Robustness of Collective Decision Dynamics

Naomi Ehrich Leonard

Mechanical and Aerospace Engineering
Princeton University
naomi@princeton.edu
www.princeton.edu/~naomi

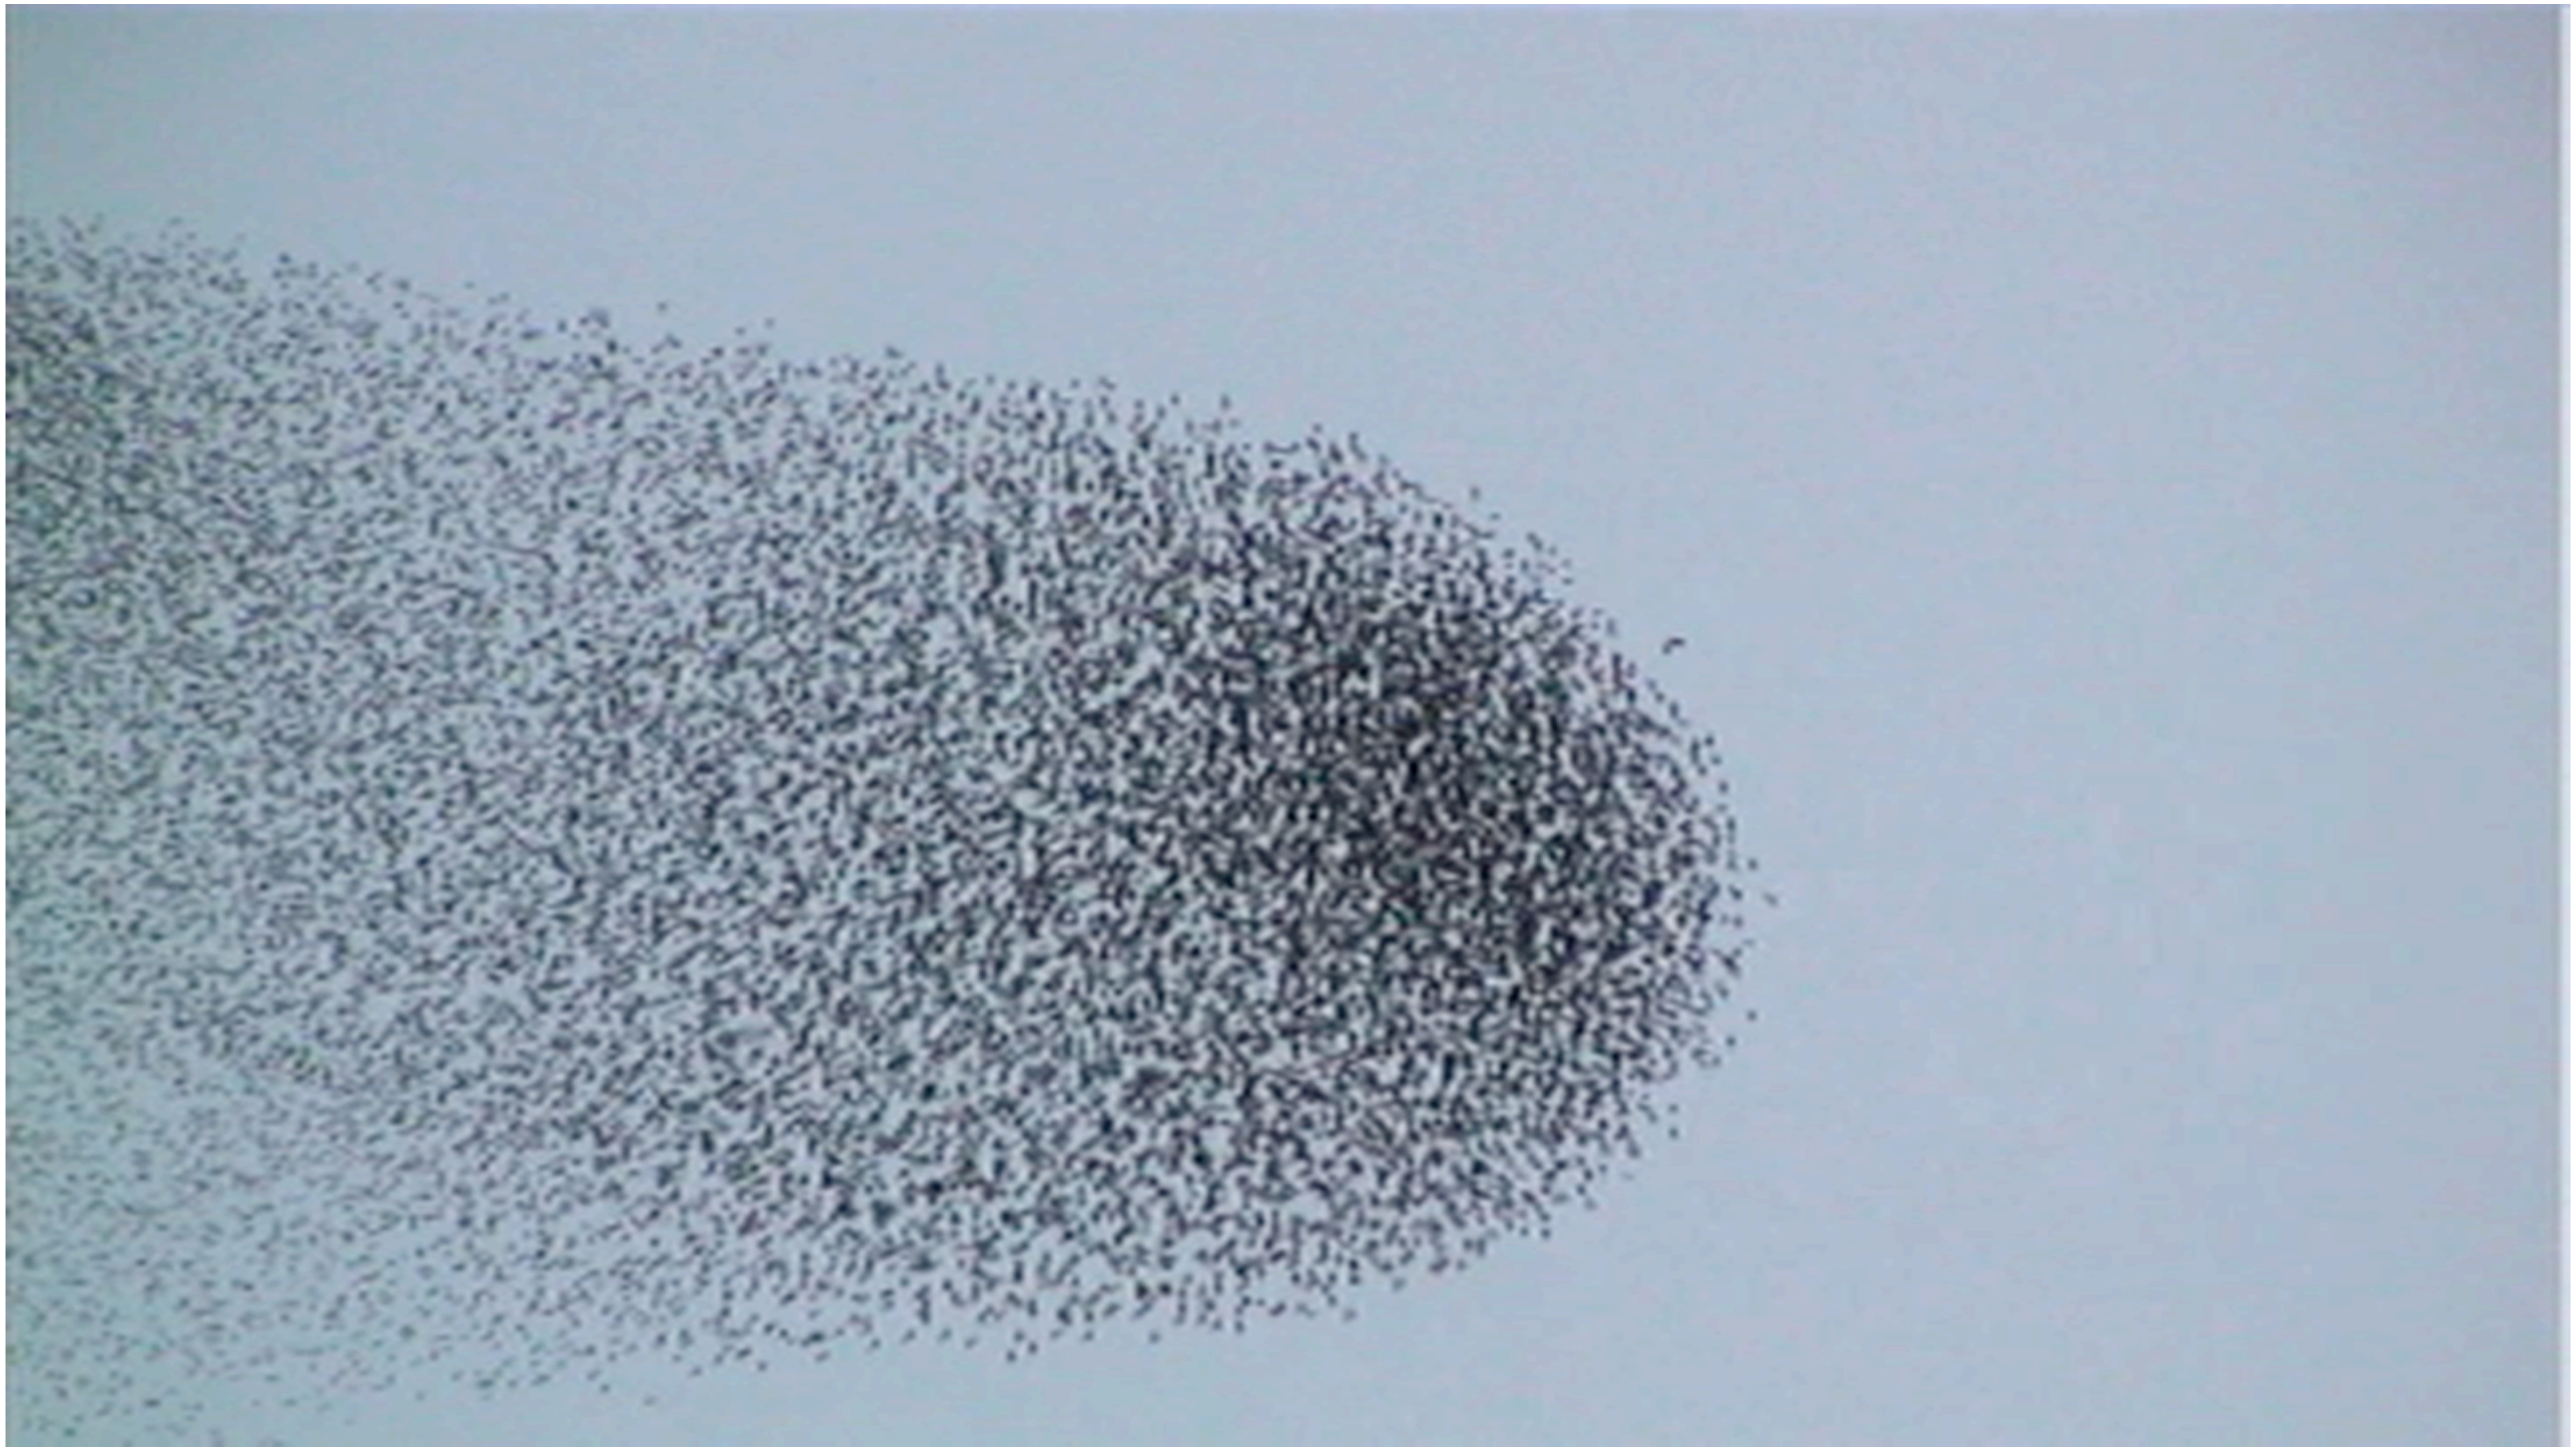

Starling flocks in Rome.
Credit: Claudio Carere, Andrea Cavagna, Irene Giardina et al

Robustness of Collective Decision Making

Dependence of robustness on (directed) graph describing the sensing topology.

1. Role of interconnection topology on L_2 gain in robustness defined by L_2 stability:

L. Scardovi, N.E. Leonard, RoboComm 2009

$$y_k = H\left(u_k + \sum_{j=1}^{N} a_{kj}(y_j - y_k)\right)$$

$H : L_{2e}^m \to L_{2e}^m$ is relaxed co-coercive with constant γ

$$\lambda = \frac{1}{2}\lambda_{\min}[Q(L+L^T)Q^T]$$

$$\|\tilde{y}\|_T \leq \alpha\|\tilde{u}\|_T, \quad \alpha \leq \frac{1}{\lambda+\gamma}$$

2. Role of interconnection topology on H_2 robustness of noisy consensus dynamics.

G. Young, L. Scardovi, N.E. Leonard, ACC 2010

3. Role of interconnection topology on uncertainty in network of decision-making units, each represented by a Drift-Diffusion Model (DDM), accumulating evidence toward a decision.

I. Poulakakis, L. Scardovi, N.E. Leonard, ACC 2010

4. Role of interconnection on uncertainty in steady-state distributions for human decisions in two-alternative choice tasks.

A. Stewart, M. Cao, N.E. Leonard, ACC 2010

H_2 Robustness: Consensus Dynamics with White Noise

G. Young, L. Scardovi and N.E. Leonard, ACC 2010

$$\dot{x}(t) = -Lx(t) + \xi(t)$$

$$x(t), \xi(t) \in \mathbb{R}^N, \;\; E[\xi(t)] = 0, \;\; E[\xi(t)\xi^T(t)] = \tfrac{1}{2} I_N \delta(t-\tau), \;\; E[x(0)\xi^T(\tau)] = 0$$

Let $\tilde{x} = Qx$. Agents make same decision when $x = \bar{x}\mathbf{1} \;\; \Leftrightarrow \;\; \tilde{x} = 0.$

$Q \in \mathbb{R}^{(N-1)\times N}$ where $Q\mathbf{1}_N = 0, \; QQ^T = I_{N-1}, \; Q^T Q = I_N - \frac{1}{N}\mathbf{1}_N^T \mathbf{1}_N.$

Consider reduced dynamics on space orthogonal to span of $\mathbf{1}_N$

$$\dot{\tilde{x}} = -\bar{L}\tilde{x} + Q\xi, \quad \bar{L} = QLQ^T$$

If graph is connected, $-\bar{L}$ is Hurwitz.

H_2 norm as measure of robustness

$$\lim_{t\to\infty} E[\|\tilde{x}(t)\|] = \mathrm{tr}(\Sigma_{ss})^{\frac{1}{2}} \qquad \Sigma_{ss} := \lim_{t\to\infty} E[(\tilde{x}(t)\tilde{x}^T(t))]$$

Xiao, Boyd, Kim 2007

This measure is H_2 norm of reduced dynamics with output $z(t) = I_{N-1}\tilde{x}(t)$,

$$\bar{L}\Sigma_{ss} + \Sigma_{ss}\bar{L}^T = I$$

If Q^TQL is normal then $H_2 = \mathrm{tr}(\Sigma_{ss})^{\frac{1}{2}} = \left(\sum_{i=2}^{N} \frac{1}{2\mathrm{Re}\{\lambda_i\}}\right)^{\frac{1}{2}}$

$0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_N$ eigenvalues of L

Applies to directed graphs:

Directed cycle

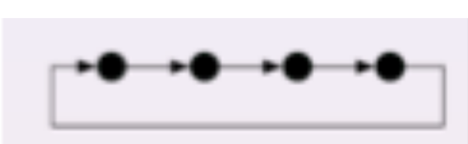

$$H_2 = \sqrt{\frac{N^2-1}{12}}$$

Directed star

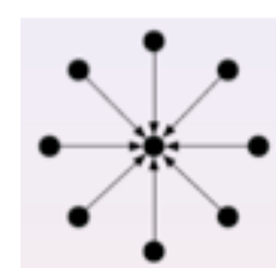

$$H_2 = \sqrt{\frac{N-1}{2}}$$

Comparison

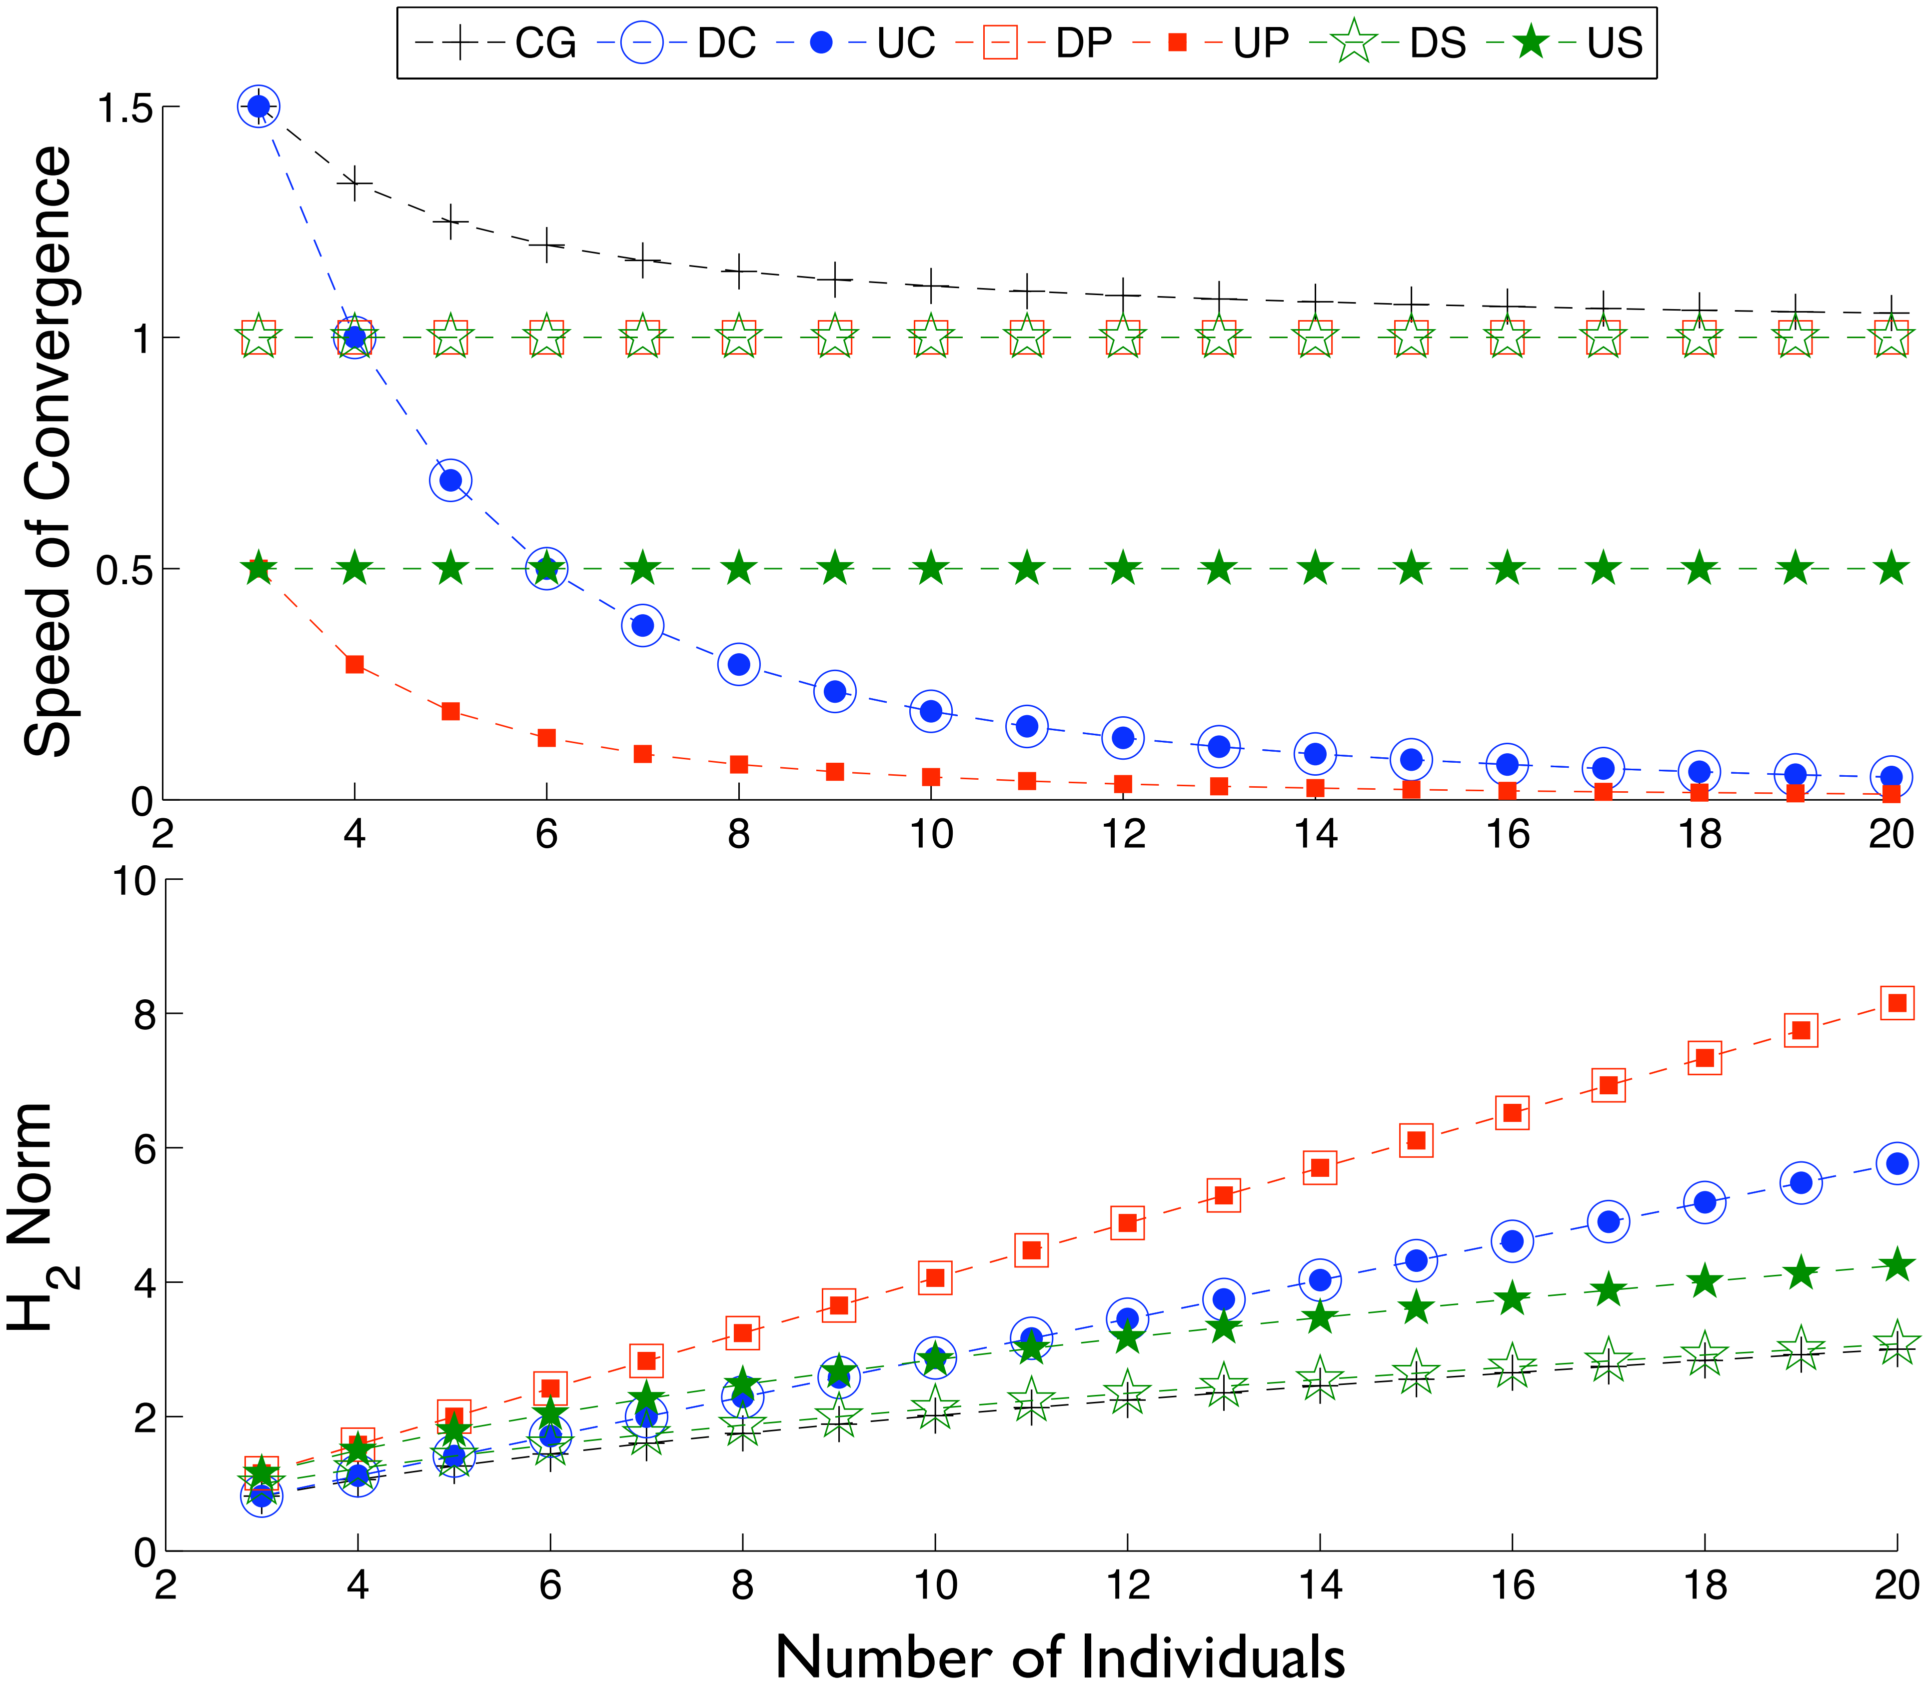

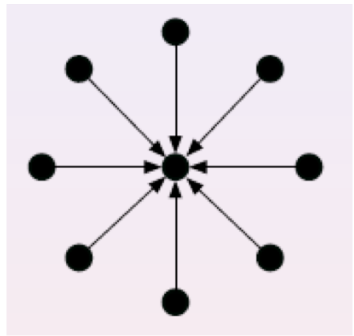

Directed Star (Imploding)

