The Bayesian Brain:
the timing of perceptual decisions

Jan Drugowitsch

jan_drugowitsch@hms.harvard.edu

Department of Neurobiology
Harvard Medical School

CoSMo 2017

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-modal decision-making

Source code

Get code/data from

`https://github.com/DrugowitschLab/CoSMo2017`

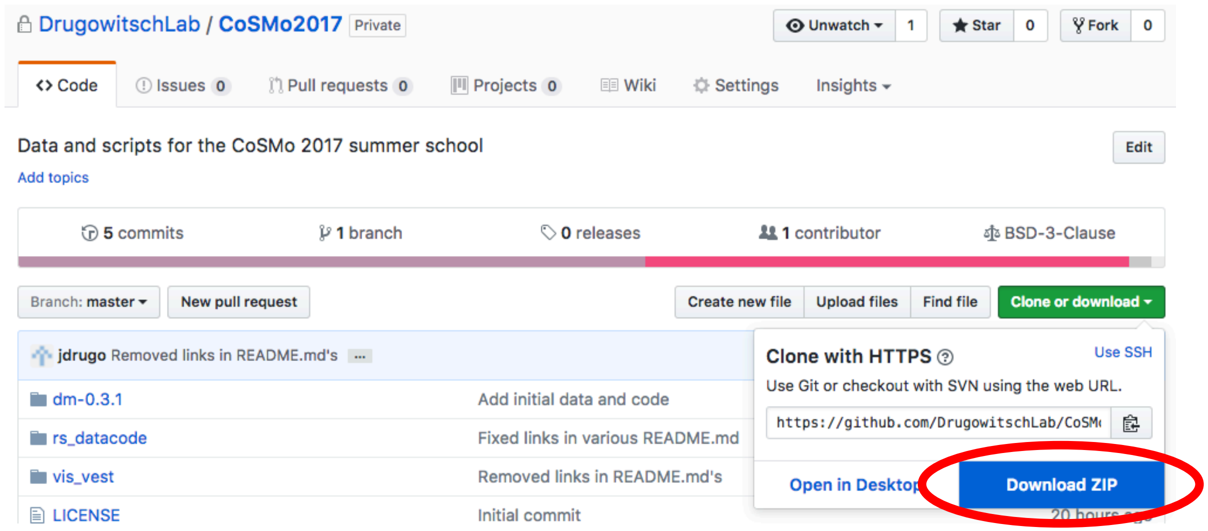

Extract & open folder in Matlab, try `load('phs_ah.mat')`

Add dm library to path

```
>> addpath('dm-0.3.1/matlab/')
>> ddm_fpt_example
```

Road map

→ Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Decisions are endemic

Every action is a decision

Requires: identification of choice options

e.g., *should I stay, or should I go?*

gather knowledge (external/internal) about either option

evaluate choices with respect to expected outcome

e.g., *if I stay there will be trouble*
if I go there will be double

Main focus today: **perceptual decisions**
(decisions based on what we observe)

speed? accuracy? underlying process?

Uncertain information

Information we have about the world is **uncertain**

Uncertainty due to *noise* and *ambiguity*

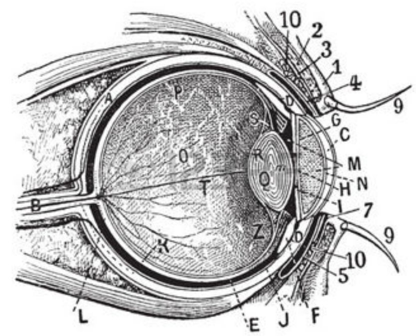

Noisy sensory noise (physical limitations)

discretization (spatial limitations)

noise in the environment

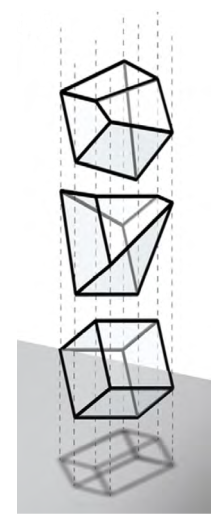

Ambiguous no unique reconstruction of environment

e.g. visual 3D to 2D mapping

mixture of odors

(Little) time contributes to uncertainty

There is no such a thing as an instantaneous percept

Yabus (1967)

Uncertain evidence is accumulated across time / space

Perceptual decisions (at least) require evidence accumulation across time

How much evidence should we accumulate?

More evidence is expected to lead to better decisions → why ever stop?

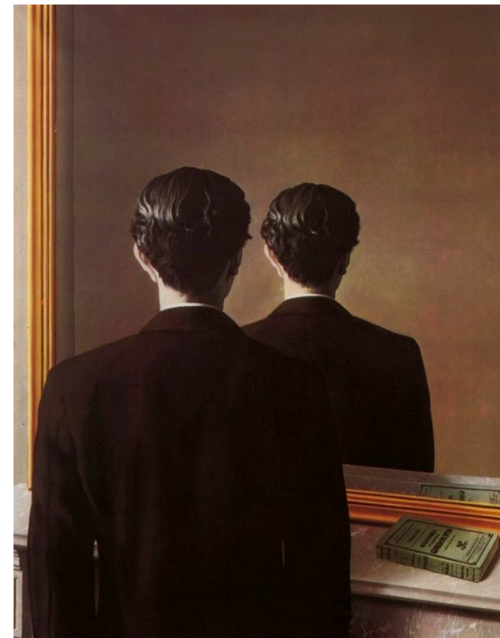

("Not to be reproduced", Magritte, 1937)

Reasons to stop accumulating: evidence/time is costly

world is volatile

evidence "flow" is limited

Costly evidence introduces speed/accuracy trade-off

accumulate evidence over time

⇩

commit to / execute choice

fast choices ← **speed/accuracy trade-off** → **slow choices**

might be inaccurate

come at low evidence cost

should be accurate

come at high cost

The speed/accuracy trade-off in experiments

Forced choice paradigm

- show two simuli
  (sequentially or simultaneously)

- choice is always A or B (or A and notA)

- choice is made (forced) on each trial

- difficulty might vary across blocks or trials

- record reaction time (RT)
  choice

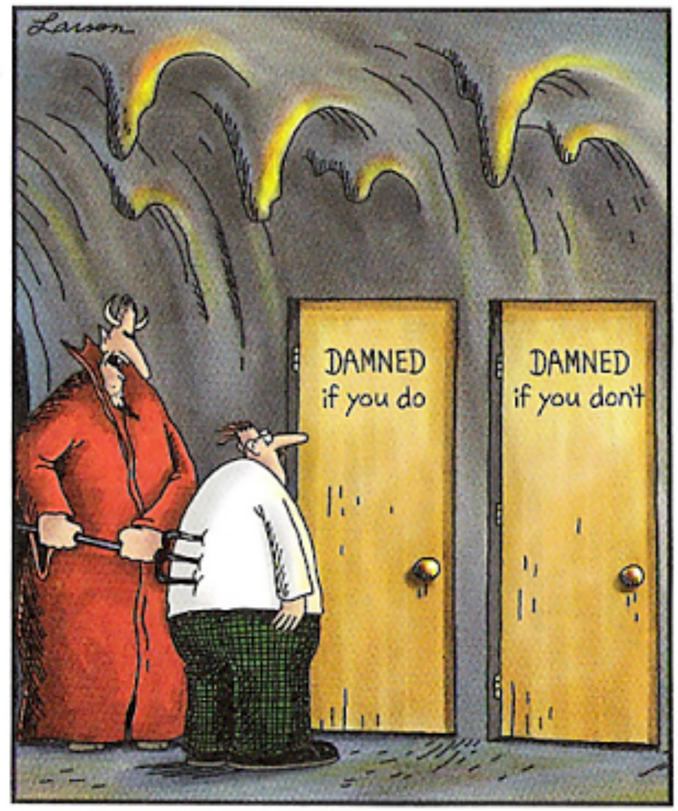

"C'mon, c'mon—it's either one or the other."

Examples

- *word* vs. *non-word* decisions

- numerosity judgments

- random dot motion task

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

stay

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

slan

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

gohm

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

goon

Word vs. *non-word* decisions

(e.g., Ratcliff, Gomez & McKoon, 2004)

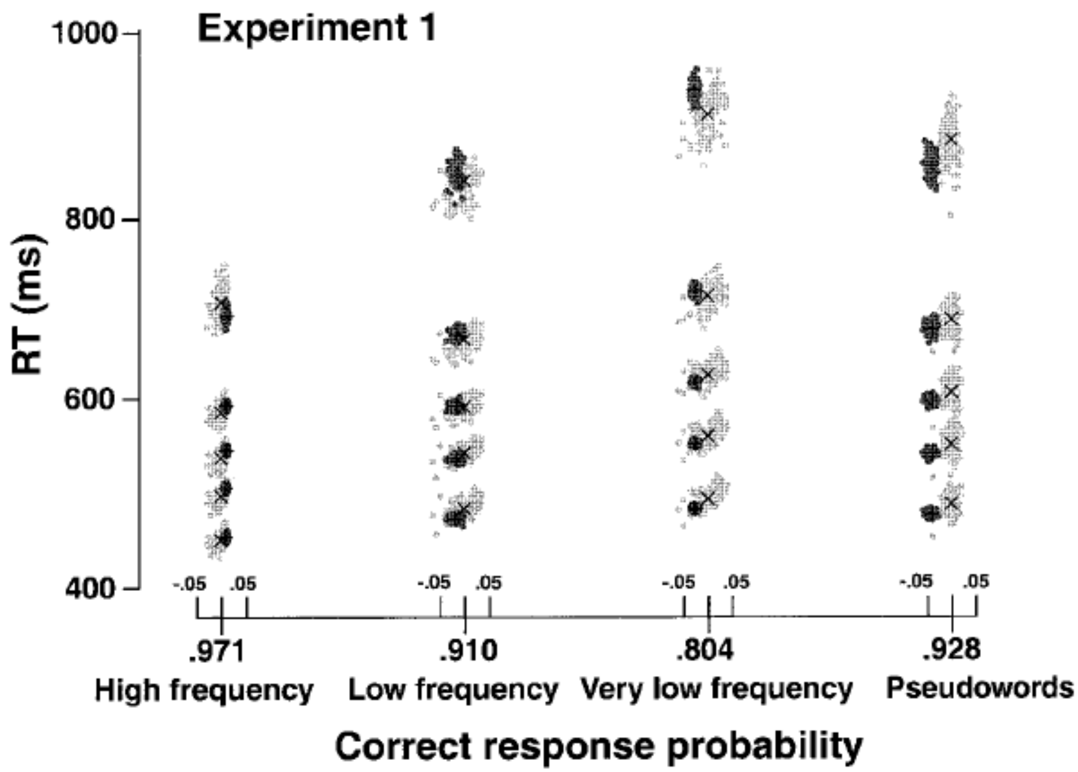

RT quantile plots

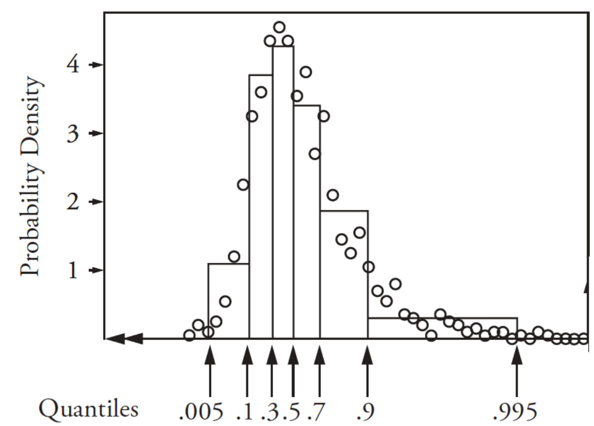

Ratcliff & Smith (2015)

Uncertainty: processing words / memory

Difficulty: word frequency / phonetic/lexical similarity / …

Usual findings: decisions faster and more accurate for high-frequency words

Numerosity judgments

(e.g., Ratcliff, 2006)

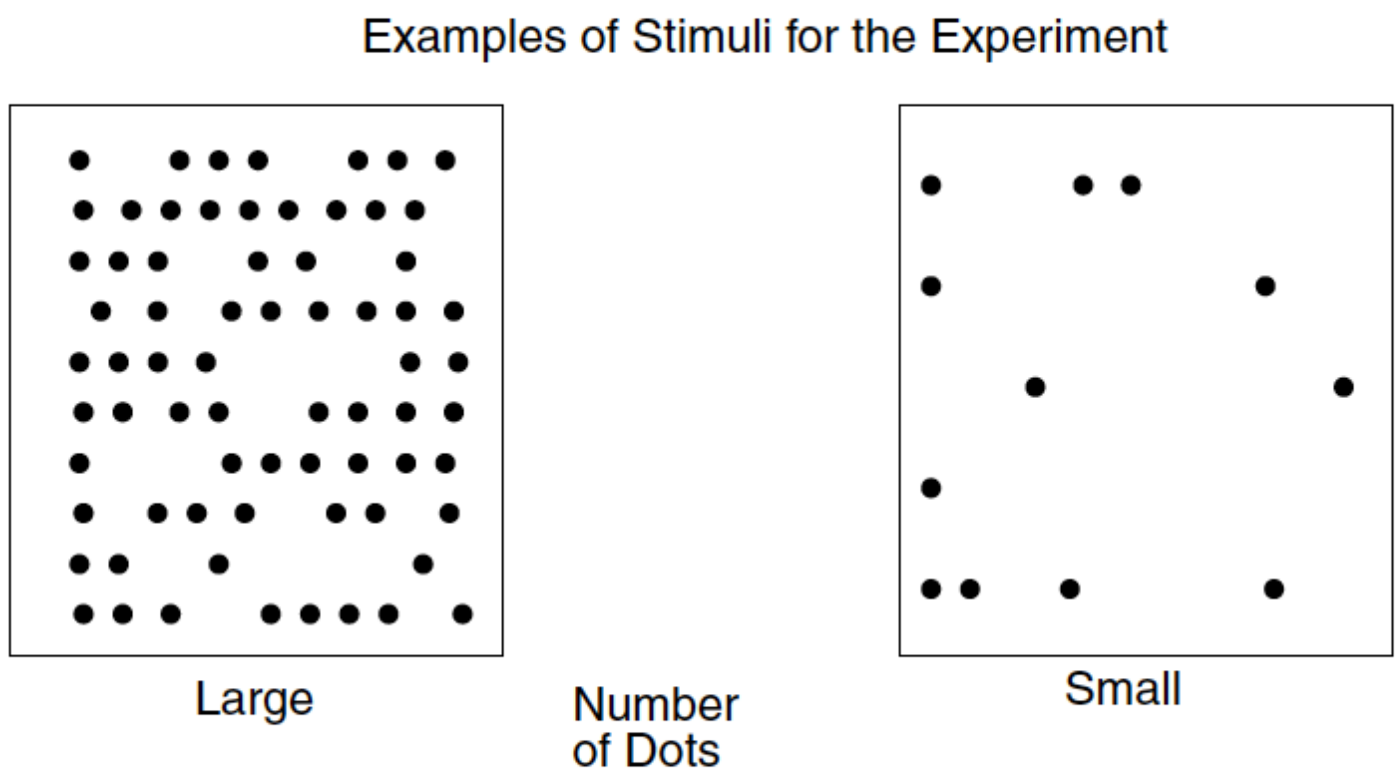

More/less than 50 dots?

Displays closer to 50-dot threshold: slower and less accurate

The random-dot motion task (RDM)

(e.g., Newsome, Britten, Movshon & Shadlen, 1989; Roitman & Shadlen, 2002)

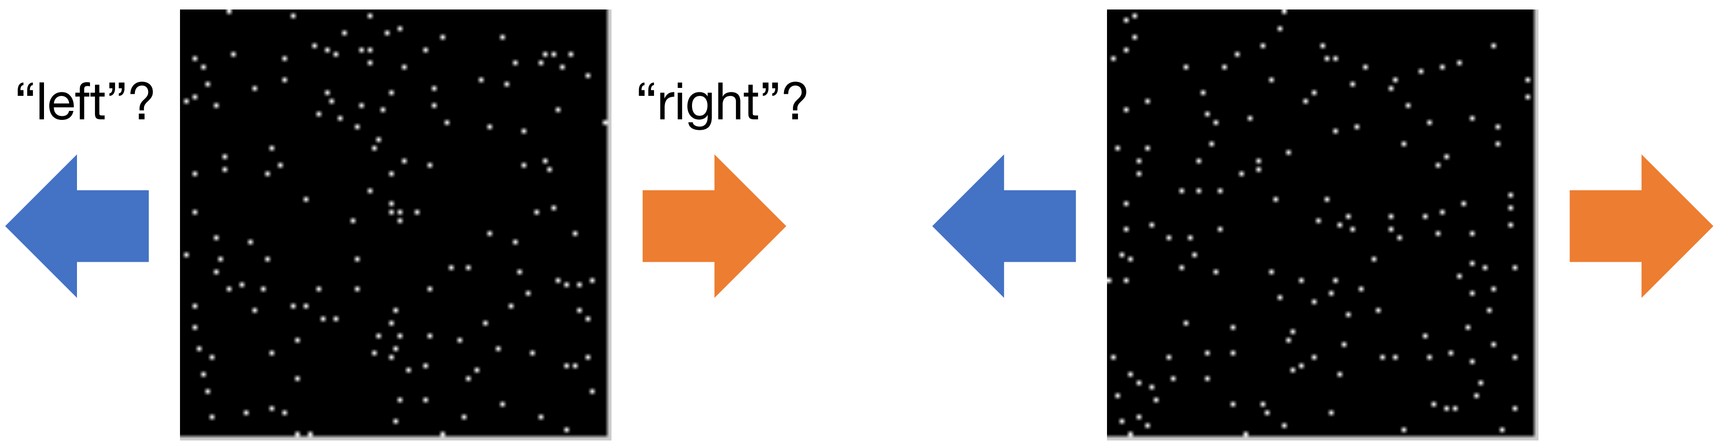

51.2% coherence

12.8% coherence

“respond as quickly and accurately as possible”

Uncertainty: stimulus is inherently ambiguous

Difficulty: coherence

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- → characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Behavior in the random dot motion task

Palmer, Huk & Shadlen (2005) dataset: 6 human subjects performing RDM task

`load('phs_[subj_id].mat')`

(`subj_id` ∈ {'ah', 'eh', 'jd', 'jp', 'mk', 'mm'})

Contains three vector, one element per trial:

`choice`	`0` - "left" / `1` - "right"
`rt`	reaction time in seconds
`cohs`	signed coherence, positive/negative – rightwards/leftwards motion

To become familiar with dataset:

- open `plot_psych_chron.m` in editor
- update line 17 to compute vector `corr_choice` (0 = incorrect, 1 = correct)
  Hint: choice is correct if "right" for rightward motion, "left" for leftward motion

Behavior in the random dot motion task

Computing correct choices

```
corr_choice = 0.5 * (sign(cohs + 1e-6) + 1) == choice;
```

0/1 for leftward/rightward motion

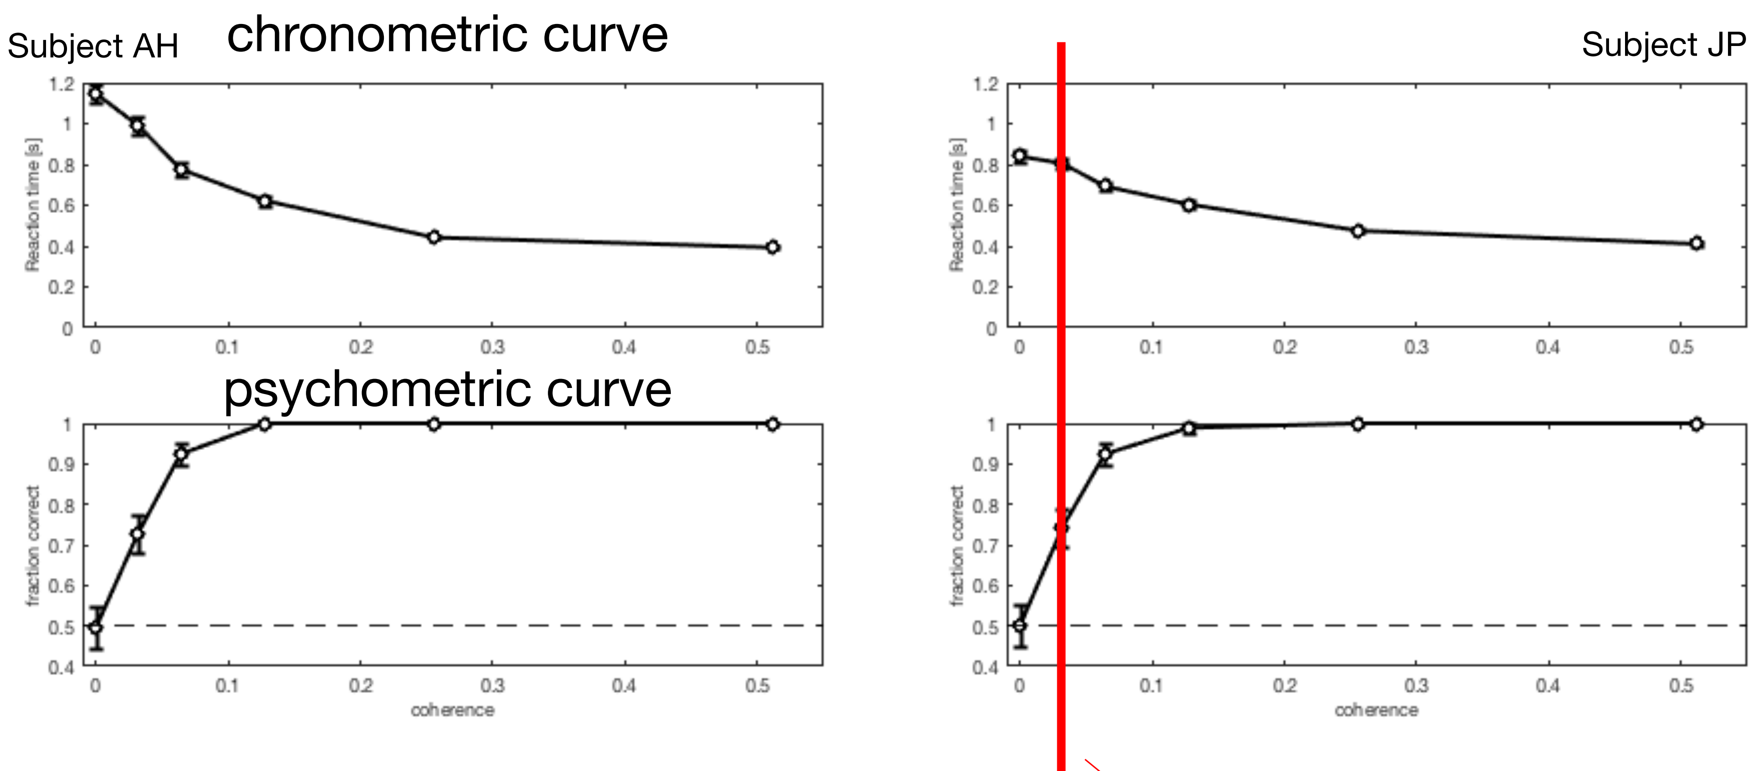

Higher coherence → faster, better choices

speed/accuracy trade-off
for *fixed* coherence

Speed/accuracy trade-off in the PHS dataset?

```
load('phs_[subj_id].mat')
plot_speed_accuracy
```

per-coherence RT median split

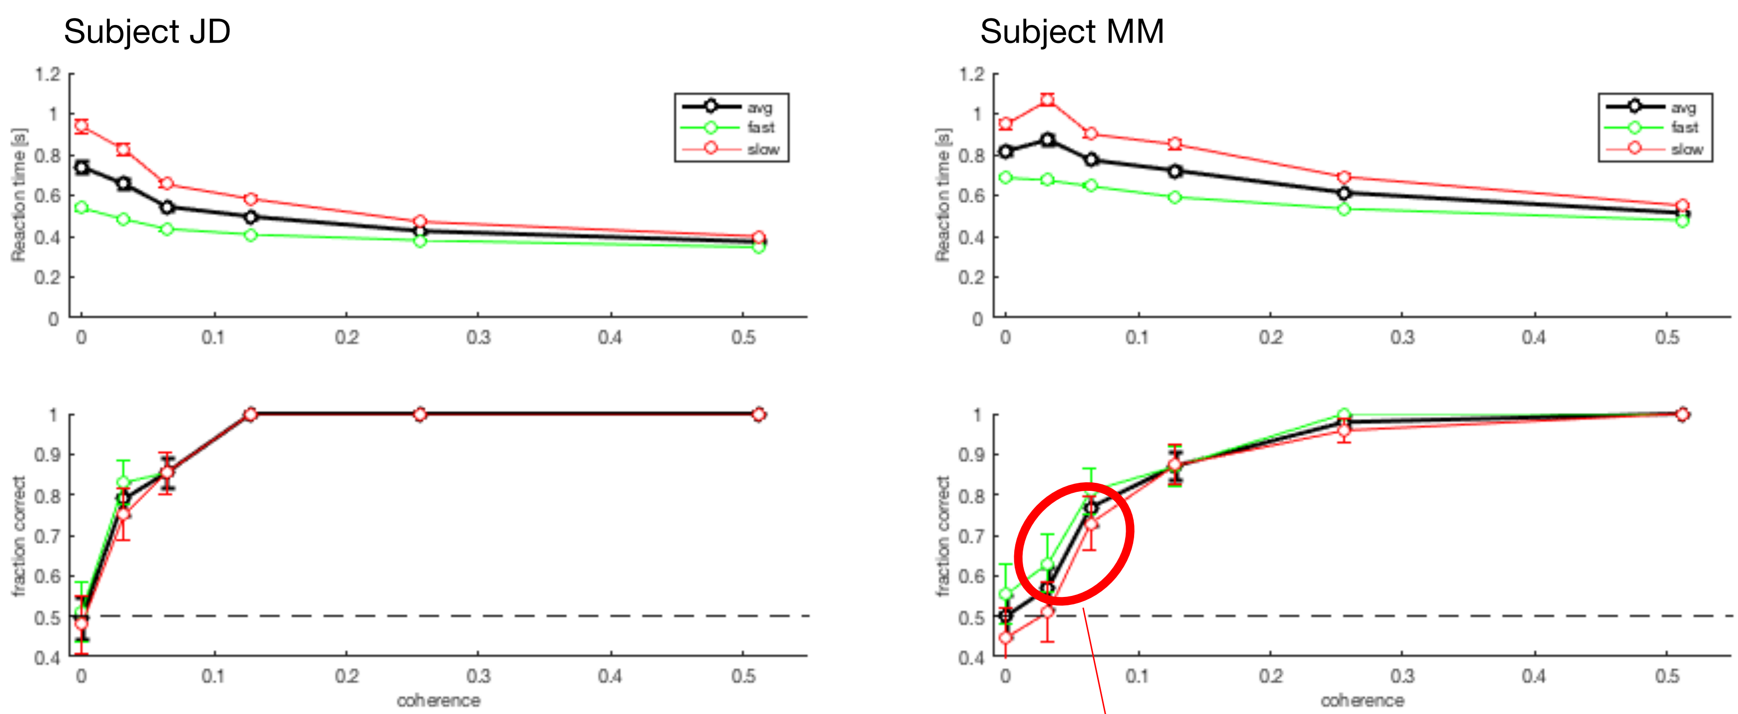

faster choices also more accurate?

