A Rational Decision-Making Framework for Inhibitory Control (Poster T7)

Pradeep Shenoy *Rajesh Rao* *Angela J. Yu*

University of California, San Diego

University of Washington

What is Inhibitory Control?

What is Inhibitory Control?

⇒

⇒

What is Inhibitory Control?

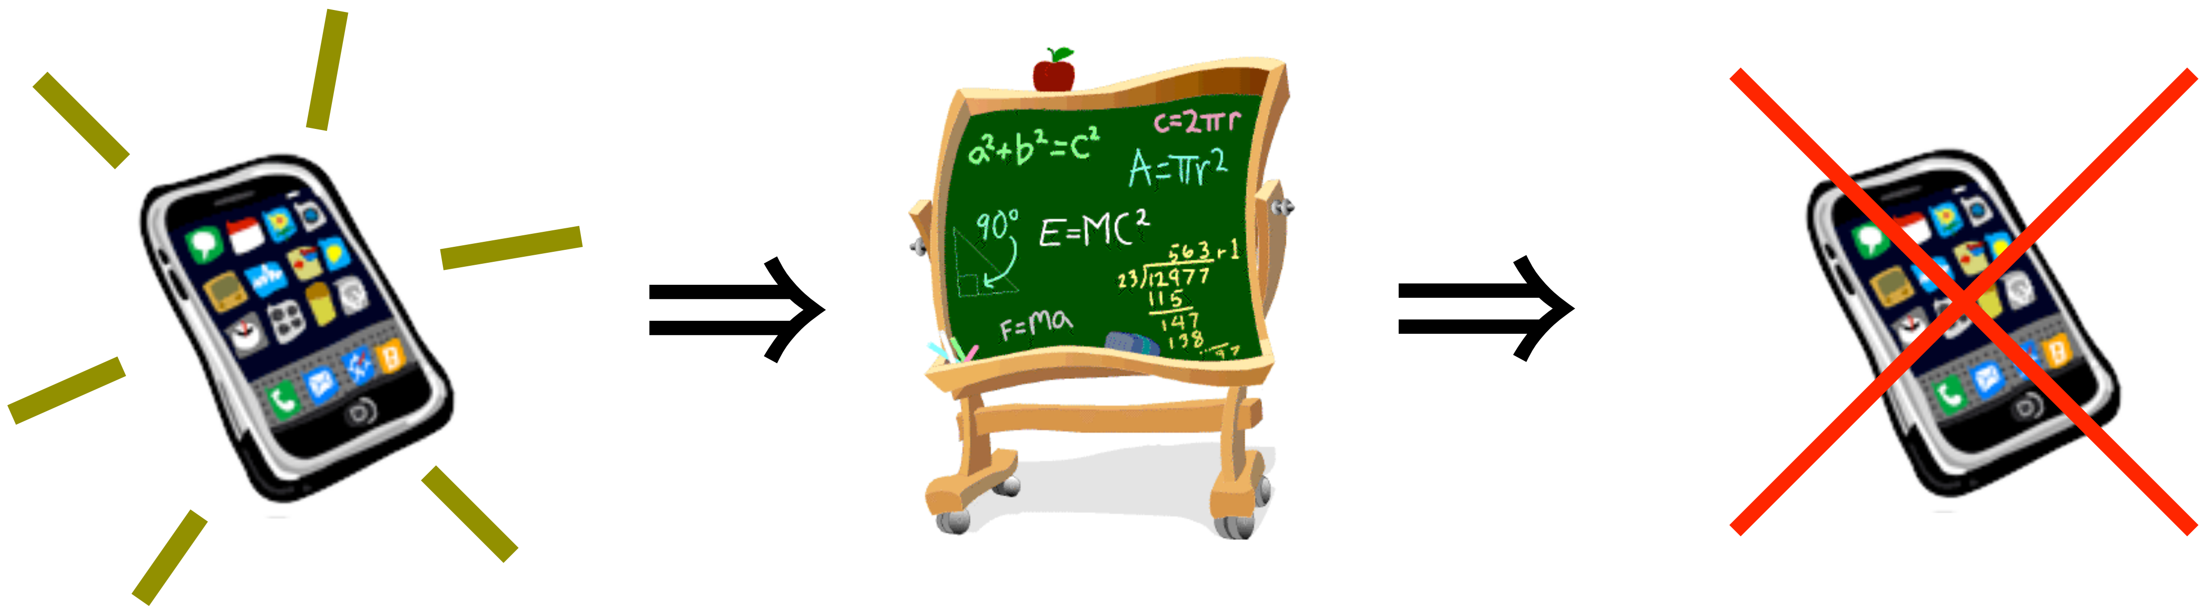

What is Inhibitory Control?

⇒

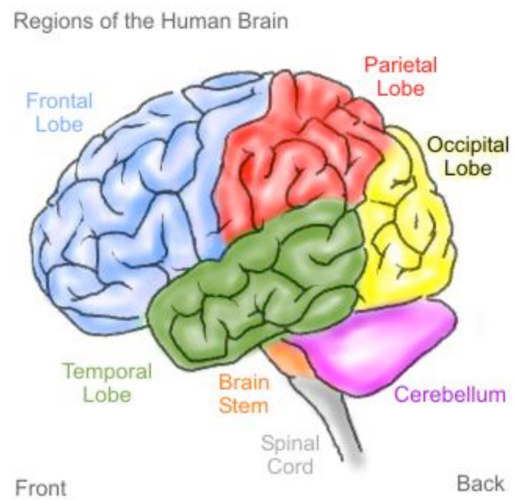

⇒

What is Inhibitory Control?

Inhibitory Control

the ability to inhibit a prepotent response or modify a planned course of action according to behavioral context

Inhibitory Deficits ⇔ Patient Populations

Inhibitory Deficits ⇔ Patient Populations

- **Psychiatric conditions**
  - drug abuse (Nigg et al., 2006)
  - attention-deficit hyperactivity disorder (Alderston et al., 2007)
  - obsessive compulsive disorder (Menzies et al., 2007)
  - schizophrenia (Enticott et al., 2008)

Inhibitory Deficits ⇔ Patient Populations

- **Psychiatric conditions**
  - drug abuse (Nigg et al., 2006)
  - attention-deficit hyperactivity disorder (Alderston et al., 2007)
  - obsessive compulsive disorder (Menzies et al., 2007)
  - schizophrenia (Enticott et al., 2008)
- **Neurological diseases**
  - Parkinson’s (Gauggel et al., 2004)
  - Alzheimer’s (Amieva et al., 2004)

Inhibitory Deficits ⇔ Patient Populations

- **Psychiatric conditions**
  - drug abuse (Nigg et al., 2006)
  - attention-deficit hyperactivity disorder (Alderston et al., 2007)
  - obsessive compulsive disorder (Menzies et al., 2007)
  - schizophrenia (Enticott et al., 2008)
- **Neurological diseases**
  - Parkinson's (Gauggel et al., 2004)
  - Alzheimer's (Amieva et al., 2004)
- **Pharmacology**
  - ADHD drug (atomoxetine, NE-reuptake inhibitor) improves stopping in animals, healthy volunteers, ADHD patients (Chamberlain et al., 2008)

Experimental Paradigm: Stop Signal Task

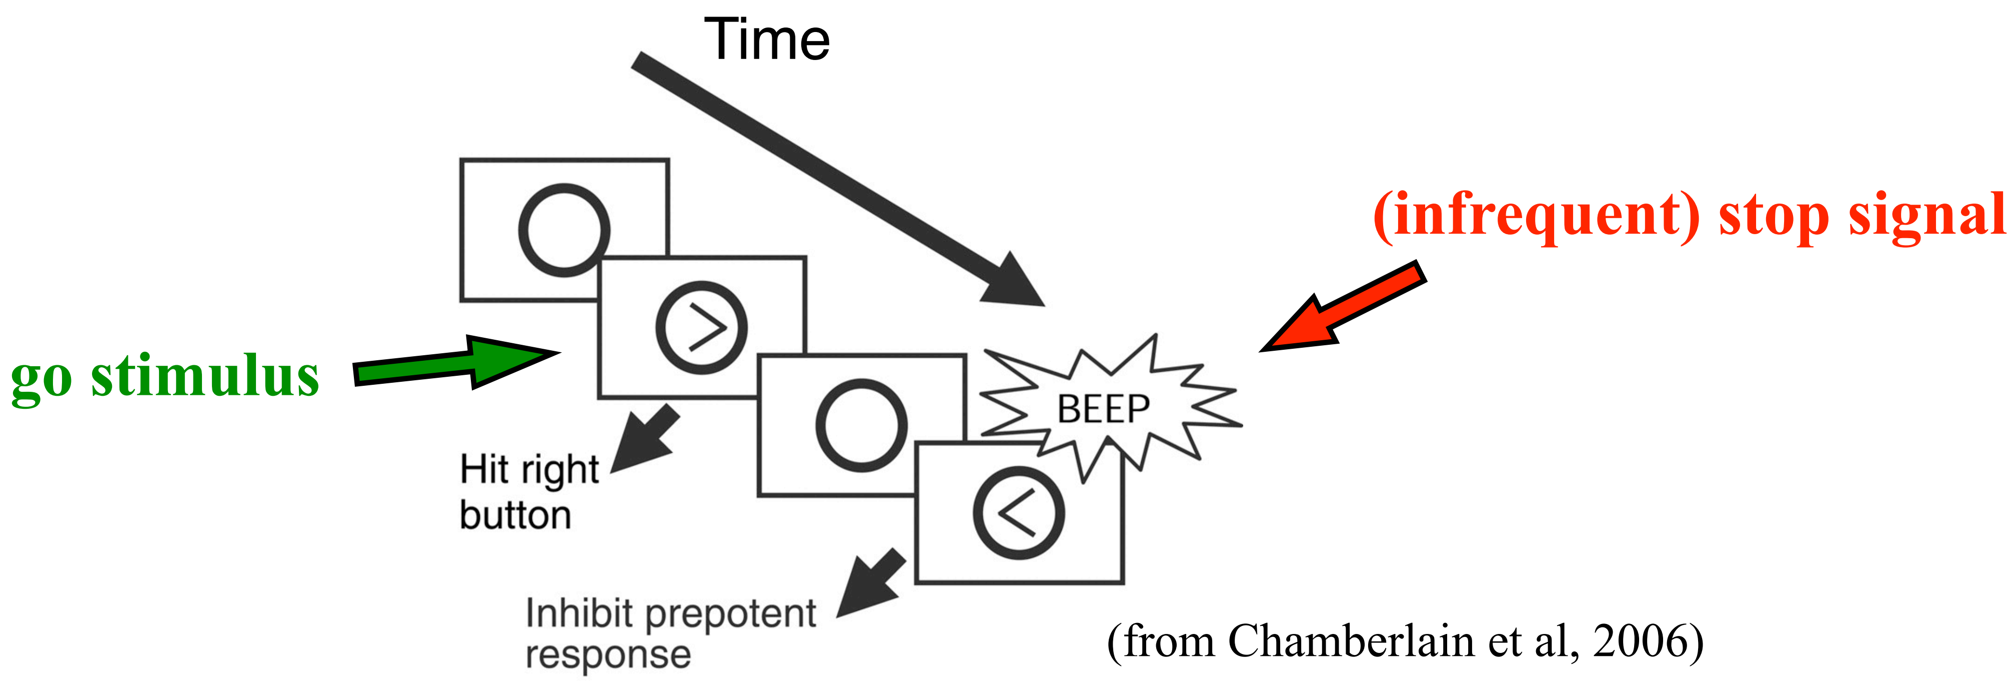

(from Chamberlain et al, 2006)

Experimental Paradigm: Stop Signal Task

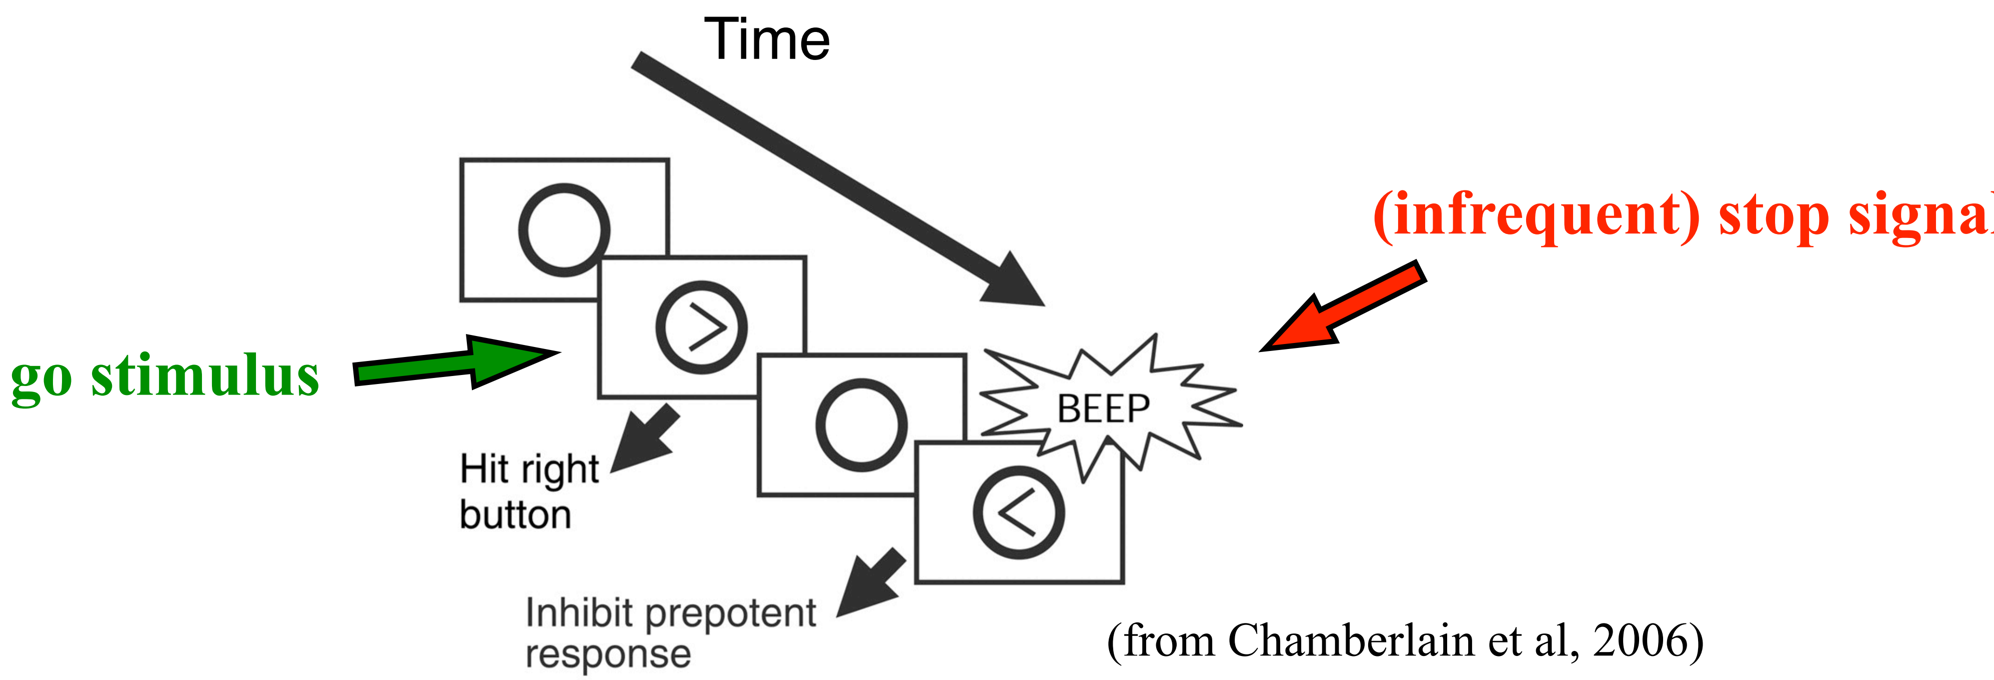

(from Chamberlain et al, 2006)

Stop trial: correct (canceled response)

Experimental Paradigm: Stop Signal Task

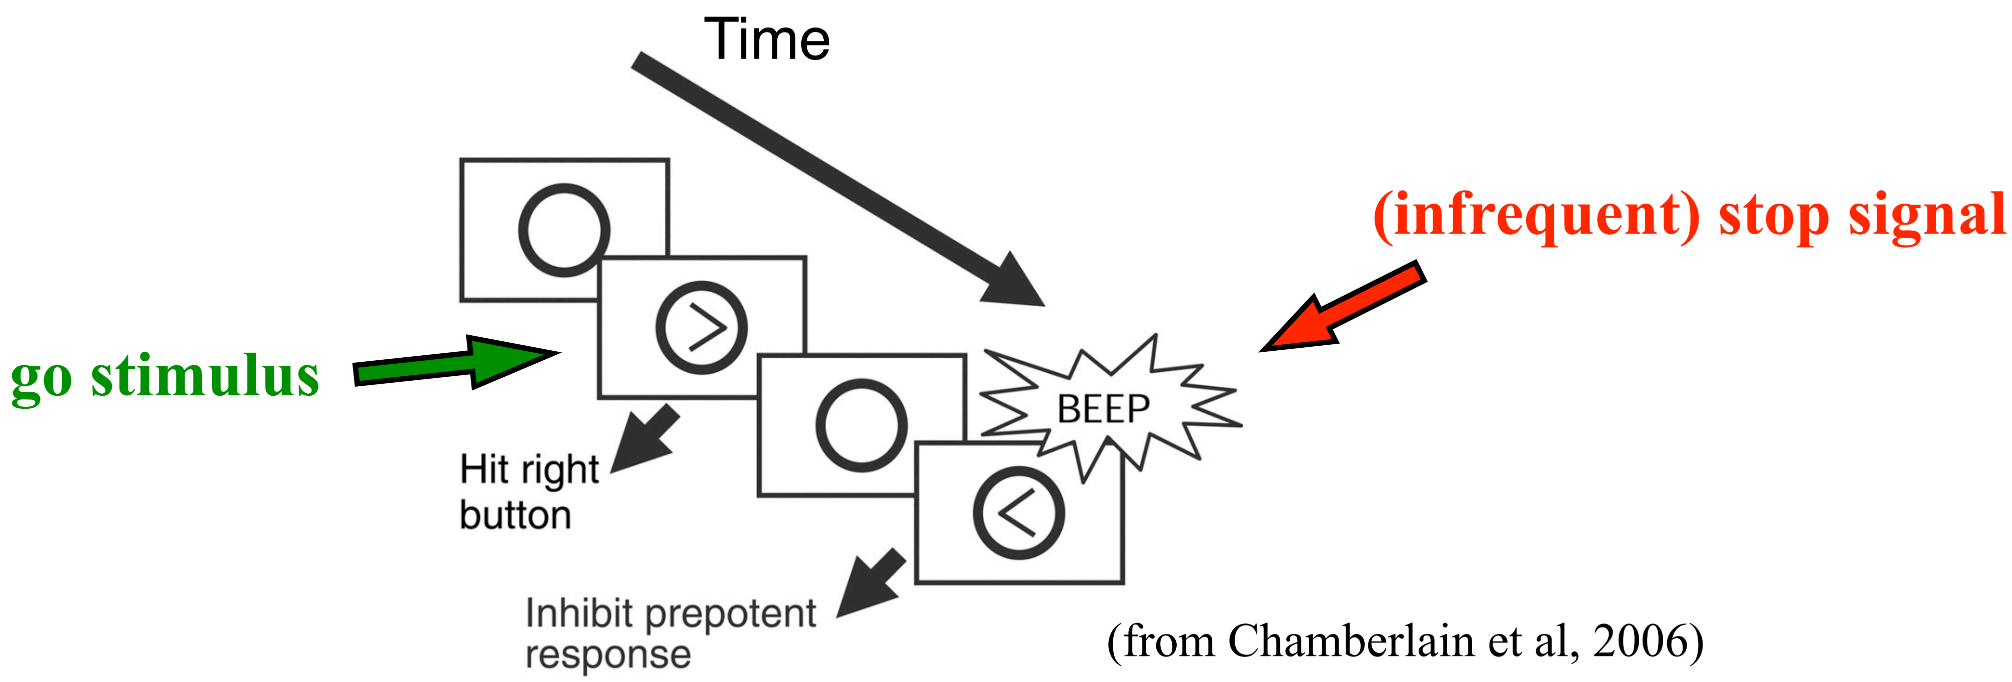

(from Chamberlain et al, 2006)

Stop trial: correct (canceled response)

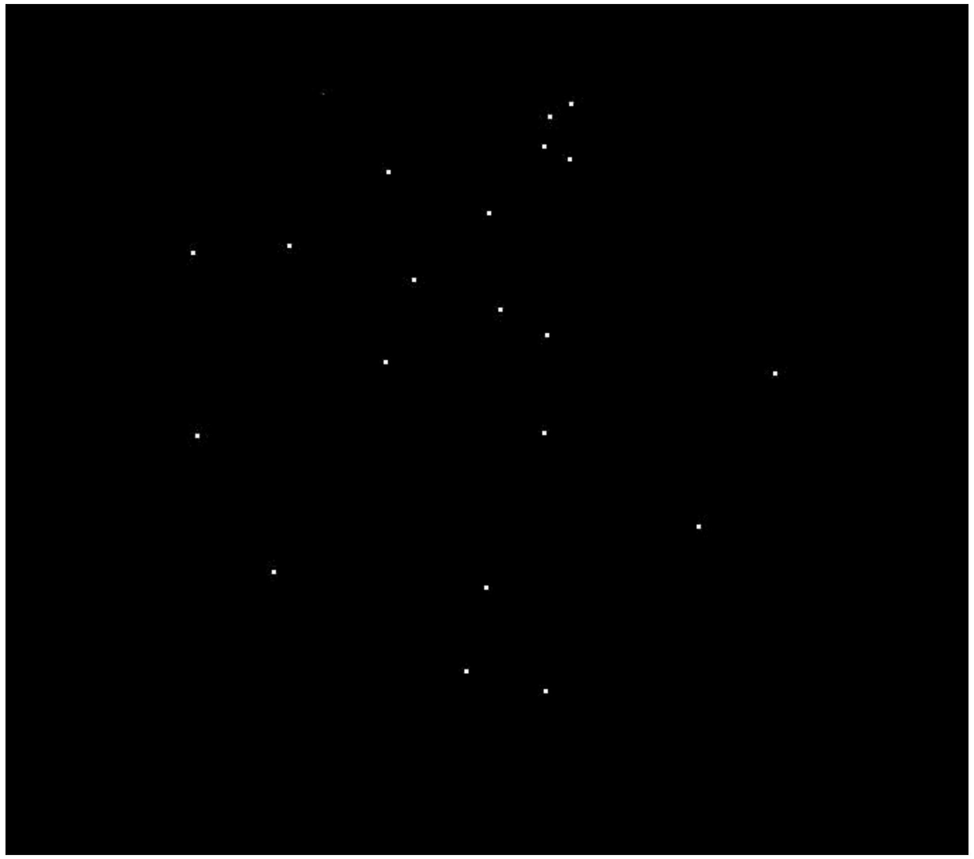

Experimental Paradigm: Stop Signal Task

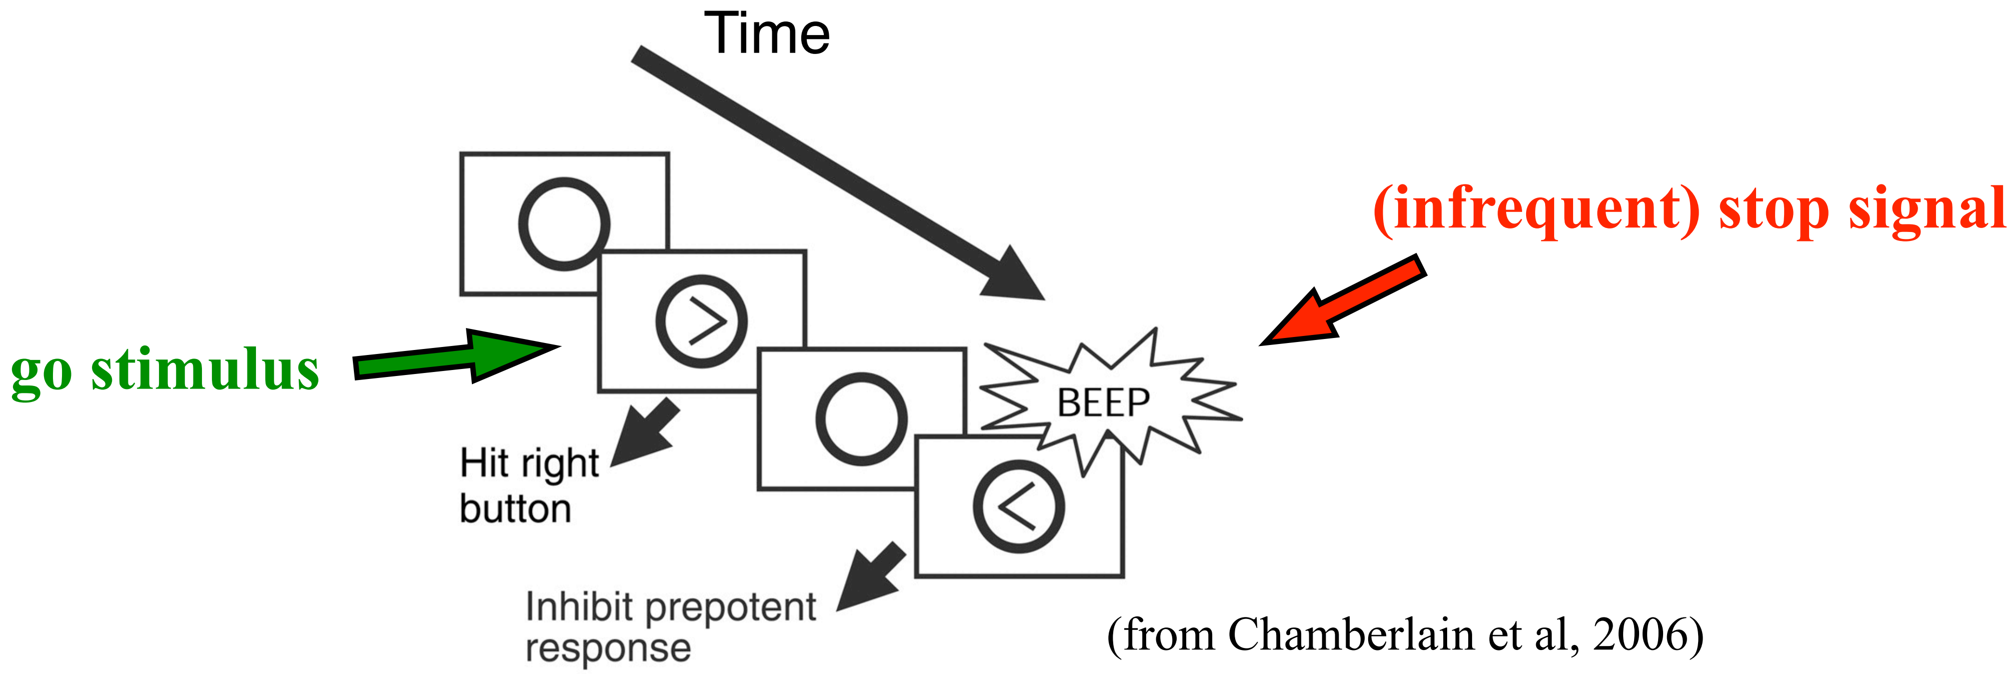

Stop trial: error (non-canceled response)

Experimental Paradigm: Stop Signal Task

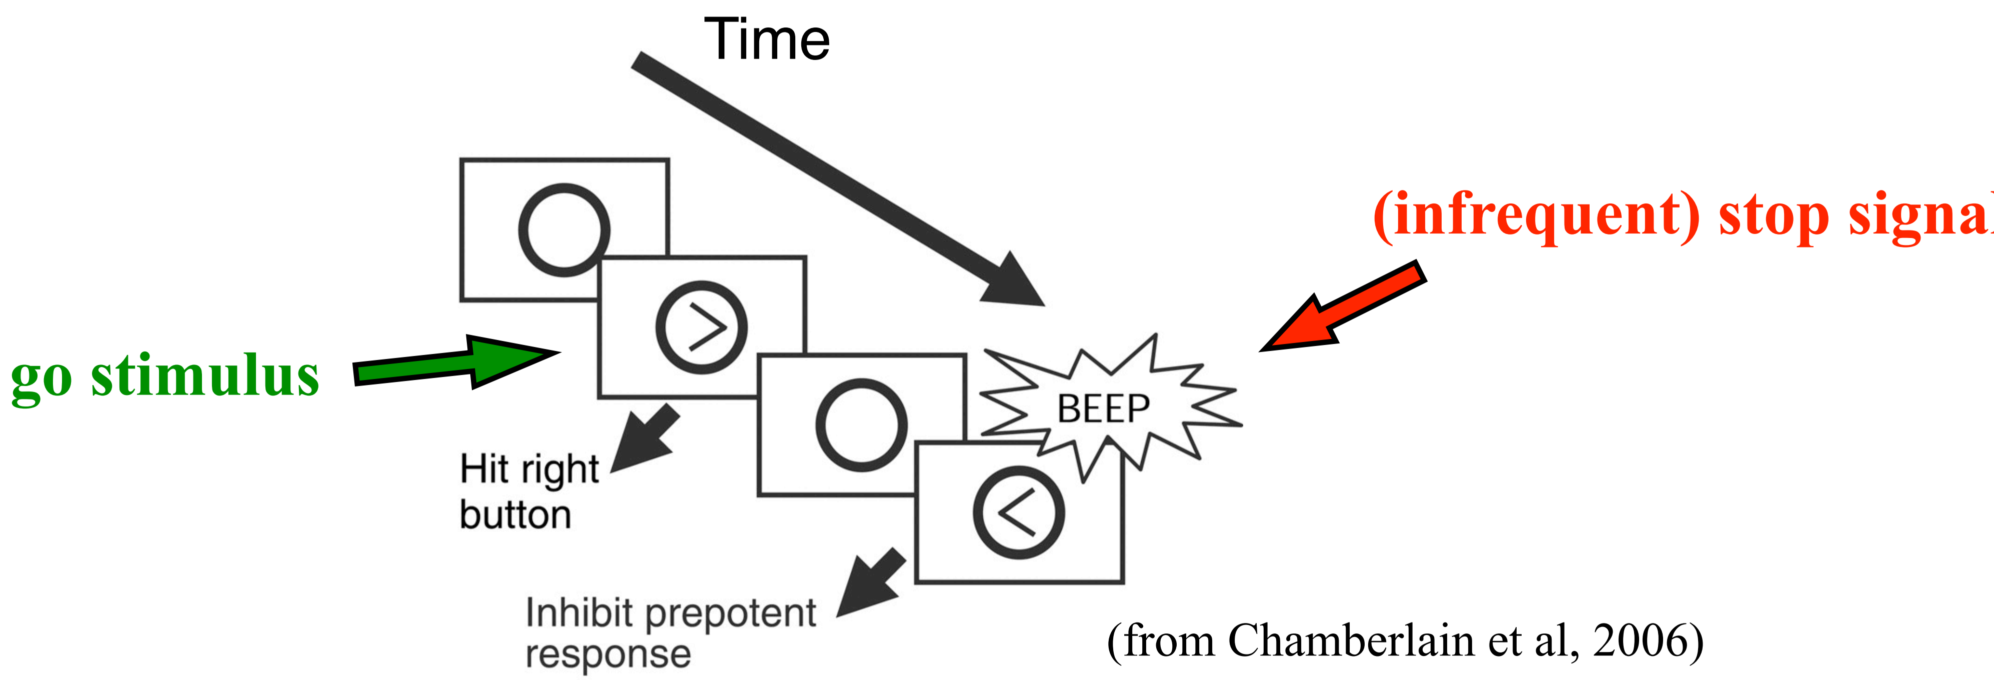

(from Chamberlain et al, 2006)

Stop trial: error (non-canceled response)

