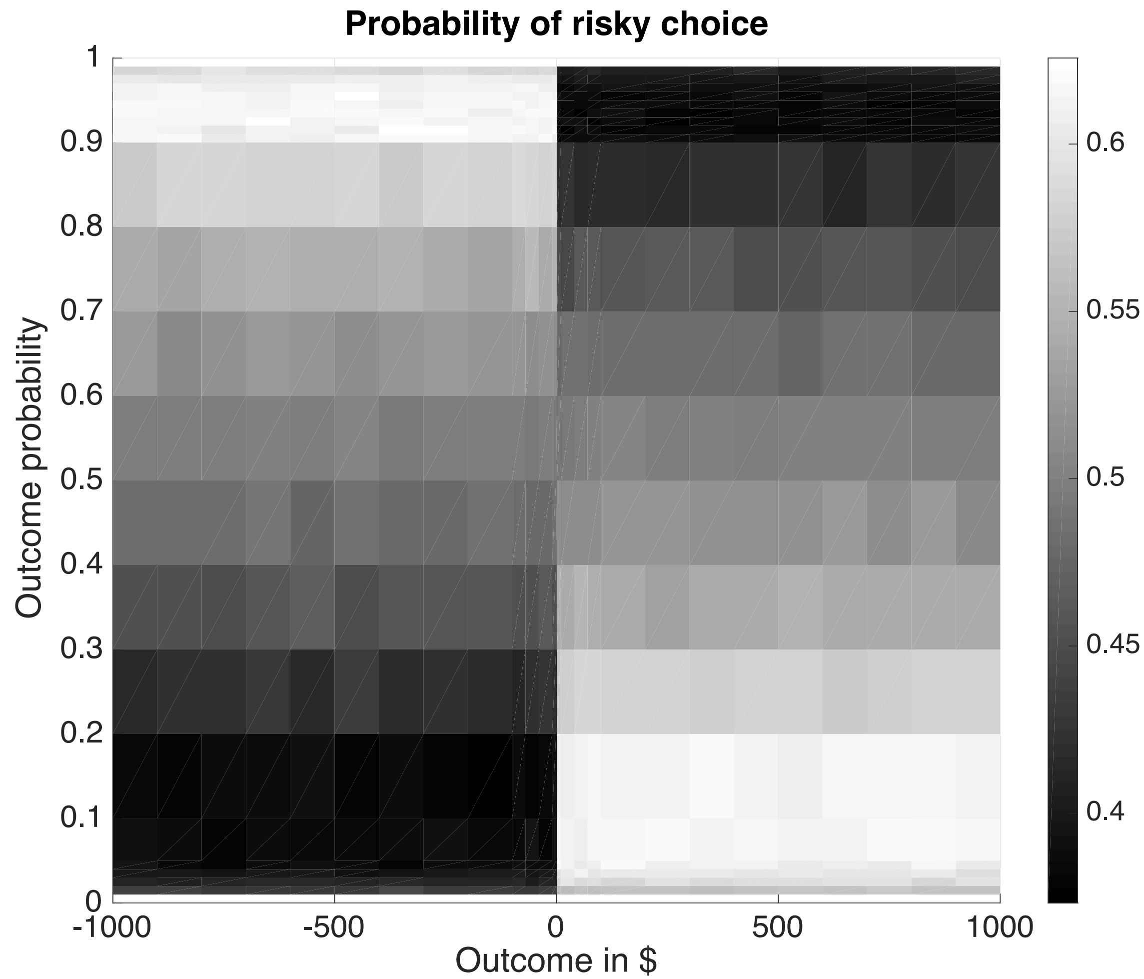

UWS captures risk preferences in Technion choice prediction competition

$MSD_{UWS} = 0.0266$ vs. $MSD_{CPT} = 0.0837$
($t(59) = -5.4$, $p < .001$)

UWS risk preference agrees with people's in 87% of the choices.