Here, most RT fluctuations driven by fluctuations in stimulus *informativeness* (would need to compare fast/slow choices for same stimulus sequence)

Usually skewed reaction time distributions

Try `plot_rt_dist` and `plot_rt_quant`

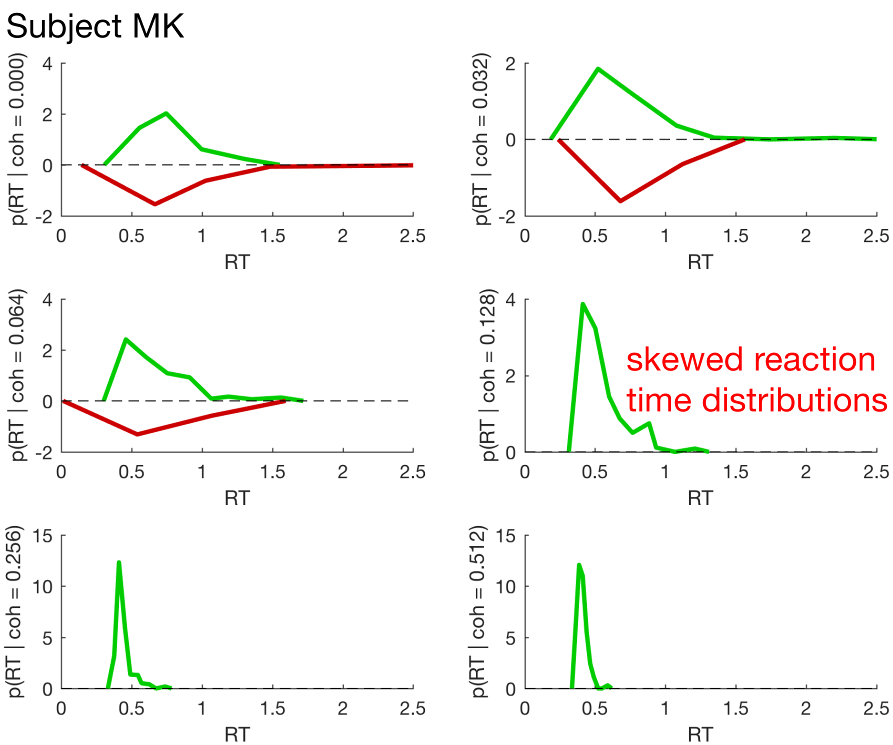

`plot_rt_dist(coh, choice, rt)`

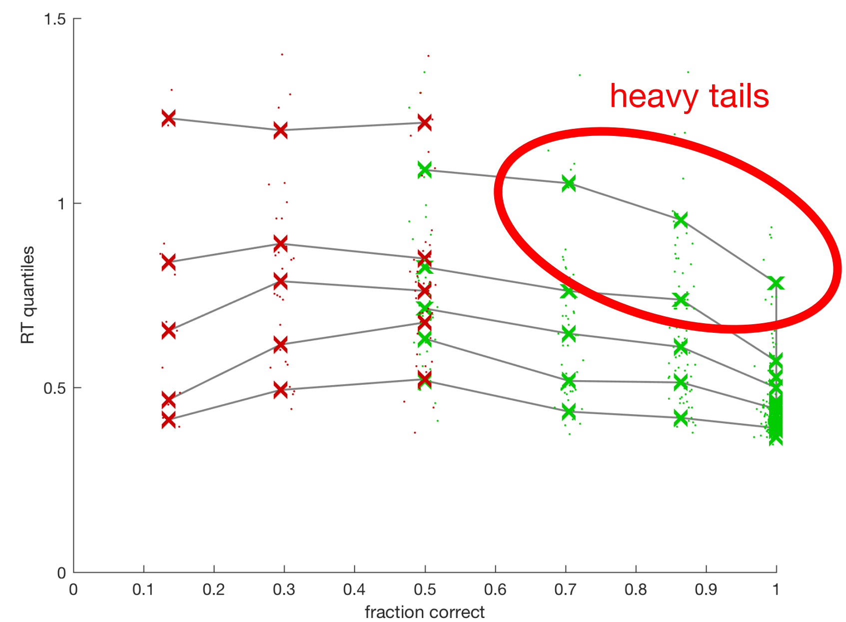

`plot_rt_quant(coh, choice, rt)`

Features of a successful decision-making model

Fits *mean reaction times* and
choice probability across conditions

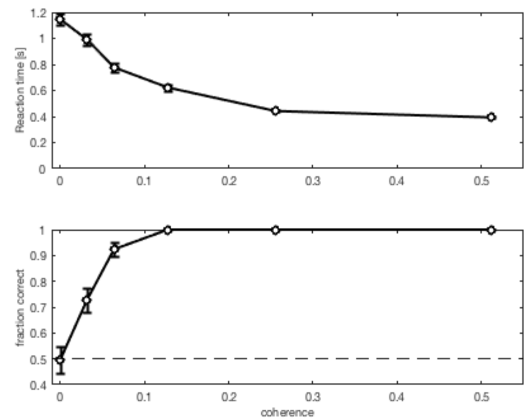

Accounts for *variability*:
reproduces RT distributions

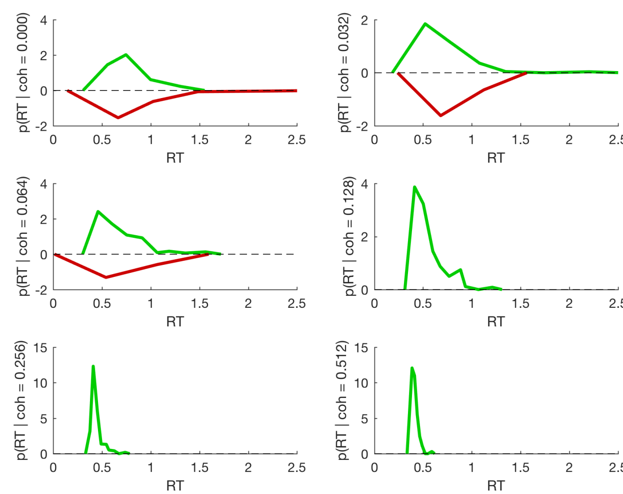

Reproduces *task difficulty influence*:
- easy task: fast choices, high accuracy
- hard task: slow choices, low accuracy
(to be revisited)

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

→ Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Accumulator models

Noisy evidence in small samples of continuous evidence stream

Accumulation to bound

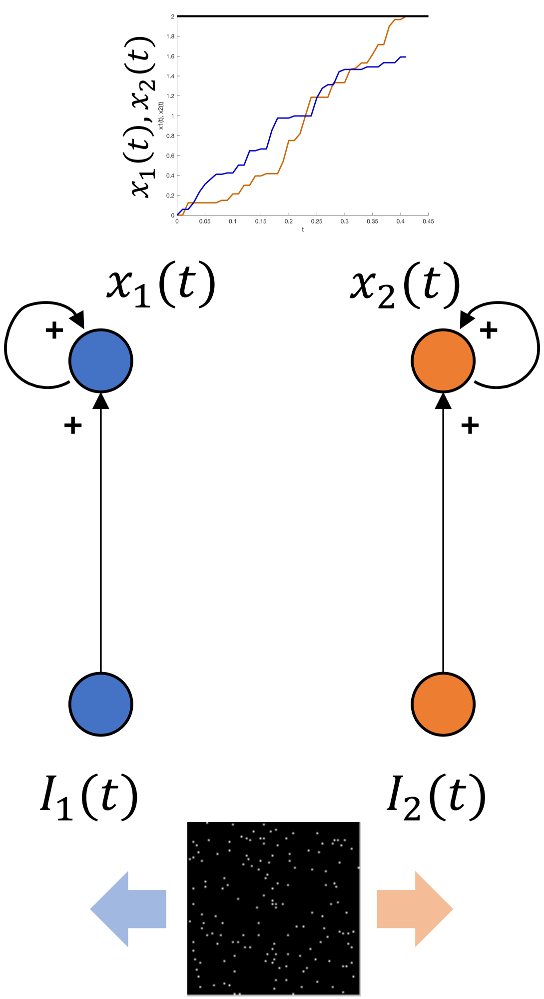

$$\frac{dx_1}{dt} = I_1$$

$$\frac{dx_2}{dt} = I_2$$

Inputs modulated by coherence, motion direction

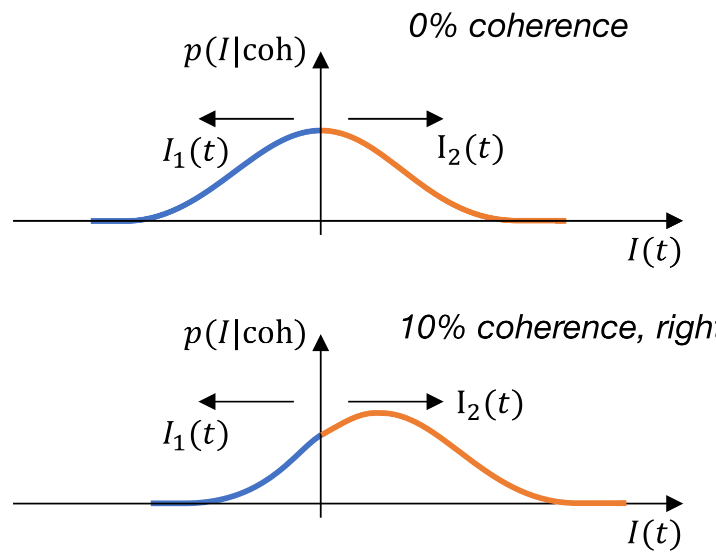

Exists in multiple variant, with discrete (Poisson) inputs, continuous (Gaussian) inputs, etc.

Accumulator model have their issues

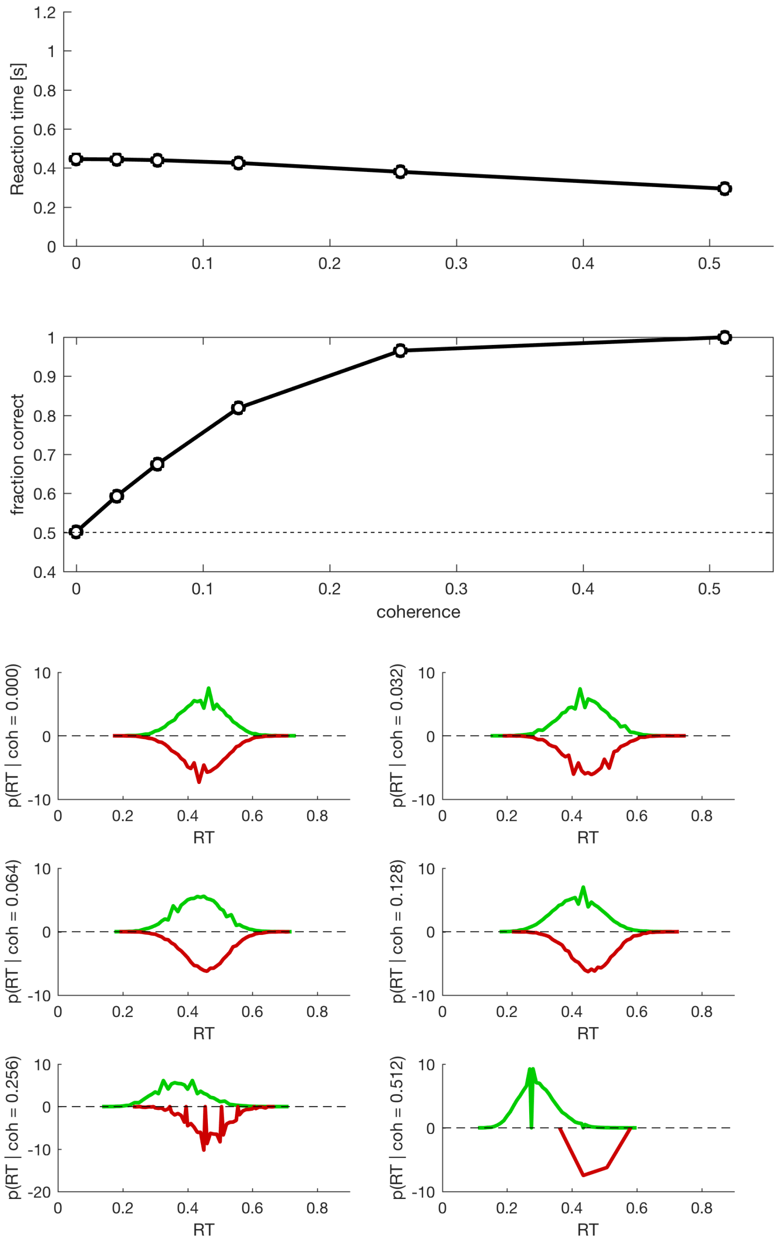

Don’t well reproduce
reaction-time modulation by difficulty

Feature variable reaction times,
but *not with with a heavy-tailed distribution*

sim_accum.m

The drift diffusion model

(or *diffusion decision model*; or *diffusion model*; Ratcliff, 1978)

Introduced by Ratcliff (1978) as model for memory recall;
one of the most successful models in neuroscience

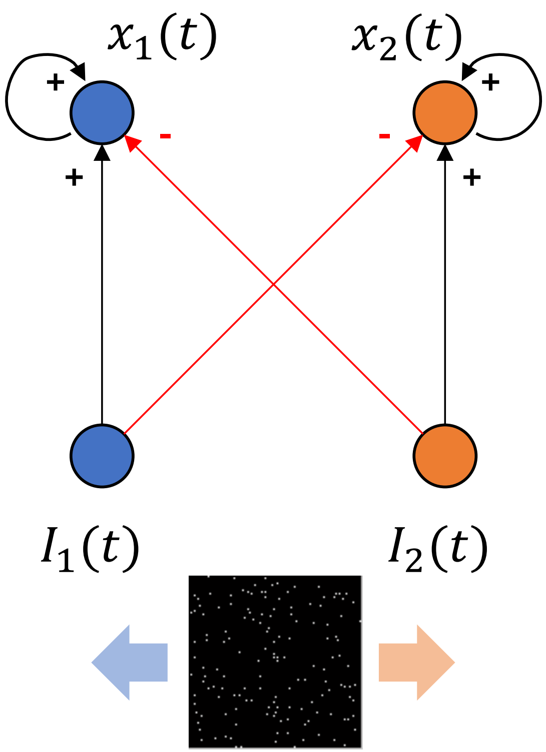

$$\frac{\mathrm{d}x_1}{\mathrm{d}t} = I_1 - I_2$$

$$\frac{\mathrm{d}x_2}{\mathrm{d}t} = I_2 - I_1 = -\frac{\mathrm{d}x_1}{\mathrm{d}t}$$

accumulators perfectly anti-correlated

single decision process

The drift diffusion model

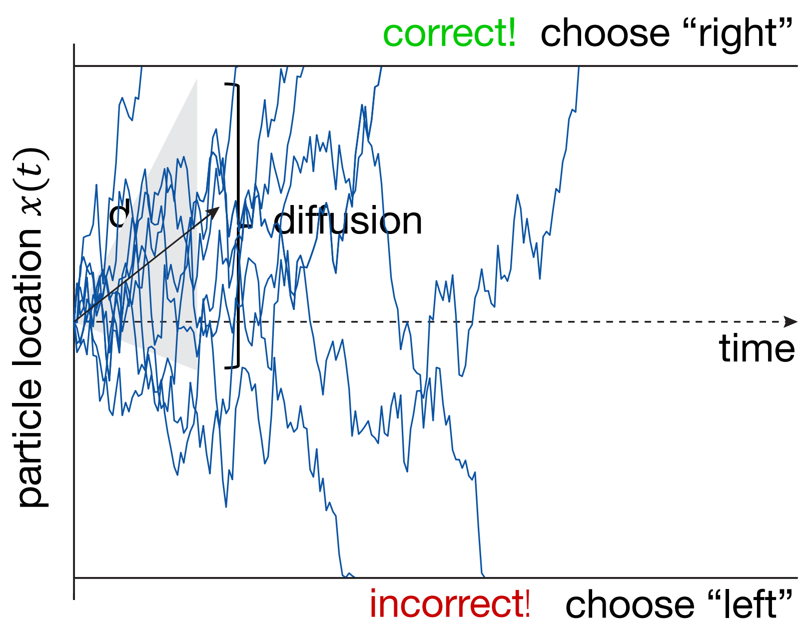

drift = K × coherence

$$\frac{\mathrm{d}x}{\mathrm{d}t} = I_1 - I_2 = \mu + \sigma\eta(t)$$

white noise process

drift

diffusion standard dev.

$|\mu|$ = mean evidence strength

$\mathrm{sign}(\mu)$ = determines correct choice

$\frac{|\mu|}{\sigma}$ = signal/noise ratio

accumulating uncertain evidence = stochastic particle motion

↓

commit to / execute choice = threshold crossing

Simulating the drift-diffusion model

Using the Euler method:

From continuous-time process…

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mu + \sigma\eta(t) \approx \frac{x(t+\delta t) - x(t)}{\delta t}$$

…to discrete-time simulation

$$x(t+\delta t) = x(t) + \mu\delta t + \sqrt{\delta t}\sigma z$$

$$z \sim \mathrm{N}(0,1)$$

(zero-mean unit-variance Gaussian random number)

Careful: too large δt cause biased first-passage time

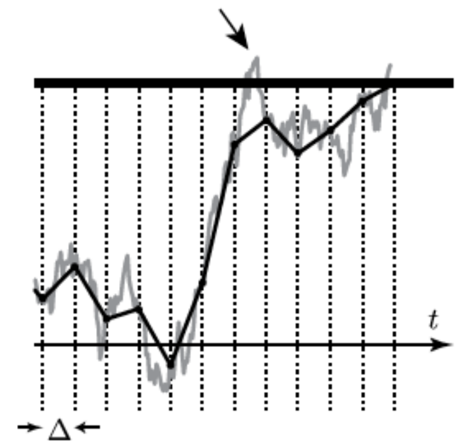

Drugowitsch (2016)

Alternatives: see `dm` library

See, for example, `sim_ddm.m`

Some diffusion model predictions

Generated with `sim_ddm.m`

$$\langle DT|\mu,\theta\rangle = \begin{cases} \theta^2, & \mu = 0 \\ \dfrac{\theta}{\mu}\tanh(\theta\mu), & \text{otherwise} \end{cases}$$

$$p(right) = \frac{1}{1+e^{-2\theta\mu}}$$

(e.g. Palmer, Huk & Shadlen, 2005)

$\mu = 0$

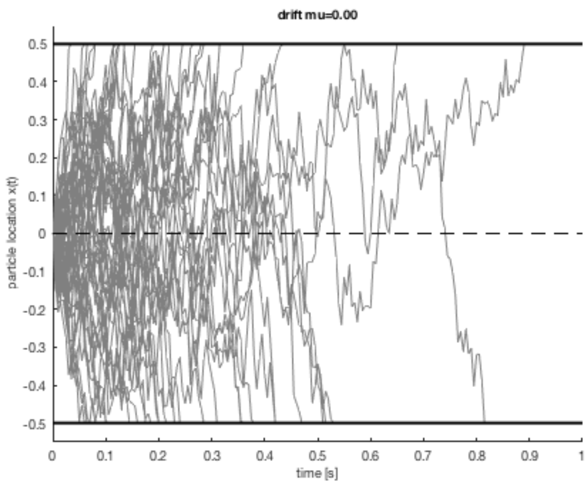

$\mu = 2$

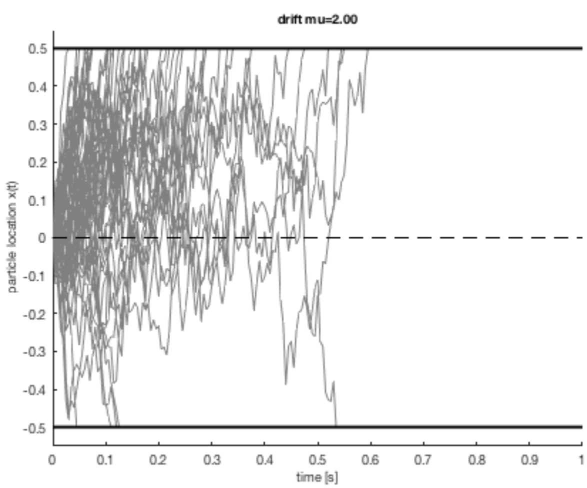

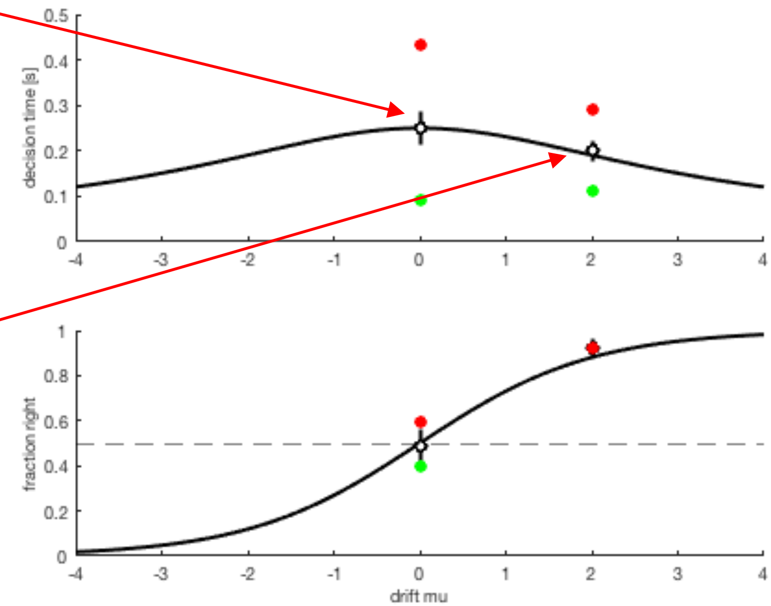

What happens for higher/lower bounds?

Try it out: `ddm_sim.m`, setting of `theta`

Adjusting drift and boundary heights

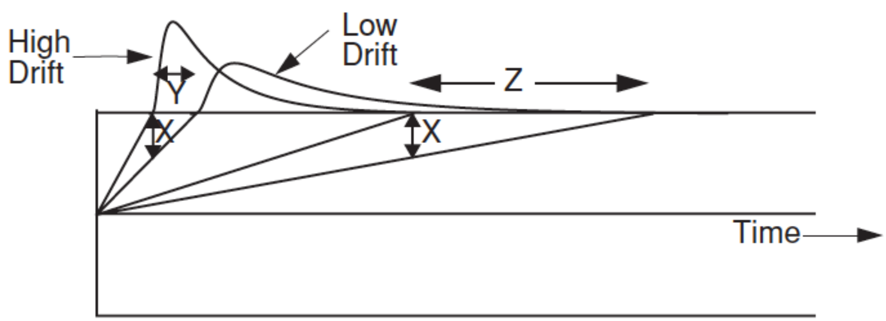

Lower drift:
slower, less accurate choices

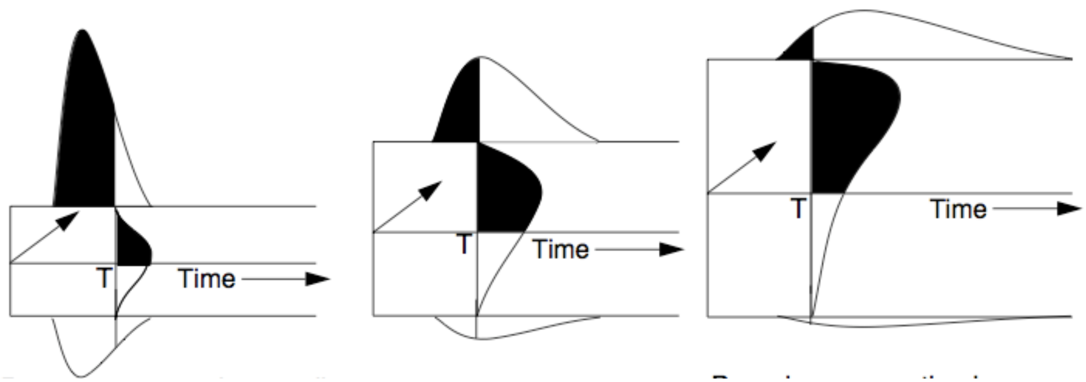

Raise bound:
Slower, more accurate choices

Speed/Accuracy tradeoff
Only **boundary separation** changes

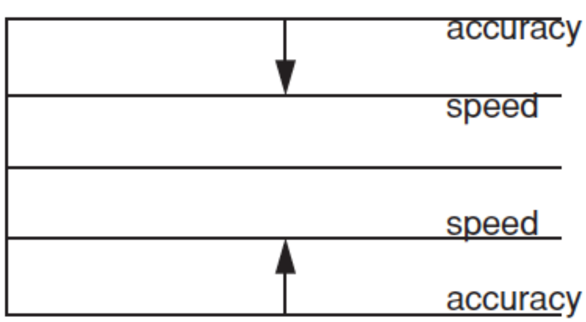

Quality of evidence from the stimulus
Only **drift rate** varies

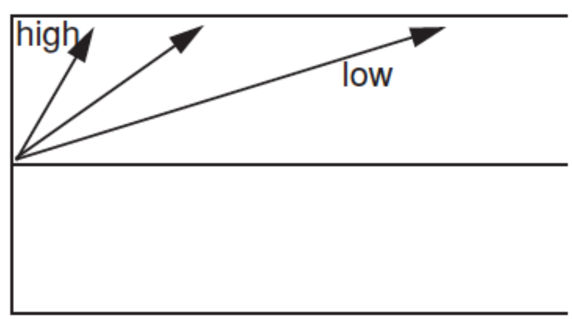

Ratcliff & McKoon (2008)

Diffusion models match well observed behavior

Assume that $\mu = k \times \text{coherence}$,

reaction time = diffusion model decision time DM + non-decision time t_{nd}.