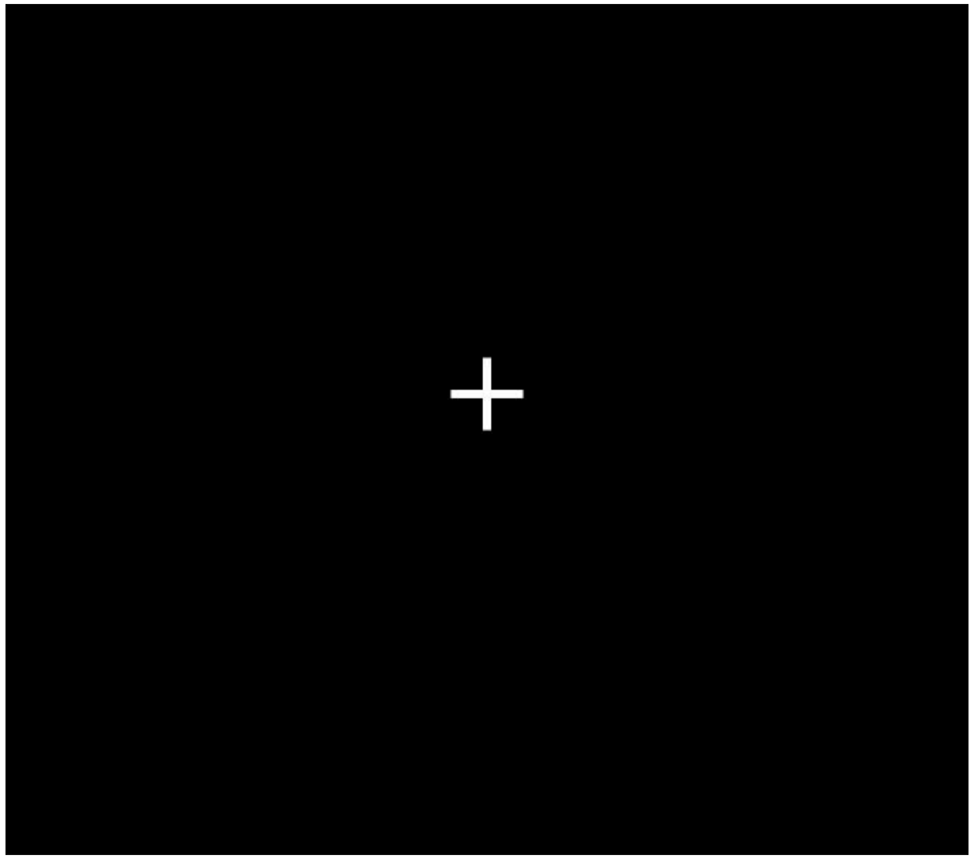

Experimental Paradigm: Stop Signal Task

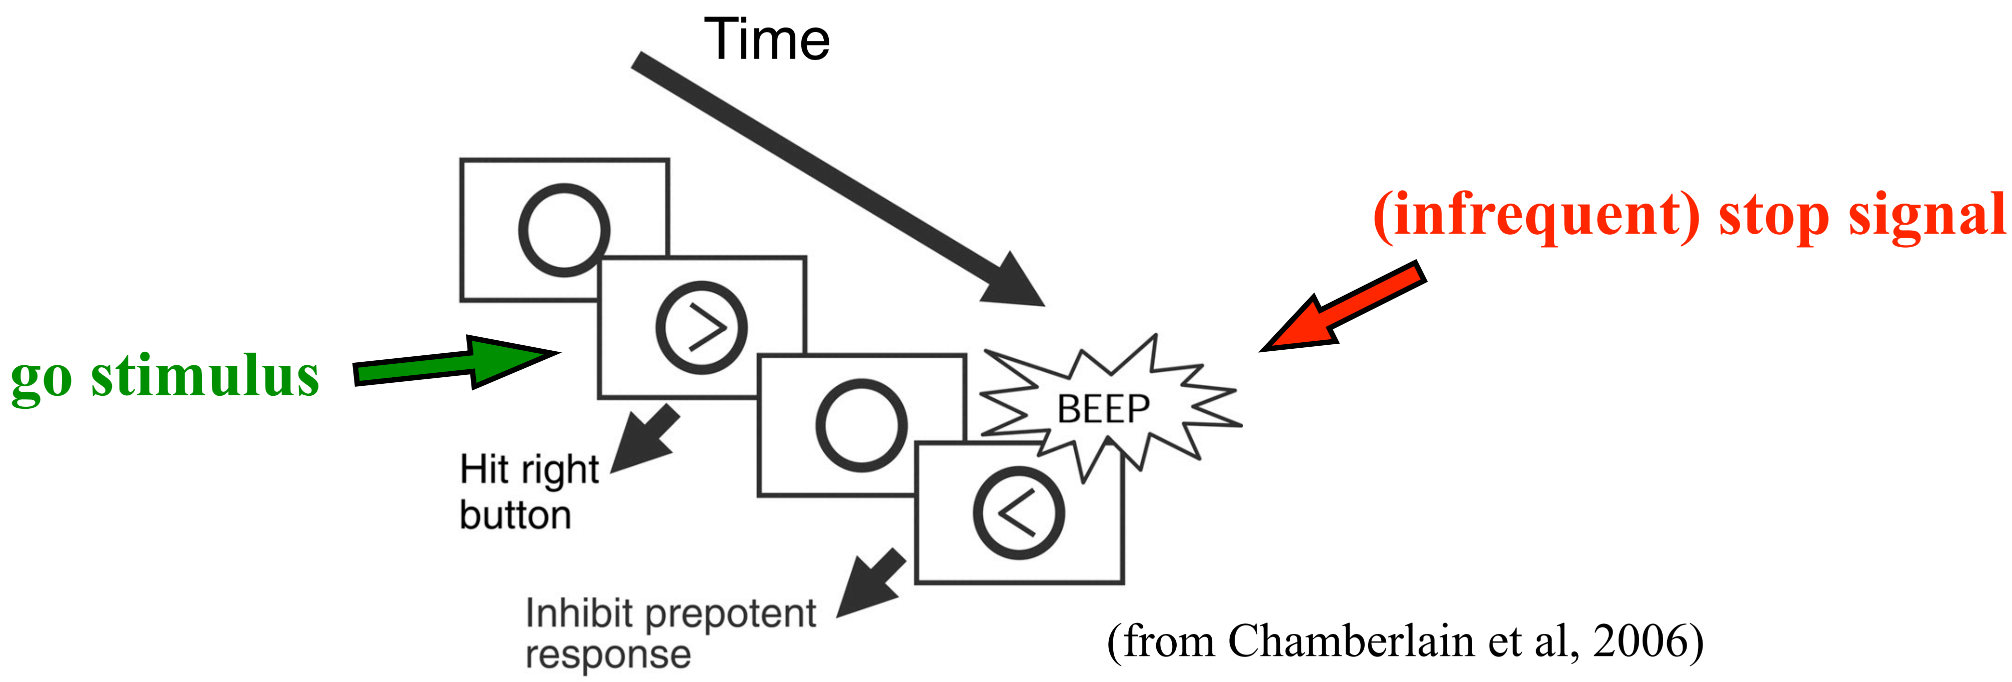

(from Chamberlain et al, 2006)

Go trial (no stop signal): error (timeout)

Experimental Paradigm: Stop Signal Task

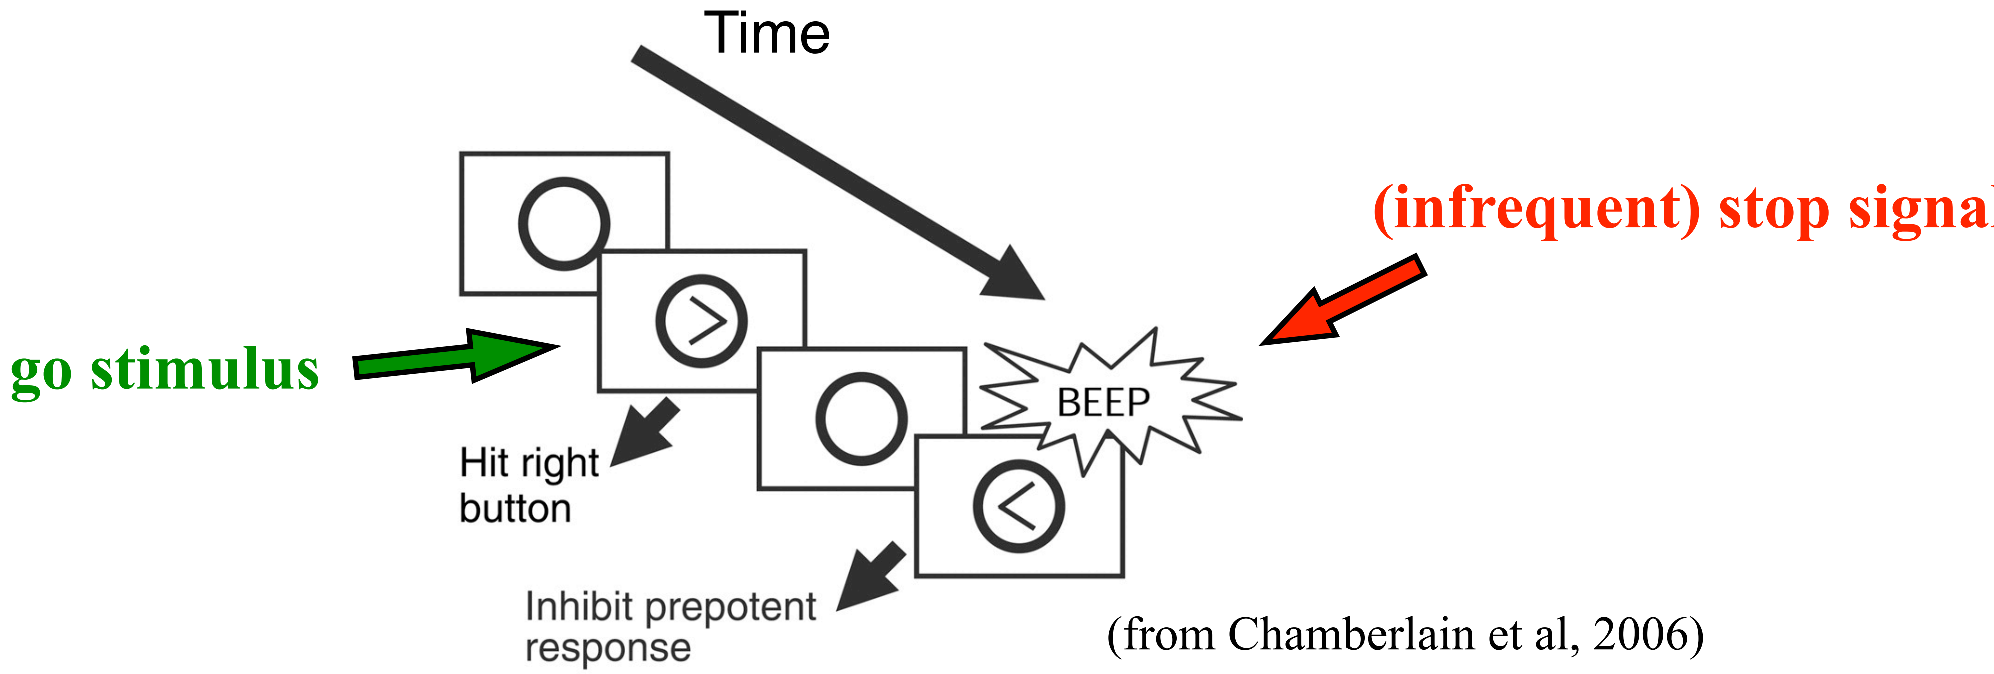

(from Chamberlain et al, 2006)

Go trial (no stop signal): error (timeout)

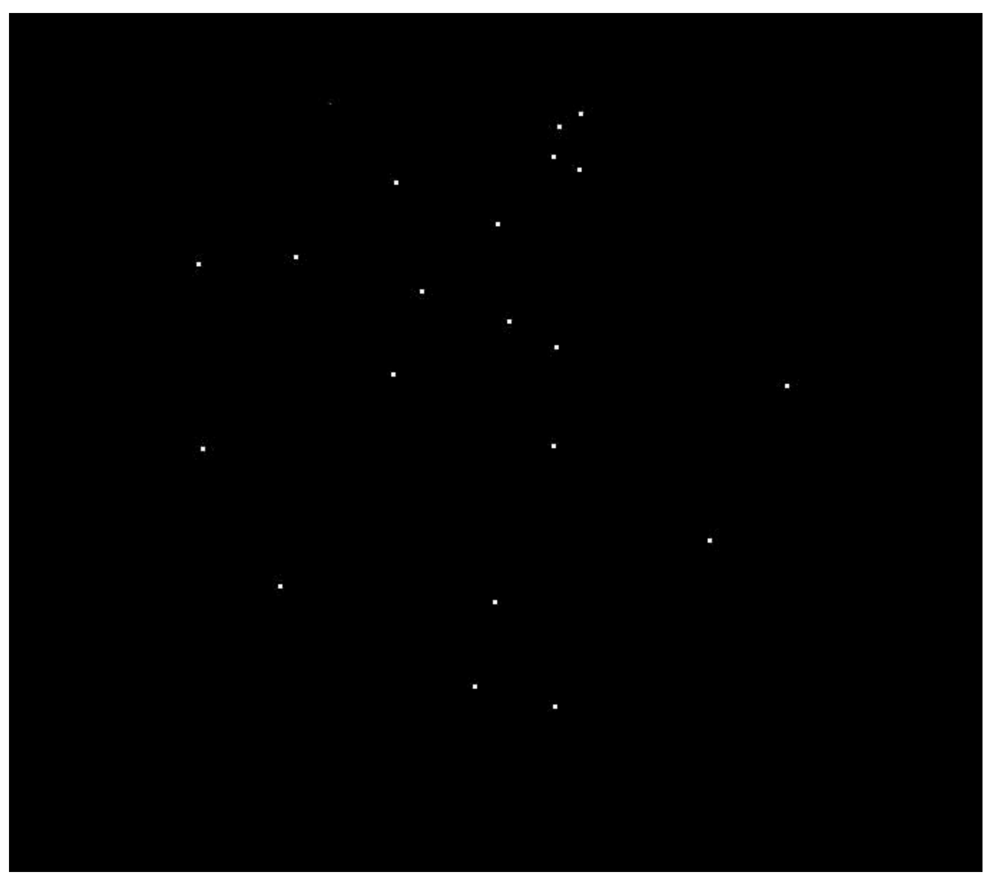

Stop Signal Task: Computational Challenges

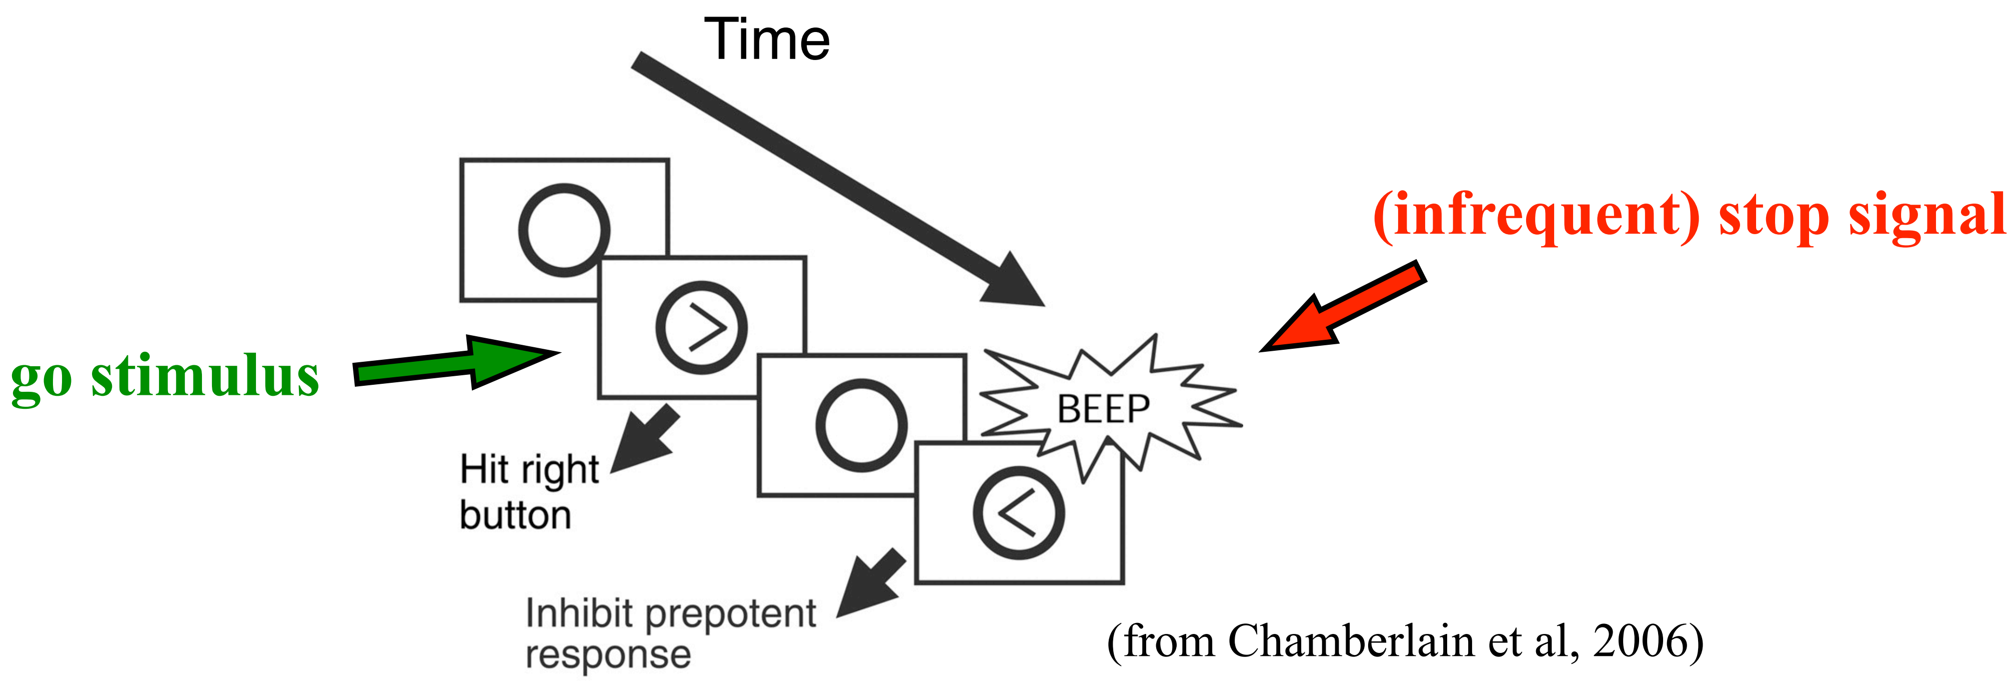

(from Chamberlain et al, 2006)

Stop Signal Task: Computational Challenges

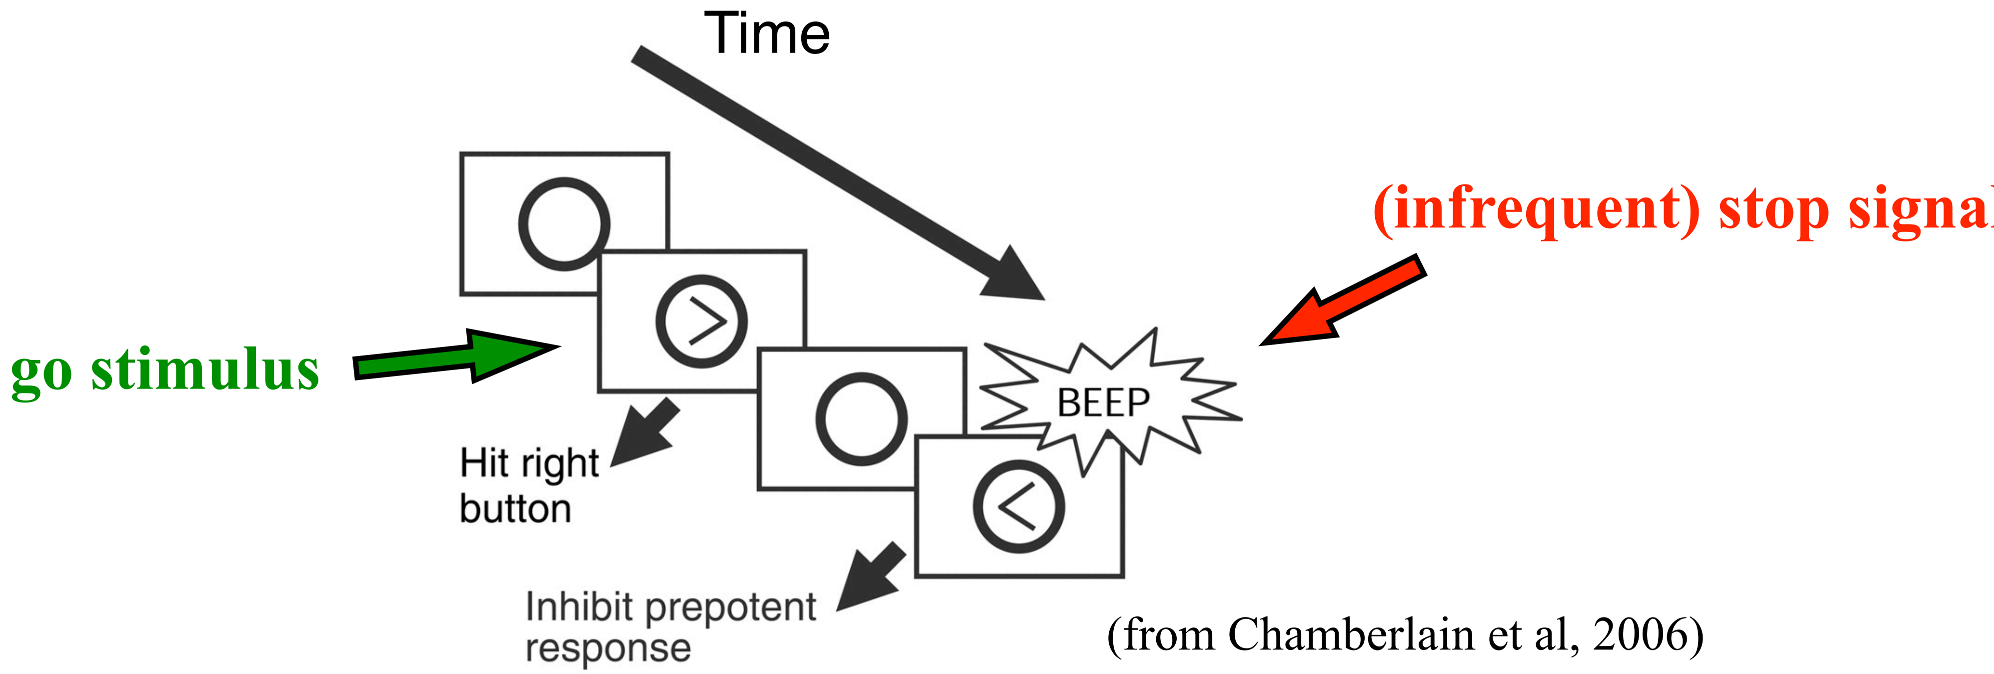

(from Chamberlain et al, 2006)

- **Sensory uncertainty**
  - go stimulus: left or right?
  - stop signal: present or absent?

Stop Signal Task: Computational Challenges

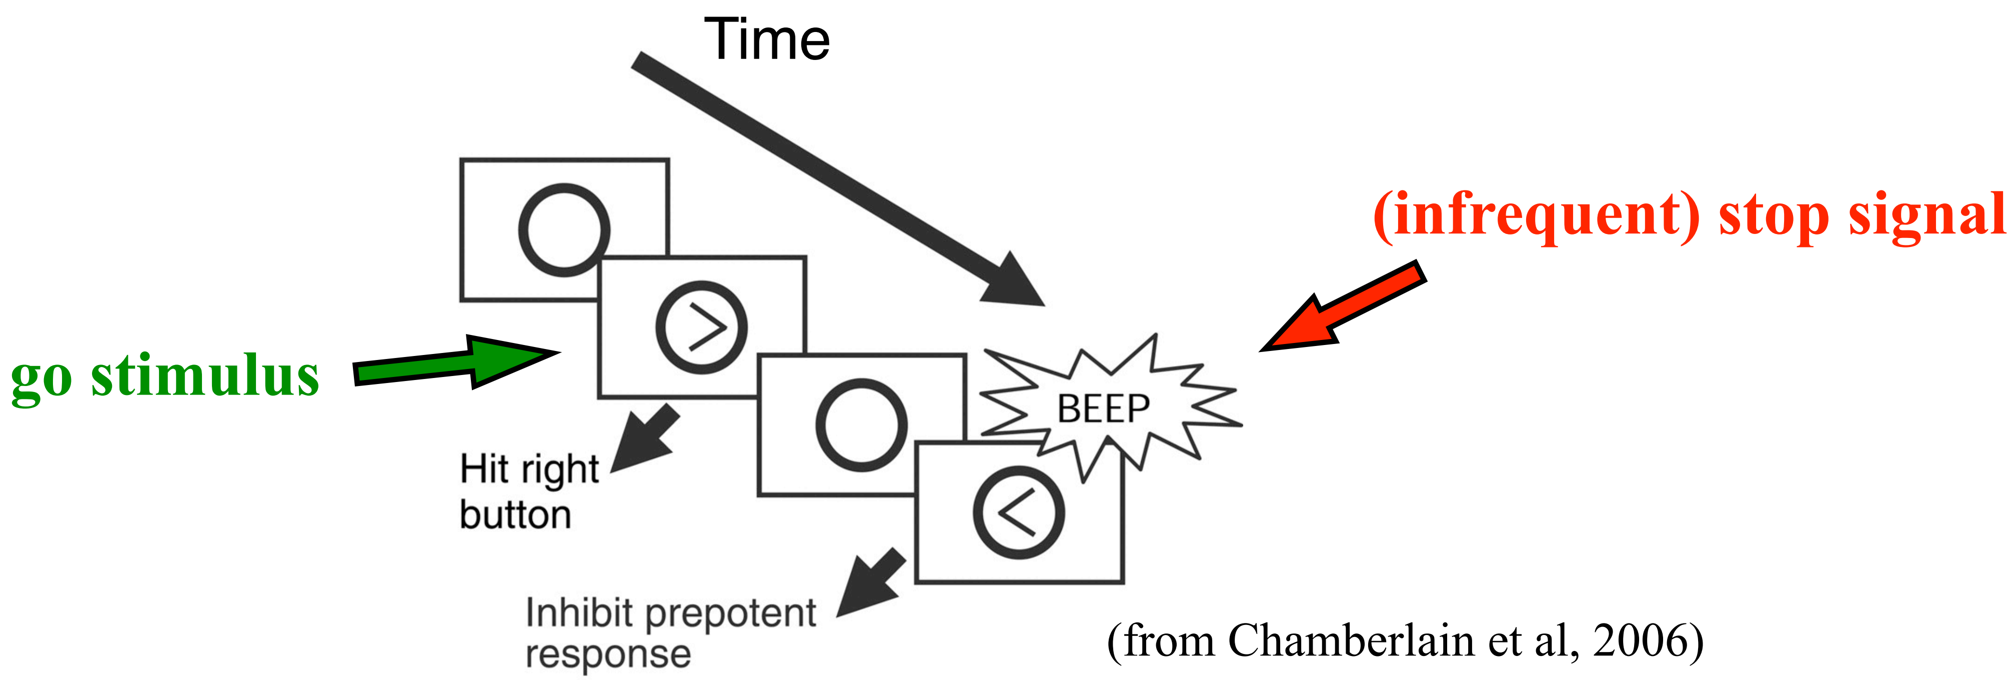

(from Chamberlain et al, 2006)

- **Sensory uncertainty**
  - **go** stimulus: left or right?
  - **stop** signal: present or absent?
- **Action uncertainty**
  - if **go** trial, would **go** response be too late?
  - if **stop** trial, would I **stop** in time?

Stop Signal Task: Computational Challenges

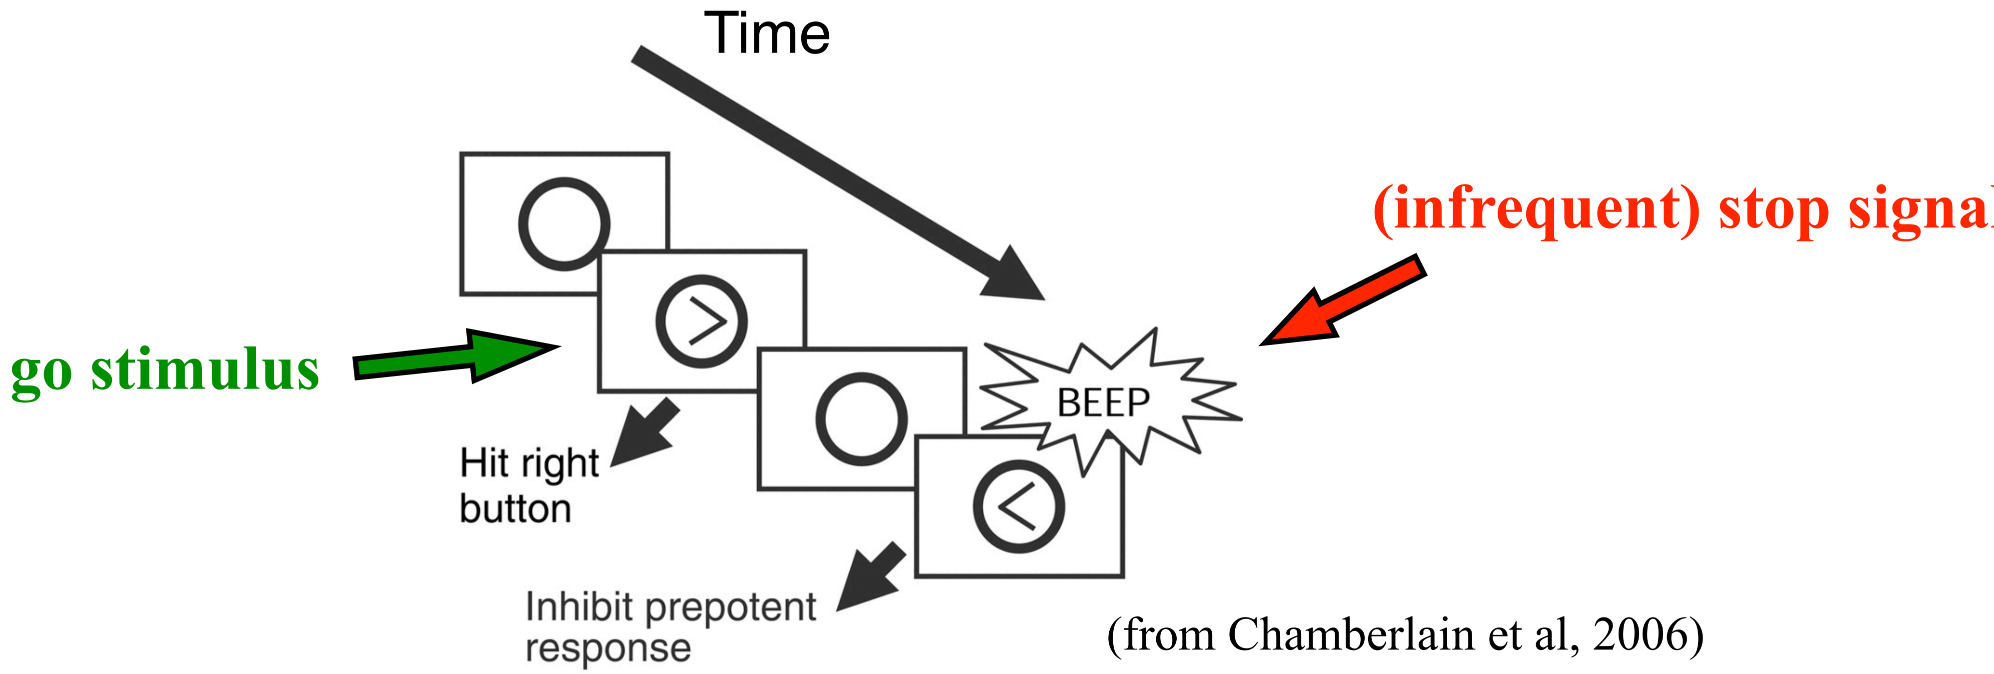

(from Chamberlain et al, 2006)

- **Sensory uncertainty**
  - **go** stimulus: left or right?
  - **stop** signal: present or absent?
- **Action uncertainty**
  - if **go** trial, would **go** response be too late?
  - if **stop** trial, would I **stop** in time?
- **Relative costs of action choices**
  - **stop** error penalty versus **go** response delay

Stop Signal Task: Computational Challenges

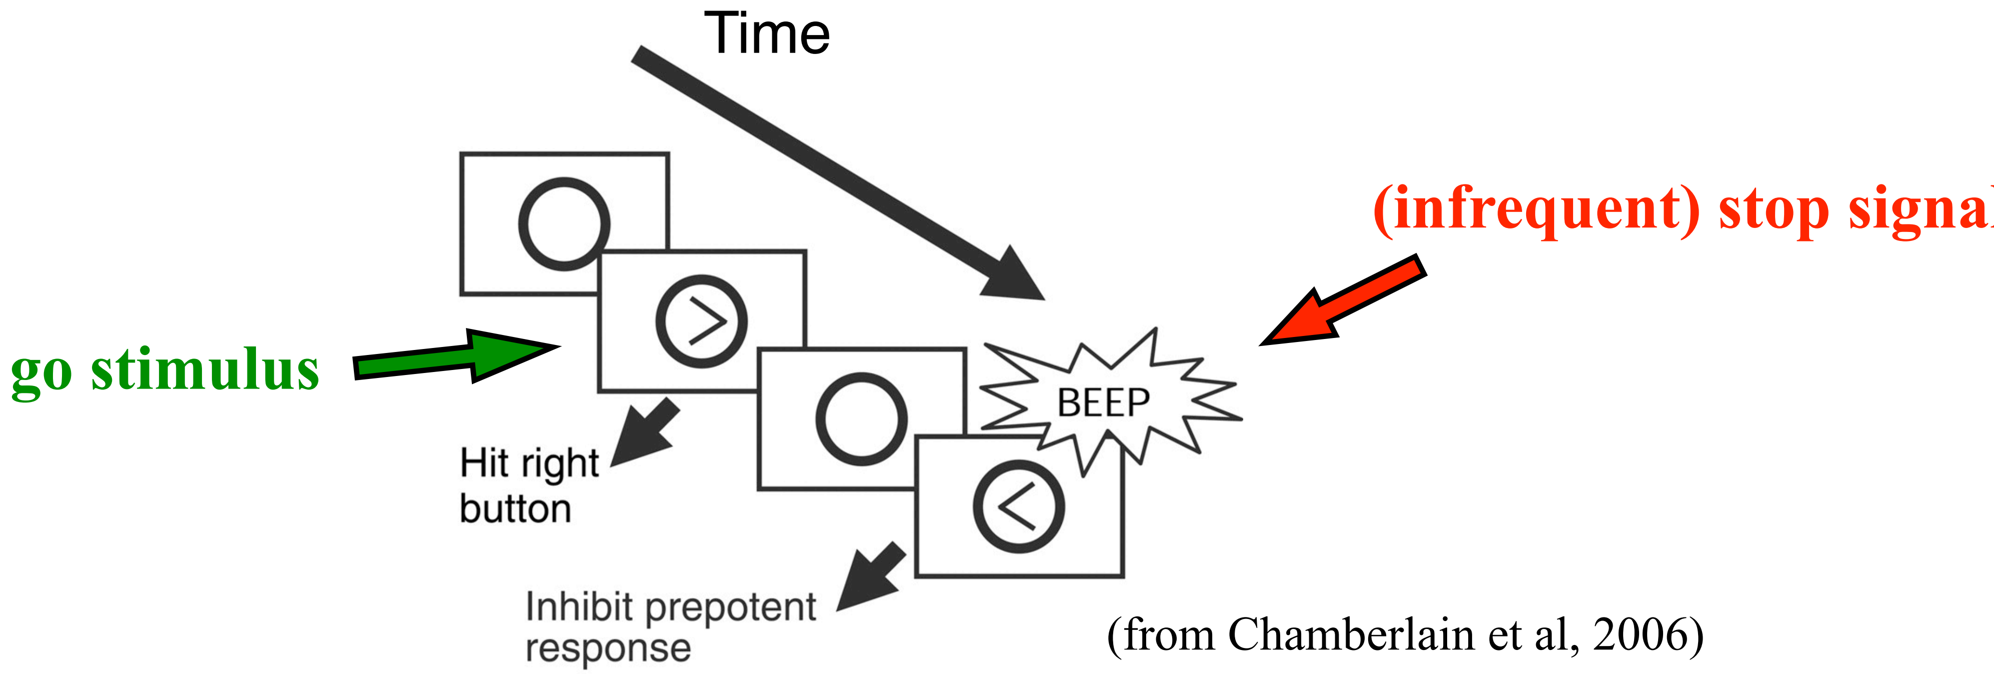

(from Chamberlain et al, 2006)

- **Sensory uncertainty**
  - **go** stimulus: left or right?
  - **stop** signal: present or absent?
- **Action uncertainty**
  - if **go** trial, would **go** response be too late?
  - if **stop** trial, would I **stop** in time?
- **Relative costs of action choices**
  - **stop** error penalty versus **go** response delay
- **Learning (prior information)**
  - frequency of **stop** trials, **stop** signal onset, penalties

Inhibitory Control: An Everyday Example

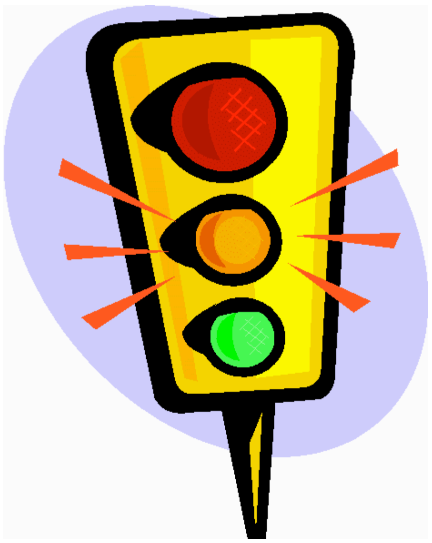

Possible actions

- stop
- go

Inhibitory Control: An Everyday Example

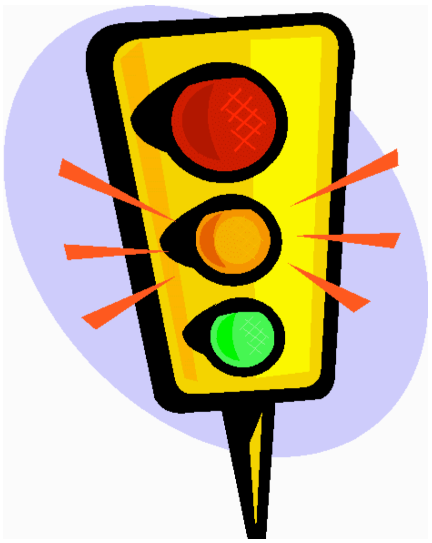

Possible actions

- stop
- go

- **Sensory uncertainty**
  - is that a yellow light or a yellow street lamp?
  - how far away is the intersection?

