NEUROECONOMICS LABORATORY

Utility-weighted sampling in decisions from experience

Falk Lieder, Tom Griffiths, Ming Hsu

UC Berkeley

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

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

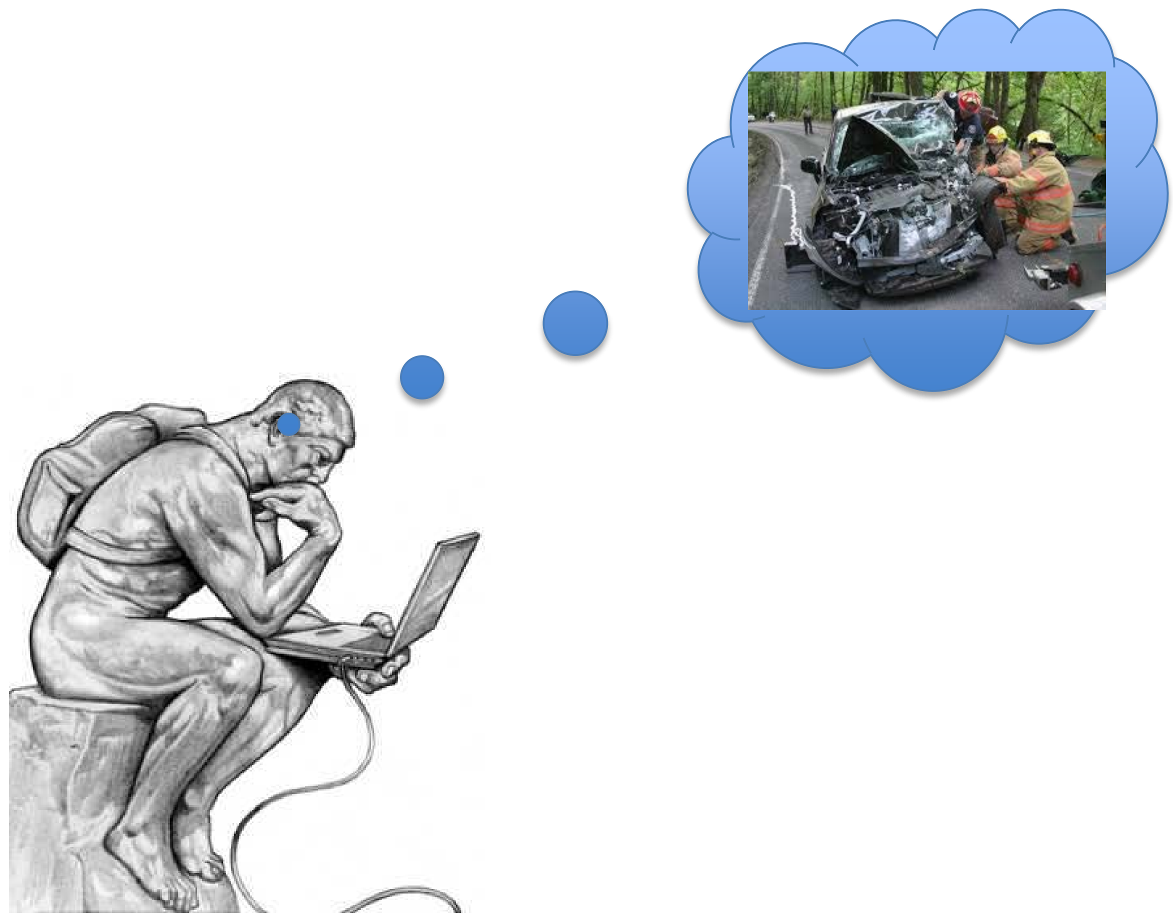

Extreme potential outcomes influence people as if they were far more likely than they really are.

Extreme potential outcomes 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.

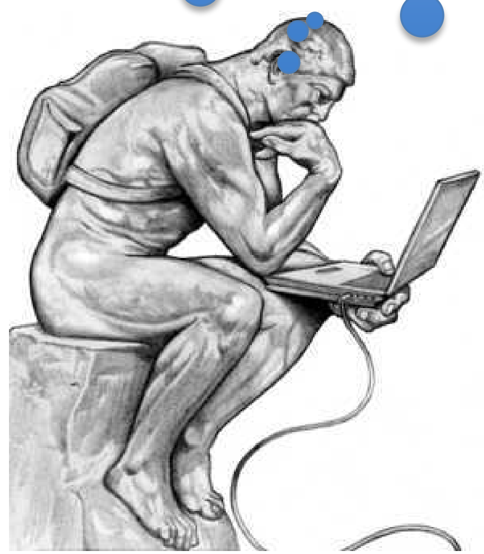

Extreme potential outcomes influence people as if they were far more likely than they really are.

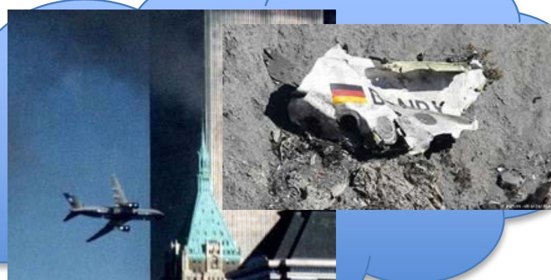

38% of Americans say they are less likely to travel overseas because of 9/11.

Expected Utility Theory

Expected Utility Theory

utility of
outcome O

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

expected
value

Expected Utility Theory

utility of
outcome O

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

expected
value

$$\square p(o \mid a) \cdot u(o) \; do$$

Expected Utility Theory

utility of
outcome O

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

expected
value

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

Intractable!

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

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

EU can be Approximated by Sampling

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

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

simulated
outcomes

EU can be Approximated by Sampling

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

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

simulated outcomes $\longrightarrow$ EU estimates

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

EU can be Approximated by Sampling

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

$$o_1, \cdots, o_s \sim p(o \mid a) \qquad \hat{a}^* = \underset{a}{\operatorname{argmax}}\ \hat{U}(a)$$

simulated outcomes → EU estimates → decision

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

EU can be Approximated by Sampling

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

$$o_1, \cdots, o_s \sim p(o \mid a) \qquad\qquad \hat{a}^* = \operatorname*{argmax}_a \hat{U}(a)$$

