BAYESIAN DECISION THEORY

Paul Schrater
University of Minnesota

Example decision: End point planning

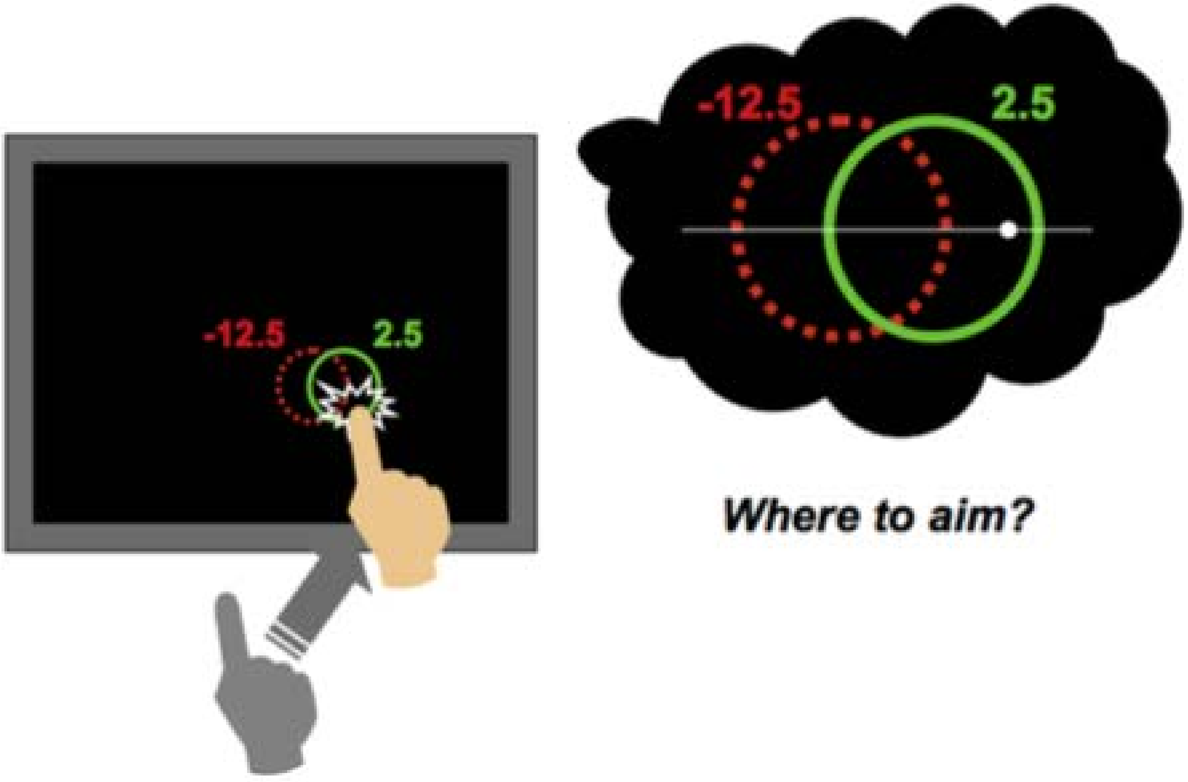

Example decision: Random Dot Coherent motion paradigm

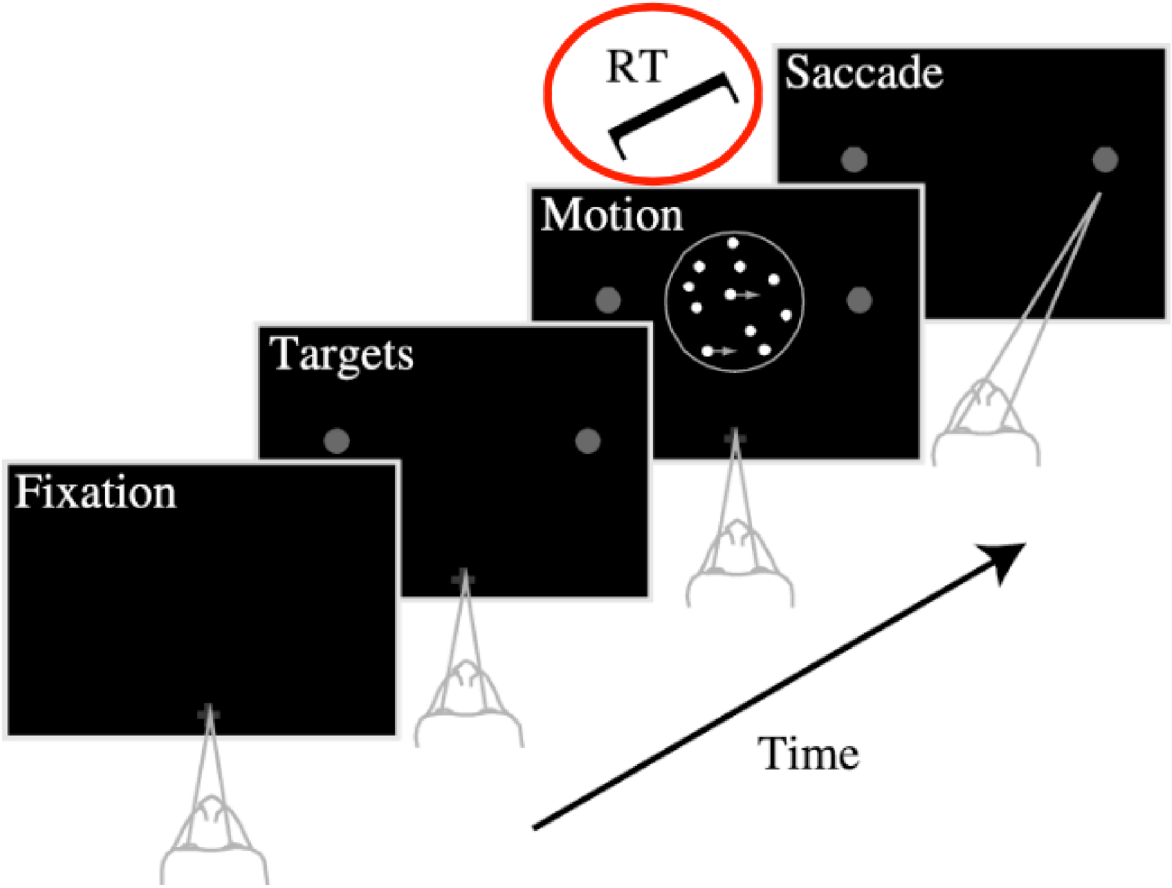

Example decision: Random Dot Coherent motion paradigm