Inhibitory Control: An Everyday Example

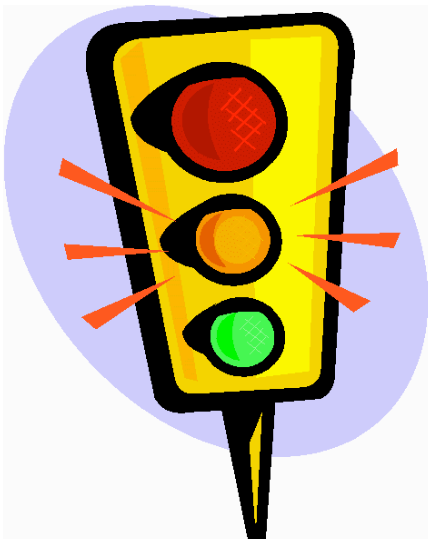

Possible actions

- stop
- go

- **Sensory uncertainty**
  - is that a yellow light or a yellow street lamp?
  - how far away is the intersection?
- **Action uncertainty**
  - if go, would light turn red before crossing intersection?
  - if stop, would the car stop in time?

Inhibitory Control: An Everyday Example

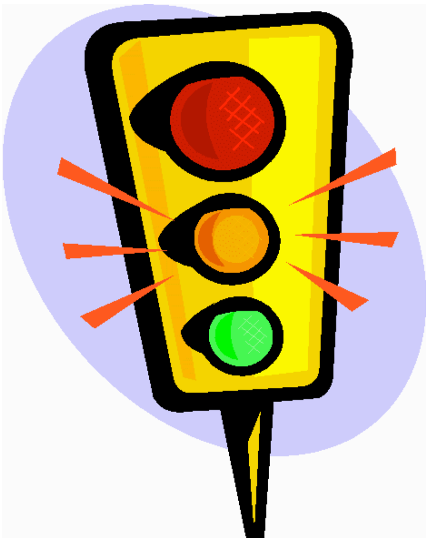

Possible actions

- stop
- go

- **Sensory uncertainty**
  - is that a yellow light or a yellow street lamp?
  - how far away is the intersection?
- **Action uncertainty**
  - if go, would light turn red before crossing intersection?
  - if stop, would the car stop in time?
- **Relative costs of action choices**
  - cops/tickets versus temporal delay

Inhibitory Control: An Everyday Example

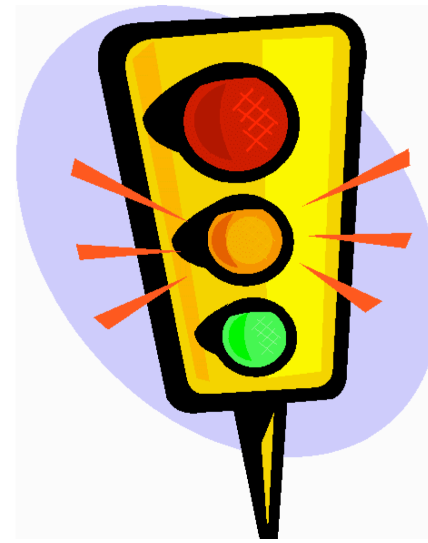

Possible actions

- stop
- go

- **Sensory uncertainty**
  - is that a yellow light or a yellow street lamp?
  - how far away is the intersection?
- **Action uncertainty**
  - if go, would light turn red before crossing intersection?
  - if stop, would the car stop in time?
- **Relative costs of action choices**
  - cops/tickets versus temporal delay
- **Learning (prior information)**
  - duration of yellow light, P(cop), $ ticket

Outline

- **Model: brain implements rational (optimal) computations**
  - Sensory processing ⇔ Bayesian inference
  - Action selection ⇔ optimal stochastic control
- **Model captures a range of behavioral results**
  - Classical results
  - Reward/motivation
  - Contextual effects, sequential effects
- **Neural implementation**
  - Race (drift-diffusion) model as neurally plausible approximation

Rational Behavior in Stop Signal Task

Fundamental decision: when (whether) to **go**?

Rational Behavior in Stop Signal Task

Fundamental decision: when (whether) to **go**?

Sensory Processing

Bayesian inference

Track beliefs over time about

- *go* stimulus identity (L/R)
- *stop* signal presence (Y/N)
- frequency of stop trials

based on ***noisy sensory inputs***

Rational Behavior in Stop Signal Task

Fundamental decision: when (whether) to **go**?

Sensory Processing

Bayesian inference

Track beliefs over time about

- *go* stimulus identity (L/R)
- *stop* signal presence (Y/N)
- frequency of stop trials

based on ***noisy sensory inputs***

Action Selection

Stochastic control

Choose action given belief state

- *go* (L/R), or
- *wait*
- *stop = wait, wait, wait...*

based on ***expected consequences***

Sensory Processing = Bayesian Inference

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1 - p_z^{t-1}) h(t)) f_z(y^t)$$

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1 - p_z^{t-1}) h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

Simulation: Belief State Trajectories

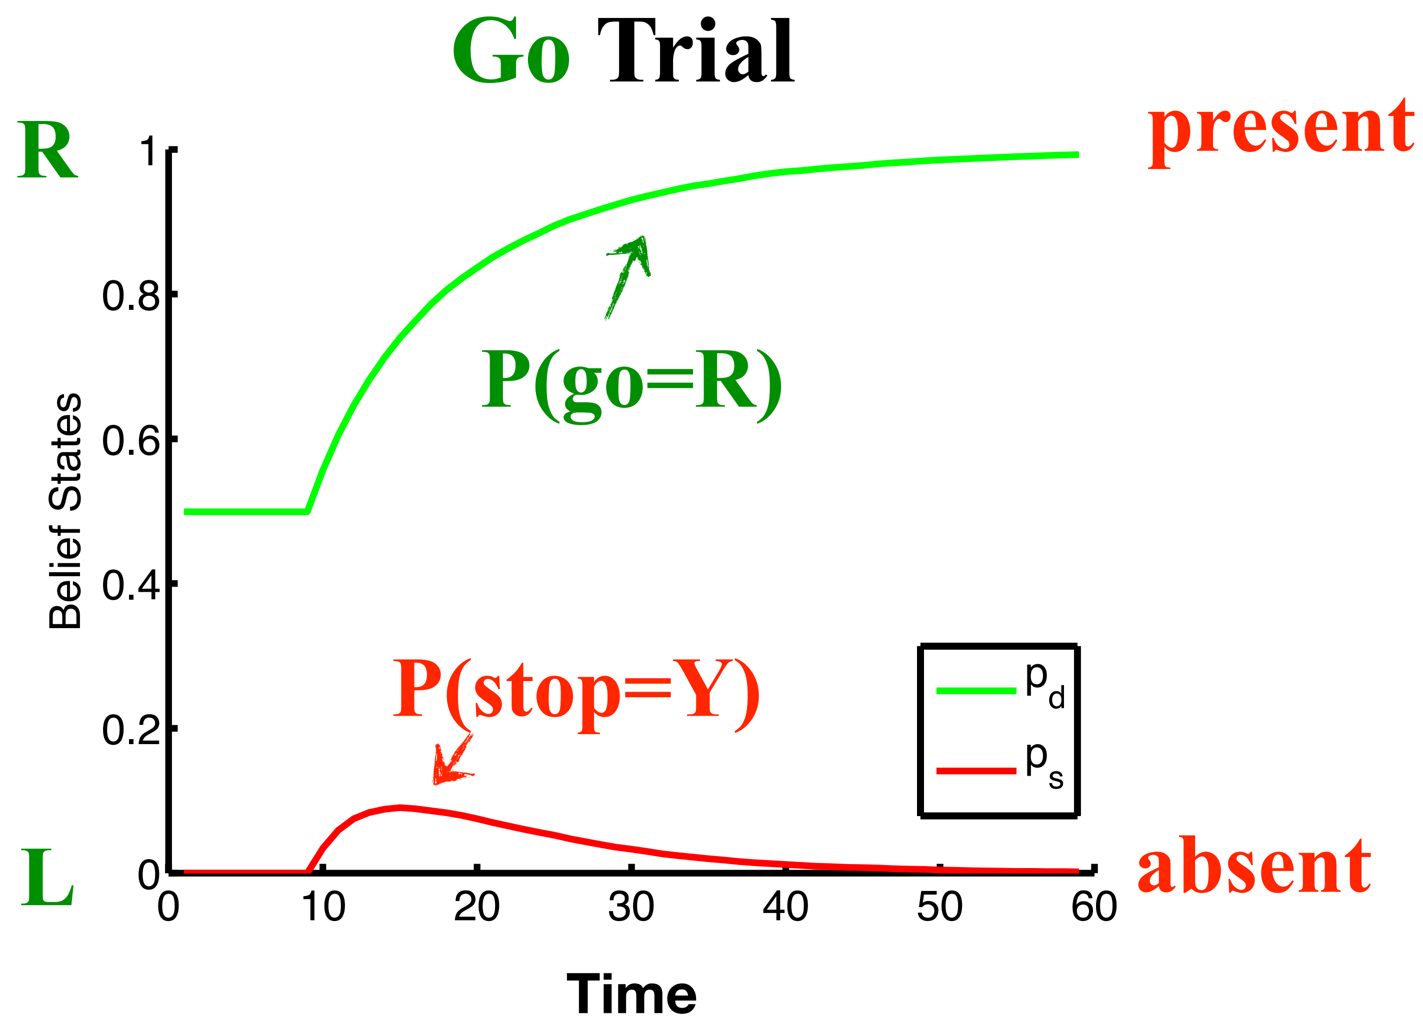

Simulation: Belief State Trajectories

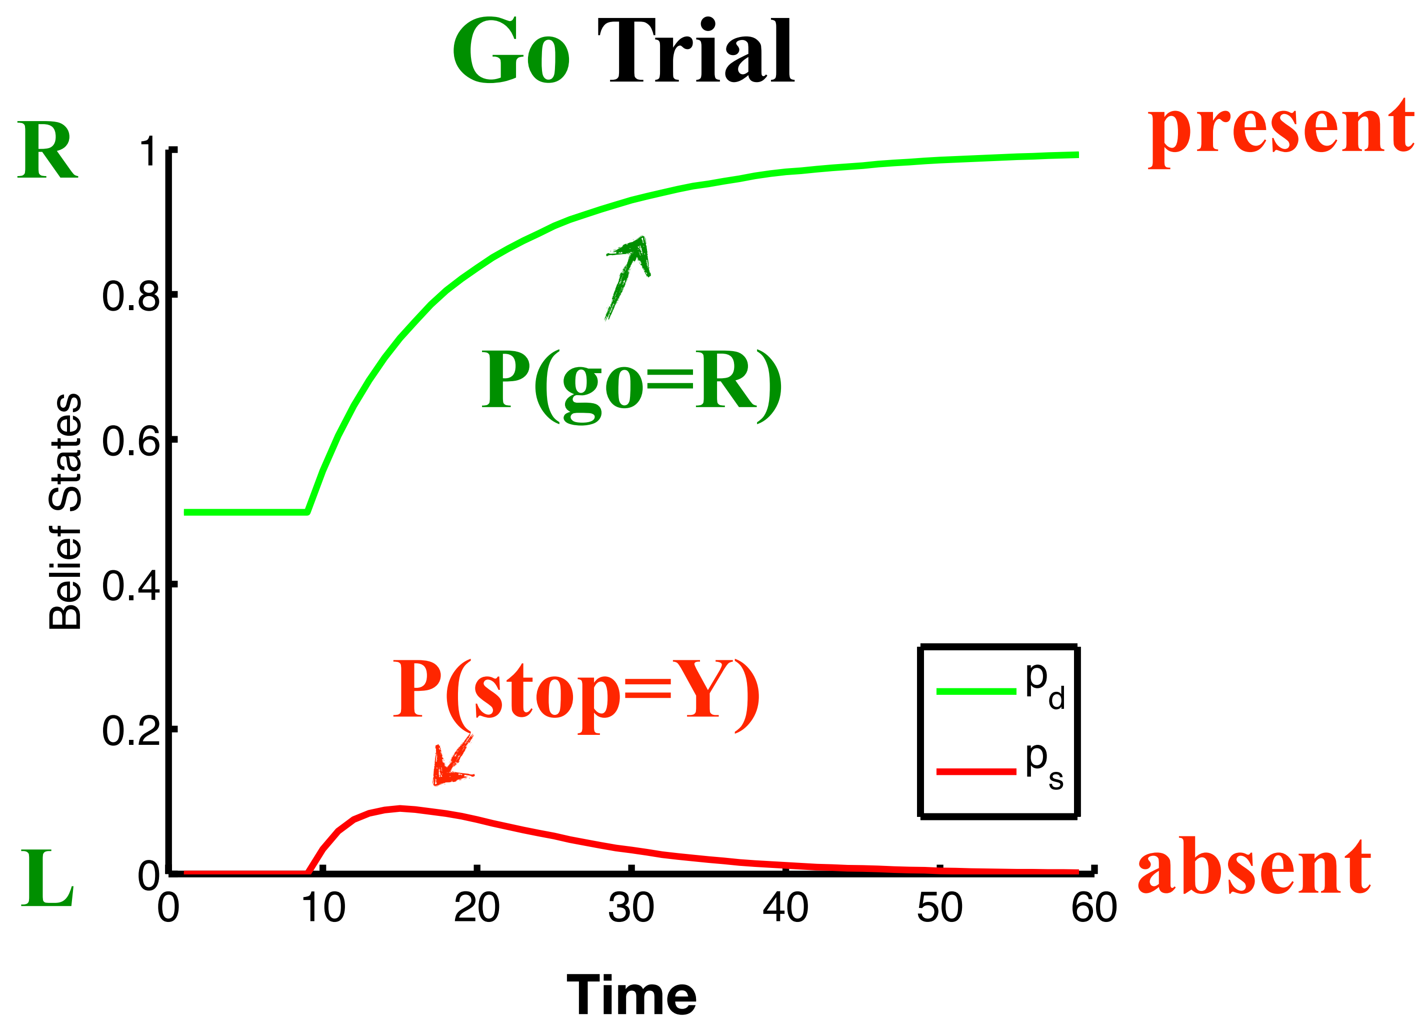

- evidence accumulates for go stimulus identity (p_d)

Simulation: Belief State Trajectories

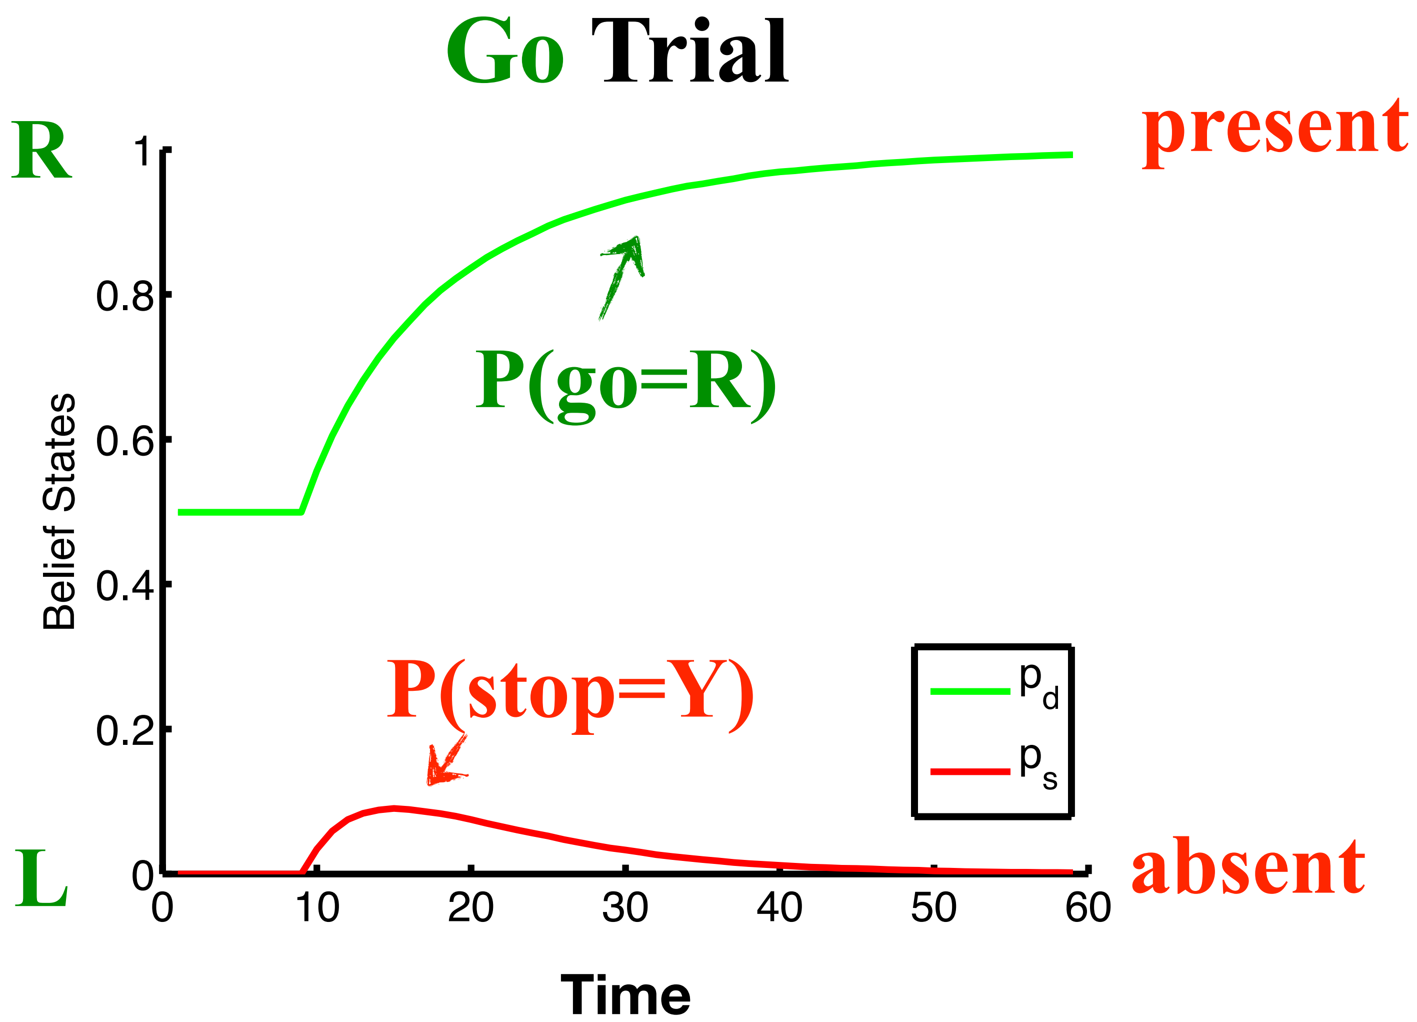

- evidence accumulates for go stimulus identity (p_d)
- stop trial probability (p_s) rises (prior expectation) then falls