Gives 3 parameters: k, θ, t_{nd}

Minimizing parameter log-likelihood
given mean RTs and choice probabilities (Palmer, Huk & Shadlen, 2005)

```
fit_psych_chron(cohs, choice, rt)
```

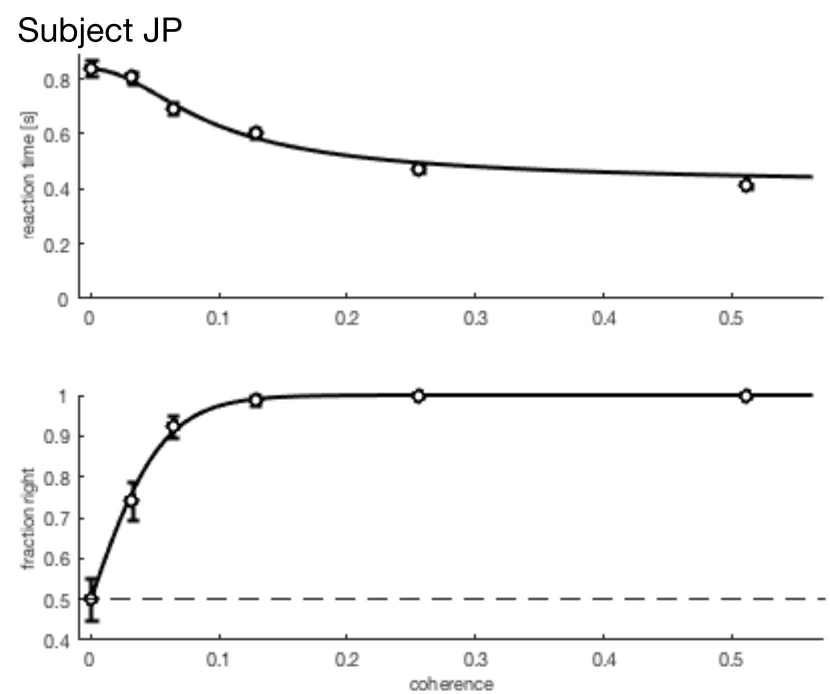

…but there are issues: #1 symmetry

Incorrect choices are frequently slower than correct choices

Uncomment relevant lines
in `plot_psych_chron.m`

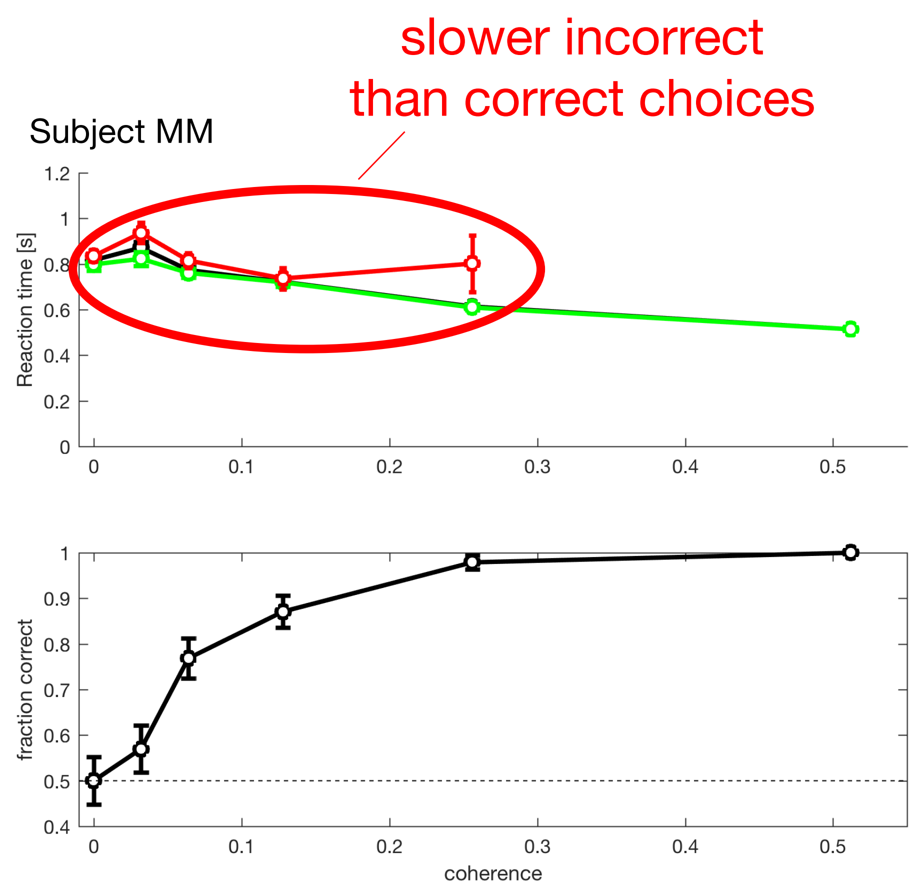

(but see subj JP)

Vanilla diffusion models predict symmetric RT distributions

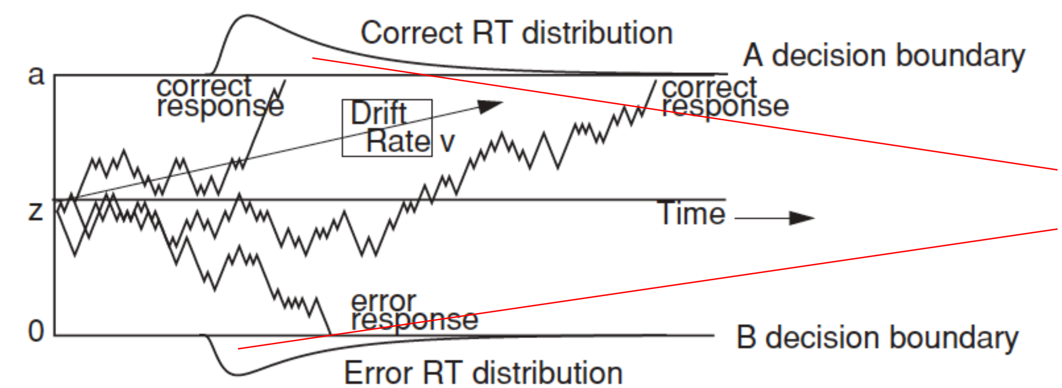

reaction time distributions are symmetric!

Ratcliff & McKoon (2008)

Reason: flipping path scales its probability by a constant

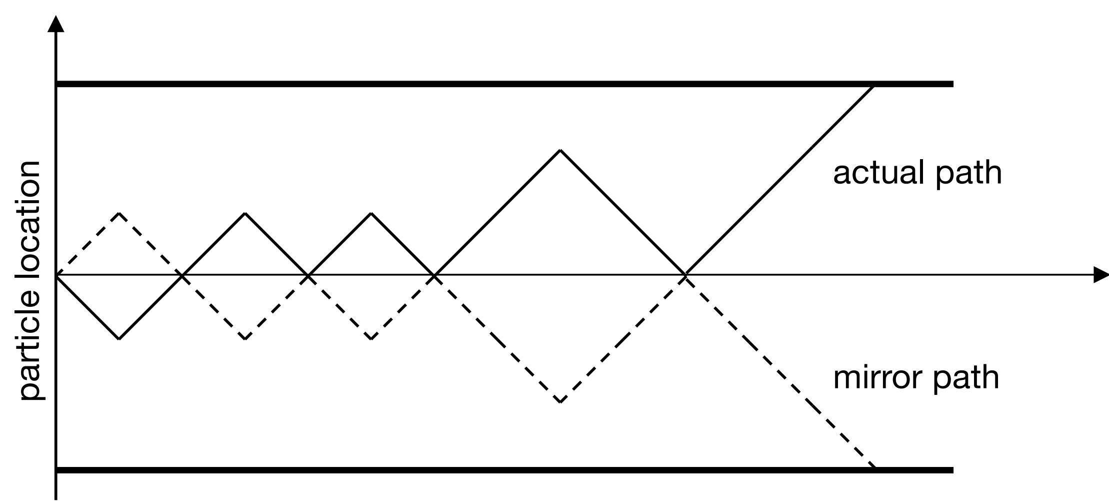

n_{up} up-steps with probability p_{up}

n_{down} down-steps with probability $1 - p_{up}$

$$p(\text{actual path}) = p_{up}^{n_{up}} (1 - p_{up})^{n_{down}}$$

n_{down} up-steps with probability p_{up}

n_{up} down-steps with probability $1 - p_{up}$

$$p(\text{mirror path}) = p_{up}^{n_{down}} (1 - p_{up})^{n_{up}}$$

$$p(\text{mirror path}) = \left(\frac{p_{up}}{1 - p_{up}} \right)^{\theta} p(\text{actual path})$$

…but there are issues: #2 long-tail predictions

Observed reaction time distributions don't always have a long tail

Try `plot_fitted_rt_dists(cohs, choice, rt)`

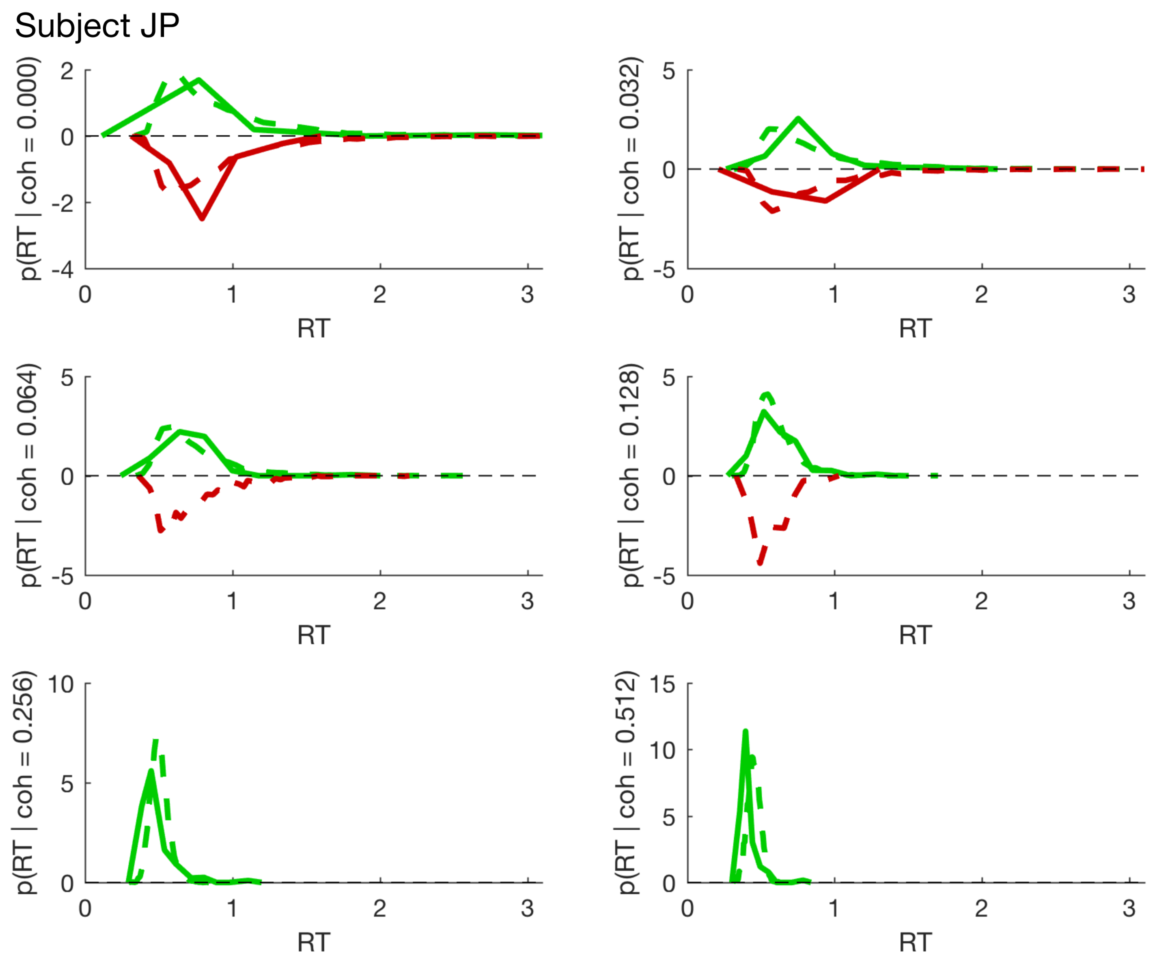

Monkeys are even less patient

Roitman & Shadlen (2002) dataset: 2 monkeys performing RDM task

```
load('rs_[monkey_id].mat')
plot_fitted_rt_dists(cohs, choice, rt)
```

(monkey_id ∈ {'b', 'n'})

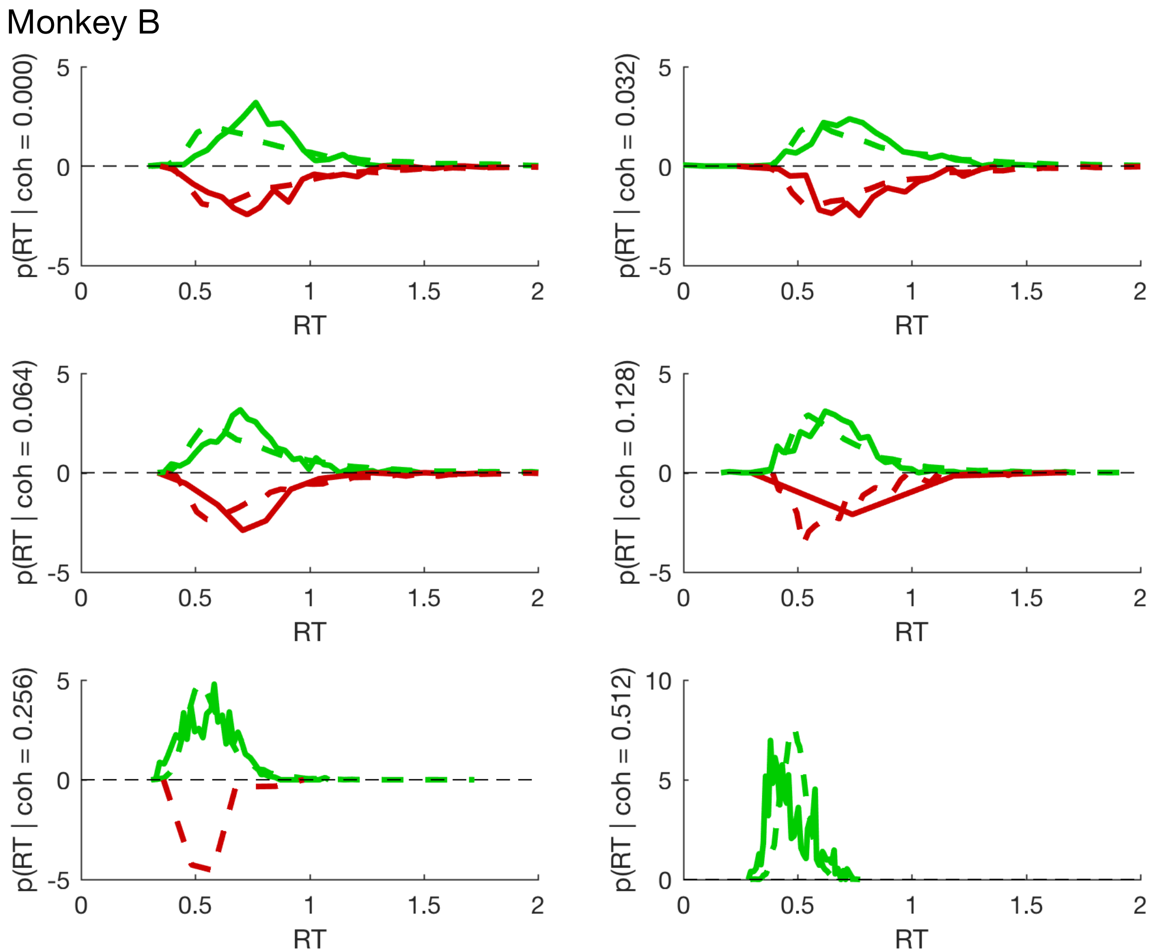

Heuristic “fix”: the Ratcliff diffusion model

\+ diffusion models implement both, and fit mean RTs and choice probabilities

\- predict same correct/incorrect RTs

\- don’t match reaction time distributions

How to fix: add more parameters!

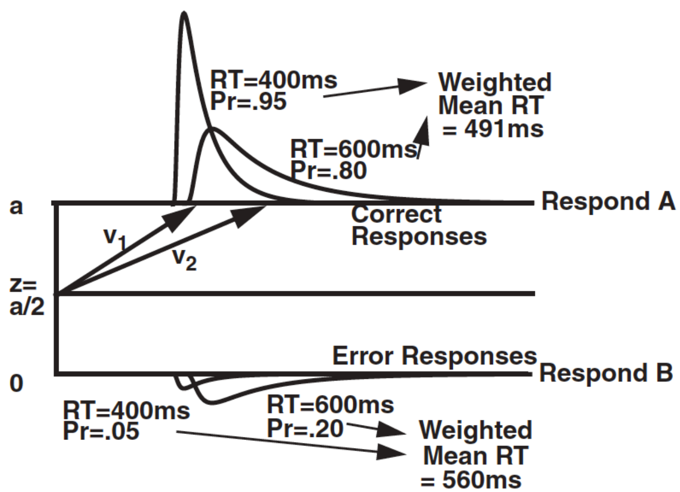

Ratcliff & McKoon (2008)

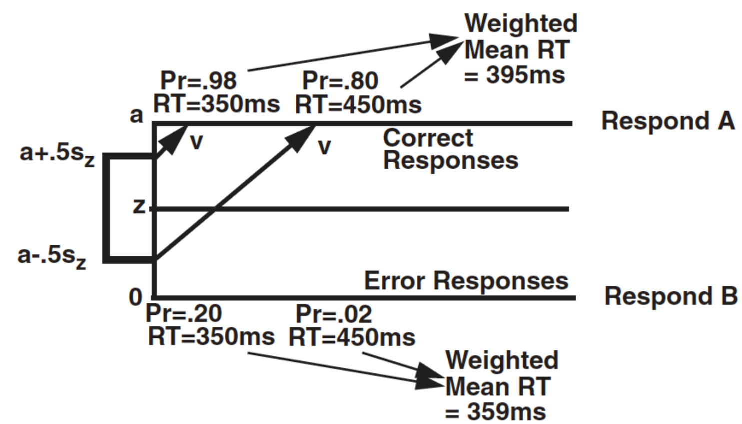

Variable drift rates: slower errors

Variable starting point: faster errors

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

→ Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Normative approach: how ought we make decisions?

accumulating evidence

handling uncertain information
using Bayesian statistics

Rev. Thomas Bayes
(1701-1761)

deciding when to decide

trading of benefits with costs
using Dynamic programming

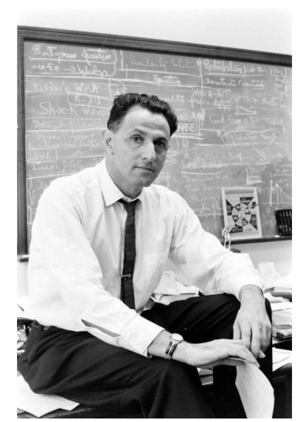

Richard E. Bellman
(1920-1984)

A model for the momentary evidence

Assume: fixed coherence μ_0, two motion directions, $\mu \in \{-\mu_0, \mu_0\}$.

uniform prior, $p(\mu = -\mu_0) = p(\mu = \mu_0) = \frac{1}{2}$

At any point n in time: noisy observation x_n of μ,

$$p(x_n|\mu) = \underbrace{\mathrm{N}(x_n|\mu, 1)}_{\text{“}x_n\text{ is Gaussian/Normal with mean }\mu\text{ and variance 1”}}$$

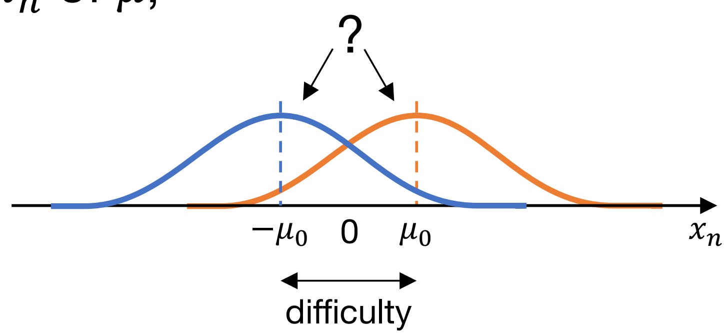

Observe $x_1, x_2, \ldots$; identify if they came from blue or orange distribution

$$p(\mu = \mu_0 | x_{1:n}) = ?$$

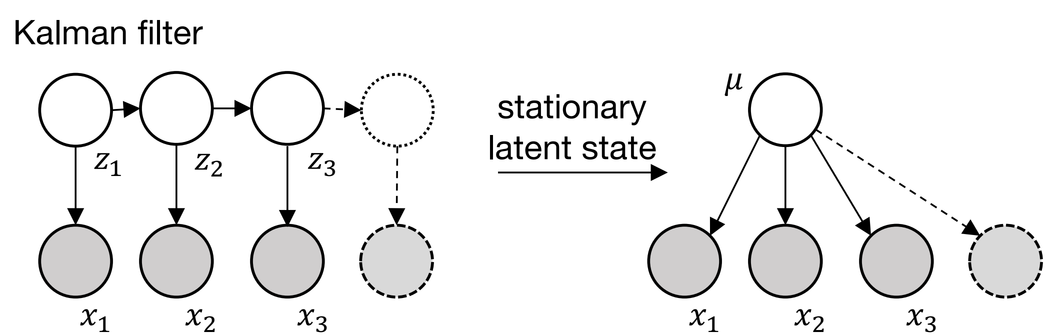

Why not use Kalman fiter? Explicit derivations provide further insight

Deriving the posterior

$$
\begin{aligned}
p(\mu = \mu_0 | x_{1:N}) &= \frac{p(x_{1:N}|\mu = \mu_0)p(\mu = \mu_0)}{p(x_{1:N})} \\
&\propto_\mu p(x_{1:N}|\mu = \mu_0)p(\mu = \mu_0) \\
&= p(\mu = \mu_0) \prod_n N(x_n|\mu = \mu_0, 1) \\
&\propto_\mu \prod_n \frac{1}{2\pi} e^{-\frac{(x_n-\mu)^2}{2}} \\
&\propto_\mu e^{-\frac{N\mu^2}{2} + \mu \sum_n x_n} \\
&= e^{-\frac{N\mu_0^2}{2} + \mu_0 \sum_n x_n}
\end{aligned}
$$

