Affective Cognitive Learning and Decision Making : the role of emotions

Hyungil Ahn and Rosalind Picard
MIT Media Lab

Inspirations and Goals

The role of affect in learning and decision making

A Computational Framework of Affective Cognitive Learning and Decision Making

- Covert biases related to previous emotional experiences assist the reasoning process (Bechara & Damasio)
- Loss aversion (Prospect Theory)
- Affective and Deliberative Processes (Loewenstein & Lerner)
- Effects of mood on Decision making (Lerner & Keltner, Mita)

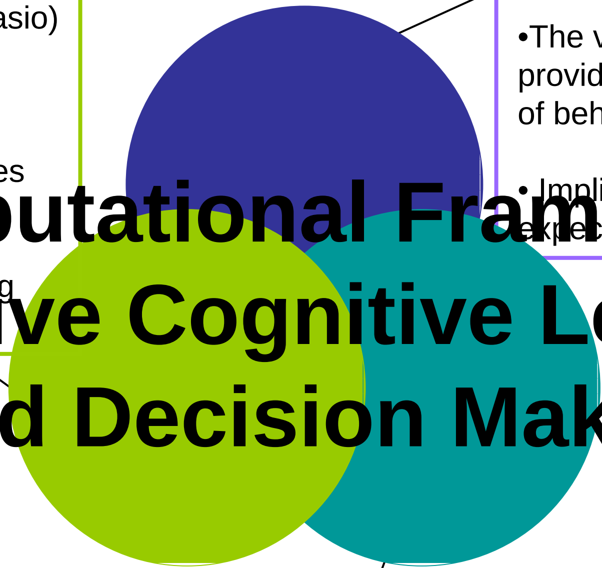

- The basic emotional systems such as SEEKING, RAGE, FEAR, PANIC and PLAY circuits generate distinct affective/neurodynamic states (Panksepp)
- The valenced affective feeling states provide fundamental values for guidance of behavior (Panksepp)
- Implicit “wanting” vs. Explicit cognitive expectation (Berridge & Robinson)

- Most learning and decision-making models are still purely cognitive

Ultimate goals of the computational framework

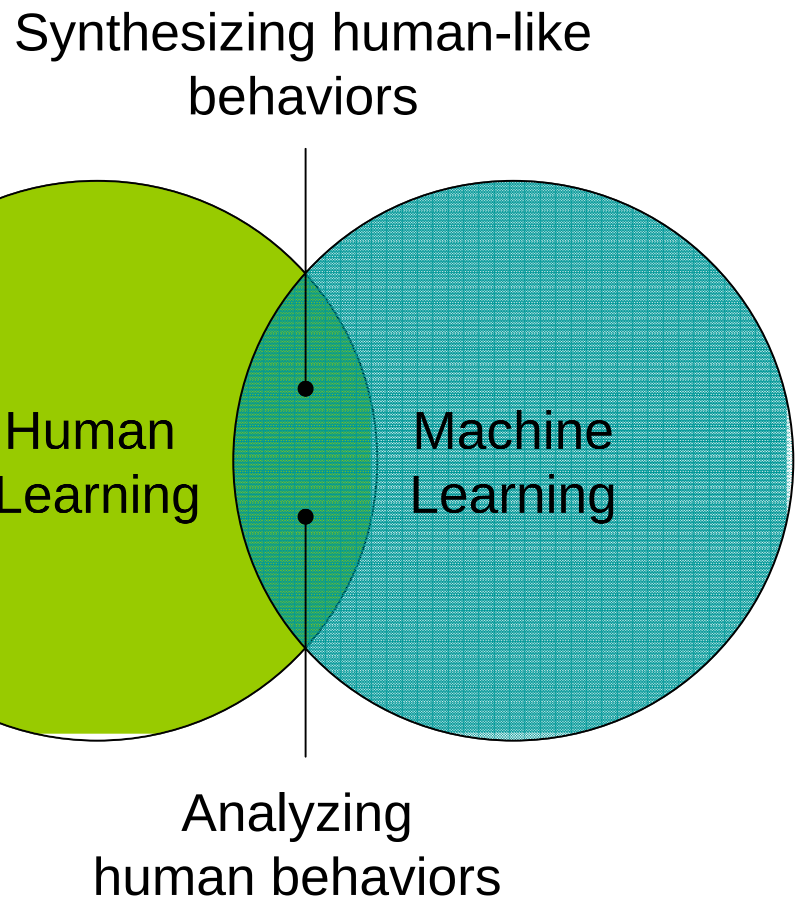

Backgrounds

- **Affective biases**
- **Loss aversion**
- **The Prospect theory (PT) value function**
- **Effects of mood on decision making**

(Example 1) Two-armed bandit gambling tasks

Inspired by Bechara & Damasio's IOWA gambling tasks (Bechara et al. 1997)