Simulation: Belief State Trajectories

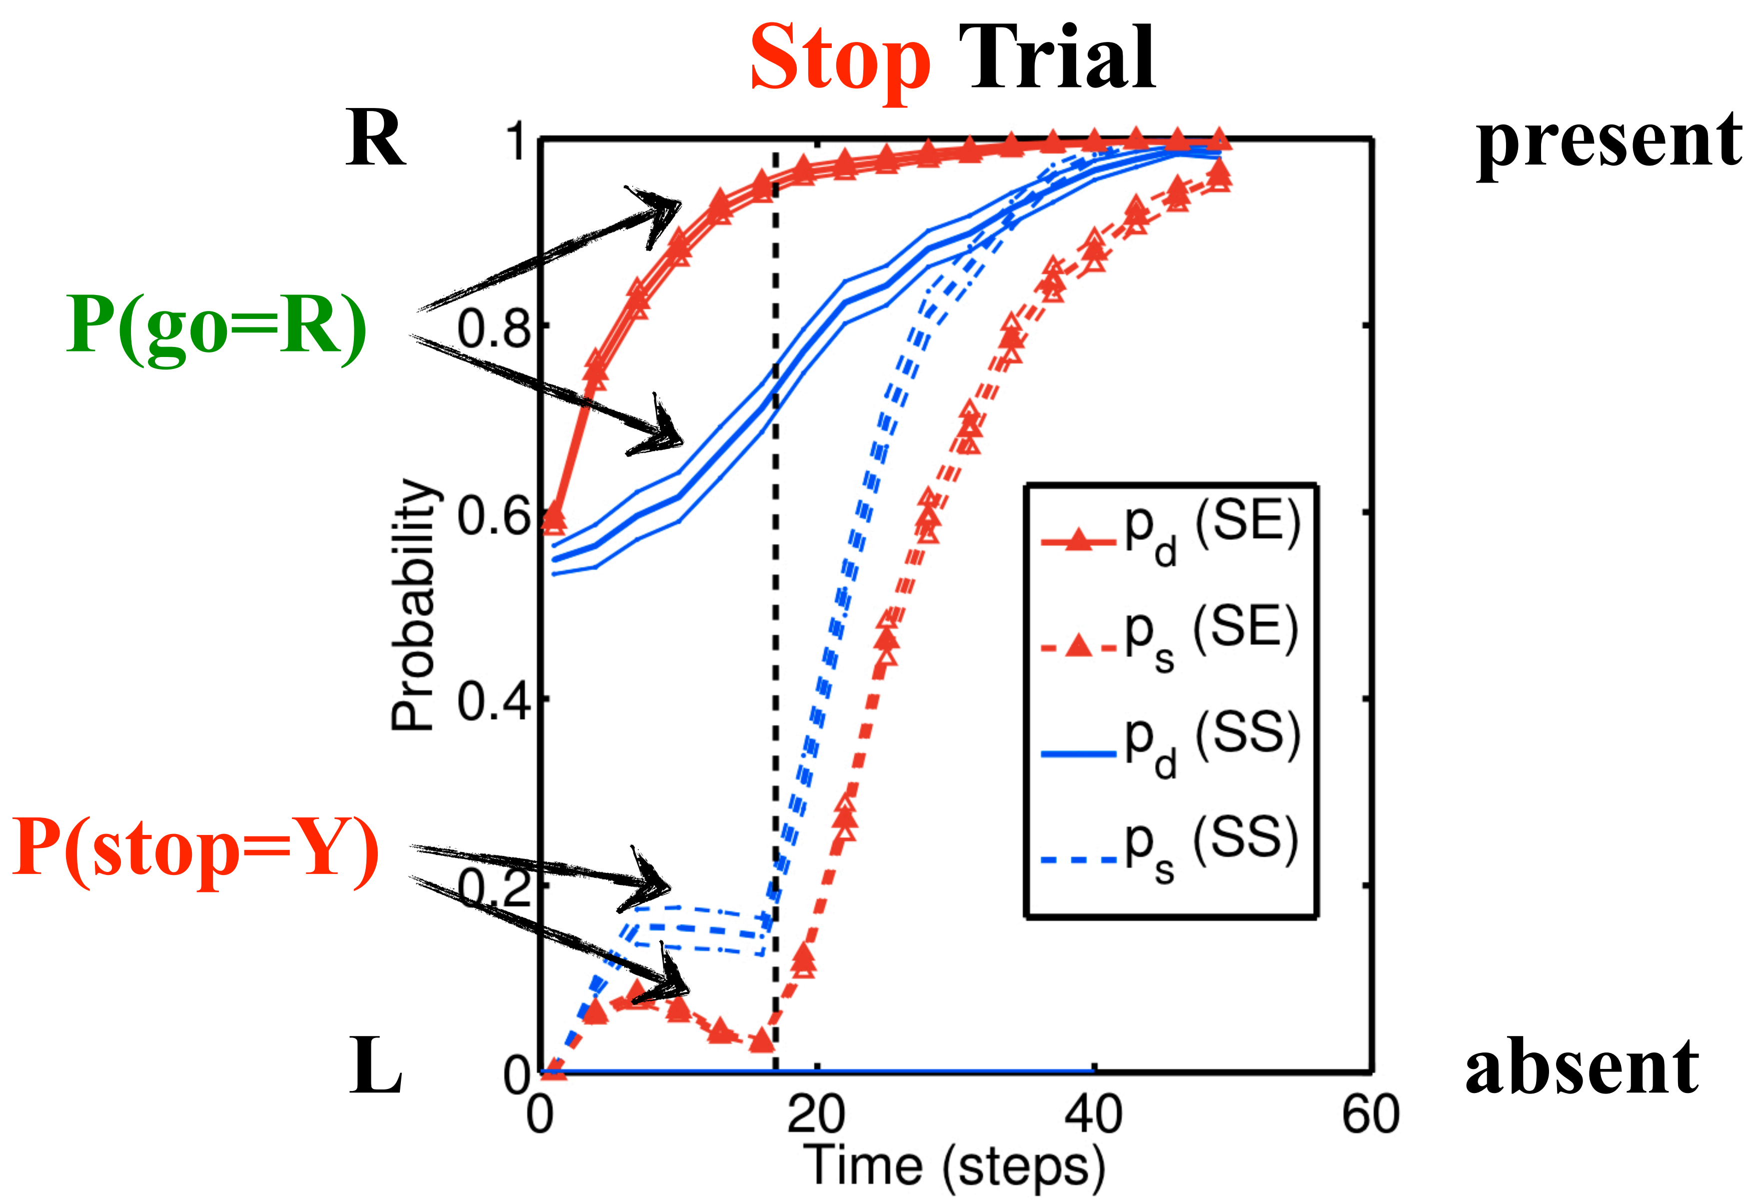

Simulation: Belief State Trajectories

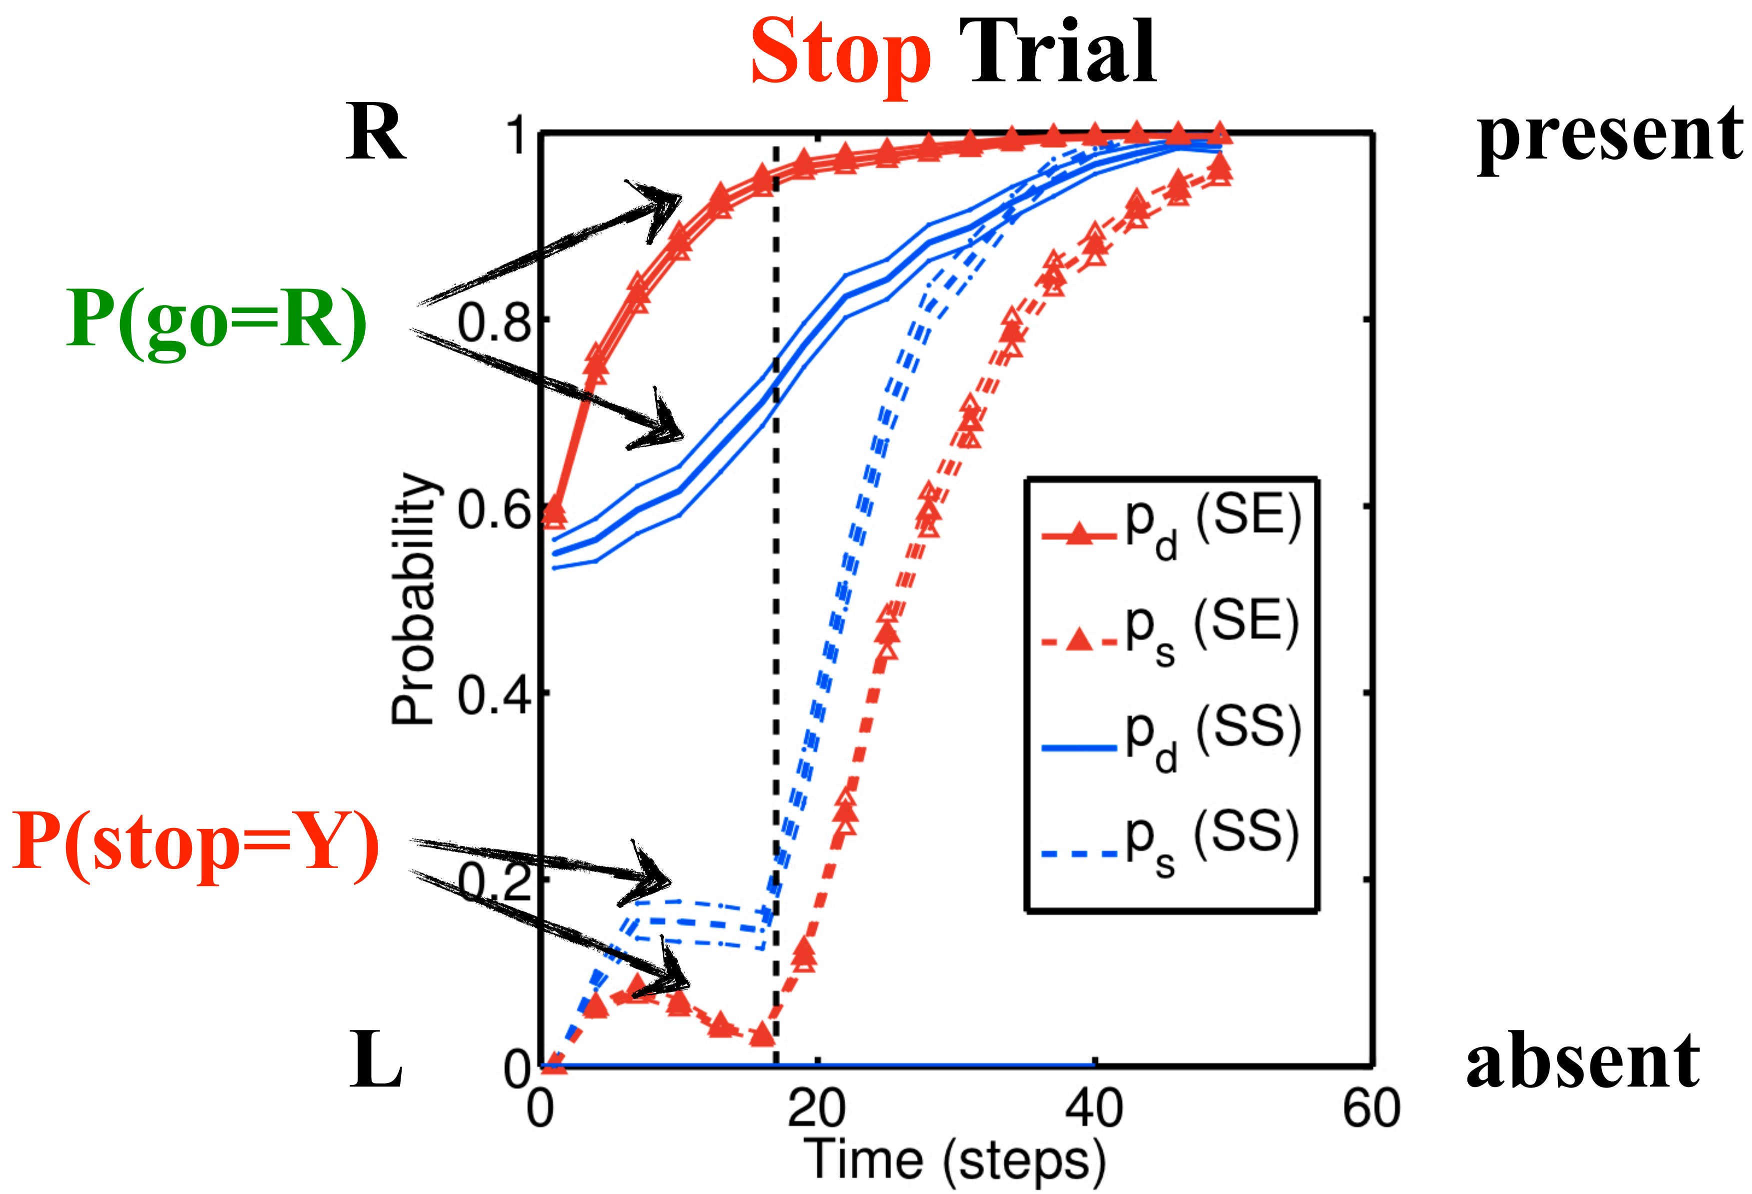

- stop-success: $p_d \uparrow$ slowly, $p_s \uparrow$ quickly

Simulation: Belief State Trajectories

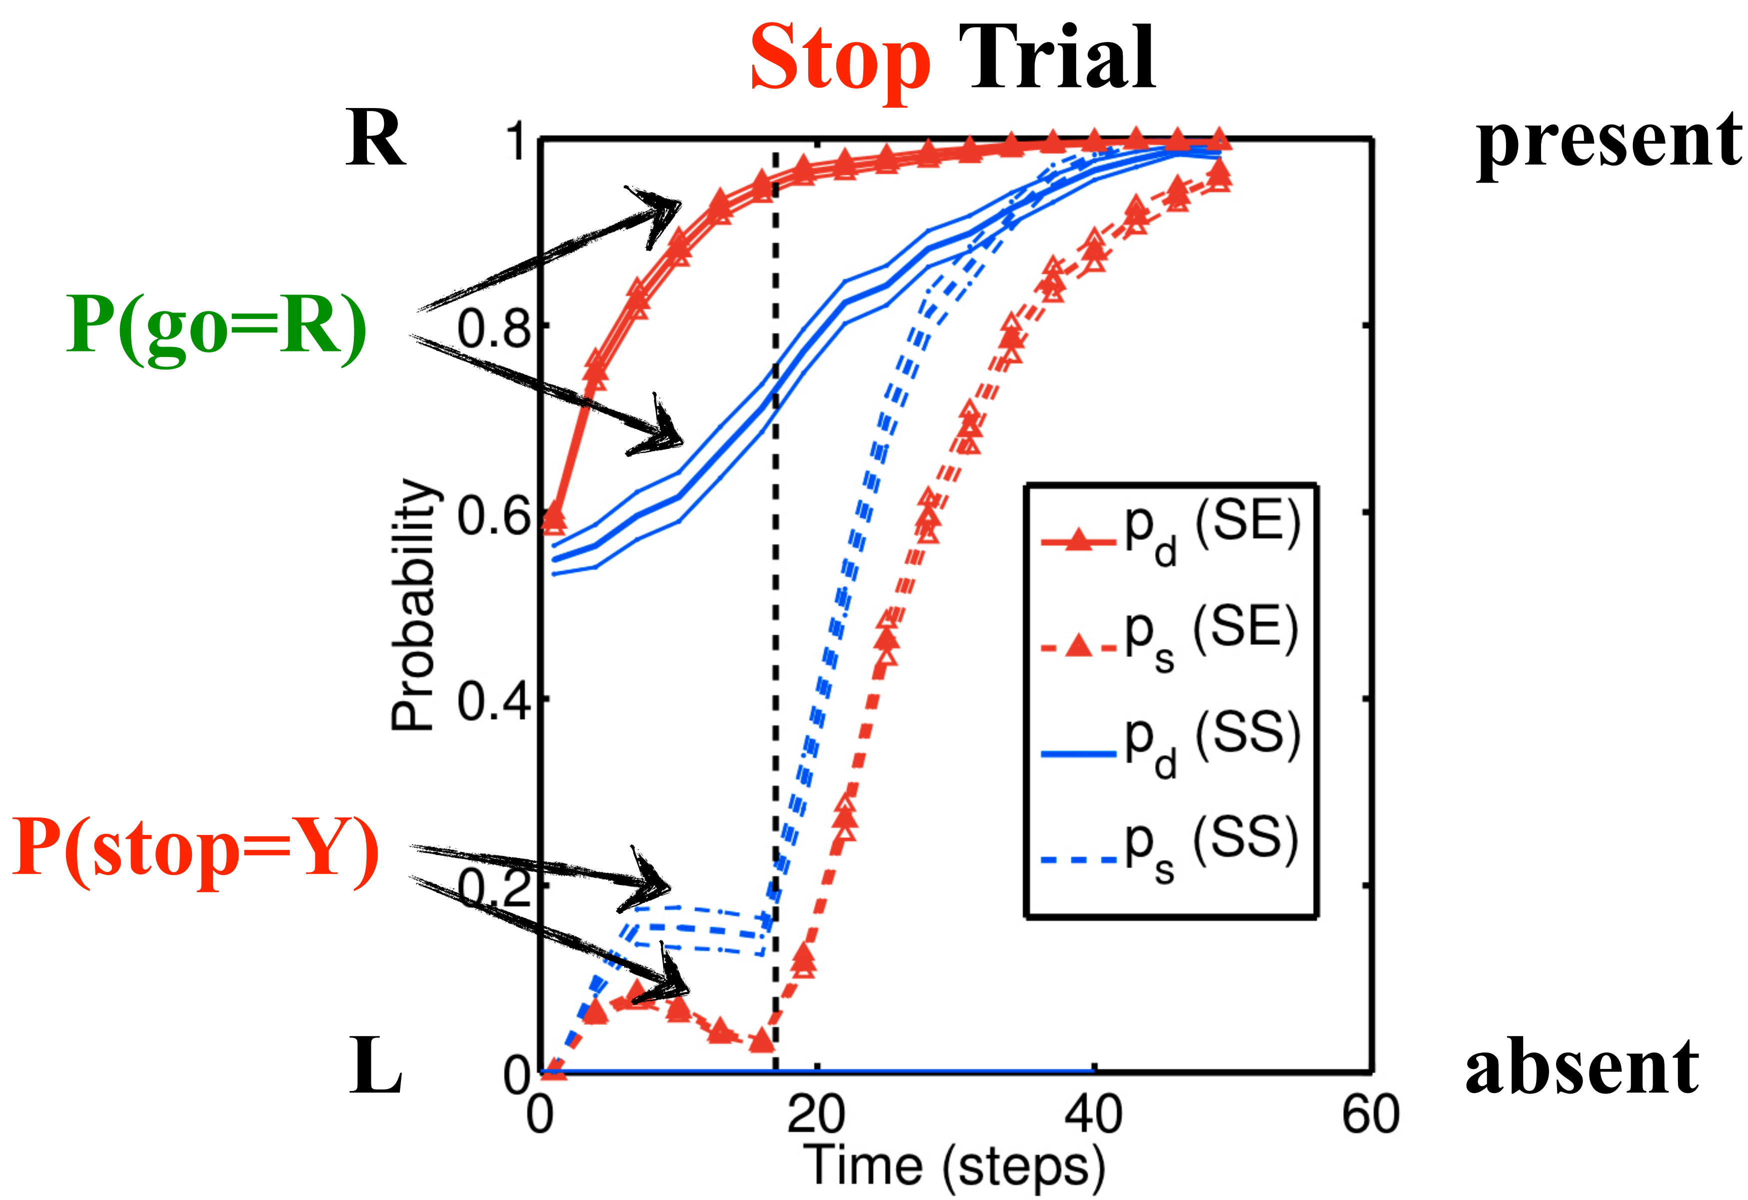

- stop-success: $p_d \uparrow$ slowly, $p_s \uparrow$ quickly
- stop-error: $p_d \uparrow$ quickly, $p_s \uparrow$ slowly

Simulation: Belief State Trajectories

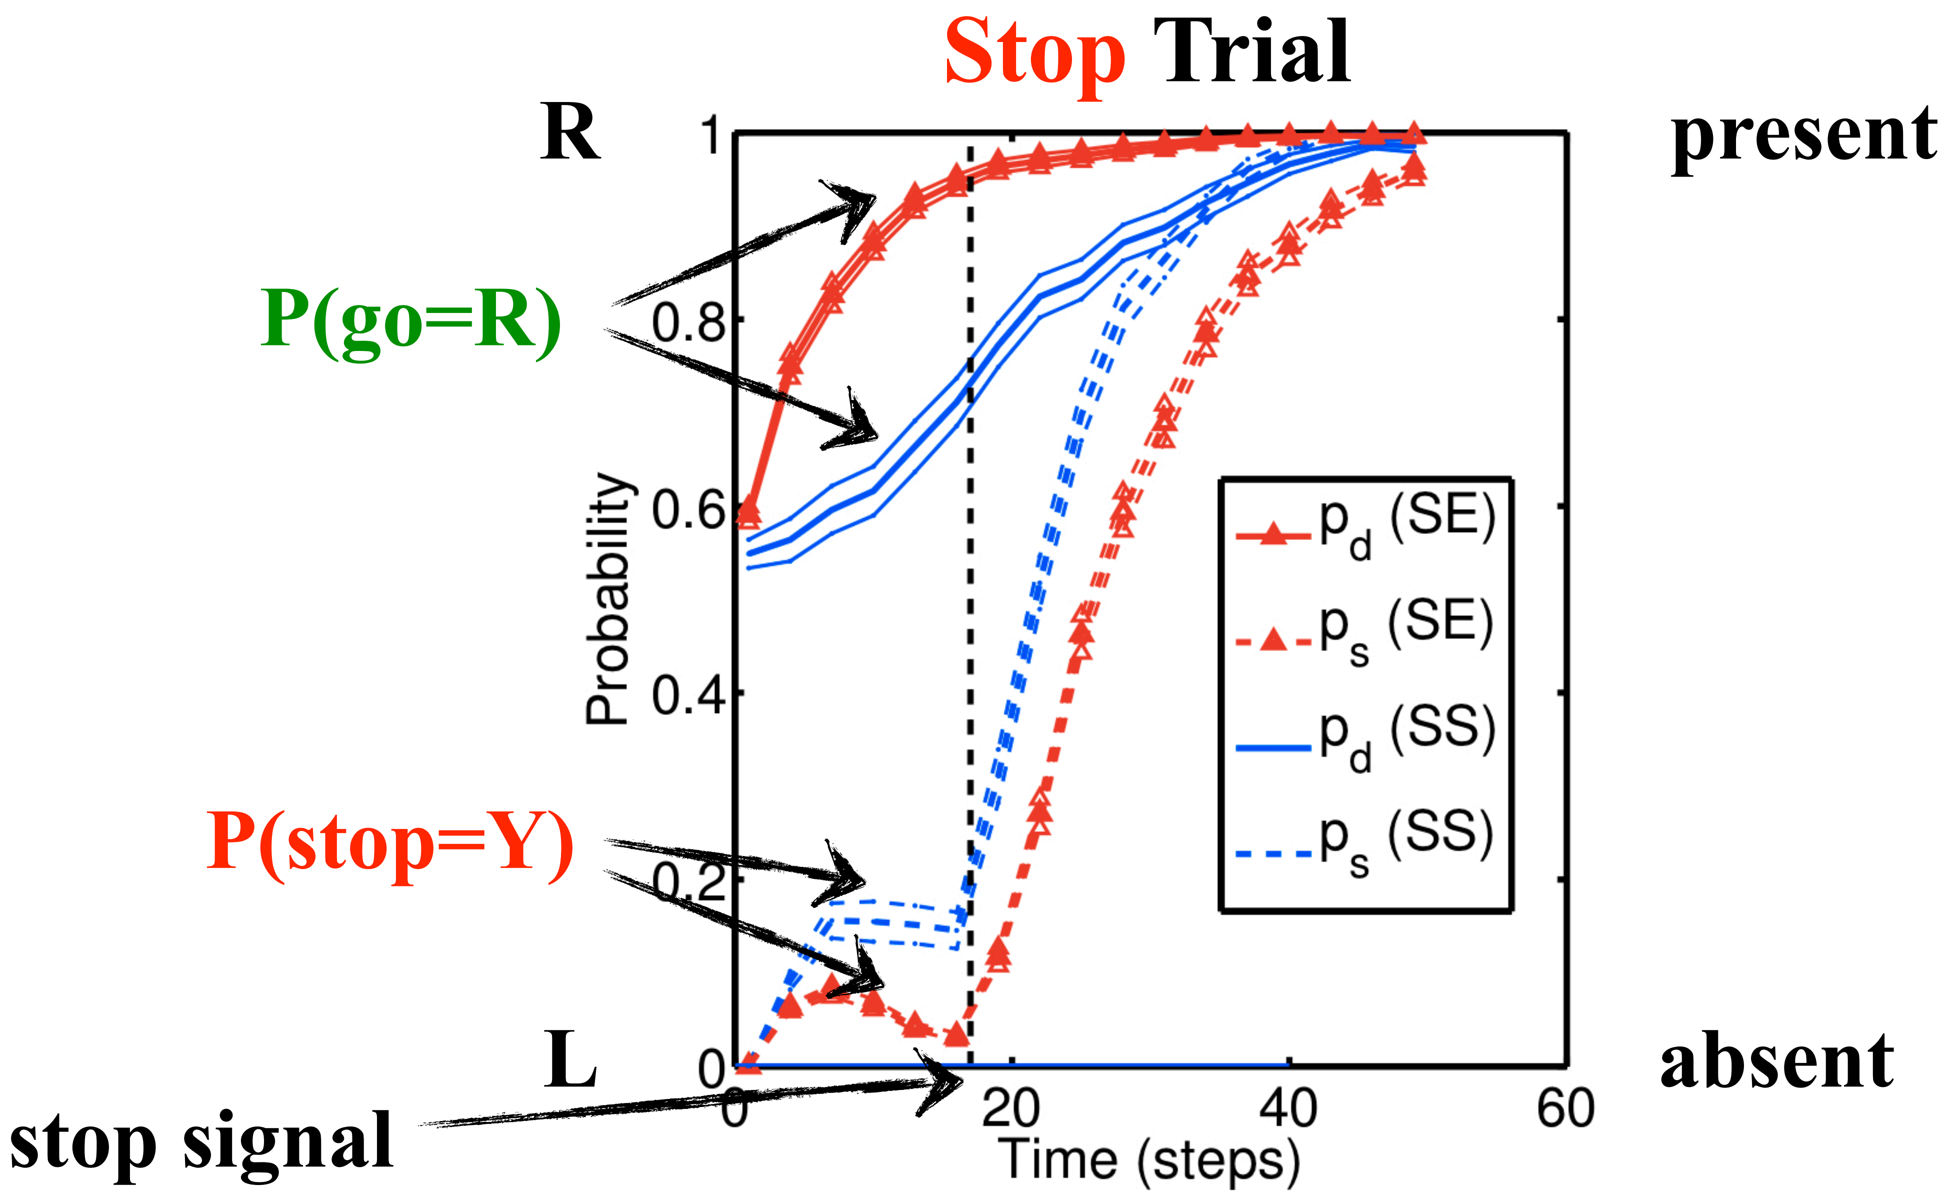

- stop-success: $p_d \uparrow$ slowly, $p_s \uparrow$ quickly
- stop-error: $p_d \uparrow$ quickly, $p_s \uparrow$ slowly

Action Selection = Stochastic Control

What is optimal? Define global cost function

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s = 1) + (1-r)P(\tau < D, \delta \neq d | s = 0) + (1-r)P(\tau = D | s = 0)$$

τ: response time *s*: stop trial
δ: chosen target *d*: true target
r: freq(stop trials) *D*: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s=1) + (1-r) P(\tau < D, \delta \neq d | s=0) + (1-r) P(\tau = D | s=0)$$

τ: response time   *s*: stop trial
δ: chosen target   *d*: true target
r: freq(stop trials)   *D*: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost

time cost

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s=1) + (1-r)P(\tau < D, \delta \neq d | s=0) + (1-r)P(\tau = D | s=0)$$

τ: response time   *s*: stop trial

δ: chosen target   *d*: true target

r: freq(stop trials)   *D*: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost

time cost

stop error (non-canceled)

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s = 1) + (1-r) P(\tau < D, \delta \neq d | s = 0) + (1-r) P(\tau = D | s = 0)$$

τ: response time	s: stop trial
δ: chosen target	d: true target
r: freq(stop trials)	D: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost   *time cost*   *stop error (non-canceled)*   *go error (wrong response)*

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s = 1) + (1 - r)P(\tau < D, \delta \neq d | s = 0) + (1 - r)P(\tau = D | s = 0)$$

τ: response time   *s*: stop trial
δ: chosen target   *d*: true target
r: freq(stop trials)   *D*: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost ***time cost*** ***stop error (non-canceled)*** ***go error (wrong response)*** ***go error (deadline)***

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s = 1) + (1-r)P(\tau < D, \delta \neq d | s = 0) + (1-r)P(\tau = D | s = 0)$$

τ: response time	s: stop trial
δ: chosen target	d: true target
r: freq(stop trials)	D: deadline

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost ***time cost*** ***stop error (non-canceled)*** ***go error (wrong response)*** ***go error (deadline)***

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D|s=1) + (1-r)P(\tau < D, \delta \neq d|s=0) + (1-r)P(\tau = D|s=0)$$

τ: response time   *s*: stop trial
δ: chosen target   *d*: true target
r: freq(stop trials)   *D*: deadline

Policy: $x_1, \ldots, x_t \Rightarrow \{$*left*, *right*, *wait*$\}$

Action Selection = Stochastic Control

What is optimal? Define global cost function

expected cost *time cost* *stop error (non-canceled)* *go error (wrong response)* *go error (deadline)*

$$L_\pi = c\langle\tau\rangle + c_s r P(\tau < D | s = 1) + (1-r)P(\tau < D, \delta \neq d | s = 0) + (1-r)P(\tau = D | s = 0)$$

τ: response time   *s*: stop trial
δ: chosen target   *d*: true target
r: freq(stop trials)   *D*: deadline

Policy: $x_1, \ldots, x_t \Rightarrow$ {*left*, *right*, *wait*}

Objective: minimize expected (average) cost

Optimal Policy: Compare *Go* & *Wait* Costs

$$\mathbf{b}^t = (p_d^t, p_s^t)$$

Optimal Policy: Compare *Go* & *Wait* Costs

Bellman's Dynamic Programming Principle (Bellman, 1952)

Optimal policy: repeatedly choose best (least costly) action

$$V^t(\mathbf{b}^t) = \min\left(Q_g^t(\mathbf{b}^t), Q_w^t(\mathbf{b}^t)\right) \qquad \mathbf{b}^t = (p_d^t, p_s^t)$$

Optimal Policy: Compare *Go* & *Wait* Costs

Bellman's Dynamic Programming Principle (Bellman, 1952)

Optimal policy: repeatedly choose best (least costly) action

$$V^t(\mathbf{b}^t) = \min\left(Q_g^t(\mathbf{b}^t), Q_w^t(\mathbf{b}^t)\right) \qquad \mathbf{b}^t = (p_d^t, p_s^t)$$

Cost of Go action = time cost + stop error + go error (wrong response)

$$Q_g^t(\mathbf{b}^t) = ct + c_s p_s^t + (1 - p_s^t)\min(p_d^t, 1 - p_d^t)$$

Optimal Policy: Compare *Go* & *Wait* Costs

Bellman's Dynamic Programming Principle (Bellman, 1952)

Optimal policy: repeatedly choose best (least costly) action

$$V^t(\mathbf{b}^t) = \min\left(Q_g^t(\mathbf{b}^t), Q_w^t(\mathbf{b}^t)\right) \qquad \mathbf{b}^t = (p_d^t, p_s^t)$$

Cost of Go action = time cost + stop error + go error (wrong response)

$$Q_g^t(\mathbf{b}^t) = ct + c_s p_s^t + (1 - p_s^t)\min(p_d^t, 1 - p_d^t)$$

Cost of wait action = expected future cost or deadline penalty

$$Q_w^t(\mathbf{b}^t) = \begin{cases} \langle V^{t+1}(\mathbf{b}^{t+1}|\mathbf{b}^t\rangle, & D > t+1 \\ c(t+1) + (1 - p_s^t), & D = t+1 \end{cases}$$