Easy

30% coherence

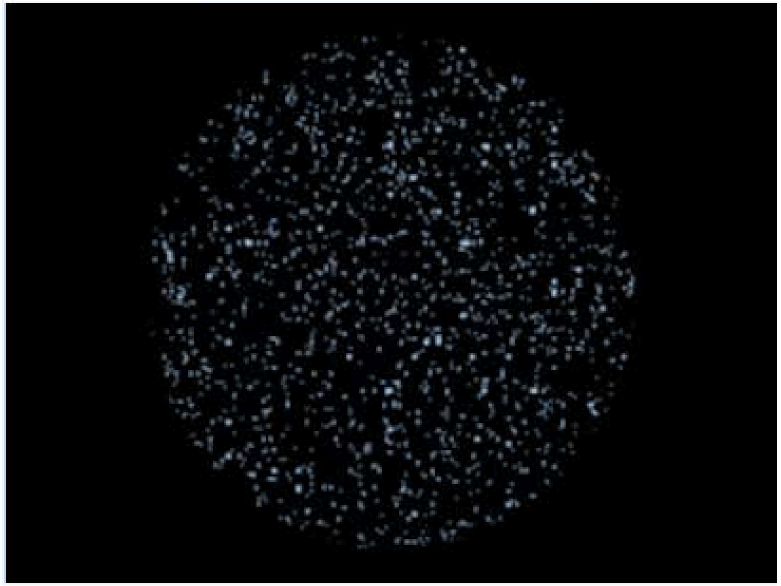

Difficult

5% coherence

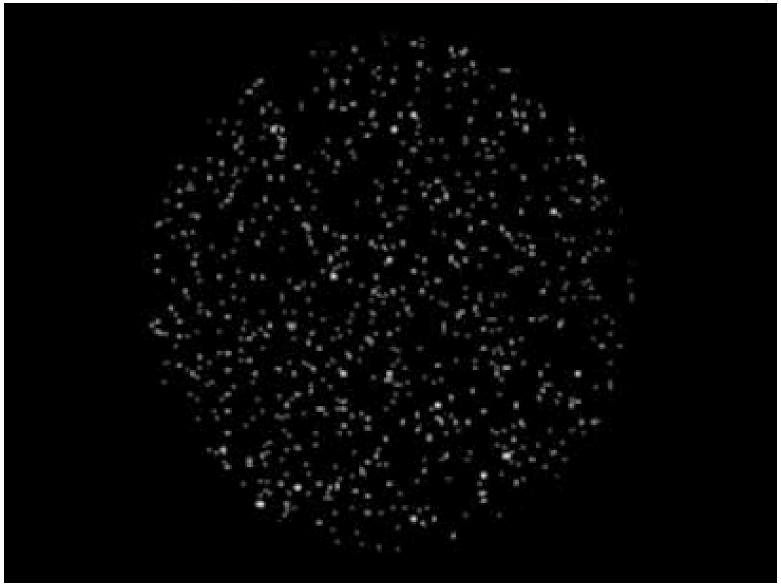

Left or Right?

Example decision: Random Dot Coherent motion paradigm

Speed **vs.** **Accuracy**

When to initiate movement?