simulated outcomes → EU estimates → decision

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

finite time → finitely many simulated outcomes

Representative sampling is dangerous

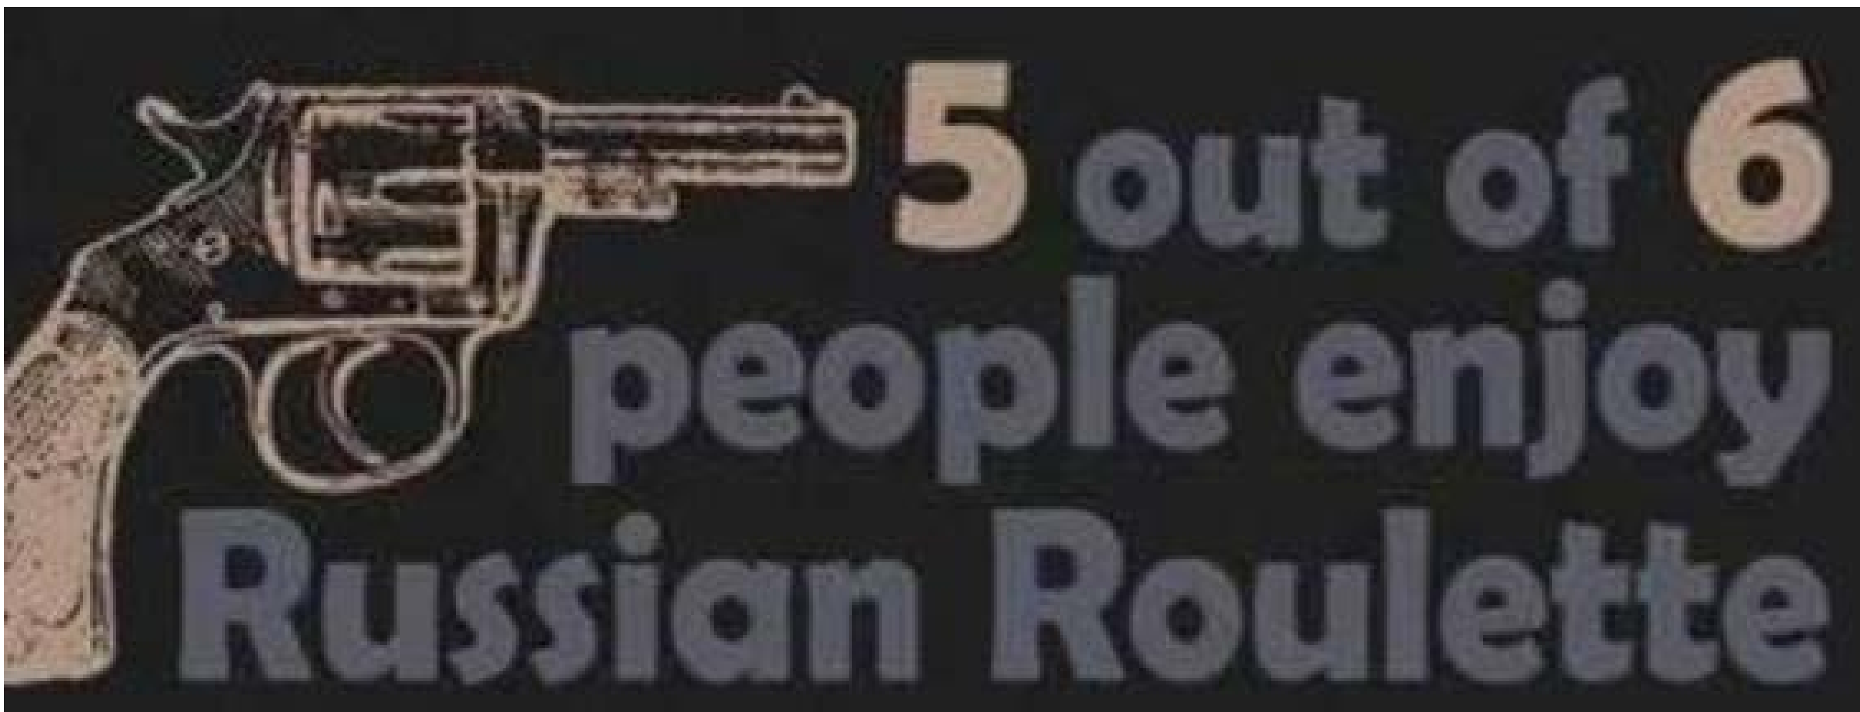

Representative sampling is dangerous

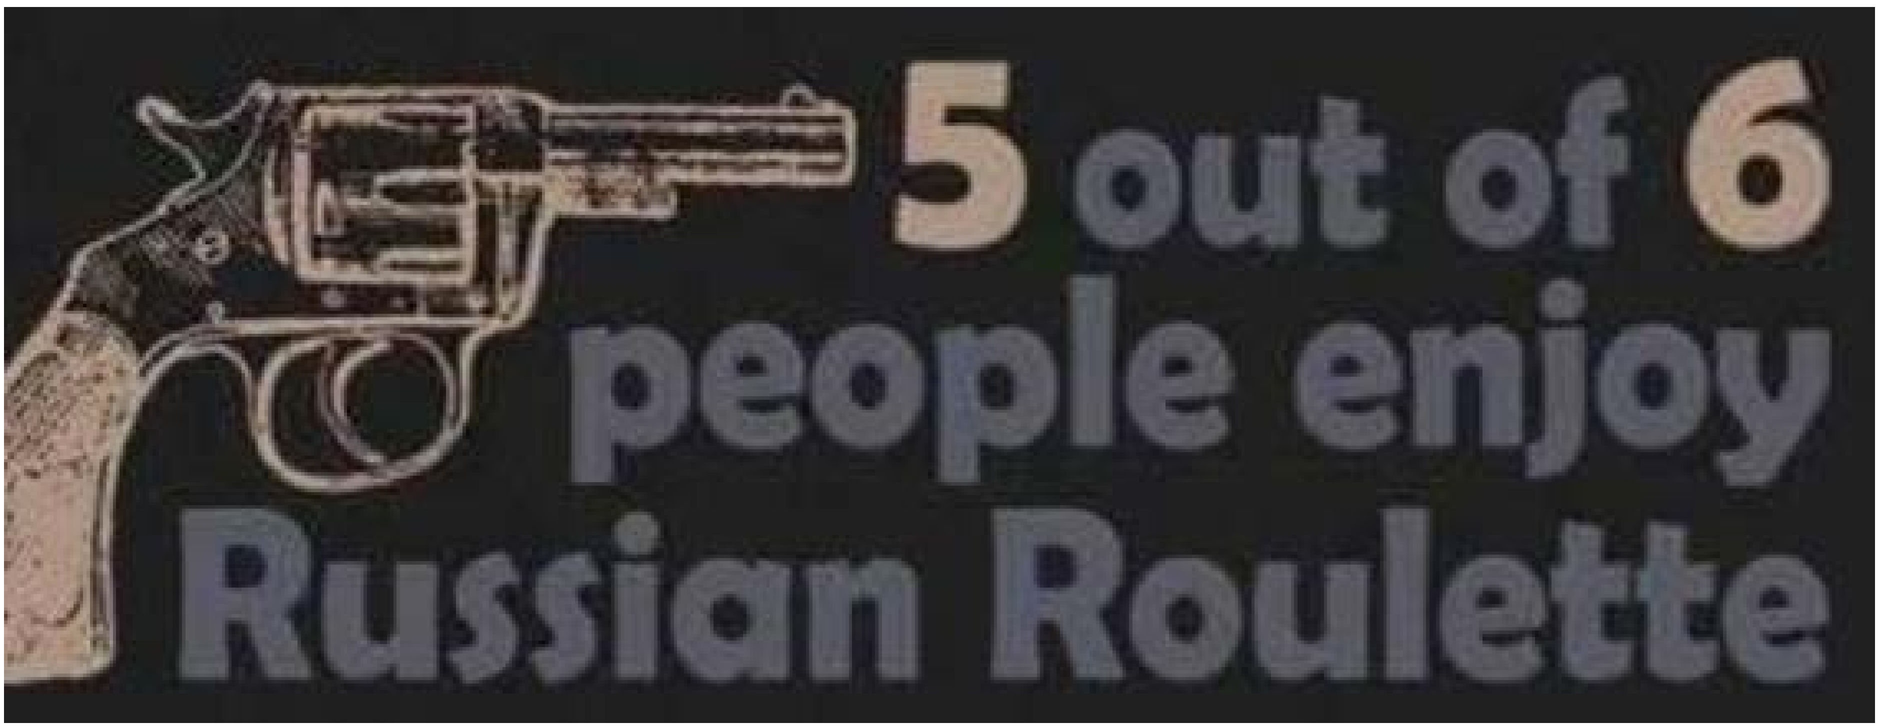

Representative sampling is dangerous

variance

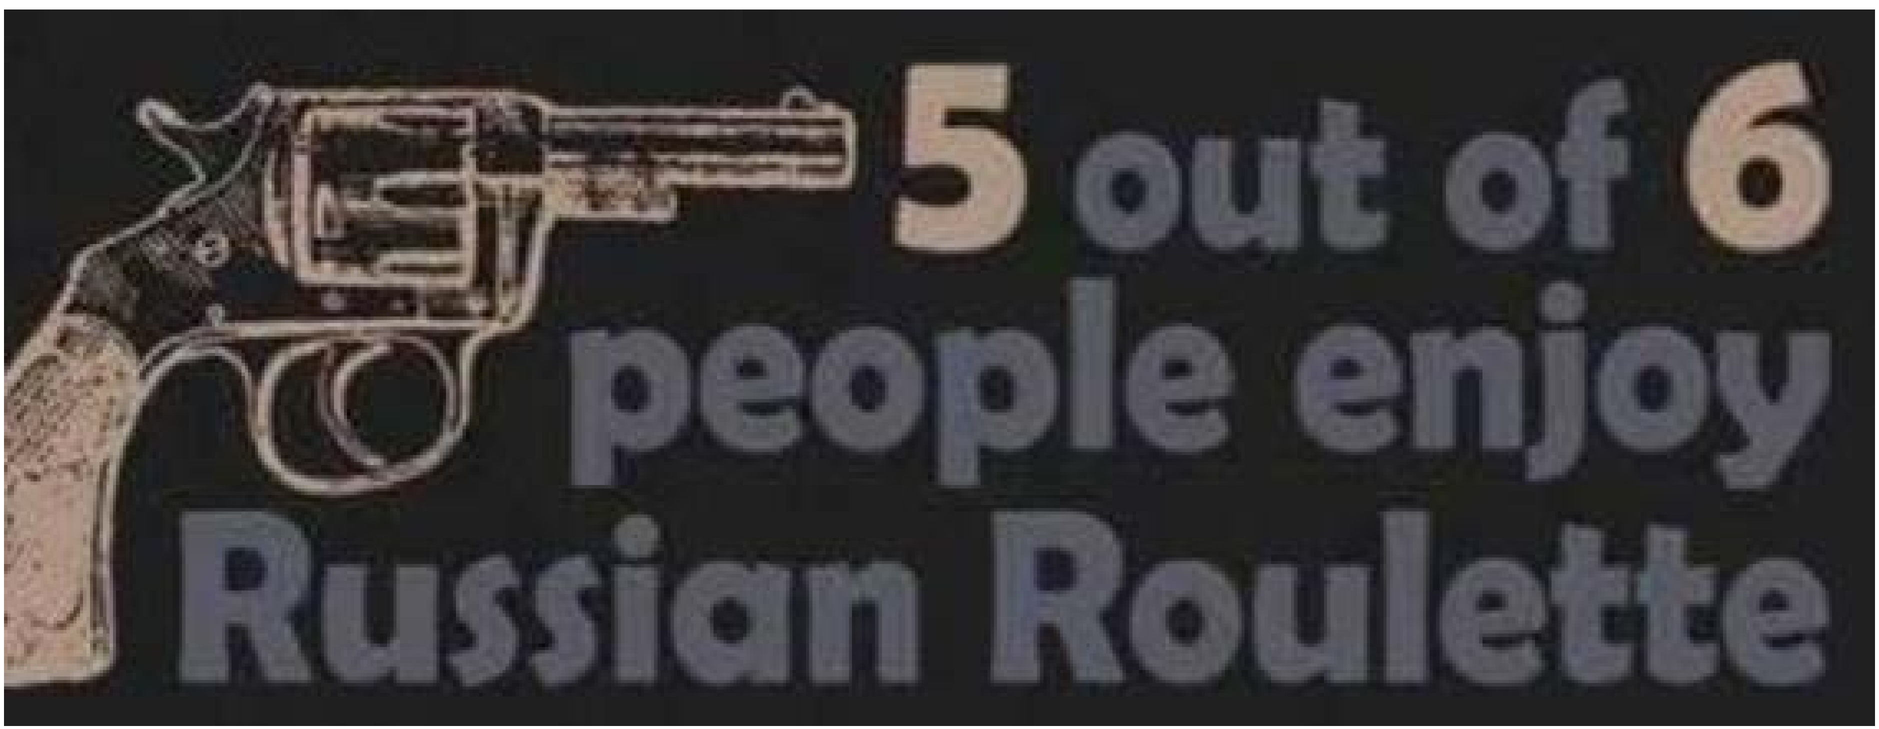

Representative sampling is dangerous

variance

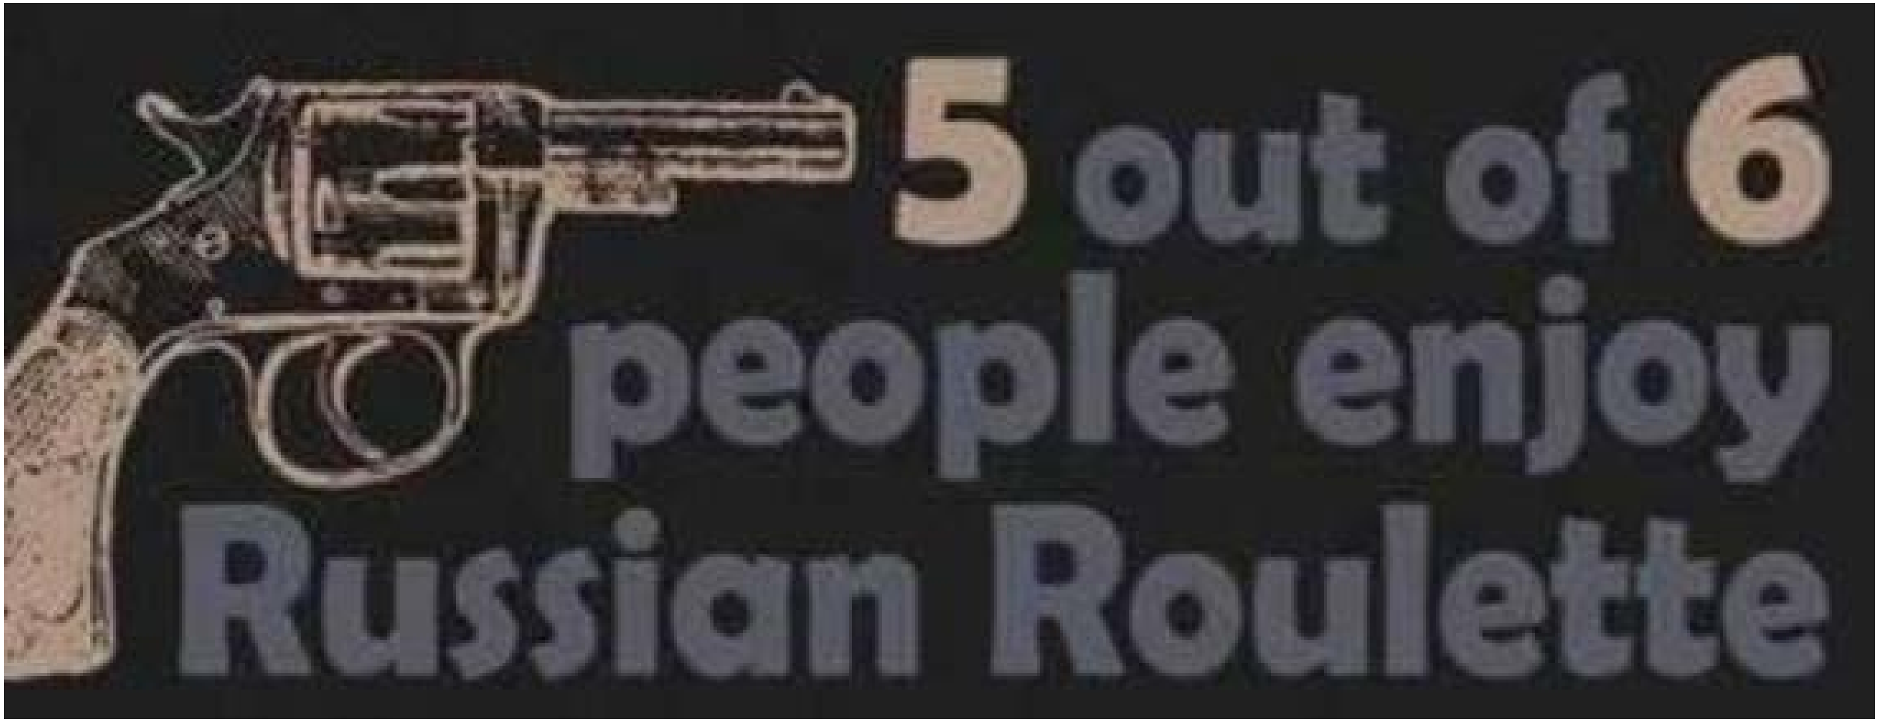

Representative sampling is dangerous

bias

variance

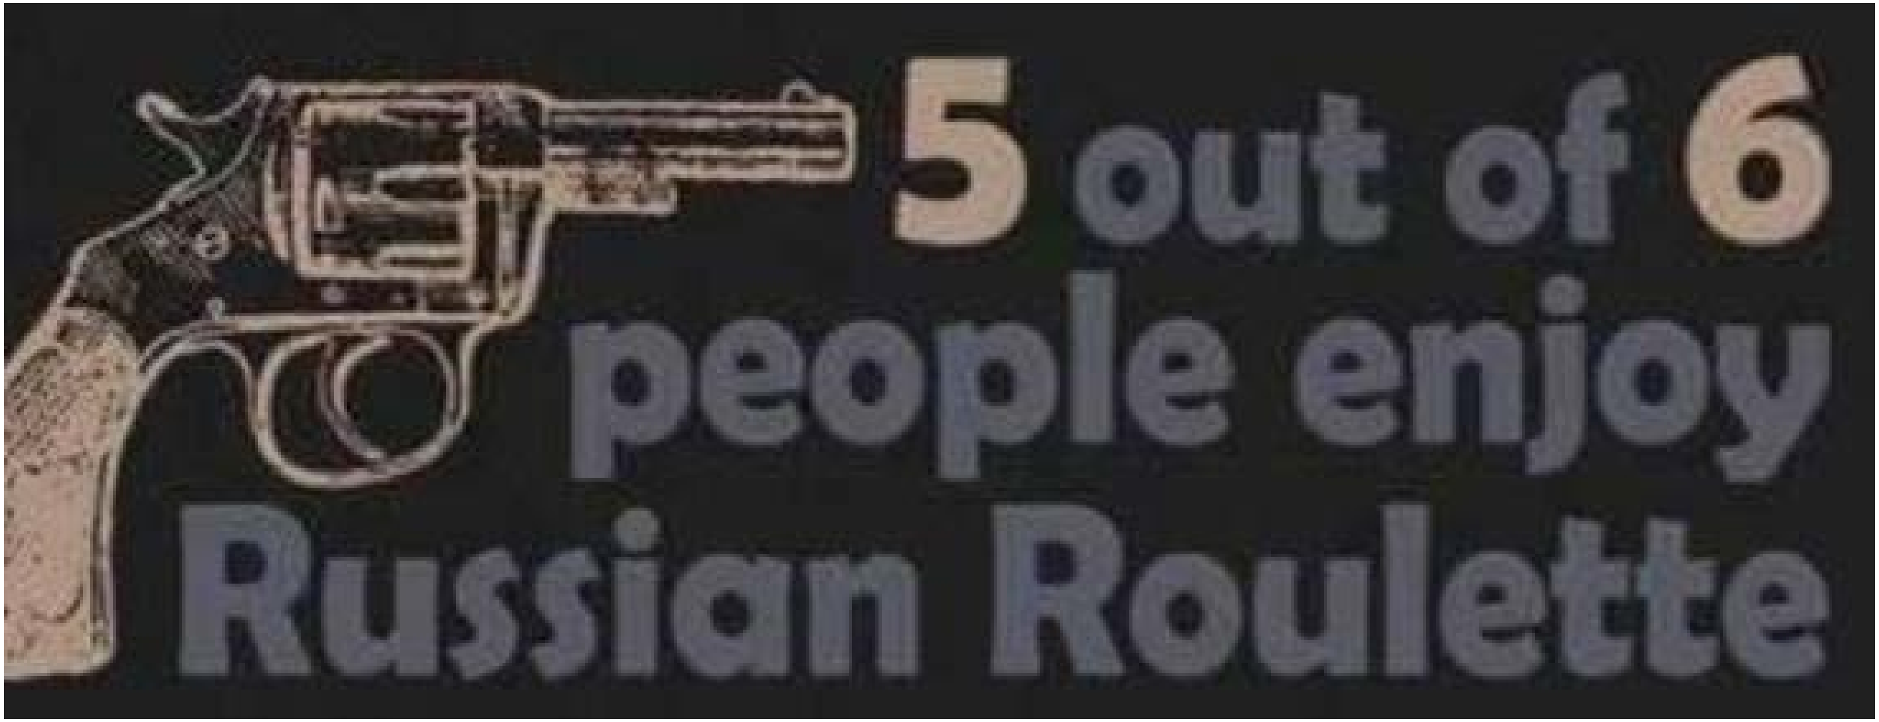

Representative sampling is dangerous

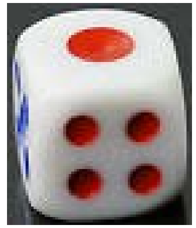

vs.