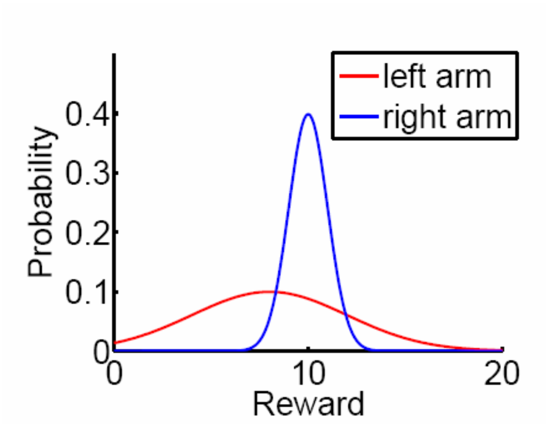

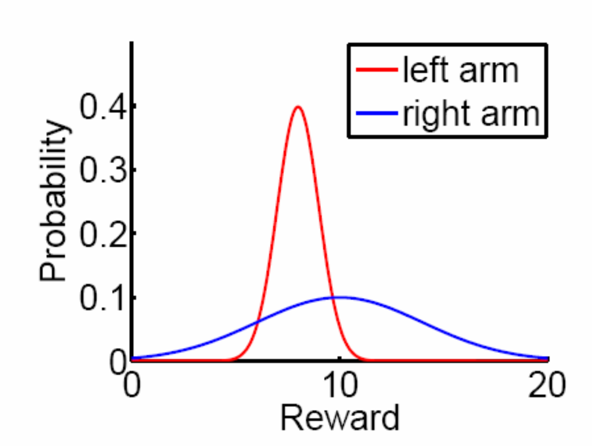

The left arm has 'Negative Valence'
Arousal (uncertainty) as 'feeling uneasy'

The right arm has 'Positive Valence'
Arousal (uncertainty) as 'feeling lucky'

(Example 2) Loss aversion

- Loss aversion: People strongly prefer avoiding losses than acquiring gains

- ‘*Risk-Averse*’ choices in the domain of ‘Likely Gains’

Option 1: $3000 (Pr=1) **>** Option 2: $4000 (Pr=0.8), $ 0 (Pr=0.2)

Expected value = $3000 (Gain) < Expected value = $4000 * 0.8 + $0 * 0.2 = $3200 (Gain)

- ‘*Risk-Seeking*’ choices in the domain of ‘Likely Losses’

Option 1: - $3000 (Pr=1) **<** Option 2: - $4000 (Pr=0.8), $ 0 (Pr=0.2)

Expected value = - $3000 (Loss) > Expected value = - $4000 * 0.8 + $0 * 0.2 = - $3200 (Loss)

The PT (Prospect Theory) value function

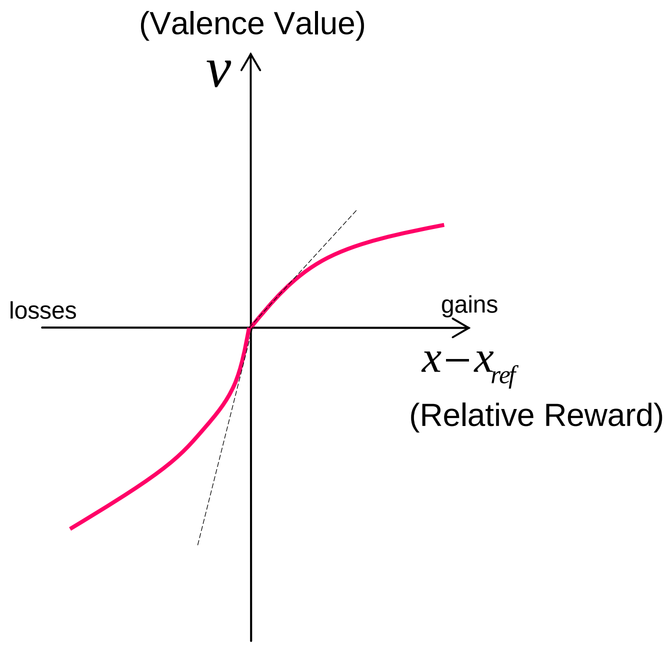

$$v(x - x_{ref}) = \begin{cases} (x - x_{ref})^{\alpha}, & x - x_{ref} \geq 0 \\ -\lambda(-(x - x_{ref}))^{\beta}, & x - x_{ref} < 0 \end{cases}$$

$$0 < \alpha < 1, \quad 0 < \beta < 1, \quad \lambda > 1$$

- Reference Dependence: gains and losses are defined relative to the reference point (x_{ref})

- Concave above the reference point
- Convex below the reference point

(Tversky & Kahneman)

- Diminishing sensitivity: less sensitive to outliers for both gains and losses

- Loss aversion: the function is steeper in the negative (loss) domain

(Example 3) Effects of mood on decision making

(Lerner & Keltner 2001, Mittal 1998)

Happiness

Anger

Optimistic about the present

Optimistic about the future

Pessimistic about the future

Pessimistic about the present

Fear

Sadness

A Computational Framework for Affective Cognitive Learning and Decision Making

- **Introduction**
- **A simple decision-making problem**
- **Decision values**
- **The decision-making model**
- **The affective seeking value**
- **Affective Heating and Cooling**

Affective Cognitive Learning and Decision Making

- A new computational framework for learning and decision making inspired by the neural basis of motivations and the role of emotions in human behaviors

- A motivational value (reward)-based learning theory:

decision value = extrinsic (cognitive) value + intrinsic (affective) value

extrinsic value from the cognitive (deliberative and analytic) systems
intrinsic value from multiple affective systems such as Seeking, Fear, Rage, and other circuits.