Directed Path

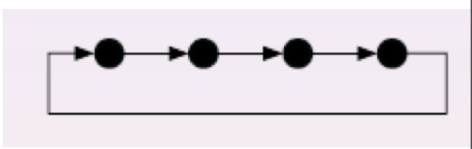

Directed Cycle

H_2 Robustness for Tree Graphs

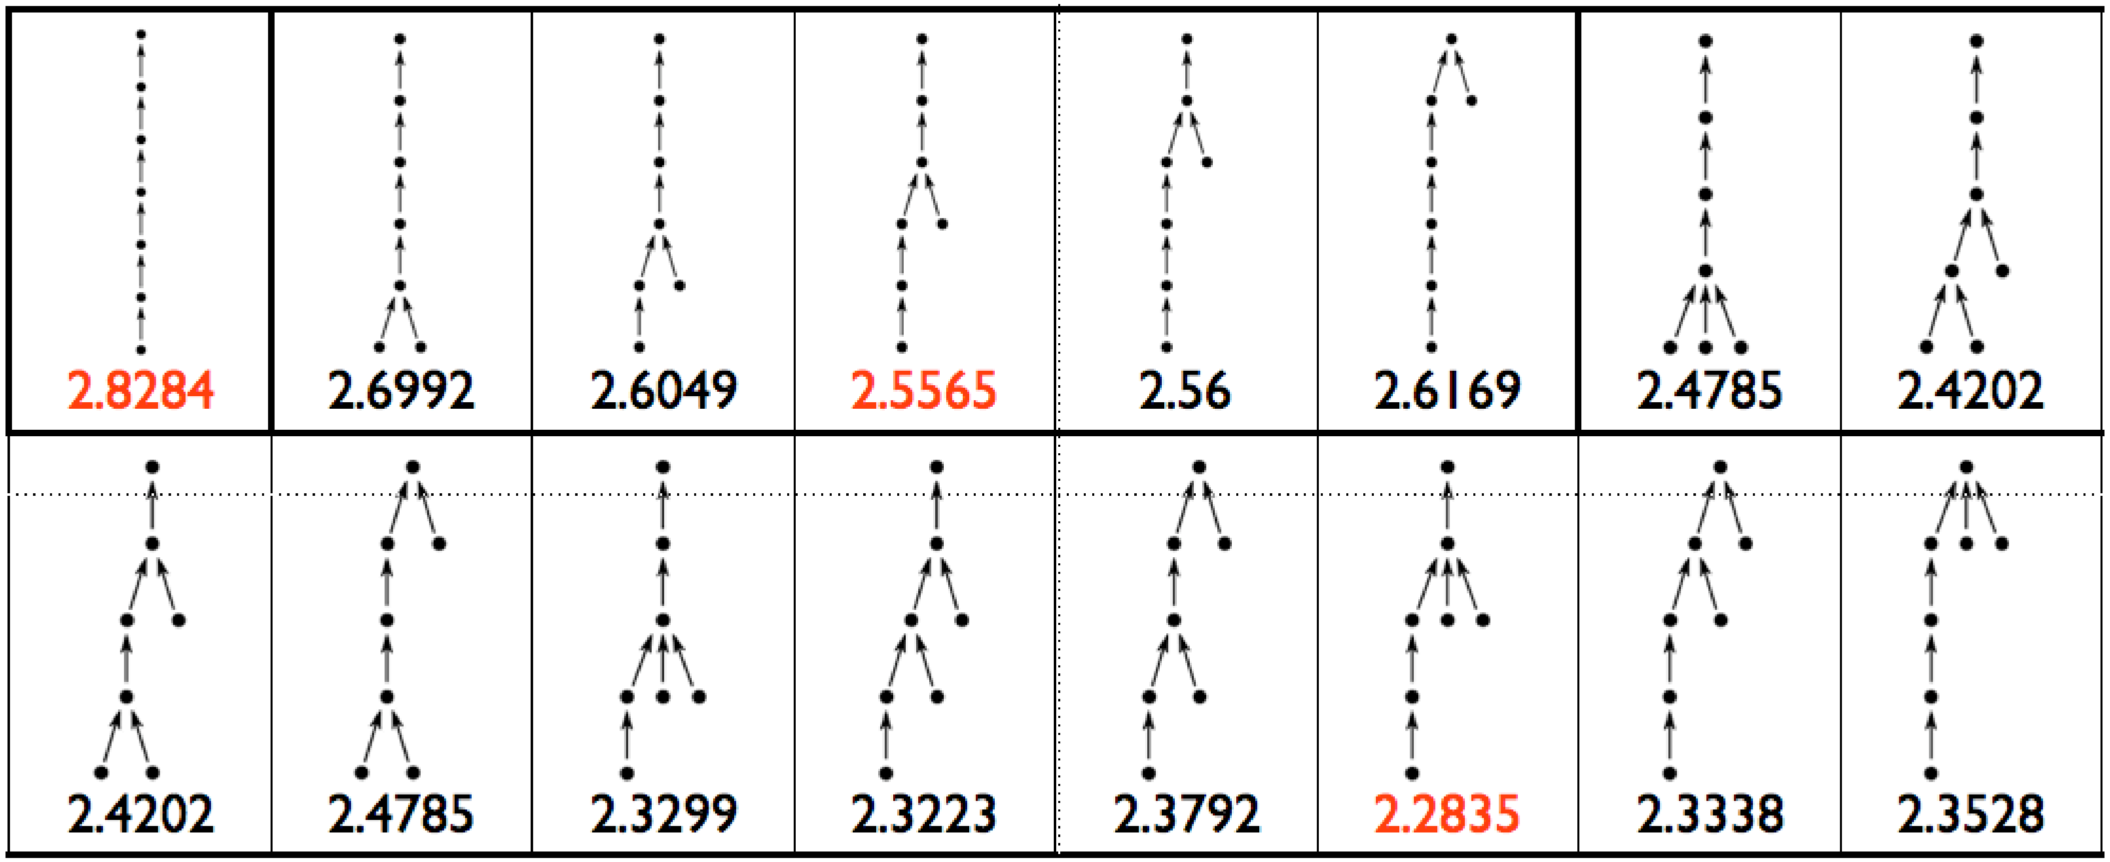

H_2 Robustness in Starling Flocks

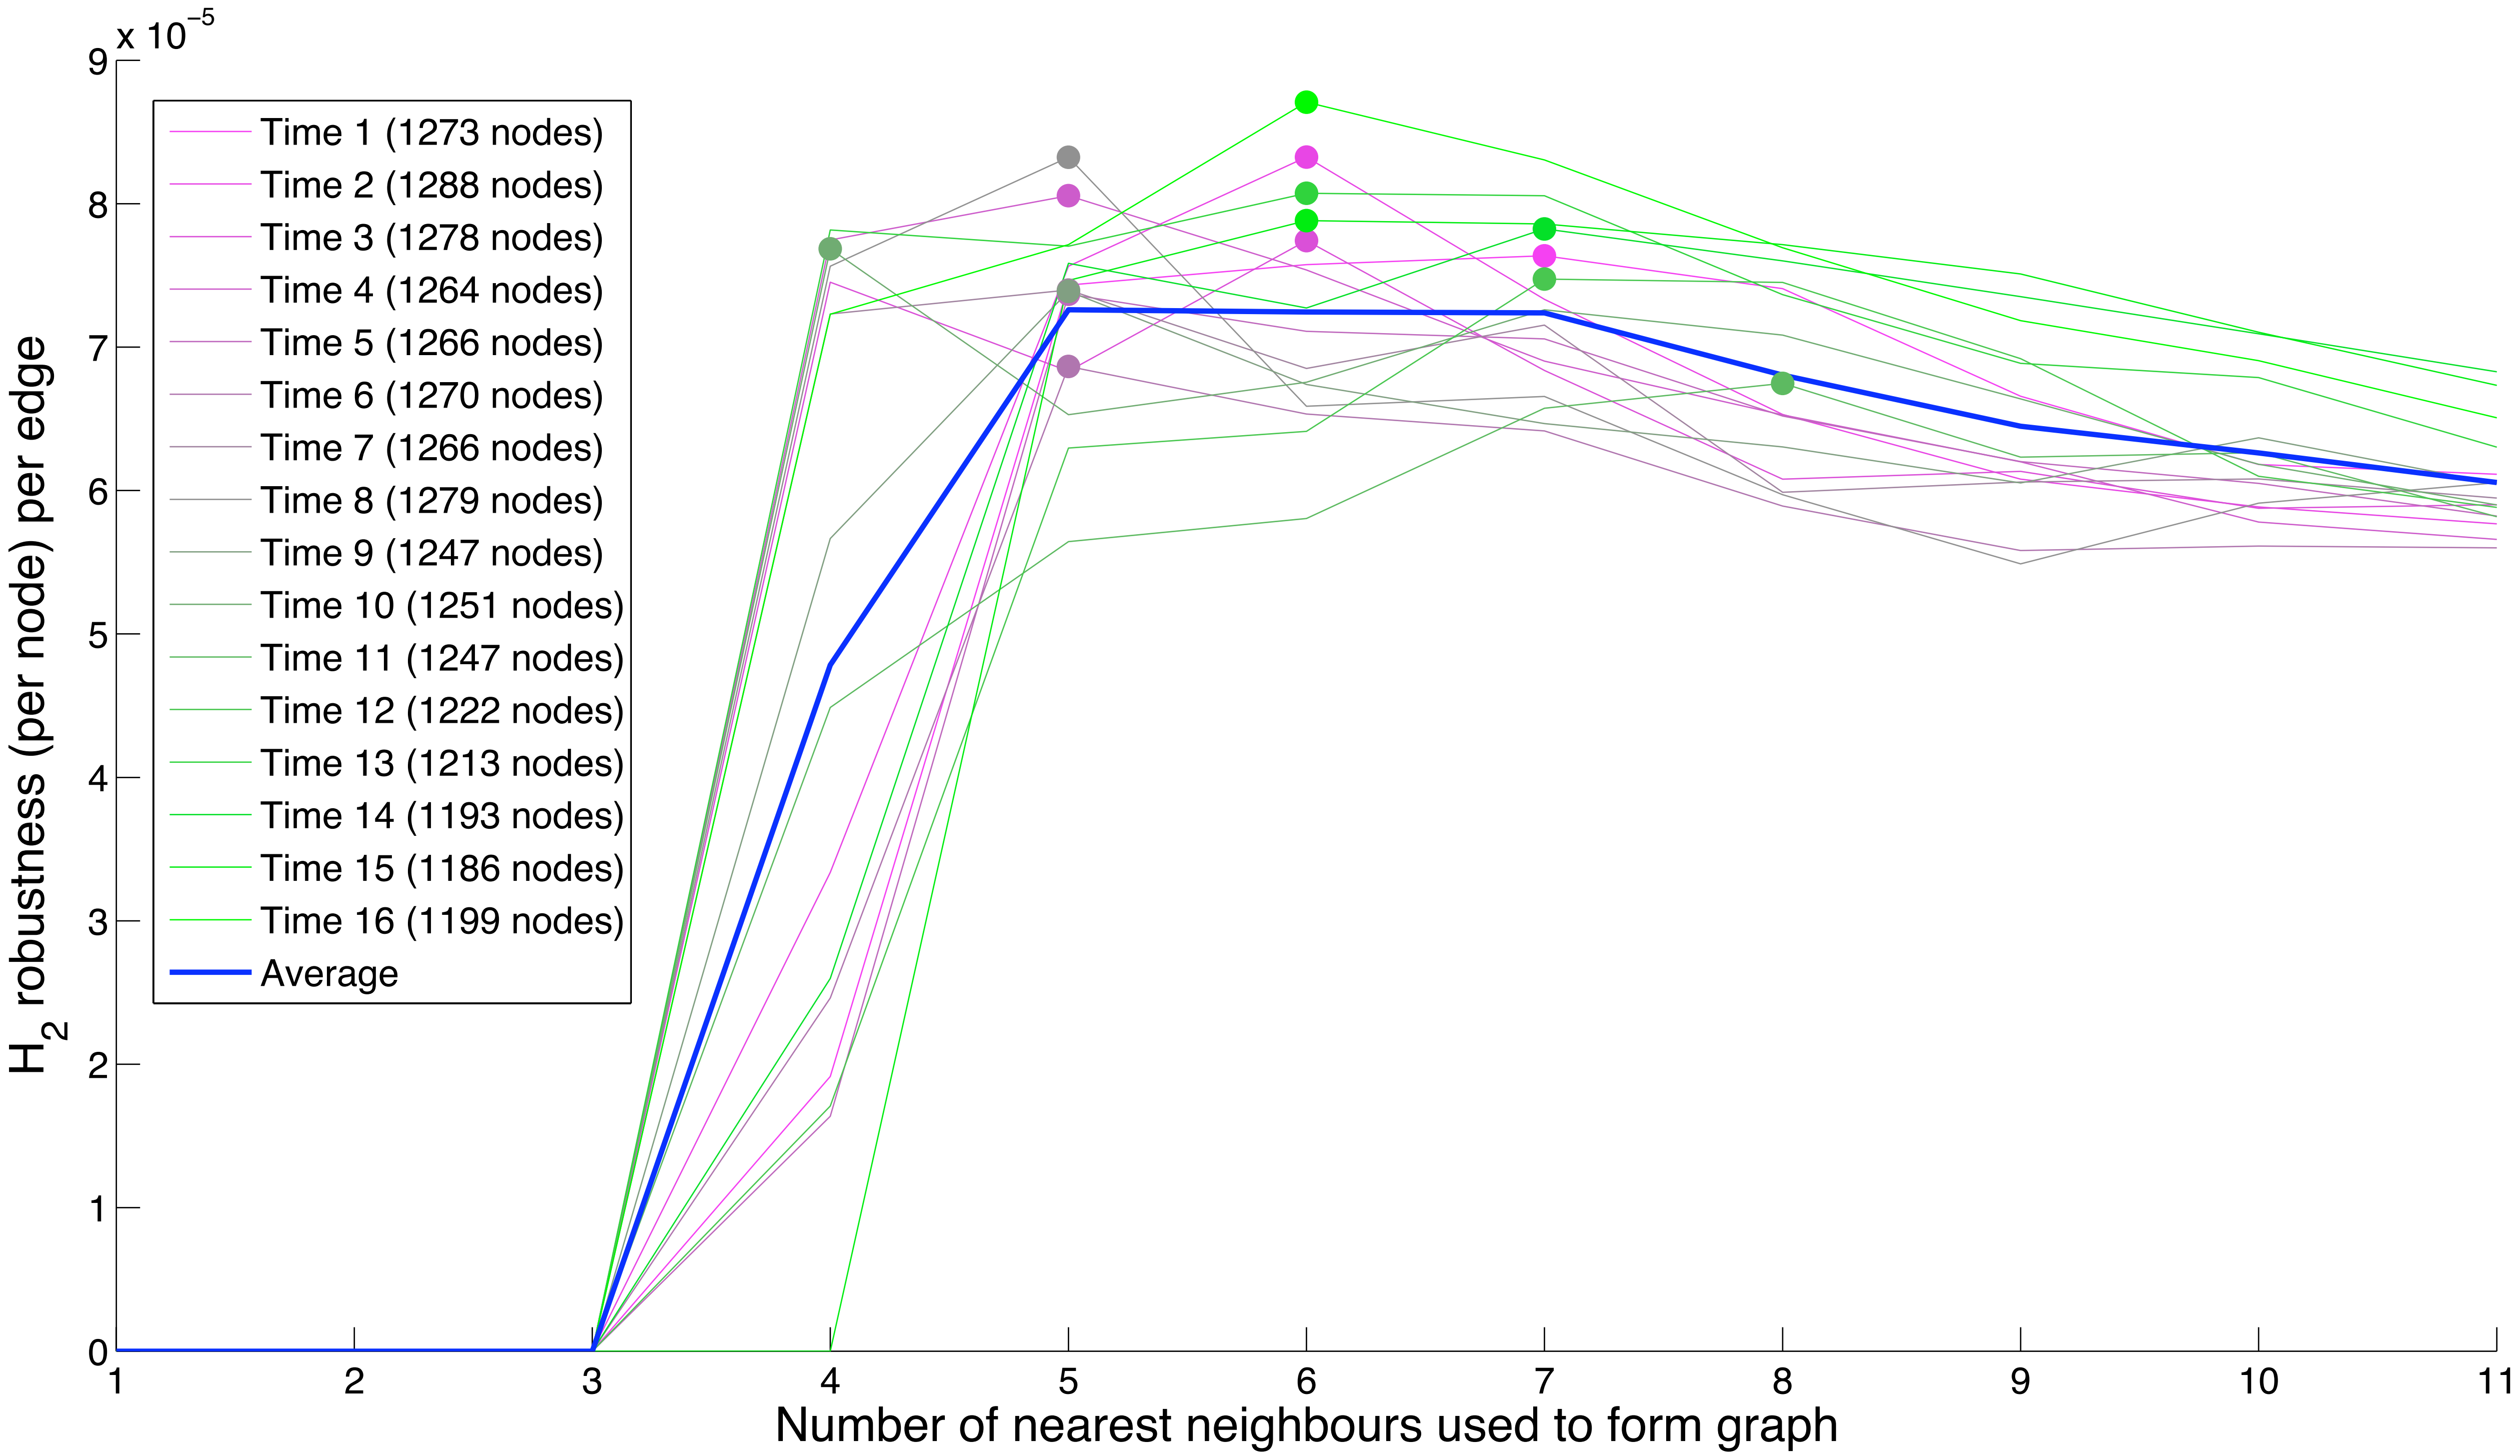

G. Young, L. Scardovi, A. Cavagna, I. Giardina, N. Leonard

Consensus Dynamics with Decision

Consider the case when there is a "correct" decision x^* or choice to be made

$$\dot{x}_k = x^* - x_k + \frac{K}{N}\sum_{j=1}^{N} \alpha_{kj}(x_j - x_k) + \xi_k$$

Example: Agents steering to align with different information about "correct" direction:

$$\dot{\theta}_j = \sin\left(\bar{\theta}_1 - \theta_j\right) + \frac{K_1}{N}\sum_{l=1}^{N} a_{jl}\sin\left(\theta_l - \theta_j\right), \qquad j \in \mathcal{N}_1$$

$$\dot{\theta}_j = \sin\left(\bar{\theta}_2 - \theta_j\right) + \frac{K_1}{N}\sum_{l=1}^{N} a_{jl}\sin\left(\theta_l - \theta_j\right), \qquad j \in \mathcal{N}_2$$

$$\dot{\theta}_j = \frac{K_1}{N}\sum_{l=1}^{N} a_{jl}\sin\left(\theta_l - \theta_j\right), \qquad j \in \mathcal{N}_3$$

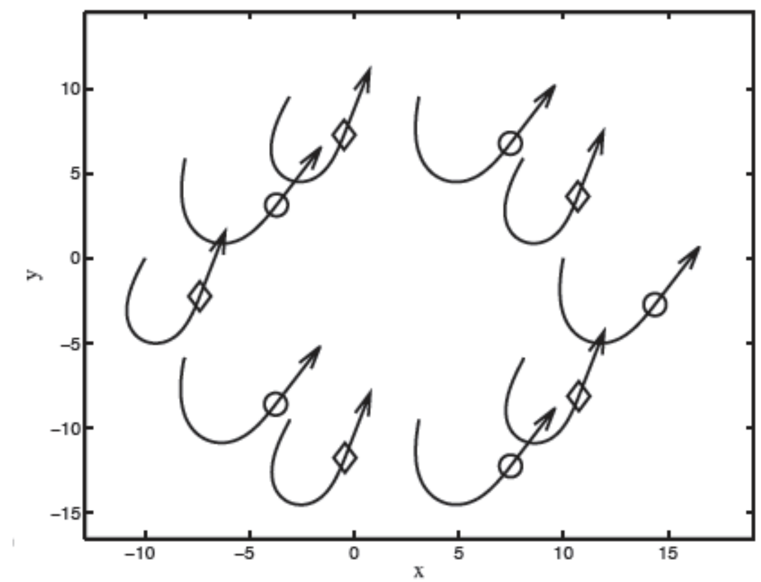

$$\dot{\eta}_{lj} = K_2(\rho_{lj} - r), \qquad \rho_{lj} = \left|\frac{e^{i\theta_l} + e^{i\theta_j}}{2}\right| \qquad l \in \{1,\ldots,N\}, j \in \{1,\ldots,N\}$$

$$a_{lj} = \frac{1}{1 + e^{-\eta_{lj}}}, \qquad l \in \{1,\ldots,N\}, j \in \{1,\ldots,N\},$$

Nabet, Leonard, Couzin, Levin, *J. Nonlinear Science*, 2009
Nabet, PhD Thesis, 2009
Leonard, Nabet, Scardovi, Couzin, Levin, 2010

Role of Interconnection on Uncertainty in Network of Drift-Diffusion Models

Task is to correctly identify a noisy stimulus drawn at random between two known alternatives.