variance

bias

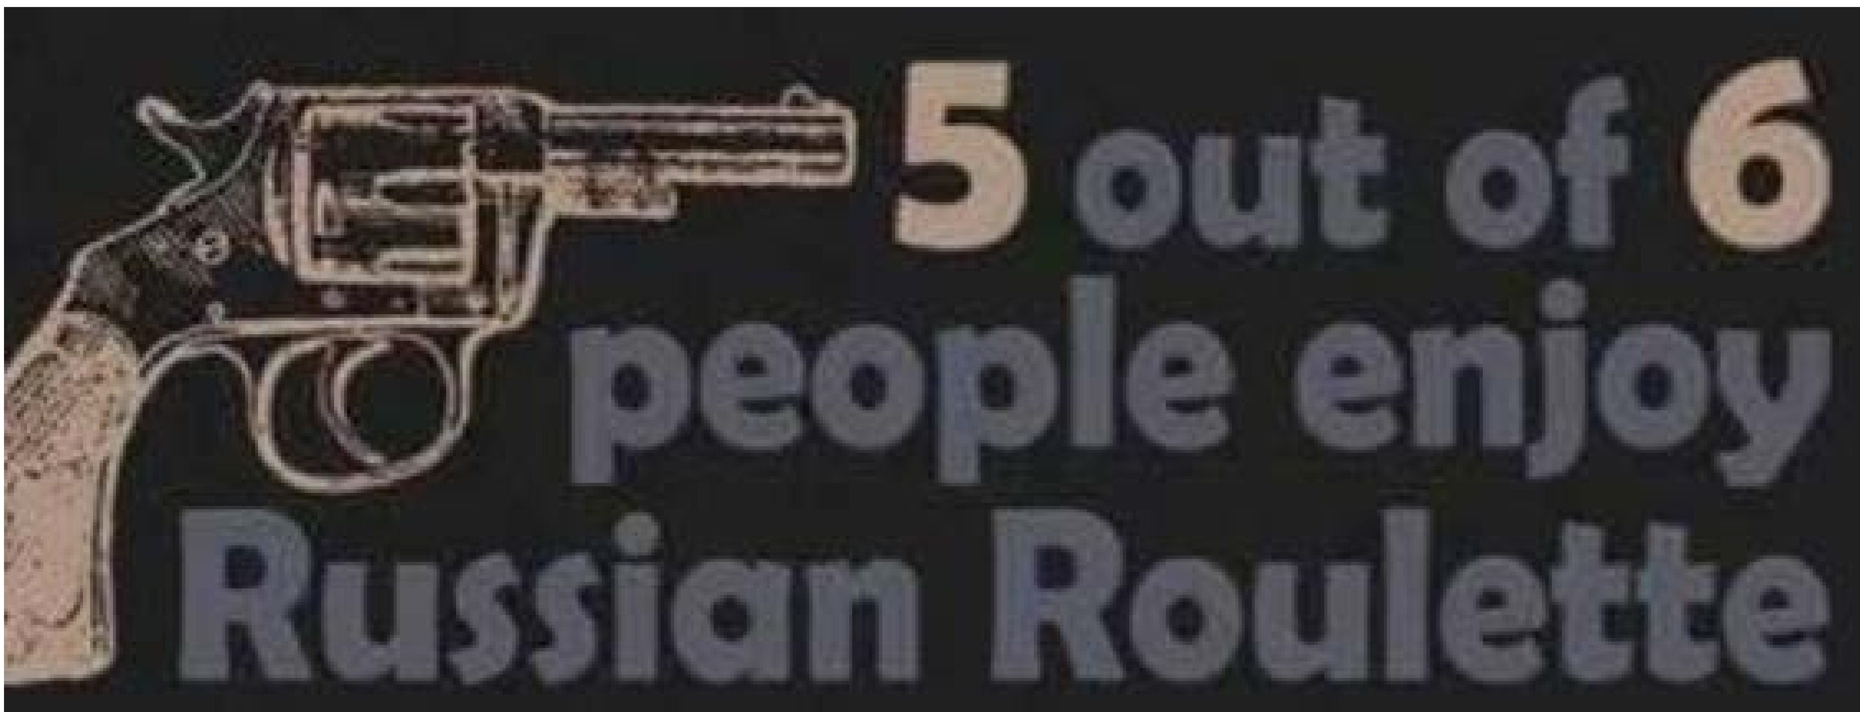

Utility estimation by importance sampling

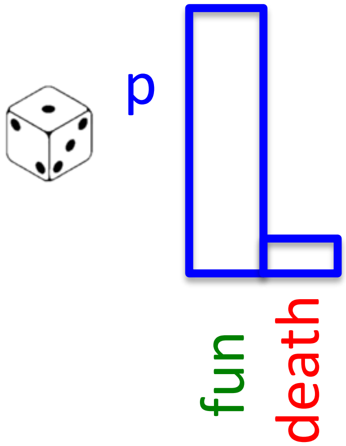

Utility estimation by importance sampling

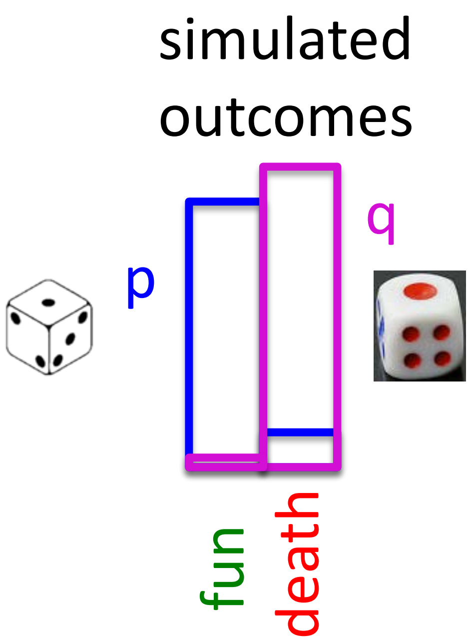

Utility estimation by importance sampling

$o_1, o_2, o_3 \sim q$

simulated outcomes

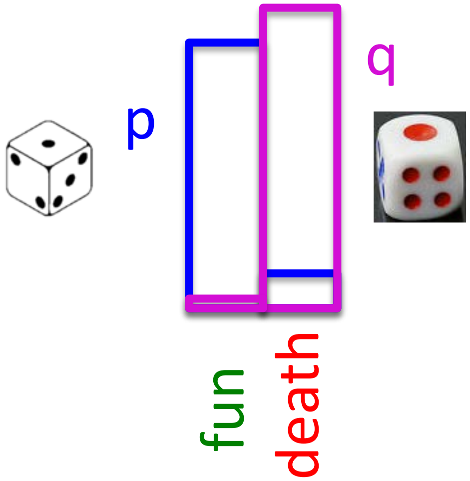

Utility estimation by importance sampling

$o_1, o_2, o_3 \sim q$

simulated outcomes → EU estimates

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

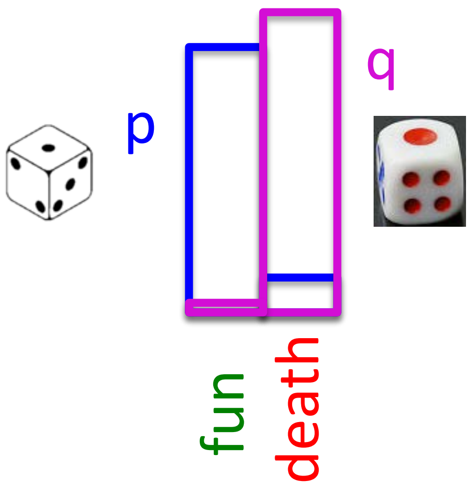

Utility estimation by importance sampling

$o_1, o_2, o_3 \sim q$

simulated outcomes → EU estimates

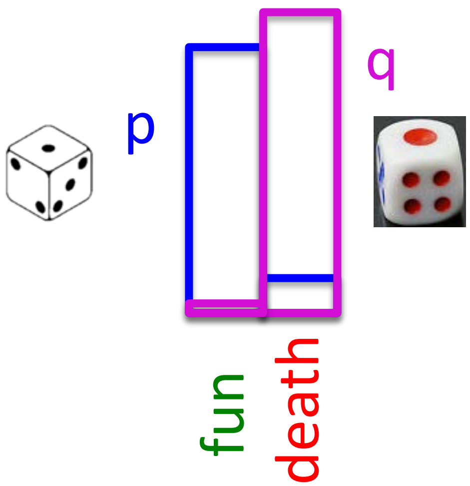

$$\hat{U}^{\mathrm{IS}}_{q,s}(a) = \frac{1}{\sum_{j=1}^{s} w_j} \sum_{j=1}^{s} w_j \cdot u(o_j)$$

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

Utility estimation by importance sampling

$o_1, o_2, o_3 \sim q$

simulated outcomes → EU estimates

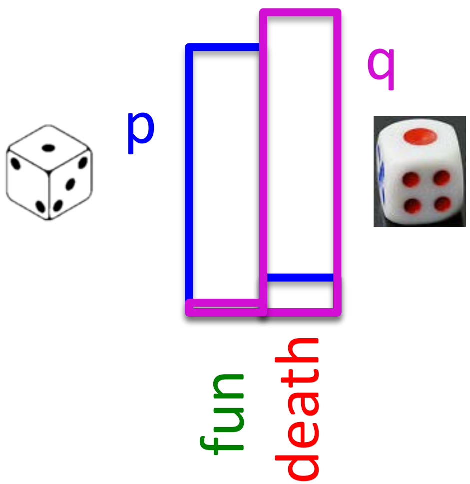

$$\hat{U}^{\mathrm{IS}}_{q,s}(a) = \frac{1}{\sum_{j=1}^{s} w_j} \sum_{j=1}^{s} w_j \cdot u(o_j)$$

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

$$\cdots$$

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

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

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

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

$\ldots$

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

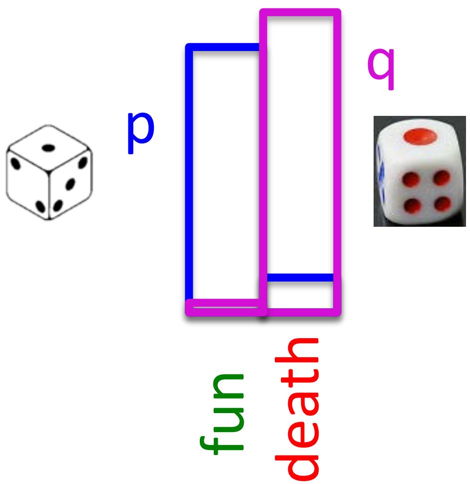

Utility estimation by importance sampling

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

$o_1, o_2, o_3 \sim q$

simulated outcomes → EU estimates → decision

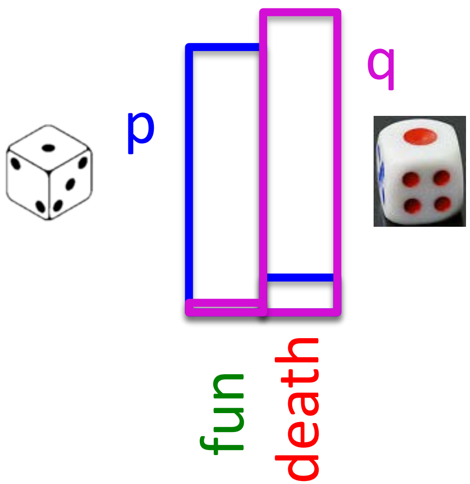

$$\hat{U}^{\mathrm{IS}}_{q,s}(a) = \frac{1}{\sum_{j=1}^{s} w_j} \sum_{j=1}^{s} w_j \cdot u(o_j)$$

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

$$\ldots$$

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

Which distribution should the brain sample from?

Answer: Utility-Weighted Sampling (UWS)

probability

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

simulation frequency

extremity

Answer: Utility-Weighted Sampling (UWS)

probability

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

simulation frequency

extremity

Lieder, Hsu, Griffiths (2014)

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

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

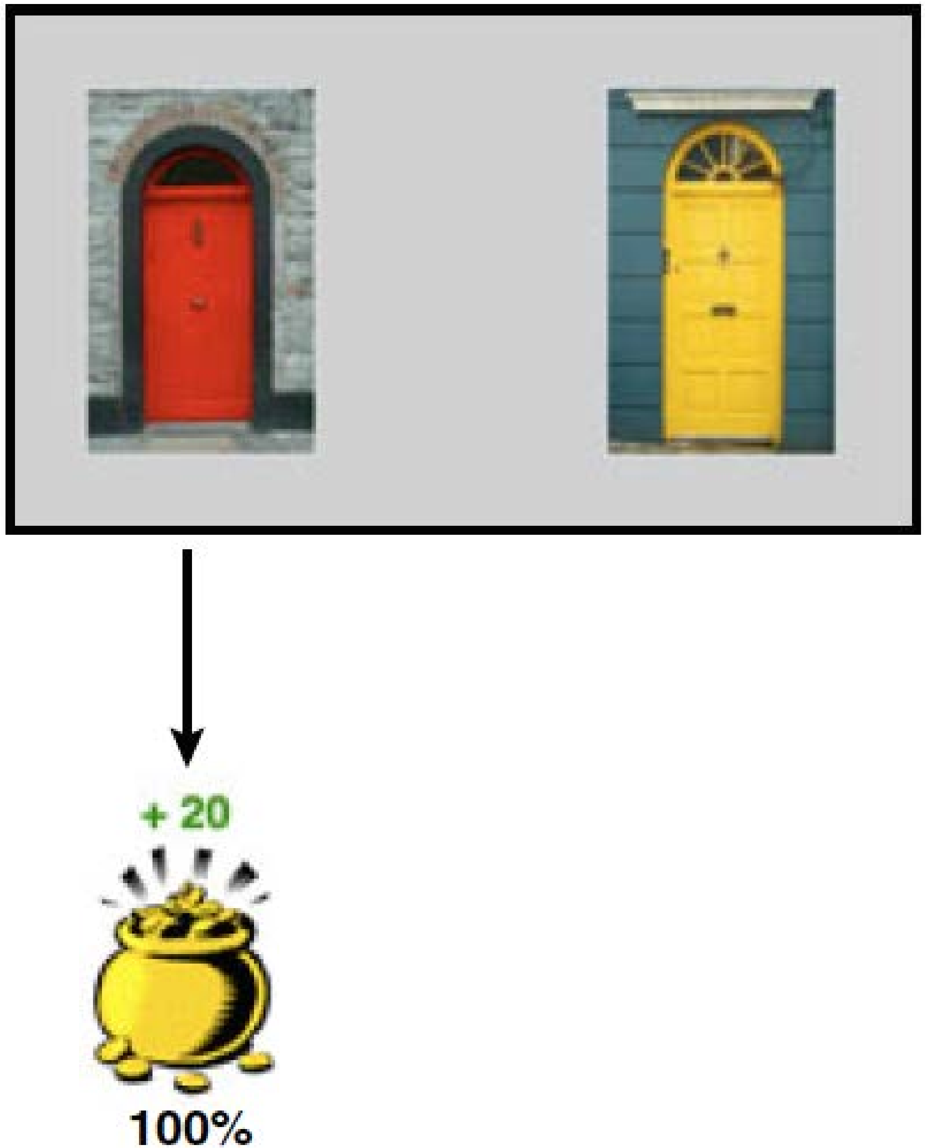

Decisions from Experience (Ludvig, et al., 2014)

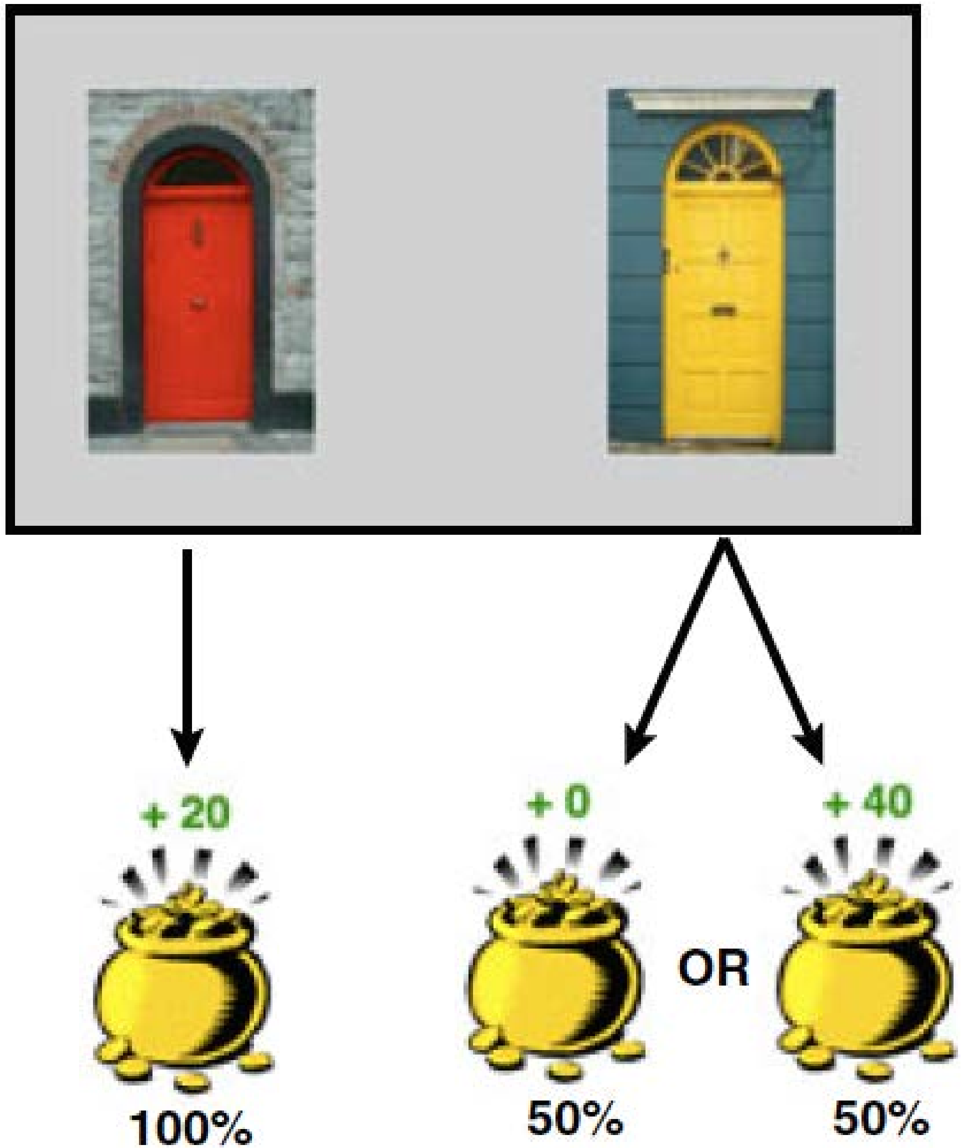

Decisions from Experience (Ludvig, et al., 2008)