Optimal Policy: Belief State ⇒ *Go* & *Wait* Regions

Policy: $x_1, \ldots, x_t \Rightarrow \{$*left*, *right*, *wait*$\}$

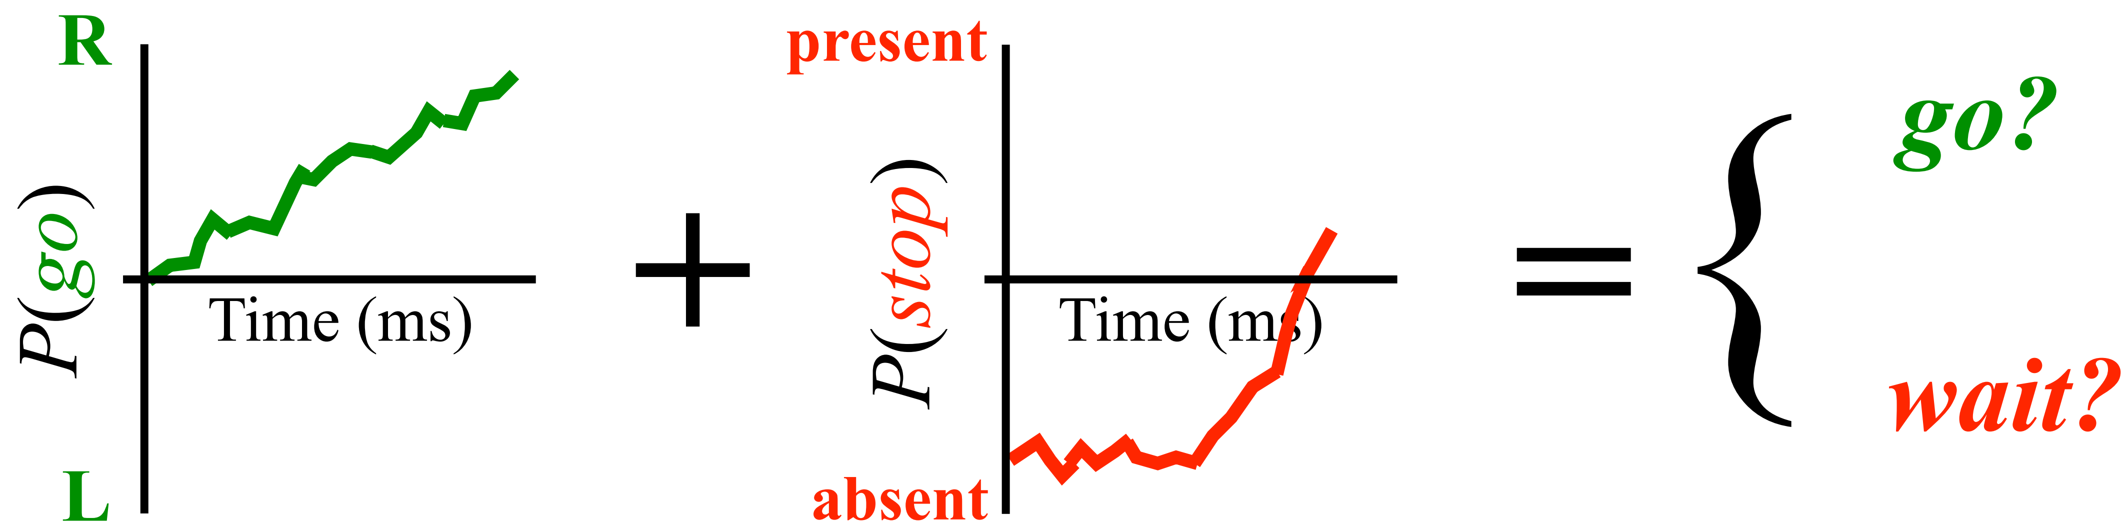

Optimal Policy: Belief State ⇒ *Go* & *Wait* Regions

Policy: $x_1, \ldots, x_t \Rightarrow \{left,\ right,\ wait\}$

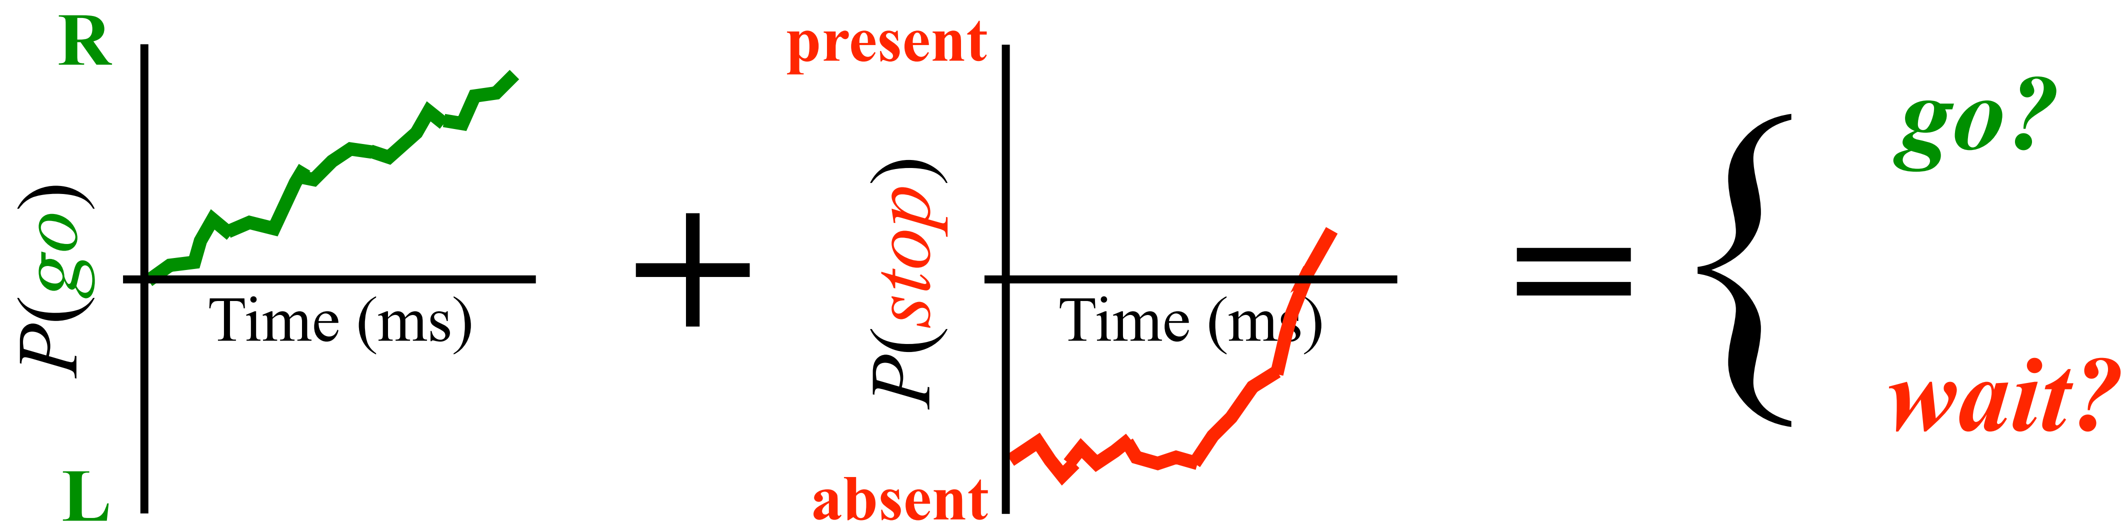

***Go* & *wait* regions**

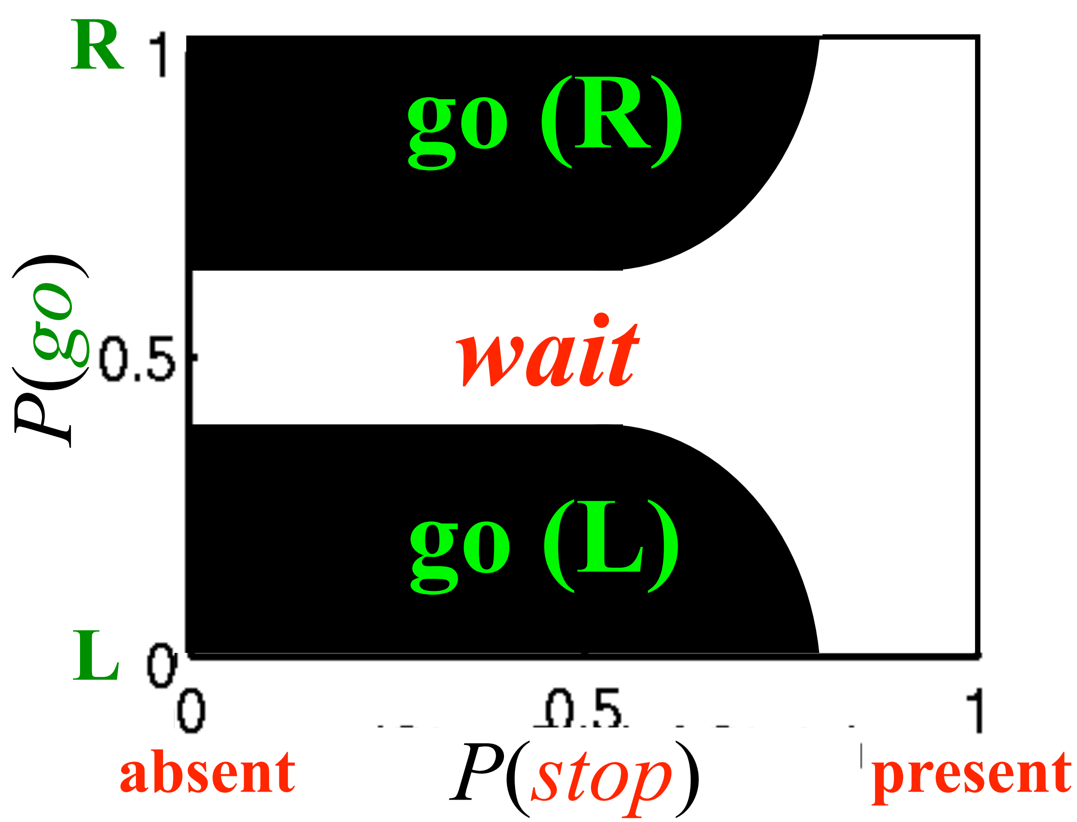

Optimal Policy: Belief State ⇒ *Go* & *Wait* Regions

Policy: $x_1, \ldots, x_t \Rightarrow \{left, right, wait\}$

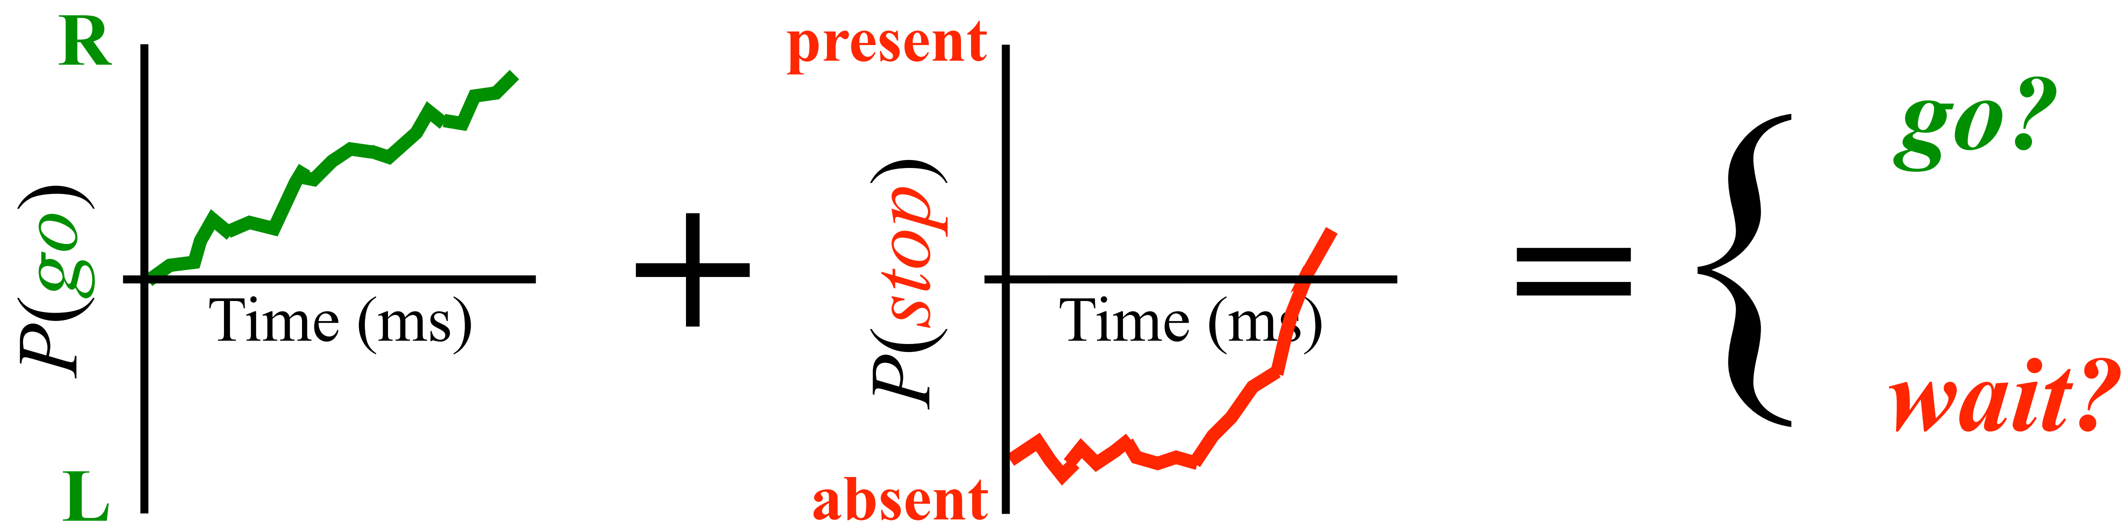

***Go* & *wait* regions**

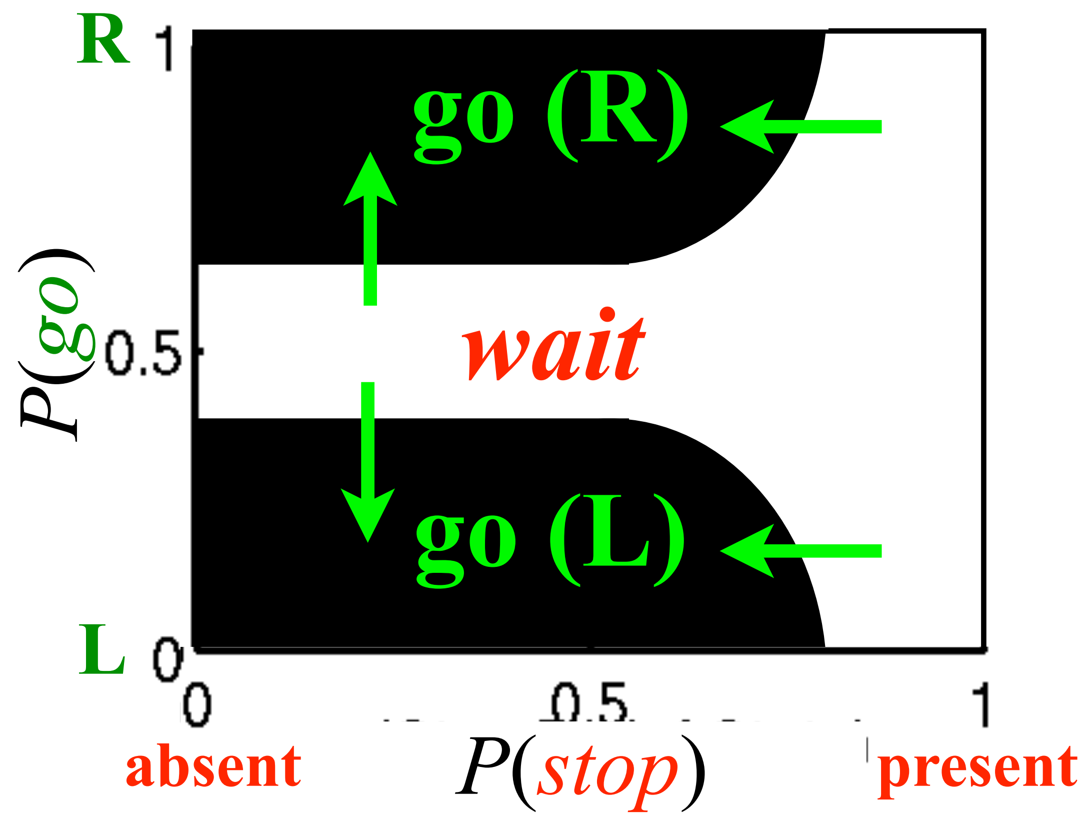

Simulation: *Go* Cost vs. *Wait* Cost

stop signal

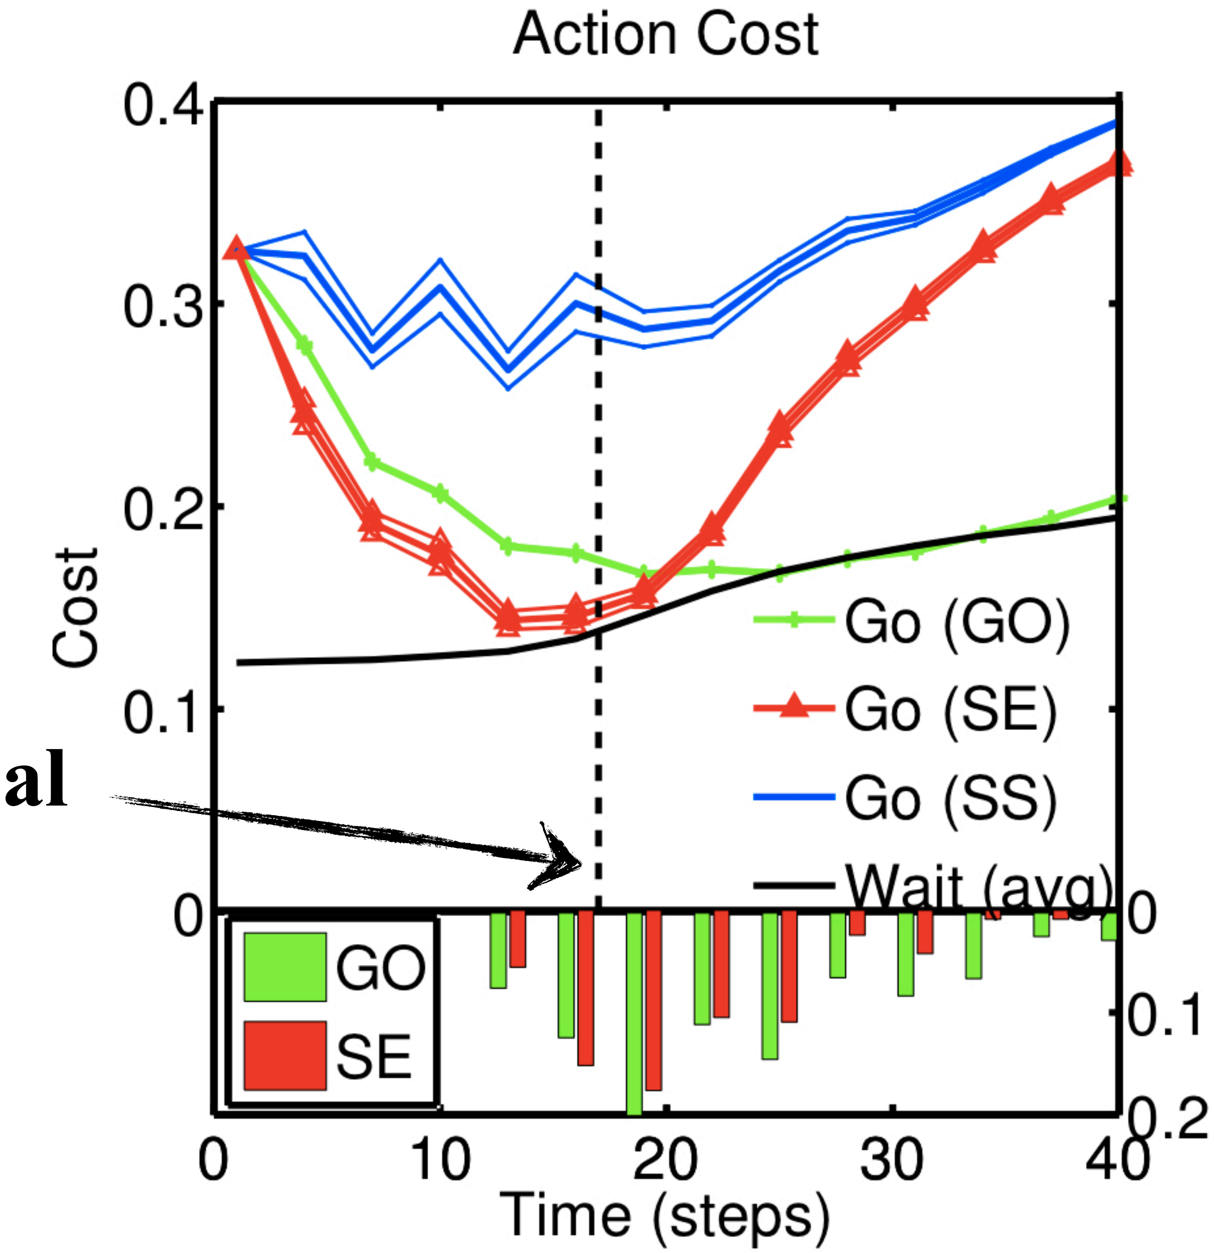

Simulation: *Go* Cost vs. *Wait* Cost

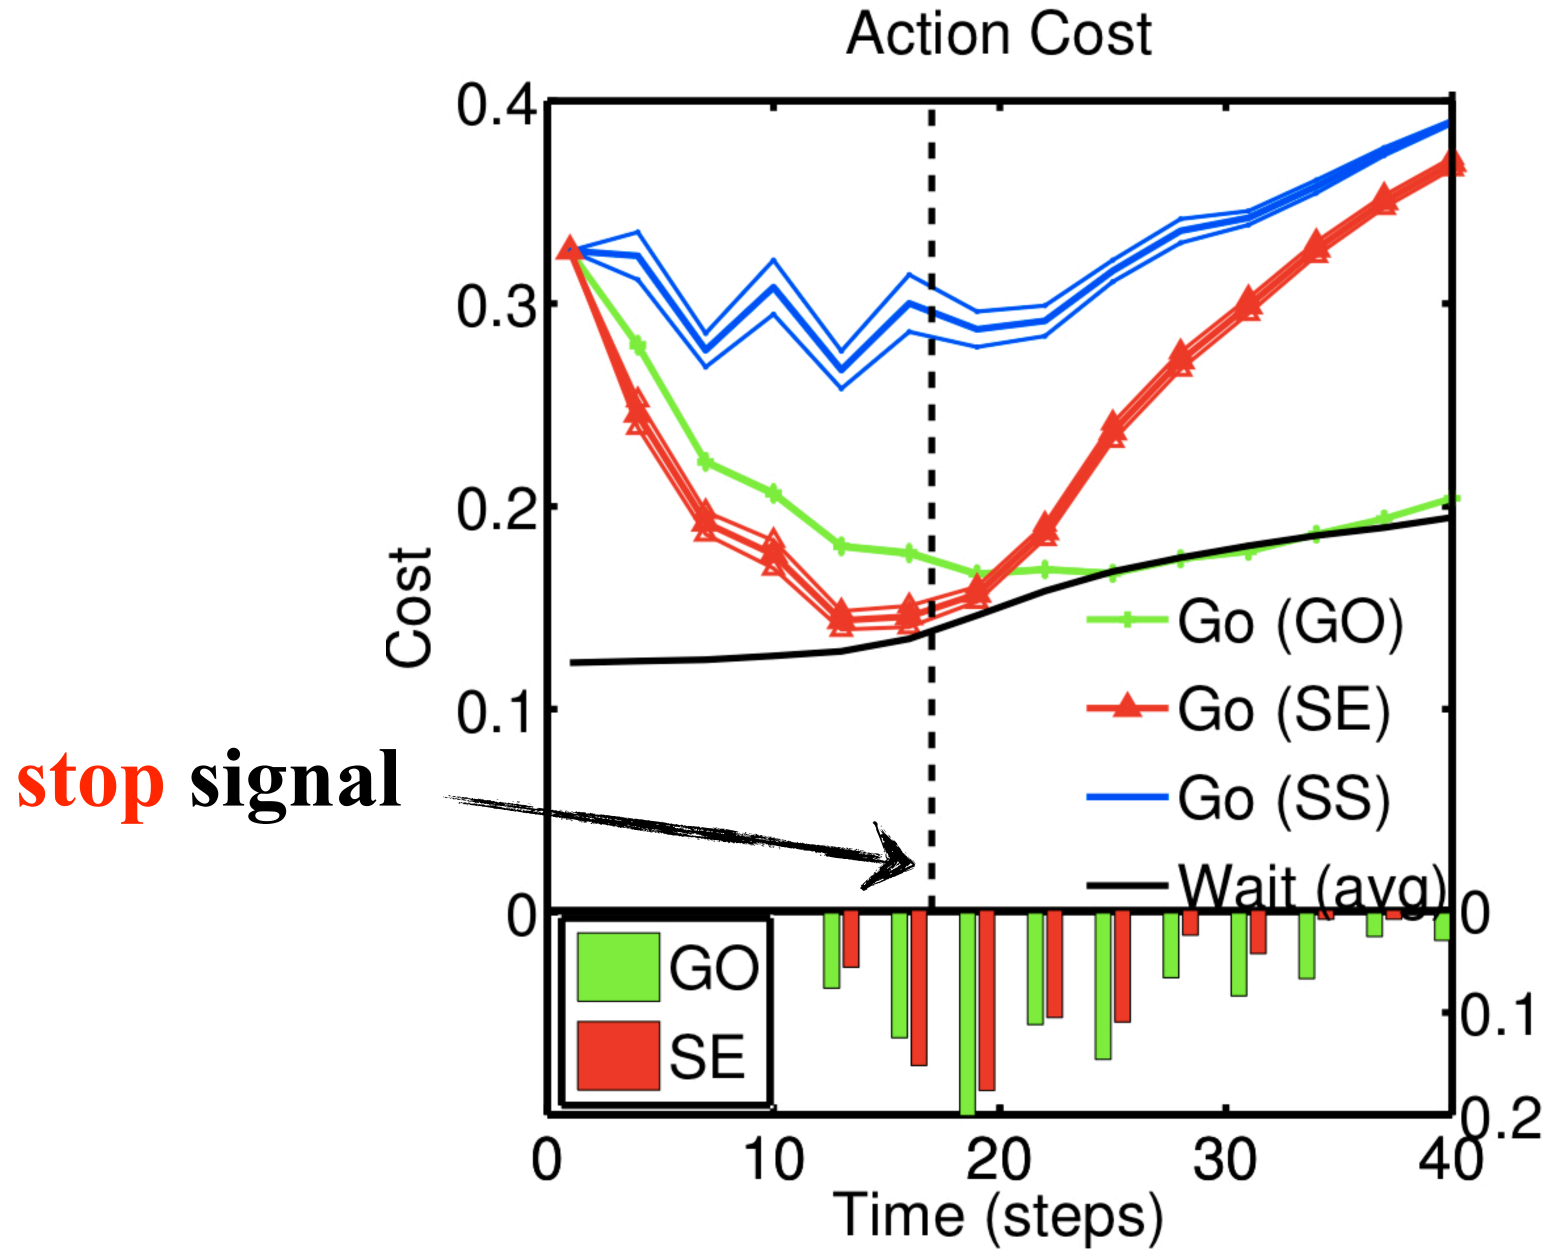

- Q(go) decreases as go stimulus becomes less ambiguous
- Q(go) increases after stop-signal onset
- Q(go) dips below Q(wait) $\Rightarrow$ go response, otherwise wait

Outline

- **Model: brain implements rational (optimal) computations**
  - Monitoring process - Bayesian inference
  - Decision process - optimal stochastic control
- **Model captures a range of behavioral results**
  - Classical results
  - Reward/motivation
  - Contextual effects, sequential effects
- **Neural implementation**
  - Race (drift-diffusion) model neurally plausible approximation

Classical Behavioral Results

Data: RT

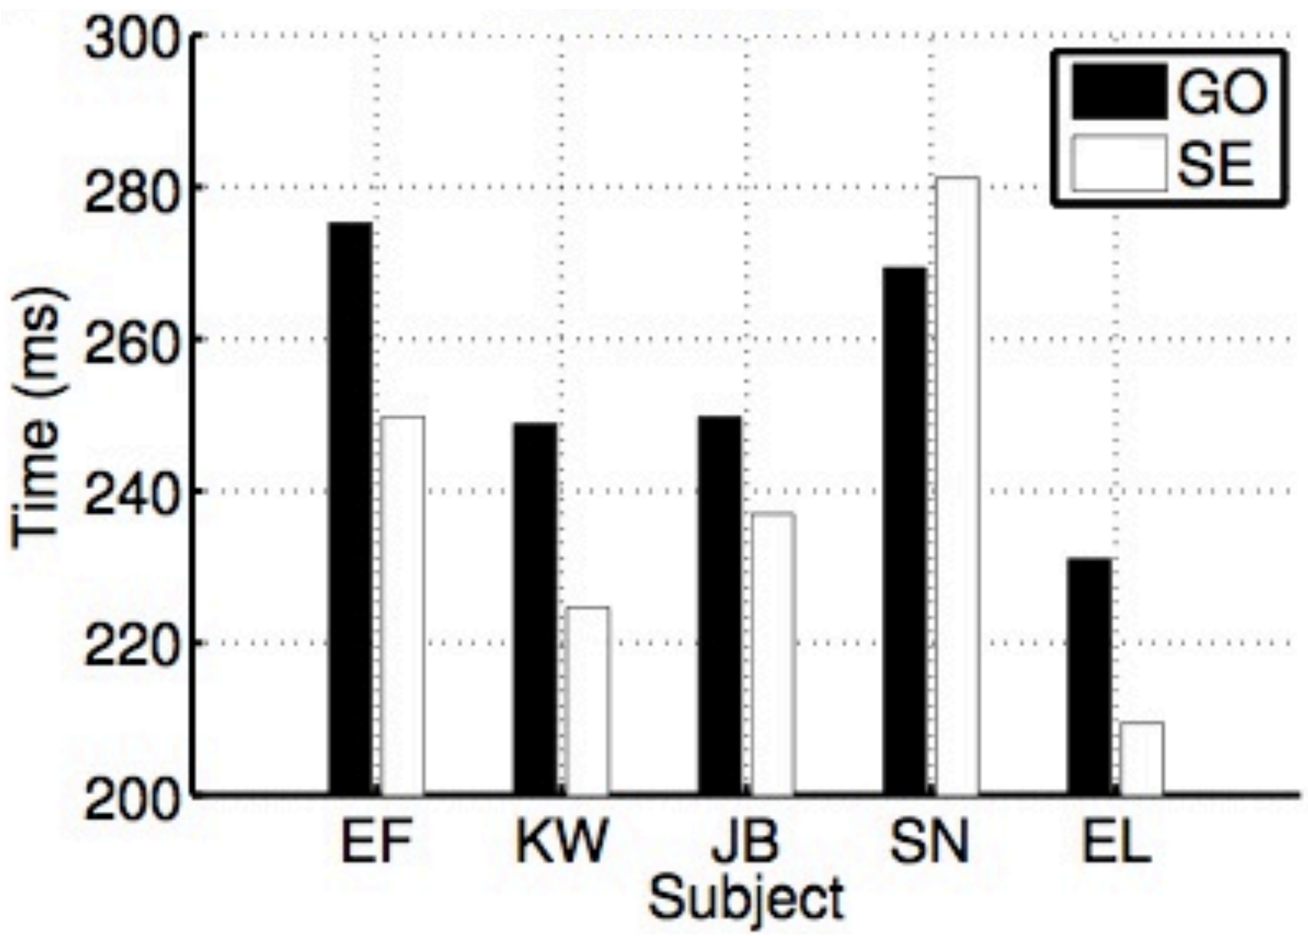

(from Emeric *et al*., 2007)

Model: RT

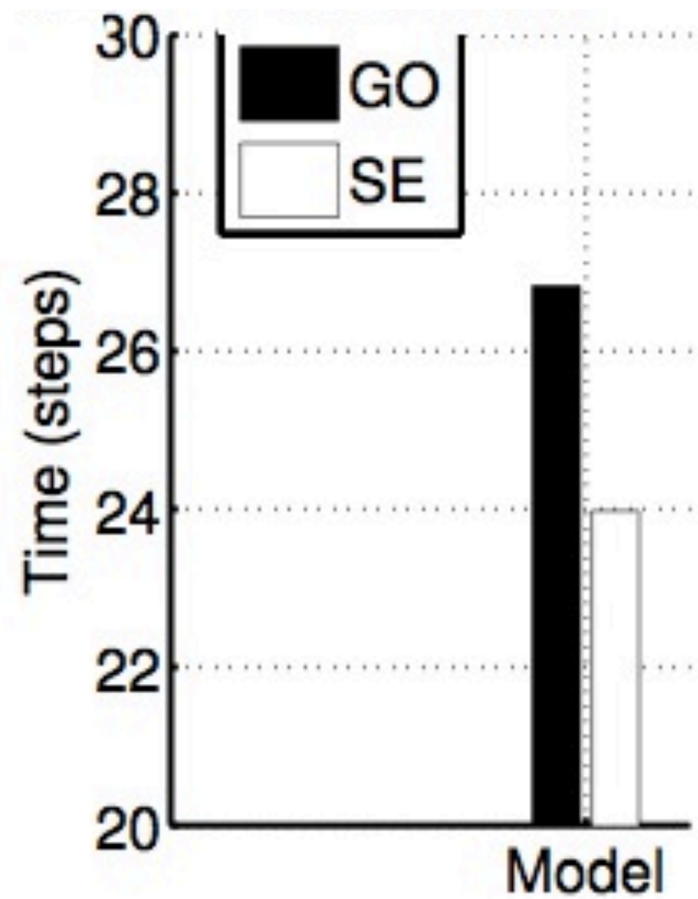

Classical Behavioral Results

Data: RT

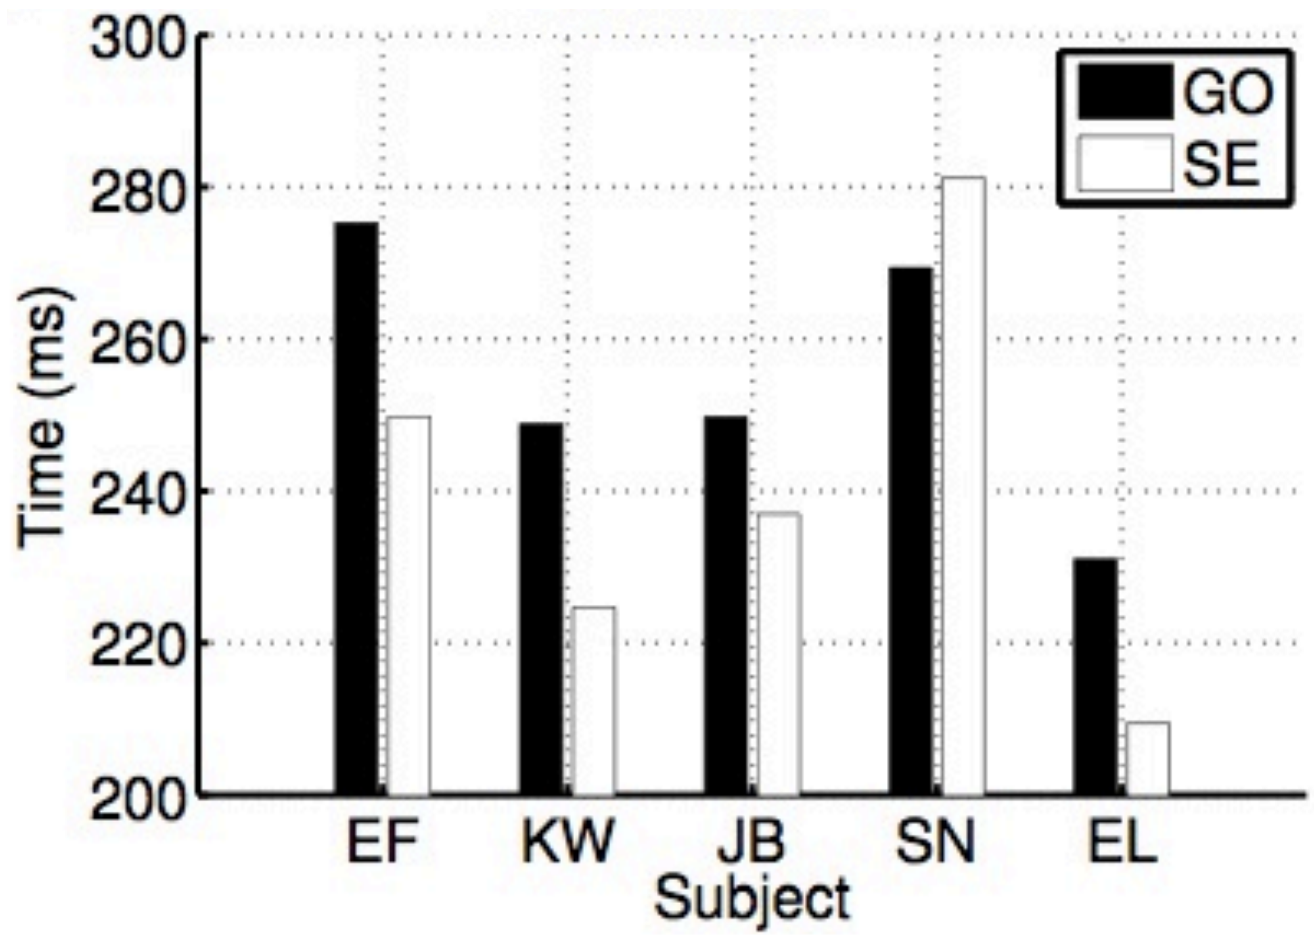

(from Emeric *et al*., 2007)

Model: RT

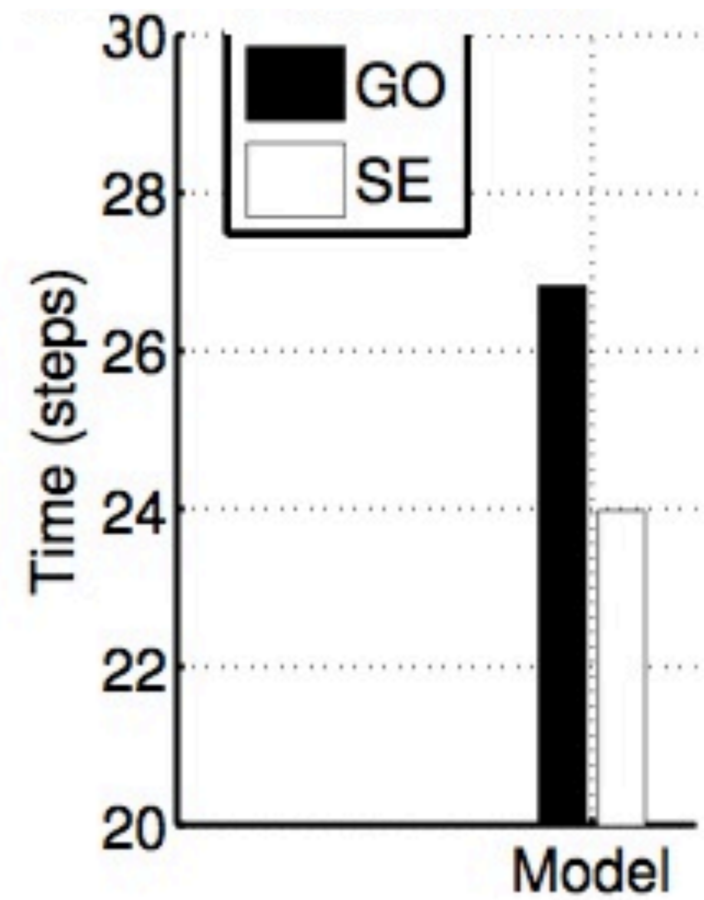

- non-canceled SE RT shorter than go RT
- Q(go) needs to dip below Q(wait) early enough to elicit response

Classical Behavioral Results

Data: error vs. SSD

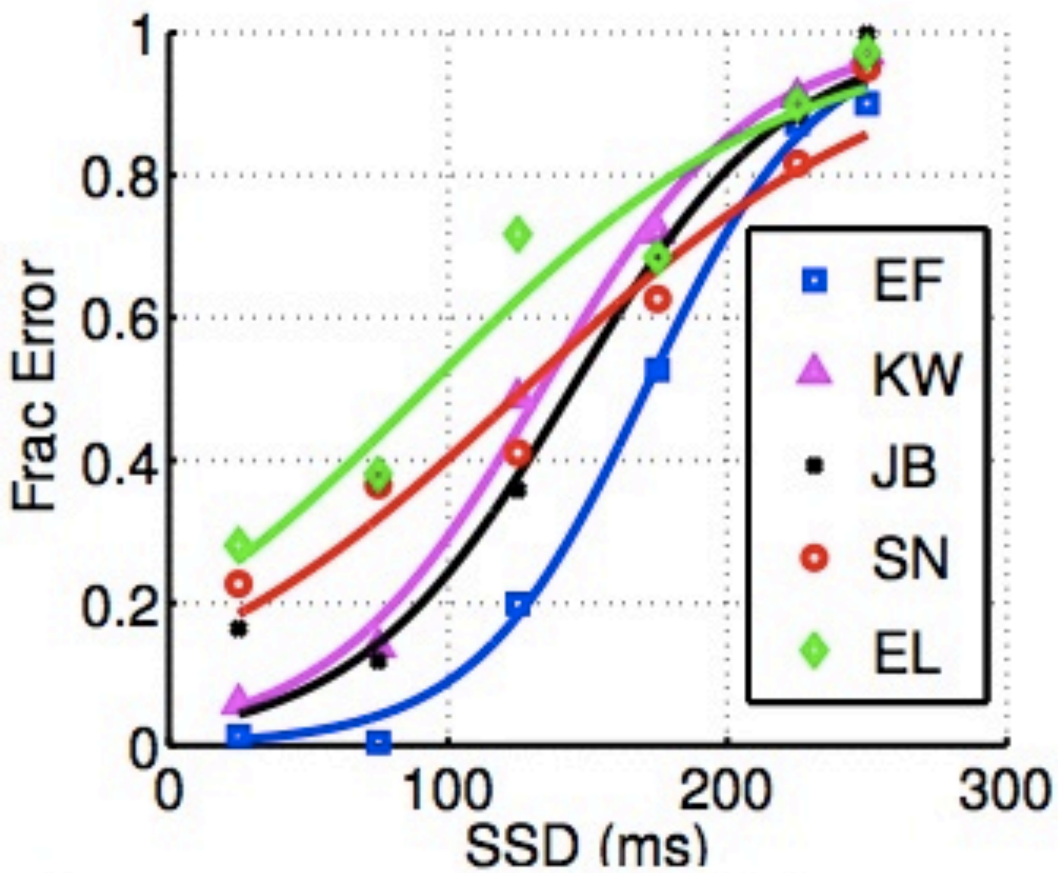

(from Emric *et al.*, 2007)

Model: error vs. SSD

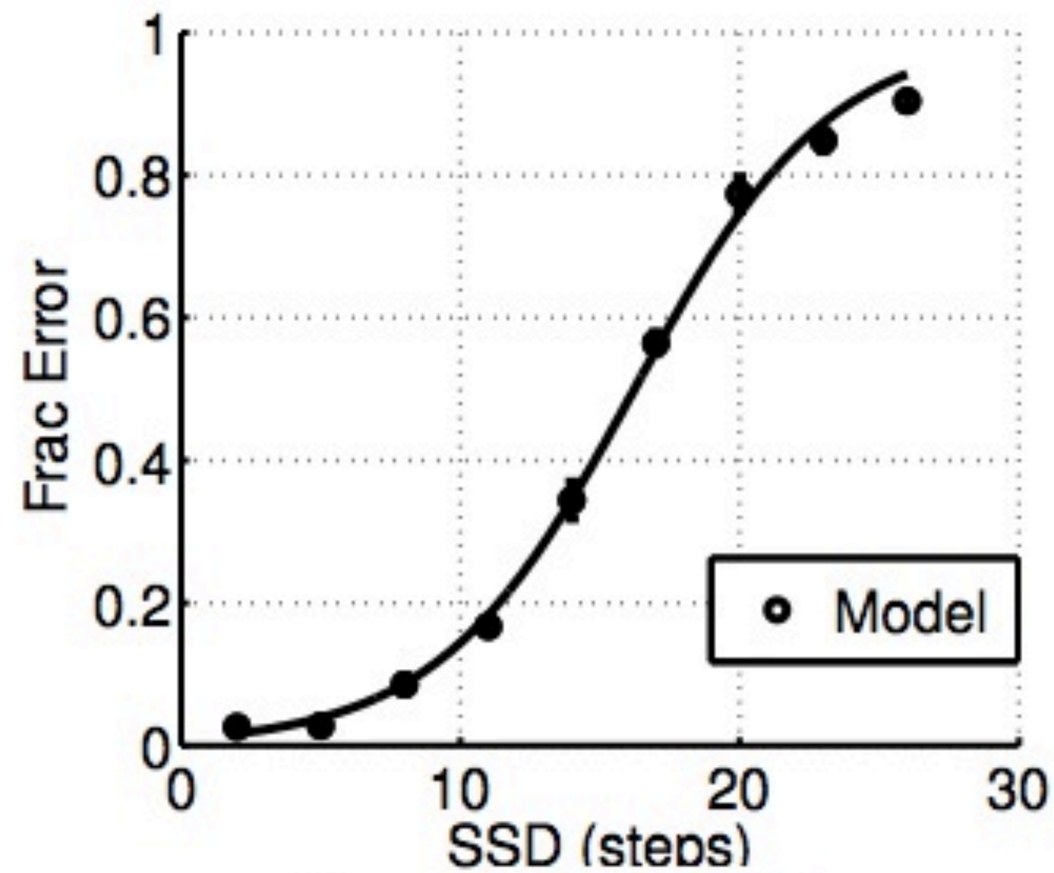

Classical Behavioral Results

Data: error vs. SSD

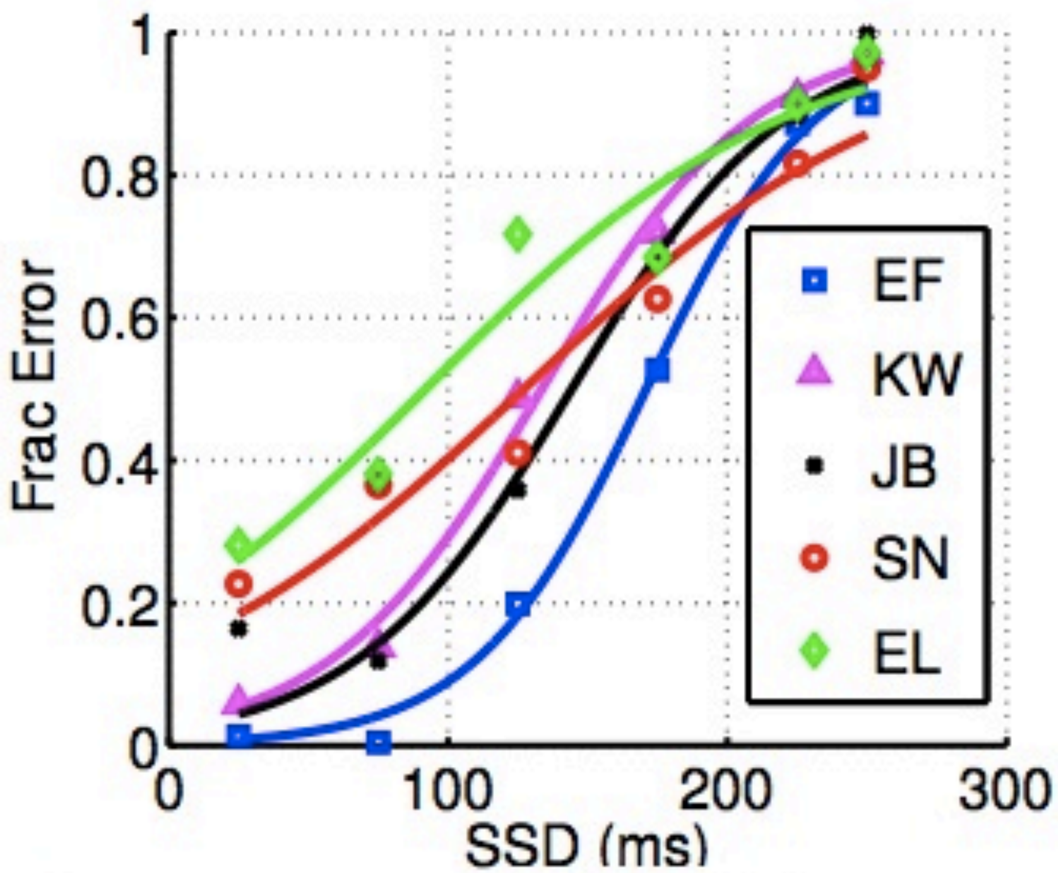

(from Emric *et al.*, 2007)