$r_{UWS}(59) = 0.88$, $p < 10^{-15}$
vs. $r_{CPT} = 0.86$ and $r_{priority} = 0.65$

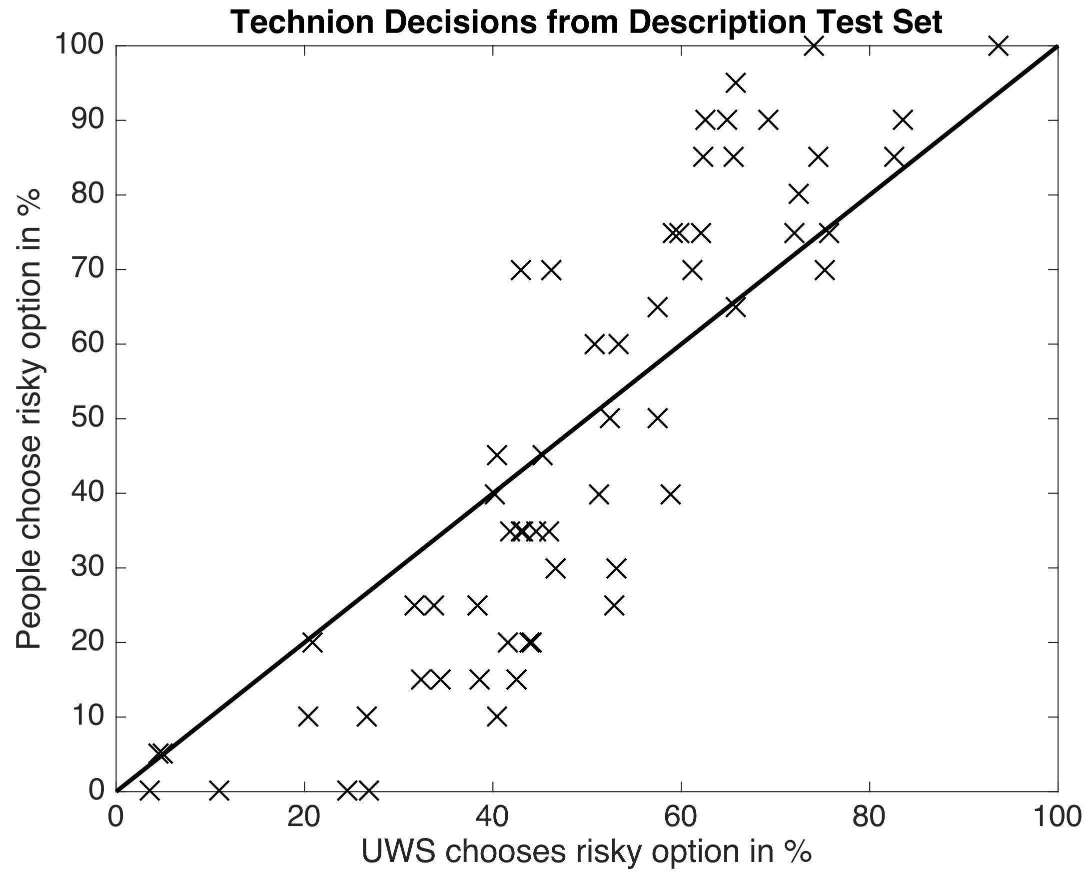

Conclusions

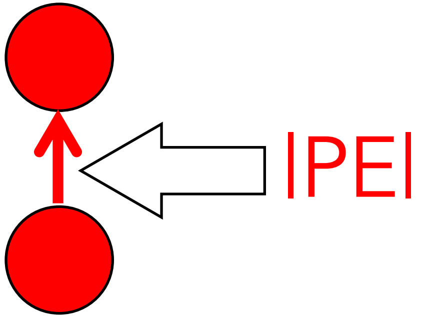

1. Utility-weighted sampling provides a unifying explanation for biases in memory, judgment, and decision making.
2. Utility-weighted sampling can emerge from reward-modulated associative learning.
3. People overweight extreme events, because it is rational to focus on the most important eventualities.
4. Some cognitive biases may serve or reflect the rational allocation of finite cognitive resources.