Decision-making unit, represented by a DDM, accumulates evidence according to

Bogacz et al, 2006

$$dx = \beta dt + \sigma dW, \;\; x(0) = 0$$

(continuum limit of Sequential Probability Ratio Test (SPRT))

$x(t)$ is difference in evidence favoring each choice

βdt is increase in evidence supporting correct decision

σdW increments drawn from a Wiener process with standard deviation σ

$$\mathbb{E}[x(t)] = \beta t$$

$$\mathrm{Var}(x(t)) = \sigma^2 t$$

Forced-response protocol: at cue time T_{dec}

If $\beta > 0$ (resp. $\beta < 0$) and $x(T_{\mathrm{dec}}) > 0$ (resp. $x(T_{\mathrm{dec}}) < 0$) correct choice is made

Error rate for $\beta > 0$

$$\mathrm{ER} := \mathbb{P}[x(T_{\mathrm{dec}}) < 0] = \int_{-\infty}^{0} p(x(T_{\mathrm{dec}}), T_{\mathrm{dec}}) dx$$

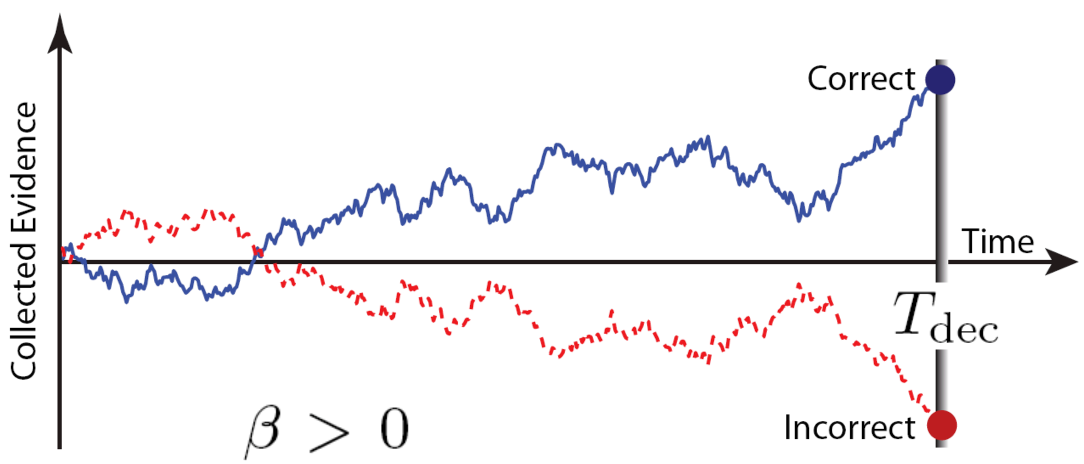

Network of Drift-Diffusion Models

I. Poulakakis, L. Scardovi and N.E. Leonard, ACC 2010

$$dx_k = \left[\beta + \sum_{j=1}^{n} \alpha_{kj}(x_j - x_k)\right] dt + \sigma dW_k \qquad k = 1, 2, \ldots, n$$

Equivalently $$dx = [b - Lx]dt + CdW \qquad b := \beta \mathbf{1} \qquad C := \sigma I$$

$$\mathbb{E}[x(t)] = \beta t \mathbf{1} \qquad \mathrm{Cov}(x(t), x(t)) = \sigma^2 \int_0^t e^{-L(t-\tau)} e^{-L^{\mathrm{T}}(t-\tau)} d\tau$$

In general, each node has different uncertainty in accumulating evidence.
Let decision at cue time be determined by the single node with least uncertainty.
Error rate is marginal probability that this node makes incorrect choice.

$$\mathrm{ER} := \frac{1}{2}\left[1 - \mathrm{erf}\left(\frac{\beta t}{\sqrt{2 \min_k [\mathrm{Var}(x_k(t))]}}\right)\right] \qquad \mathrm{erf}(x) := \frac{2}{\pi} \int_0^x e^{-u^2} du$$

Only the variance influences the error rate.

Undirected Path Graphs

Diagonal elements of the covariance matrix

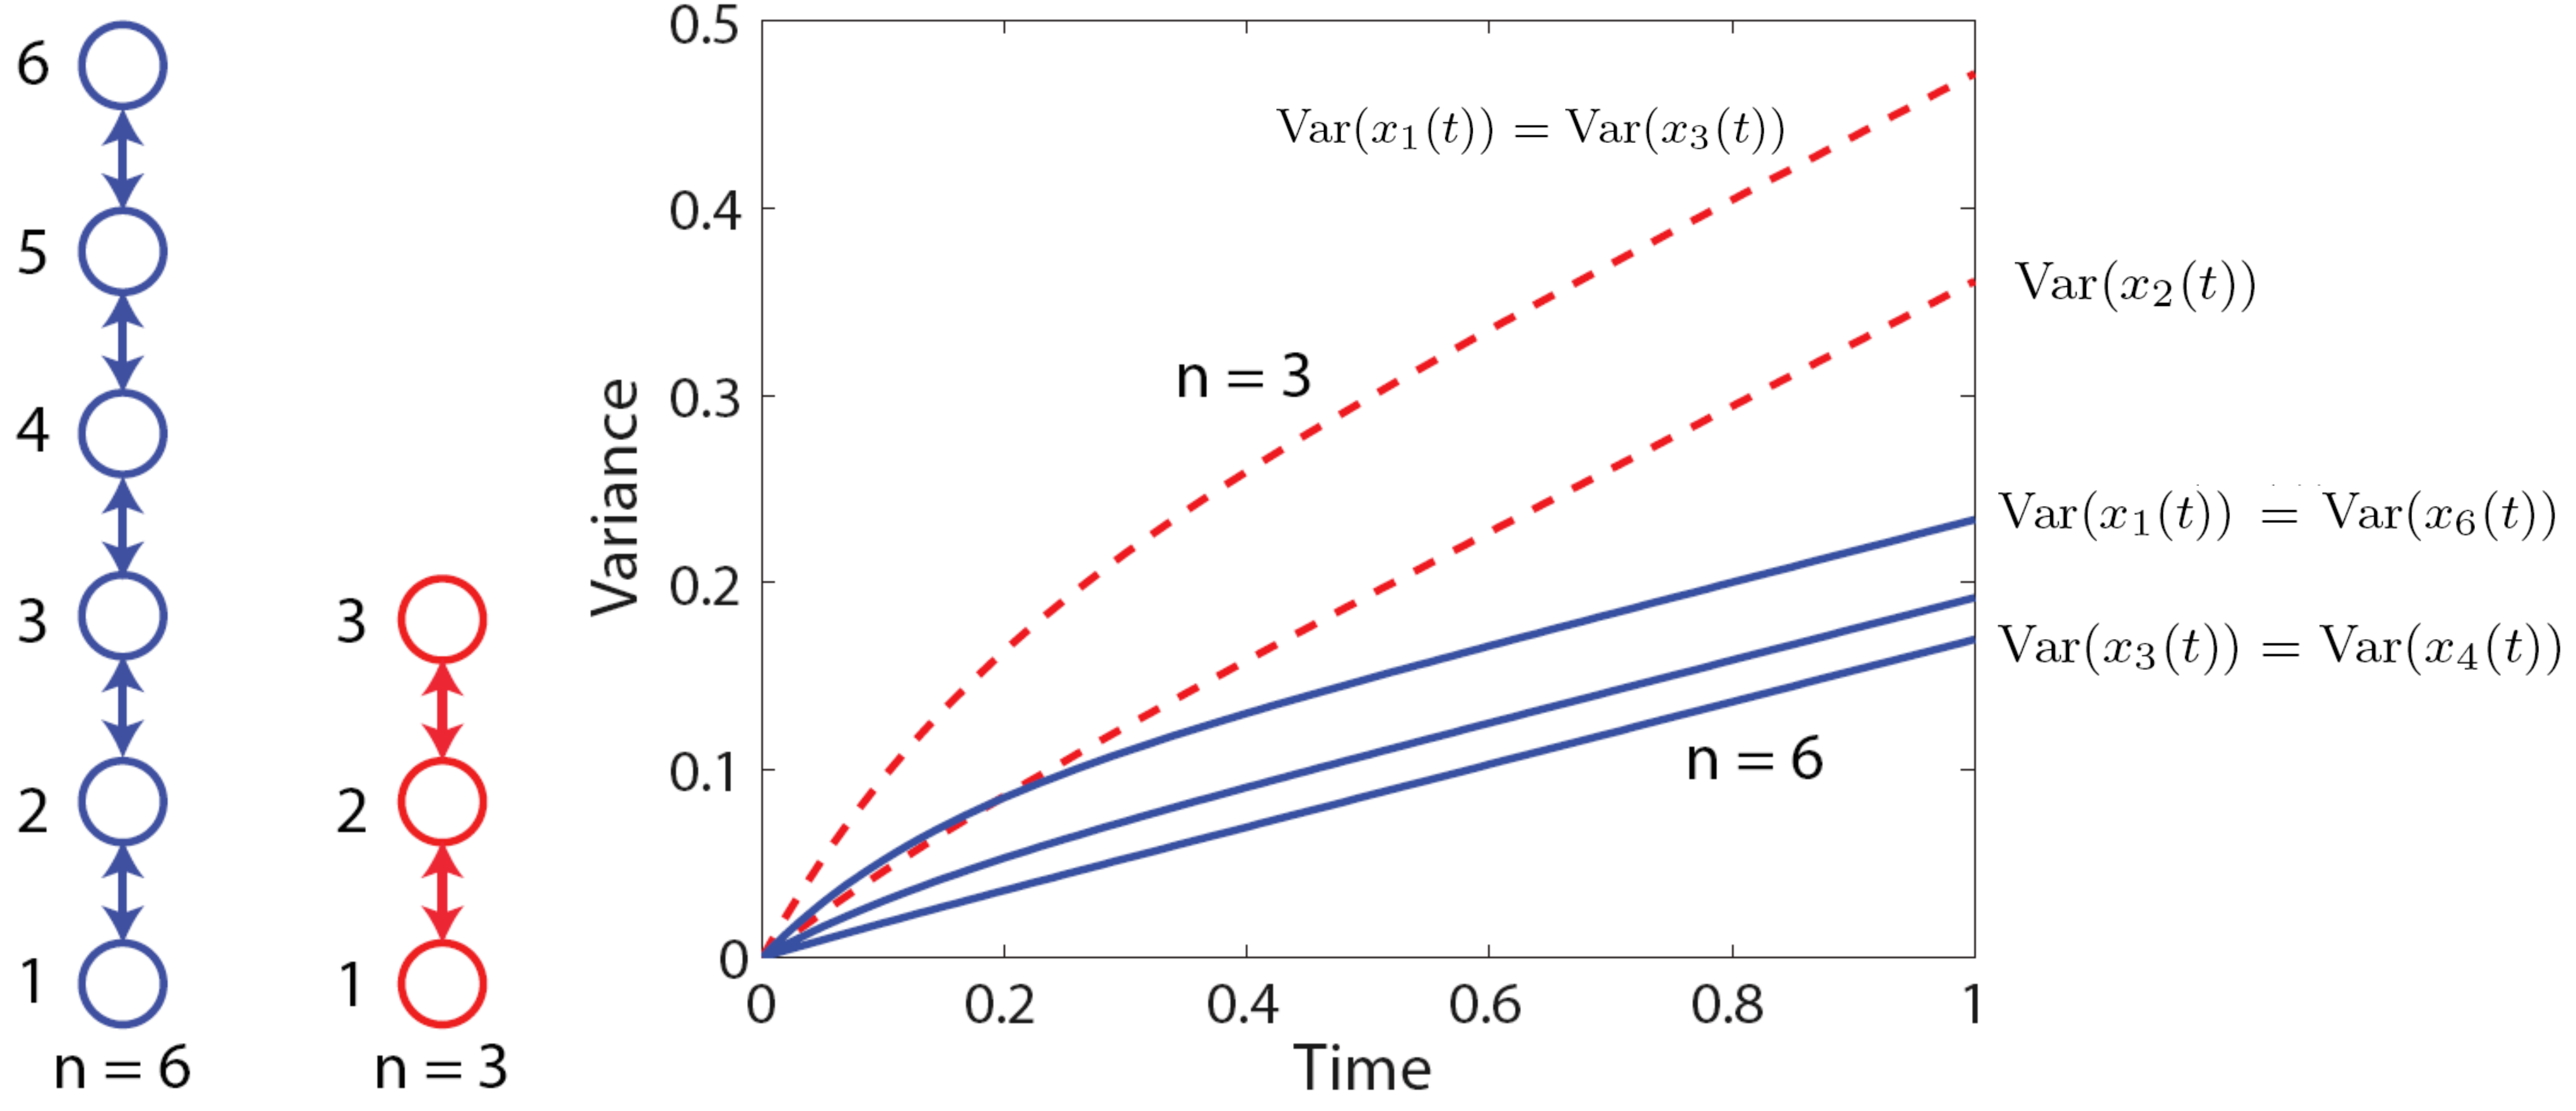

Graph Comparison

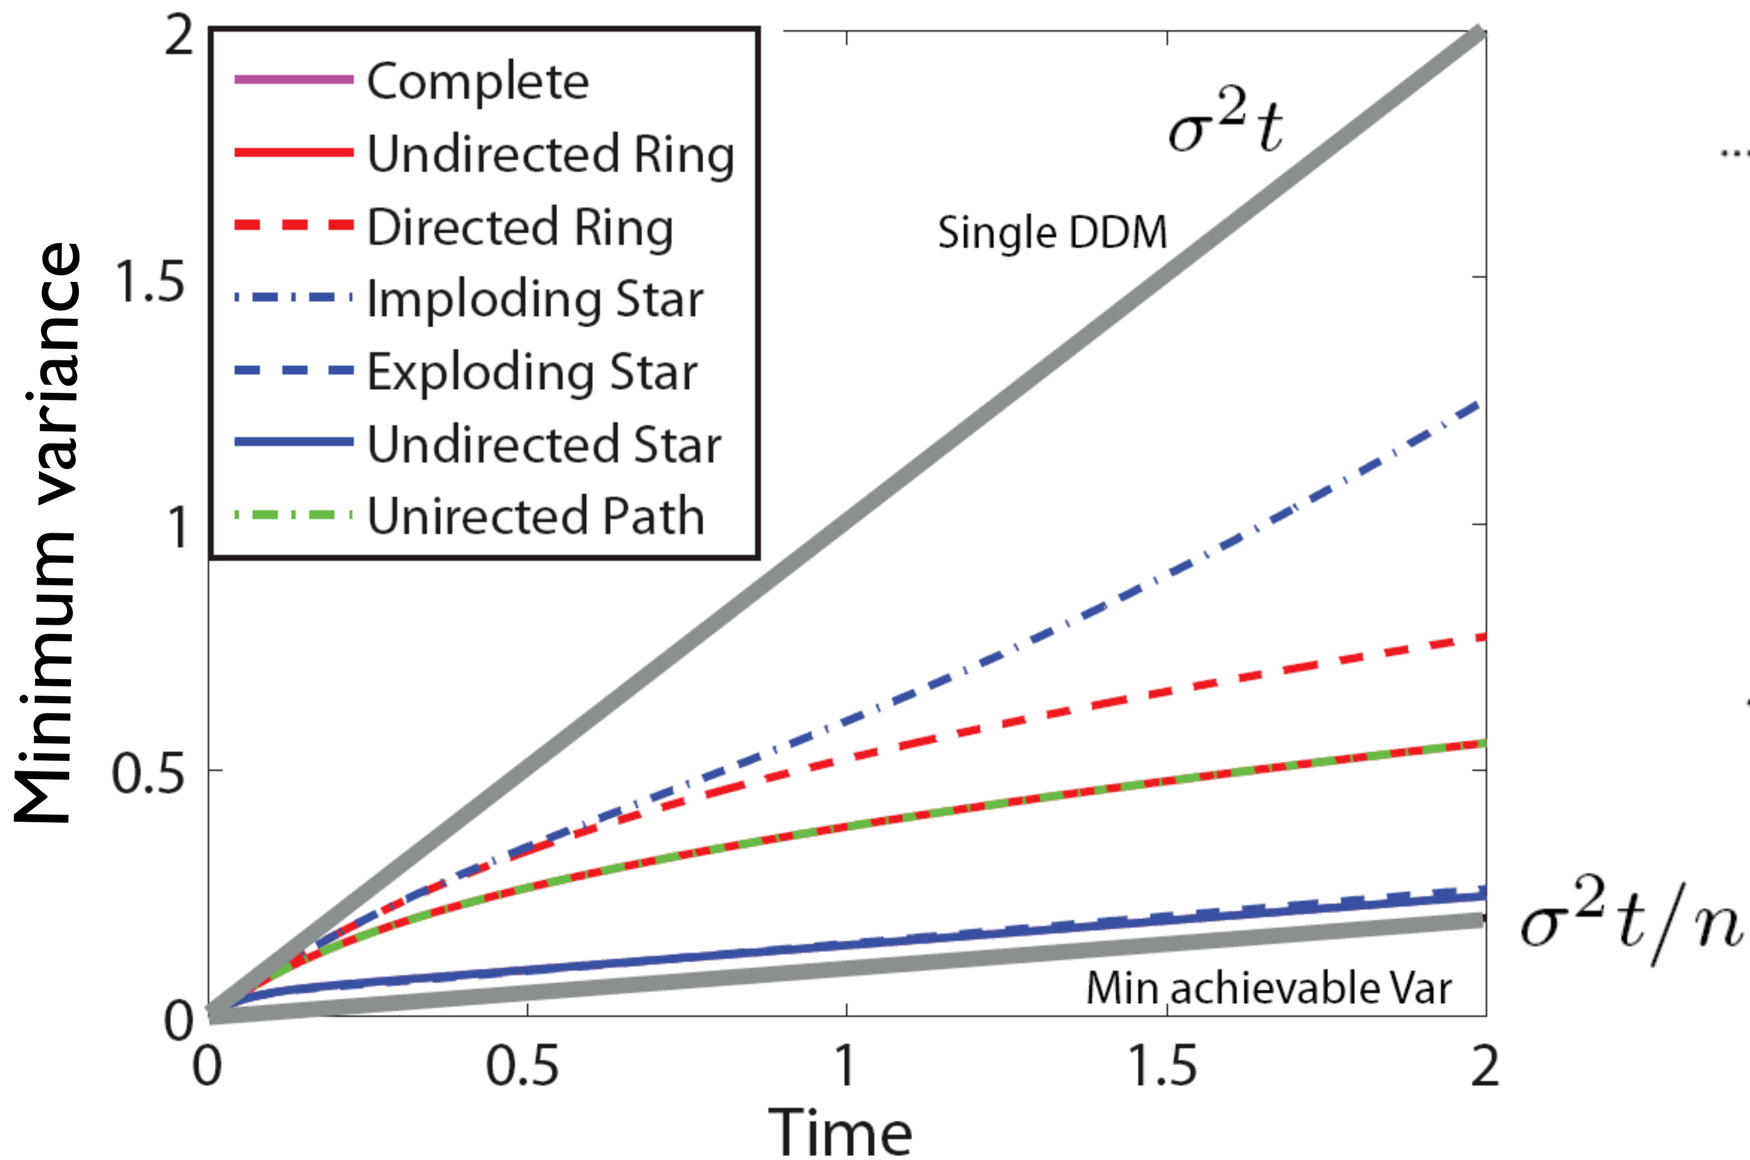

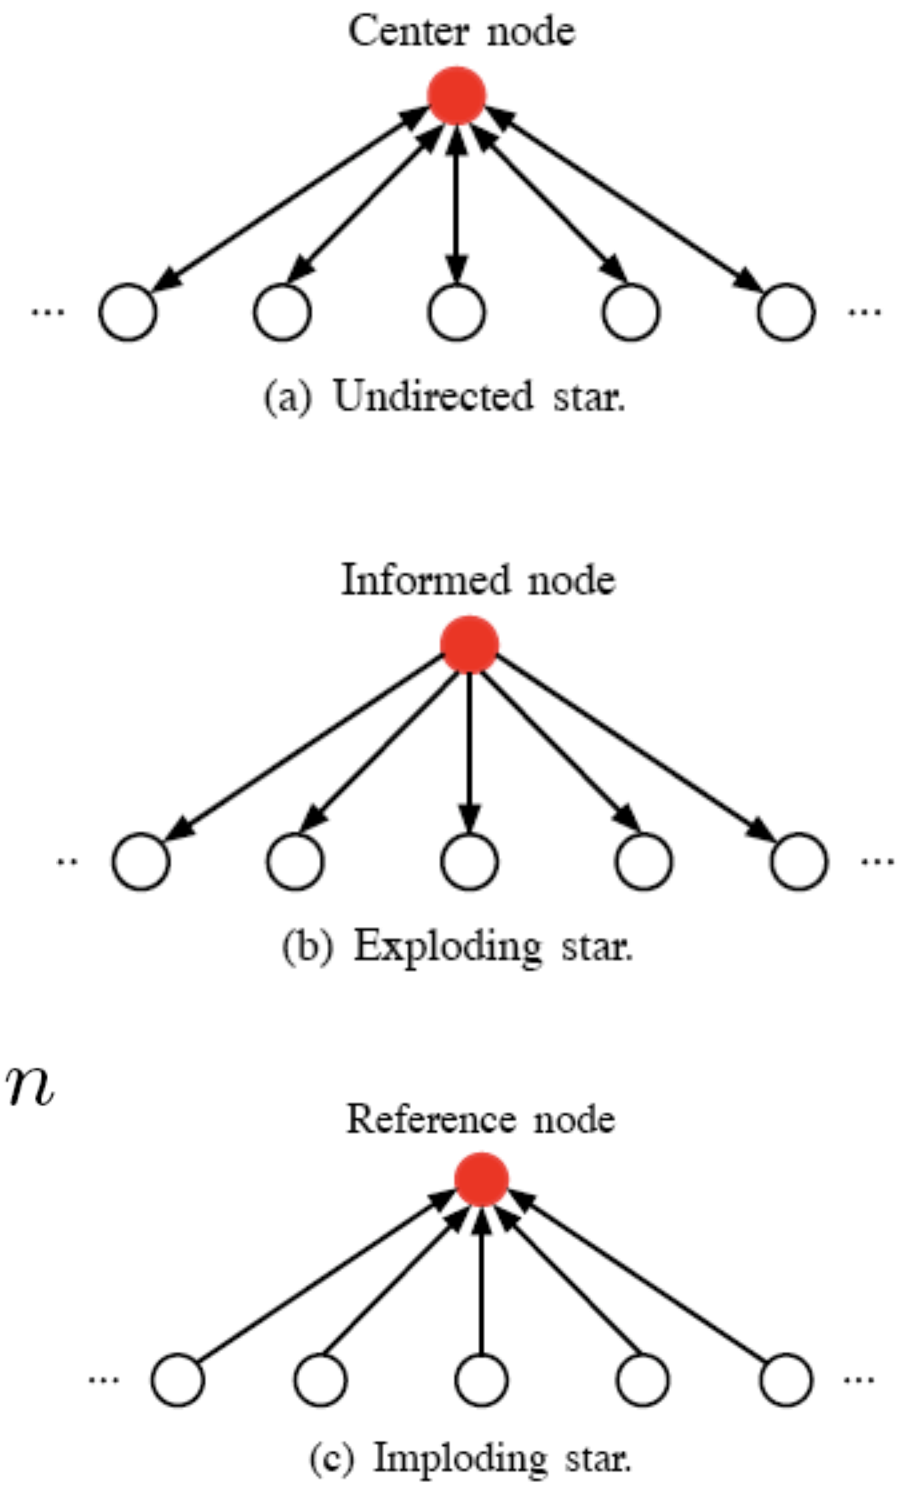

$n = 10$

Human Decision-Making Models

In human-in-the-loop systems, humans confronted repeatedly with decision-making problems in which, having observed system performance, they must choose between two or more alternatives to maintain or improve performance.

Two-alternative forced choice (TAFC) tasks have been used extensively in psychology literature to investigate human decision-making behavior. Montague and Berns, Neuron 2002

The drift diffusion model (DDM) enjoys widespread use to model human decision-making TAFC tasks: DDM has been successfully fit to behavioral and neural data.

Egelman et al 1998, Ratcliff et al 1999, Gold and Shadlen 2001, Schall 2001.

As part of a multi-disciplinary team to investigate decision dynamics in mixed teams of humans and robots, colleagues at Princeton are running social TAFC tasks.

Nedic, Tomlin, Holmes, Prentice, Cohen, 2009

We study role of information passing after every choice on steady distribution of choice sequences.