Model: error vs. SSD

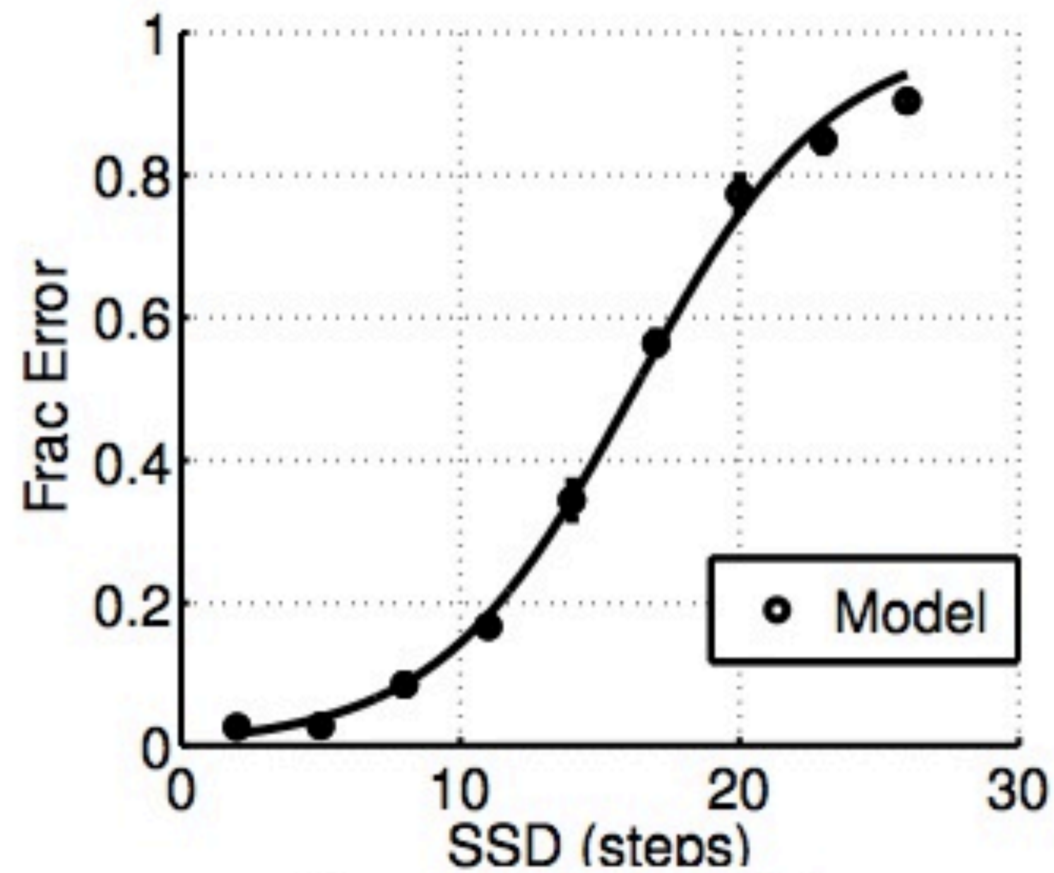

- Longer stop signal delay results in more errors
- More likely Q(go) has already dipped below Q(wait)

Reward/Motivation ⇒ Stopping Behavior

Data: (from Leotti & Wager, 2009)

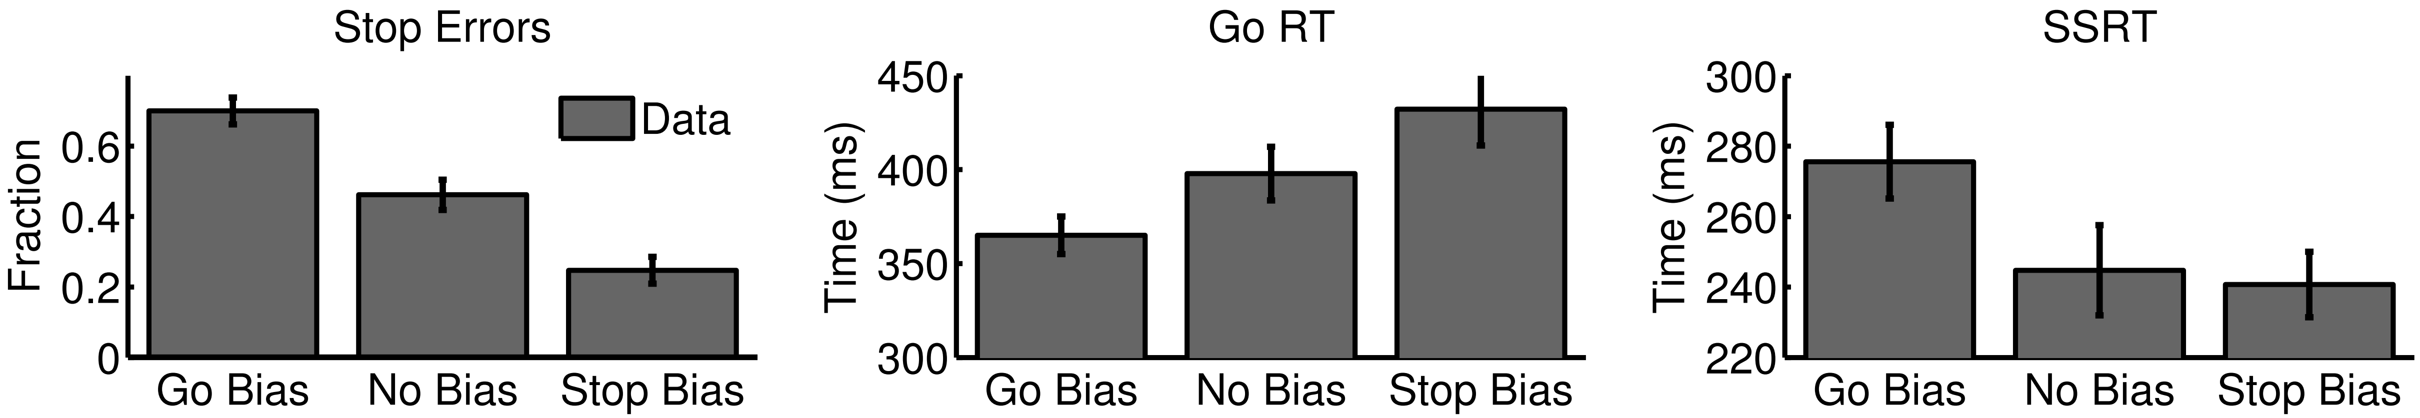

Reward/Motivation ⇒ Stopping Behavior

Data: (from Leotti & Wager, 2009)

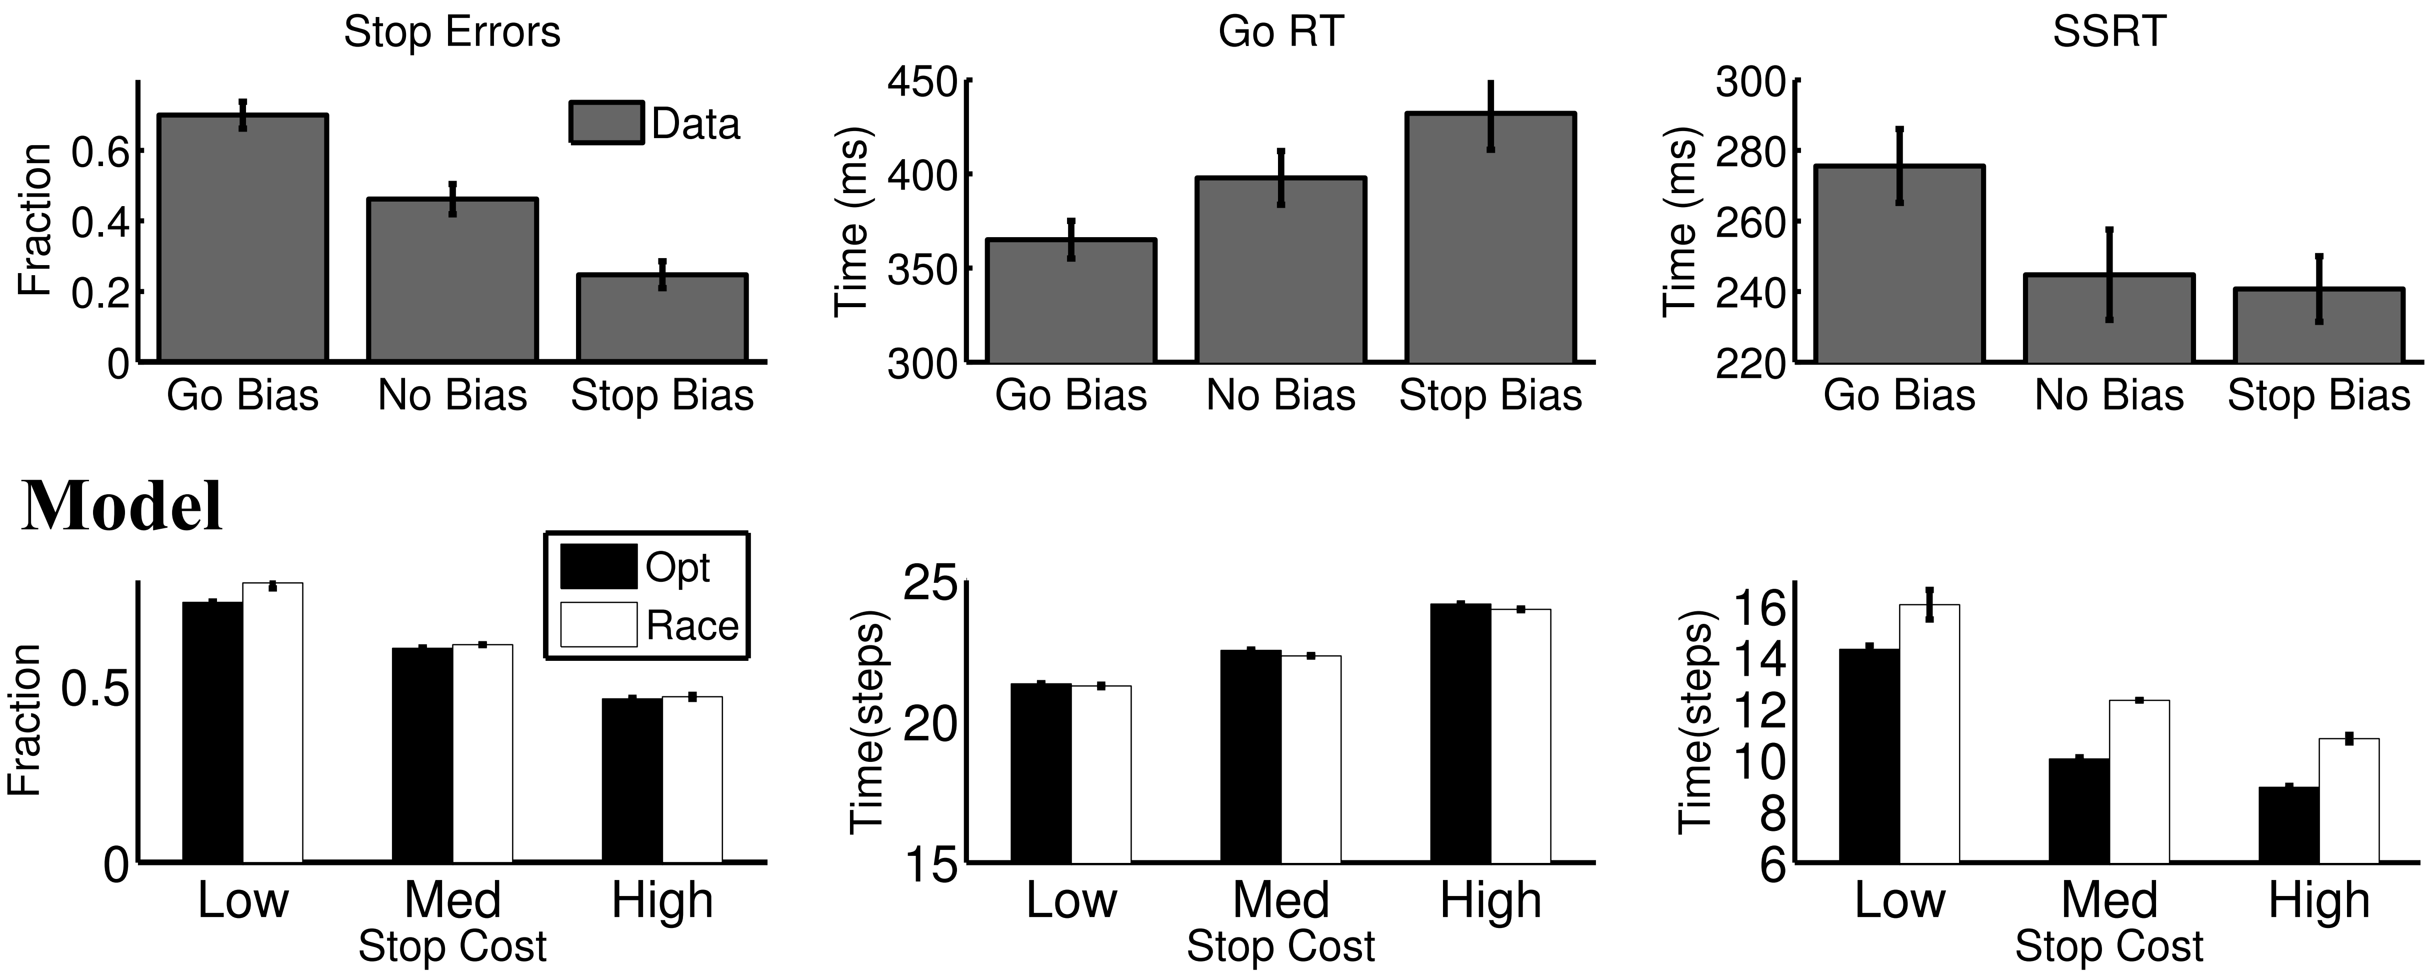

Reward/Motivation ⇒ Stopping Behavior

Data: (from Leotti & Wager, 2009)

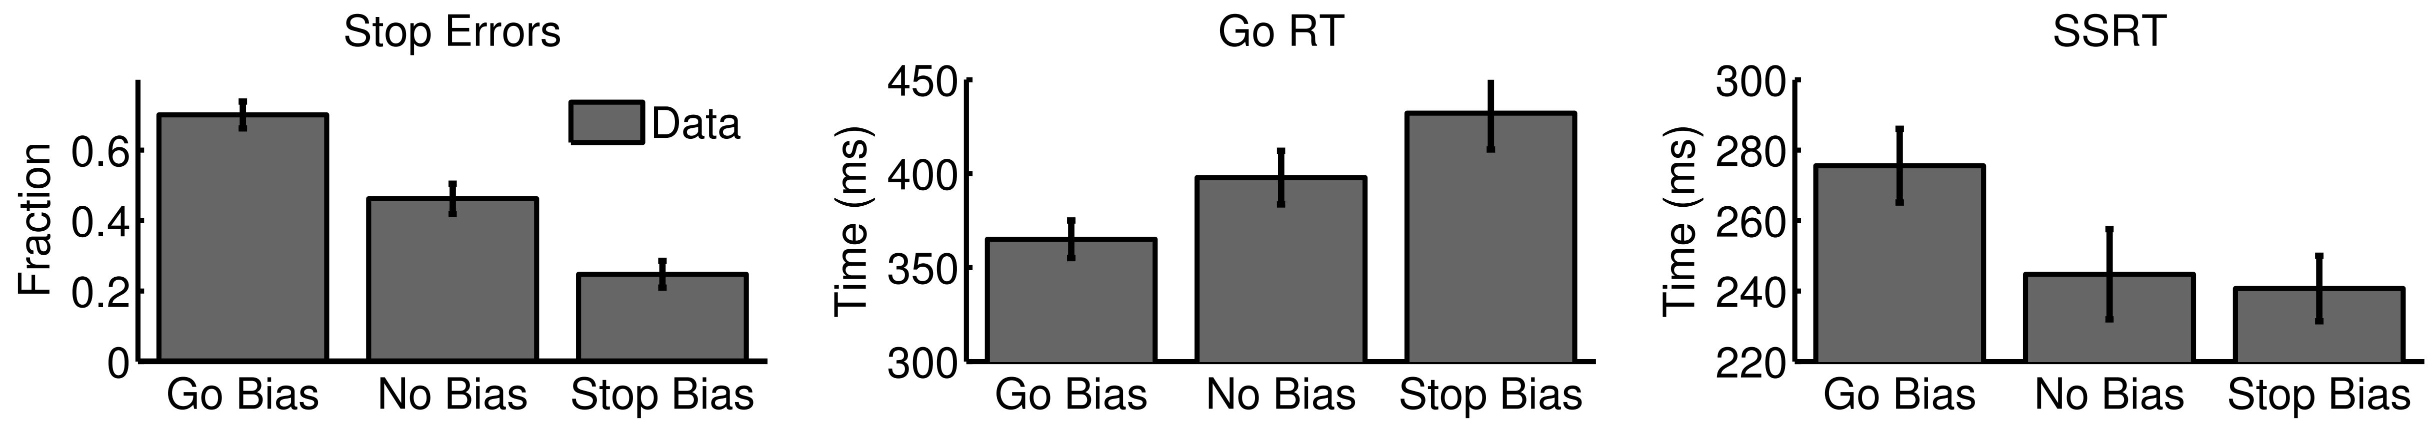

Model

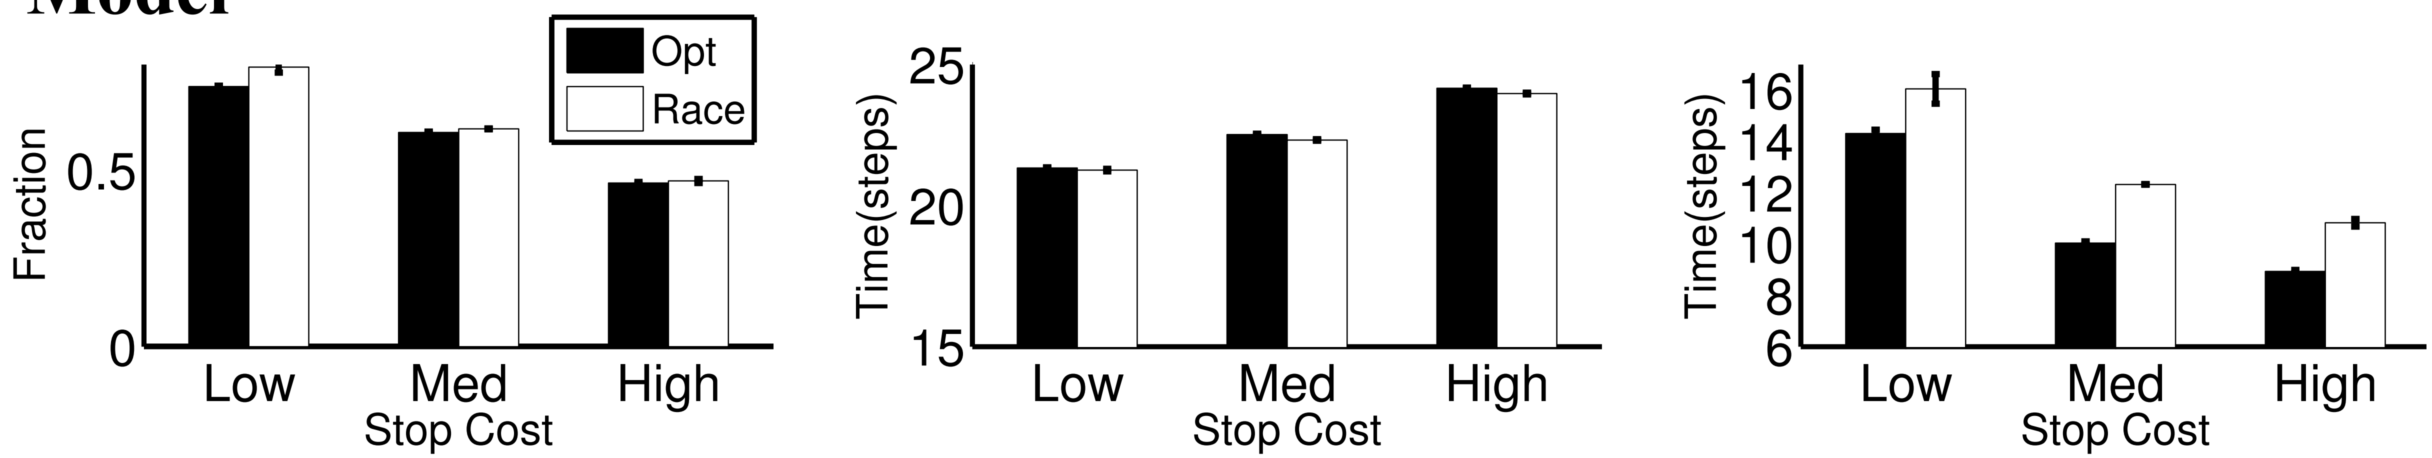

Race model as approximation to optimal decision-making

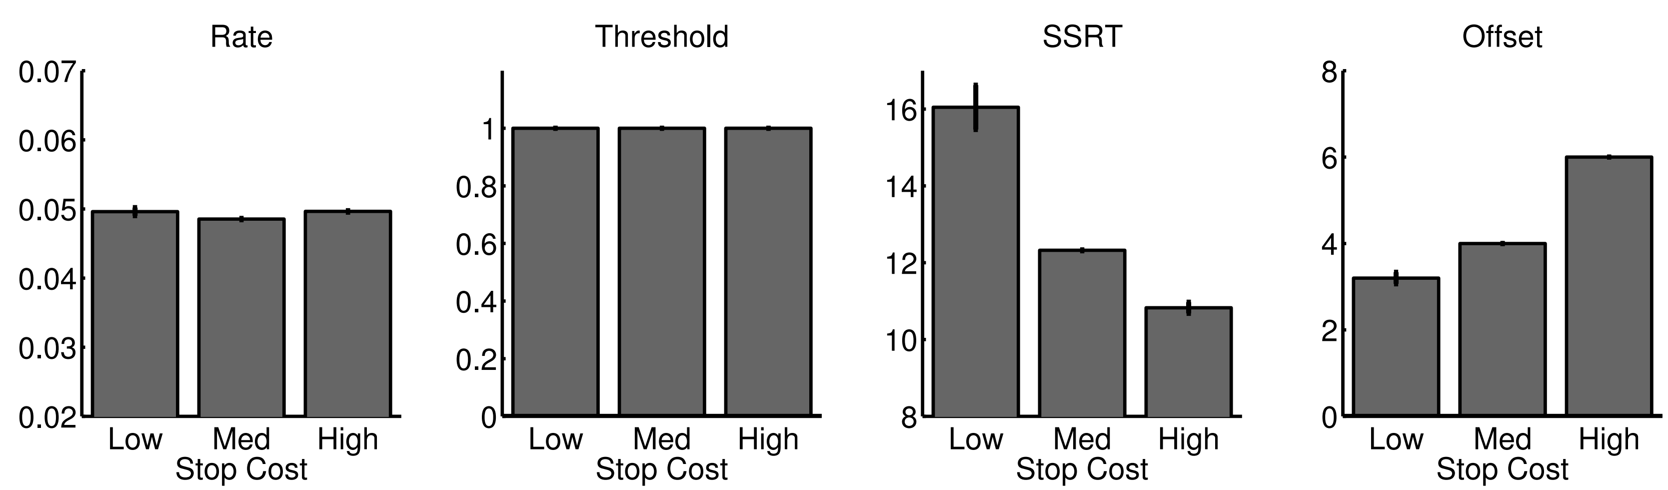

Stimulus Statistics ⇒ Stopping Behavior

(Emeric *et al.*, 2007)

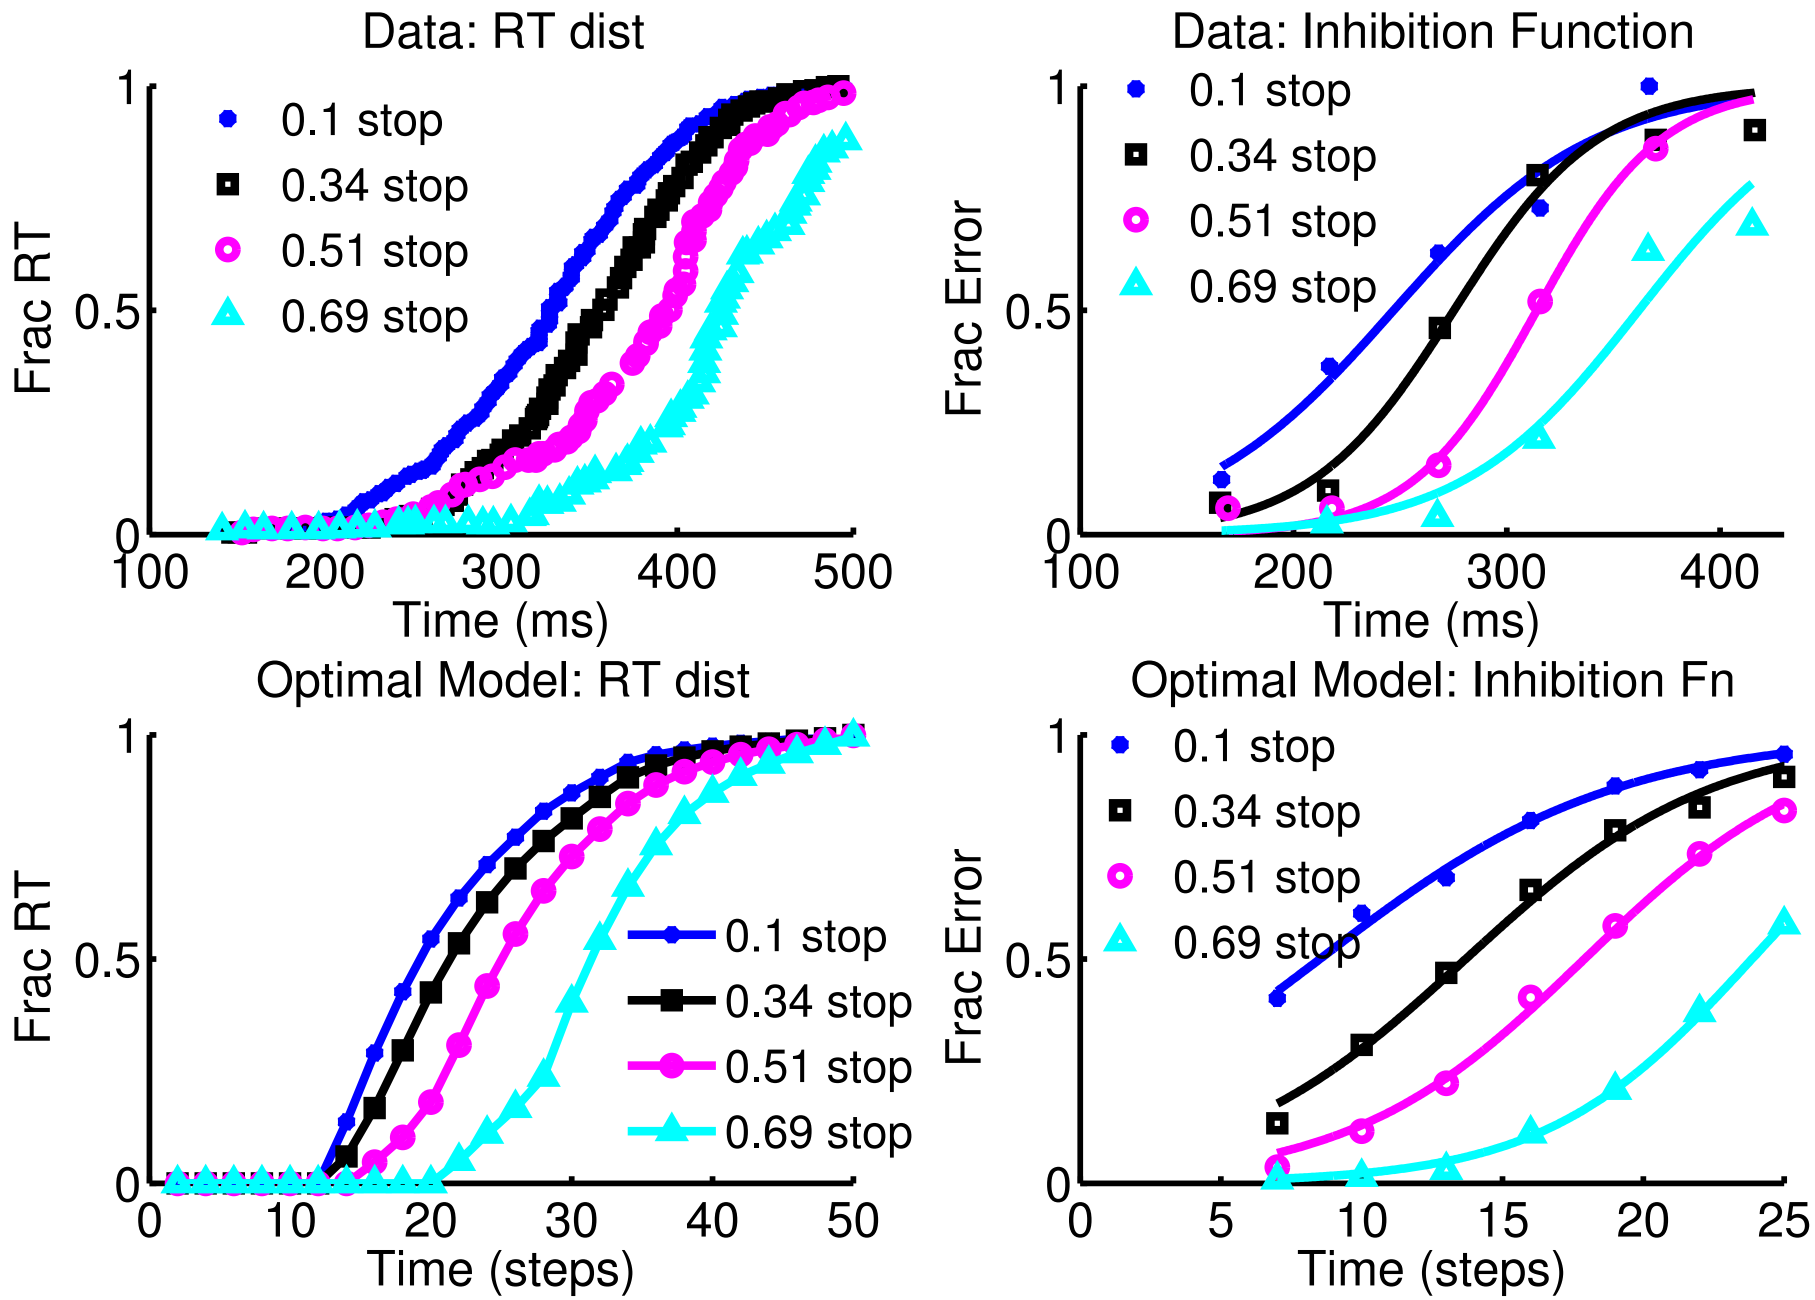

Stimulus Statistics ⇒ Stopping Behavior

(Emeric *et al.*, 2007)