Basic Decision Components

- *Sensory Space (Observable variables)*

  $$\boldsymbol{x} = \{\boldsymbol{x}_1, \ldots., \boldsymbol{x}_t\}$$

- *State Space (hidden variables)*

  $$\boldsymbol{s} = \{\boldsymbol{s}_1, \ldots., \boldsymbol{s}_t\}$$

- *Action Space (controllable variables)*

  $$\boldsymbol{a} = \{\boldsymbol{a}_1, \ldots., \boldsymbol{a}_t\}$$

- *Outcome Space (variables affecting costs/rewards)*

  $$\boldsymbol{o} = \{\boldsymbol{o}_1, \ldots., \boldsymbol{o}_t\}$$

Model components

- *Sensory Likelihood*

$$P(x|s)$$

- *Prior*

$$P(s)$$

- *Outcomes*

$$P(o|a,s)$$

- *Rewards/Costs*

$$L(o,a)$$

Goal

- Find Policy: Map between observations and actions

$$a = \pi(x) \quad \text{OR} \quad P(a|x)$$

- That minimizes Expected Reward/Costs

$$a^* = \pi(x)$$
$$= \arg\min_a L(a|x)$$

- Where

$$L(a|x) = \sum_o \sum_s L(o,a)\, P(o|s,a)\, P(s|x)$$

Decision Steps

- Infer current Bayesian inference:

$$P(s|x) = \frac{P(x|s)\,P(s)}{\sum_s P(x|s)\,P(s)}$$

- **Forecast** Outcome probabilities for each action

$$P(o|a,x) = \sum_s P(o|s,a)\,P(s|x)$$

- Compute Expected Reward/Costs

$$\mathbf{L}(a|x) = \sum_o L(o,a)\,P(o|a,x)$$

- Optimize the expected cost function to find best policy

Policy: Sensory history -> choice

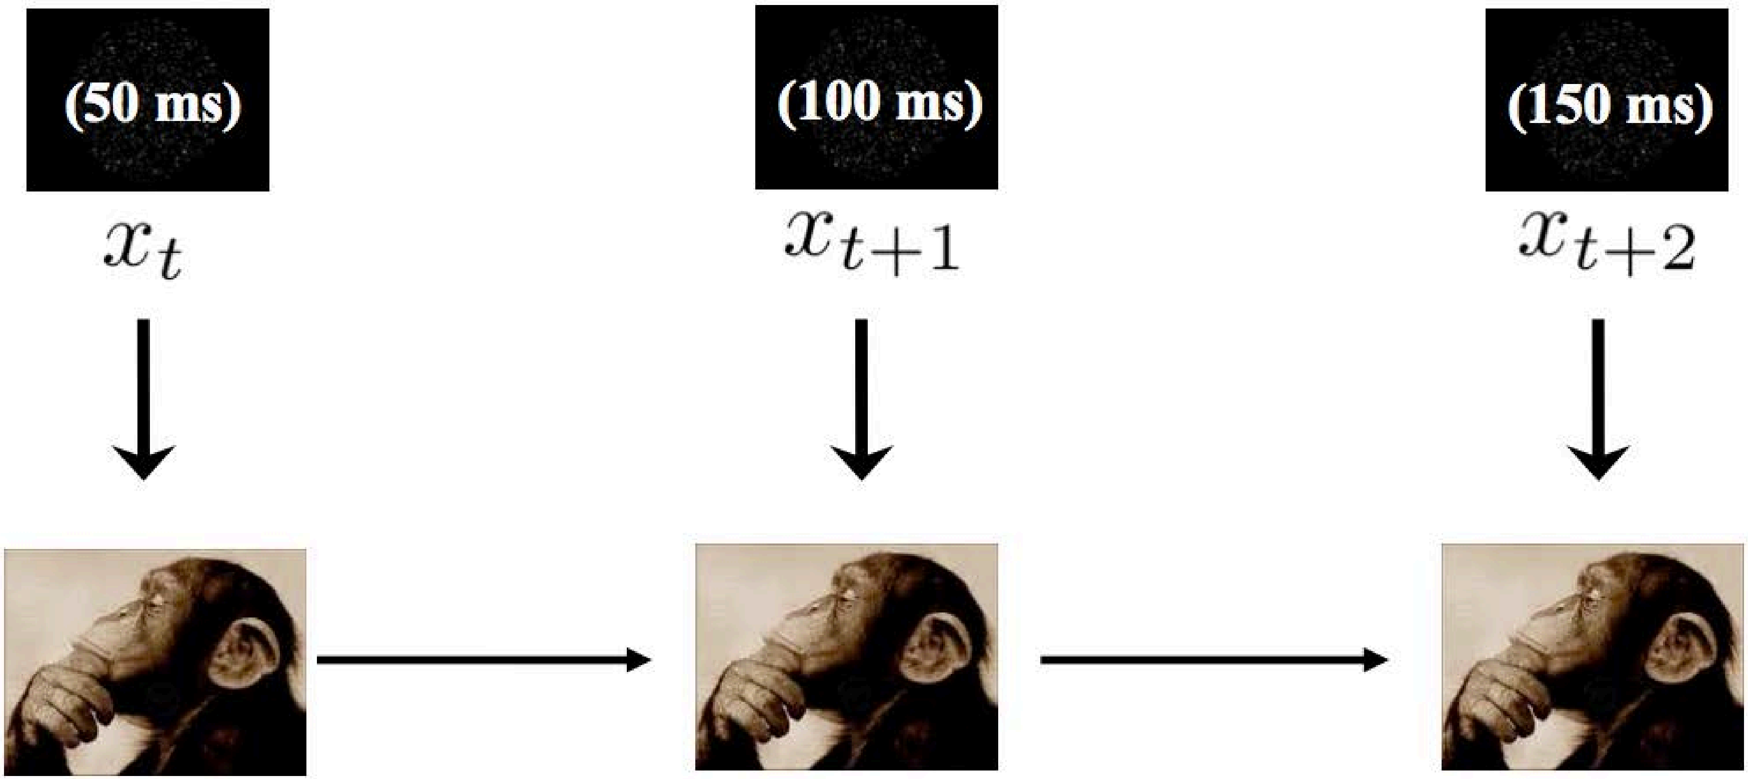

Left or Right?