Stewart, Cao, Leonard, ACC 2010

The Two-Alternative Forced-Choice Task

Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

The Two-Alternative Forced-Choice Task

Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

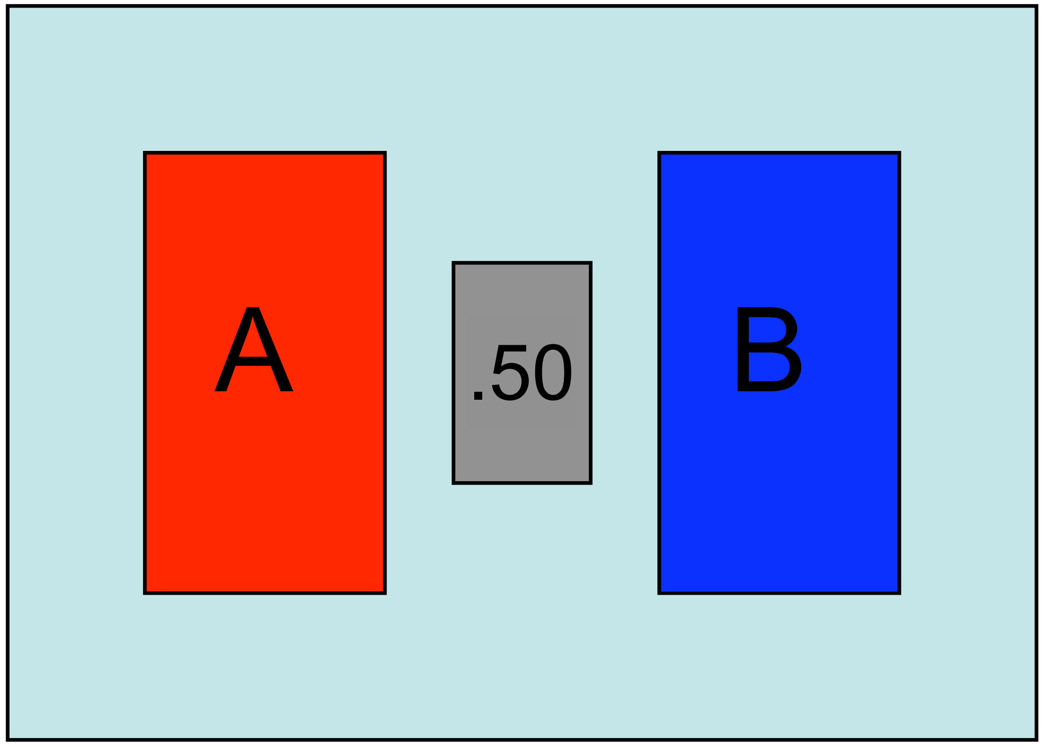

The Two-Alternative Forced-Choice Task

Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

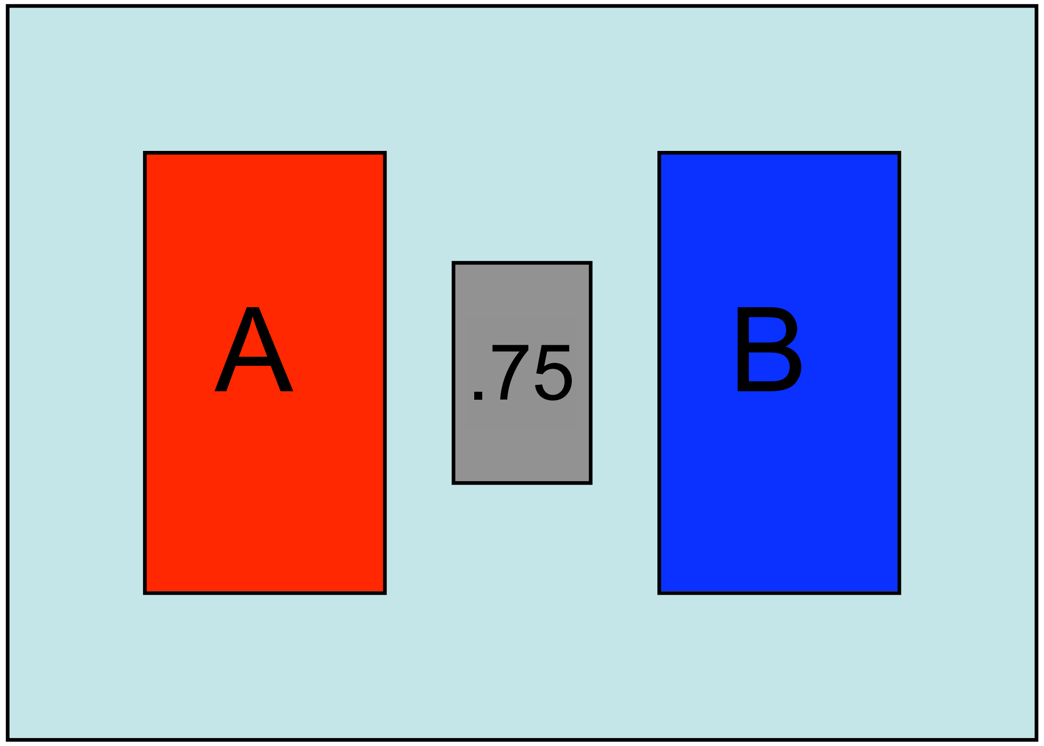

The Two-Alternative Forced-Choice Task

Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

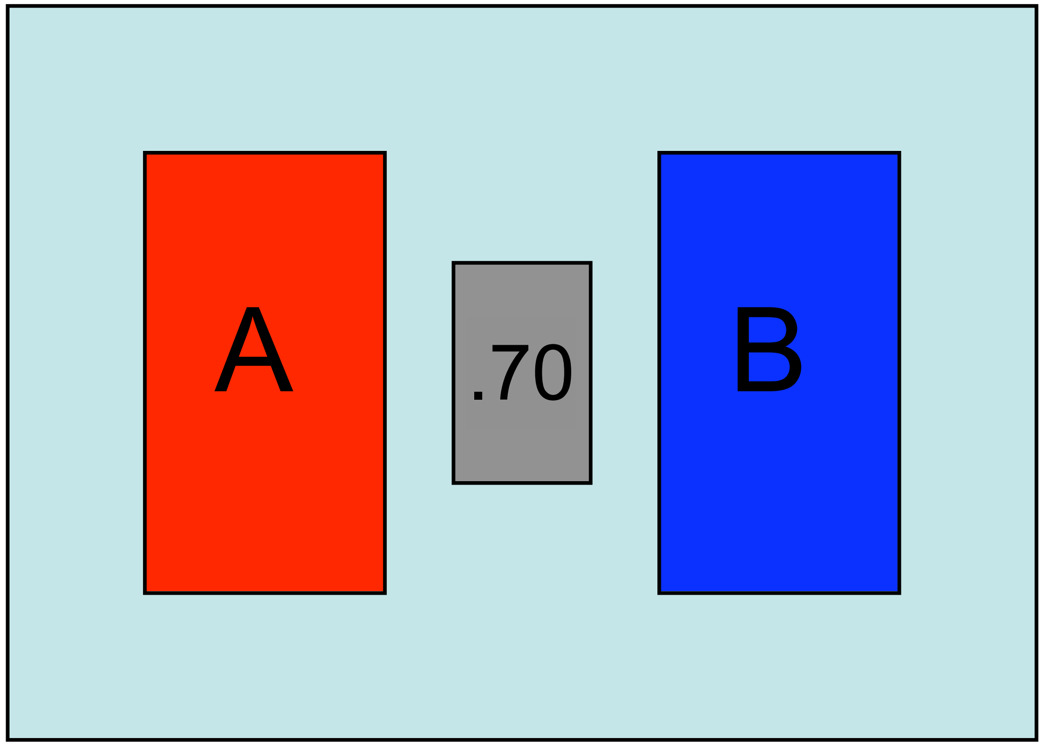

The Two-Alternative Forced-Choice Task

Experimental Studies

Human subject

1. Picks choice “A” or “B”
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

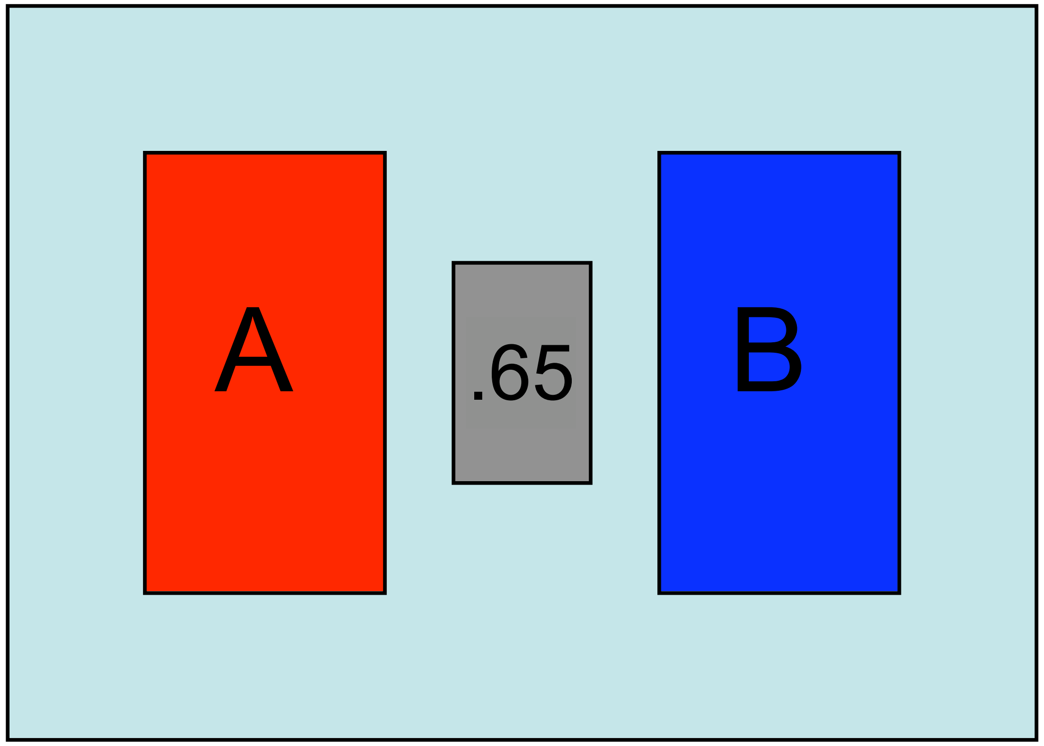

The Two-Alternative Forced-Choice Task
Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

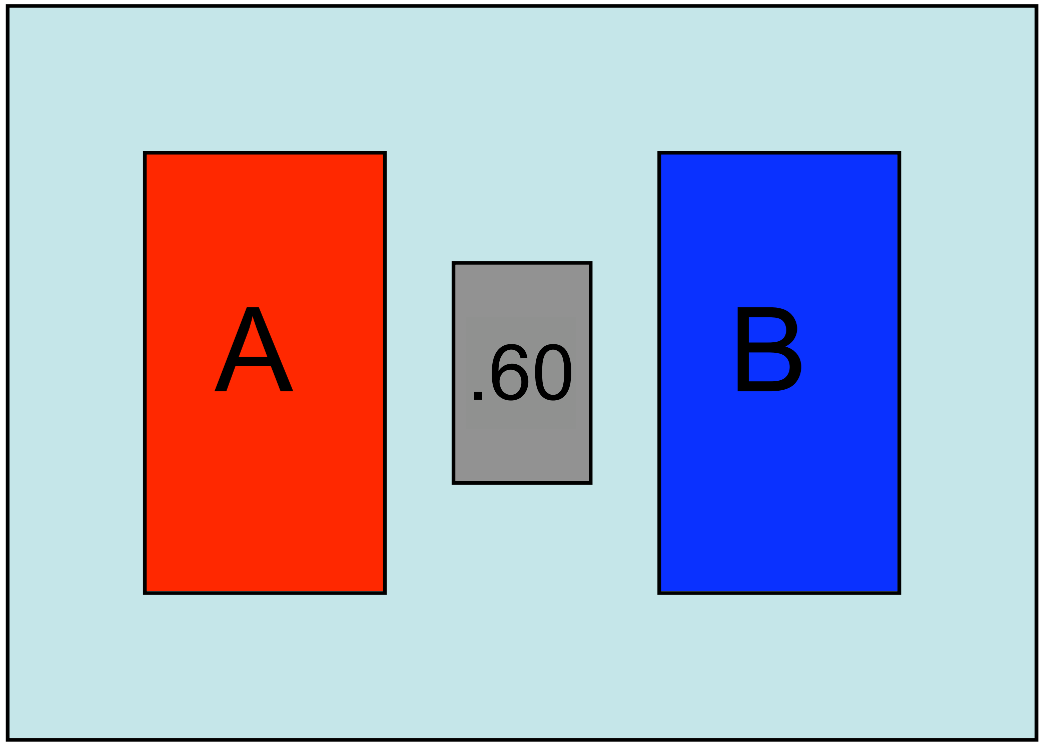

The Two-Alternative Forced-Choice Task
Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