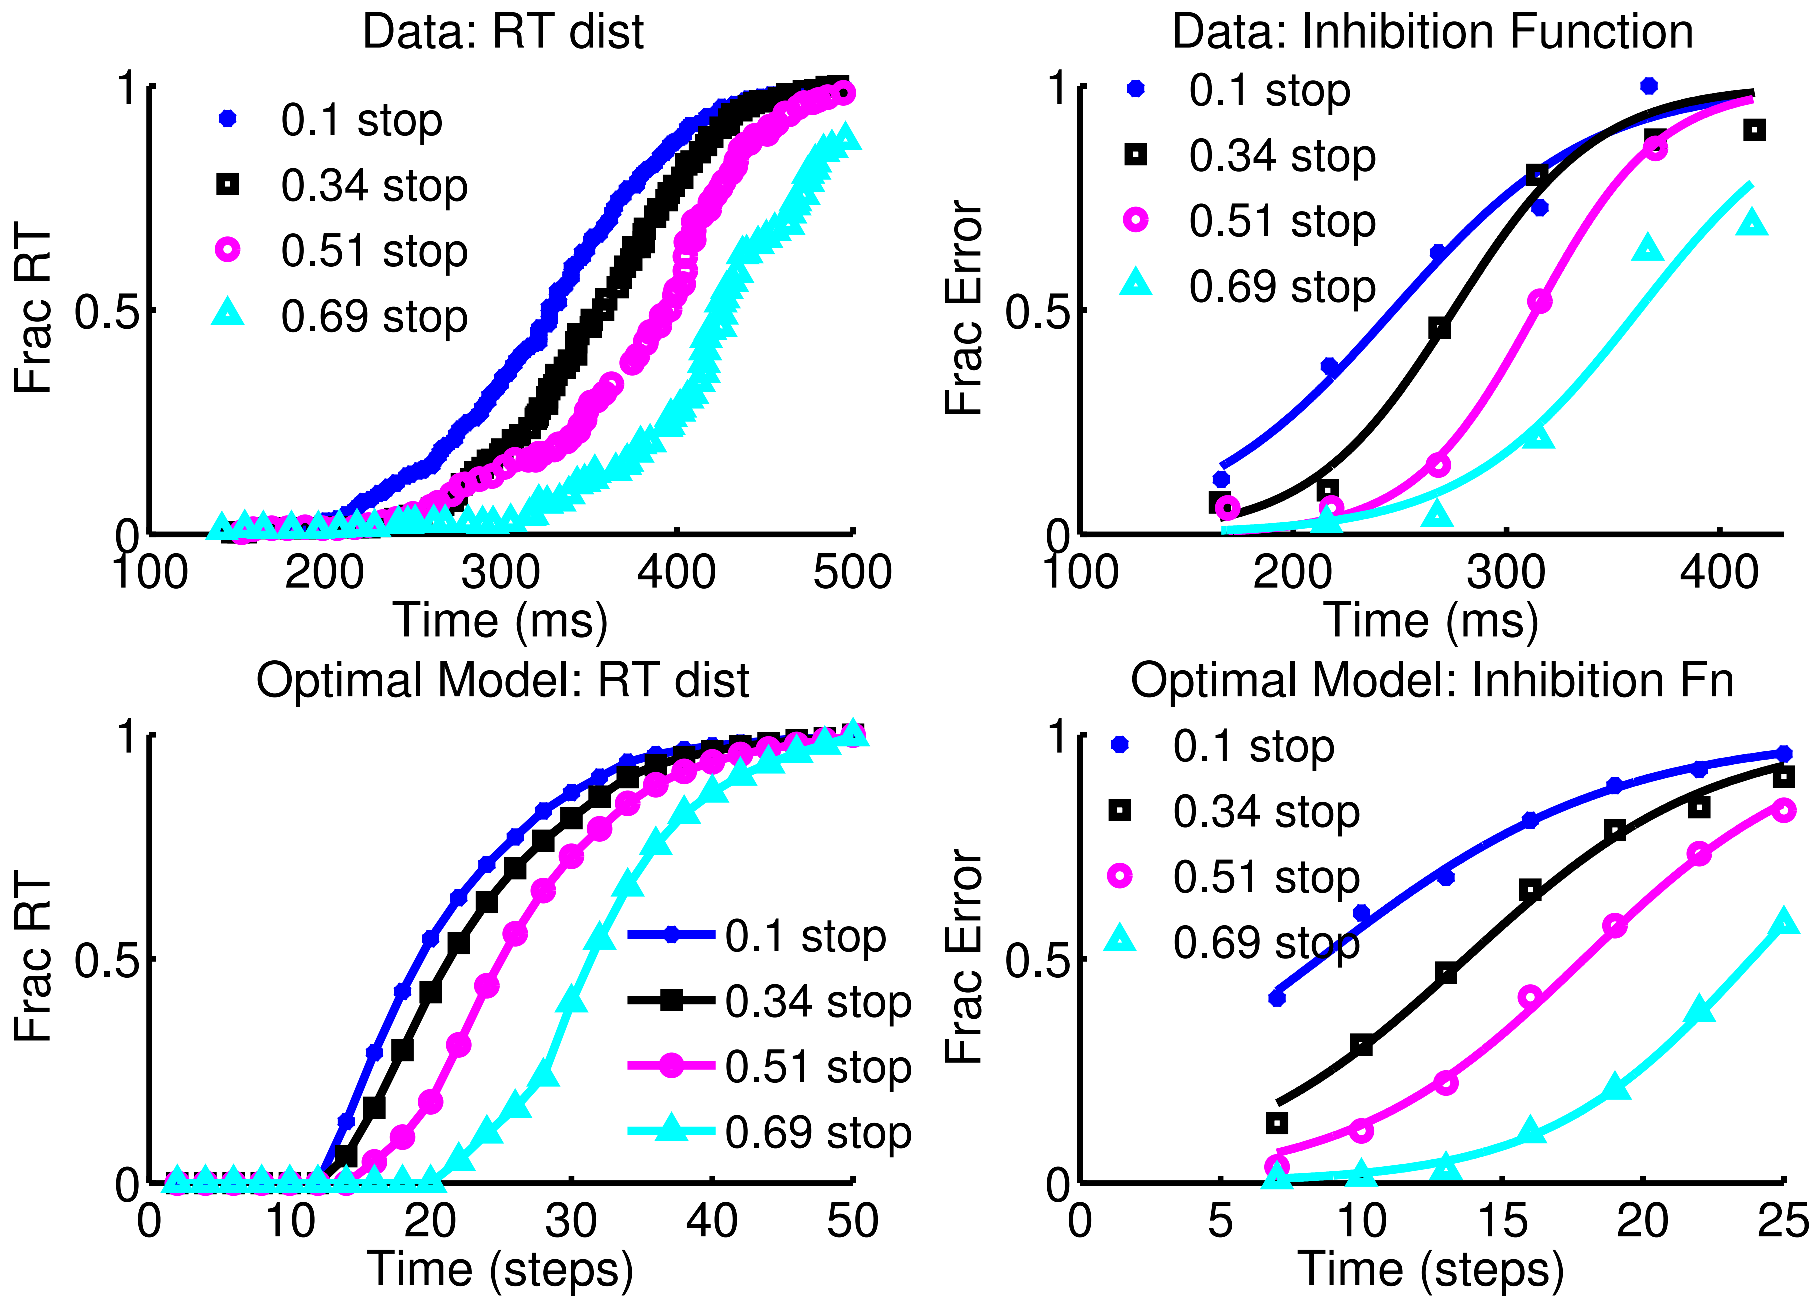

More stop trials ⇒

- ↑ go RT, ↓ stop errors

Stimulus Statistics ⇒ Stopping Behavior

(Emeric *et al.*, 2007)

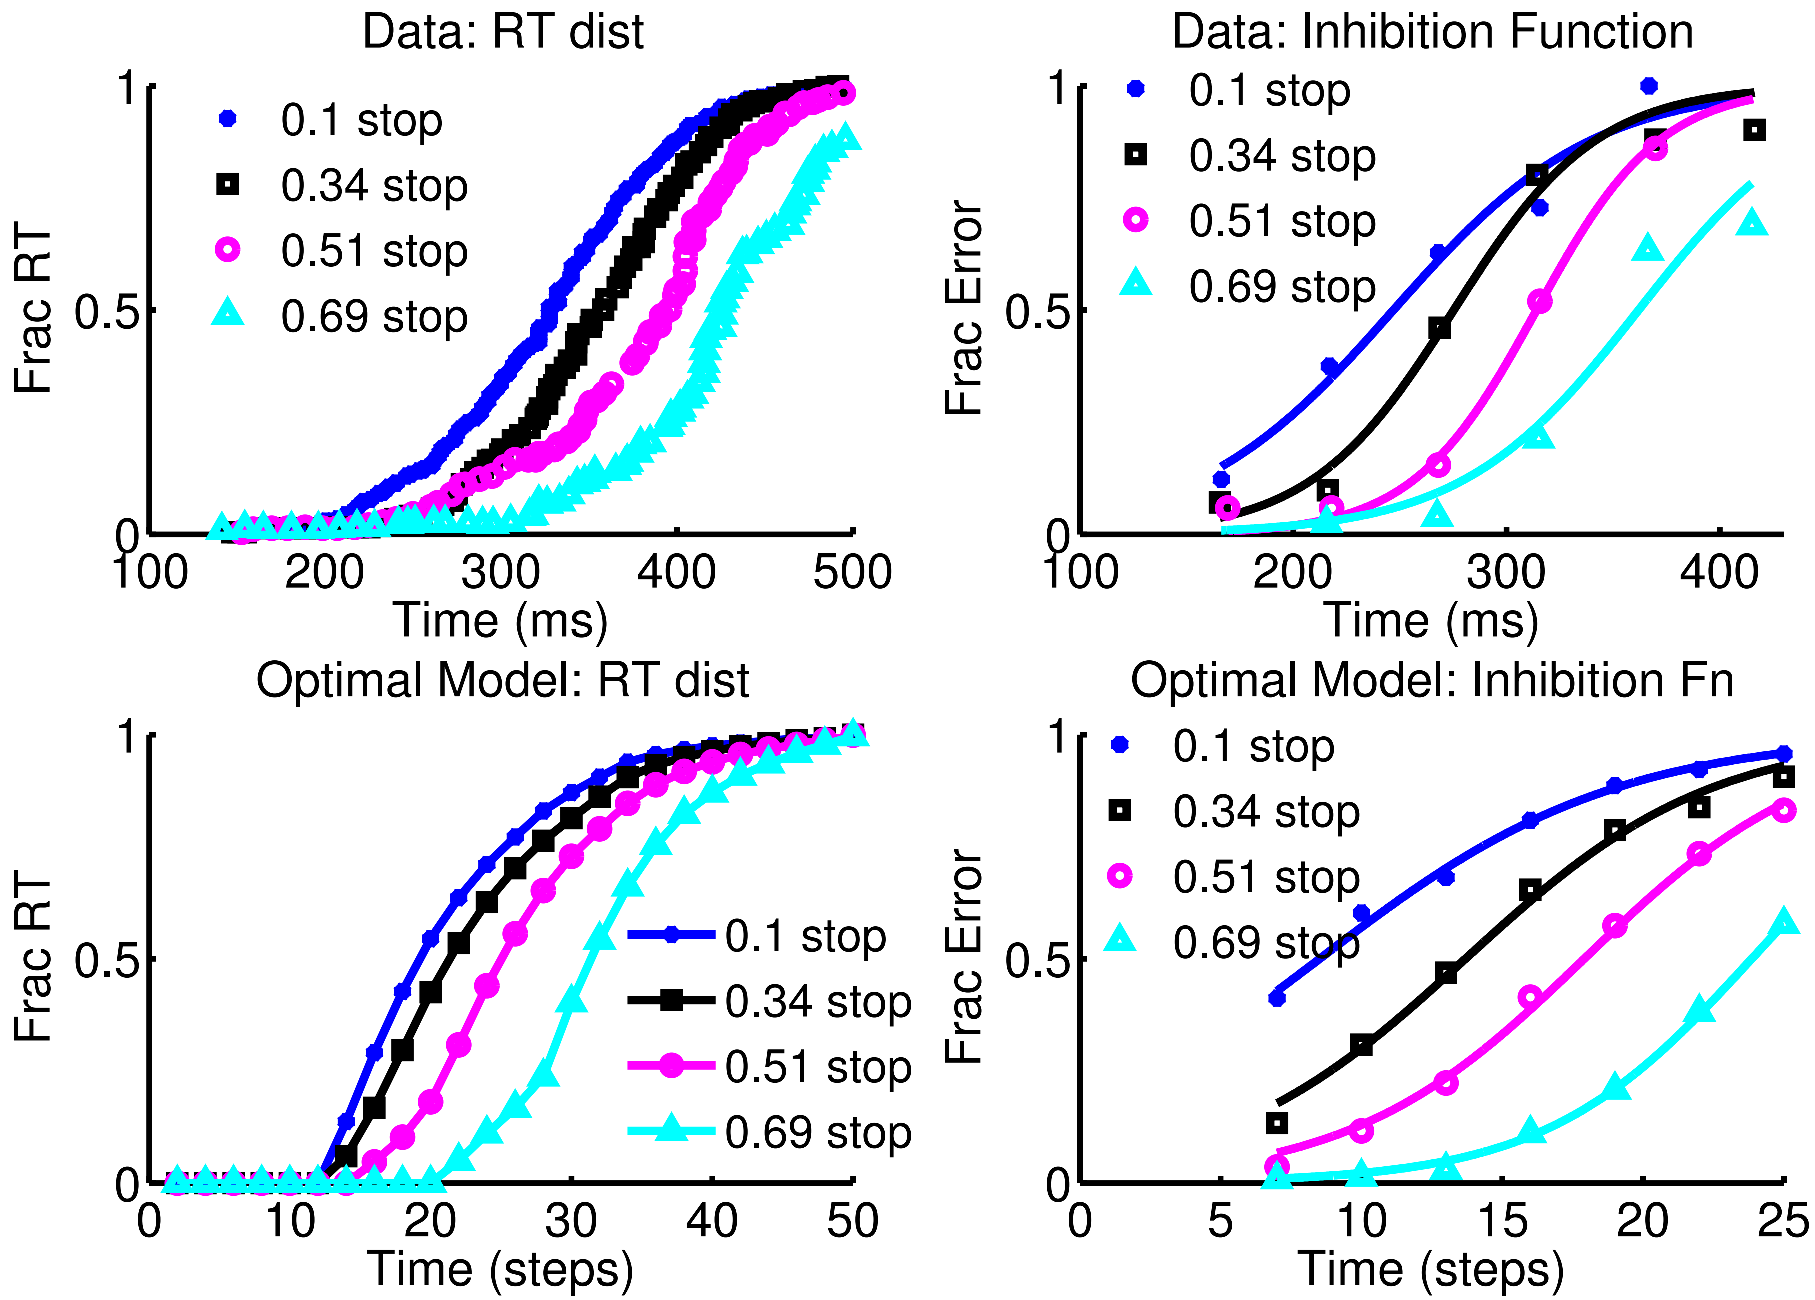

(Leotti & Wager, 2009)

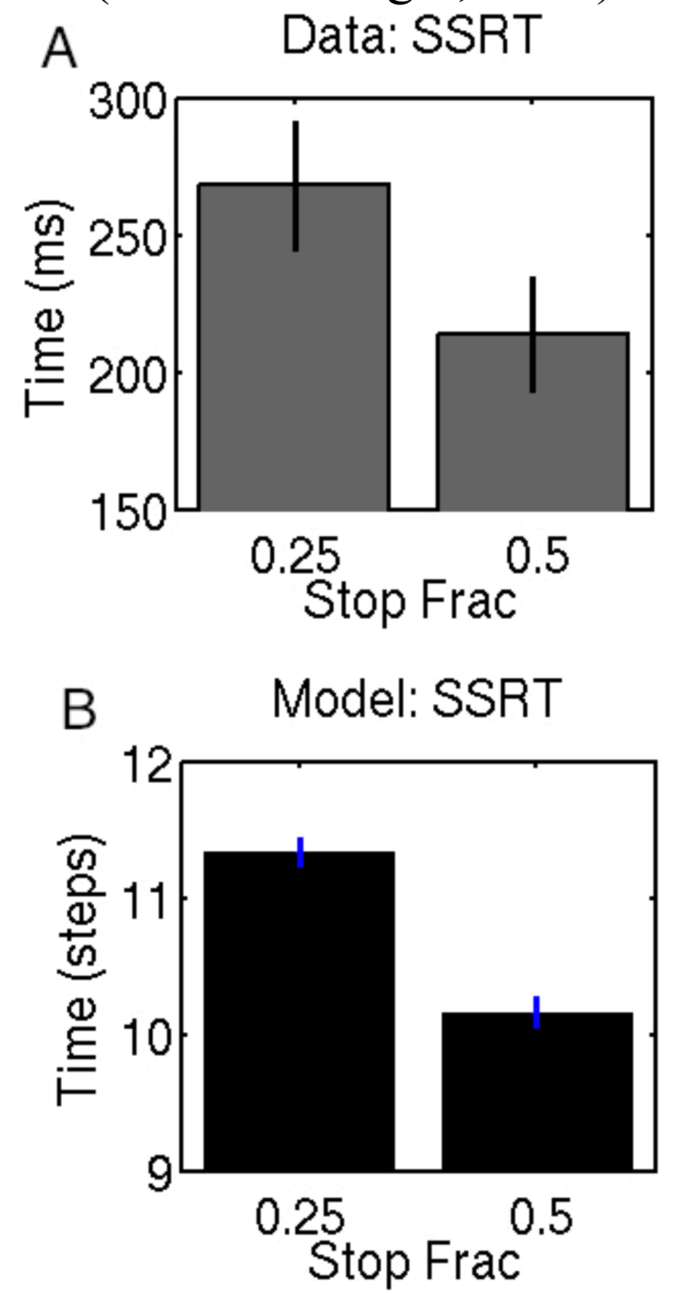

More stop trials ⇒

- ↑ go RT, ↓ stop errors
- ↓ stopping latency (SSRT)

Immediate Context ⇒ Stopping Behavior

Data (from Emric *et al.*, 2007)

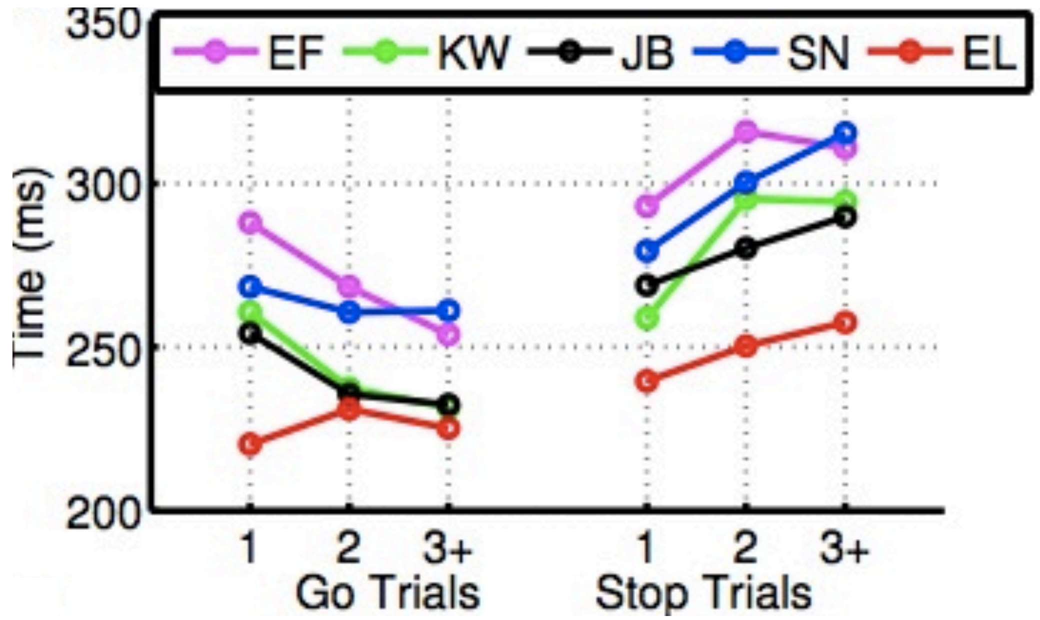

Immediate Context ⇒ Stopping Behavior

Data (from Emric *et al.*, 2007)

- **Data: dependence on trial history**
  - faster RT after go trials
  - slower after stop trials

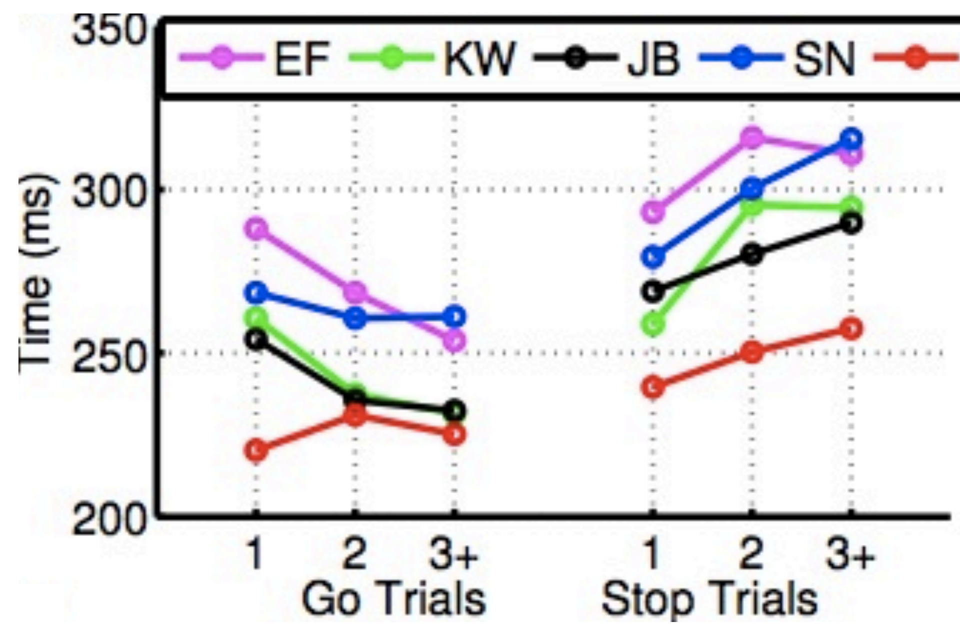

Immediate Context ⇒ Stopping Behavior

Data (from Emric *et al.*, 2007)

- **Data: dependence on trial history**
  - faster RT after go trials
  - slower after stop trials

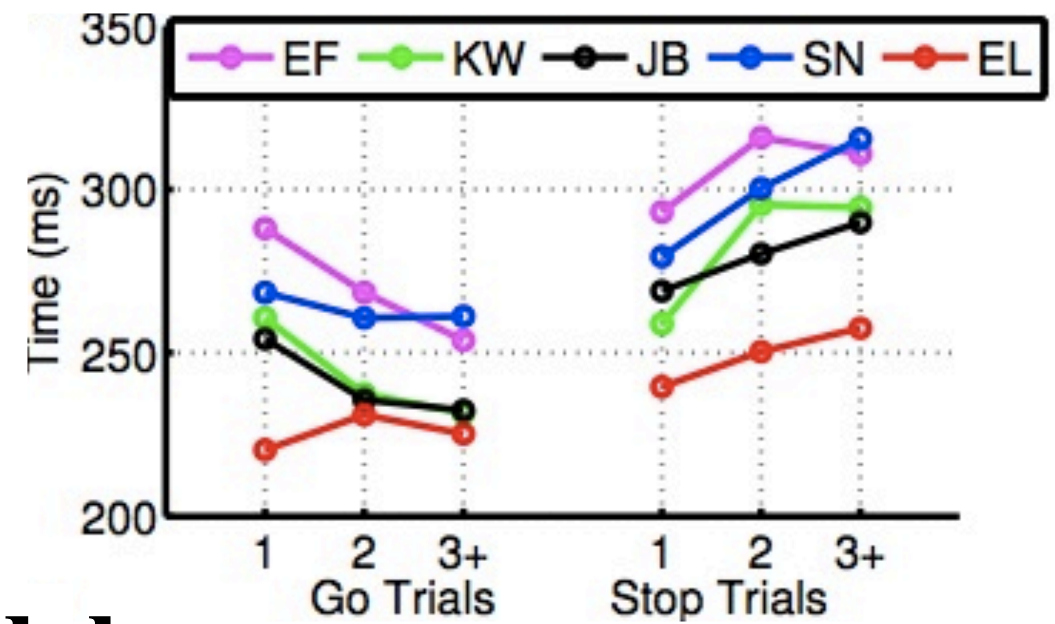

- **Model: estimating P(stop)**

Model

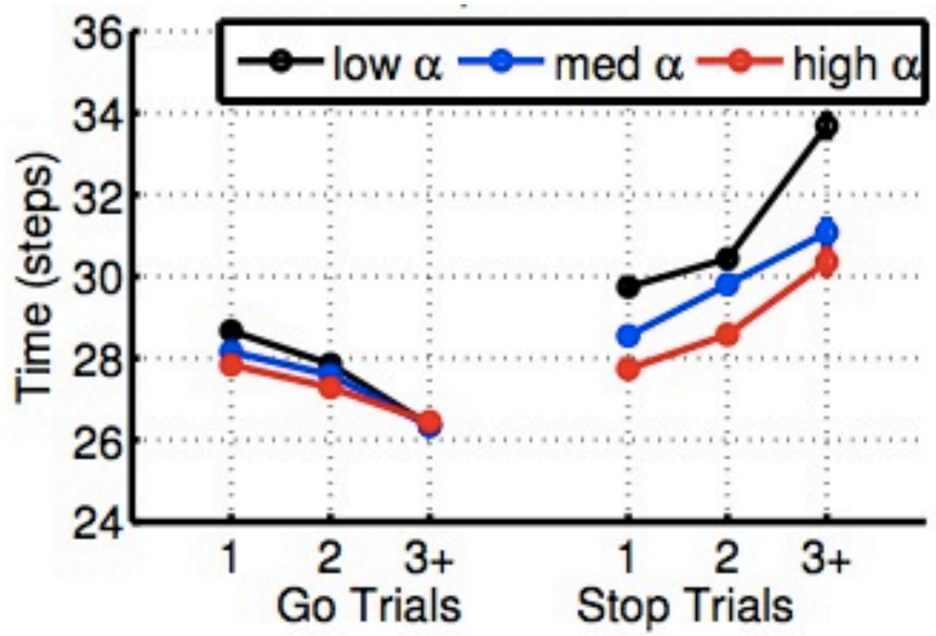

Immediate Context ⇒ Stopping Behavior

Data (from Emric *et al.*, 2007)

- **Data: dependence on trial history**
  - faster RT after go trials
  - slower after stop trials

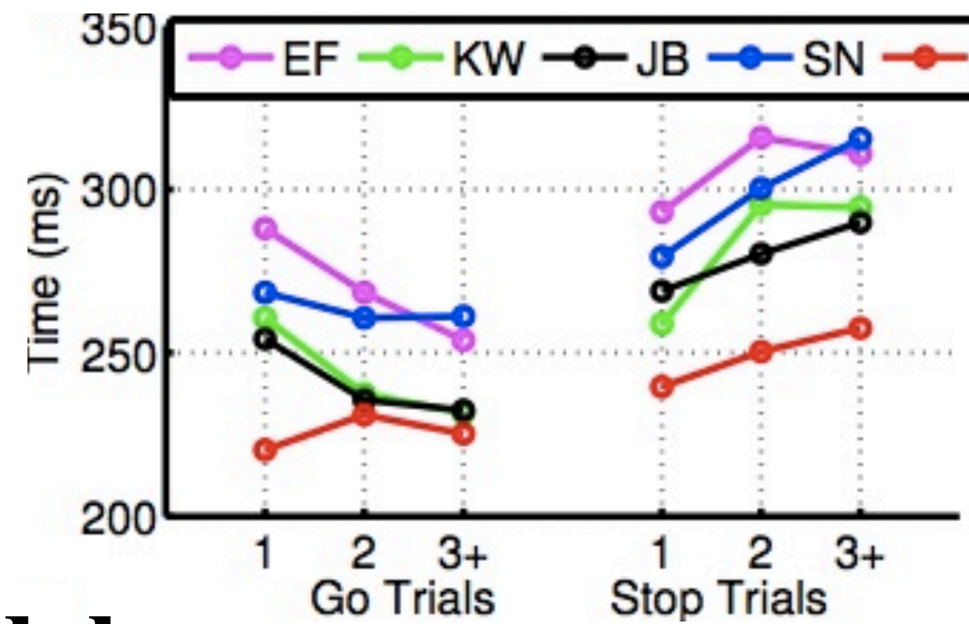

- **Model: estimating P(stop)**
  - tracking P(stop) ⇒ sequential effects

Model

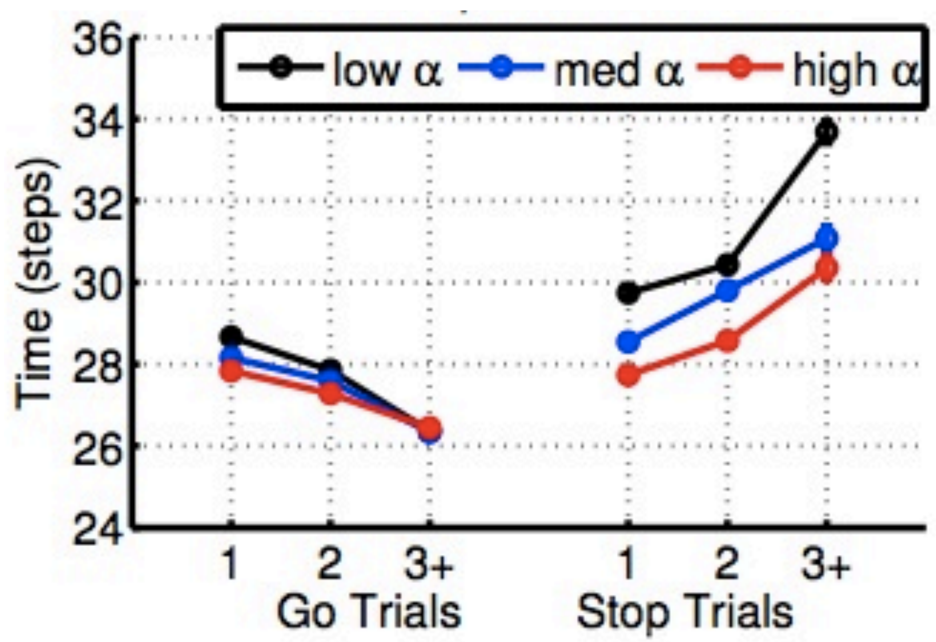

$$P(r_k|\mathbf{s}_k) \propto P(s_k|r_k)((1-\alpha)P(r_{k-1}|\mathbf{s}_{k-1}) + \alpha P_0(r-k))$$

(α = volatility = learning rate)

Immediate Context ⇒ Stopping Behavior

Data (from Emric *et al.*, 2007)

- **Data: dependence on trial history**
  - faster RT after go trials
  - slower after stop trials

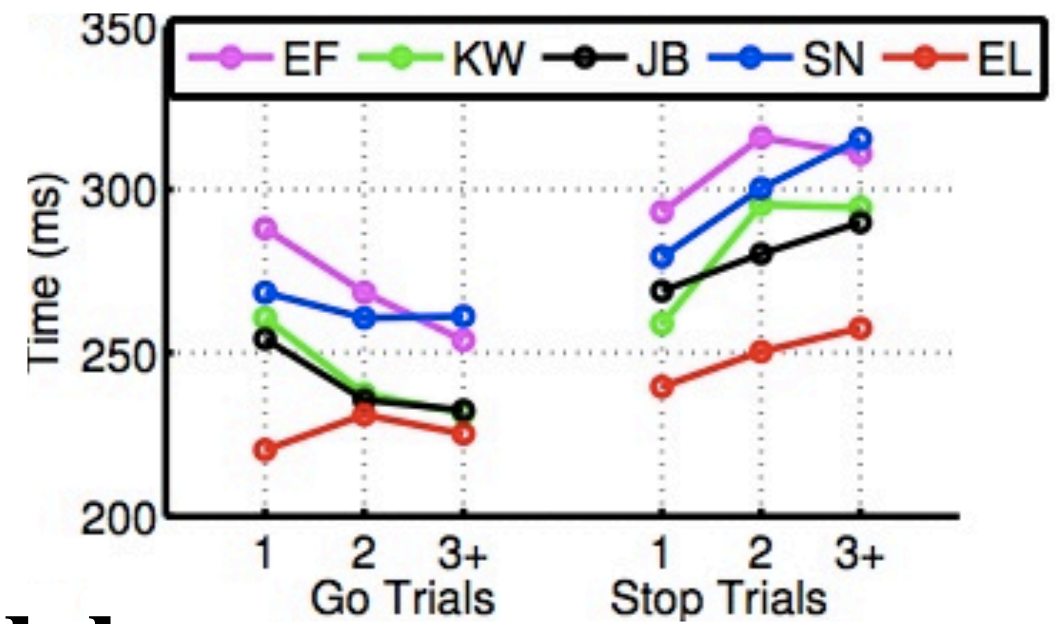

Model

- **Model: estimating P(stop)**
  - tracking P(stop) ⇒ sequential effects
  - inter-subject variability due to memory/learning rate (α)?