Decisions from Experience (Ludvig, et al., 2008)

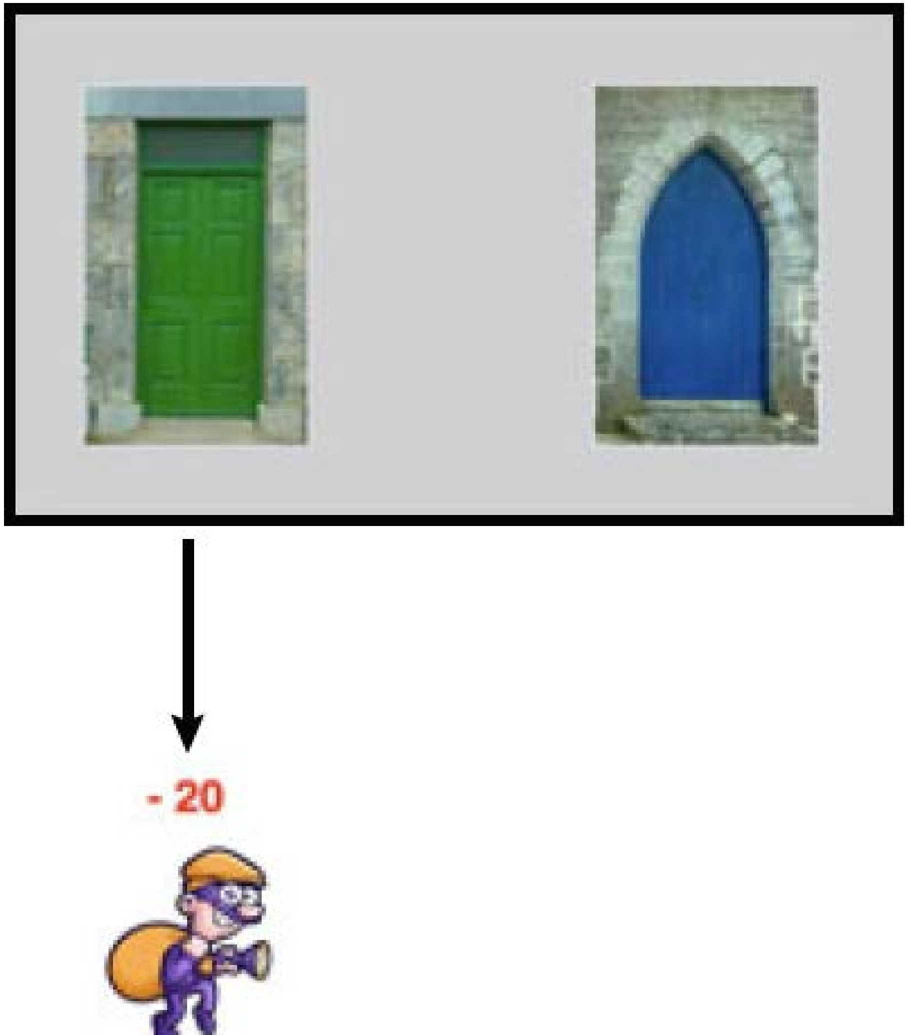

Decisions from Experience (Ludvig, et al., 2008)

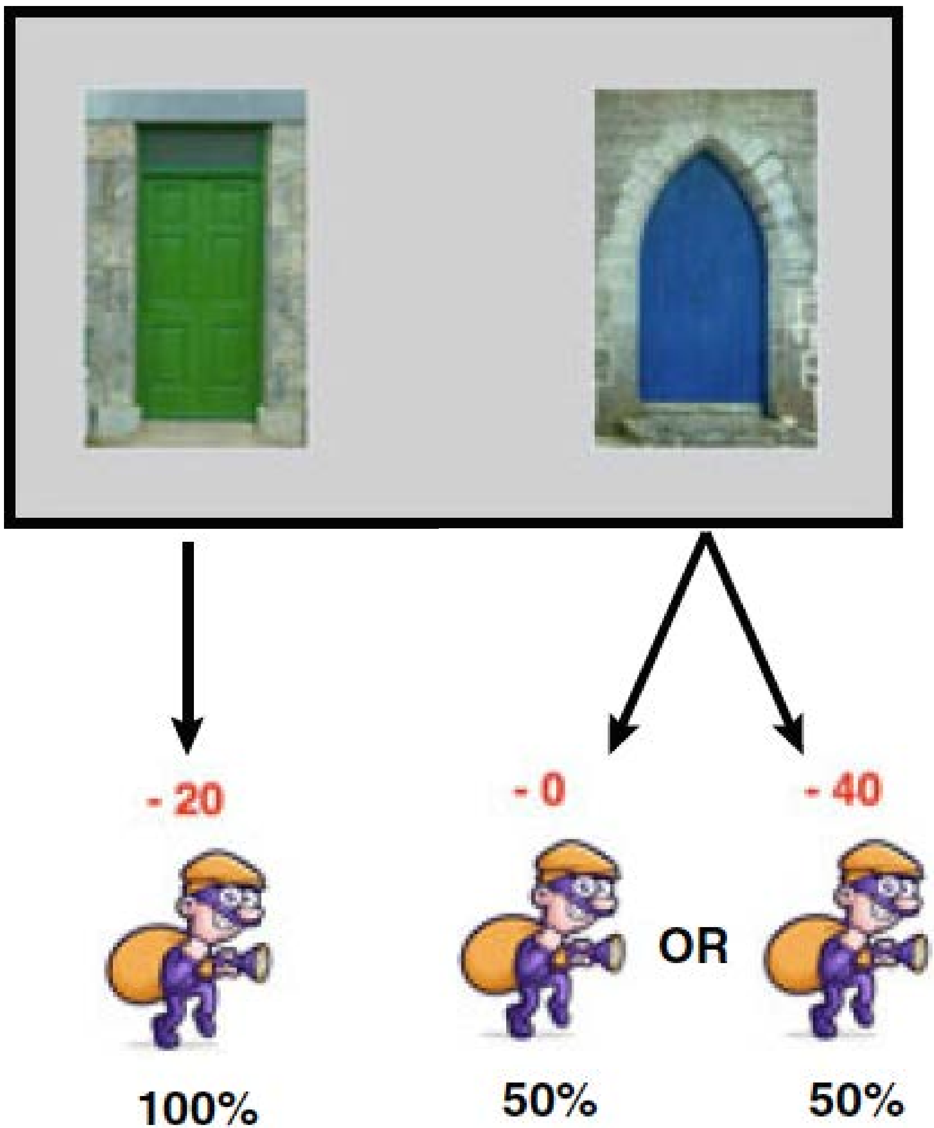

Inconsistent Risk Preferences Emerge from Learning

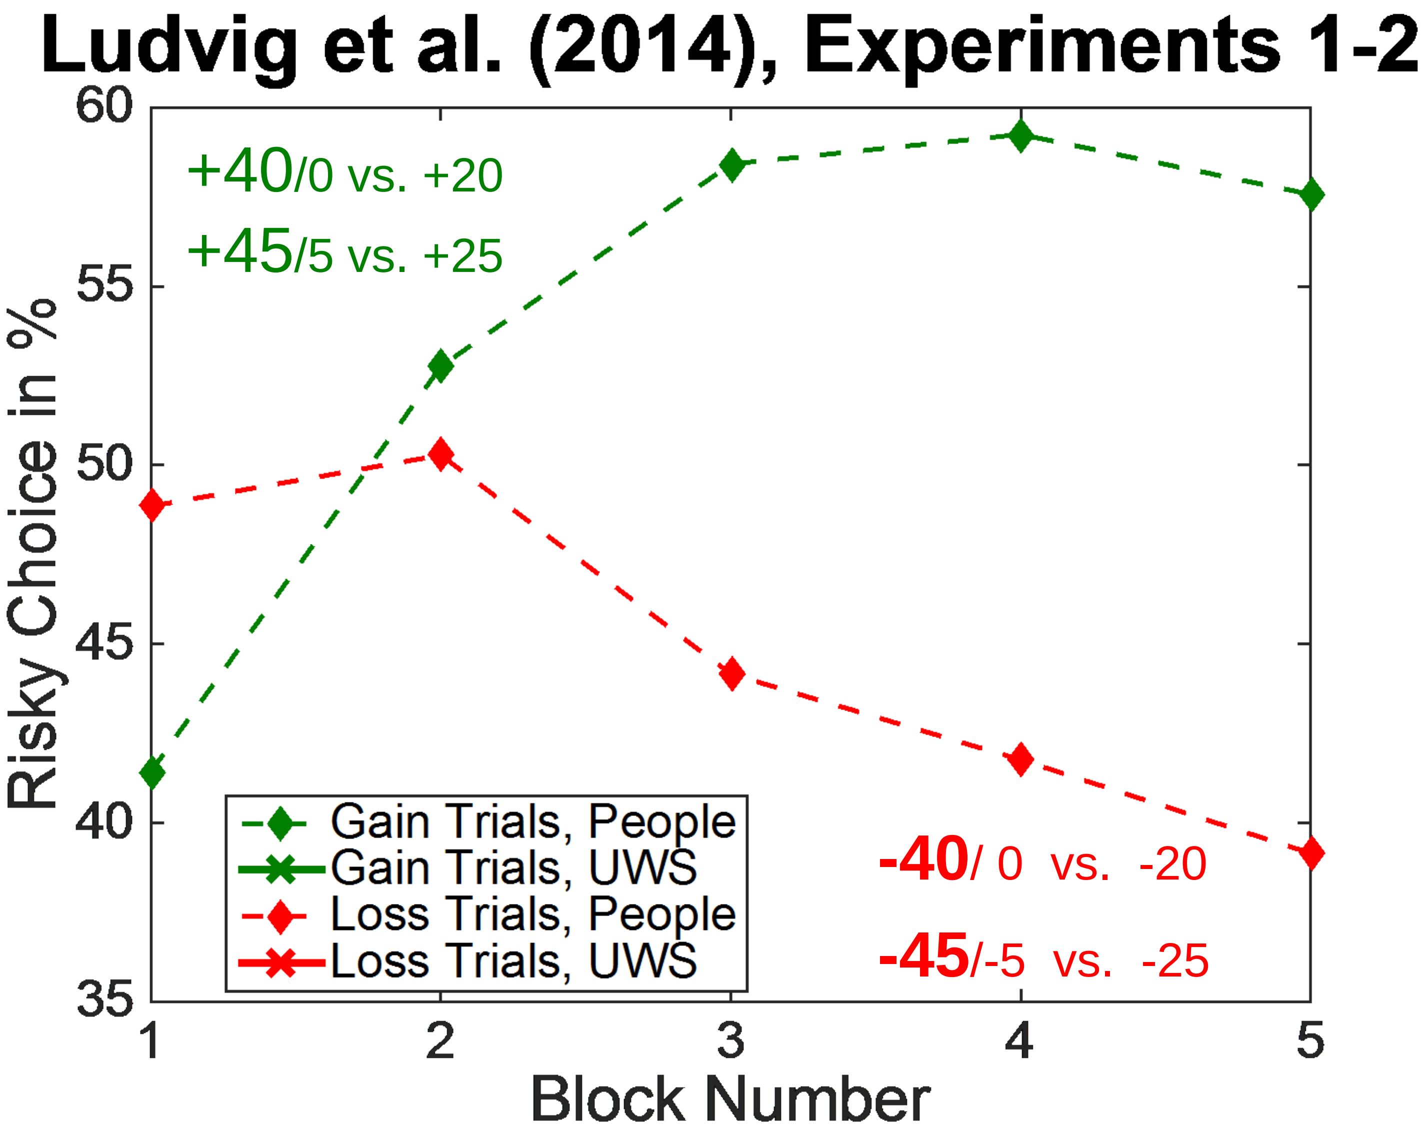

UWS Can Emerge from Reward-Modulated Associative Learning

Actions:  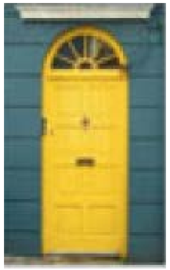

UWS Can Emerge from Reward-Modulated Associative Learning

Actions:

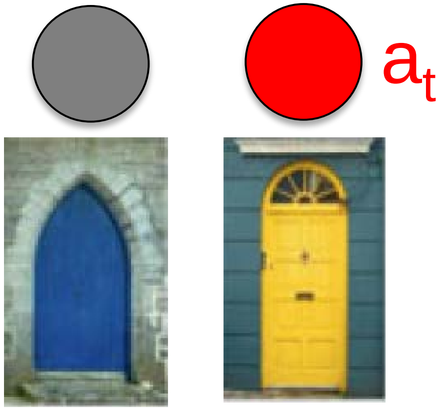

UWS Can Emerge from Reward-Modulated Associative Learning

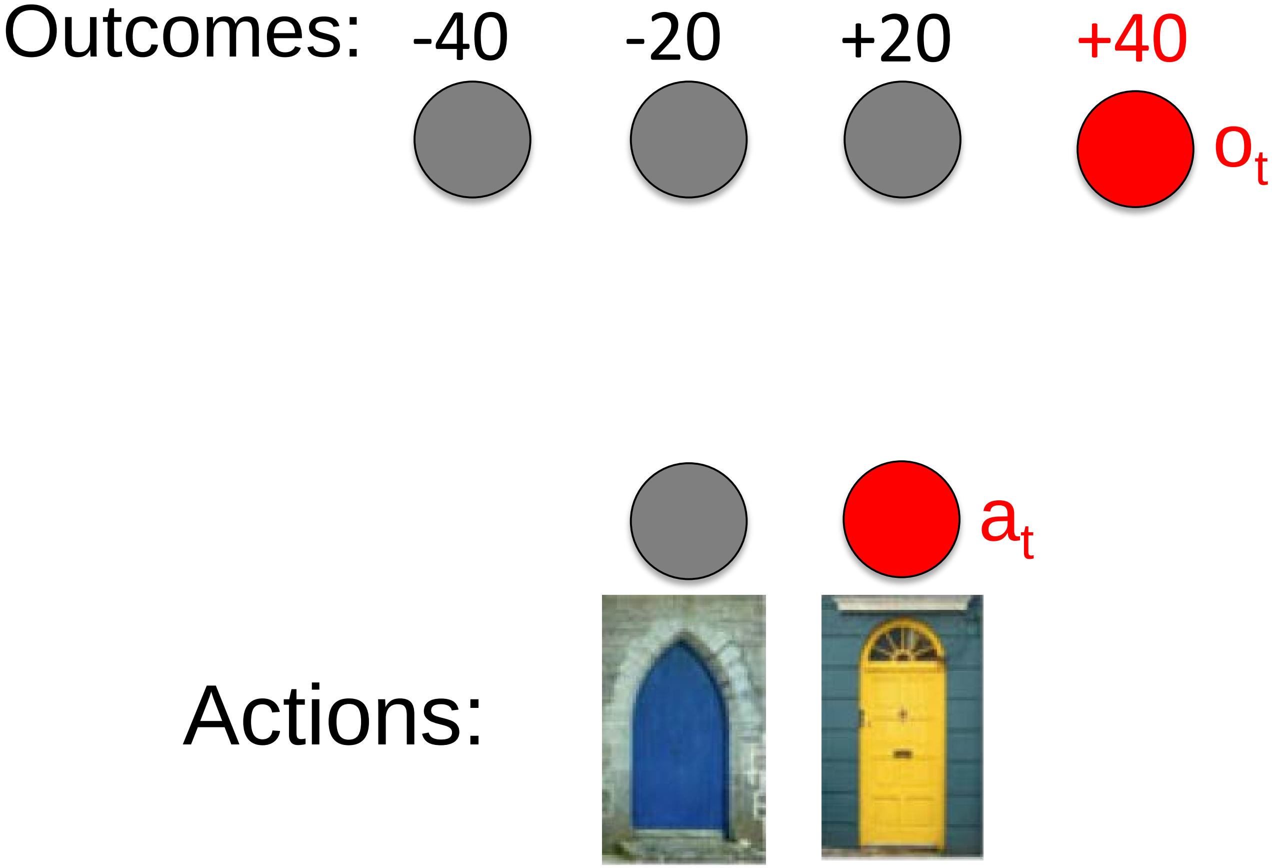

UWS Can Emerge from Reward-Modulated Associative Learning

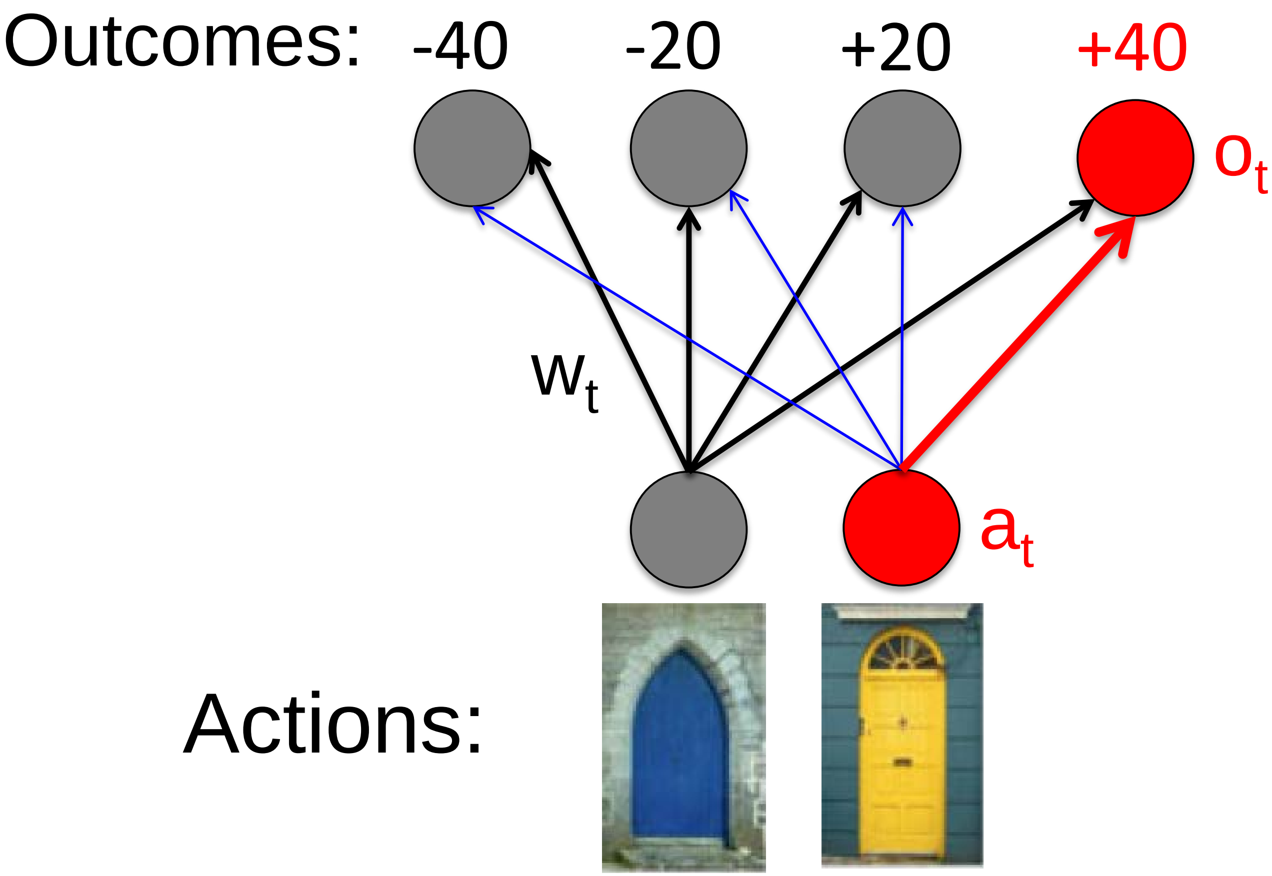

UWS Can Emerge from Reward-Modulated Associative Learning

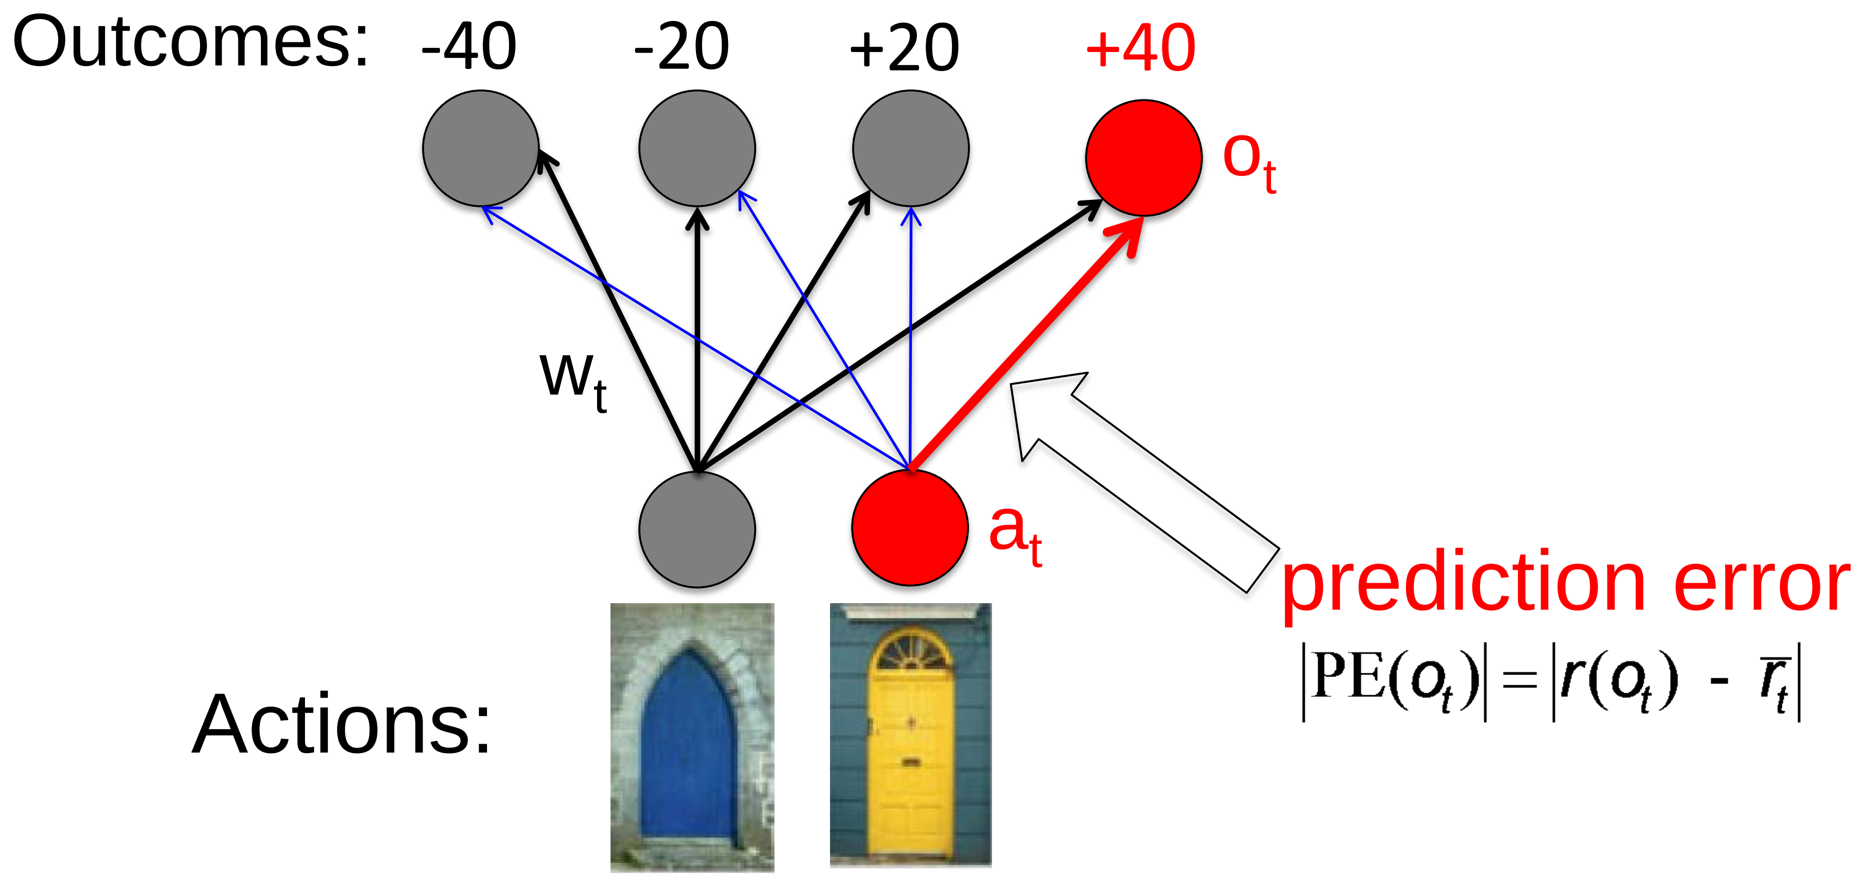

UWS Can Emerge from Reward-Modulated Associative Learning

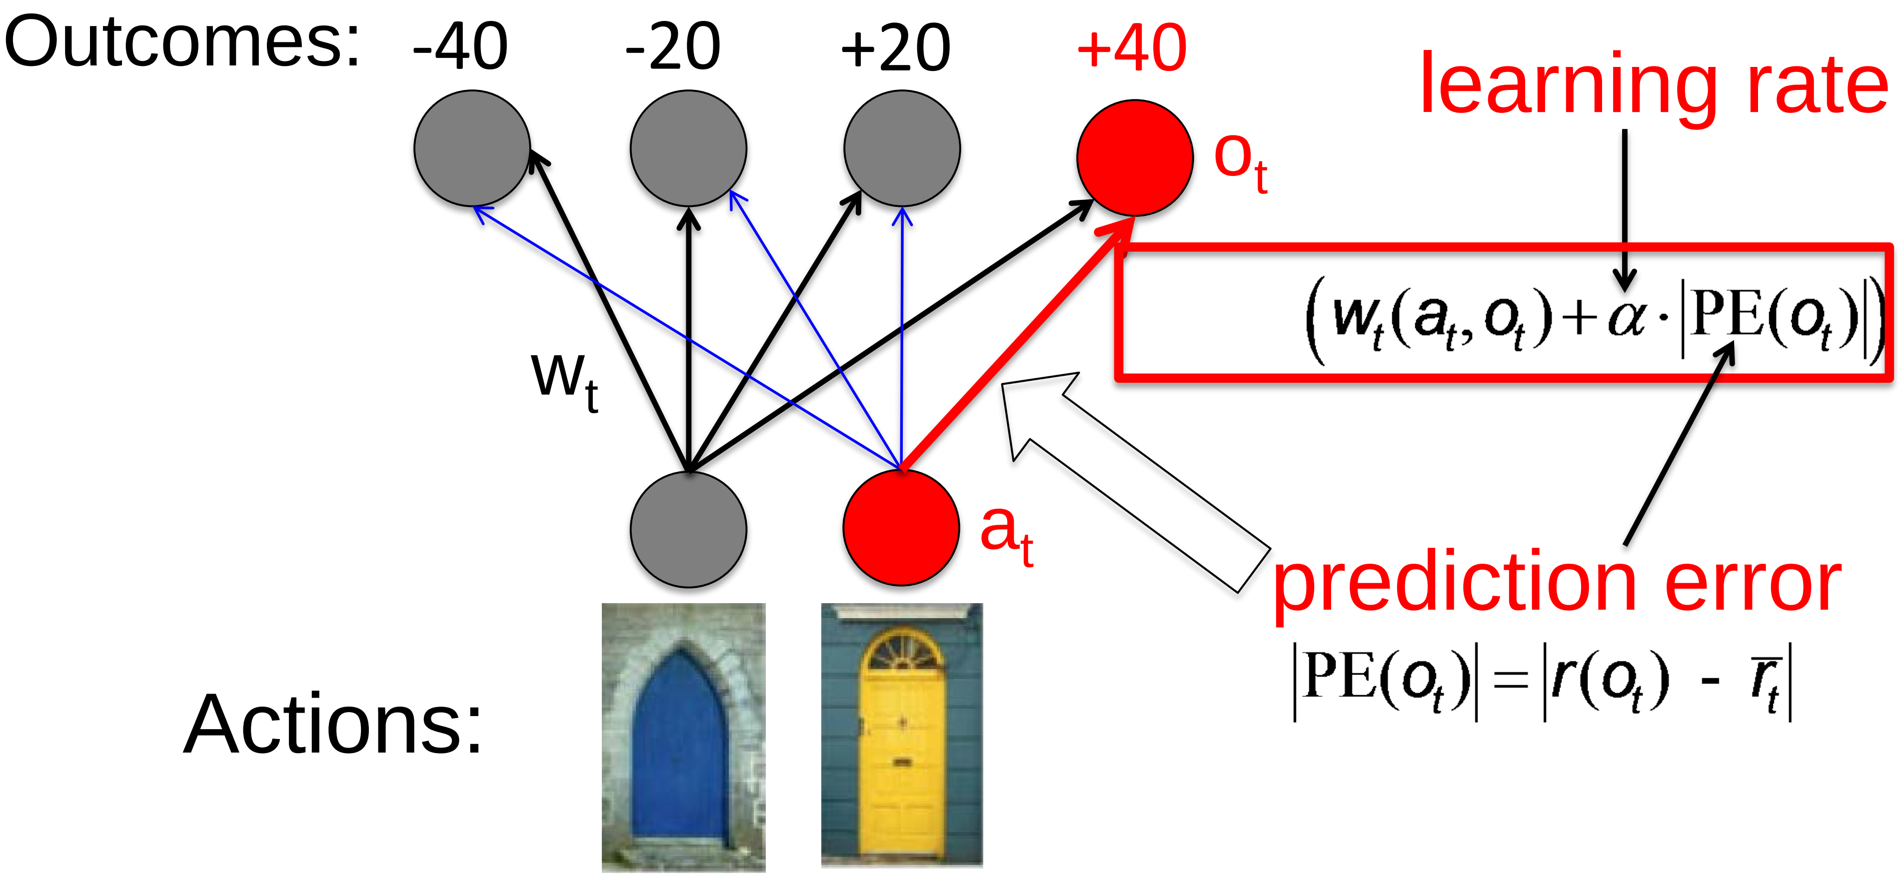

UWS Can Emerge from Reward-Modulated Associative Learning

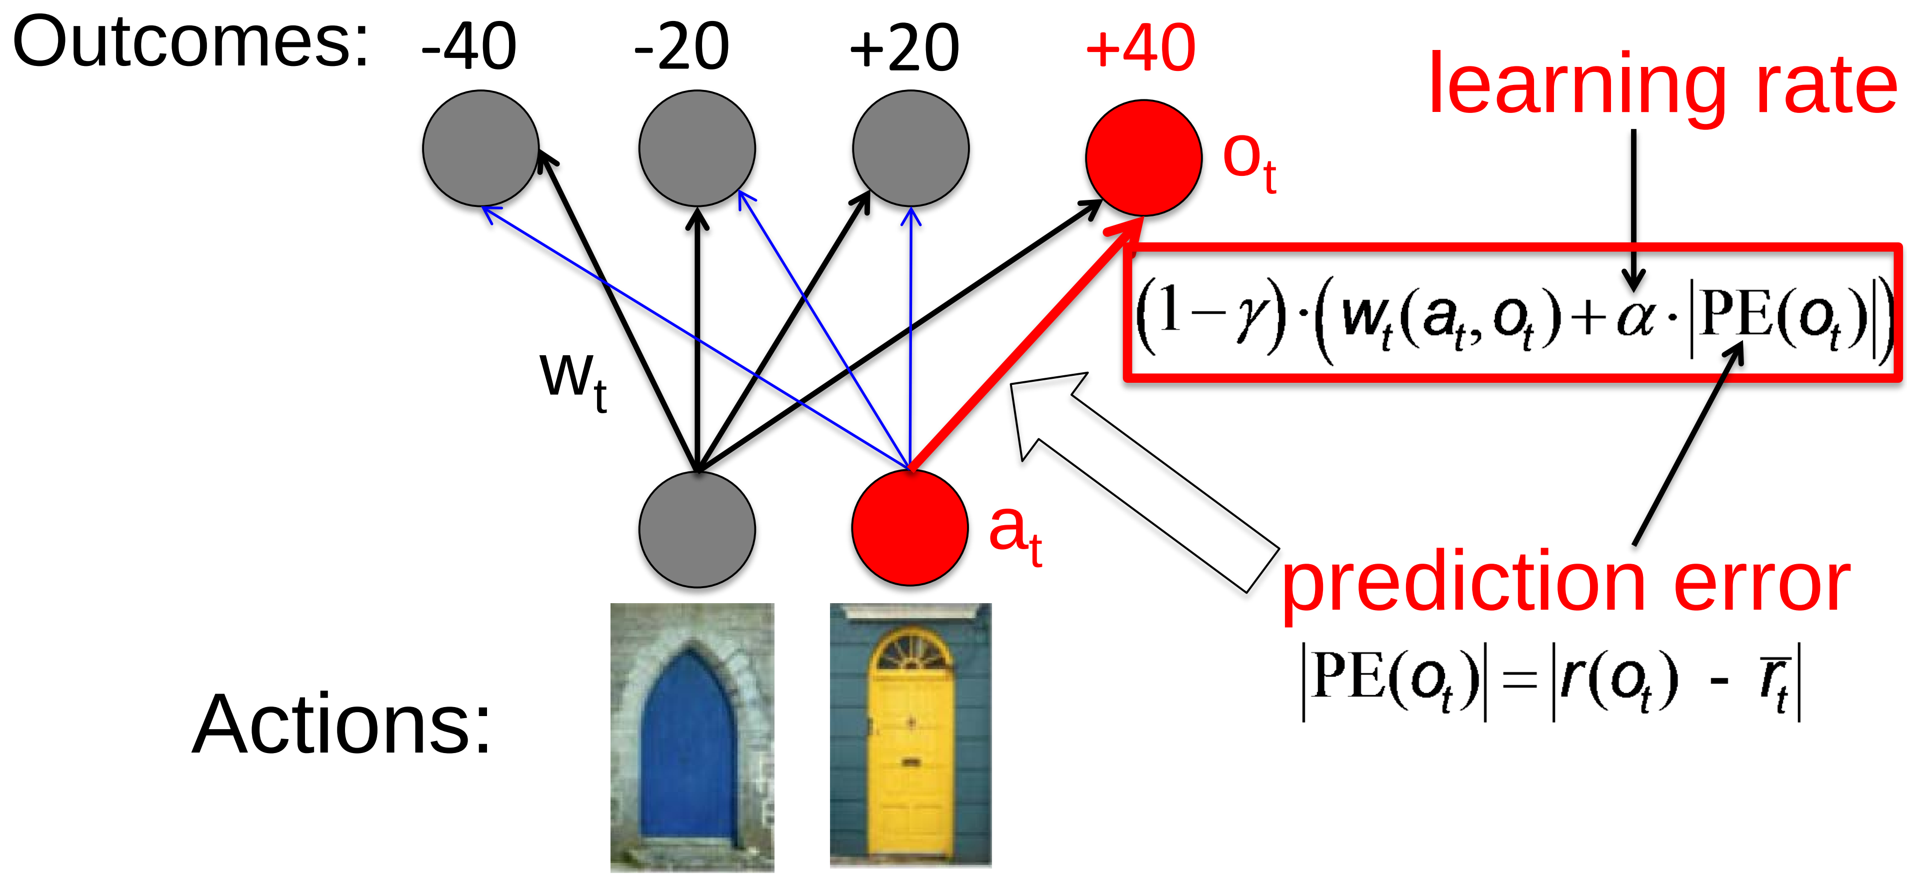

UWS Can Emerge from Reward-Modulated Associative Learning

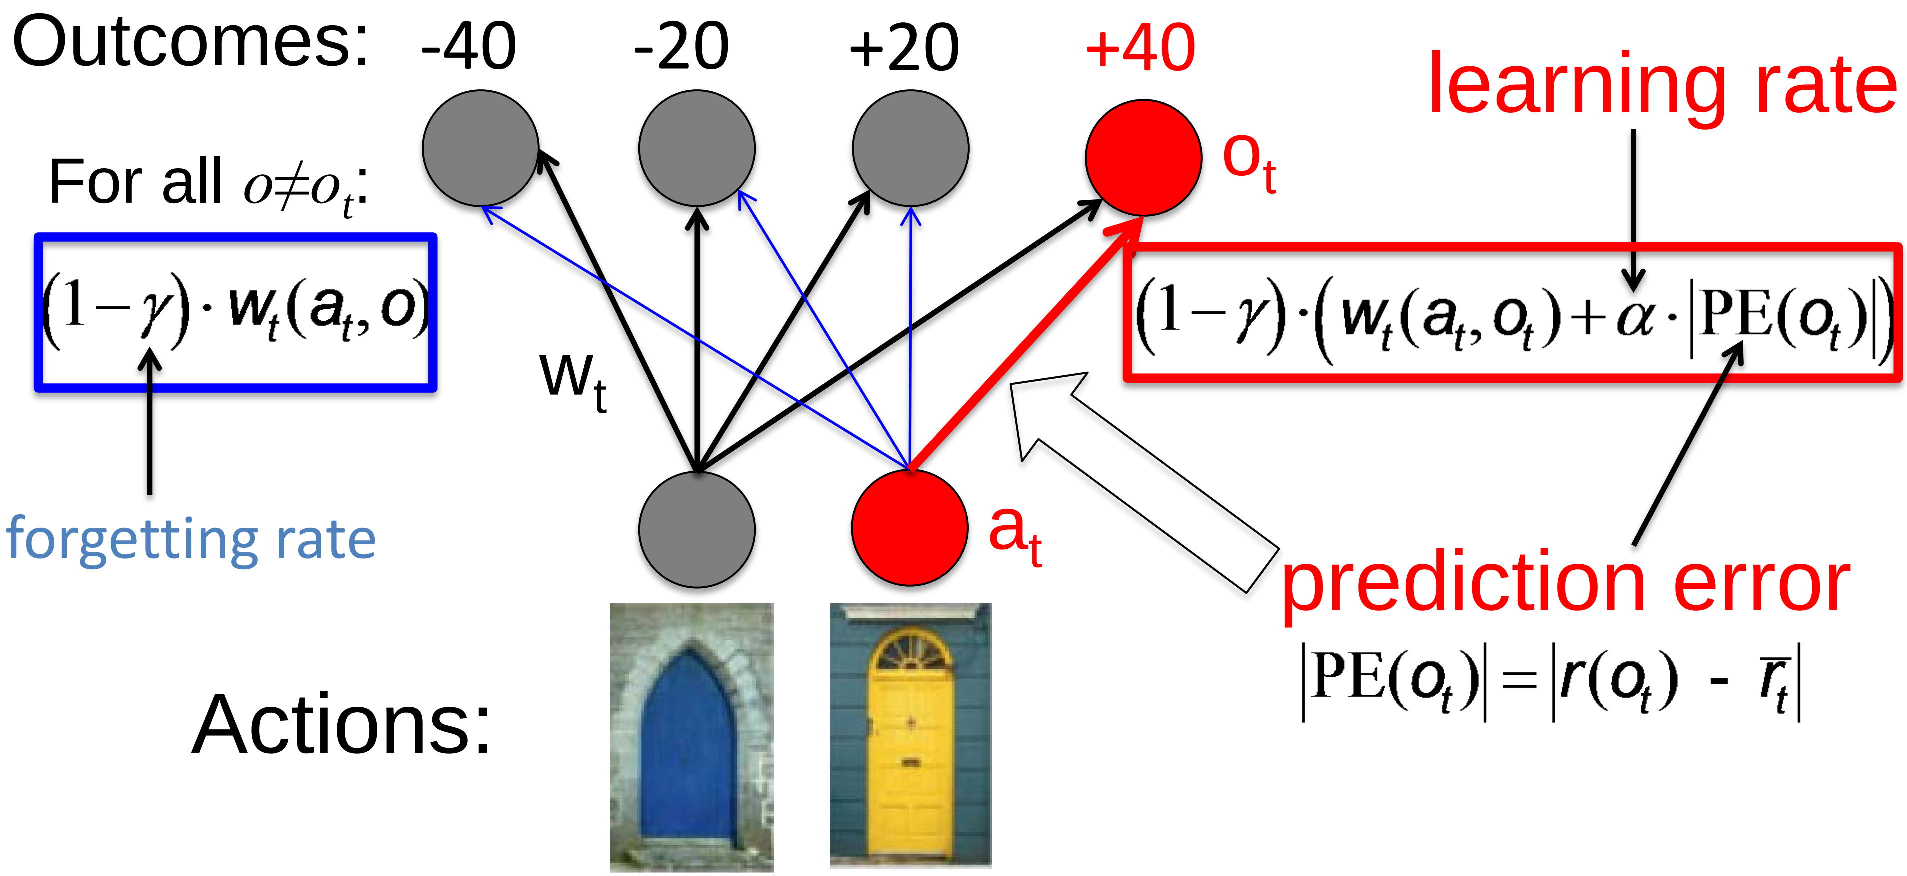

Learning Rule Convergences to Utility-Weighted Sampling

Utility-weighted learning converges to

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

Learning Rule Convergences to Utility-Weighted Sampling

Utility-weighted learning converges to

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

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)

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

Efficient coding (Summerfield & Tsetsos, 2015)

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

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

Efficient coding (Summerfield & Tsetsos, 2015)

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

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

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

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

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.