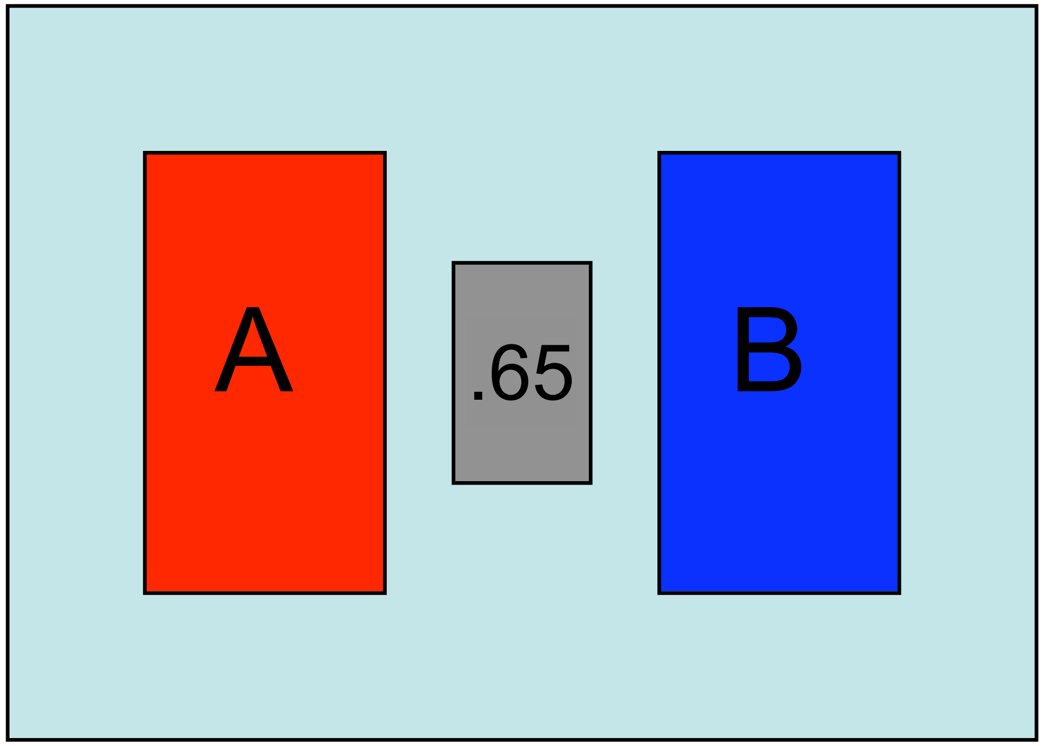

The Two-Alternative Forced-Choice Task
Experimental Studies

Human subject

1. Picks choice "A" or "B"
2. Receives reward
3. Returns to Step 1.

Payment proportional to sum of rewards

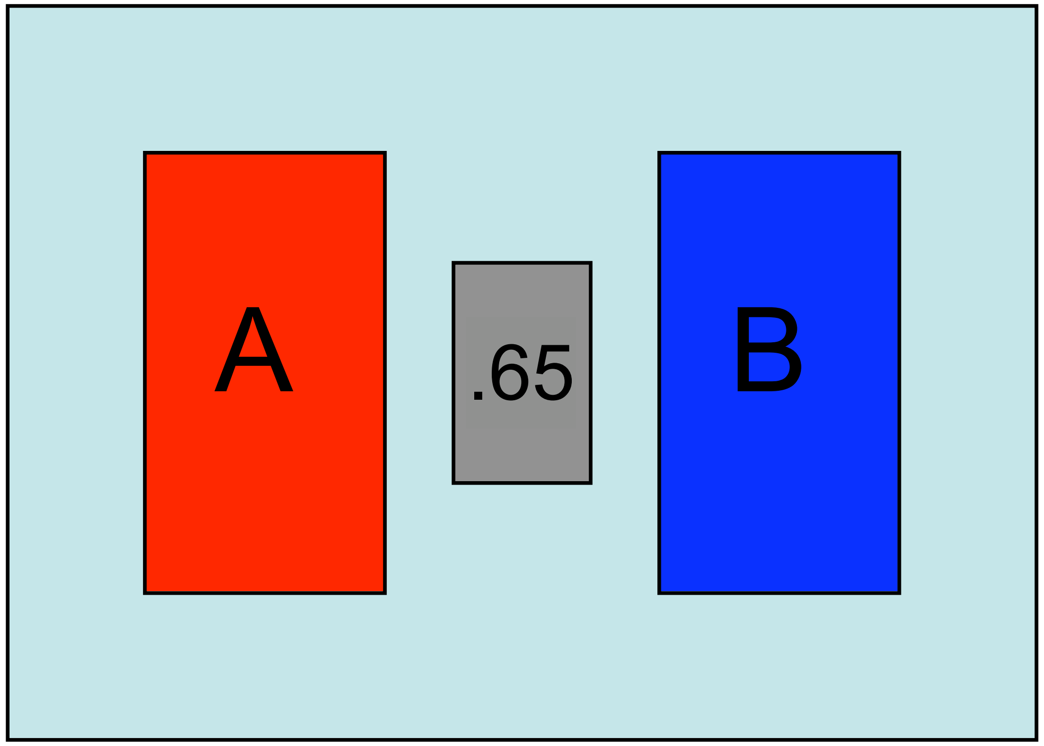

With social feedback

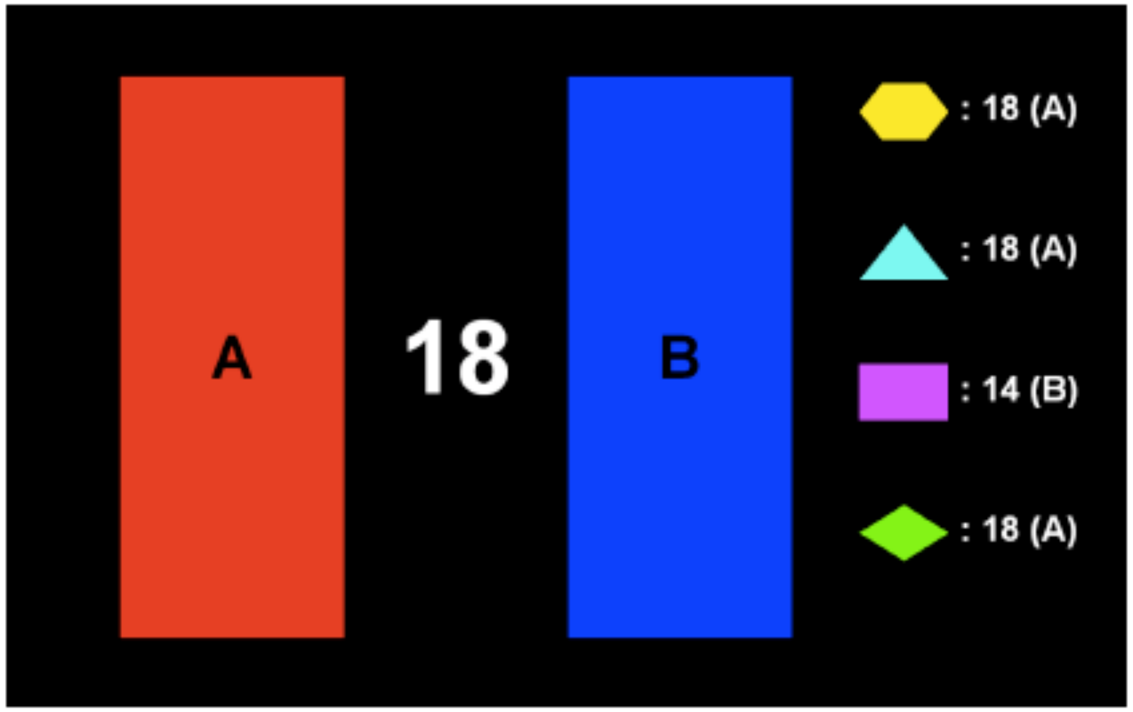

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

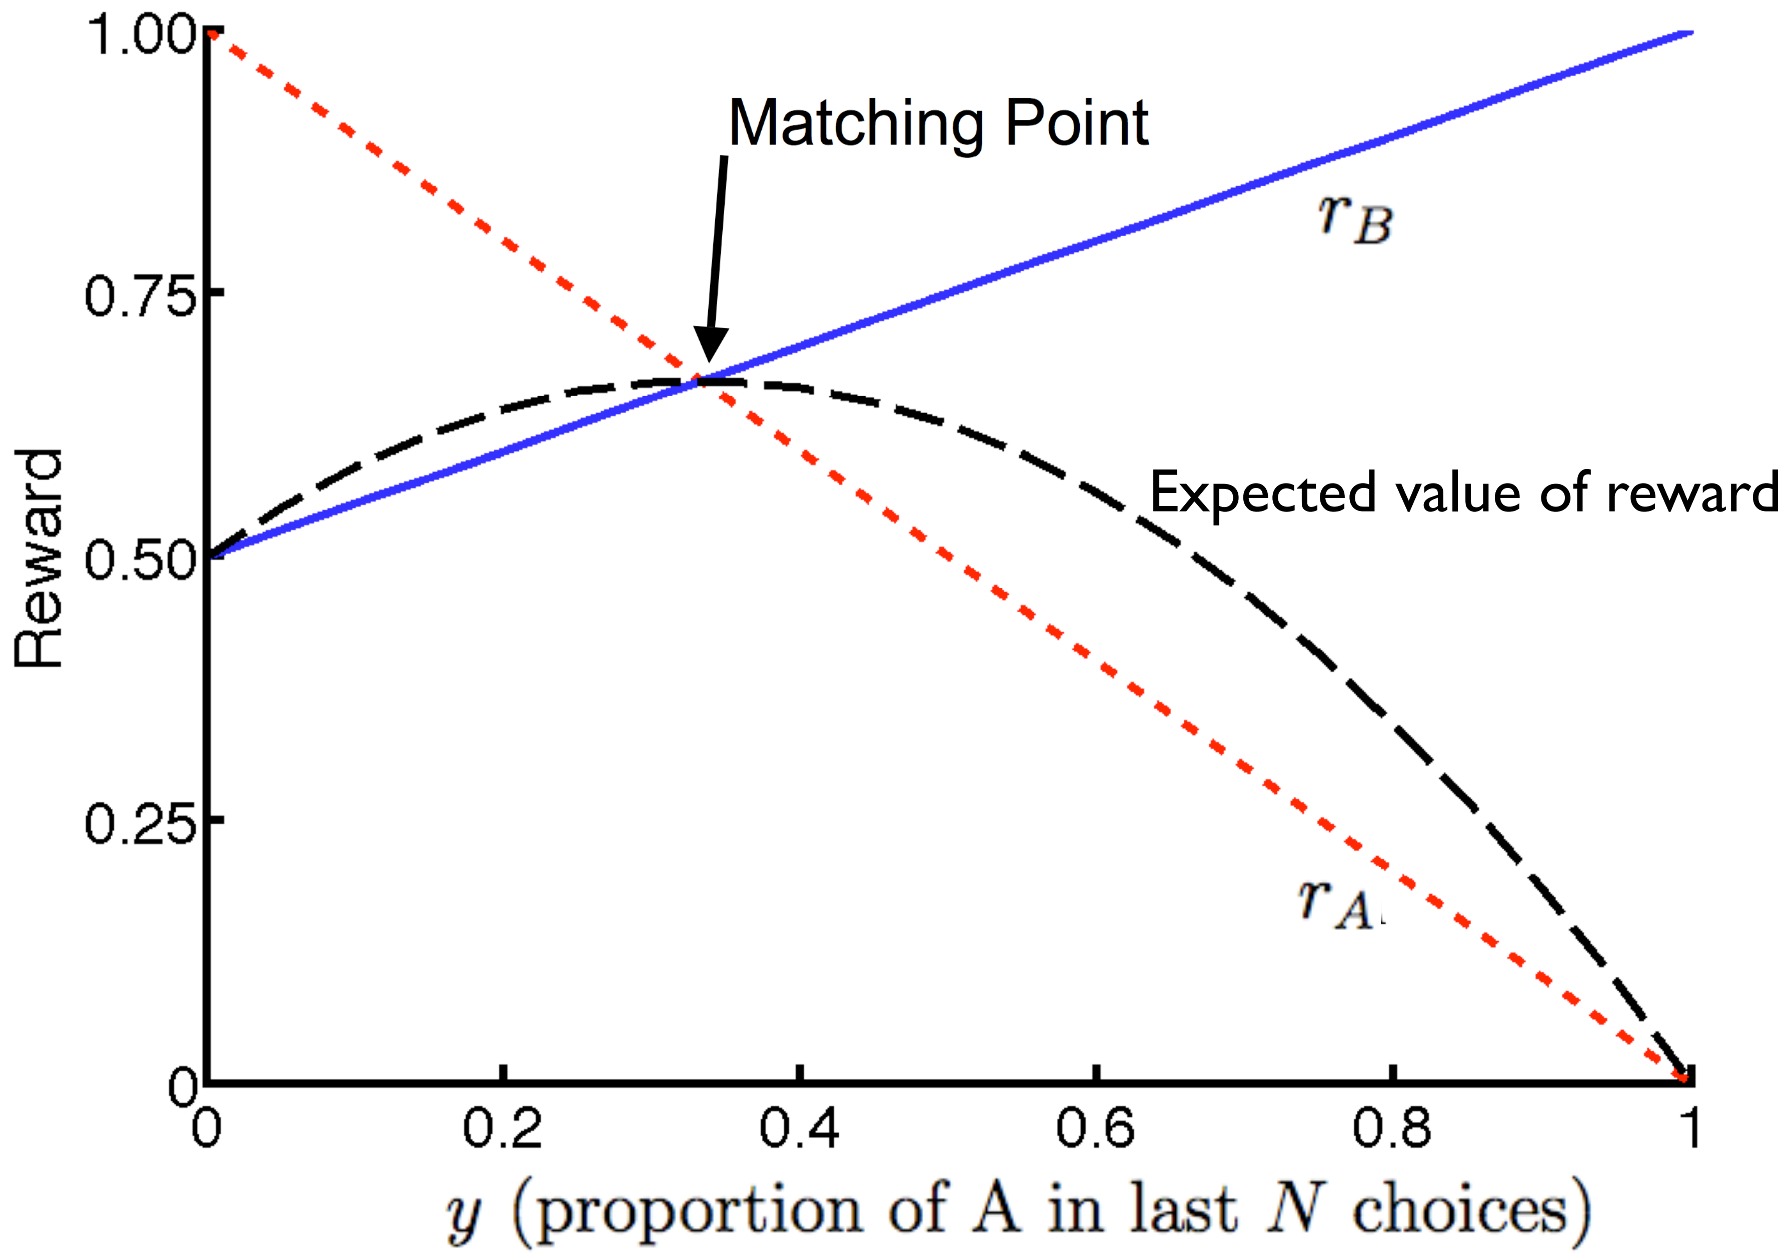

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

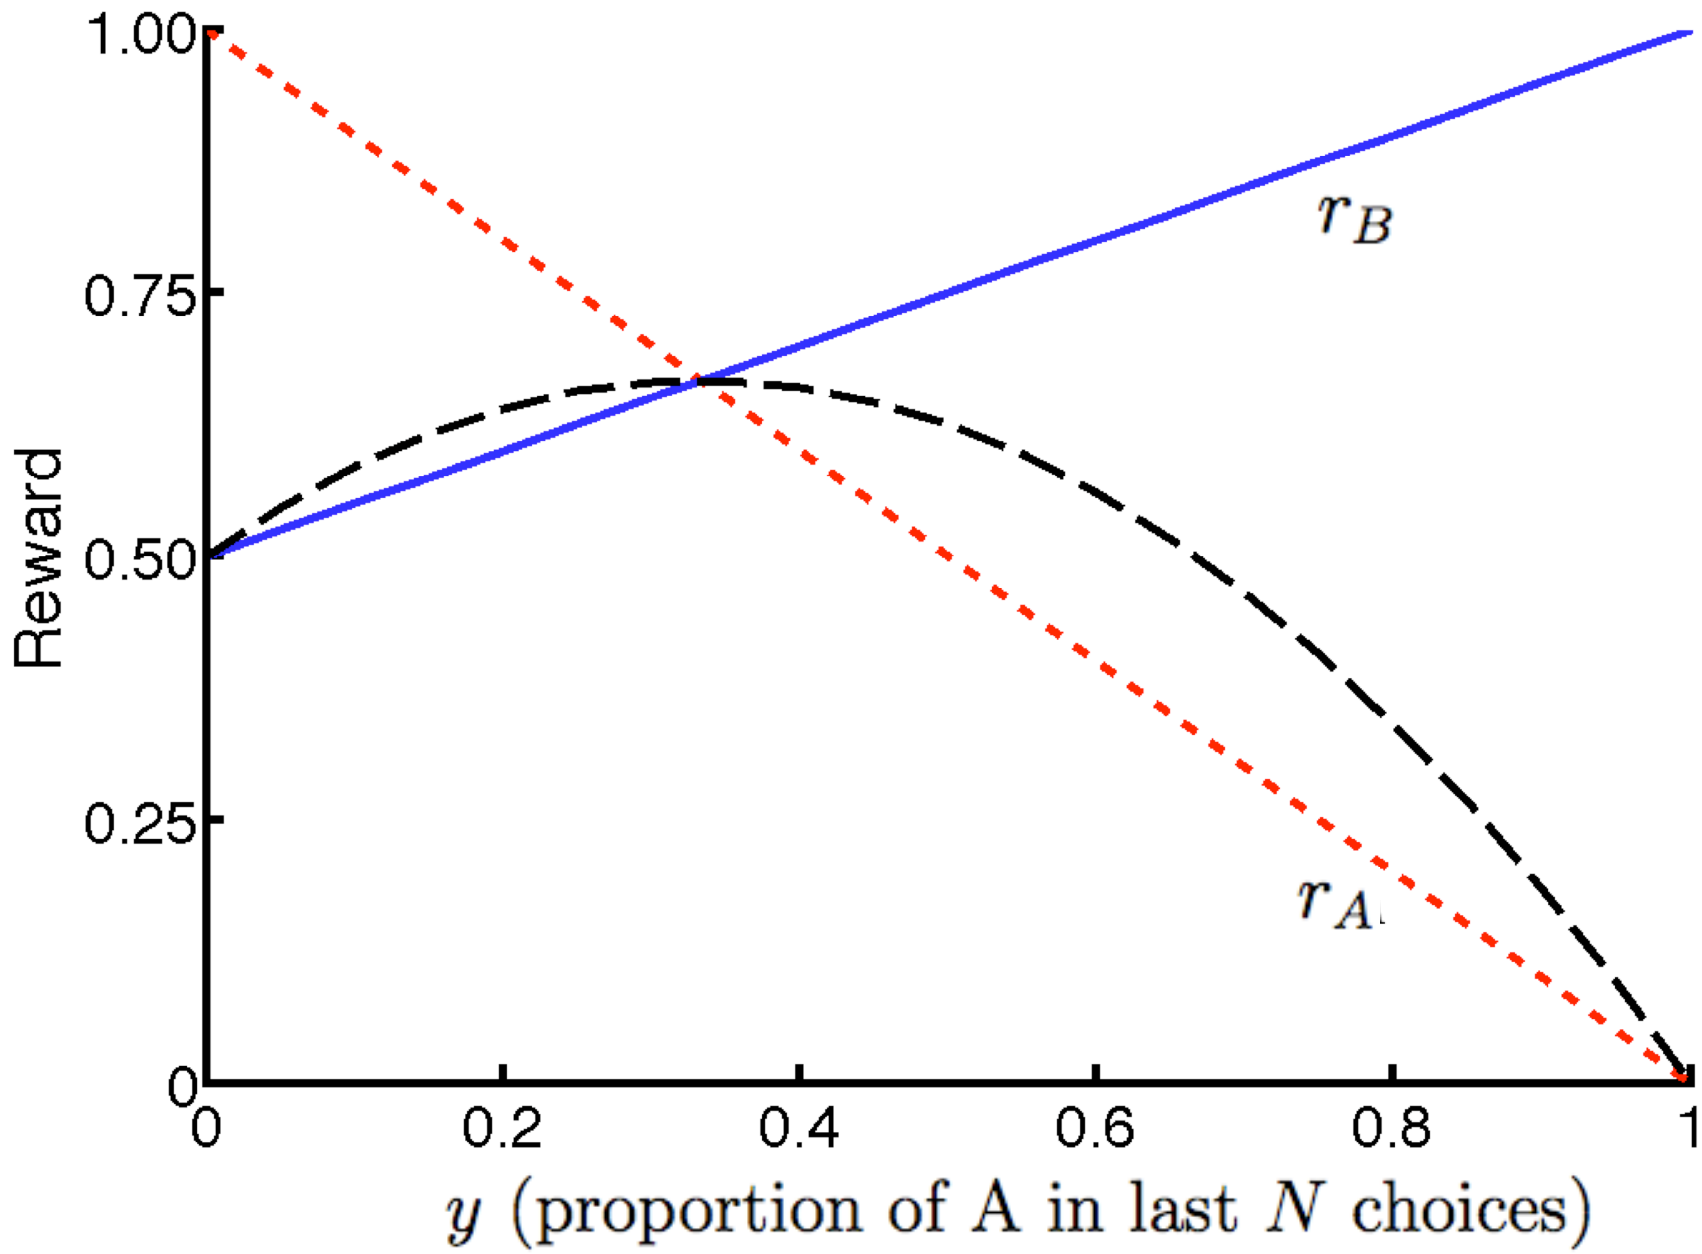

Decision History:

A B B A B B B B A B B A B A B A A A A B

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

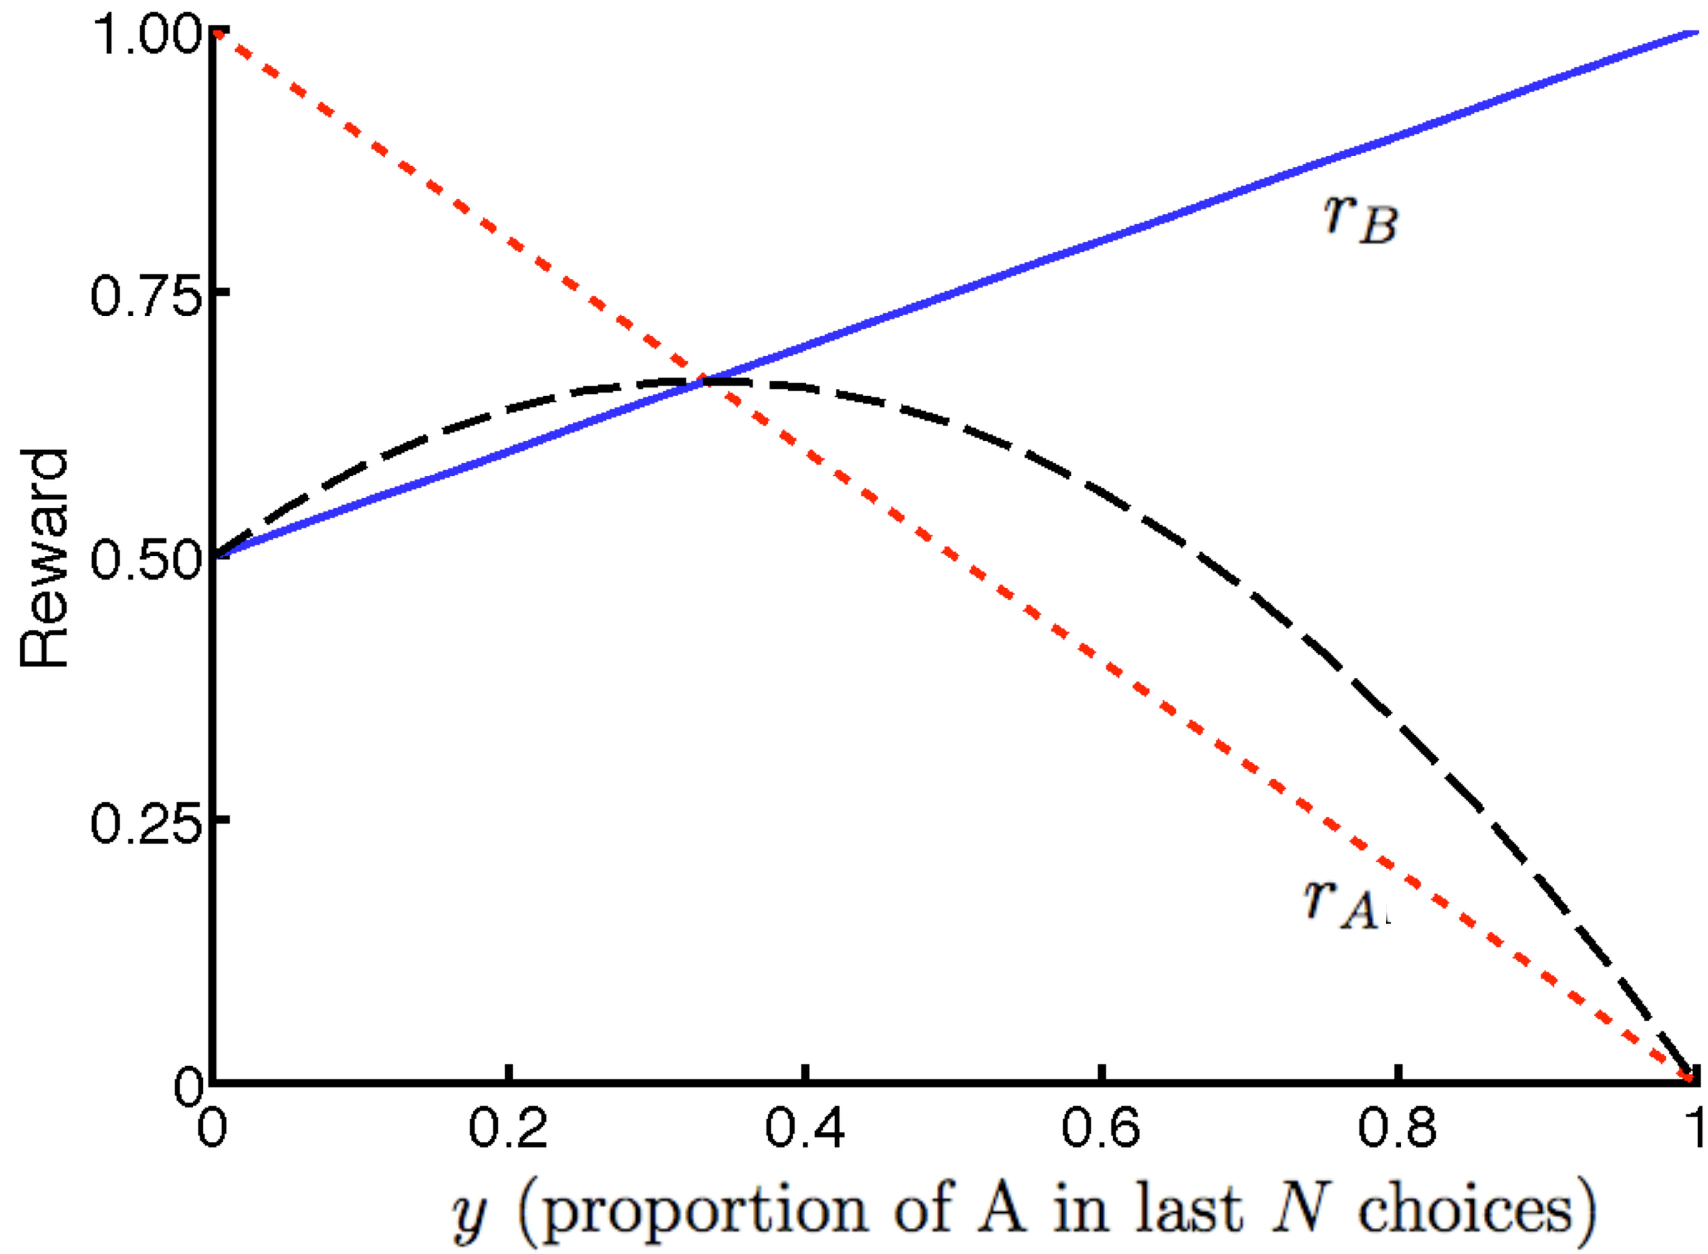

Decision History:

A B B A B B B B A B B A B A B A A A A B

↑ $x_1(t)$ ... ↑ $x_{20}(t)$

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

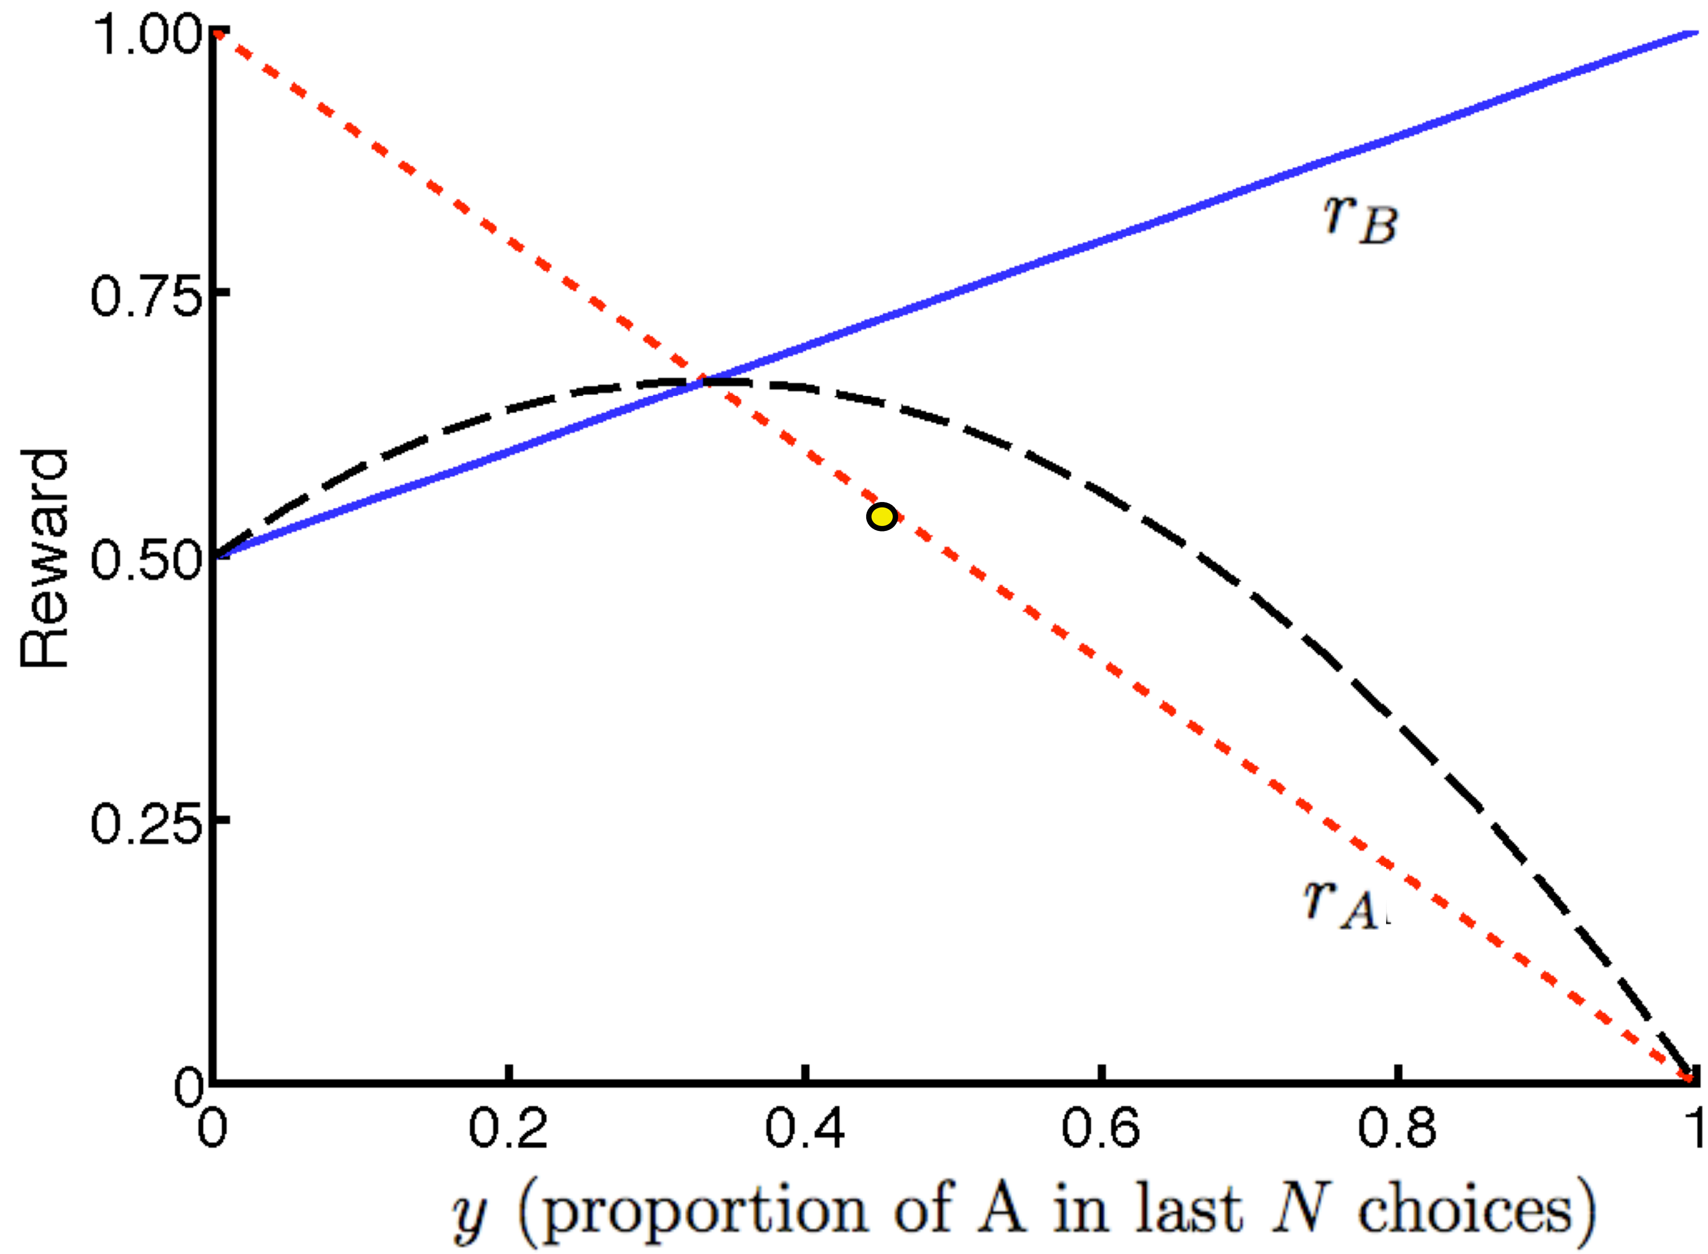

Decision History:

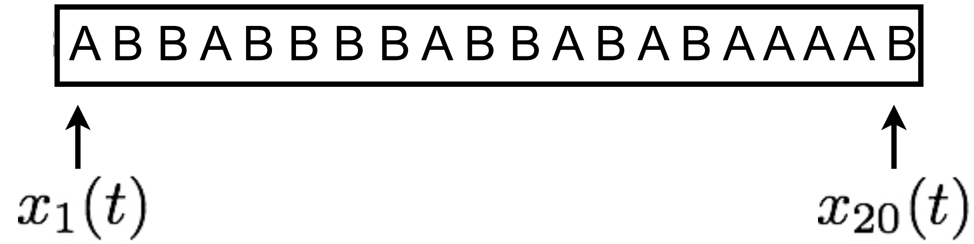

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

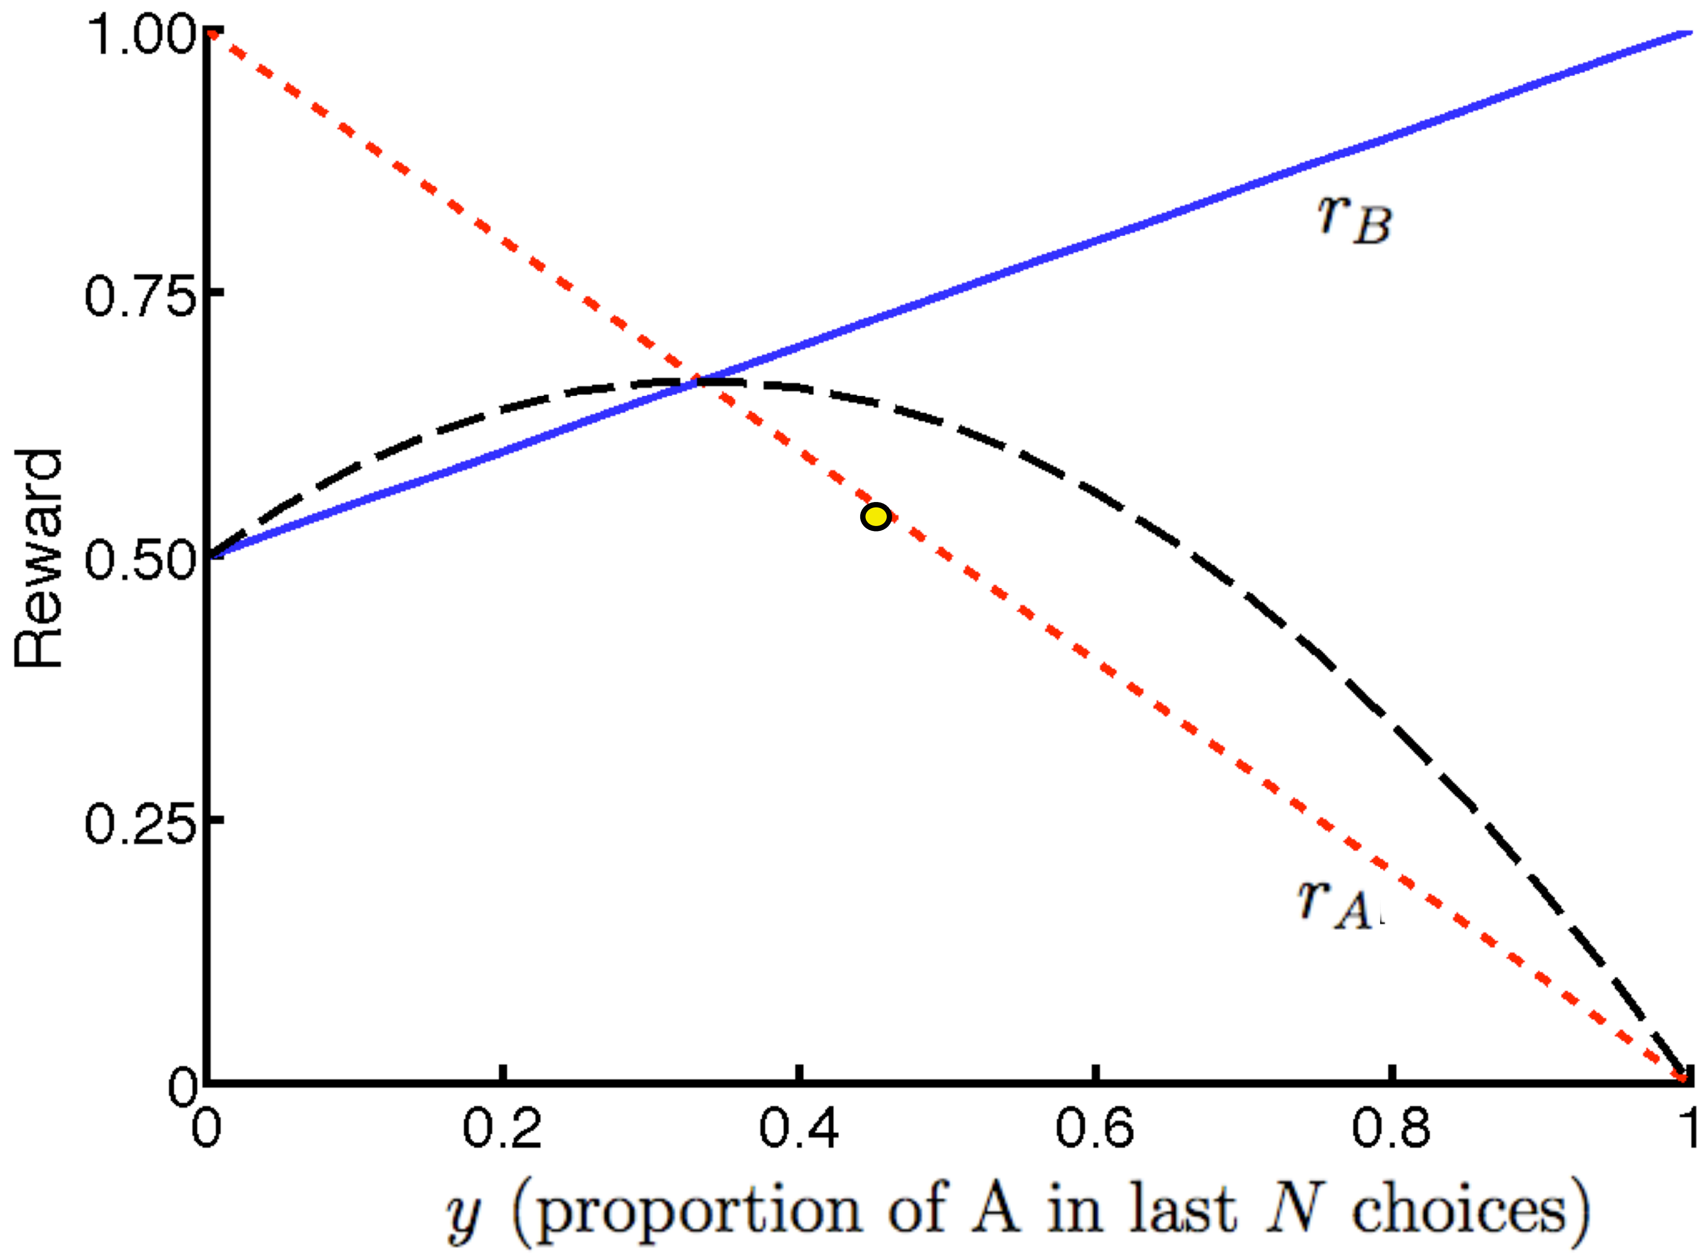

Decision History:

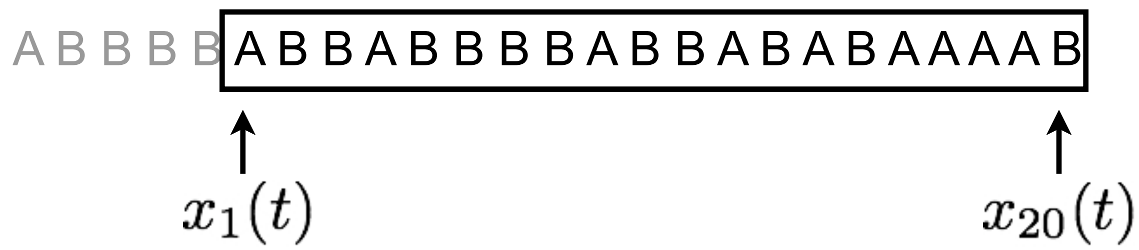

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

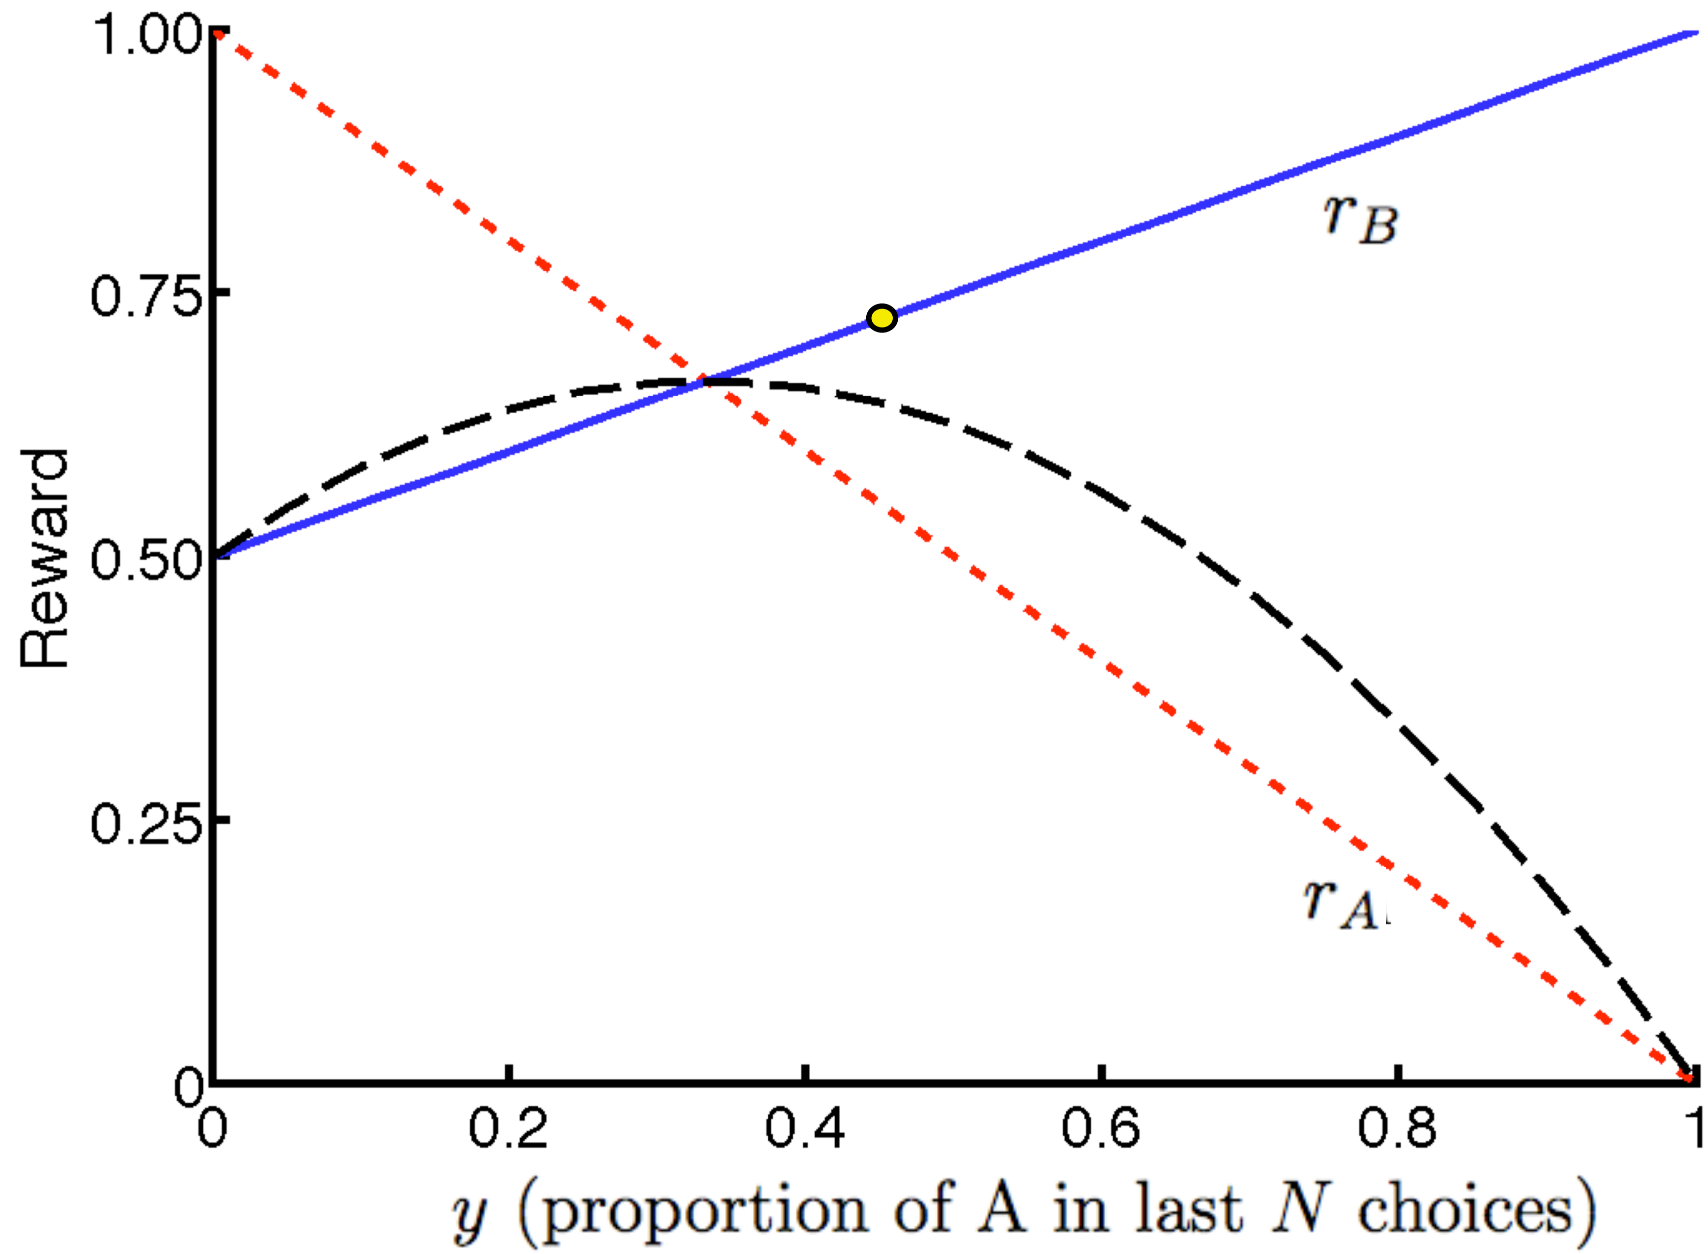

Decision History:

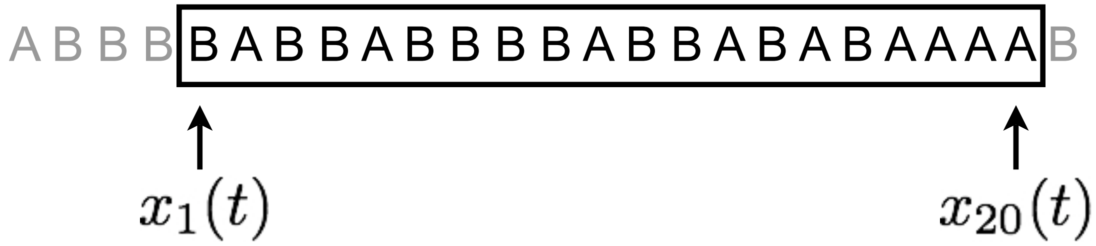

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

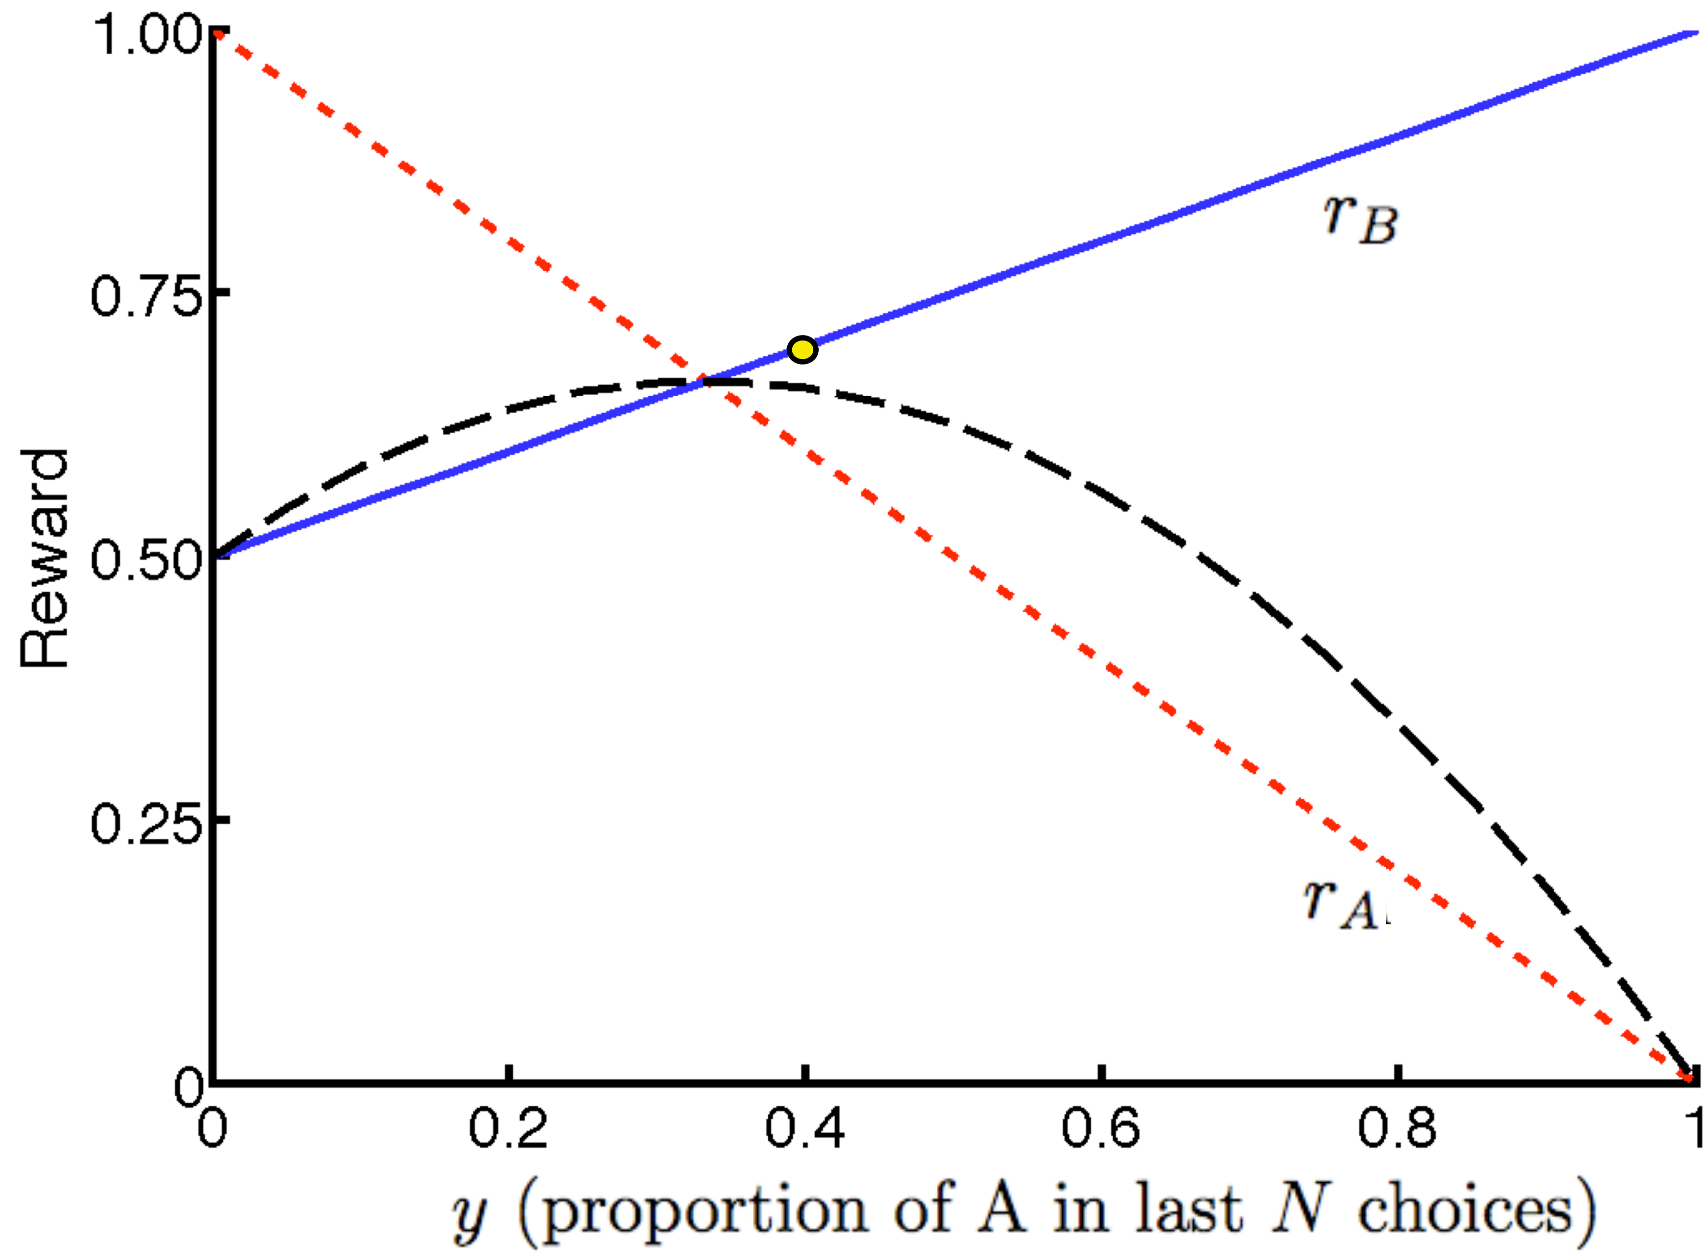

Decision History:

A B B | B B A B B A B B B B A B B A B A B A A A | A B

$x_1(t)$ ↑ ... ↑ $x_{20}(t)$

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

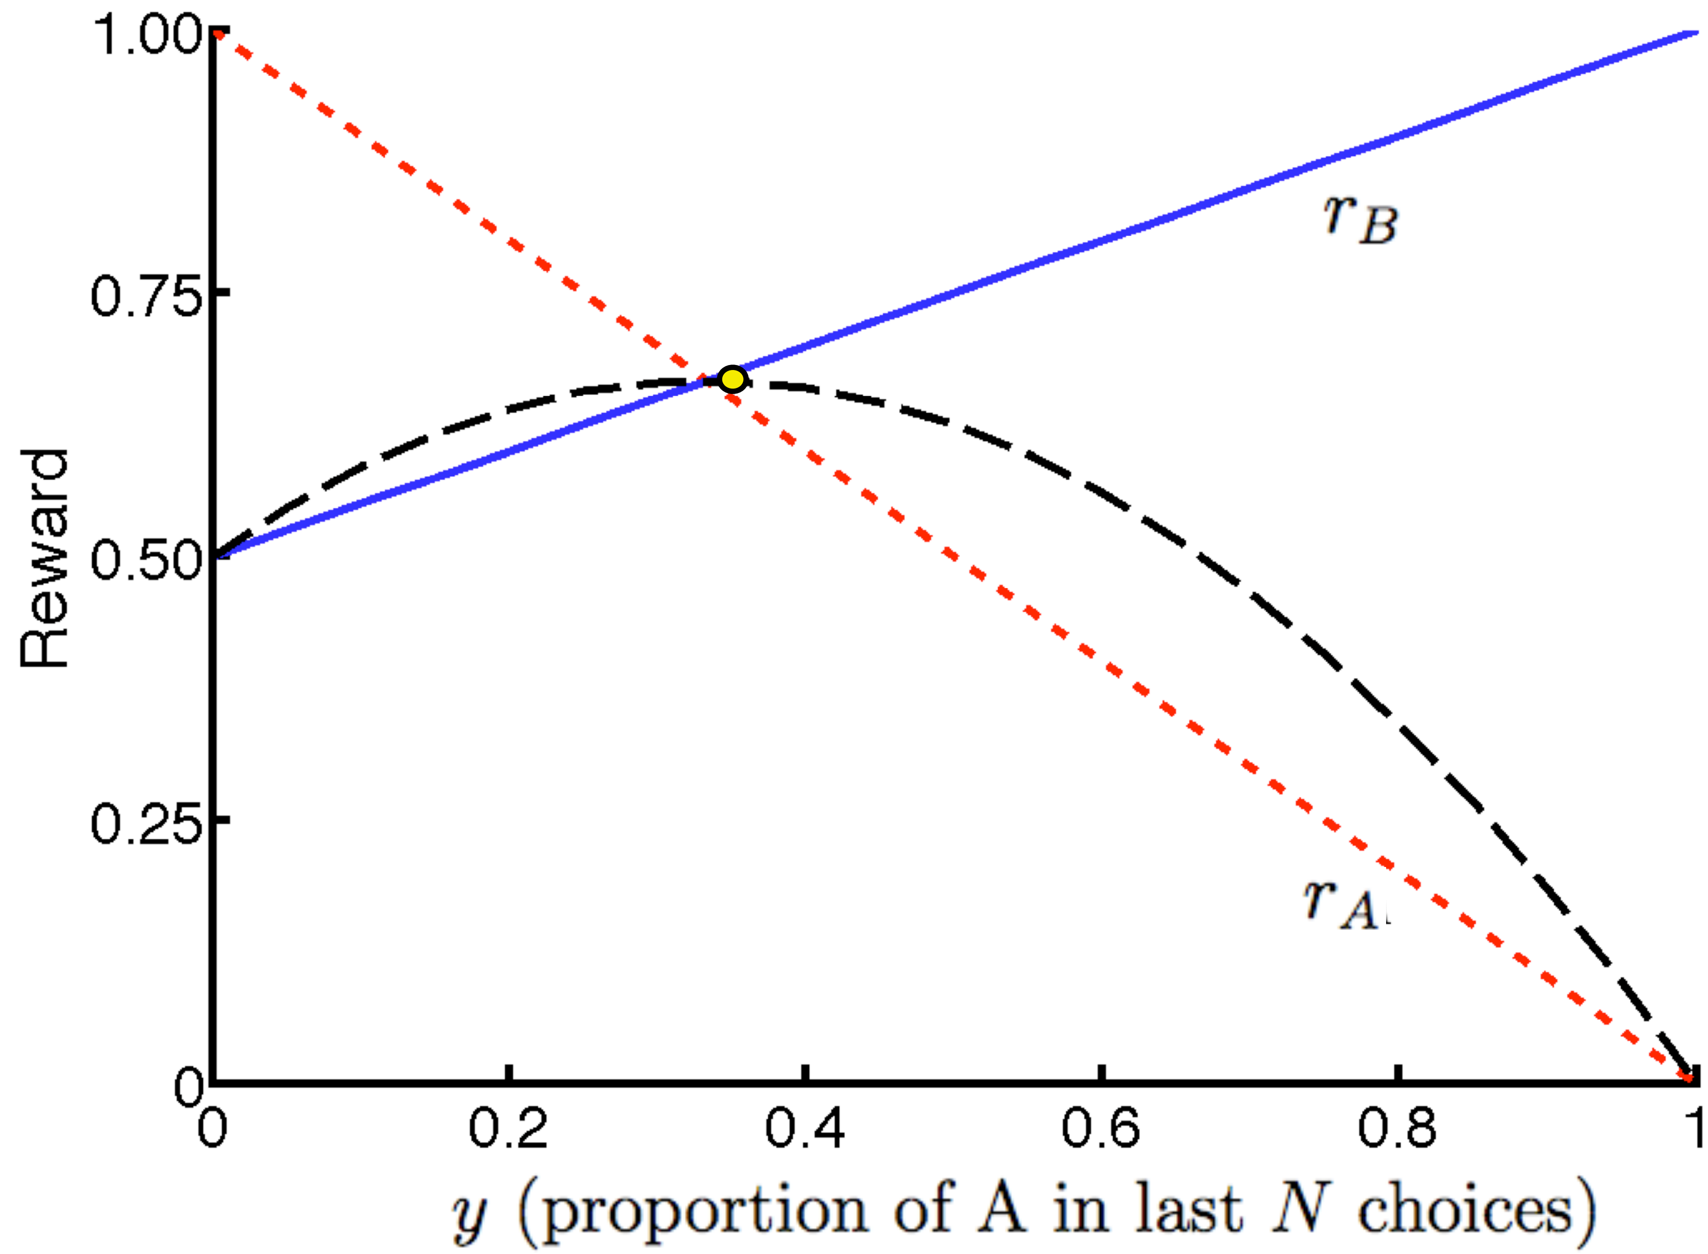

Decision History:

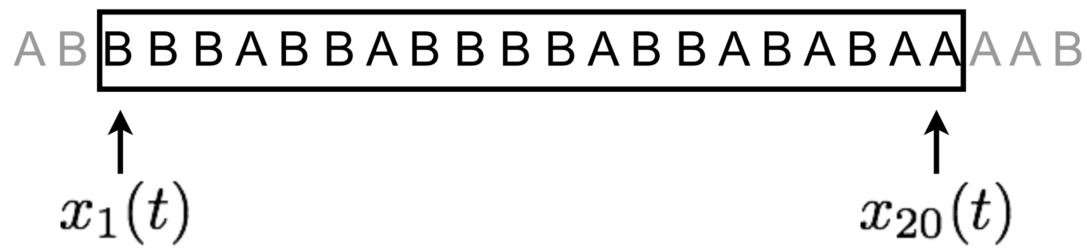

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

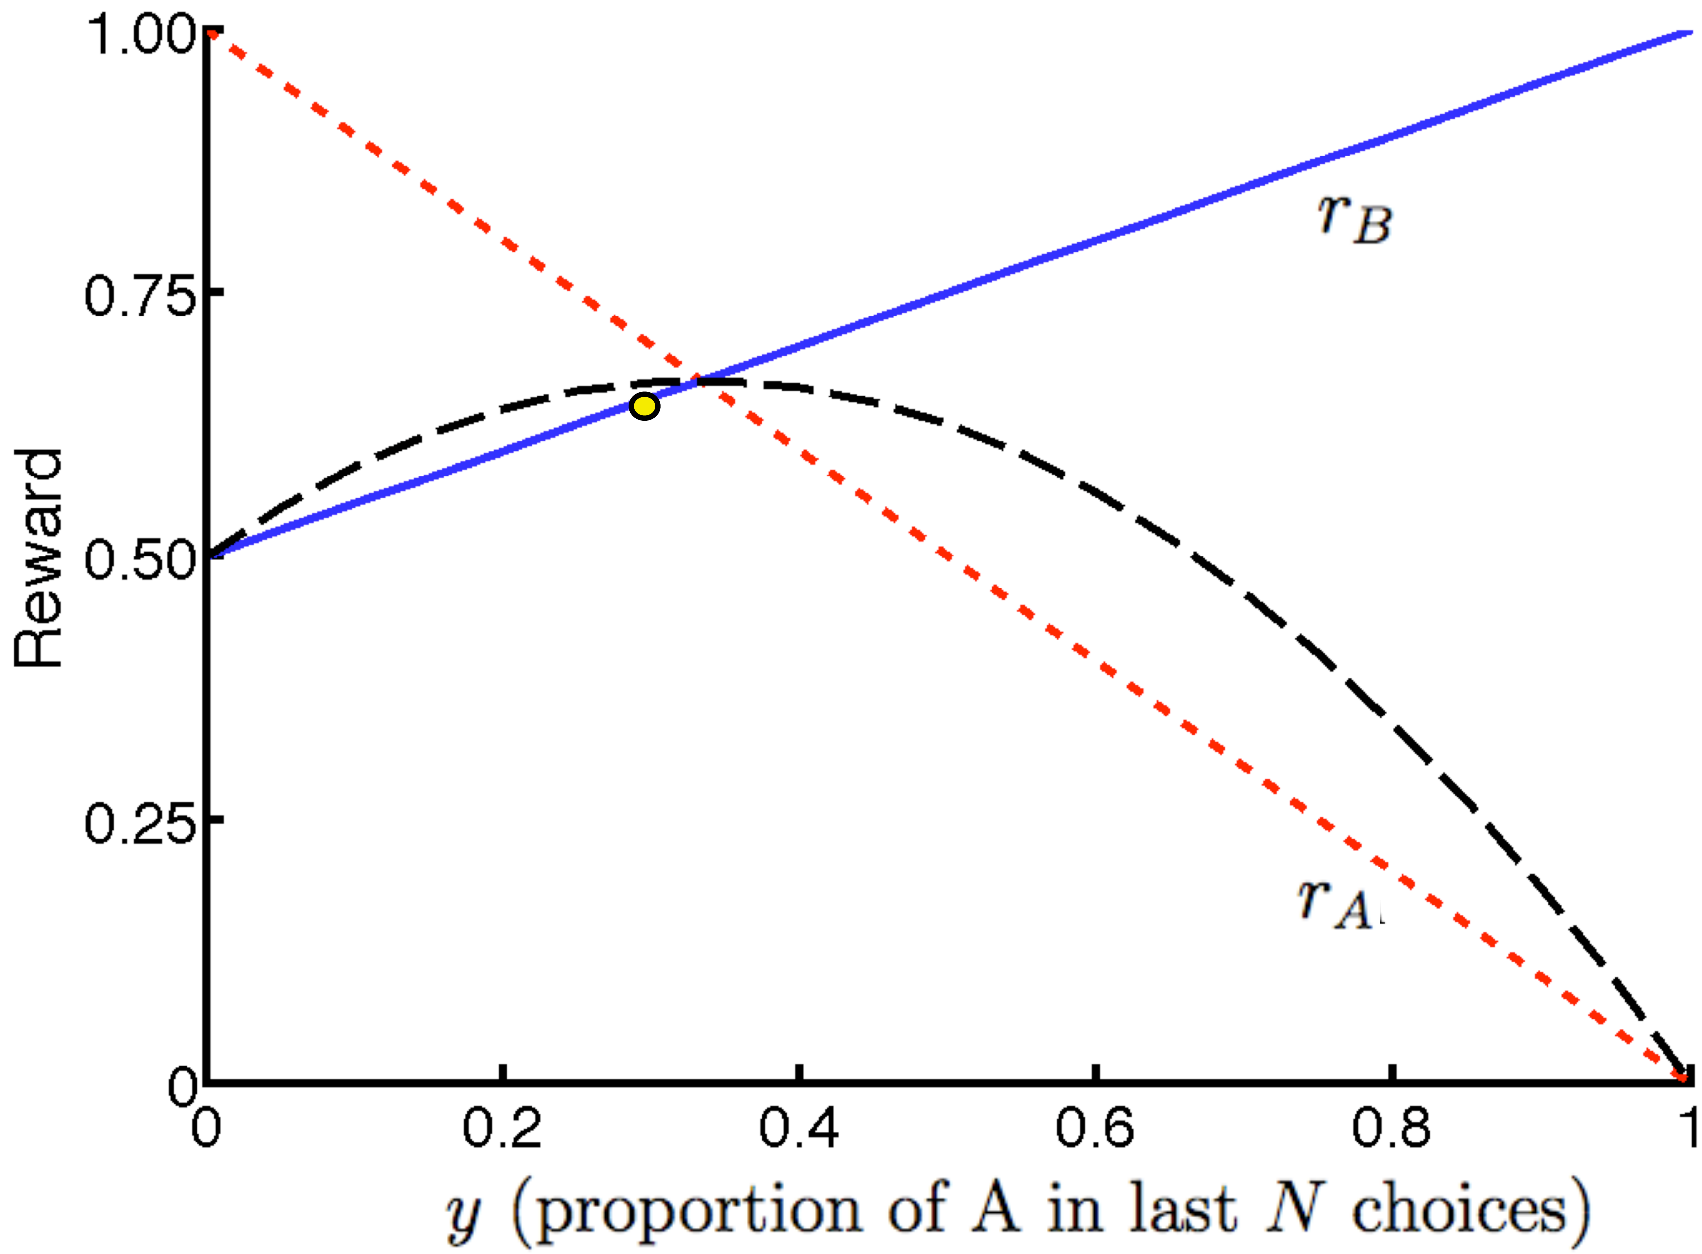

Decision History:

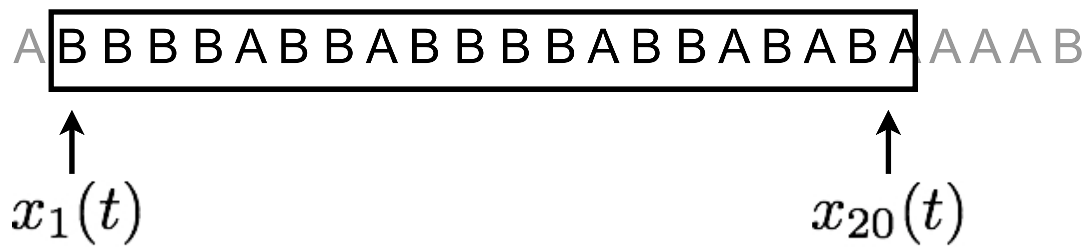

Example Reward Structure: Matching Shoulders

Egelman, Person & Montague 1998
Montague & Berns 2002
Herrnstein 1990

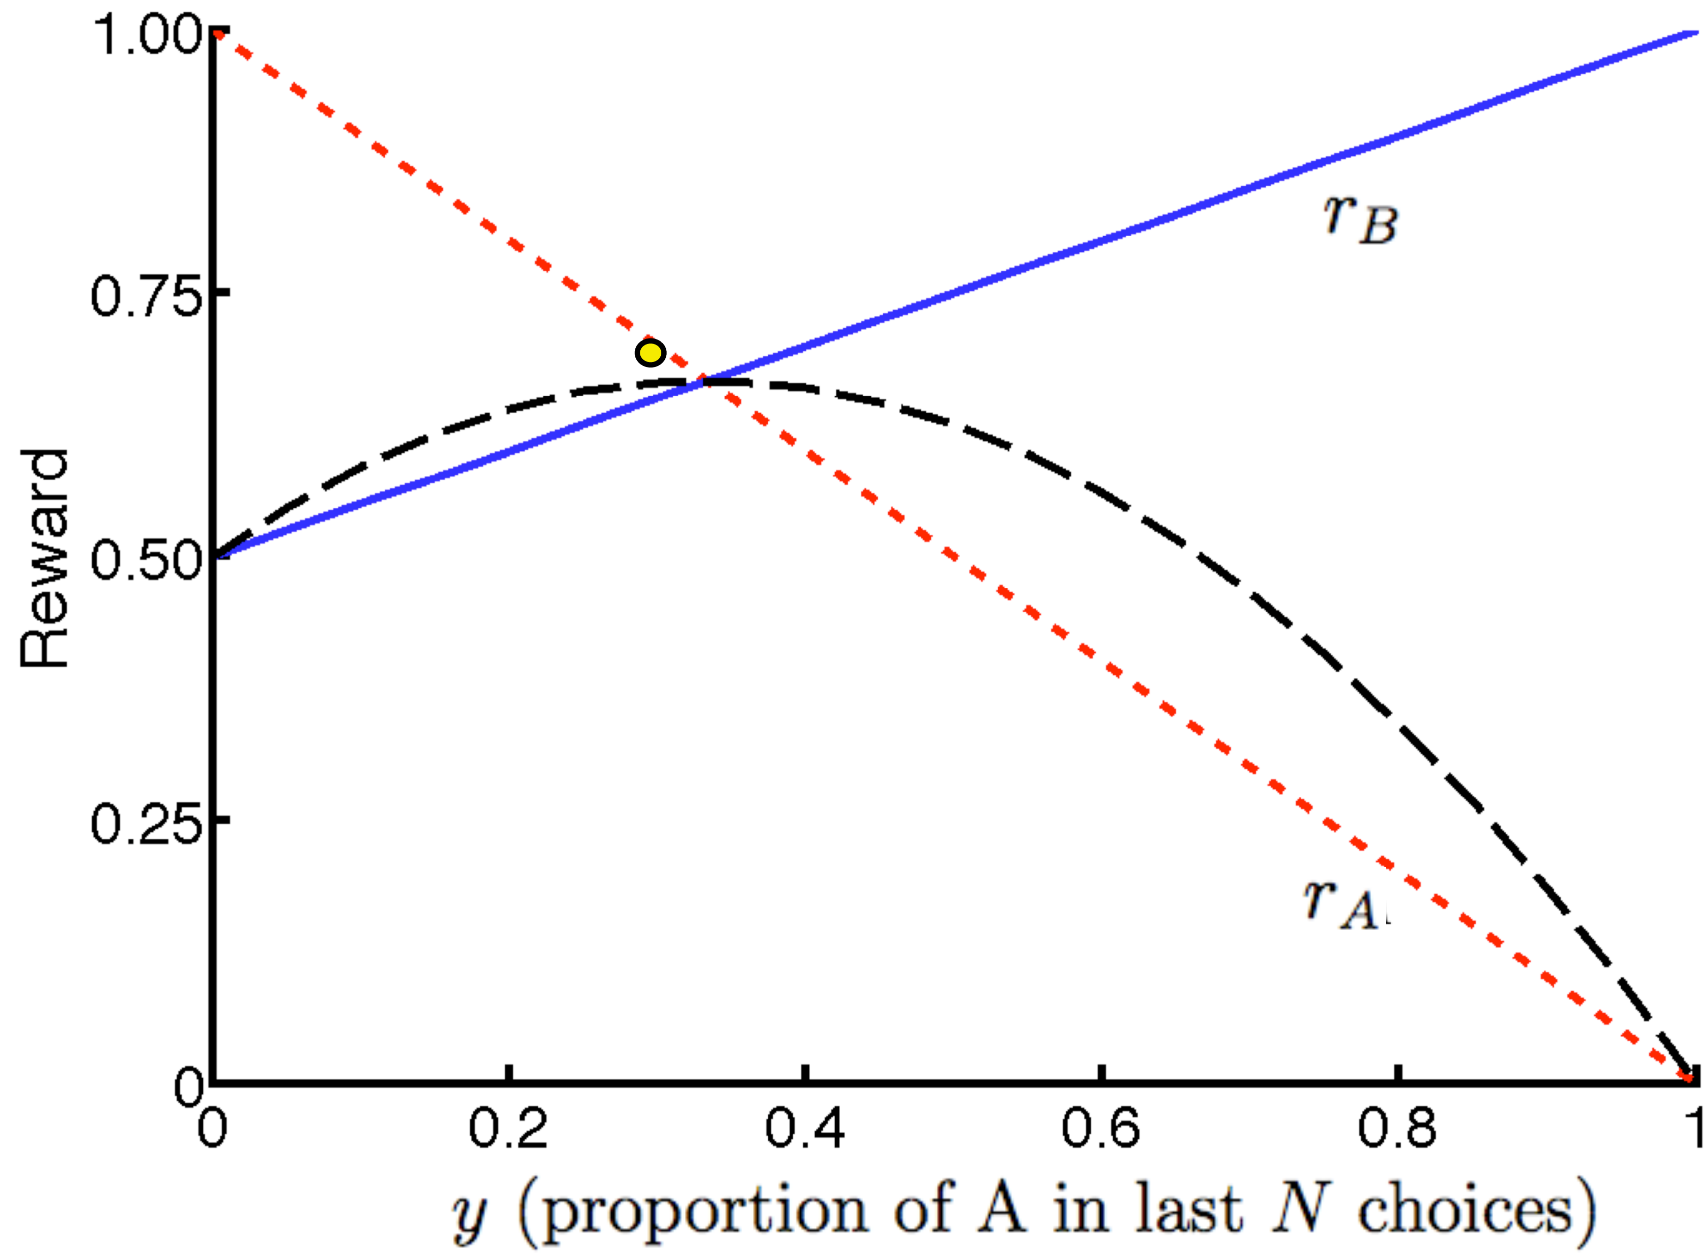

Decision History:

A B B B B A B B A B B B B A B B A B A B A A A A B

$x_1(t)$ $x_{20}(t)$

Egelman, Person and Montague and DDM Model

$$P(A) = \frac{1}{1 + e^{-\mu(w_A - w_B)}}$$

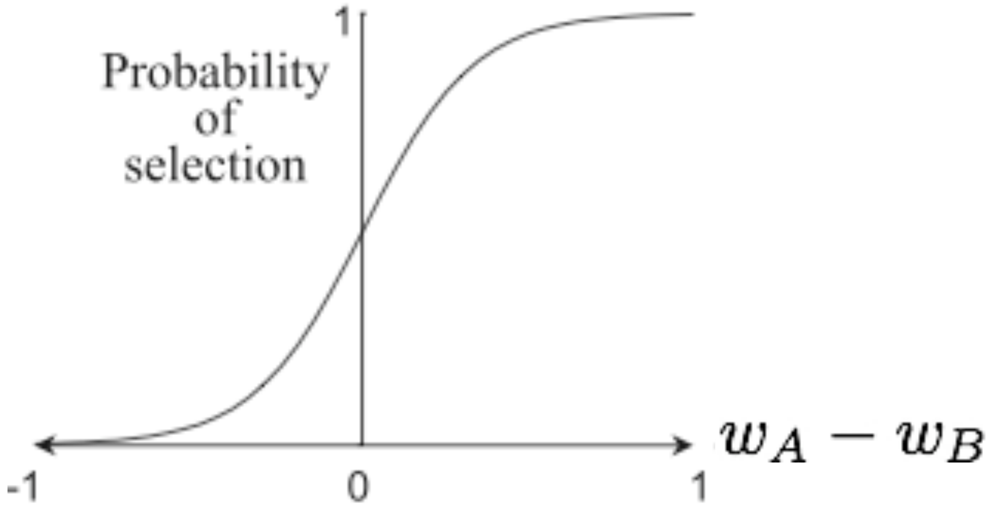

ω_A and ω_B are human subject's anticipated rewards for choices A and B.

Motivated by role of dopamine neurons in coding for reward prediction error (Montague, Dayan, Sejnowski, 1996) and temporal difference learning theory (Sutton and Barto, 1998):

$Z(t) \in \{A, B\}$ is choice at time t
λ is learning rate

$$\begin{aligned} w_Z(t+1) &= (1-\lambda) w_Z(t) + \lambda r(t) \\ w_{\bar{Z}}(t+1) &= w_{\bar{Z}}(t), \qquad t = 0, 1, 2, \ldots \end{aligned}$$

Markov Model of Decision Making

Assumption 1: $\Pr\{x_k(t) = A | x(t)\} = y(t)$

Assumption 2: $w_B(t) - w_A(t) = r_B(t-1) - r_A(t-1)$

Montague and Berns, 2002
Nedic and Holmes

Proposition: Suppose Assumptions 1 & 2 hold, then the DDM for the TAFC is a Markov Process with state $y(t)$ and transition probabilities can be computed explicitly:

$$\mathbf{P}_{ij} = \Pr\{y(t+1) = \frac{j}{N} | y(t) = \frac{i}{N}\}, i, j \in \{0, 1, \ldots, N\}$$

Markov Model of Decision Making

Assumption 1: $\Pr\{x_k(t) = A | x(t)\} = y(t)$

Assumption 2: $w_B(t) - w_A(t) = r_B(t-1) - r_A(t-1)$

Montague and Berns, 2002
Nedic and Holmes

Proposition: Suppose Assumptions 1 & 2 hold, then the DDM for the TAFC is a Markov Process with state $y(t)$ and transition probabilities can be computed explicitly:

$$\mathbf{P}_{ij} = \Pr\{y(t+1) = \frac{j}{N} | y(t) = \frac{i}{N}\}, i, j \in \{0, 1, \ldots, N\}$$

as can the unique steady-state distribution for $y(t)$ satisfying: $\pi \mathbf{P} = \pi$ and $\sum_{i=0}^{N} \pi_i = 1$

$$\pi_i = \frac{\alpha_i (1 + e^{\mu \Delta r(\frac{i}{N})}) e^{-\mu \beta_i}}{\sum_{j=0}^{N} \alpha_j e^{-\mu \beta_j} (1 + e^{\mu \Delta r(\frac{j}{N})})}$$

$$\Delta r(y) = r_B(y) - r_A(y)$$

$$y \in \{\frac{i}{N},\ i = 0, \ldots, N\}$$

$$\alpha_i = \frac{N!}{(N-i)!i!} \quad \text{and} \quad \beta_i = \sum_{j=1}^{i} \Delta r(\frac{j}{N})$$

Social Feedback

Two players (one alone, one receives choice feedback)

Choice feedback model: $$P_A = \frac{1}{1 + e^{\mu(w_B - w_A - \nu u)}} \qquad u = \begin{cases} 1 & \text{if} \quad AA \\ -1 & \text{if} \quad BB \\ 0 & \text{otherwise} \end{cases}$$

(Feedback reinforces decision when same as own)

Social Feedback

Two players (one alone, one receives choice feedback)

Choice feedback model: $$P_A = \frac{1}{1+e^{\mu(w_B - w_A - \nu u)}} \qquad u = \begin{cases} 1 & \text{if} \quad AA \\ -1 & \text{if} \quad BB \\ 0 & \text{otherwise} \end{cases}$$

(Feedback reinforces decision when same as own)

Using Assumptions 1 and 2 and π to model the player in alone condition the probability of choosing an A becomes

$$\begin{aligned} \bar{P}_A = \quad & (1 - \frac{1}{1+e^{\mu(w_B - w_A - \nu)}}) y \sum_{i=0}^{N} \frac{1}{1+e^{\mu(\Delta r(\frac{i}{N}))}} \pi_i \\ + \quad & (1 - \frac{1}{1+e^{\mu(w_B - w_A + \nu)}}) (1-y)(1 - \sum_{i=0}^{N} \frac{1}{1+e^{\mu(\Delta r(\frac{i}{N}))}} \pi_i) \\ + \quad & (1 - \frac{1}{1+e^{\mu(w_B - w_A}}) \left[y(1 - \sum_{i=0}^{N} \frac{1}{1+e^{\mu(\Delta r(\frac{i}{N}))}} \pi_i) + (1-y) \sum_{i=0}^{N} \frac{1}{1+e^{\mu(\Delta r(\frac{i}{N}))}} \pi_i \right] \end{aligned}$$

One step transition matrix $\mathbf{P}$ can be computed and from it the steady state distribution.