- Probabilistic models: Cognition (cognitive state transition), Multiple affect circuits (Seeking, Joy, Anger, Fear, …), and Decision making model

- Any prior and learned knowledge can be incorporated for expecting the consequences of decisions (or computing the cognitive value)

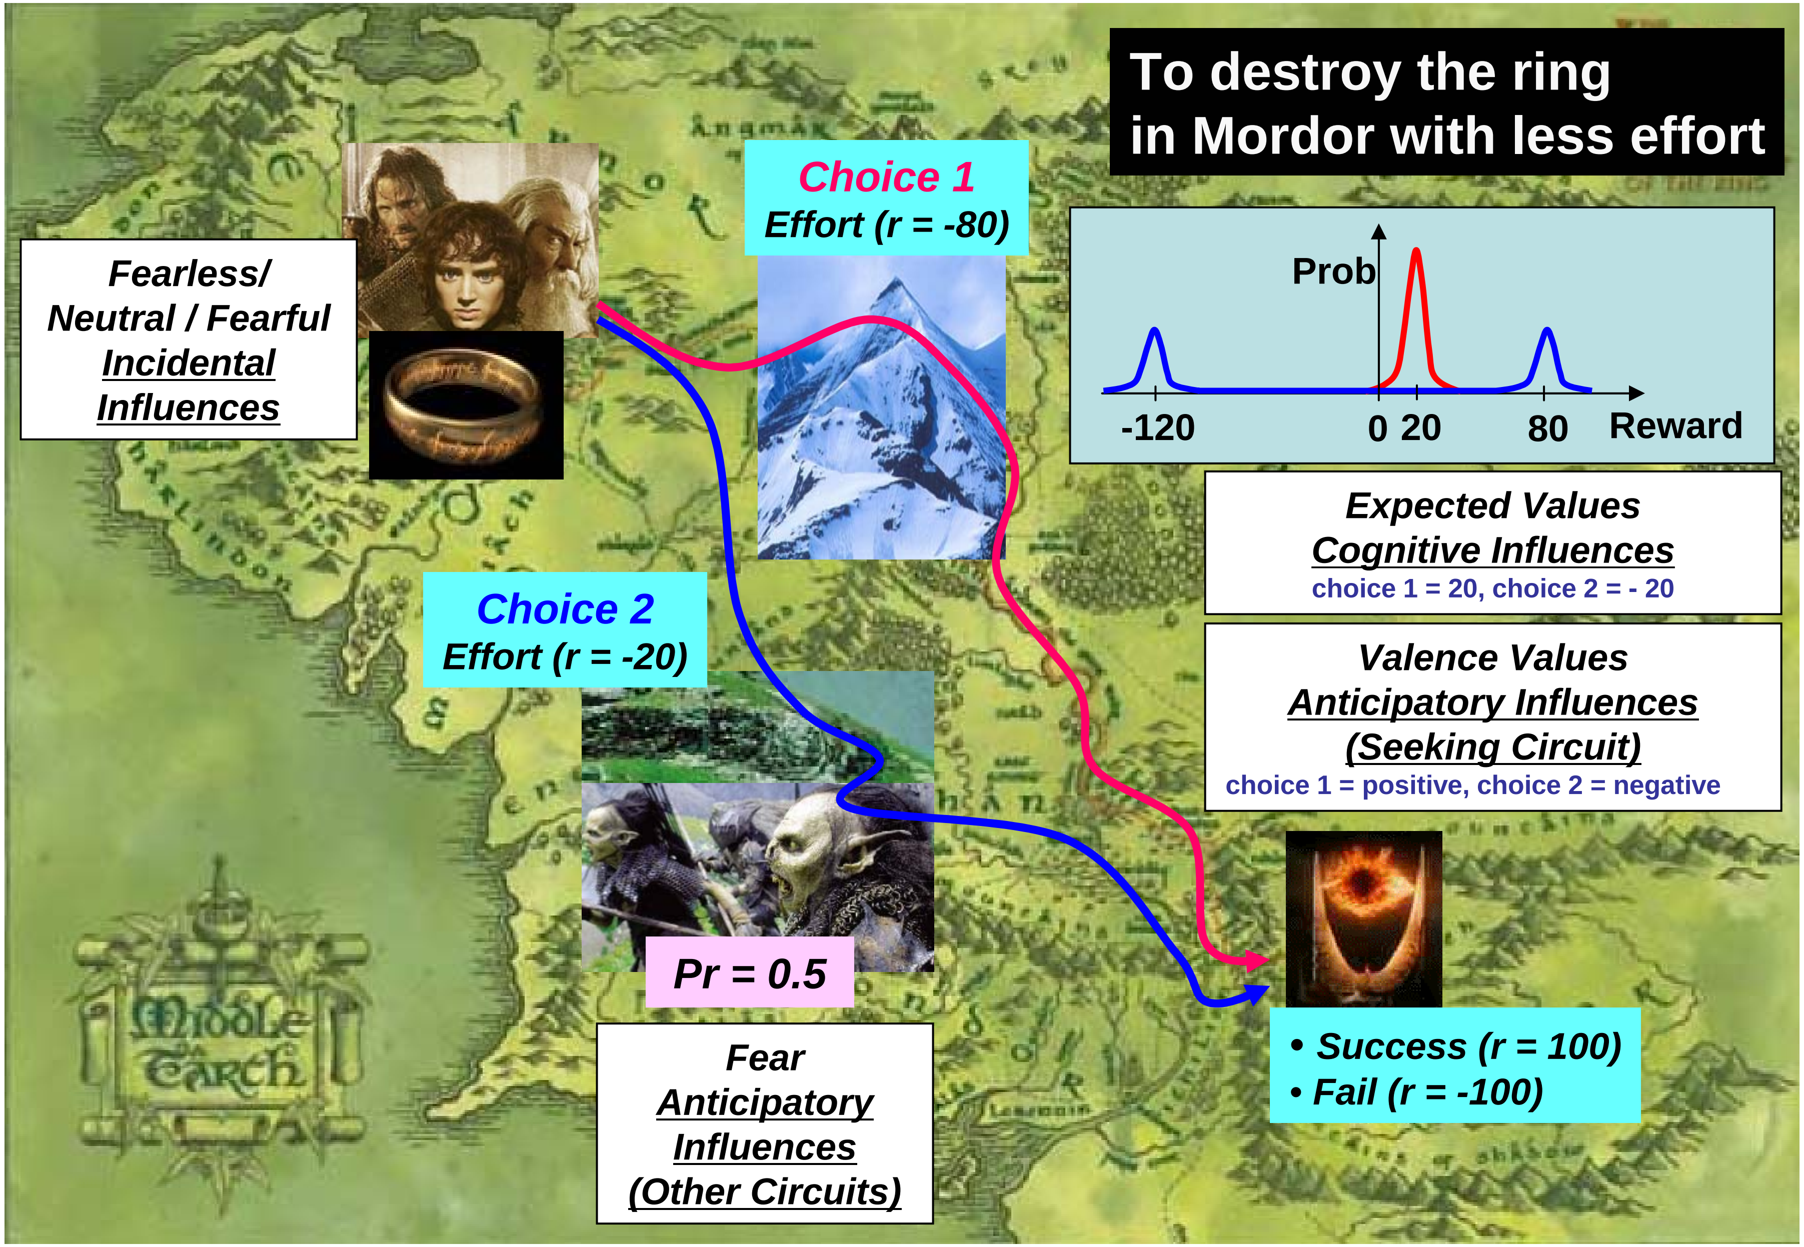

To destroy the ring
in Mordor with less effort
Choice 1
Effort (r = -80)
Fearless/
Neutral / Fearful
Incidental
Influences
Prob
-120
0 20
80
Reward
Expected Values
Cognitive Influences
choice 1 = 20, choice 2 = - 20
Choice 2
Effort (r = -20)
Valence Values
Anticipatory Influences
(Seeking Circuit)
choice 1 = positive, choice 2 = negative
Pr = 0.5
• Success (r = 100)
• Fail (r = -100)
Fear
Anticipatory
Influences
(Other Circuits)