Bayesian Inference of hidden state

hidden variable: L or R

prior $p(s;\theta)$

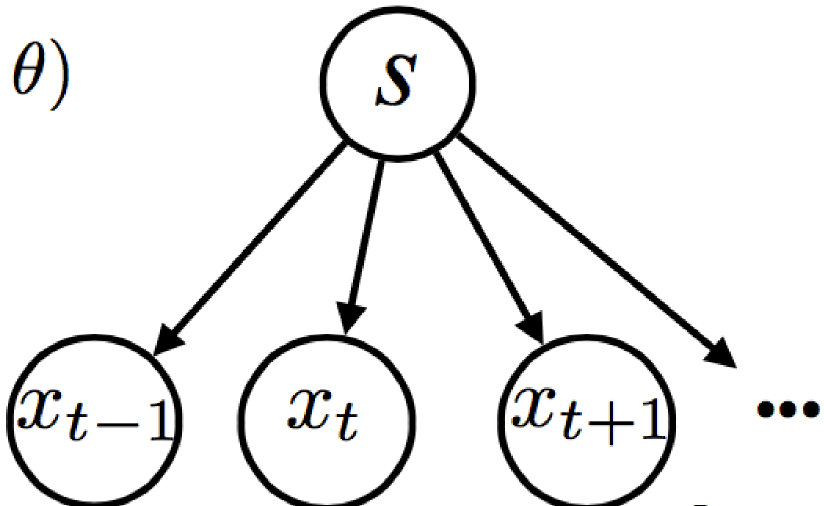

likelihood

$p(x|s;\phi)$

observed data

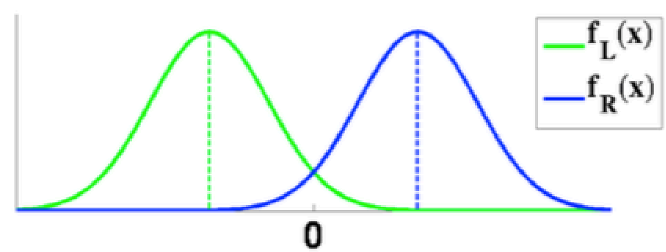

iid **noise**

$$p(\mathbf{x}_t|s;\phi) = \prod_{i=1}^{t} p(x_i|s;\phi)$$

$$\mathbf{x}_t := (x_1, \ldots, x_t)$$

Bayesian Inference

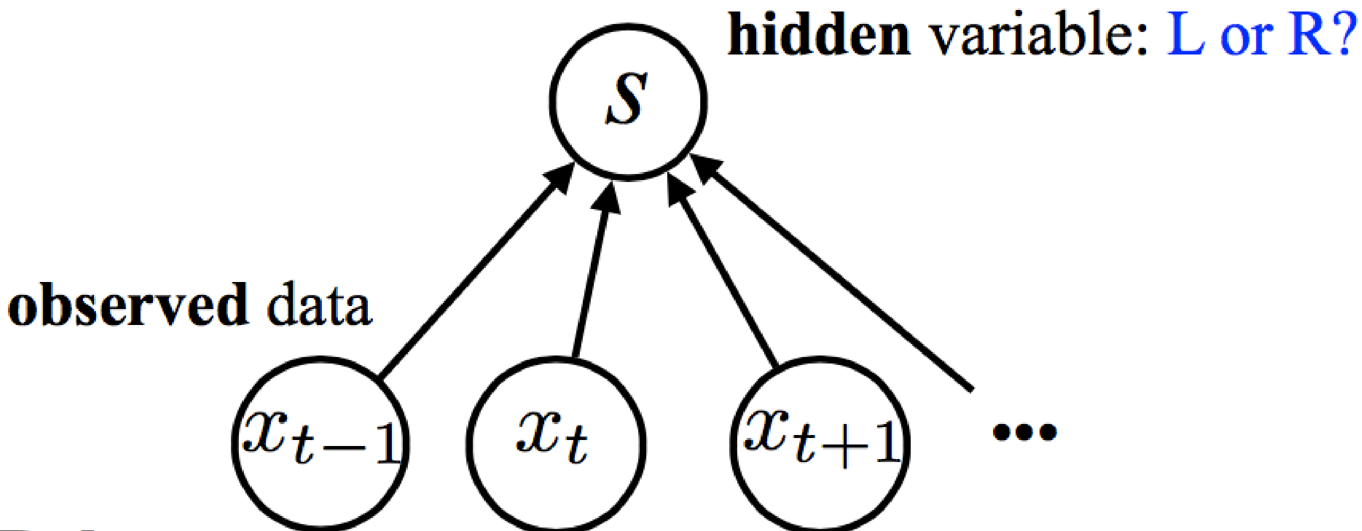

Bayes' Rule

(batch)

$$p(s|\mathbf{x}_t) = \frac{p(\mathbf{x}_t|s)p(s)}{p(\mathbf{x}_t)} = \frac{p(\mathbf{x}_t|s)p(s)}{\cancel{\int p(\mathbf{x}_t|s')p(s')ds'}\ Z} \propto p(\mathbf{x}_t|s)p(s) = p(s)\prod_{i=1}^{t} p(x_i|s)$$

$$= \frac{p(x_t|s, \cancel{\mathbf{x}_{t-1}})p(s|\mathbf{x}_{t-1})}{p(x_t|\mathbf{x}_{t-1})} = \frac{p(x_t|s)p(s|\mathbf{x}_{t-1})}{\cancel{\int p(x_t|s')p(s'|\mathbf{x}_{t-1})ds'}\ Z'} \propto p(x_t|s)p(s|\mathbf{x}_{t-1})$$