Ex 1: *Matching Shoulders*

$\mu = 5$ (both decision-makers), $\nu \in [0, 3]$

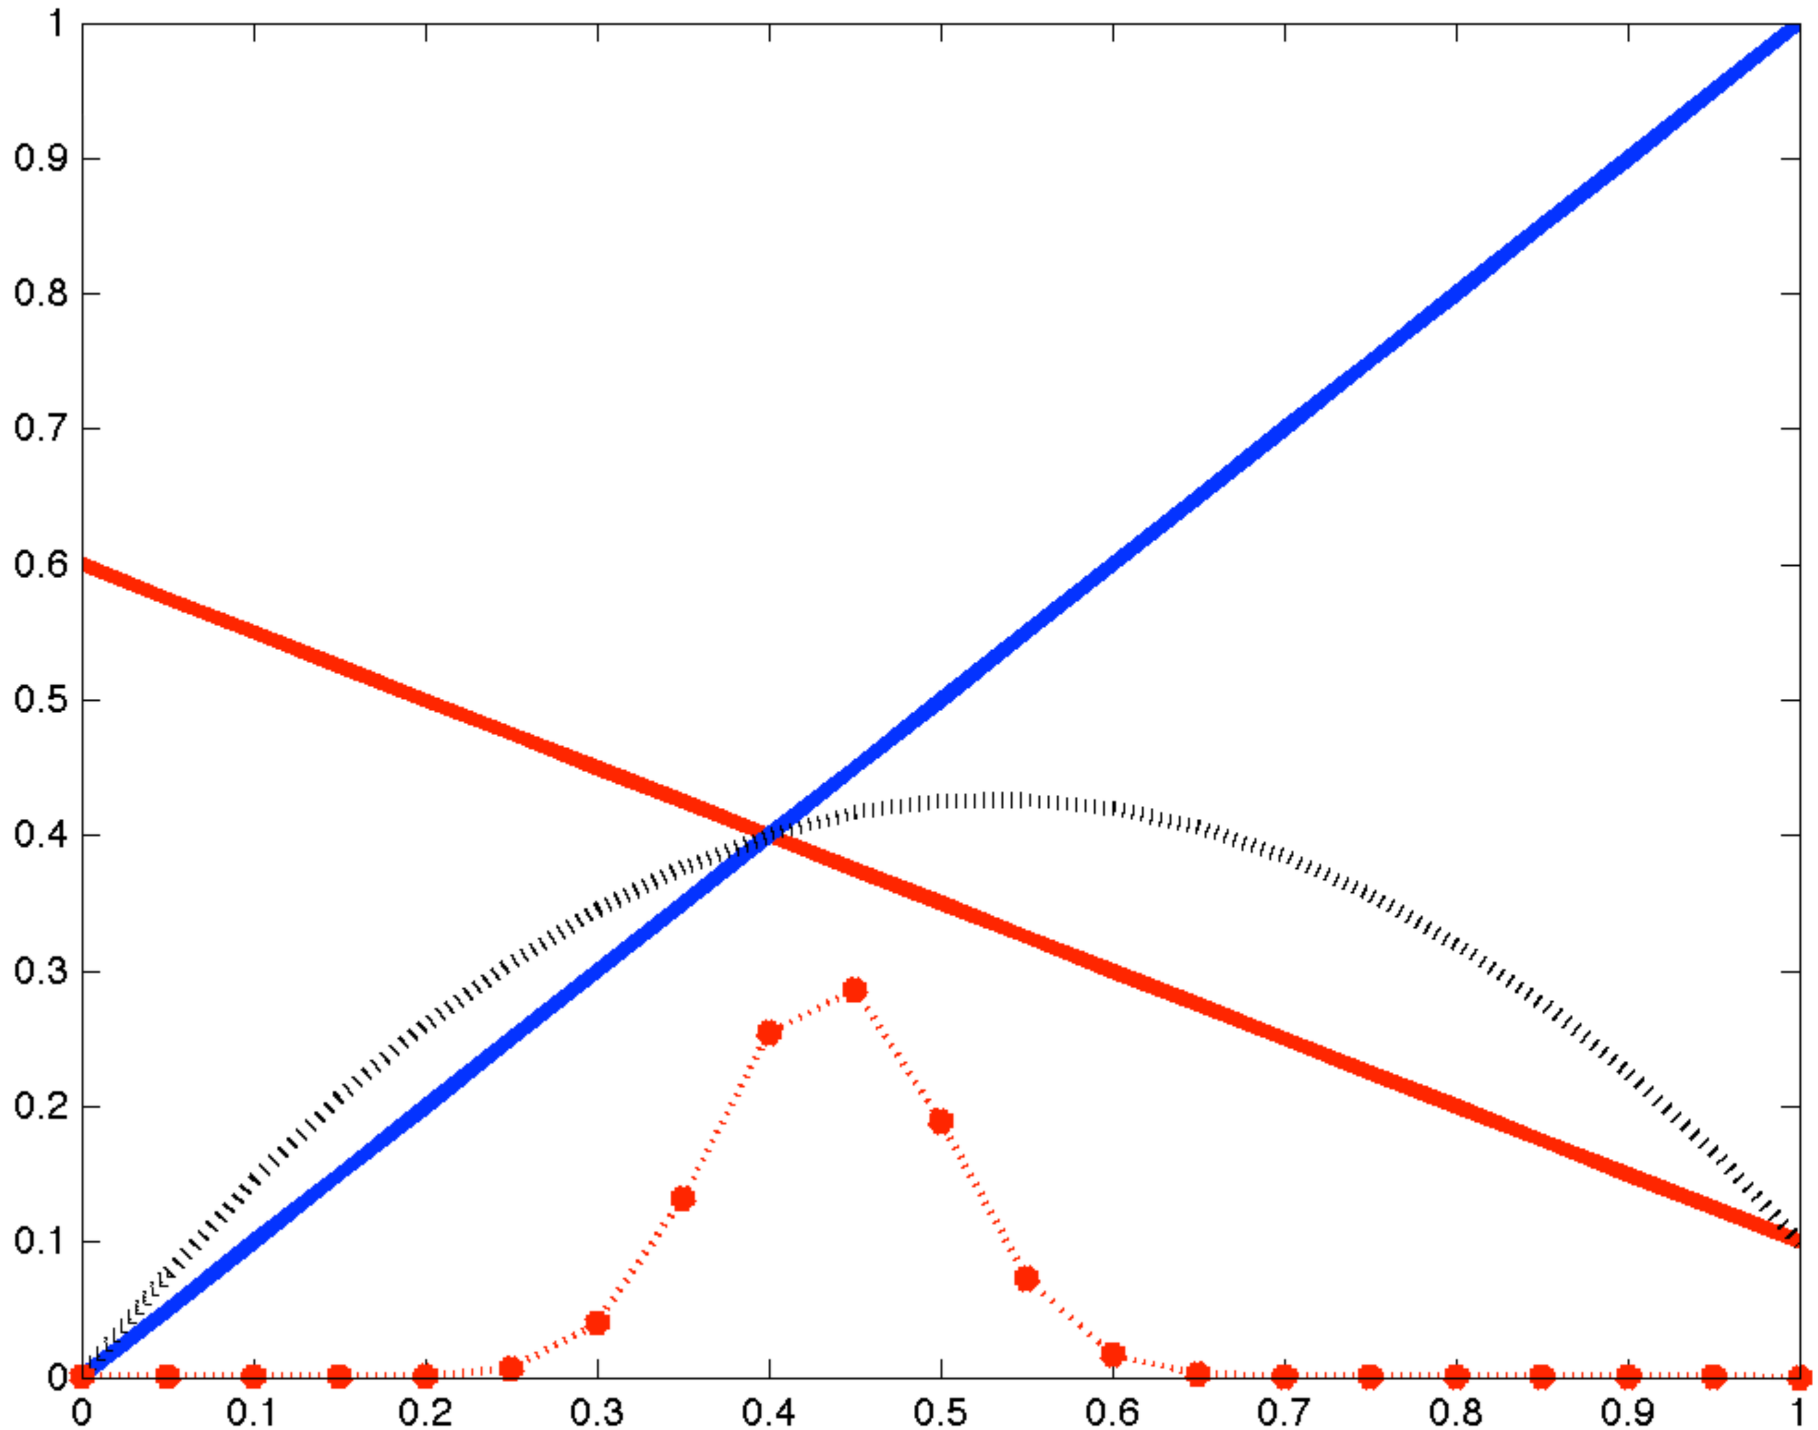

Red: individual alone; Green: individual with choice feedback

Ex 2: *Converging Gaussians*

$\mu = 5$ (both decision-makers), $\nu \in [0, 3]$

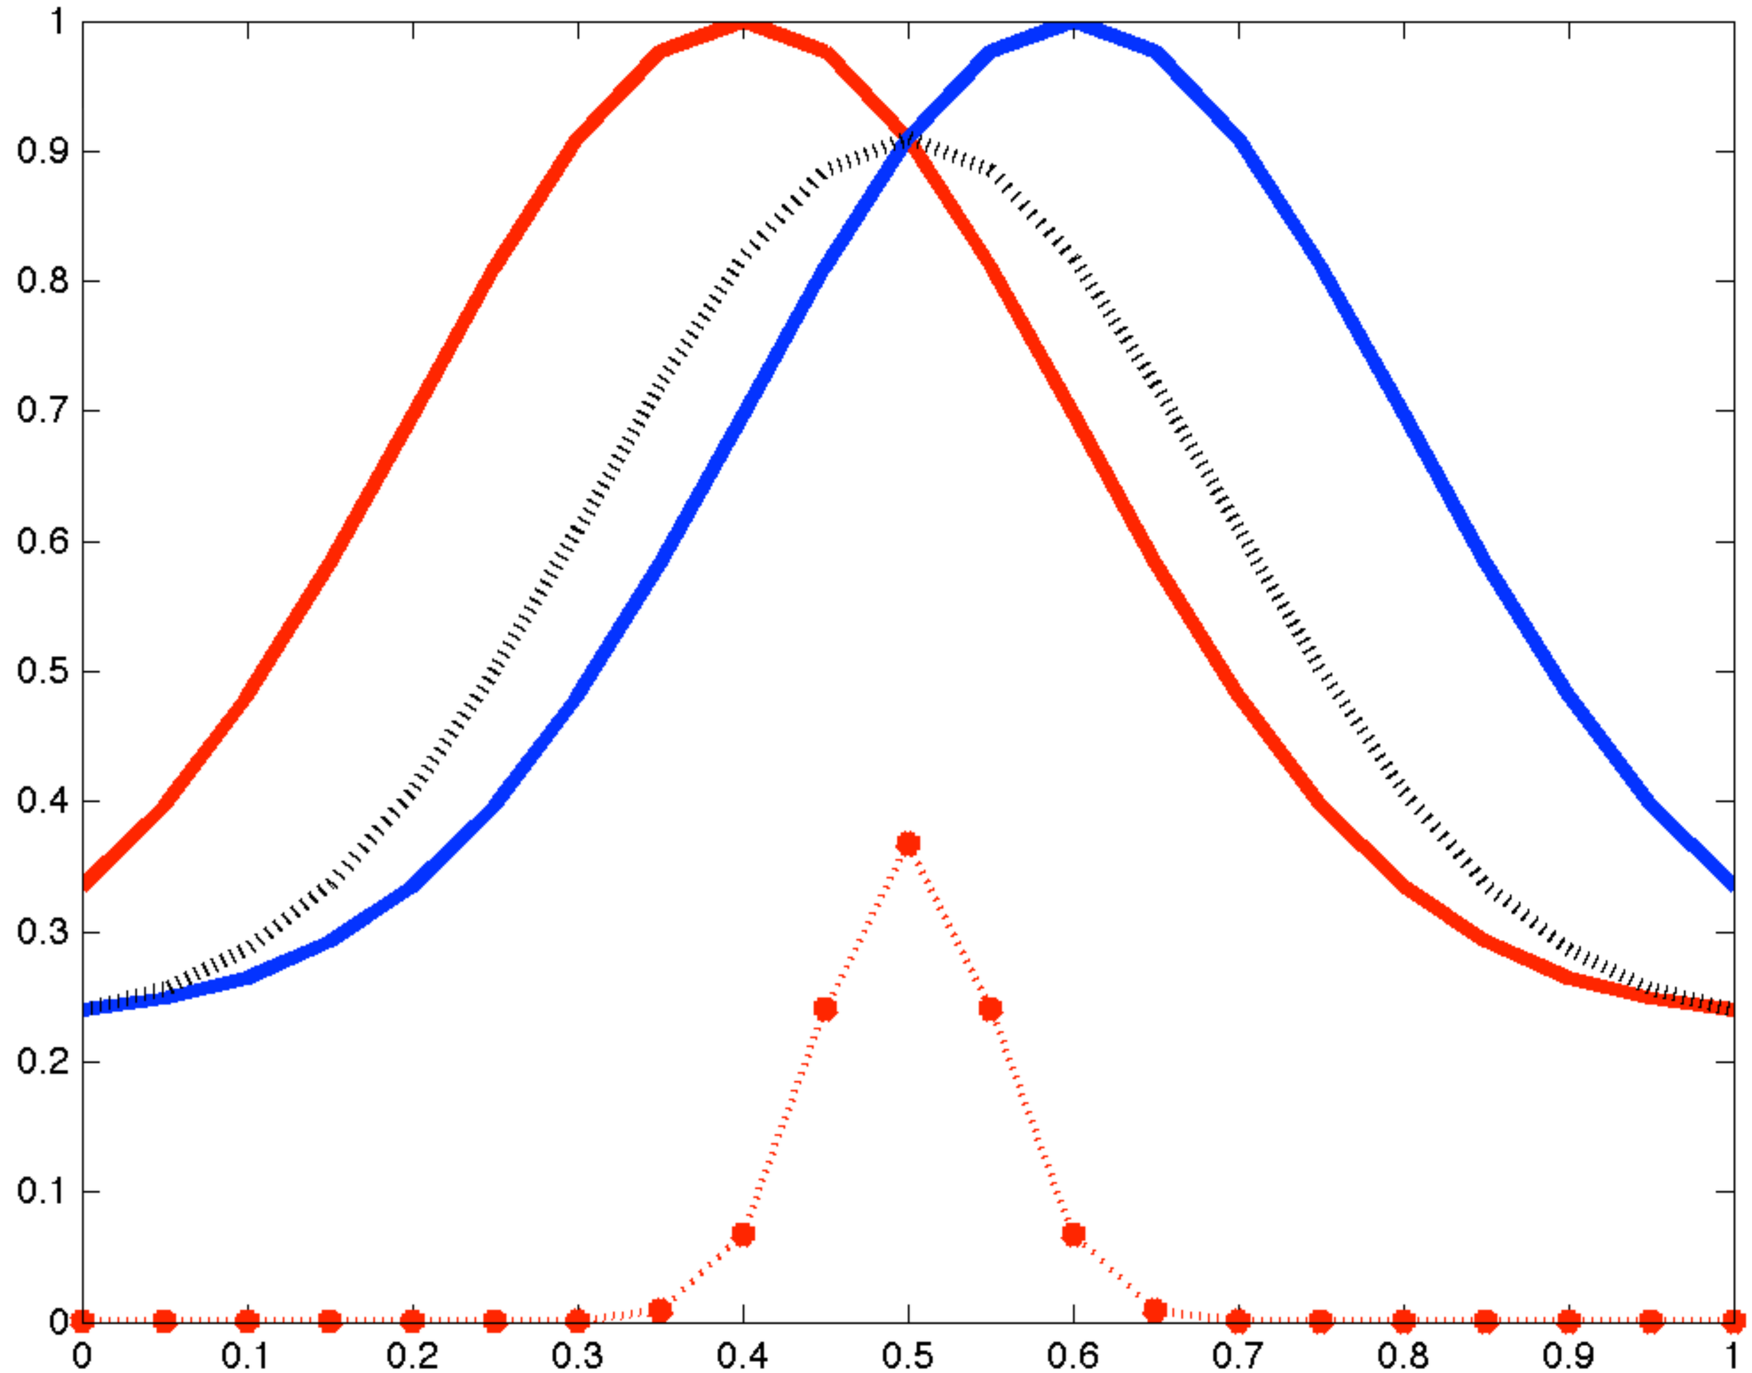

Red: individual alone; Green: individual with choice feedback

Conclusions

Interconnection topology plays important role in social decision making with uncertainty.

Uncertainty is lowest at "middle" of the group.

Results suggest that robustness does not necessarily increase with number of interconnections and some directed graphs perform better than their undirected counterparts.

Preliminary results on starling flocks suggest that with respect to robustness of consensus to white noise, there may be an optimum number of neighbors.

Preliminary studies with multiple human decision makers show that choice feedback increases the spread of steady-state choice distribution in a variety of tasks.

Robot/Fish Experiment

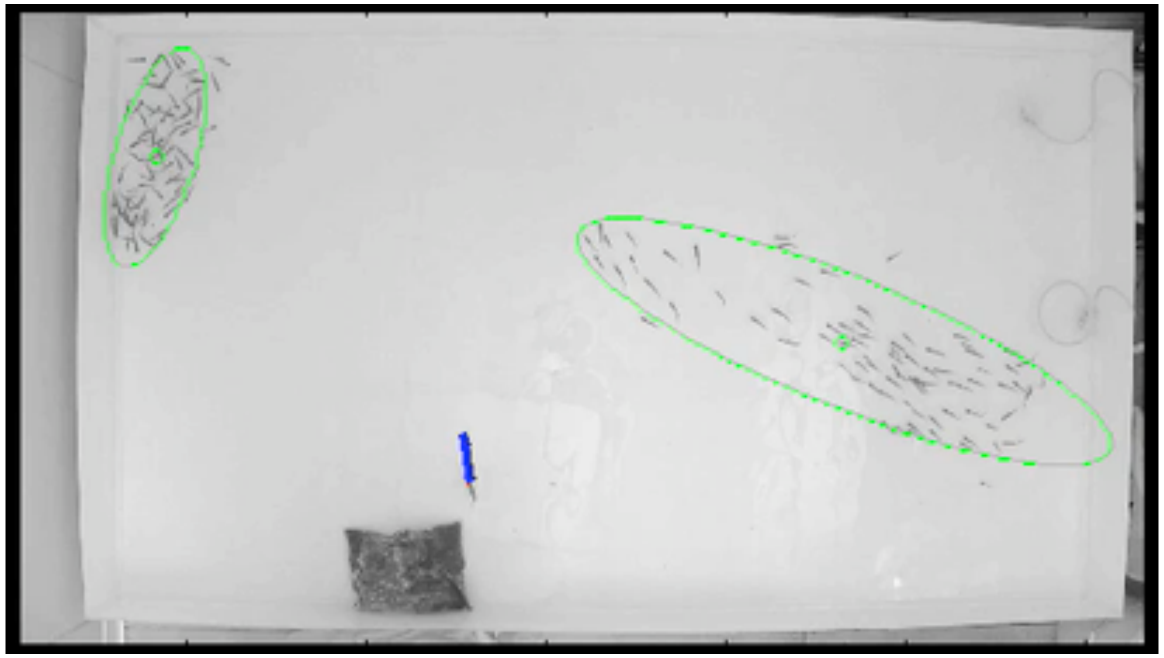

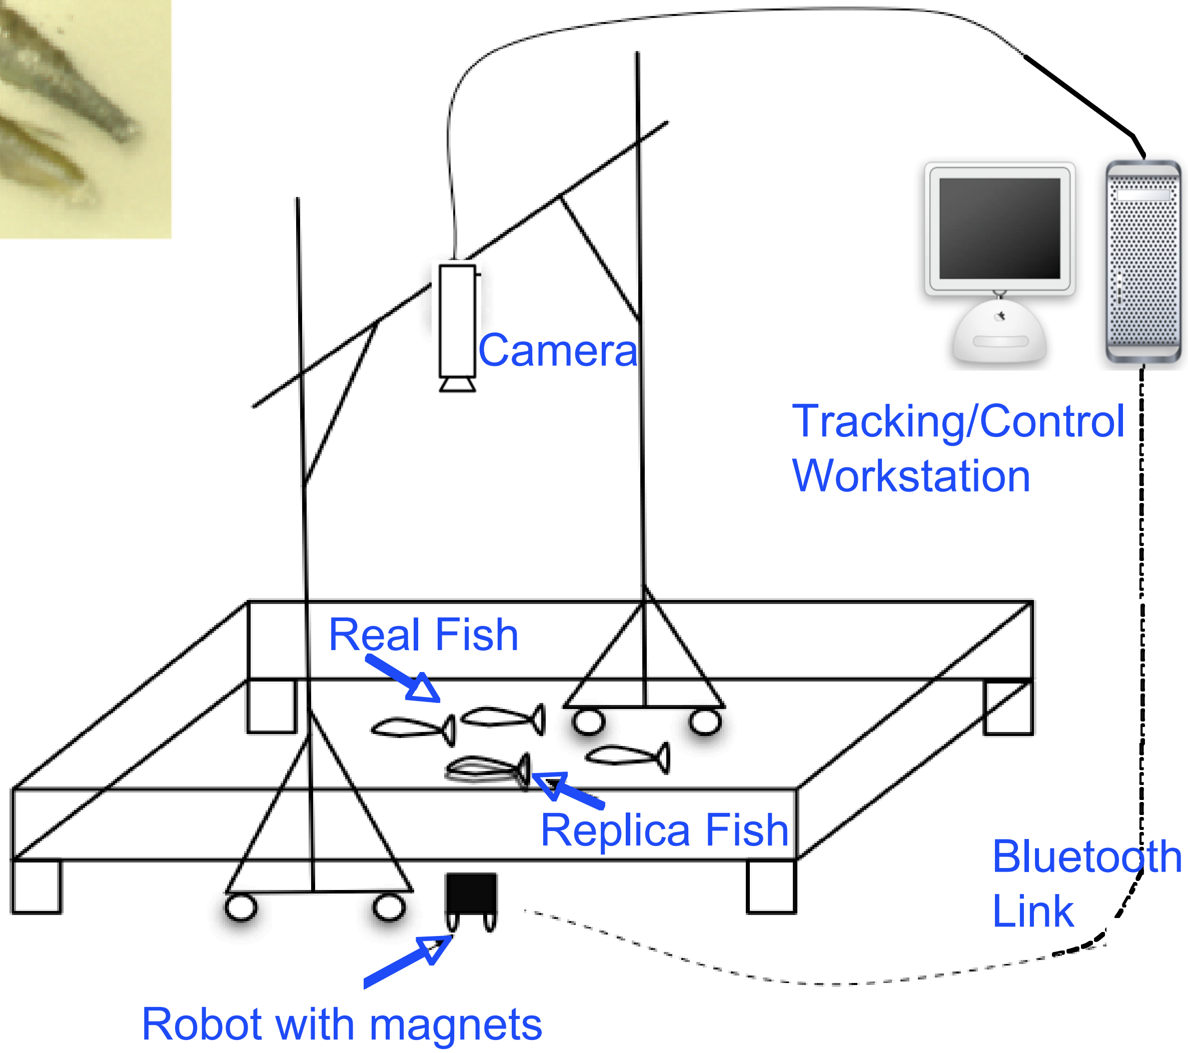

Golden Shiner school with robot predator in Princeton Robot/Fish Test-bed.
Credit: D. Swain, C. Ioannou, Y. Katz, I. Couzin, N.E. Leonard