UWS captures that people learn to overweight extreme outcomes

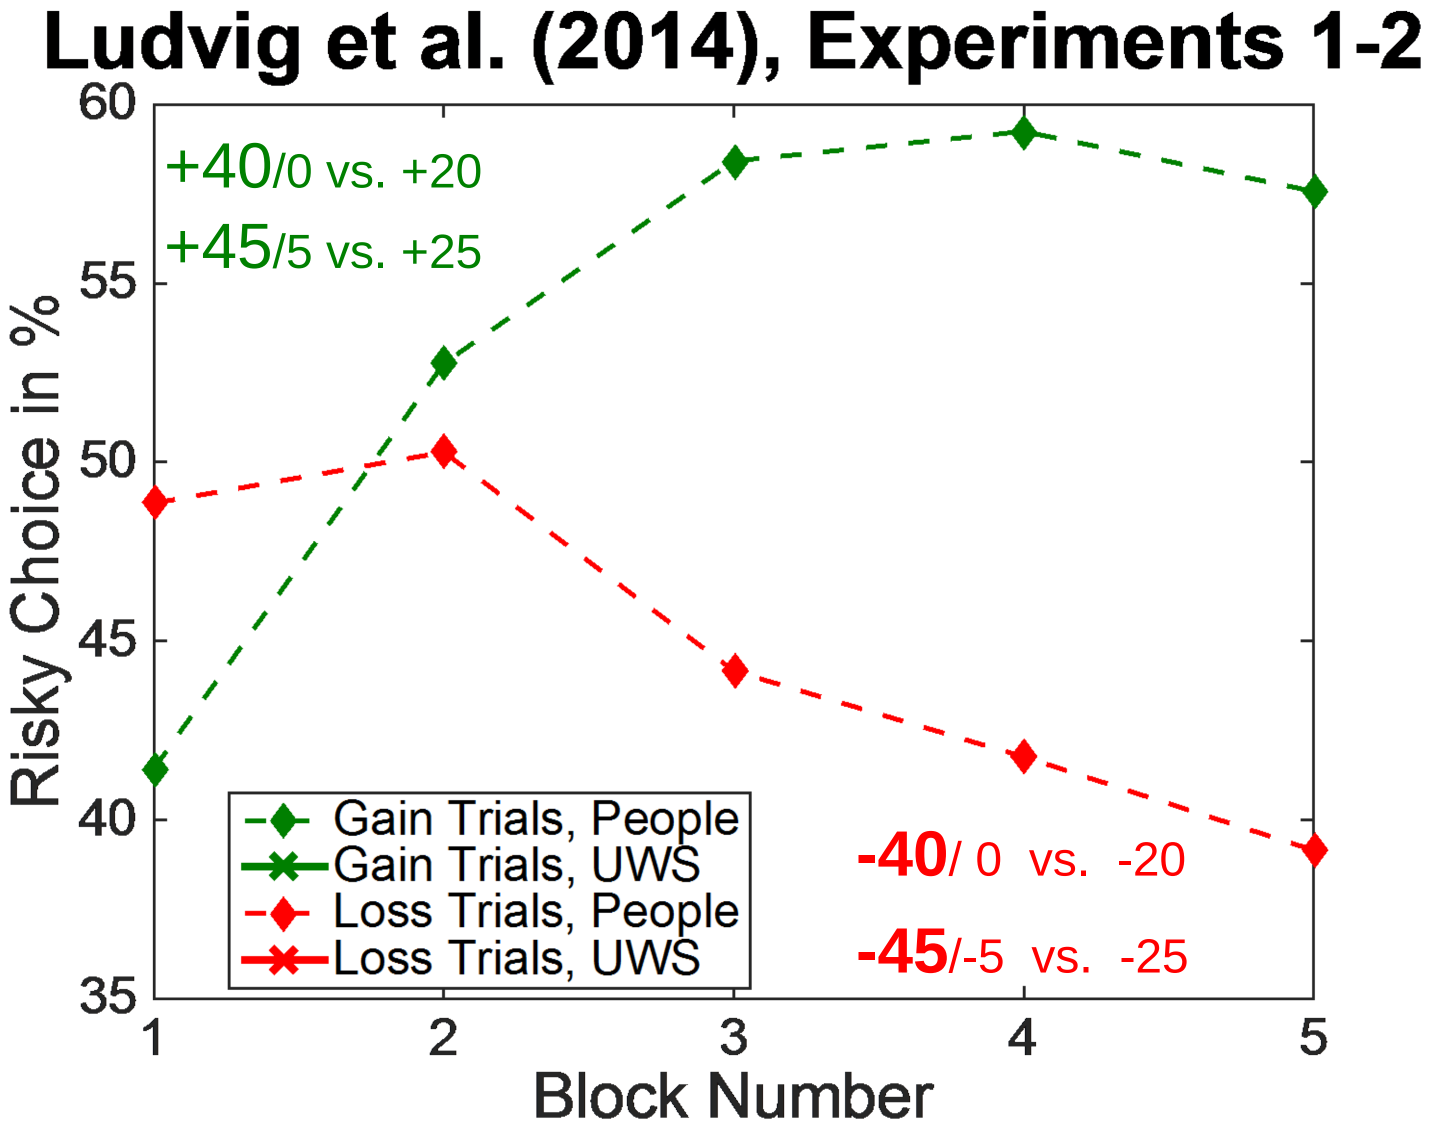

UWS captures that people learn to overweight extreme outcomes

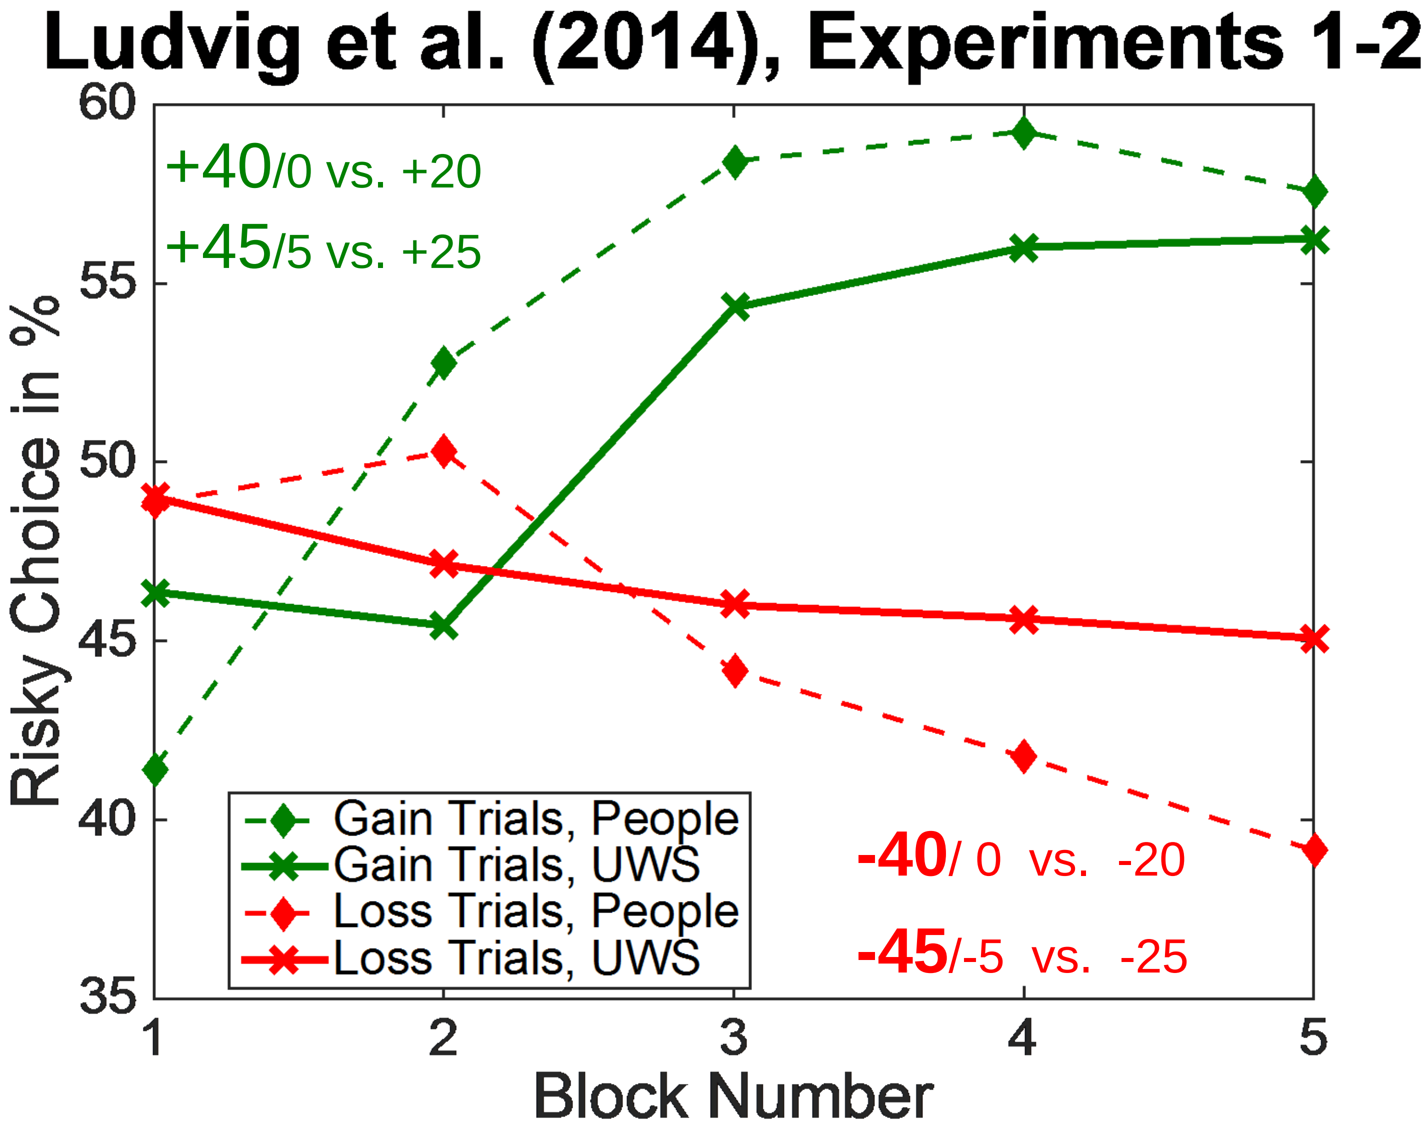

Utility-Weighted Sampling Captures Memory Biases (Madan et al. 2014)

Which outcome
comes to mind first?

Utility-Weighted Sampling Captures Memory Biases (Madan et al. 2014)

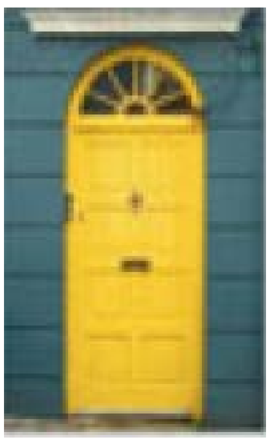

Which outcome comes to mind first?

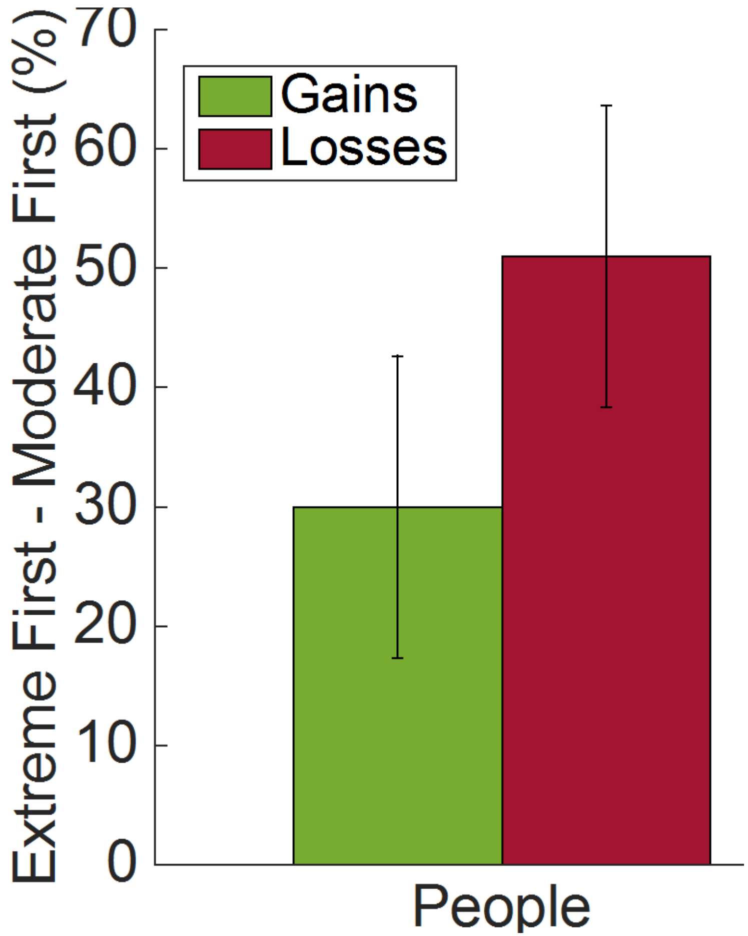

Utility-Weighted Sampling Captures Memory Biases (Madan et al. 2014)

Which outcome comes to mind first?

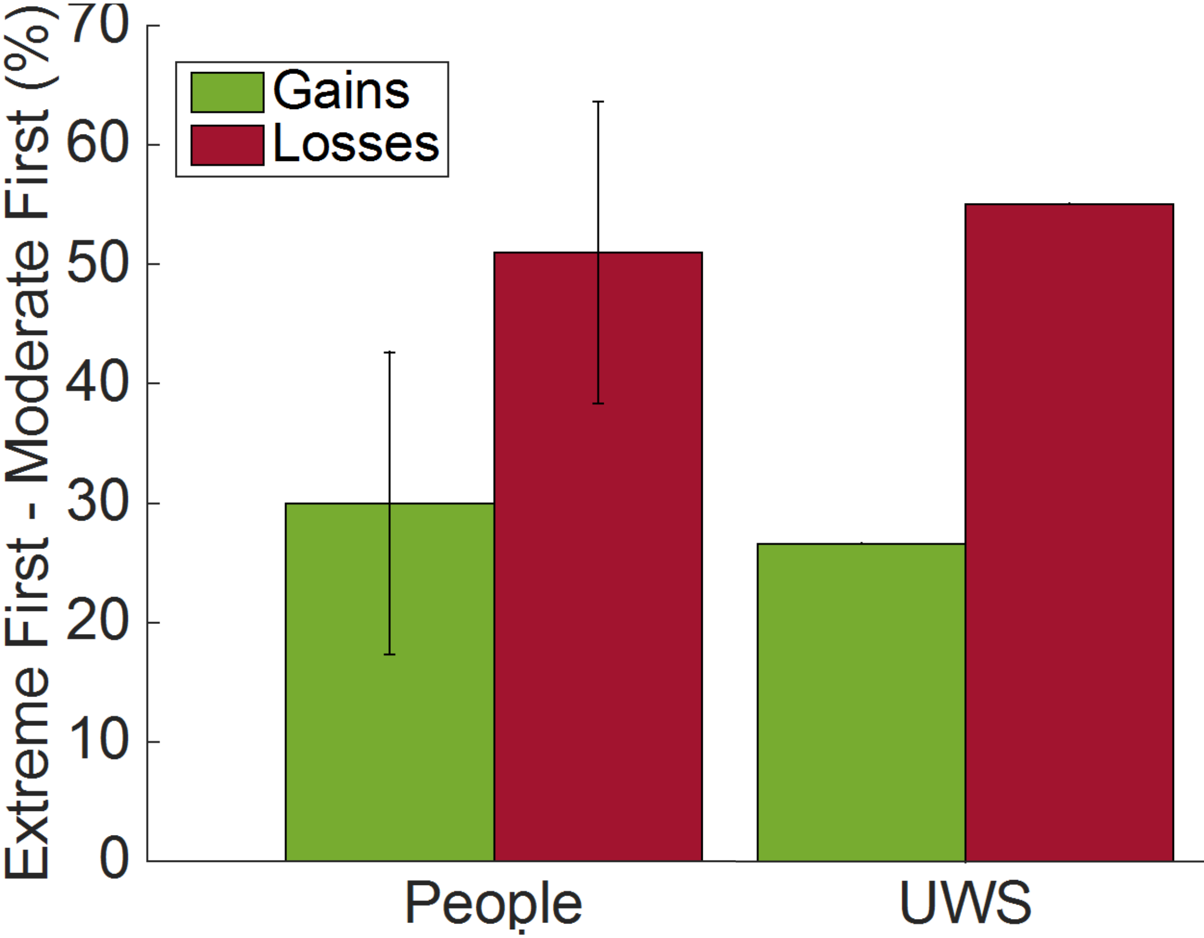

Utility-Weighted Sampling Captures Frequency Estimation Bias (Madan et al. 2014)

How often did this door lead to each outcome?

+40: ___ %
+20: ___ %
0: ___ %

Utility-Weighted Sampling Captures Frequency Estimation Bias (Madan et al. 2014)

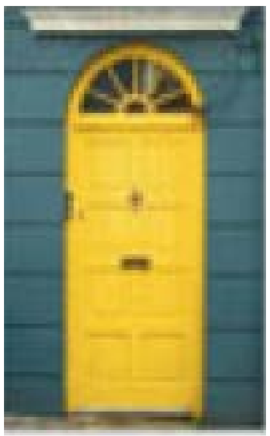

How often did this door lead to each outcome?

+40: ___ %
+20: ___ %
0: ___ %

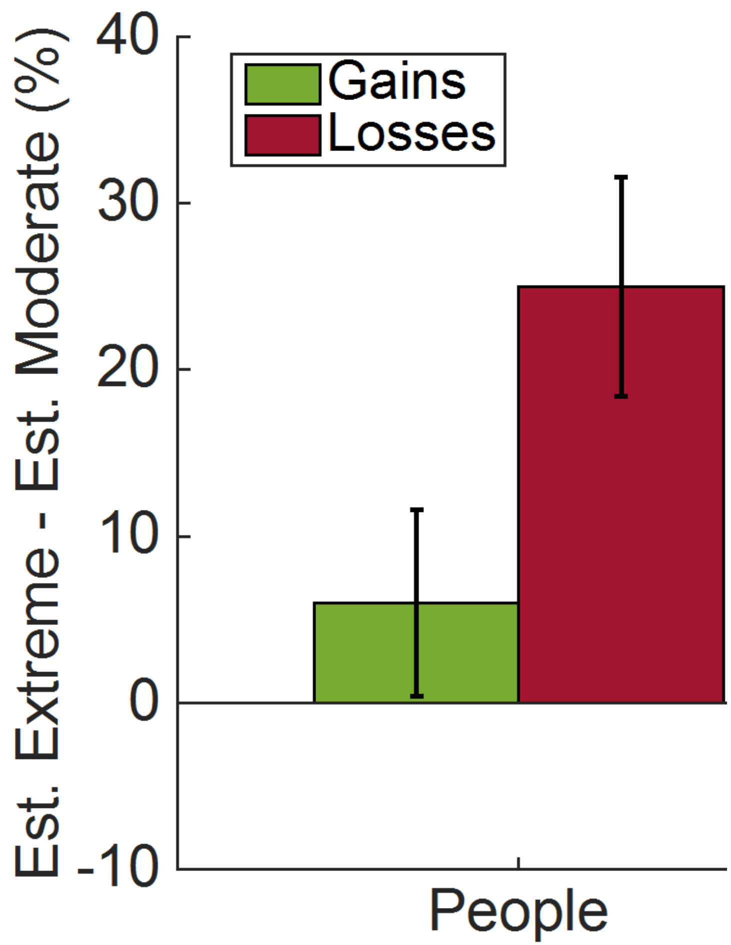

Utility-Weighted Sampling Captures Frequency Estimation Bias (Madan et al. 2014)

How often did this door lead to each outcome?