(online)

Outcome forecast

Sensory stream

Actions (stoptime(τ), L/R)

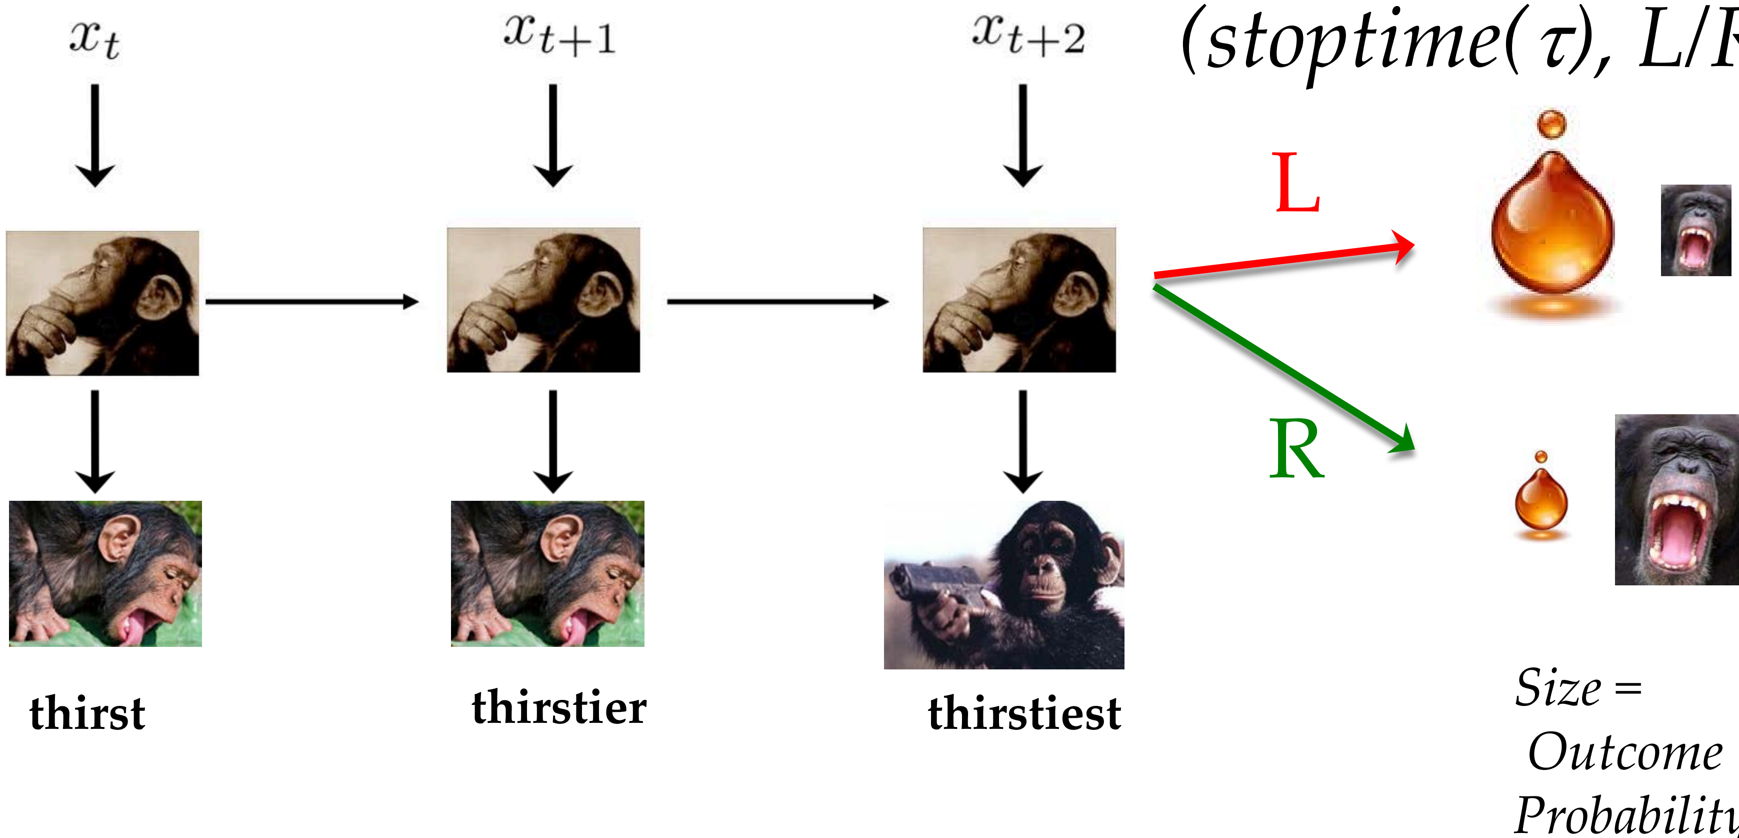

Thirst accrues with elapsed time

Outcome probability

$$P(o \mid s,a) = P(thirst = t, Juice \mid s, a_\tau, a_{L/R})$$
$$= \delta(t-a_\tau)\,\delta(s-a_{L/R})$$