Decision values

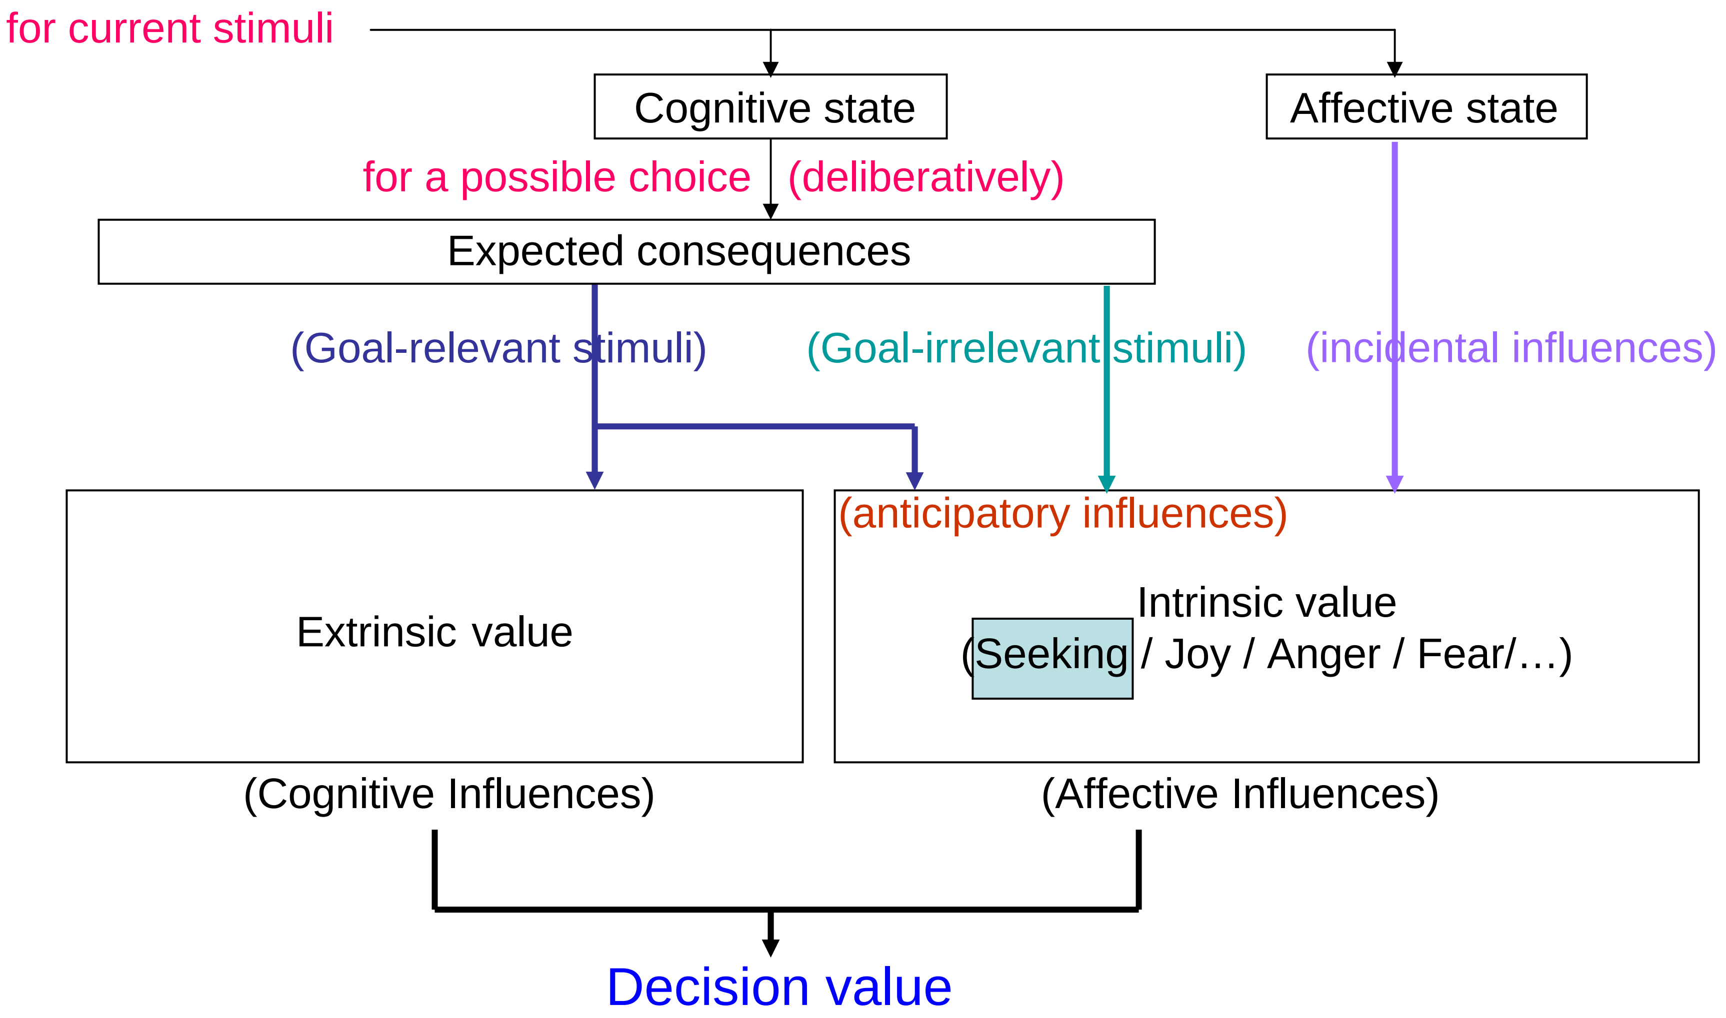

The Decision-Making Model

for a cognitive state (c), an affective state (a), and a decision (d)

$$Q_{DM}(c,a,d) := Q_{ext}(c,d) + Q_{int}(c,a,d)$$

$$\Pr(d \mid c,a) := \frac{\exp(\beta Q_{DM}(c,a,d))}{\sum_{d=1}^{|D|} \exp(\beta Q_{DM}(c,a,d))}$$

Boltzmann Selection

Inverse temperature β regulating the tradeoffs between Exploitation and Exploration

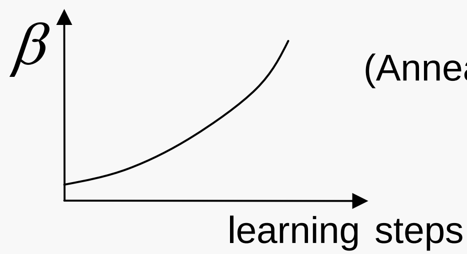

(Annealing)

The affective seeking value

- Affective seeking value = $\sigma_0(c,d)v_0(c,d)$
  (arousal) (valence)
- Valence = decided by the mean of the filtered values for the reward samples
- Arousal = uncertainty of the reward sample distribution
- Let's consider only the extrinsic value and the affective seeking value

$$\begin{aligned} Q_{DM}(c,d) &:= Q_{ext}(c,d) + Q_{int}(c,d) \\ &= Q_{ext}(c,d) + \eta_0\sigma_0(c,d)v_0(c,d) \end{aligned}$$

(the cognitive component) (the affective component)