+40: ____ %
+20: ____ %
0: ___ %

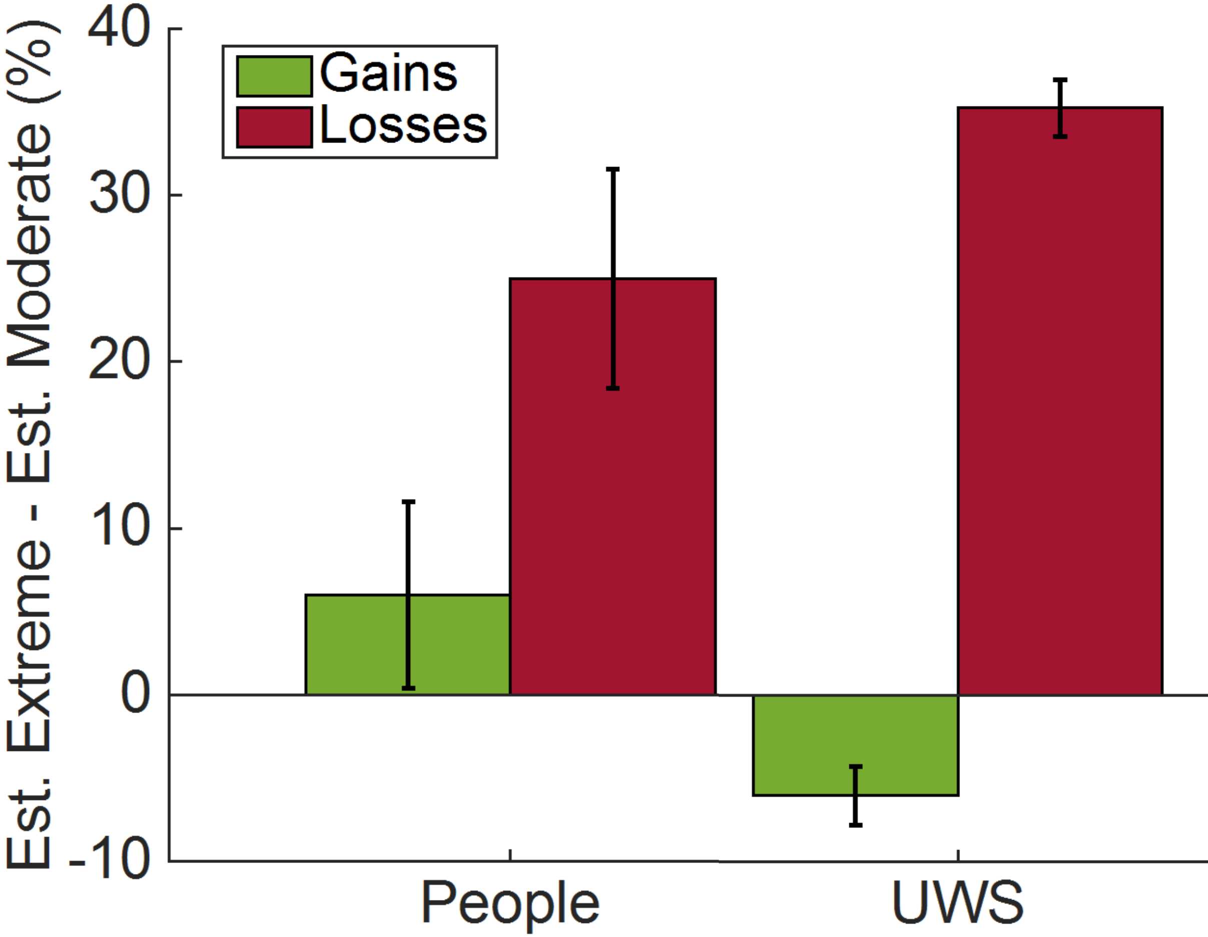

Biased Beliefs Predict Risk Seeking

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

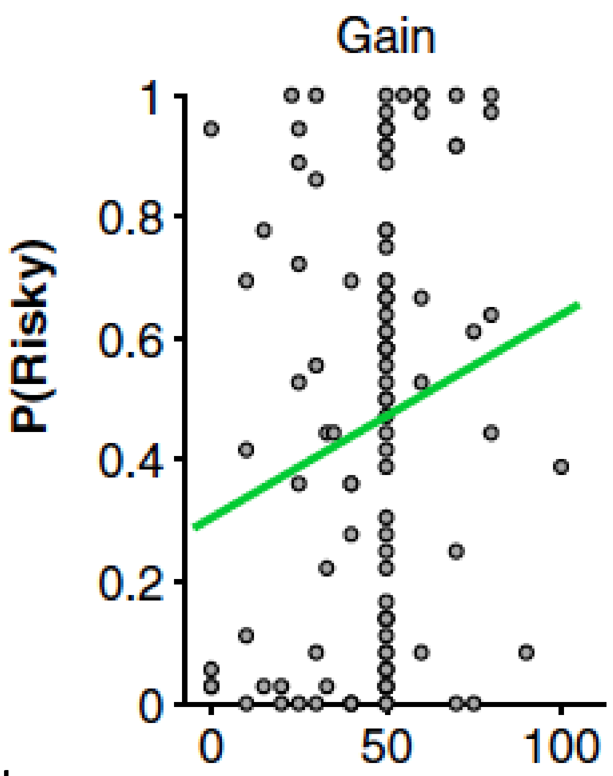

Judged Freq. of High Gain (%)

Biased Beliefs Predict Risk Seeking

r_{UWS} = +0.23

r_{people}= +0.16; *p<0.05*

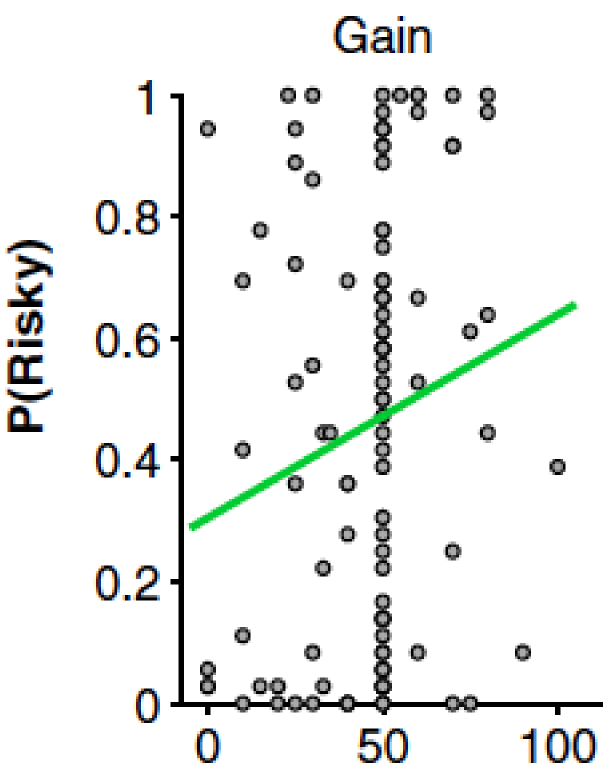

Judged Freq. of High Gain (%)

Biased Beliefs Predict Risk Seeking

$r_{UWS} = +0.23$

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

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

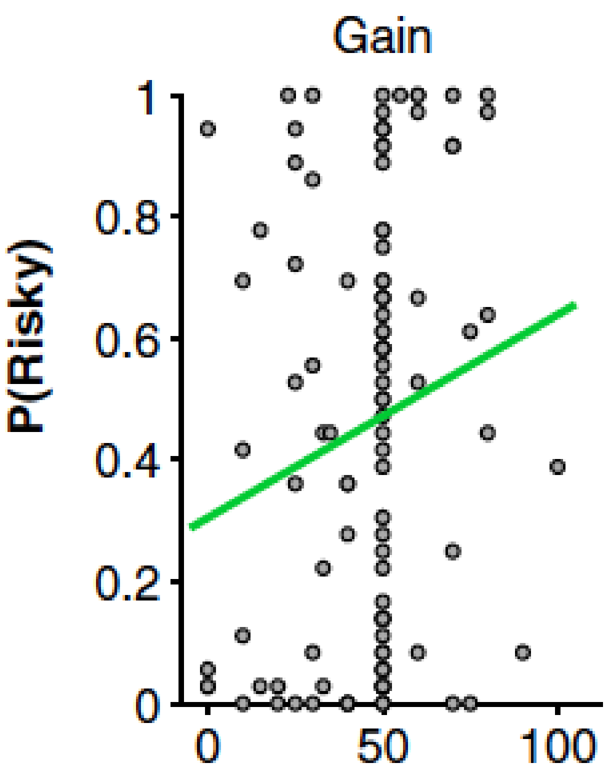

Judged Freq. of High Gain (%)

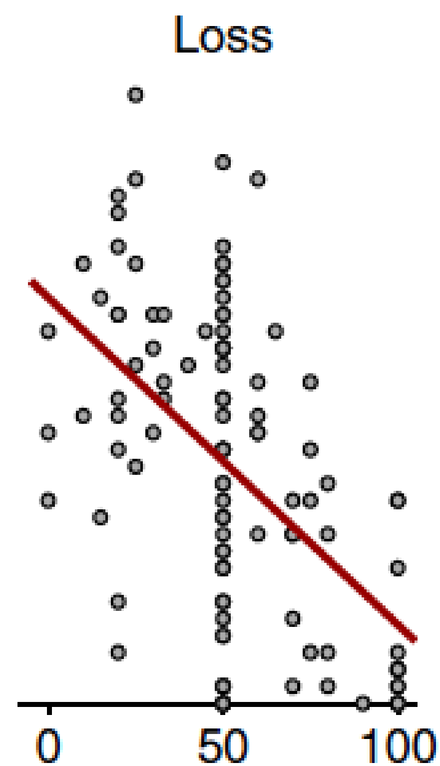

Judged Freq. of Large Loss (%)

Biased Beliefs Predict Risk Seeking

r_{UWS} = +0.23

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

r_{UWS} = -0.44

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

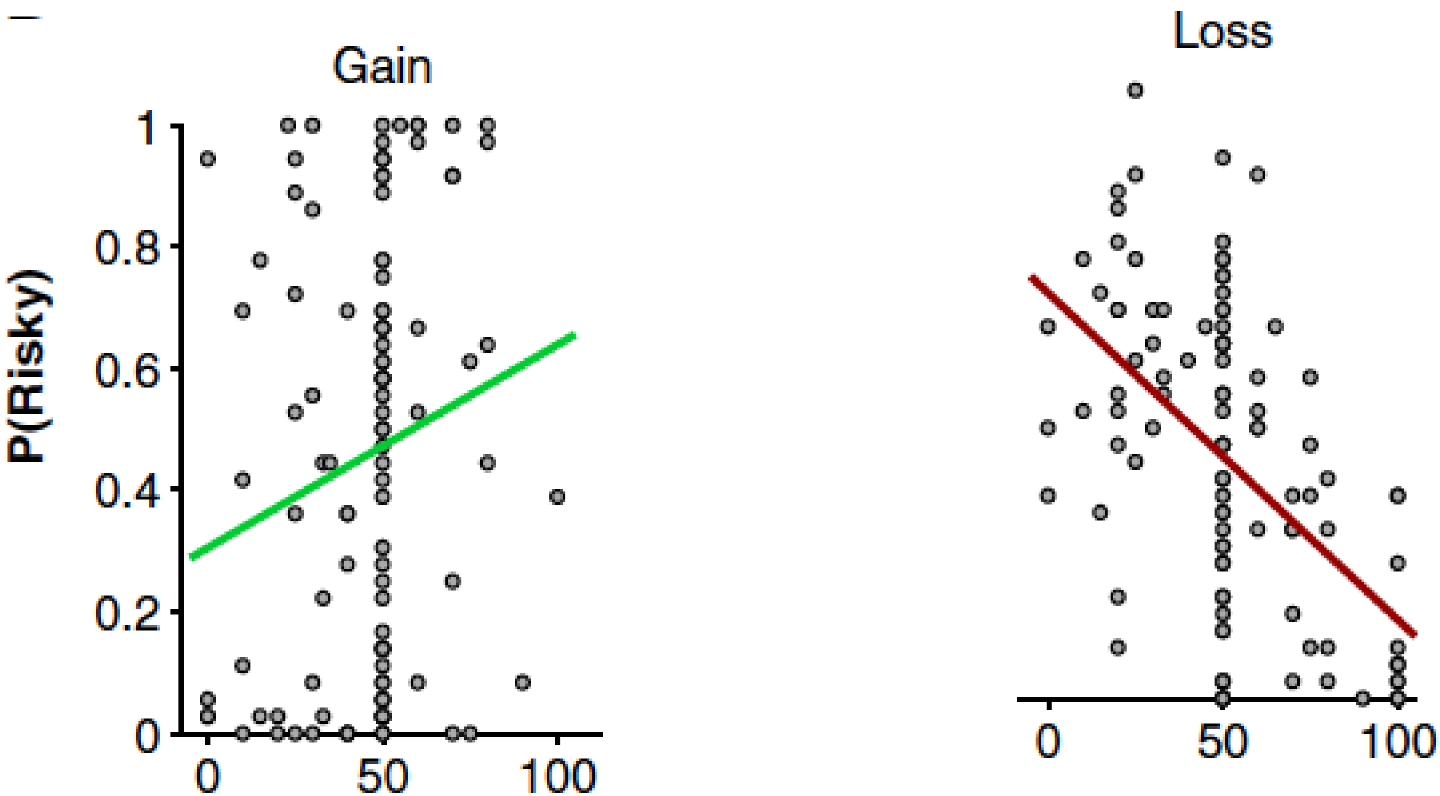

Judged Freq. of High Gain (%)

Judged Freq. of Large Loss (%)

Conclusions

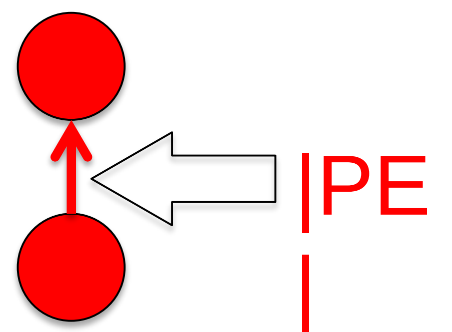

Conclusions

1. Utility-weighted sampling provides a unifying explanation for biases in memory, judgment, and decision making.

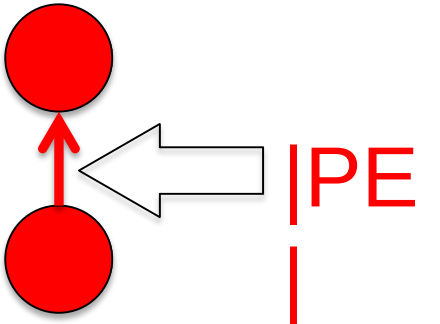

Conclusions

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.

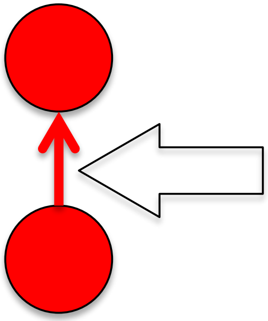

|PE|

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.

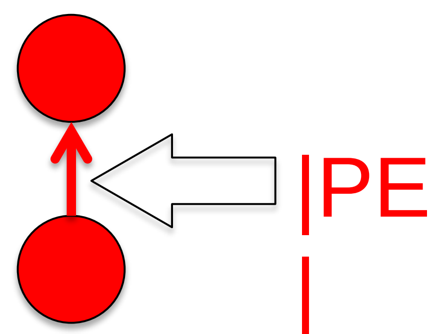

Conclusions

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.

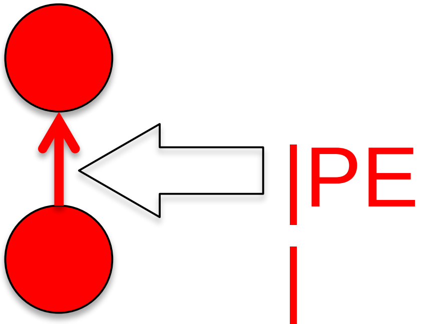

Poster T28

Thank you!

This work was supported by grant number N00014-13-1-0341 from the Office of Naval Research to Thomas L. Griffiths, grant number RO1 MH098023 from the National Institutes of Health to Ming Hsu, and a Berkeley Graduate Fellowship to Falk Lieder.

We thank Elliot Ludvig, Quentin Huys, and Mike Pacer for their suggestions and feedback.