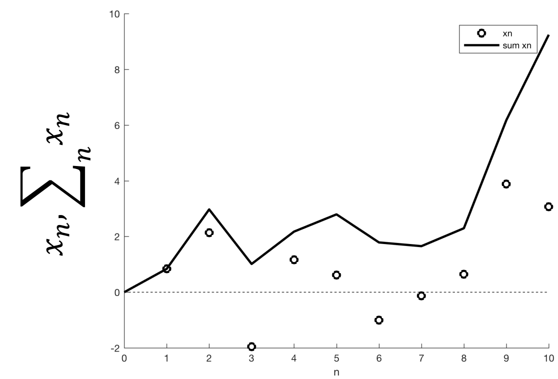

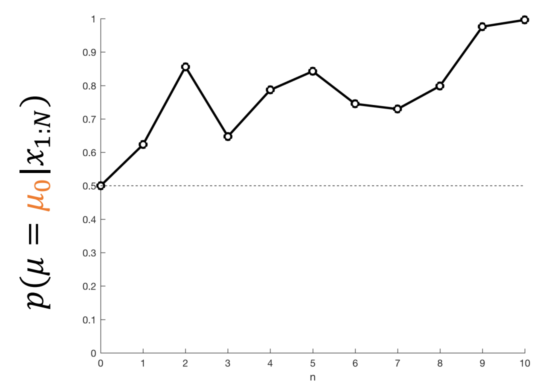

accum_evidence_discrete.m

$$
p(\mu = -\mu_0 | x_{1:N}) \propto_\mu e^{-\frac{N(-\mu_0)^2}{2} - \mu_0 \sum_n x_n}
$$

$$
\begin{aligned}
p(\mu = \mu_0 | x_{1:N}) &= \frac{e^{-\frac{N\mu_0^2}{2} + \mu_0 \sum_n x_n}}{e^{-\frac{N\mu_0^2}{2} + \mu_0 \sum_n x_n} + e^{-\frac{N\mu_0^2}{2} - \mu_0 \sum_n x_n}} \\
&= \frac{1}{1 + e^{-2\mu_0 \sum_n x_n}}
\end{aligned}
$$

logistic sigmoid

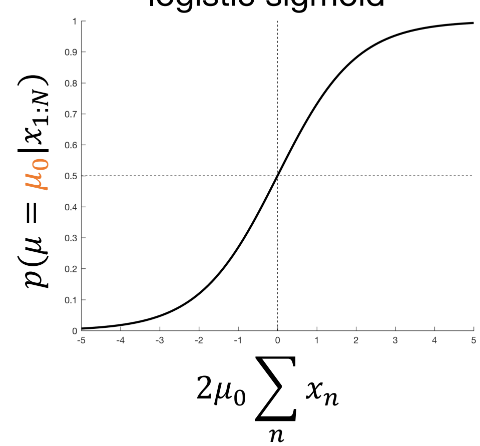

Moving to continuous time

Smaller time steps δt: less reliable evidence δx_n per time step

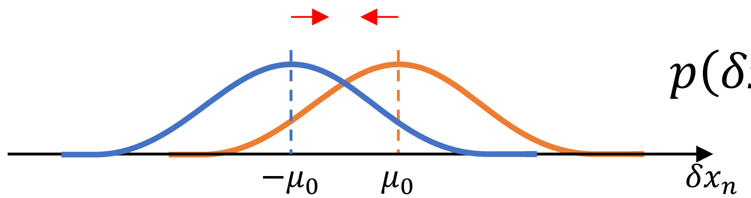

$$p(\delta x_n|\mu) = \mathrm{N}(\delta x_n|\mu\delta t, \delta t)$$

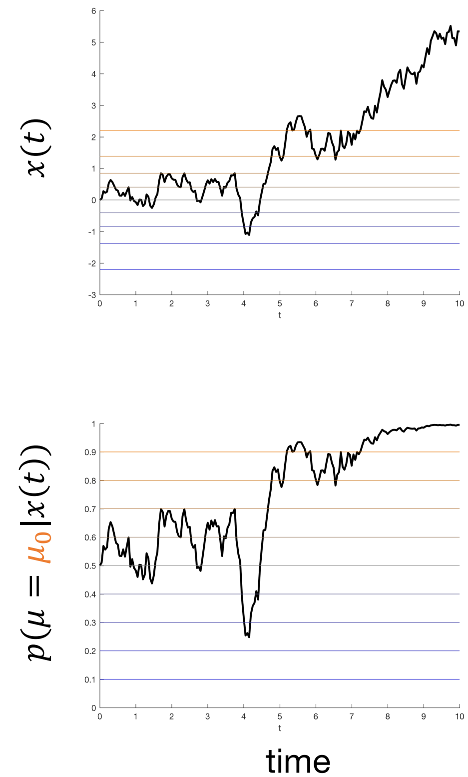

Find $p(\mu = \mu_0|x_{1:N})$, using $N\delta t = t$ and $\sum_n \delta x_n = x(t)$

$$p(\mu = \mu_0|x_{1:N}) \propto_\mu e^{-\frac{\mu_0^2}{2}N\delta t + \mu_0 \sum_n \delta x_n}$$

$$= e^{-\frac{\mu_0^2}{2}t + \mu_0 x(t)}$$

$$p(\mu = \mu_0|x_{1:N}) = \frac{1}{1 + e^{-2\mu_0 x(t)}}$$

`accum_evidence_continuous.m`

Shows why diffusion models are useful

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mu + \eta(t)$$

$x(t) > 0$ implies $p(\mu = \mu_0|x_{1:N}) > \frac{1}{2}$

$x(t) < 0$ implies $p(\mu = \mu_0|x_{1:N}) < \frac{1}{2}$

Normative approach: how ought we make decisions?

accumulating evidence

deciding when to decide

handling uncertain information
using Bayesian statistics

Rev. Thomas Bayes
(1701-1761)

trading of benefits with costs
using Dynamic programming

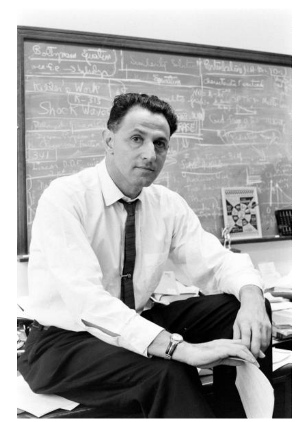

Richard E. Bellman
(1920-1984)

When to stop accumulating evidence?

Assume: aim is to maximize reward
(reward 1/0 for correct/incorrect choices)

more momentary evidence
↓
higher expected reward
↓
accumulate forever!

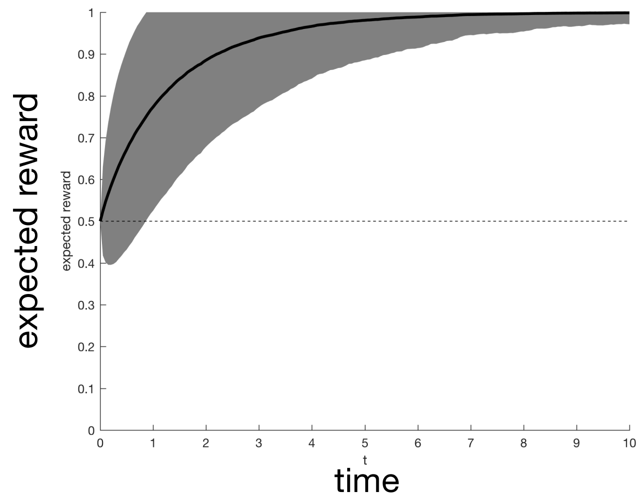

accum_evidence_reward.m

Stopping to accumulate is only rational *in presence of cost*

- Motivational/effort cost
- Cost of attention/computation
- Opportunity cost; less time on future choices

(can be internal & external)

Objective functions

Maximizing expected reward for single choice

Payoff 1 for correct choice, 0 for incorrect choice, cost c per second accumulation

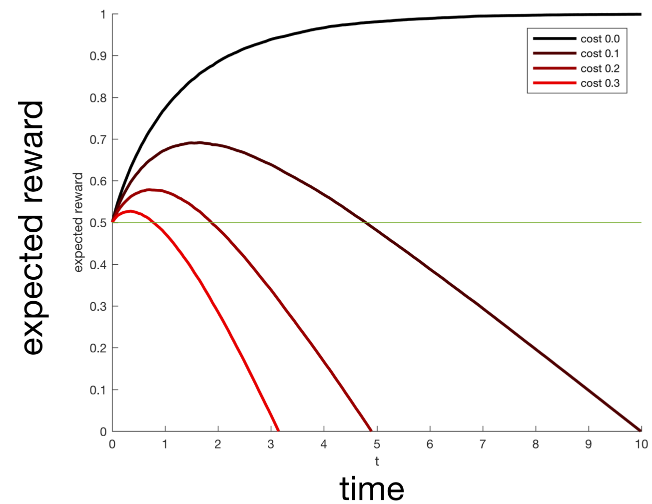

$$ER = PC - c\langle t \rangle$$

Maximizing expected reward across multiple choices

Sequence of choices with inter-choice-interval t_i

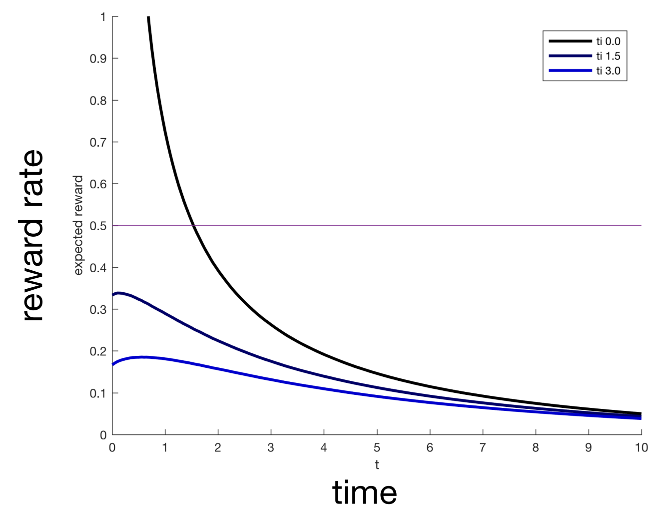

$$RR = \frac{PC - c\langle t \rangle}{\langle t \rangle + t_i}$$

Optimal stopping required closed-loop control

Interlude: dynamic programming (DP)

Markov decision process (MDP)

- set of states, $s_1, s_2, \ldots$
- set of actions, $a_1, a_2, \ldots$
- transition probabilities, $p(s'|s, a)$
- rewards, $r(s, a)$
- discount factor, $\gamma \leq 1$

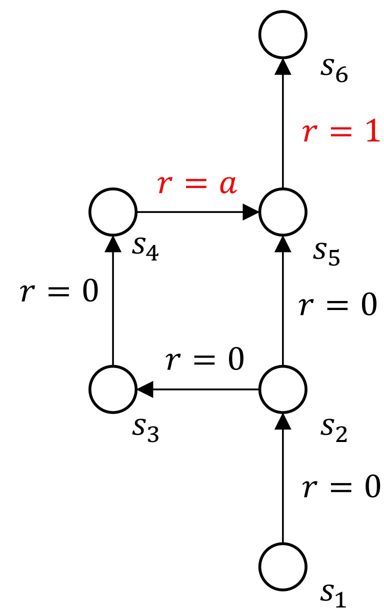

Aim: find optimal policy, $\pi(s)$ returning action for each state to maximize expected discounted future reward (or *return*)

$$V^{\pi}(s) = \left\langle \sum_{n=0}^{\infty} \gamma^n r(s_n, \pi(s_n)) \right\rangle_{p(s_1, s_2, \ldots|\pi)} = r\big(s, \pi(s)\big) + \gamma \langle V^{\pi}(s') \rangle_{p(s'|s,\pi)}$$

"*value*" of state s under policy π

Example: navigation

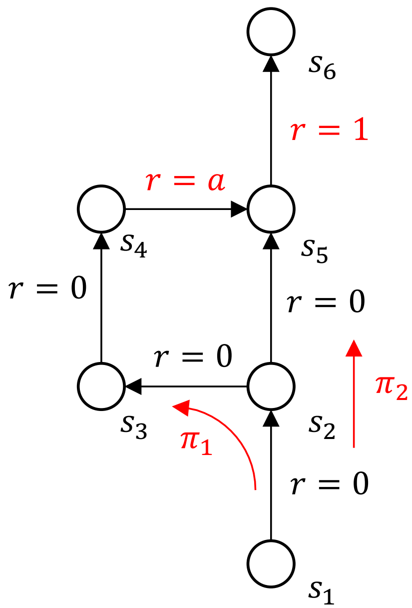

Specific solution:

$$\left.\begin{aligned} V^{\pi_1}(s_1) &= \gamma^3 a + \gamma^4 \\ V^{\pi_2}(s_1) &= \gamma^2 \end{aligned}\right\} \text{ choose } \pi_1 \text{ if } a \geq \frac{1}{\gamma} - \gamma$$

	s_5	s_4	s_3	s_2	s_1
π_1	1	$a+\gamma$	$\gamma a + \gamma^2$	$\gamma^2 a + \gamma^3$	$\gamma^3 a + \gamma^4$
π_2	1	✘	✘	γ	γ^2

Bellman's principle of optimality

"optimal policy: whatever initial state/decision, the remaining decisions must constitute an optimal policy with regard to state resulting from first decision"

Bellman's equation: $V^*(s) = \max_{\mathrm{a}} \left\{ r(s,a) + \gamma \langle V^*(s') \rangle_{p(s'|s,a)} \right\}$

the maximizing action provides the optimal policy

Dynamic programming applied to optimal stopping

- set of states, $s_1, s_2, \ldots$ ⟶ accumulated evidence/belief, $g(t) \equiv p(\mu = \mu_0 | x(t))$
- set of actions, $a_1, a_2, \ldots$ ⟶ accumulate/make choice
- transition probabilities, $p(s'|s,a)$ ⟶ change of accumulated evidence, belief transition $p(g'|g)$
- rewards, $r(s,a)$ ⟶ cost for accumulation/rewards

  choose μ_0: $r = g$

  choose $-\mu_0$: $r = 1 - g$

  accumulate another δt: $r = -c\delta t$
- discount factor, $\gamma \leq 1$ ⟶ assume $\gamma = 1$

Bellman's equation for perceptual decisions

$$V(g) = \max\left\{g, 1 - g, \langle V(g') \rangle_{p(g'|g)} - c\delta t\right\}$$

The belief transitions function

Examples for $p(g'|g)$

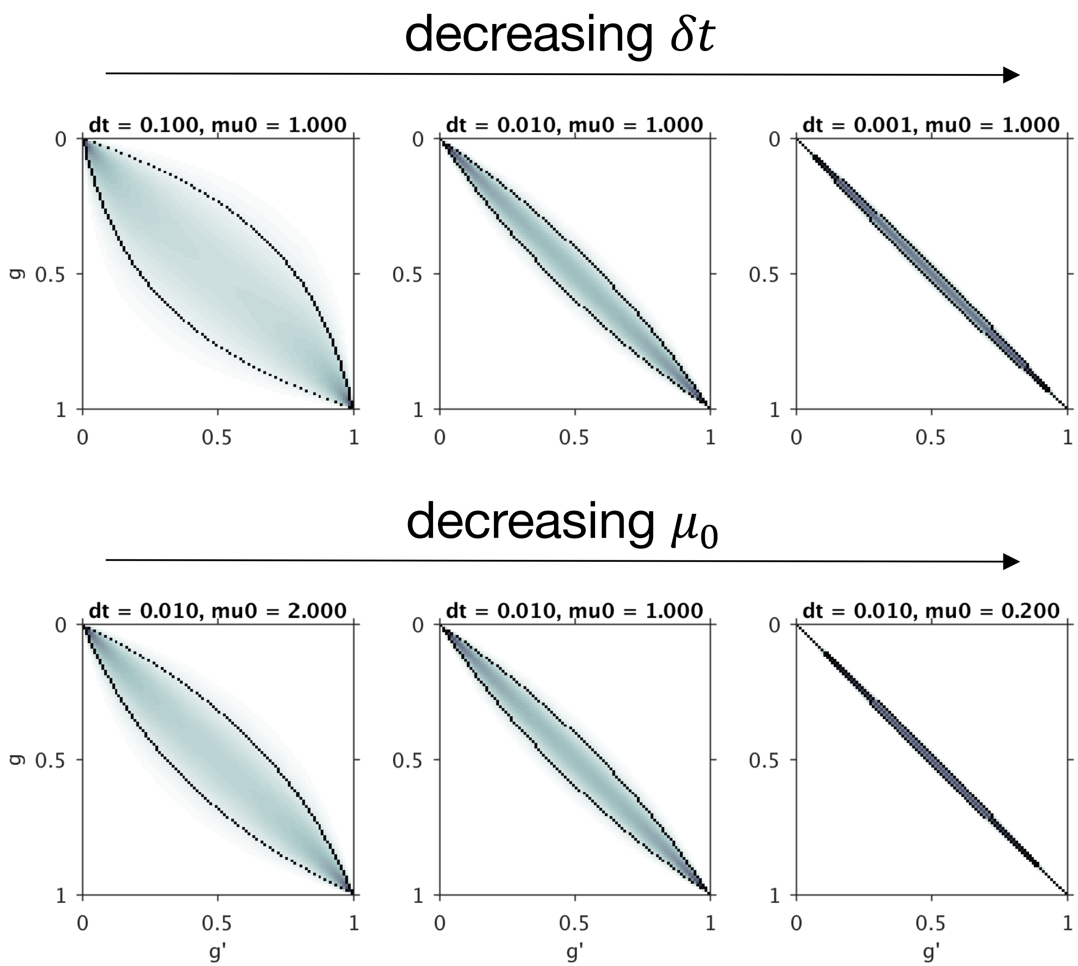

plot_g_trans_point_hyp.m

The value function for perceptual decisions

$$V(g) = \max\left\{g, 1-g, \langle V(g')\rangle_{p(g'|g)} - c\delta t\right\}$$

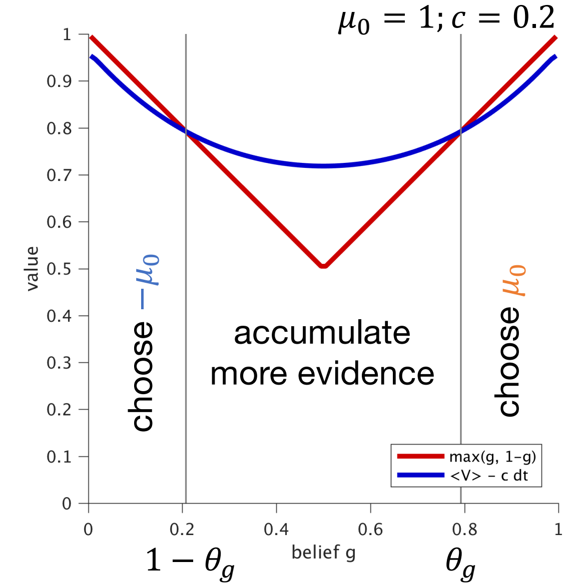

What happens if c or μ_0 changes?

Try it out:
`plot_dp_valueintersect_point(`μ_0, c`)`

`plot_dp_diffusion_point(`μ_0, c`)`:

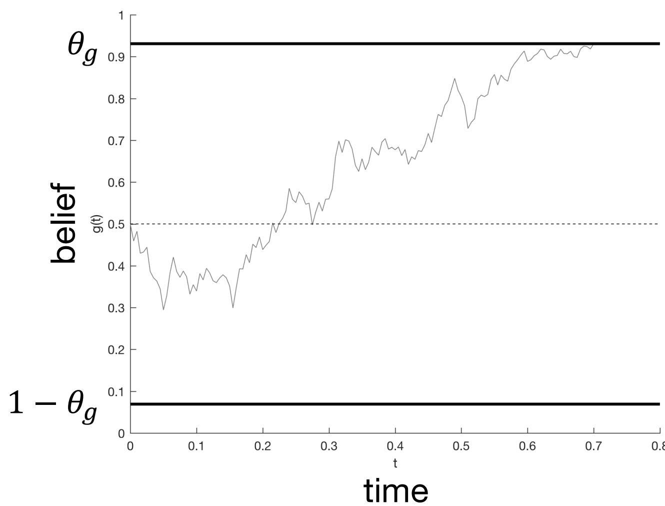

$$x(t) = \frac{1}{2\mu_0}\log\frac{g(t)}{1-g(t)}$$

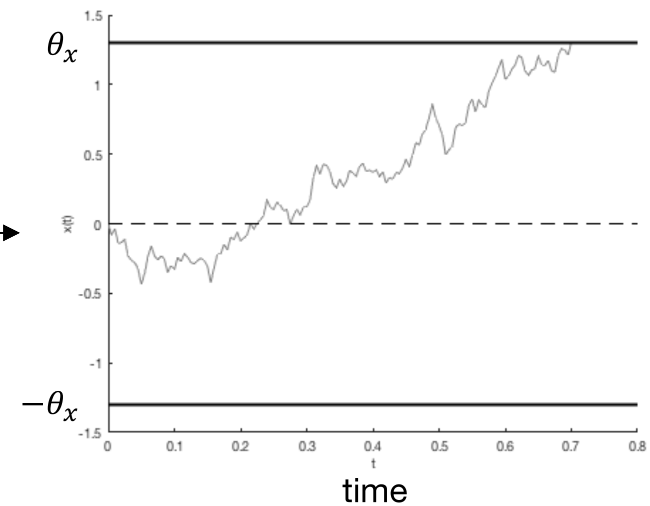

Diffusion models implement the reward-maximizing strategy

Finding the bound without dynamic programming

We now know: diffusion model with time-invariant bound is optimal

Initial aim: maximize $ER = PC - c\langle t \rangle$ ← maximize directly

$$PC = \frac{1}{1+e^{-2\mu_0\theta}} \qquad \langle t \rangle = \frac{\theta}{\mu_0}\tanh(\mu_0\theta)$$

Complete `direct_bound`(μ_0, c) in `plot_dp_bound_direct_maximization.m`

```
ER_deriv = @(theta) (mu0 - 2 * c * theta) * sech(theta * mu0)^2 / 2 - …
                    c * tanh(theta * mu0) / mu0;
theta = fzero(ER_deriv, 1);
```

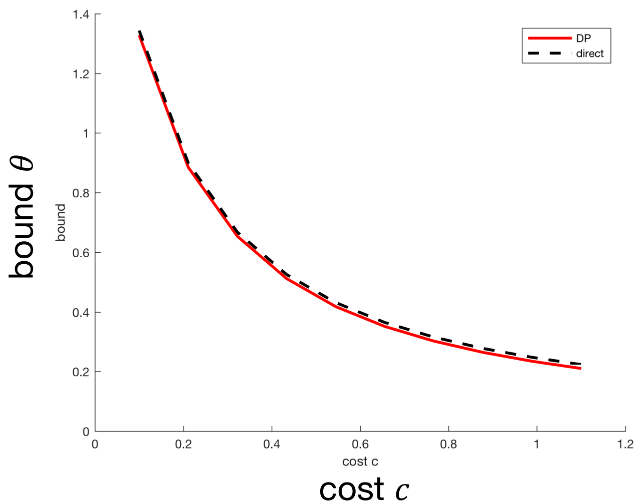

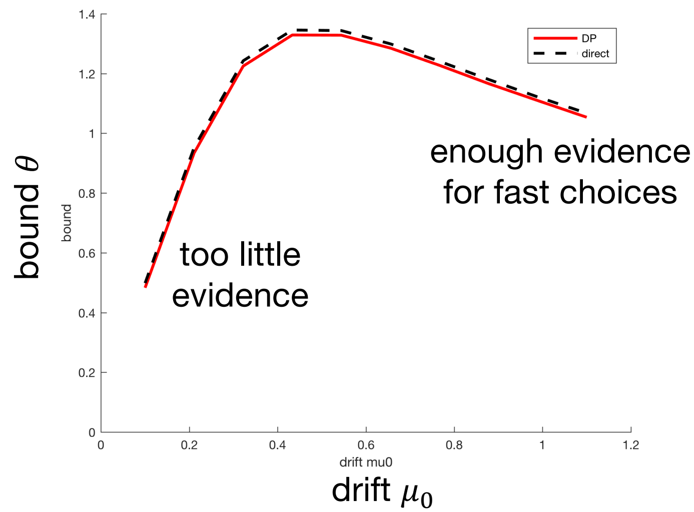

The sequential probability ratio test (SPRT)

For this simple case, the optimal policy has been known for a while.

Sequential probability ratio test (SPRT) (Wald, 1947; Wald & Wolfowitz, 1948; Turing, 194?)

Given two hypotheses H_1, H_2 with known likelihoods $p(x|H_1), p(x|H_2)$;
sequence $x_1, x_2, \ldots$ generated by which hypothesis?

Among all test with same power (type 1 error),
SPRT requires least samples on average (Wald & Wolfowitz, 1948).

SPRT accumulates evidence as long as

$$B^* \leq \frac{\prod_n p(x_n|H_1)}{\prod_n p(x_n|H_2)} \leq A^*$$

Relates to diffusion models and expected reward maximization (Bogacz et al., 2006)

Limitation: assumes known likelihood functions (e.g. known coherence)
the same applies to our derivation so far

This rarely holds in real-world decisions!