- To compare, the non-affect agent has only the cognitive component

The Affective agent vs. the Non-affect agent

$d_h = \arg\max_d Q_{ext}(c,d)$: the cognitive decision (the decision with the higher mean)

$d_l = \arg\min_d Q_{ext}(c,d)$: the other decision with the lower mean

$\Pr(d_h \mid c) := \frac{\exp(\beta Q_{DM}(c,d_h))}{\exp(\beta Q_{DM}(c,d_h)) + \exp(\beta Q_{DM}(c,d_l))}$: the probability of choosing the cognitive decision

The non-affective agent

$Q_{DM}(c,d) := Q_{ext}(c,d)$

$$\Pr(d_h \mid c) := \frac{1}{1+\exp(-\beta\left[Q_{ext}(c,d_h) - Q_{ext}(c,d_l)\right])}$$

The affective agent

$Q_{DM}(c,d) := Q_{ext}(c,d) + Q_{int}(c,d)$

$$\Pr(d_h \mid c) := \frac{1}{1+\exp(-\beta\left[Q_{ext}(c,d_h) - Q_{ext}(c,d_l) + Q_{int}(c,d_h) - Q_{int}(c,d_l)\right])}$$

$$= \frac{1}{1+\exp(-\beta'\left[Q_{ext}(c,d_h) - Q_{ext}(c,d_l)\right])}$$

$$\beta' = \left(1 + \frac{Q_{int}(c,d_h) - Q_{int}(c,d_l)}{Q_{ext}(c,d_h) - Q_{ext}(c,d_l)}\right)\beta$$ **: the "induced" inverse temperature**