The thirst state is however much time has elapsed till decision

The juice state is true if the reward location *s matches the choice action*

Expected Cost

- Simple cost function

Average amount of juice per unit time

$$L(o,a) = J/(\tau+T)$$

where T is the intertrial interval

- Expected cost

$$L(a|x) = \sum_o \sum_s L(o,a)\, P(o|s,a)\, P(s|x)$$

$$L(a|x) = \sum_o \sum_s (J/\tau)\, \delta(t-a_\tau)\, \delta(s-a_{L/R})\, P(s|x)$$

$$L(a_L, a_\tau|x) = J/(a_\tau+T)\, P_\tau(a_L|x)$$

Simulating

P(s|x) as a function of time

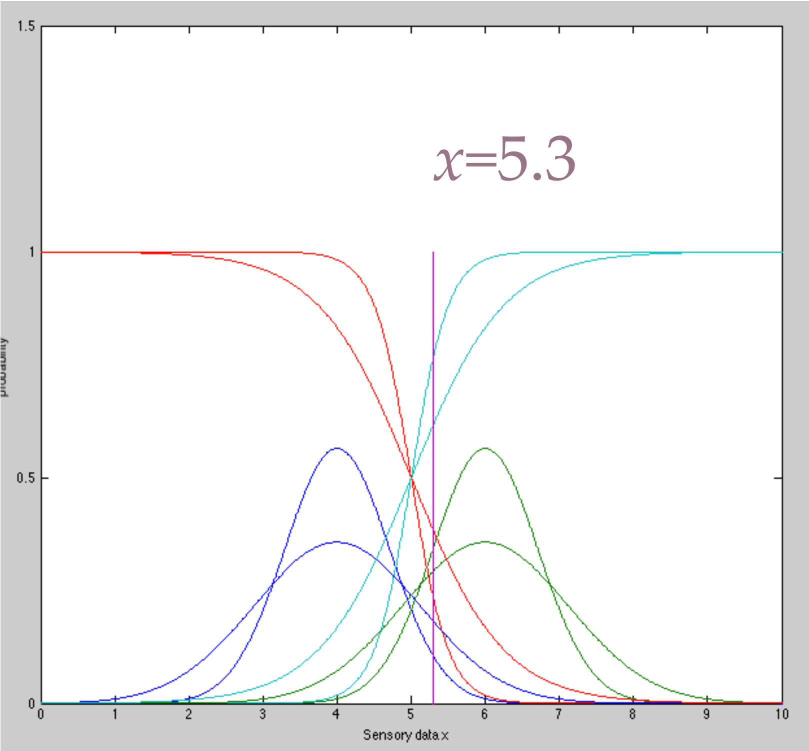

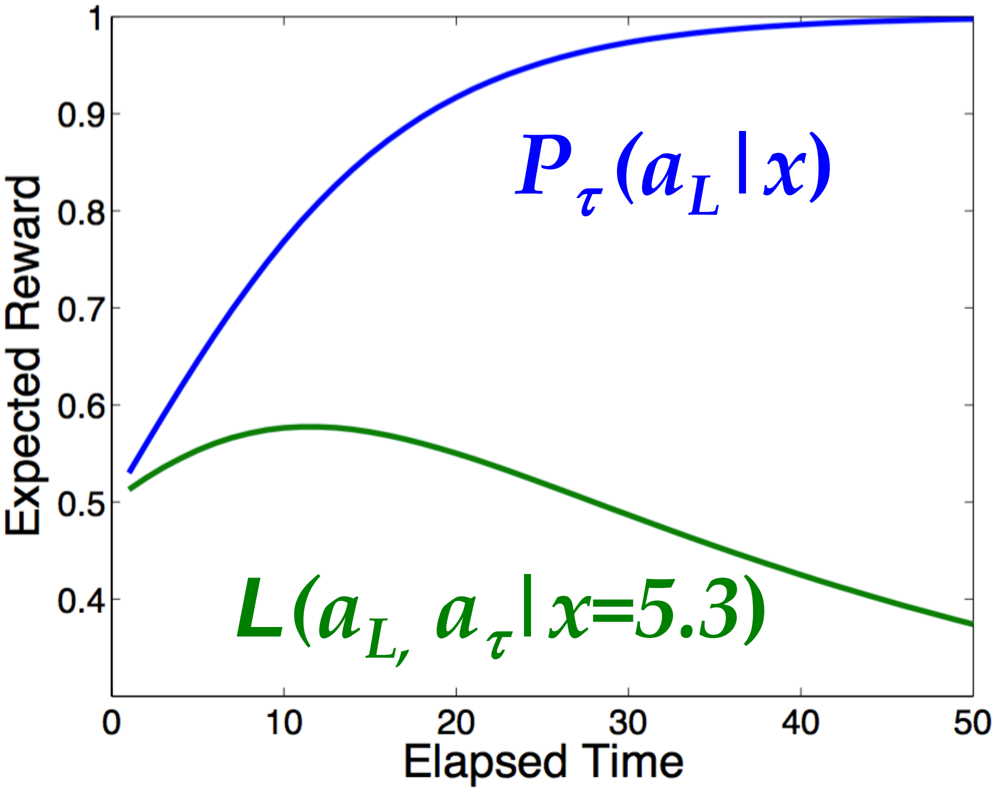

Example decision: End point planning

Action

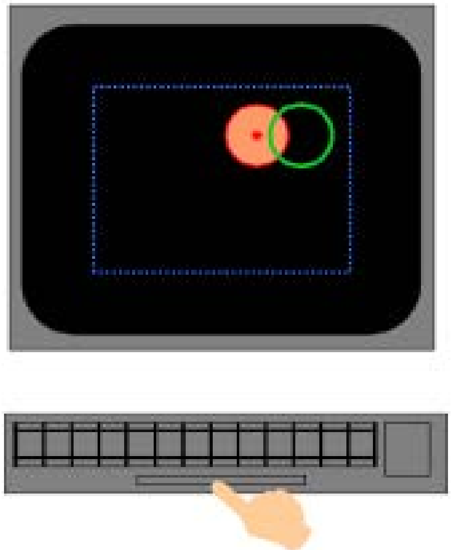

Reward/Cost

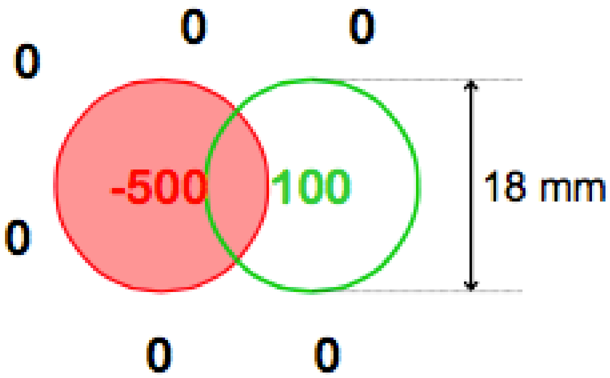

Outcome

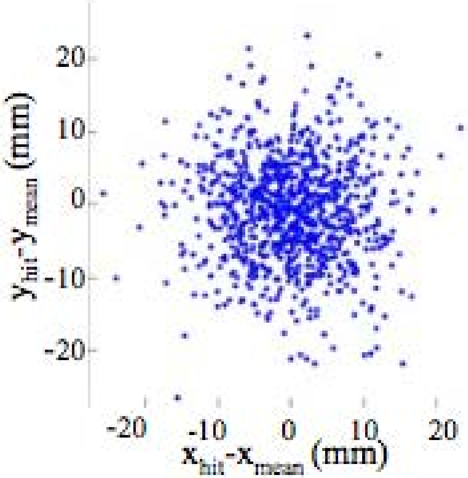

Expected Gain

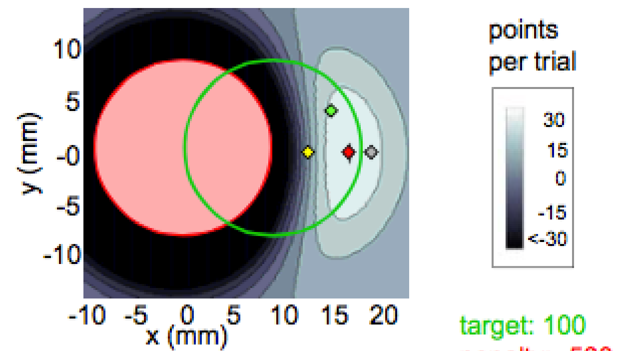

σ = 4.83 mm