Road map

Perceptual decision-making

speed/accuracy trade-off

experiments investigating perceptual decisions

characteristics of behavior

Decision-making models

accumulator / diffusion models

fit to behavior & issues

Normative analysis

simple scenario: task difficulty known

→ more complex: varying task difficulty

time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Introducing difficulty as a nuisance

Nuisance: not central to the question, but we have to deal with it

e.g., RDM stimulus: $\underbrace{\text{motion direction}}_{\text{want to know}} + \underbrace{\text{motion coherence}}_{\text{don't care}}$

latent state of interest,
$z = \operatorname{sign}(\mu) \in \{-1, 1\}$

nuisance,
determines difficulty,
$y = |\mu|$

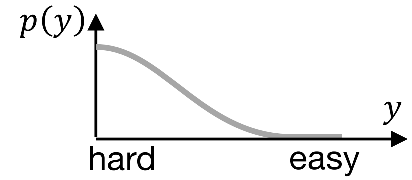

Stimulus variable: $\mu = z \times y$

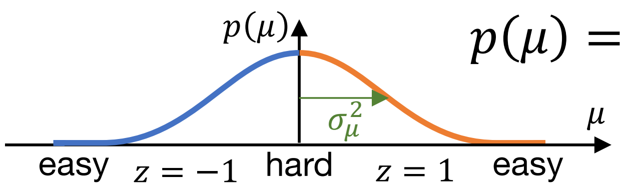

$p(\mu) = N(0, \sigma_\mu^2)$

overall difficulty

Momentary evidence: $p(\delta x_n|\mu) = \underbrace{N(x_n|\mu\delta t, \delta t)}_{\text{noisy information about } \mu}$

Aim: $\underbrace{p(z = 1|\delta x_1, \delta x_2, \dots)}_{\text{identify latent state } \textit{without} \text{ nuisance}} = \int p(z = 1, y|\delta x_1, \delta x_2, \dots)\mathrm{d}y = p(\mu \geq 0|\delta x_1, \delta x_2, \dots)$

Evidence accumulation with nuisance

Derivation in two steps: posterior over latent state and nusiance, …

$$p(\mu|\delta x_{1:N}) \propto_\mu N\left(\mu\middle|0,\sigma_\mu^2\right)\prod_n N(\delta x_n|\mu\delta t,\delta t)$$

$$\propto_\mu e^{-\frac{\mu^2}{2}\left(\frac{1}{\sigma_\mu^2}+t\right)+\mu x(t)}$$

$$\propto_\mu N\left(\mu\middle|\frac{x(t)}{\sigma_\mu^{-2}+t},\frac{1}{\sigma_\mu^{-2}+t}\right)$$

accum_evidence_gauss.m

…then averaging over nuisance

$$p(\mu \geq 0|x(t),t) = \int_0^\infty p(\mu|\delta x_{1:N})\mathrm{d}\mu = \Phi\left(\frac{x(t)}{\sqrt{\sigma_\mu^{-2}+t}}\right)$$

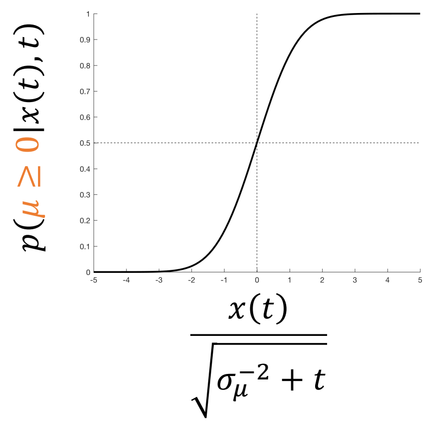

Posterior belief now depends on both $x(t)$ and t

Consequences for optimal stopping

Mapping between belief $g(t)$ and particle location $x(t)$ becomes time-dependent

$$g(t) \equiv p(\mu \geq 0 | x(t), t) = \Phi\left(\frac{x(t)}{\sqrt{\sigma_\mu^{-2} + t}}\right)$$

→ the expected change $\underbrace{p(g'|g,t)}$ also depends on time

required to compute expected return for accumulating more evidence

→ Value function depends on g (or x) and time

$$V(g,t) = \max\left\{\underbrace{g, 1-g}_{\text{deciding immediately}}, \underbrace{\langle V(g', t+\delta t)\rangle_{p(g'|g,t)} - c\delta t}_{\text{accumulating more evidence, and deciding later}}\right\}$$

→ decision boundaries depend on time

The belief transition function, unknown evidence reliability

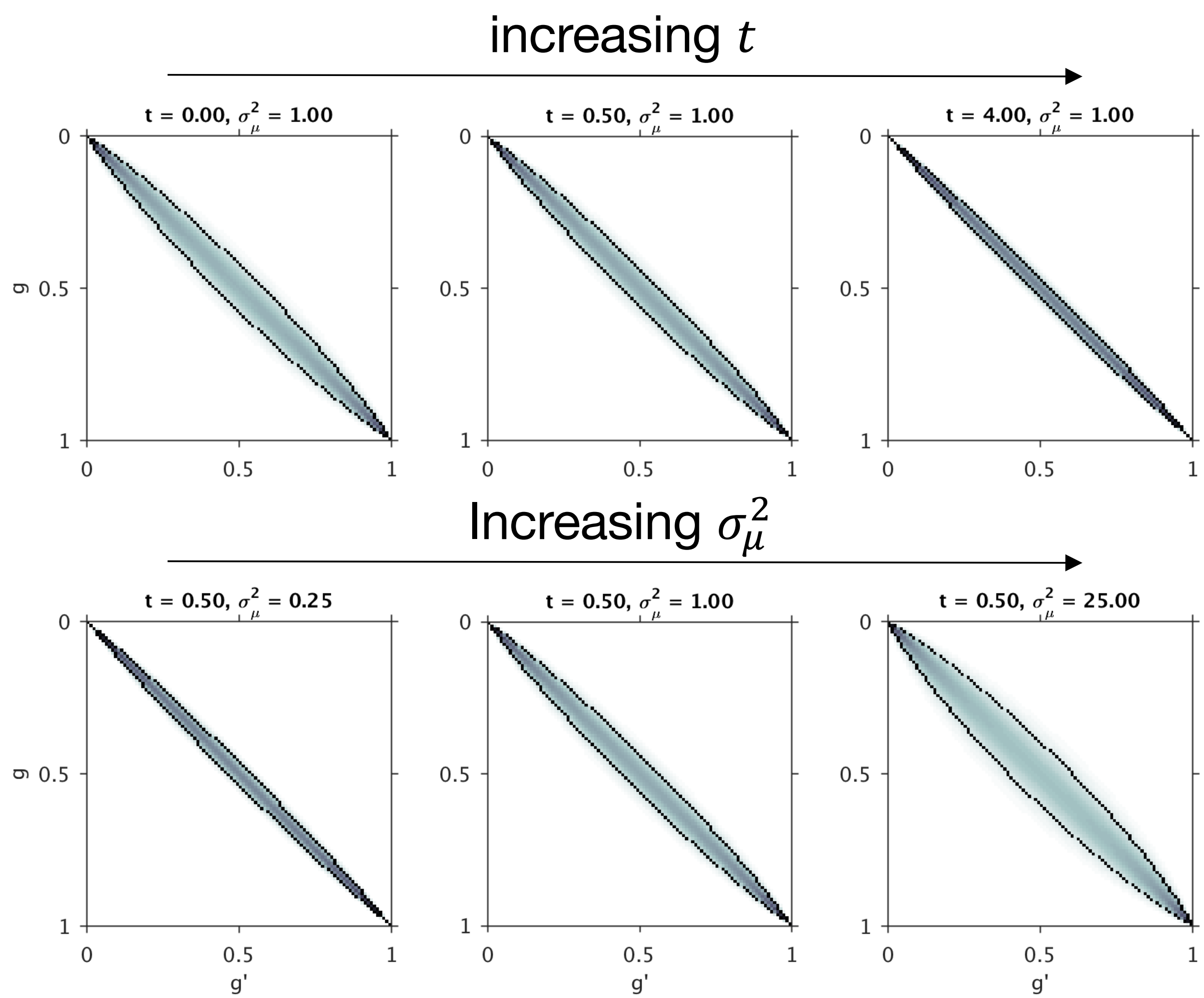

plot_g_trans_gauss_hyp.m

The value function and decision boundaries

$$V(g,t) = \max\Big\{\underbrace{g, 1-g}_{\text{deciding immediately}}, \underbrace{\langle V(g', t+\delta t)\rangle_{p(g'|g,t)} - c\delta t}_{\text{accumulating more evidence, and deciding later}}\Big\}$$

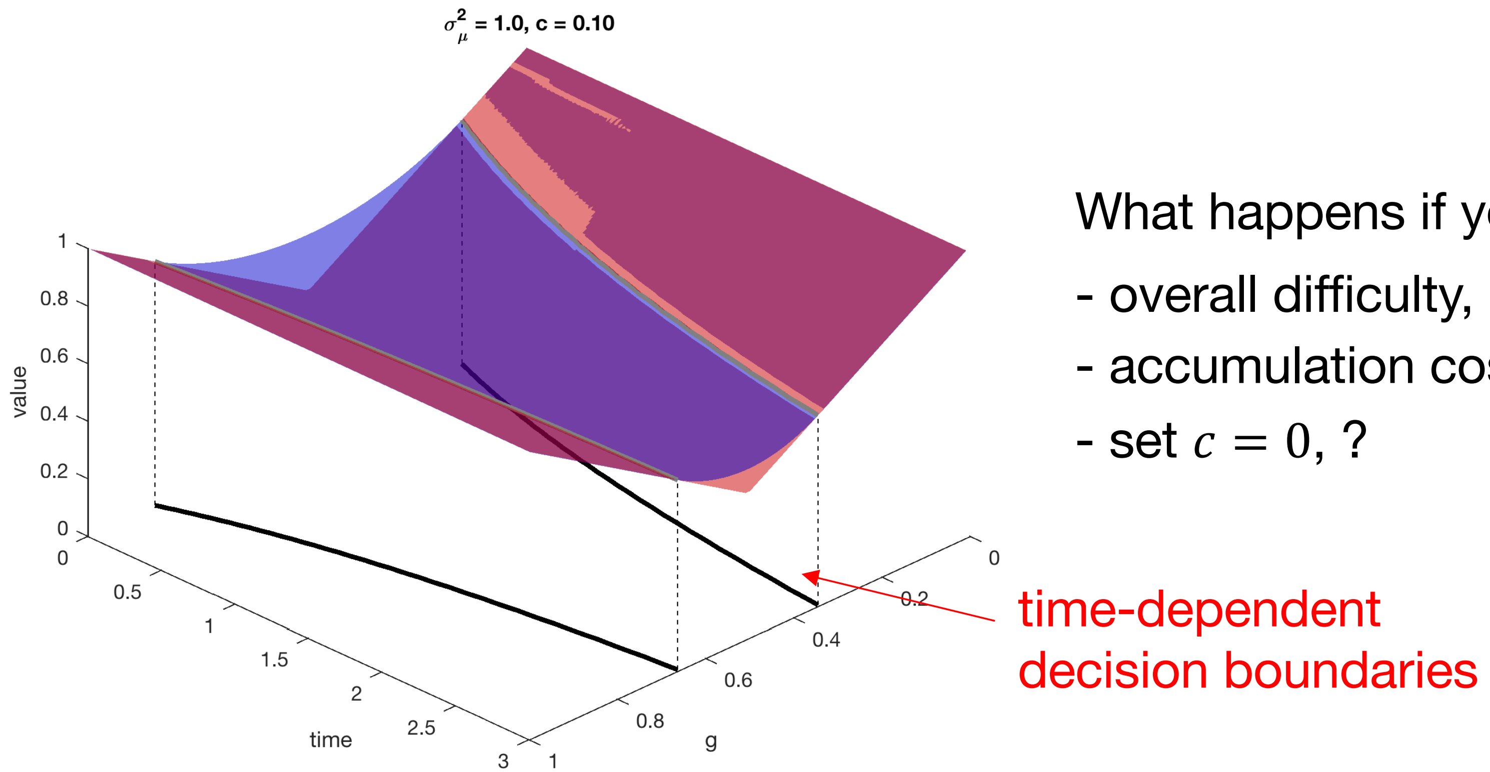

What happens if you change

- overall difficulty, σ_μ^2,
- accumulation cost, c,
- set $c = 0$, ?

plot_dp_valueintersect_gauss(σ_μ^2,c)

Diffusion models with time-dependent boundaries

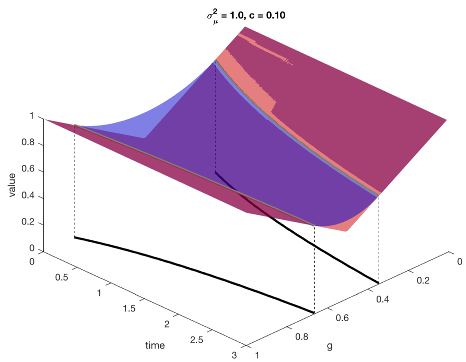

Consequences:

- SPRT is suboptimal
- No analytical RT/PC solutions
  → no direct ER optimization possible

diffusion $\frac{\mathrm{d}x}{\mathrm{d}t} = \mu + \eta(t)$

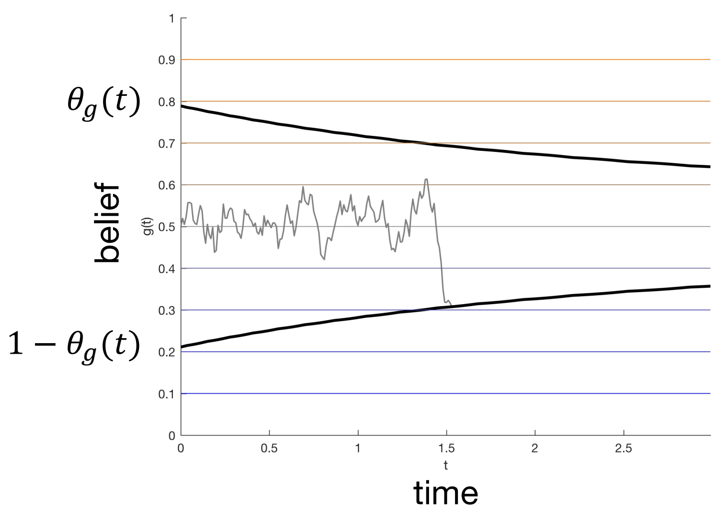

$$x(t) = \Phi^{-1}\big(g(t)\big)\sqrt{\sigma_\mu^{-2} + t}$$

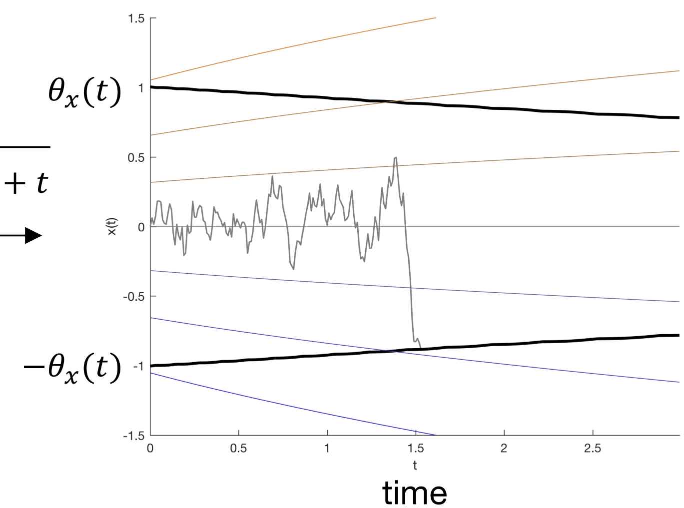

Unknown evidence reliability → collapsing boundary diffusion model optimal

Are DDMs with time-invariant bounds suboptimal?

diffusion model bounds

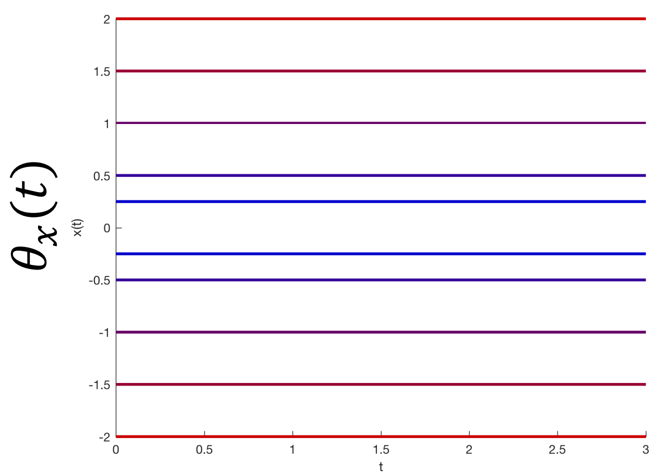

bounds in belief space

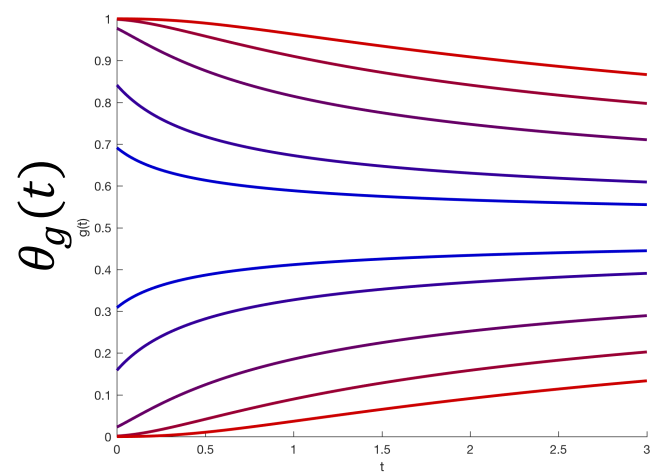

ddm_const_bound_gauss.m

$$\theta_g(t) = \Phi\left(\frac{\theta_x(t)}{\sqrt{\sigma_\mu^{-2} + t}}\right)$$

Constant diffusion model bounds implement collapsing bounds in belief
→ might be close-to-optimal (under certain circumstances)

Consequence of time-dependent boundaries

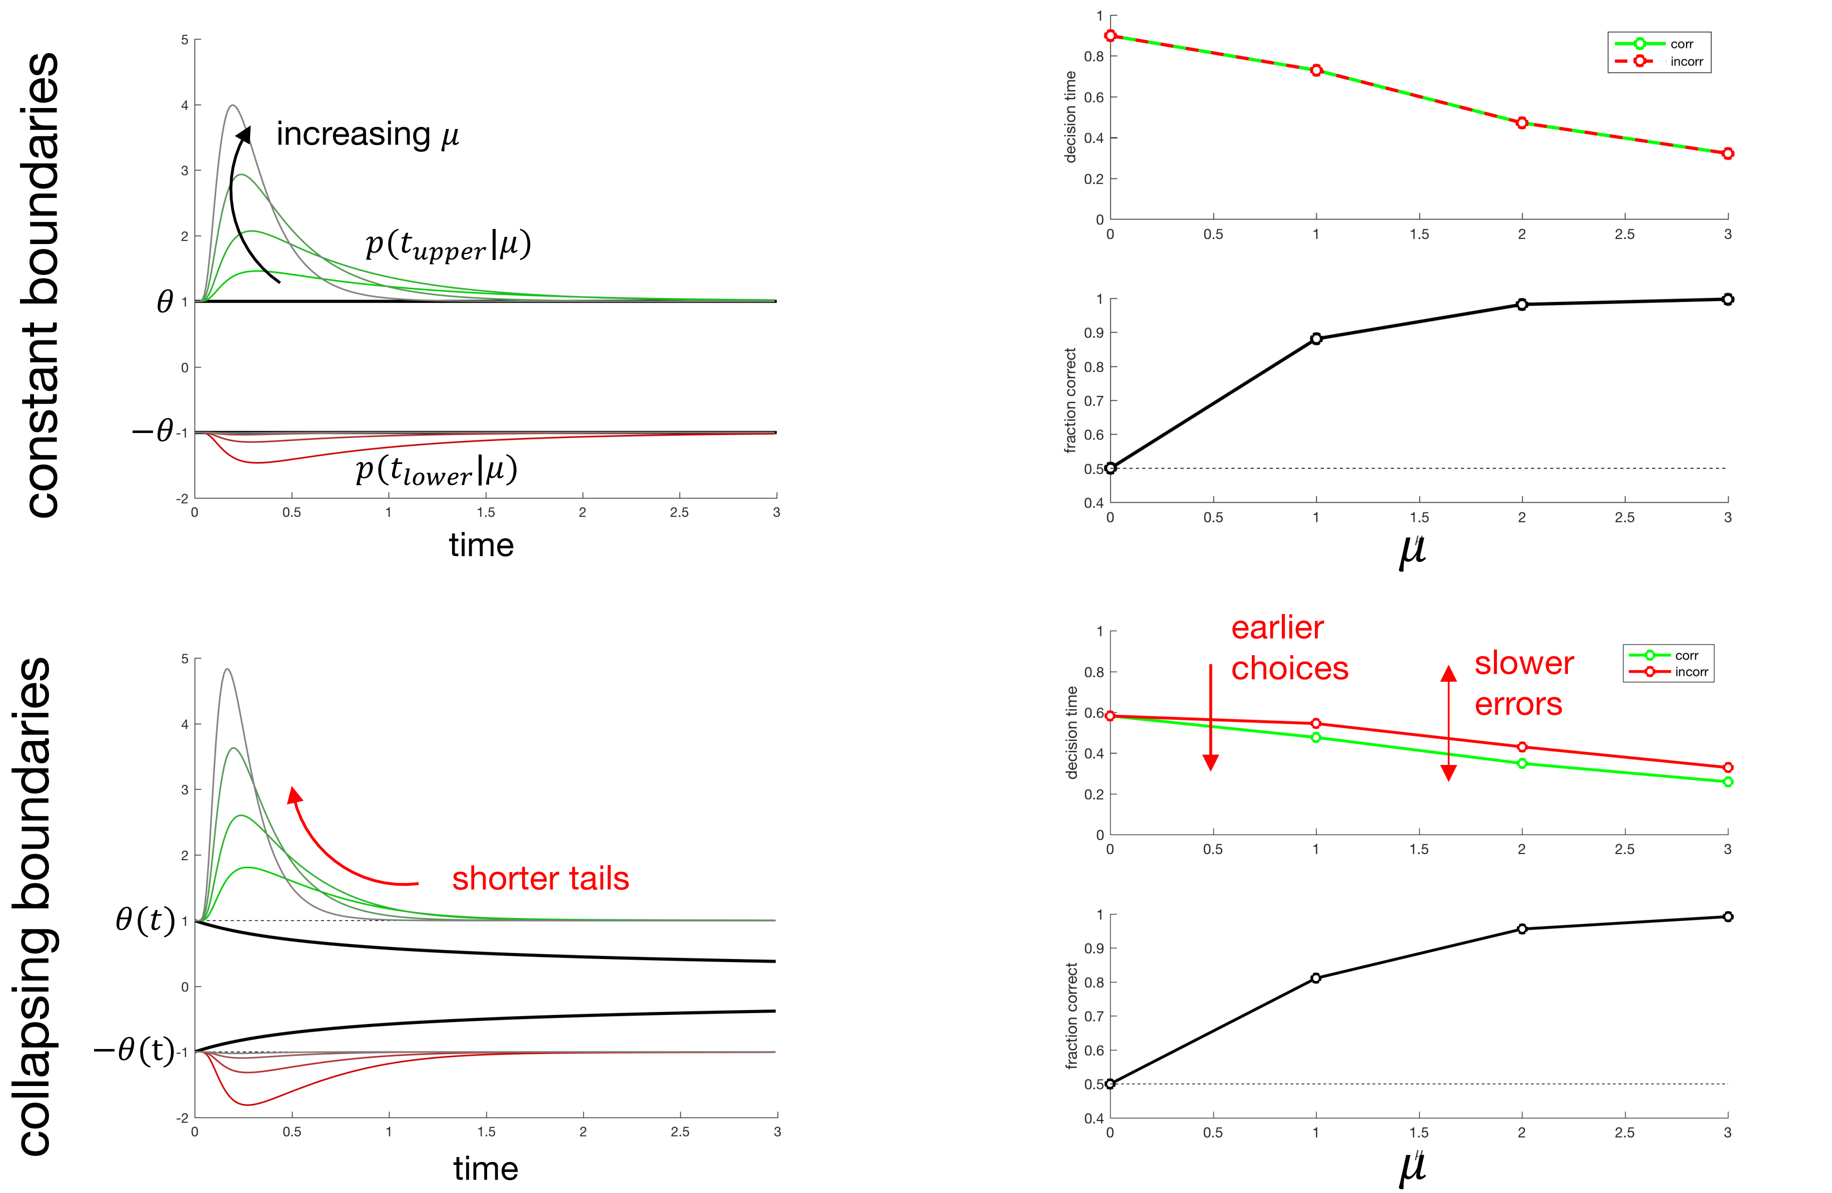

plot_fpt_vary_bound_example.m

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- → time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Evidence for bound collapse

Collapsing bound in belief → predicts dropping performance over time

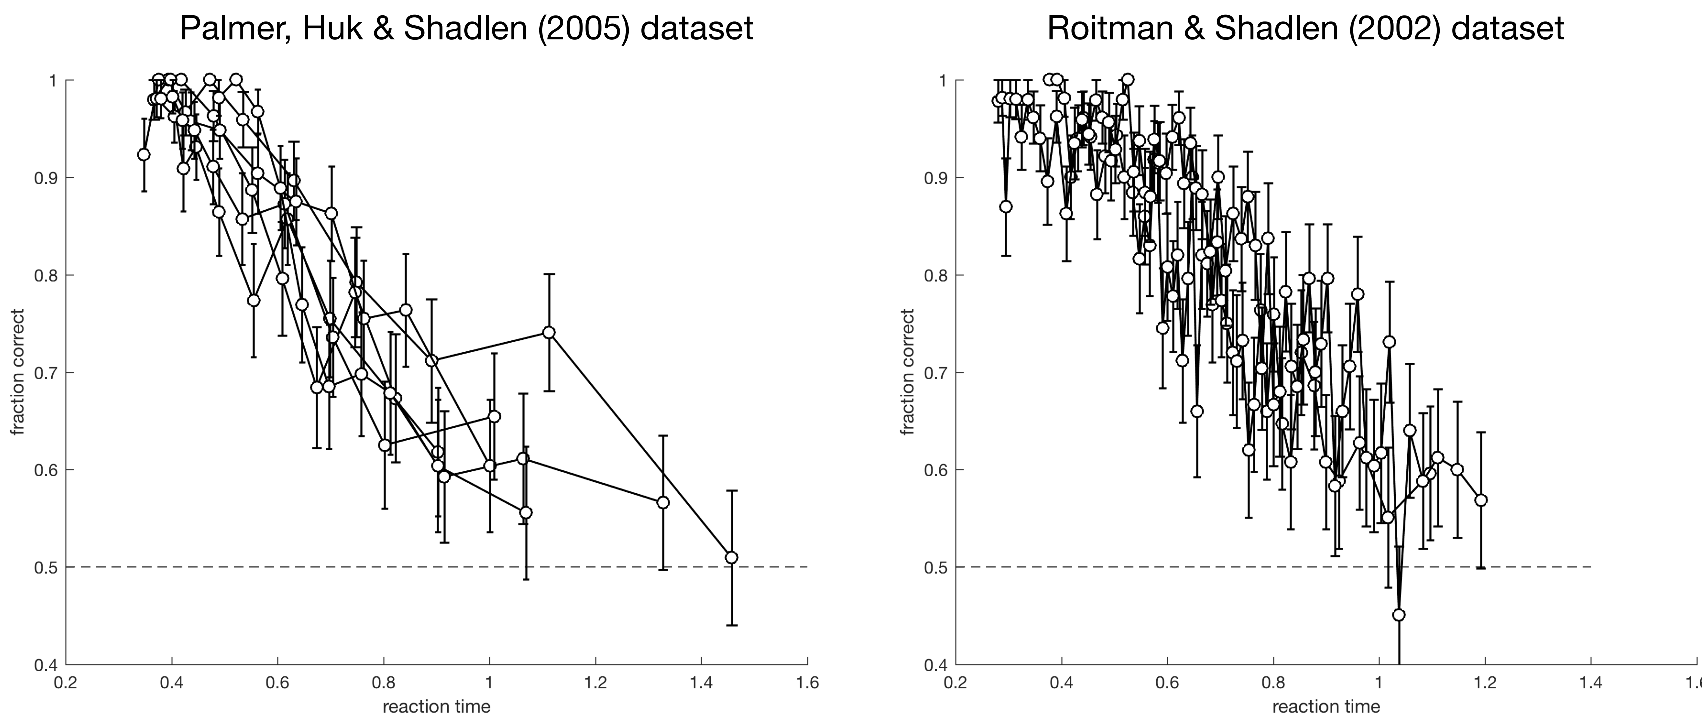

plot_pcorrect_over_time.m

In theory: we could reconstruct decision boundaries (in belief) from above plots

In practice: the non-decision time might be stochastic → prevents direct mapping

Are boundaries generally collapsing?