→ **induced by the affective component**

"Agree" (+) (Affective Cooling)
"Conflict" (-) (Affective Heating)

Valences (in a model of the seeking circuit)

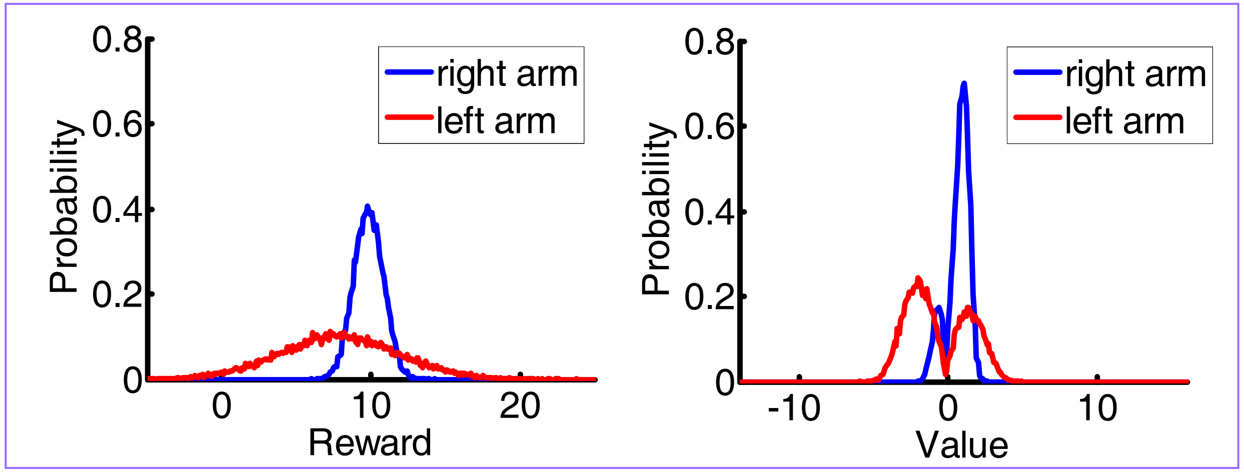

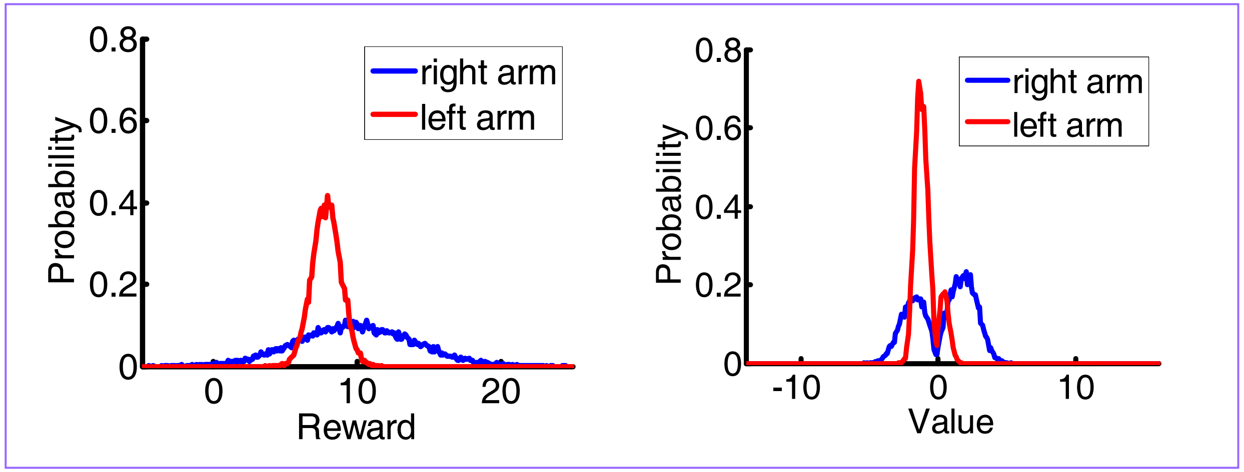

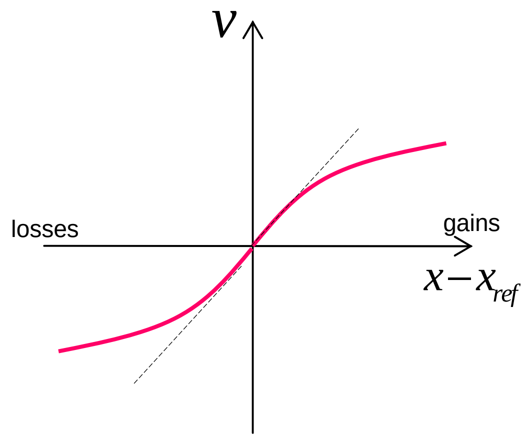

$x_{ref} \approx 9$ **: the total average of all the previous sample rewards**

- Valences are decided by the means of the PT values for given samples
- Valences are less sensitive to outliers: valence values change more slowly than the expected values, for a given outlier
- The arm with the higher mean is more likely to have positive valence (if the agent doesn't have too much loss aversion)