- 0.3 pts. per trial

30.7 pts. per trial

25.5 pts. per trial

22.6 pts. per trial

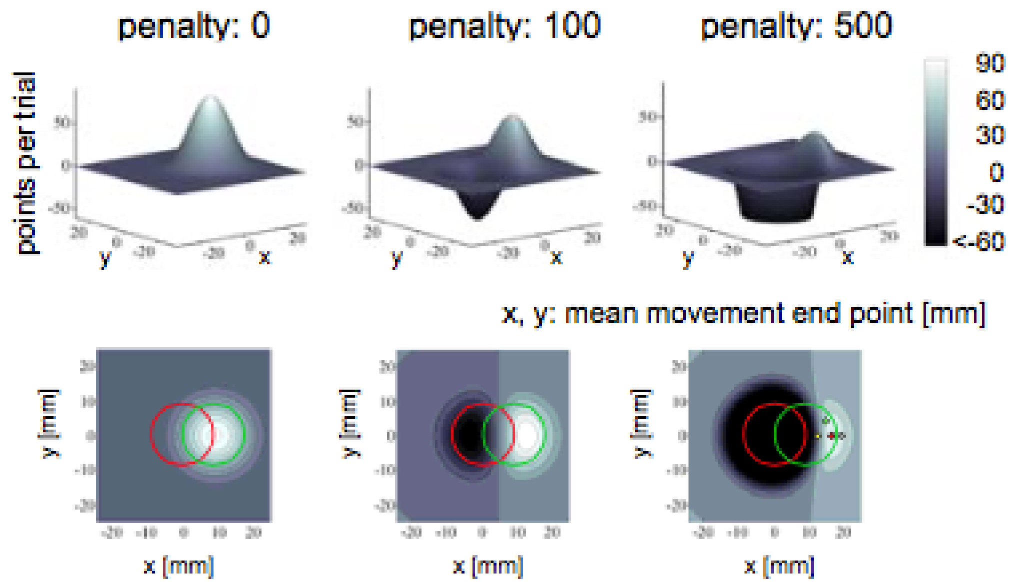

σ = 4.83 mm

Behavior

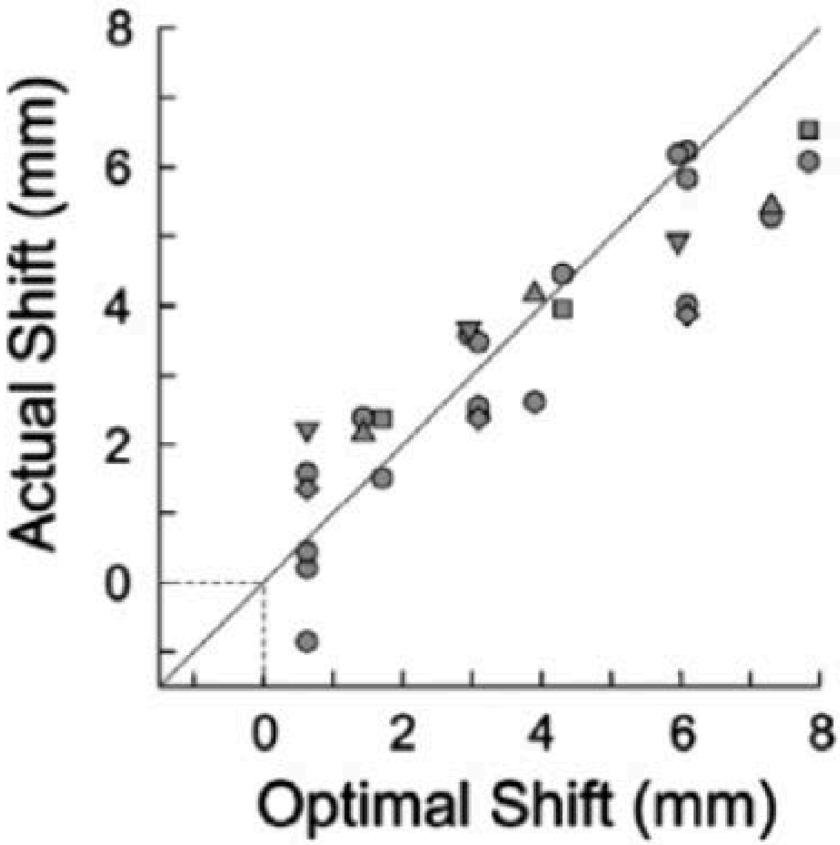

Let’s derive this

Battaglia and Schrater

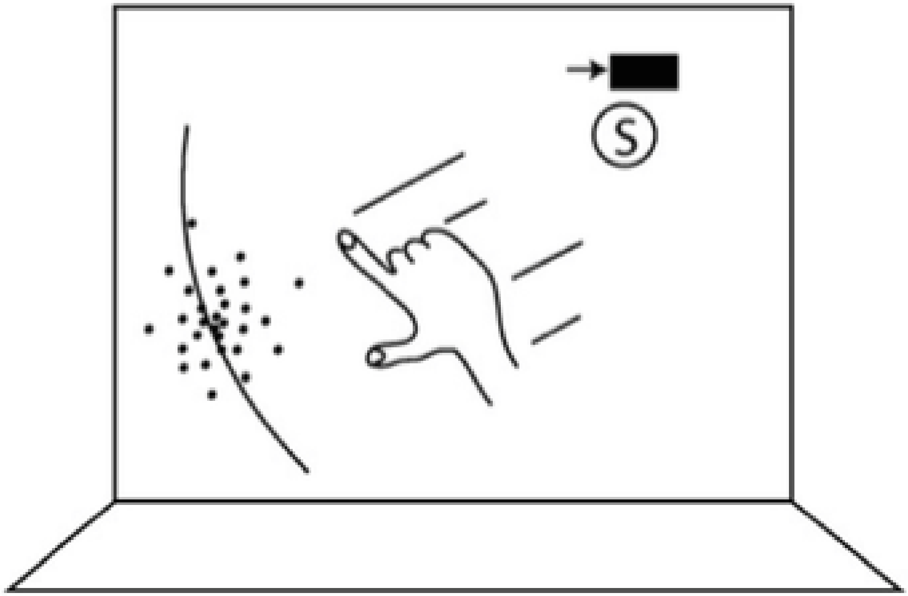

Deadline = 1400 msec

Dots accumulate while finger on button

Reach to intercept centroid before deadline

Battaglia and Schrater

State?
Actions?
Outcomes?