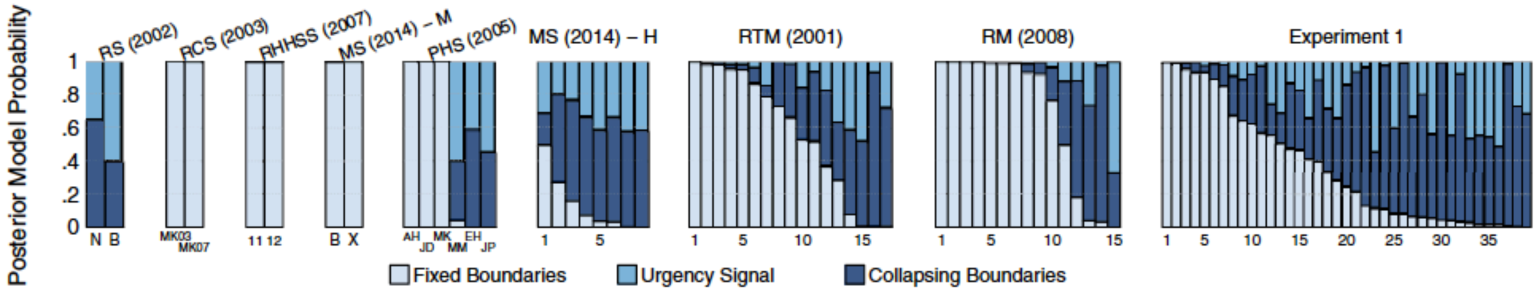

Figure 6. Approximations to posterior model probabilities in favor of the fixed bounds model with between trial variability parameters and the urgency signal and collapsing bounds models without between trial variability parameters. All details are as described for the top row of Figure 5.

(Hawkins et al., 2015)

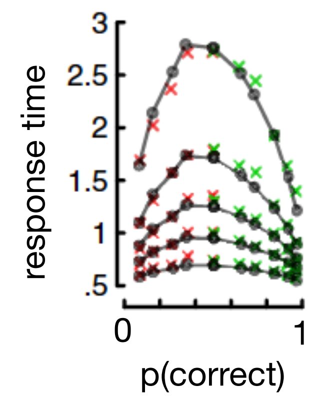

- collapse in particle space, not belief space
- fitting quantile plots, that might miss tail information (which are affected by bound collapse)
- does it matter?

How much do we gain from a collapsing boundary?

When do we expect such gains?

Hands on: benefit of collapsing boundaries

Aim: compare expected reward from optimal policy
and that arising from diffusion model with tuned constant boundary

Follow instructions in `collapse_gain.m`

Hints: Value function $V(g,t)$ returns expected reward when holding belief g at time t and behaving optimally thereafter.
$\rightarrow V\left(g=\frac{1}{2}, t=0\right)$ is expected reward for whole decision. See `plot_dp_diffusion_gauss.m` for how to find $V(g,t)$.

For given μ, we know probability correct and expected decision time for diffusion model with constant boundary. To compute expected reward, we can average these across multiple μ that well-represent $p(\mu) = N\left(\mu \middle| 0, \sigma_\mu^2\right)$. See `fixedbound_er(.)` in `collapse_gain.m`

Hands on: benefit of collapsing boundaries

Finding expected reward for optimal strategy:

```
gs = dp_discretized_g(dp_ng);
[~, Ve] = dp_getvalues_gauss_hyp(gs, dp_dt, dp_maxt, …
                                 sigmu2s(isigmu2), c);
opter_sigmu2(isigmu2) = Ve(1,ceil(dp_ng / 2));
```

Completing `fixedbound_er(.)` to return expected reward for fixed bound:

```
pcs = 1 ./ (1 + exp(-2 * theta * abs(mus)));
dts = theta ./ mus .* tanh(theta * mus);
dts(mus == 0) = theta^2;
er = mean(pcs) - c * mean(dts);
```

Finding bound height that maximizes expected reward:

```
[~,er] = fminsearch(…
    @(theta) -fixedbound_er(theta, cs(ic), sigmu2, fb_nmu),…
    1);
conster_c(ic) = -er;
```

Hands on: benefit of collapsing boundaries

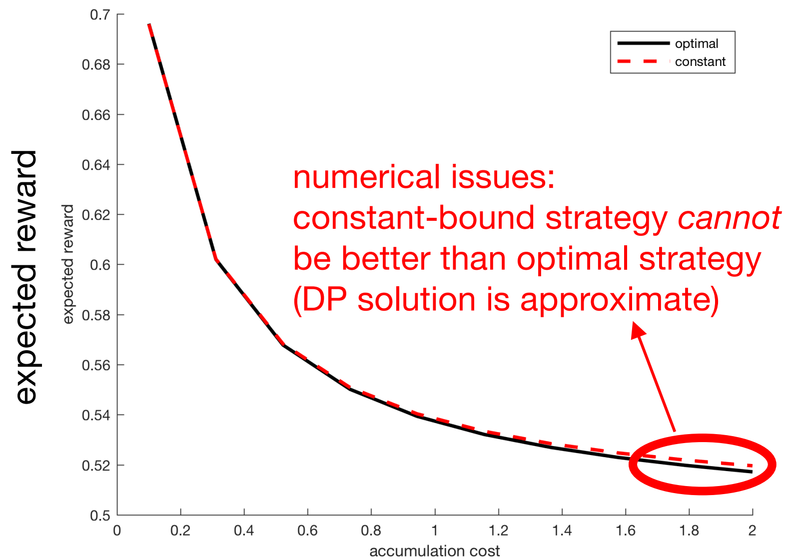

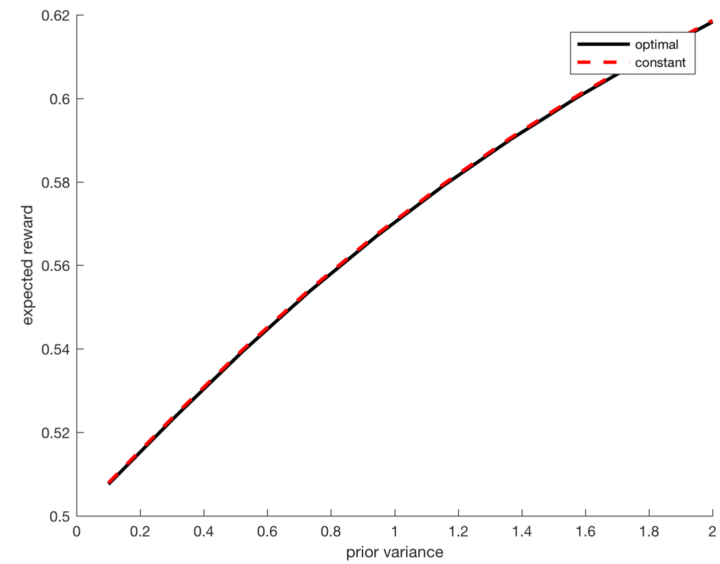

increasing c ⟶

increasing σ_μ^2 ⟶

For these scenarios, optimal solution barely better than constant boundary
(Recall: still collapsing boundary in belief)

Might change for stronger boundary collapse
e.g., accumulation cost that increases over time (e.g., Drugowitsch et al., 2012)

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

→ Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making

Neural signatures of perceptual decisions in macaque

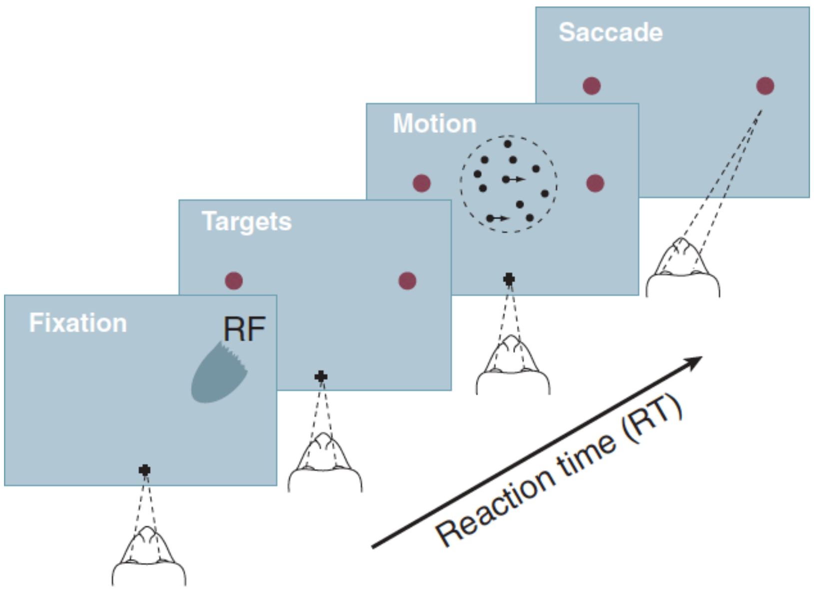

Gold & Shadlen (2007)

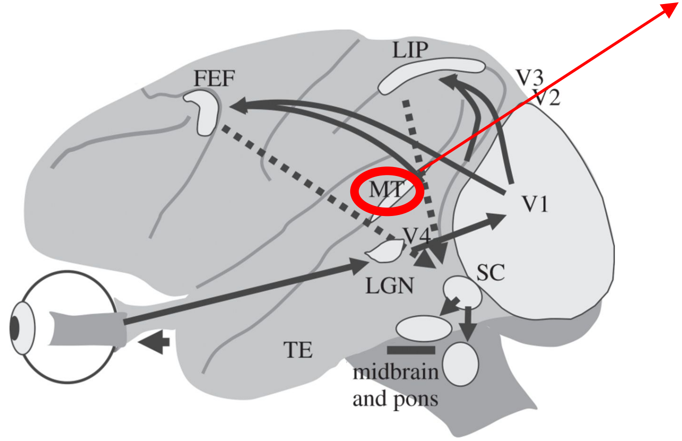

Wurtz (2015)

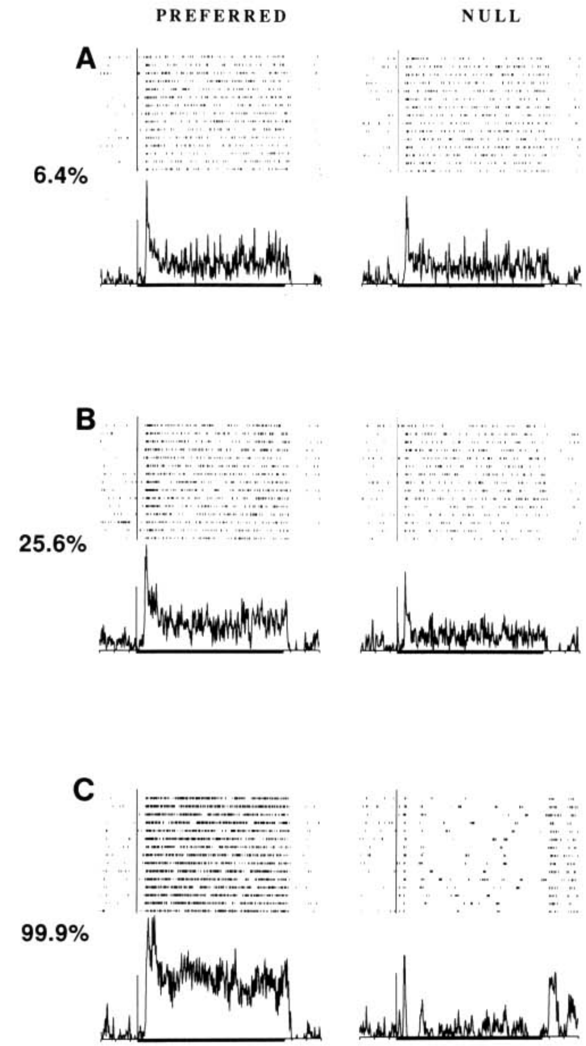

Britten et al. (1993)

Memory-guided saccade coding in macaque LIP

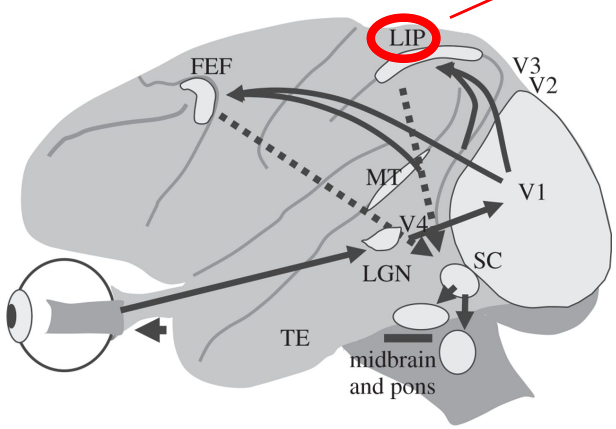

Wurtz (2015)

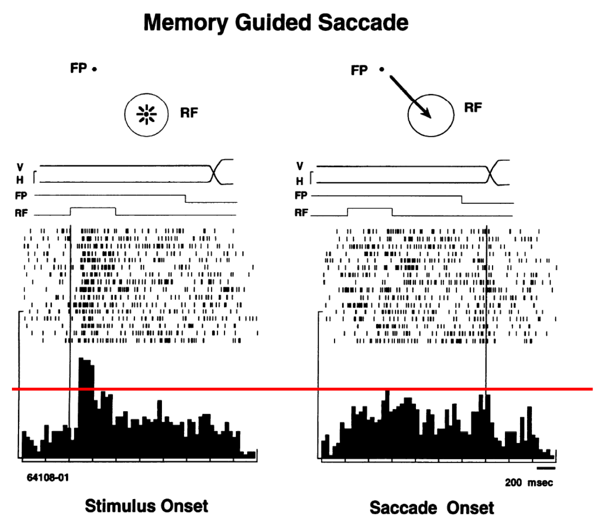

Colby, Duhamel & Goldberg (1996)

sustained activity
In memory-guided
saccades

Evidence accumulation coding in macaque LIP

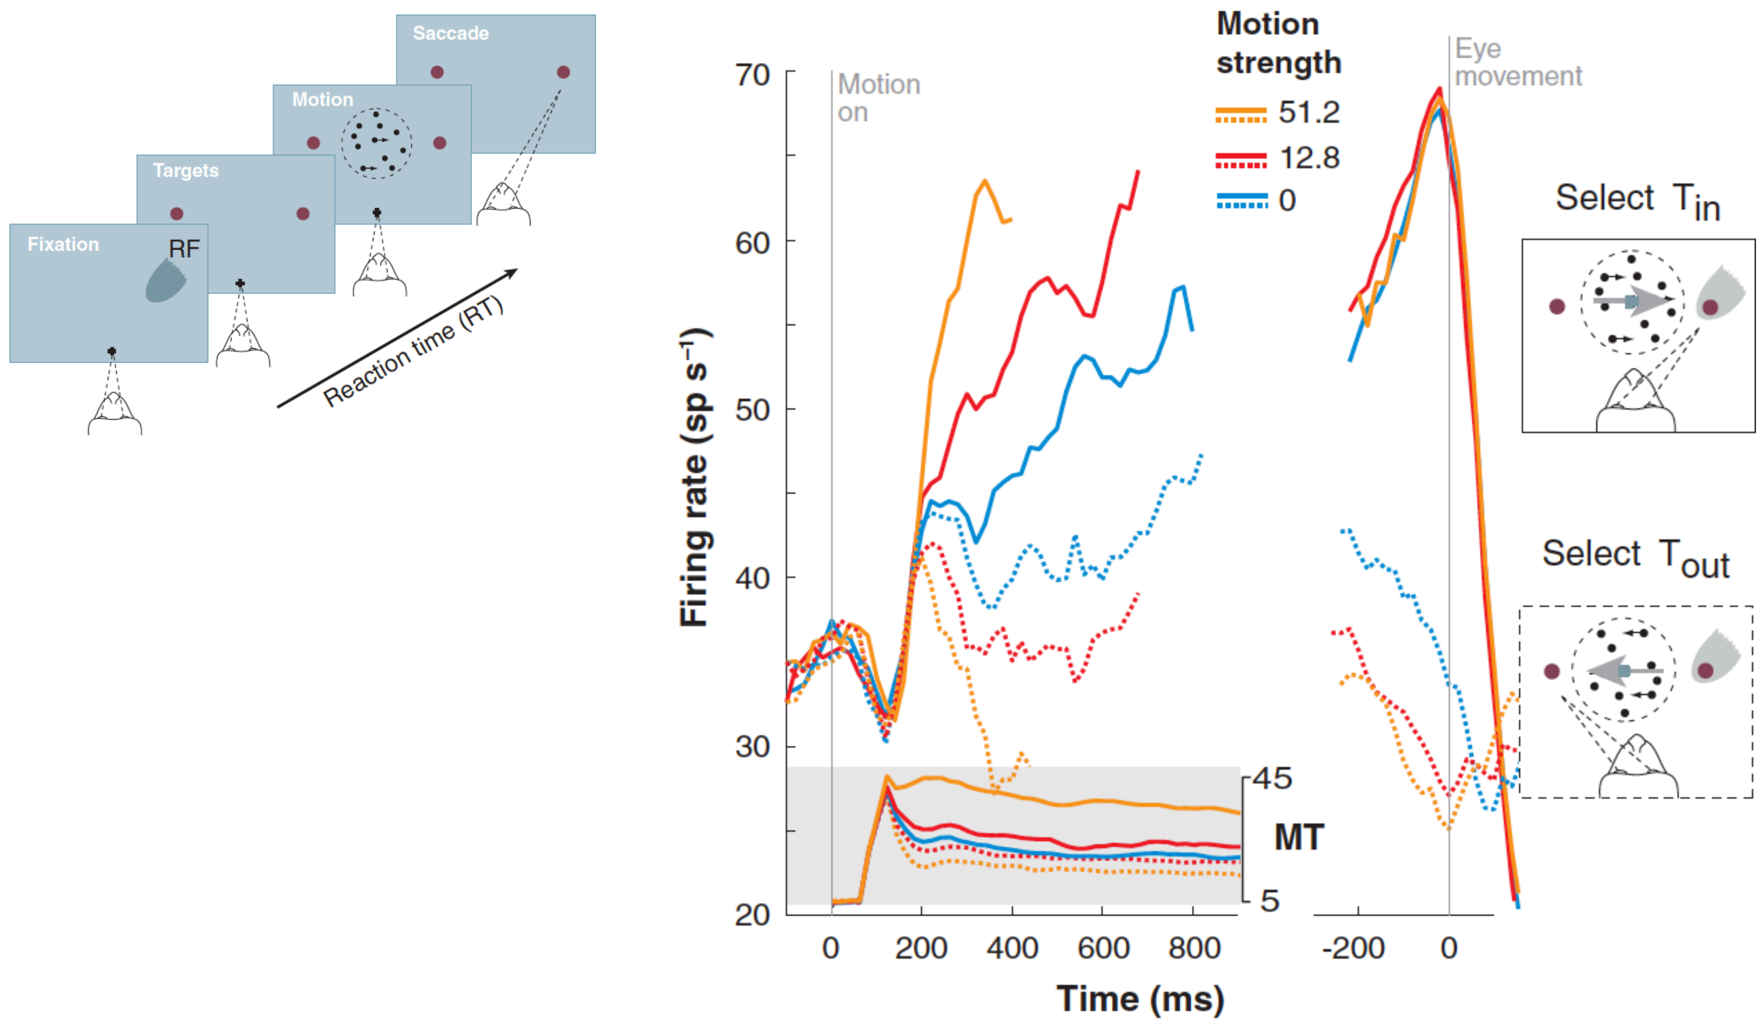

Gold & Shadlen (2007);
LIP data from Roitman & Shadlen (2002);
MT data from Britten (1992)

Does area LIP implement a diffusion model?

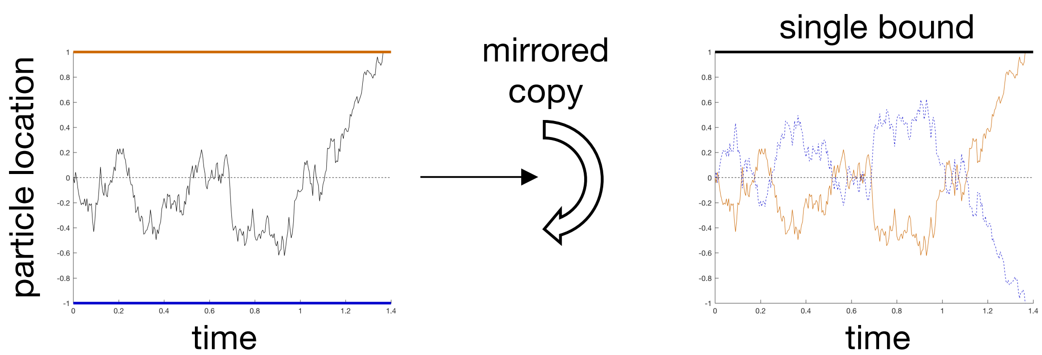

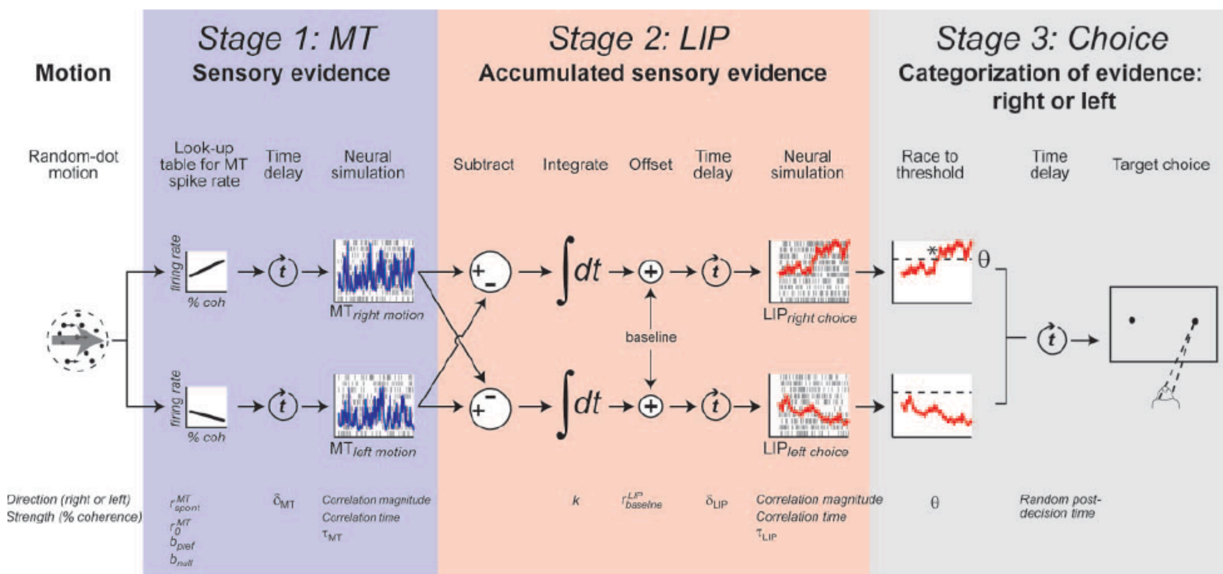

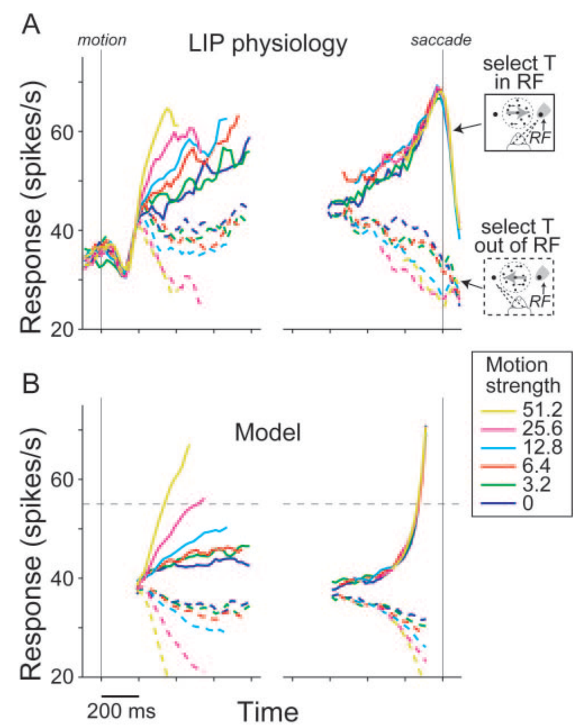

Mazurek et al. (2003)

Are LIP traces symmetric around common mean?