The influence of an outlier on the cognitive values and the valence values

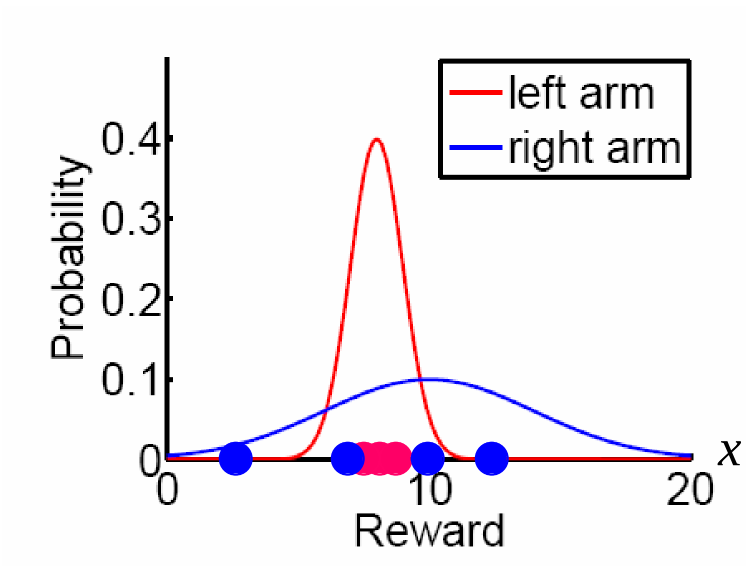

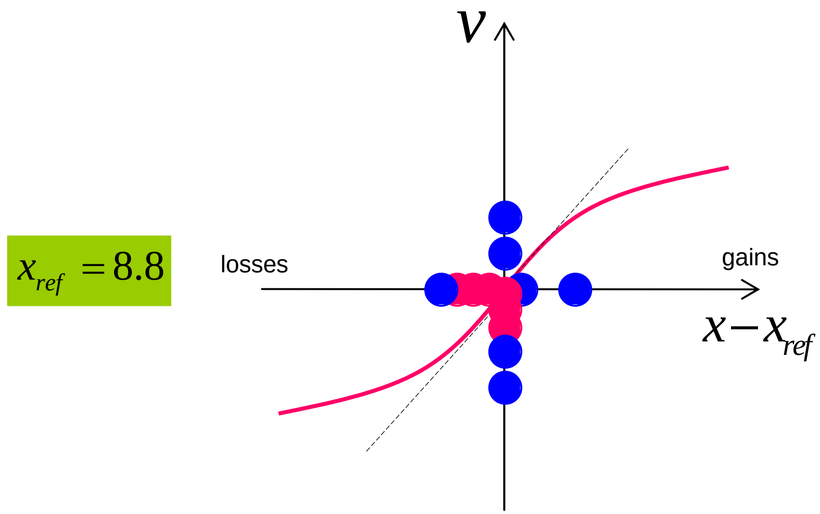

CogValue(left) = 8.27 **<** CogValue(right) = 9.33

Valence(left) = -0.68 **<** Valence(right) = 0.42

The affective component agrees with the cognitive component;
More likely to follow the cognitive decision, compared with the non-affect agent

After the outlier,

CogValue(left) = 8.27 **>** CogValue(right) = 7.75

After the outlier,

Valence(left) = -0.68 **<** Valence(right) = -0.40

The affective component conflicts with the cognitive component;
Less likely to follow the cognitive decision, compared with the non-affect agent

Example 2

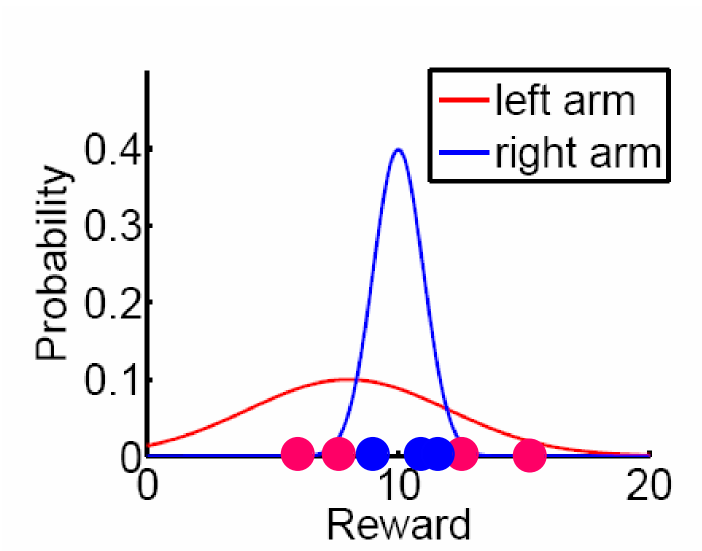

$x_{ref} = 9.5$

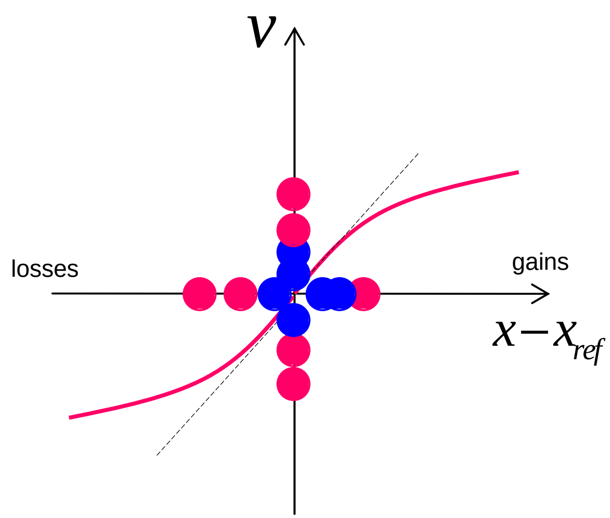

CogValue(left) = 8.33 < CogValue(right) = 10.17

Valence(left) = -0.46 < Valence(right) = 0.54

The affective component agrees with the cognitive component;
More likely to follow the cognitive decision, compared with the non-affect agent

After the outlier,

CogValue(left) = 10.38 > CogValue(right) = 10.17

After the outlier,

Valence(left) = 0.35 < Valence(right) = 0.54

The affective component conflicts with the cognitive component;
Less likely to follow the cognitive decision, compared with the non-affect agent

Affective Heating and Cooling

- The affective component is less sensitive to outliers than the cognitive component

- The agreement between two components
- More likely to follow the decision by the cognitive component *[Exploitation]*
- The value of the "induced" inverse temperature parameter *increases*
- *Affective Cooling* (like humans more depend on cognition in decision making when their affect is cool)

- The conflict between two components
- Less likely to follow the decision by the cognitive component *[Exploration]*
- The value of the "induced" inverse temperature parameter *decreases*
- *Affective Heating* (like humans less depend on cognition in decision making when their affect is hot)

10-armed bandit tasks

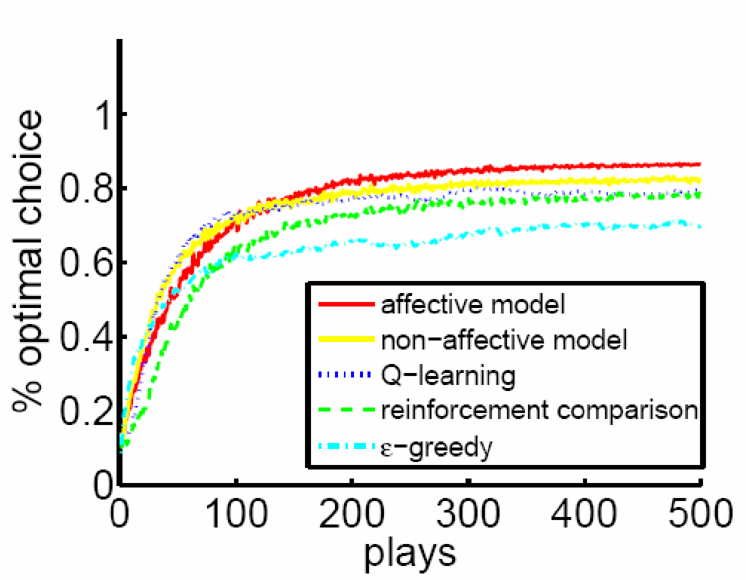

The 10-armed bandit tasks: for each arm d, the rewards were selected from a normal probability distribution with mean $m(d)$ and variance 1. The 1000 tasks were generated by reselecting the $m(d)$ 1000 times, each according to a normal distribution with mean 0 and variance 1.

The plot shows the comparison between algorithms.
Each algorithm used the optimal parameters for its own.
The lines in the plot show averages over 1000 tasks.

Too much or too little affect impairs learning

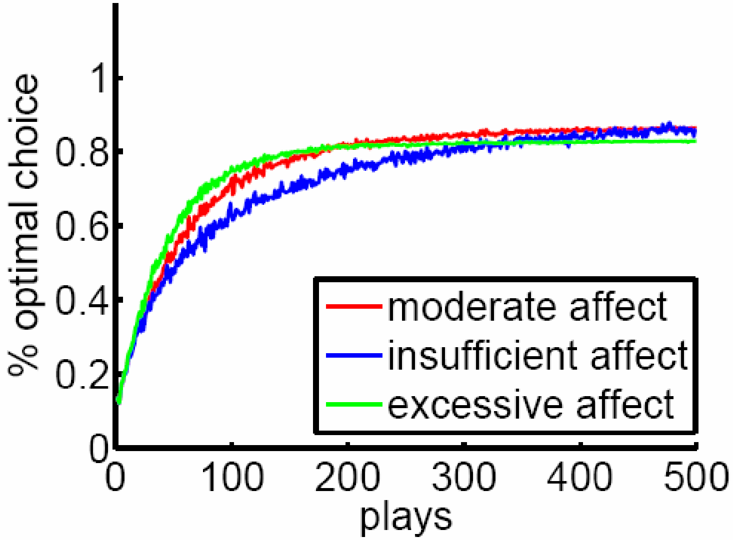

The 10-armed bandit tasks: comparison between models with different η_0: $\eta_0 = 5$ (moderate affect), $\eta_0 = 0$ (insufficient affect), and $\eta_0 = 15$ (excessive affect).

- Excessive affect is good for faster learning but not for long-run performance
- Insufficient affect reaches good performance in the long-run but is slow

Results and Conclusions

- There are lots of things to add to the framework, such as models for other affect circuits, their incidental influences on decision making, and the use of prior knowledge for expecting cognitive outcomes
- The affective bias by the affective component (the seeking circuit) helps automatically regulate exploration and exploitation
- The affective bias can speed up learning without sacrificing decision quality (e.g. 10-armed bandit tasks)
- This framework might mimic well-studied human behavior such as risk aversion, effects of mood on decision making, and self-control

Questions