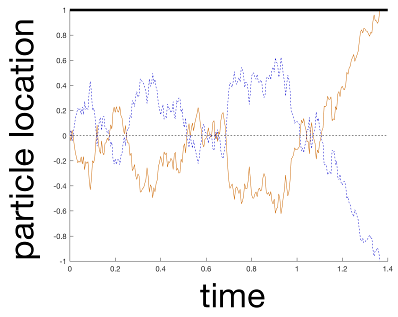

Plot $2b - r\_out(t)$:

```
avgact = mean(nanmean((m_mr1c(:,dot_ax>=200)+…
                      m_mr2c(:,dot_ax>=200))/2,2));
mirroredact = m_mr2c;
mirroredact(:,dot_ax >= 200) = …
    2*avgact - mirroredact(:,dot_ax >= 200);
plot(dot_ax, nanrunmean(mirroredact',1),'--','LineWidth',2);
```

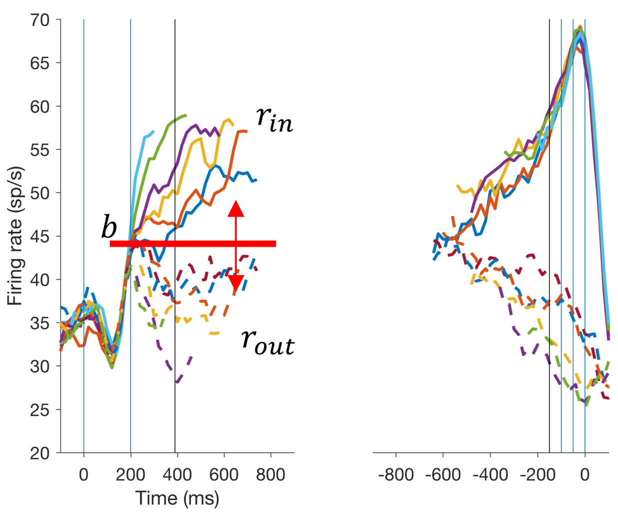

rs_datacode/lip_rt_roit_fig_7.m

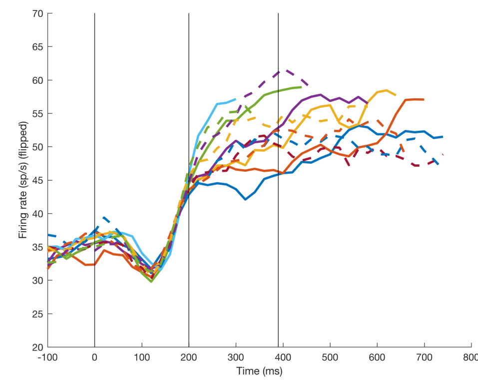

not fully symmetric

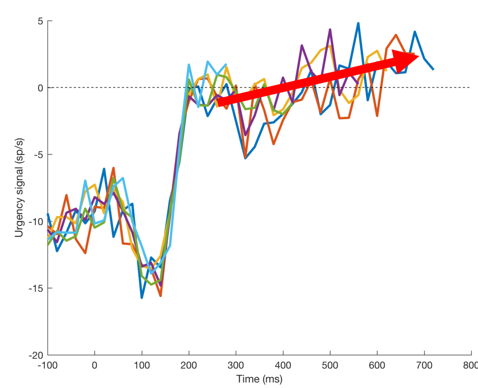

increasing "urgency" signal

$$\left.\begin{aligned} r_{in}(t) &= b + x(t) \\ r_{out}(t) &= b - x(t) \end{aligned}\right\} \; b \approx \frac{\langle r_{in}(t) + r_{out}(t)\rangle_t}{2}$$

(for $t > 200ms$)

Urgency signal implements collapsing boundary

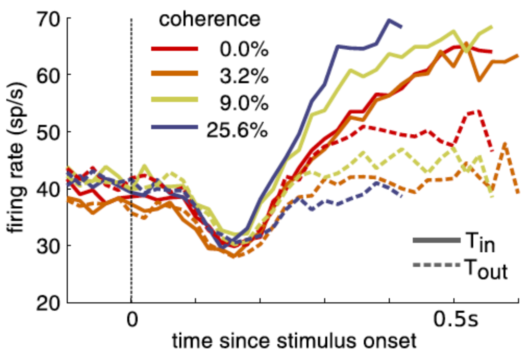

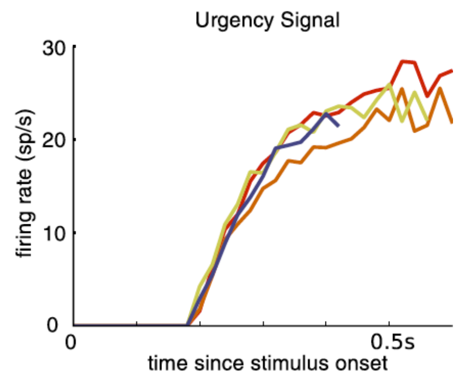

Data from Churchland et al. (2008)

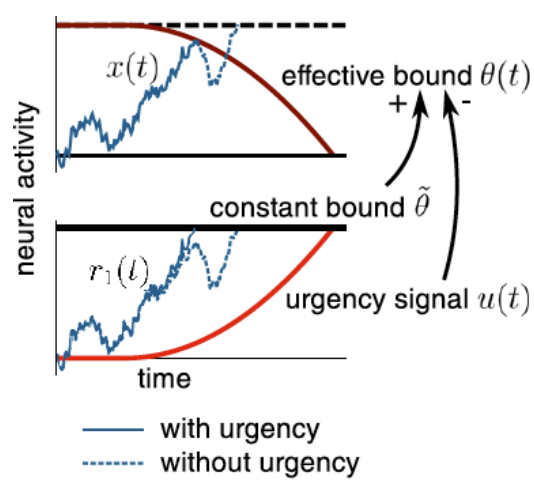

Drugowitsch et al. (2012)

Neural evidence accumulation signatures in rodents

Rat click count discrimination task

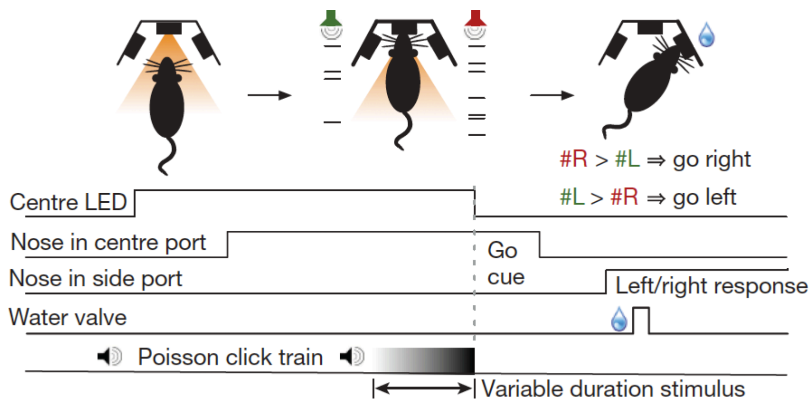

→ accumulate click difference

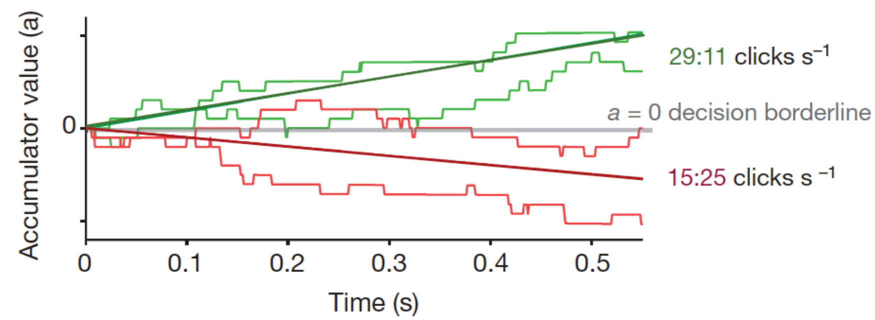

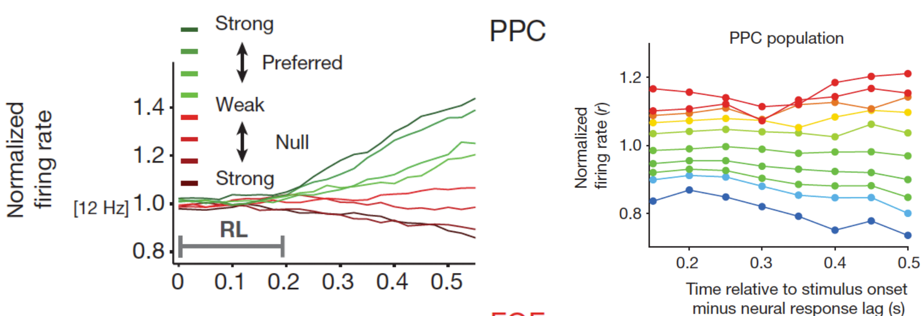

seems to reflect accumulation

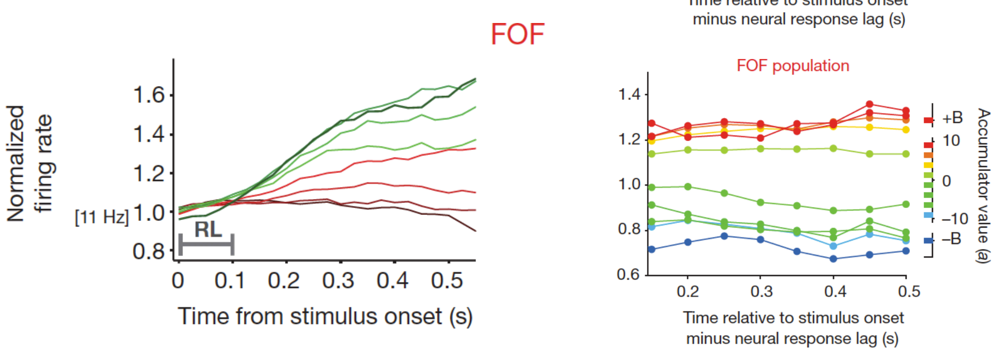

seems to reflect decision

Hanks et al. (2015)

But: inactivation studies

Rodents:

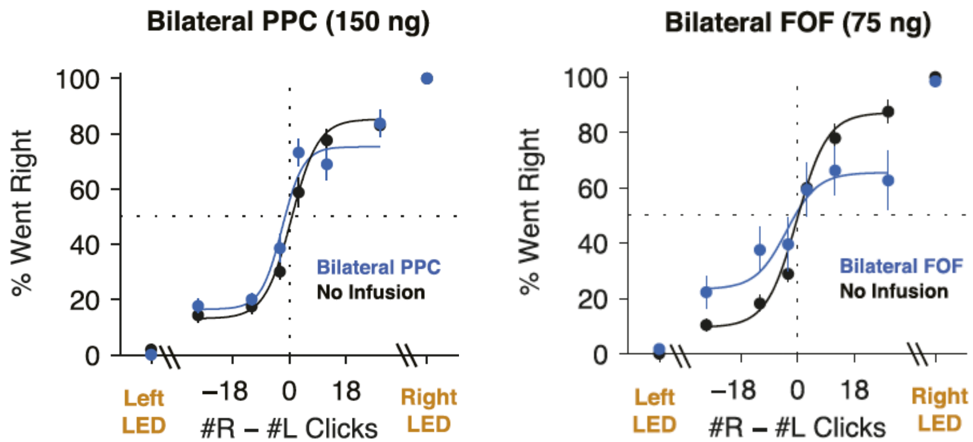

Ehrlich et al. (2015)

momentary evidence
↓
PPC
↓
FOF
↓
choice

Monkeys:

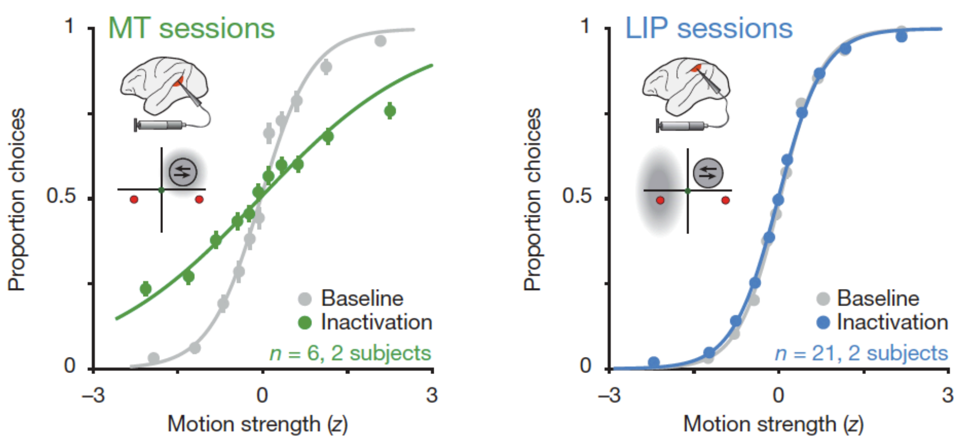

Katz et al. (2016)

momentary evidence
↓
MT
↓
LIP
↓
choice

Also: not everything that accumulates, ramps

Rodent VR cue accumulation task

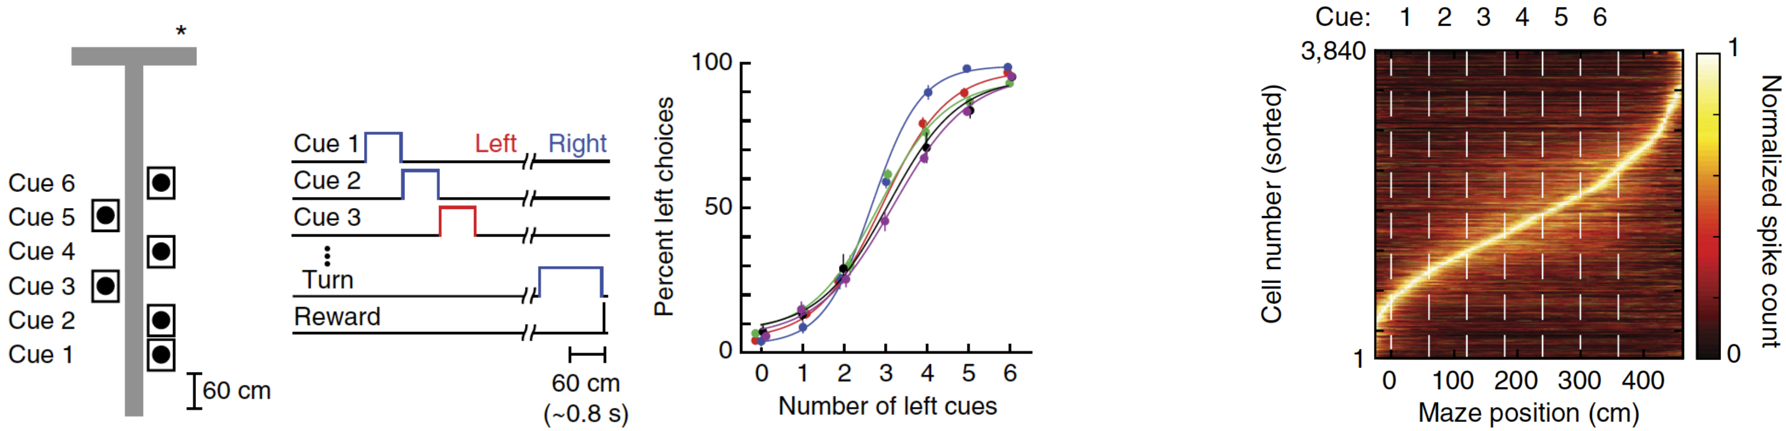

Morcos & Harvey (2016)

This does not invalidate normative approach!

Neural implementation is less clear
(there are multiple ways to implement evidence accumulation)

Road map

Perceptual decision-making
- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models
- accumulator / diffusion models
- fit to behavior & issues

Normative analysis
- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

→ Extended tutorial: multi-model decision-making

Bayesian cue combination

Frequently, evidence from multiple cues needs to be combined

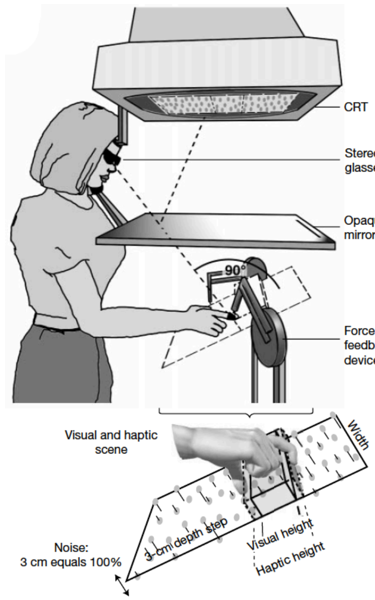

Ernst & Banks (2002)

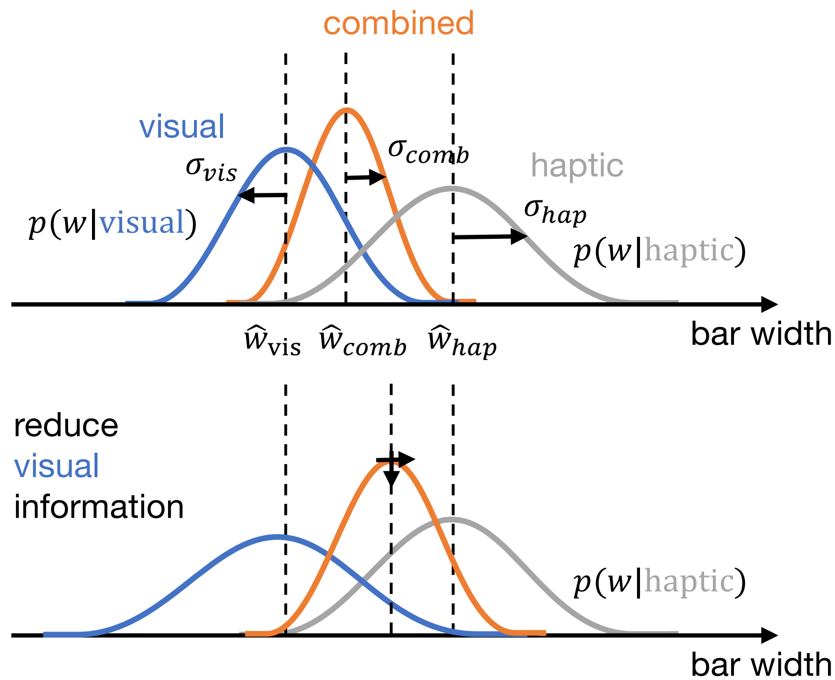

Bayesian cue integration:

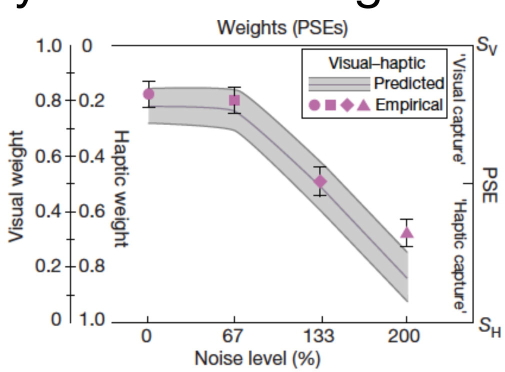

More reliable cue contributes more strongly

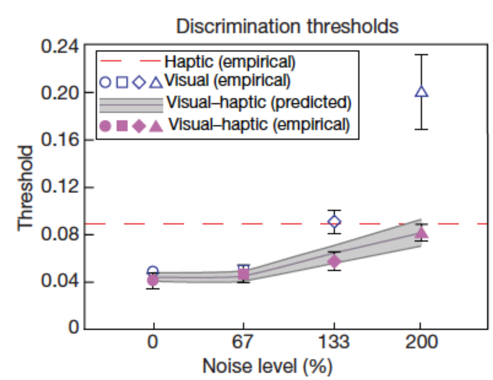

Combined reliability > individual reliability $\frac{1}{\sigma_{comb}^2} = \frac{1}{\sigma_{vis}^2} + \frac{1}{\sigma_{hap}^2}$

The speed/accuracy trade-off in mutlisensory decision-making

Standard cue combination paradigm is fixed-duration

- Ignores temporal evidence accumulation
- Frequently, decision time is under the decision-maker's control

A cue-combination reaction time task (Drugowitsch et al., 2014)

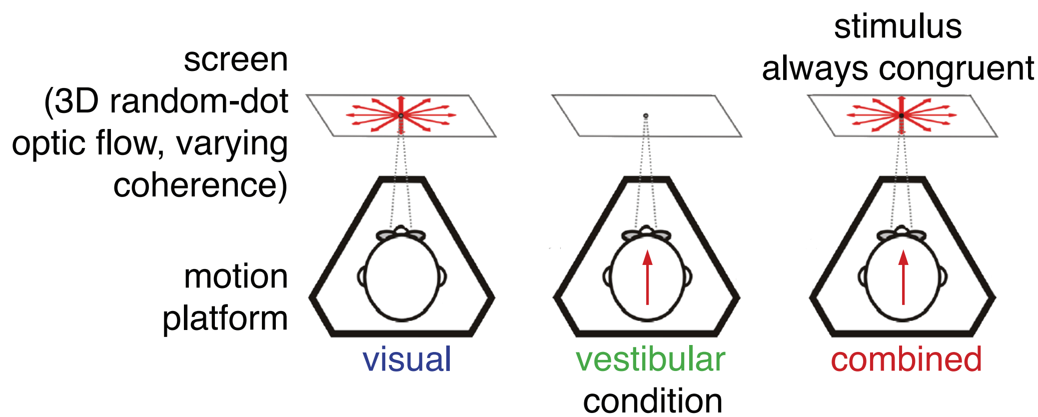

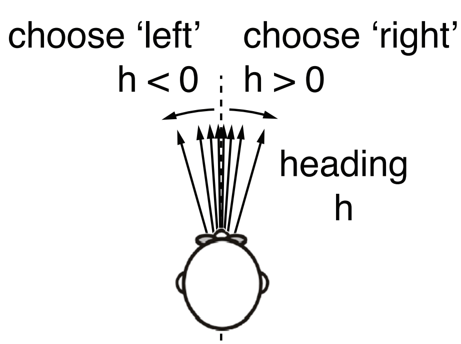

heading discrimination task

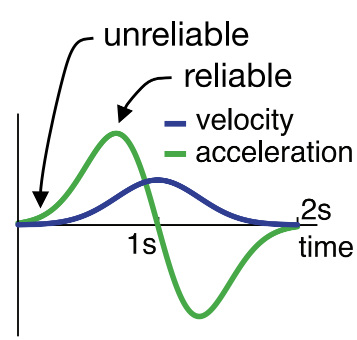

varying reliability time-course

Visual stimulus example

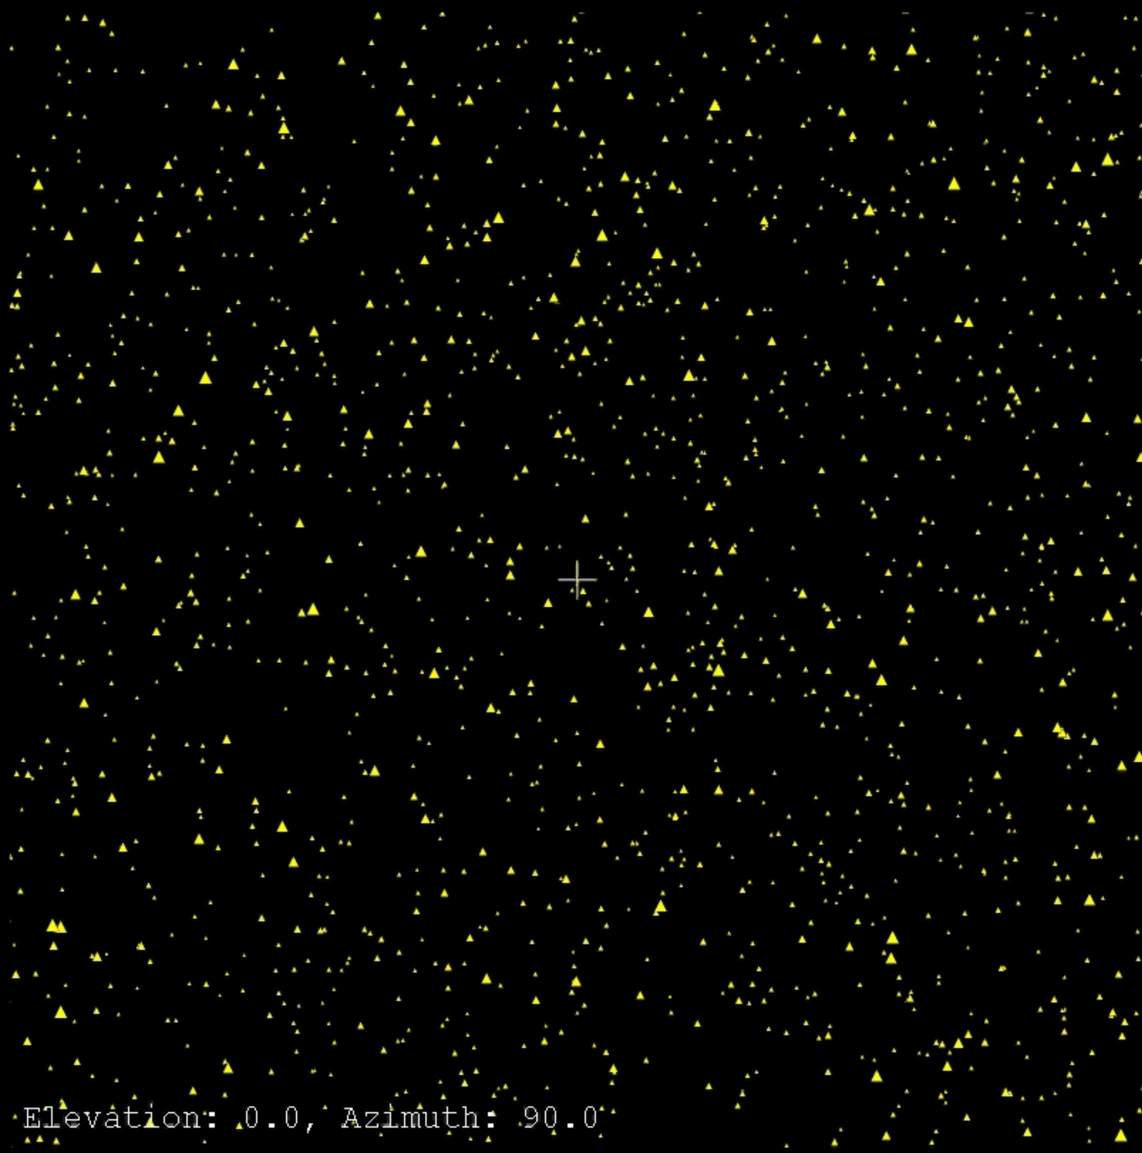

Visual reliability modulated by coherence

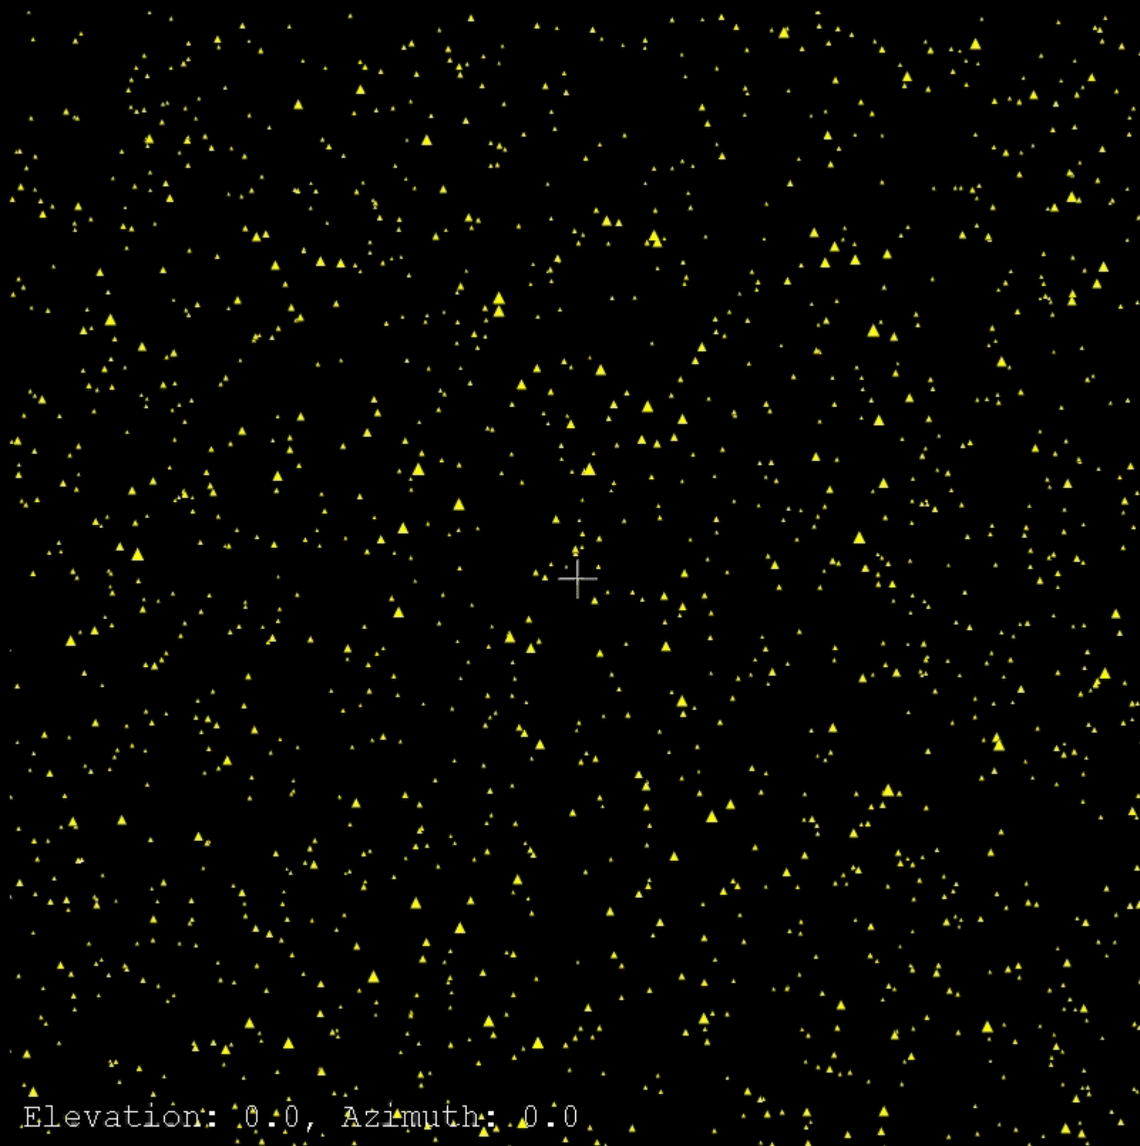

Evidence reliability modulated by four factors

heading direction (angle away from straight-ahead)

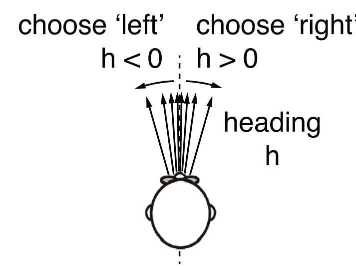

visual flow field coherence

velocity/acceleration time-course

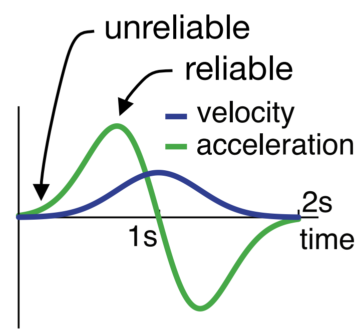

presence of multiple modalities

The vis/vest cue combination dataset

See content of `vis_vest` folder:

`vis_vest_[x].mat`:	per-trial data for single subject [x]
`vis_vest_README.txt`:	details of data format

A trial was characterized by

`oris`:	heading direction (+ve: right; -ve: left)
`mod`:	modalities present (vis/vest/comb)
`cohs`:	visual coherence, $\in \{0.25, 0.37, 0.70\}$

The subject's response consisted of

`choice`:	0 - ”left”; 1 – “right”
`rt`:	reaction time in [s], stimulus onset to choice

Further documents:

`vis_vest_tutorial.pdf`:	detailed instructions, derivations, some solutions (if you get stuck)
`Drugowitsch2014.pdf`:	paper that used this dataset

What you should do

Look at `vis_vest_tutorial.pdf`

- Become familiar with the data and behavior
- Perform standard Bayesian cue combination analysis
- Derive Bayes-optimal evidence accumulation & simulate
    - Single cue, evidence reliability that changes over time
    - Multiple cues, constant evidence reliability
    - Bonus: combination of both
- Simulate behavior in a virtual experiment & try to match human data
- Bonus: refine simulations
- Bonus: derive optimal decision boundaries

Good luck!

Behavior

vestibular
visual
combined

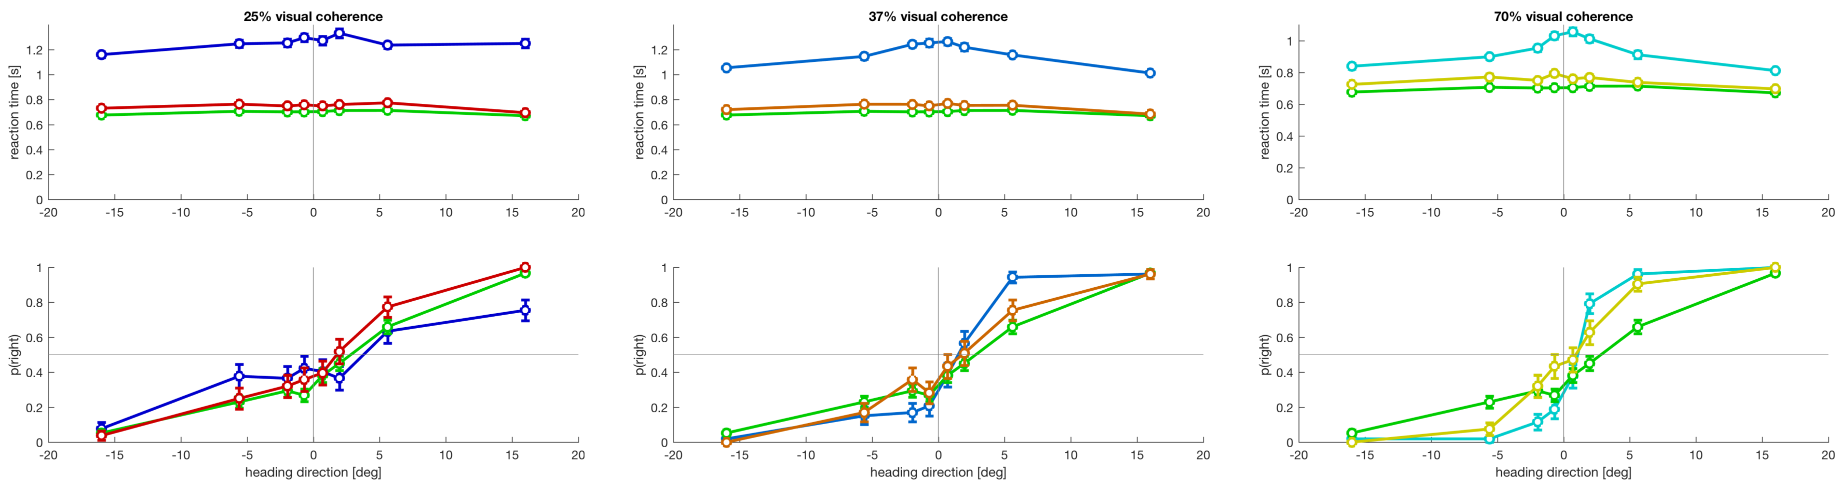

plot_psych_chron(.)

increasing coherence

drop in reaction times

increase in correct choices

Standard cue combination test

Estimating thresholds σ^2 by fitting cumulative Gaussians

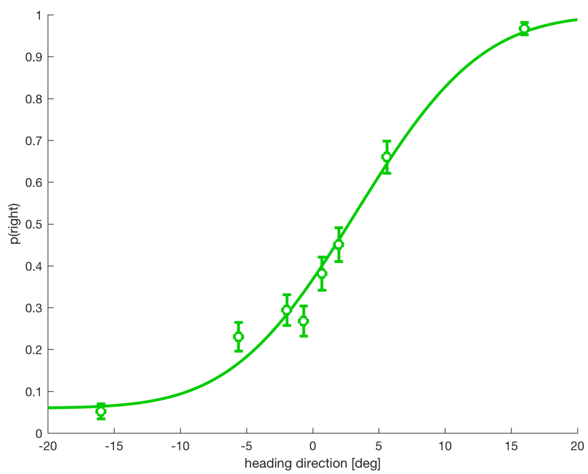

`test_fit_cumul_gauss(.)`

Complete `test_standard_cue_comb(.)`

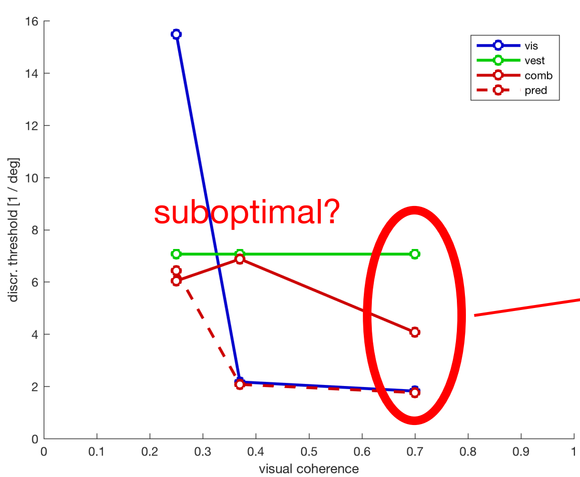

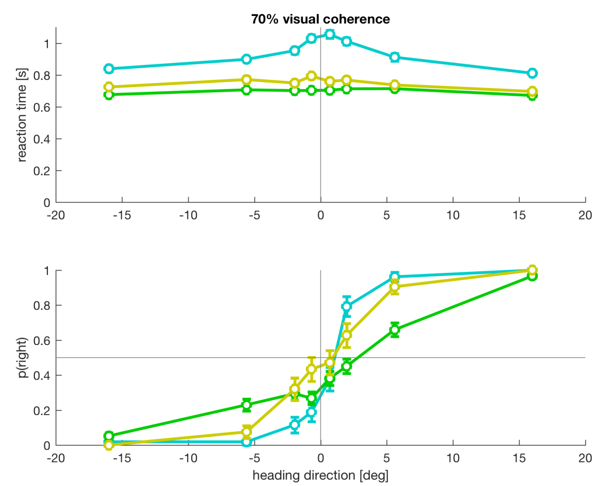

But, comb is faster than vis!

Deriving optimal evidence accumulation

Momentary evidence likelihood (visual modality)

$$p(\delta x_n | z(h), c) = \mathrm{N}(\delta x_n | z(h) v_n k(c) \delta t, \delta t)$$

information in heading direction, $z(h) = \sin(h)$

time-dependent reliability (velocity, $v_n = v(n\delta t)$)

global reliability (depends on coherence)

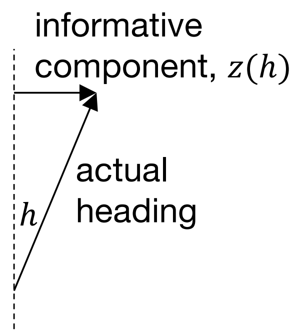

Find posterior $z(h)$ given some momentary evidence $\delta x_1, \dots, \delta x_n$

$$p(z(h)|\delta x_1, \dots, \delta x_n) \propto \prod_{j=1}^{n} p\left(\delta x_j \middle| z(h), c\right) \text{ with } \quad x_v(t) = \sum_{j=1}^{n} v_j \delta x_j \qquad V(t) = \sum_{j=1}^{n} v_n^2$$

Find posterior belief of right-ward motion,

$$p(z(h) \geq 0 | x_v(t), t) \int_0^\infty p(z(h)|x_v(t), t) \mathrm{d}z(h)$$

(use $\int_0^\infty \mathrm{N}(x|a,b)\mathrm{d}x = \Phi\left(\frac{a}{\sqrt{b}}\right)$)

Simulate weighted evidence accumulation

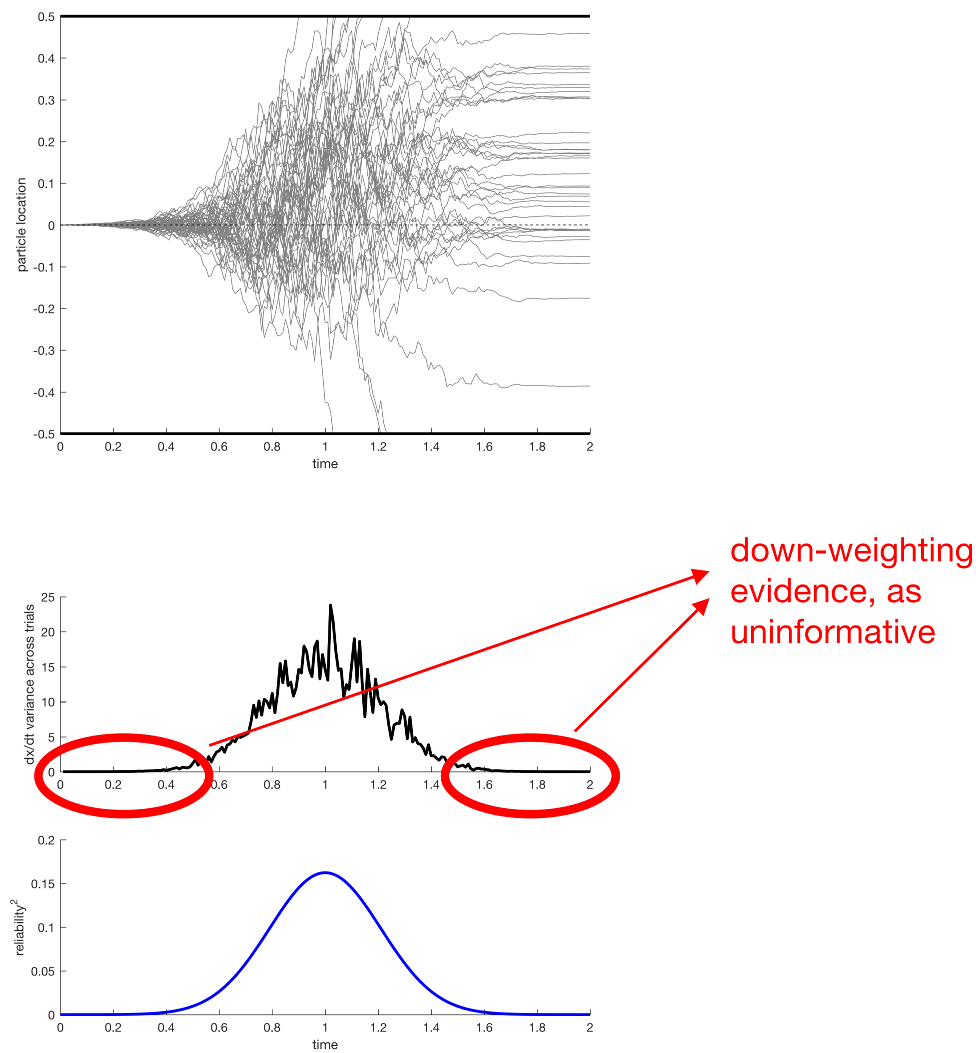

sim_weighted_diffusion.m

Simulating behavior

sim_behavior.m

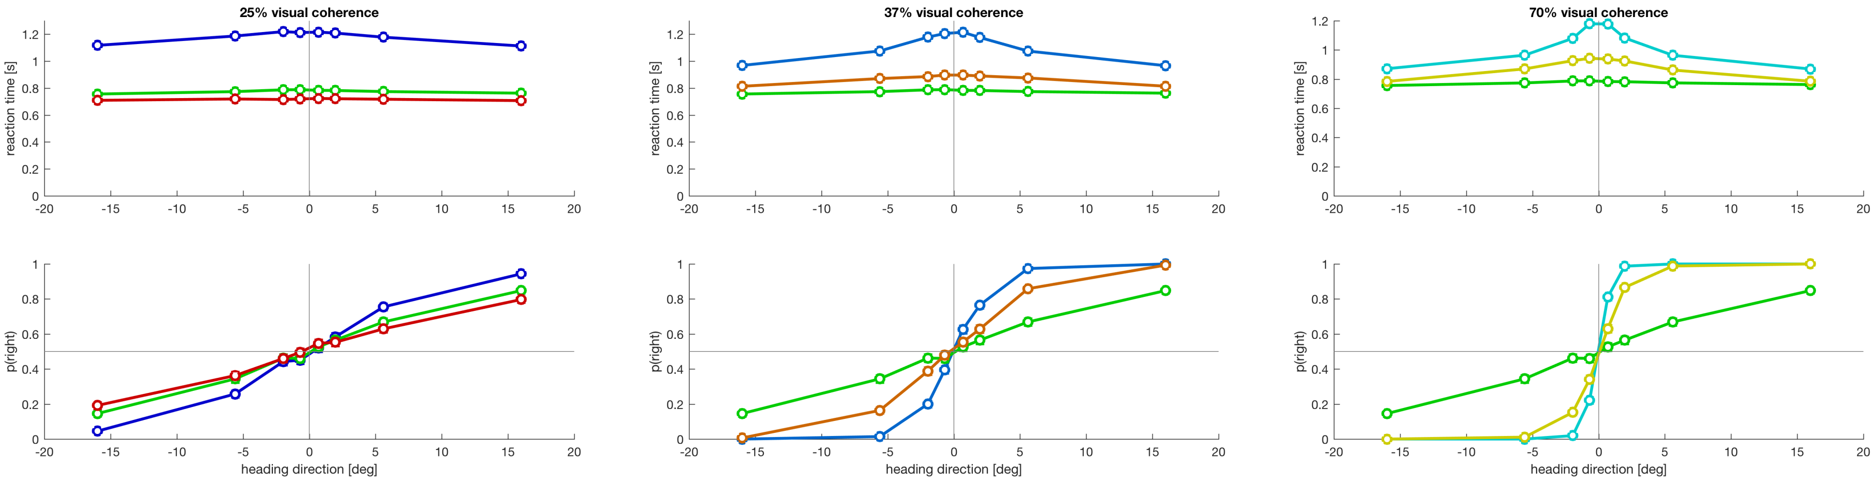

Actual behavior

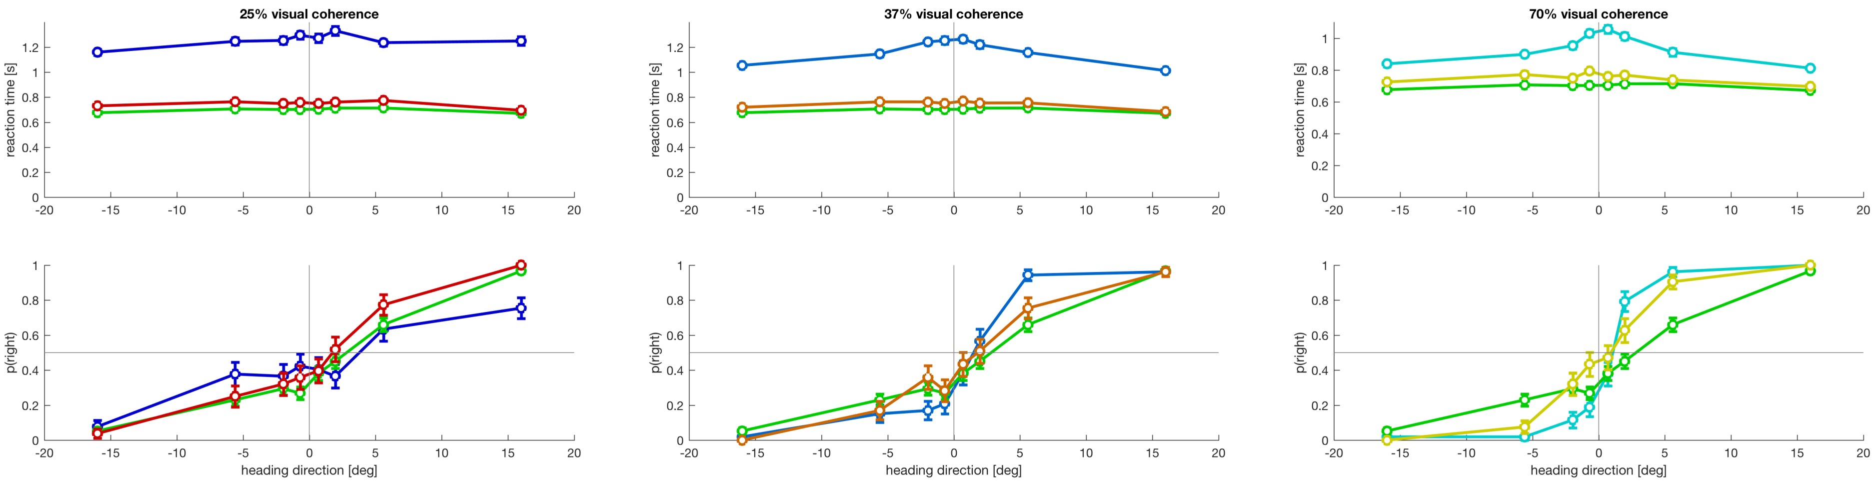

Road map

Perceptual decision-making

- speed/accuracy trade-off
- experiments investigating perceptual decisions
- characteristics of behavior

Decision-making models

- accumulator / diffusion models
- fit to behavior & issues

Normative analysis

- simple scenario: task difficulty known
- more complex: varying task difficulty
- time-varying decision boundaries: behavioral evidence

Neural correlates of perceptual decisions

Extended tutorial: multi-model decision-making