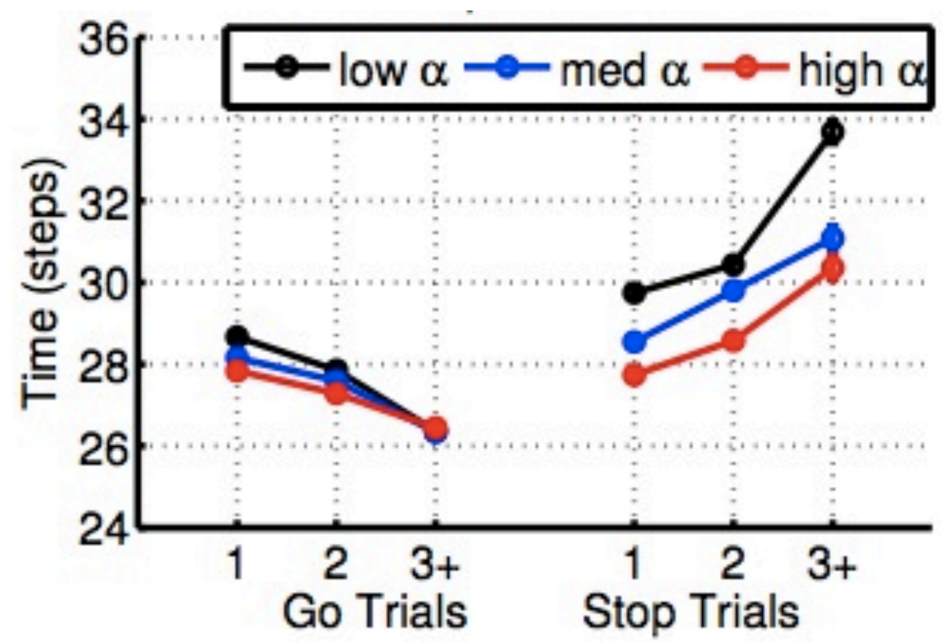

$$P(r_k|\mathbf{s}_k) \propto P(s_k|r_k)((1-\alpha)P(r_{k-1}|\mathbf{s}_{k-1}) + \alpha P_0(r-k))$$

(α = volatility = learning rate)

Prediction: Go Difficulty $\Rightarrow$ Stopping Behavior

Prediction: Go Difficulty $\Rightarrow$ Stopping Behavior

30% coherence

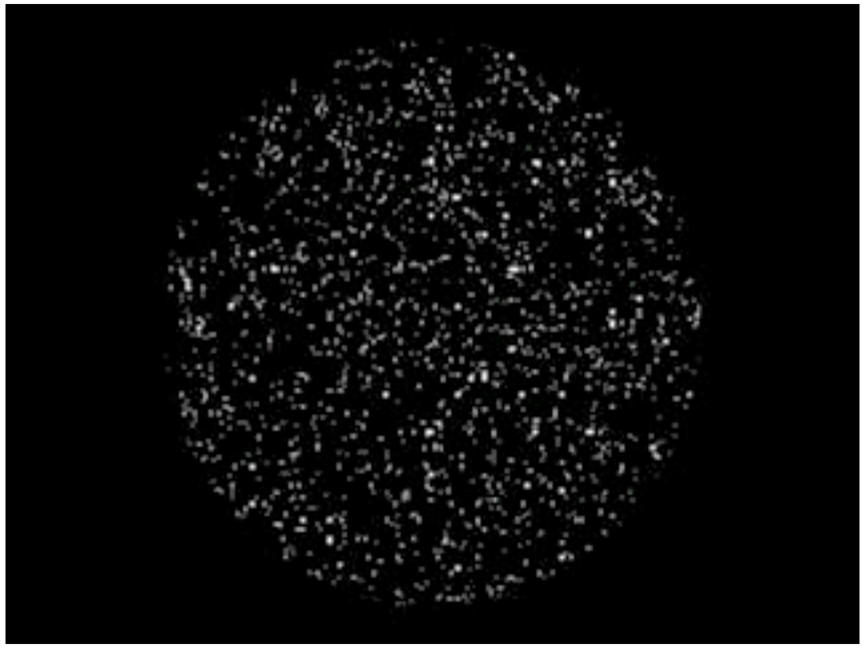

Prediction: Go Difficulty ⇒ Stopping Behavior

30% coherence

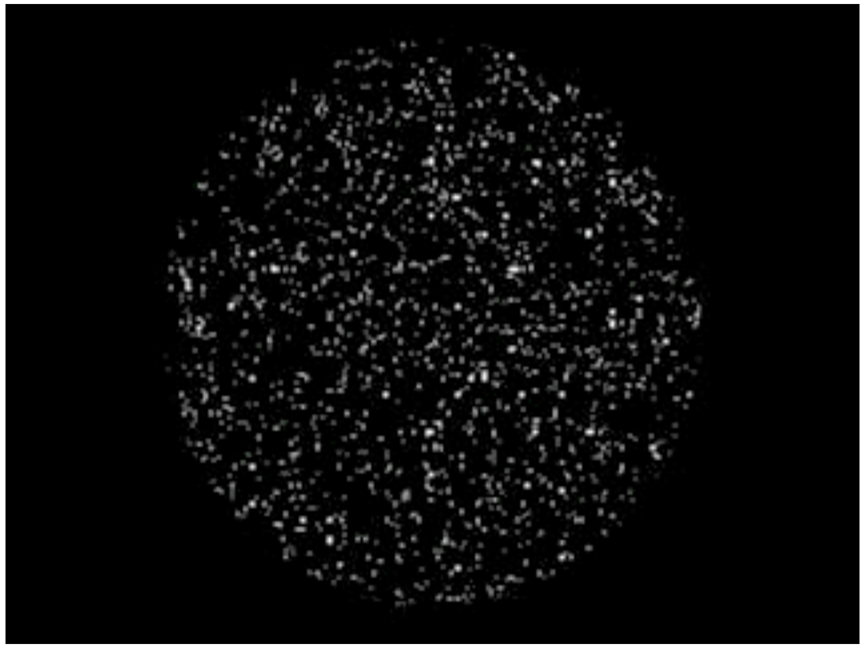

5% coherence

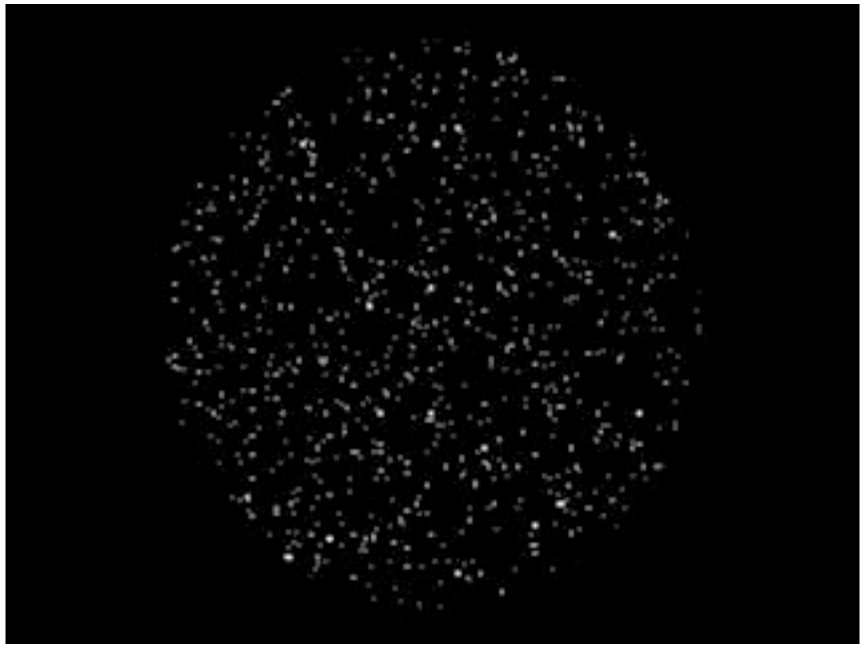

Predictions: Go Difficulty Affects Stopping

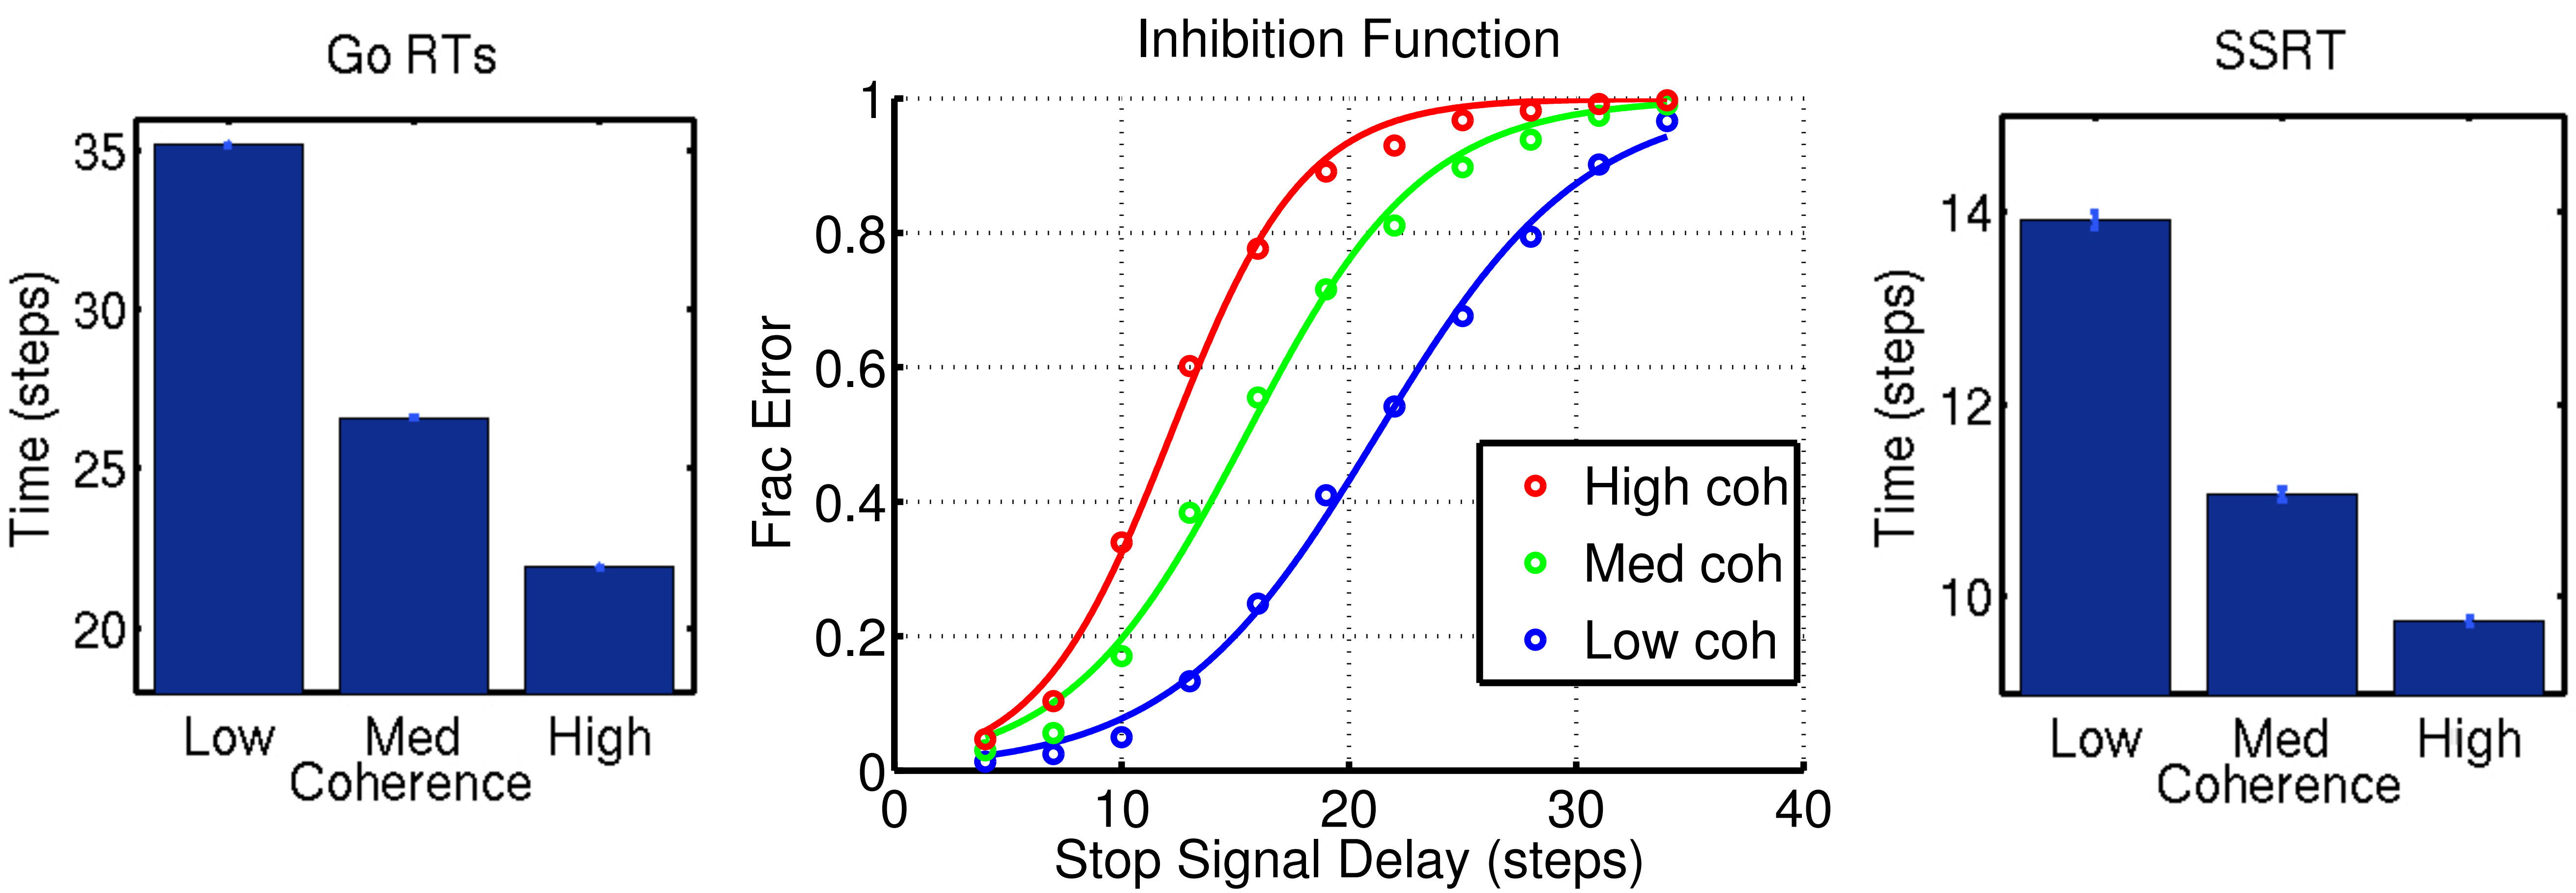

Easier go discrimination (higher coherence) ⇒

- ↓ go RT
- ↑ stop errors
- ↓ SSRT (stopping latency)

Outline

- **Model: brain implements rational (optimal) computations**
  - Monitoring process ⇔ Bayesian inference
  - Decision process ⇔ optimal stochastic control
- **Model captures a range of behavioral results**
  - Classical results
  - Reward/motivation
  - Contextual effects, sequential effects
- **Neural implementation**
  - Race (drift-diffusion) model neurally plausible approximation

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

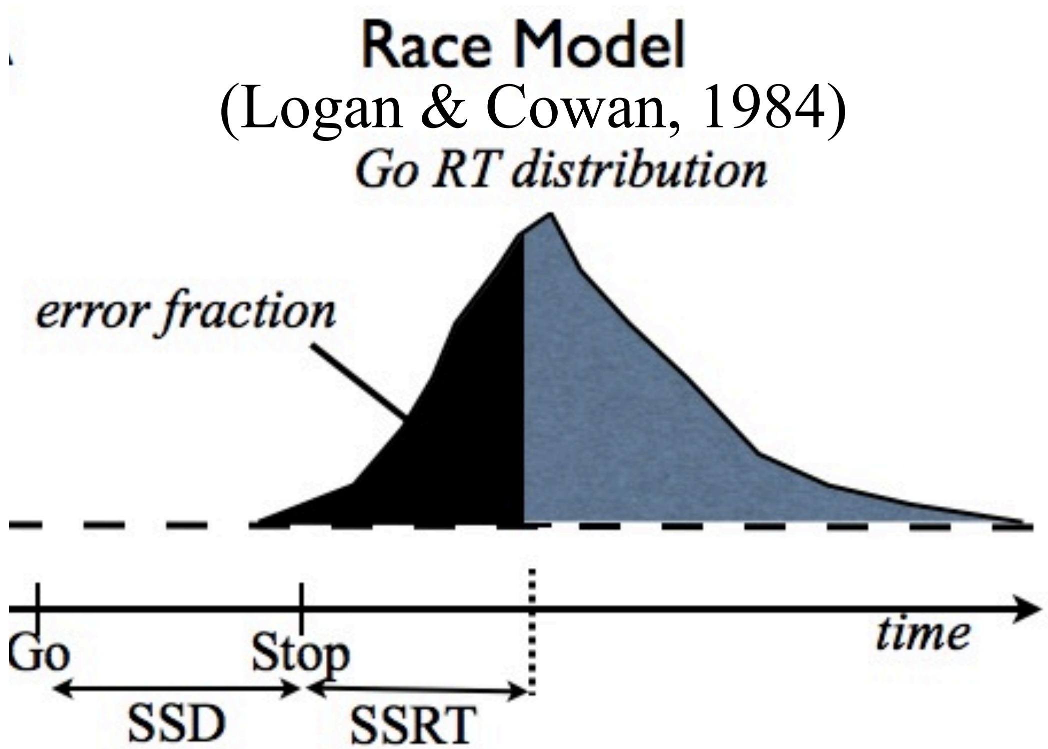

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

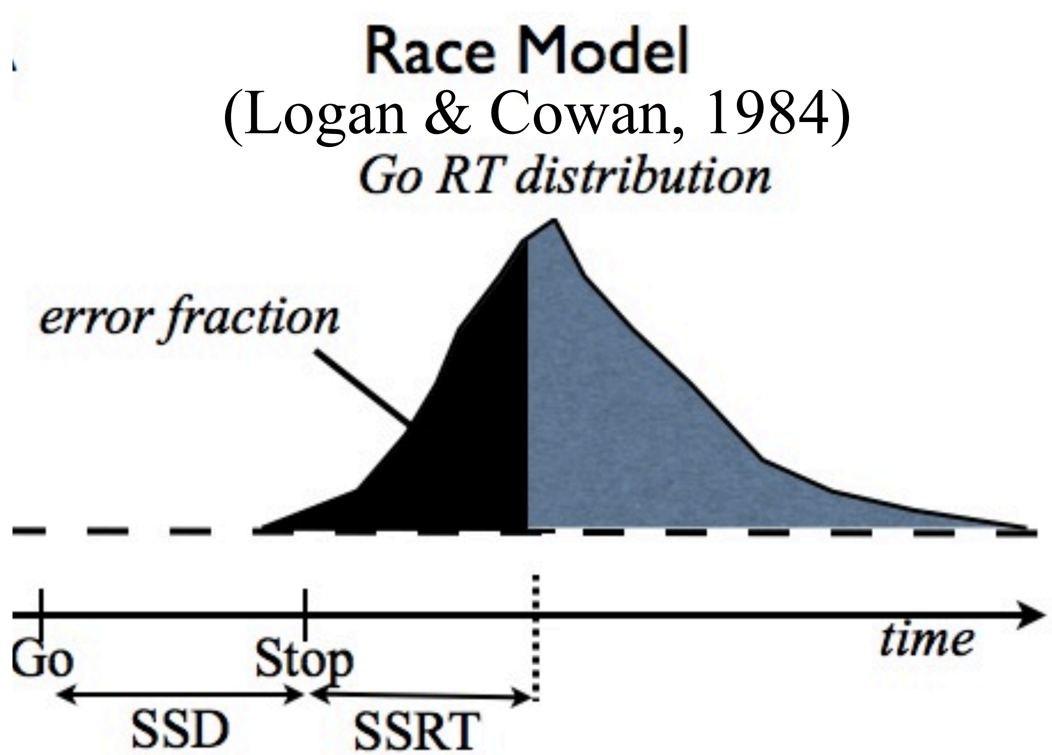

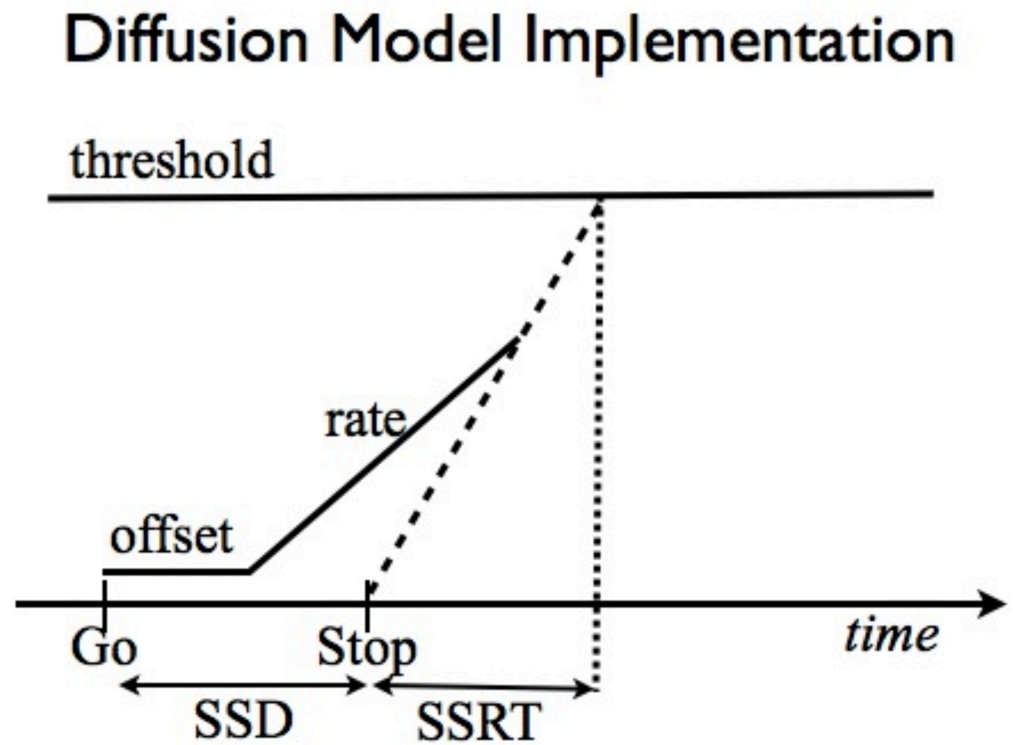

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

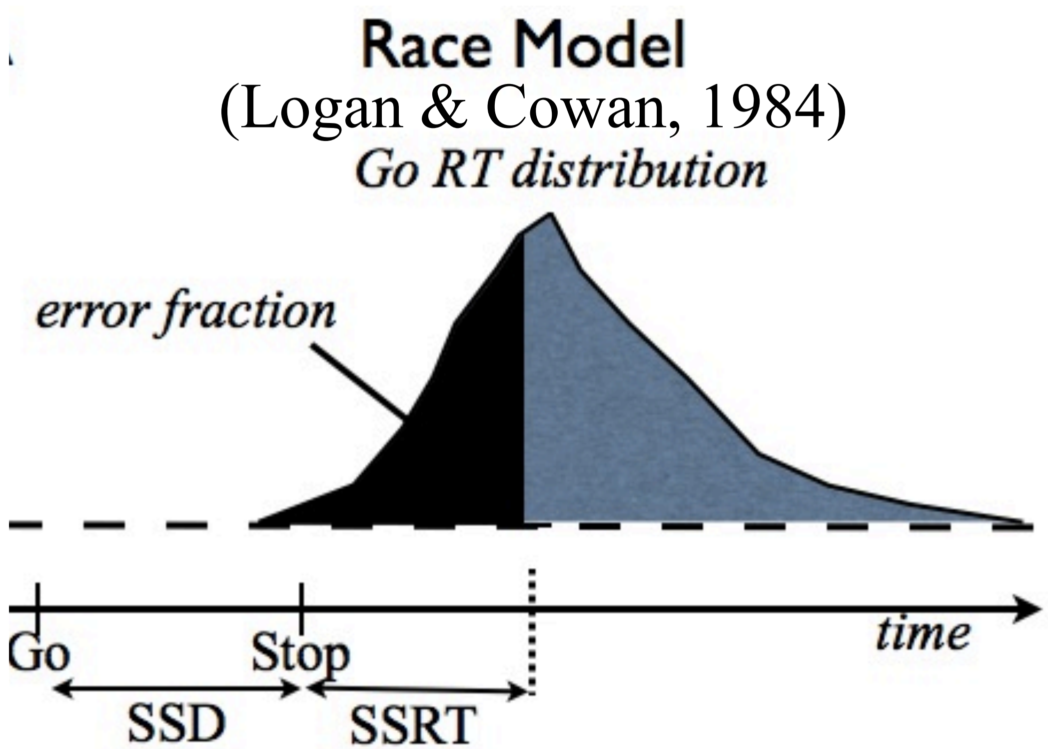

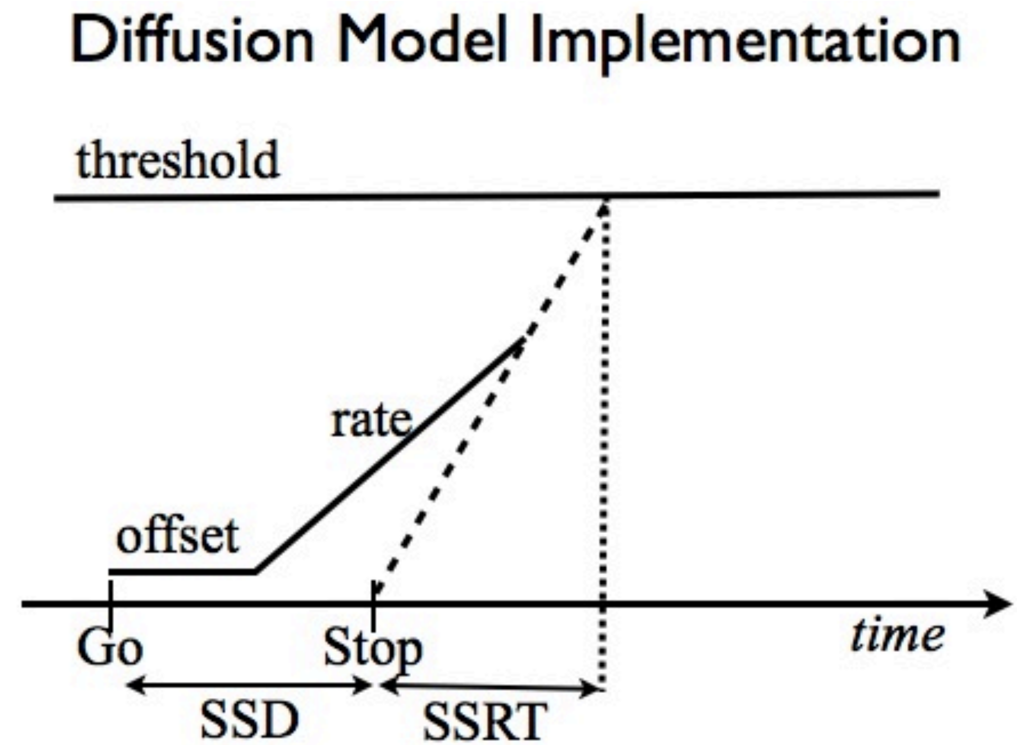

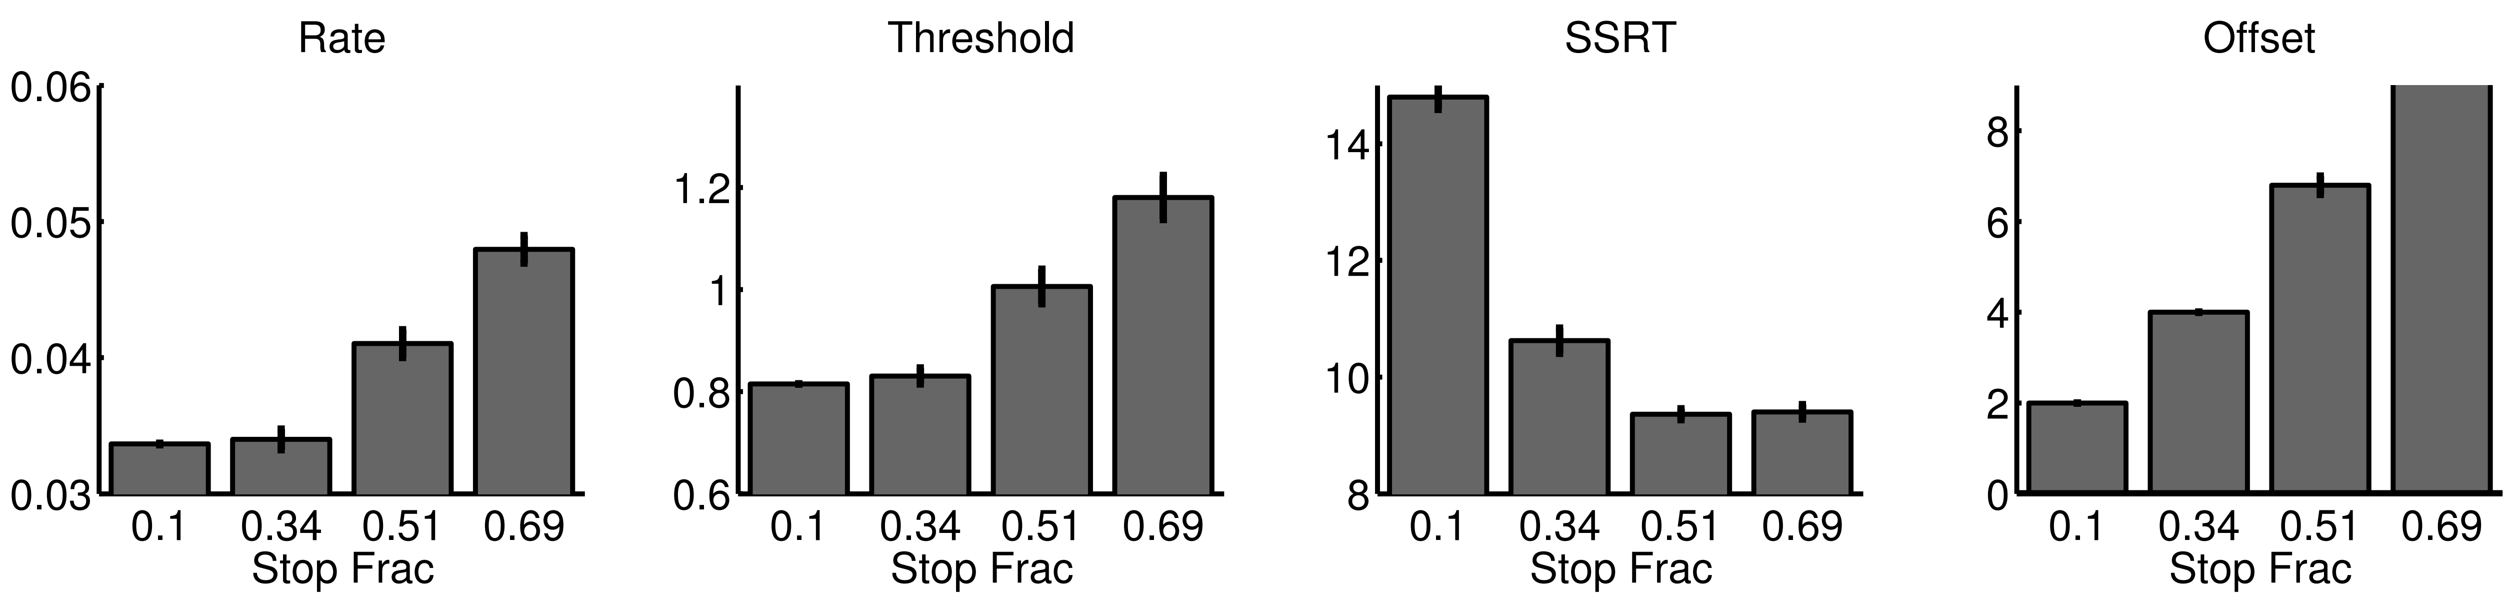

Summary
(Poster T7)

Summary
(Poster T7)

- Optimality framework for inhibitory control

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - ✤ sensory processing ⇔ Bayesian inference

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control
  - race model as (neurally plausible?) approximation

Summary

(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control
  - race model as (neurally plausible?) approximation
- Model explains wide range of behavioral data

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control
  - race model as (neurally plausible?) approximation
- Model explains wide range of behavioral data
  - classical results

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control
  - race model as (neurally plausible?) approximation
- Model explains wide range of behavioral data
  - classical results
  - reward/motivation

Summary

(Poster T7)

- Optimality framework for inhibitory control
  - ✤ sensory processing ⇔ Bayesian inference
  - ✤ decision policy ⇔ optimal stochastic control
  - ✤ race model as (neurally plausible?) approximation
- Model explains wide range of behavioral data
  - ✤ classical results
  - ✤ reward/motivation
  - ✤ contextual/sequential effects of stop trial frequency

Summary
(Poster T7)

- Optimality framework for inhibitory control
  - sensory processing ⇔ Bayesian inference
  - decision policy ⇔ optimal stochastic control
  - race model as (neurally plausible?) approximation
- Model explains wide range of behavioral data
  - classical results
  - reward/motivation
  - contextual/sequential effects of stop trial frequency
  - prediction: go difficulty affects stopping behavior

Thanks to...

Thanks to...

- **You all!**

Thanks to...

- **You all!**

Thanks to...

- You all!

- NIPS

Thanks to...

- **You all!**

- **NIPS**
- **Yu Lab**
  - Pradeep Shenoy
  - Joseph Schilz
  - Crane H Huang, Jake Olson, Katherine Naimark, Jeremy Karnowski

Thanks to...

- **You all!**

  

- **NIPS**
- **Yu Lab**
  - Pradeep Shenoy
  - Joseph Schilz
  - Crane H Huang, Jake Olson, Katherine Naimark, Jeremy Karnowski
- **Collaborators**
  - Chiang-shan Li & Jaime Ide, Martin Paulus, Veit Stuphorn, Birte Forstmann (U. Amsterdam)

Experimental Paradigm: Stop Signal Task

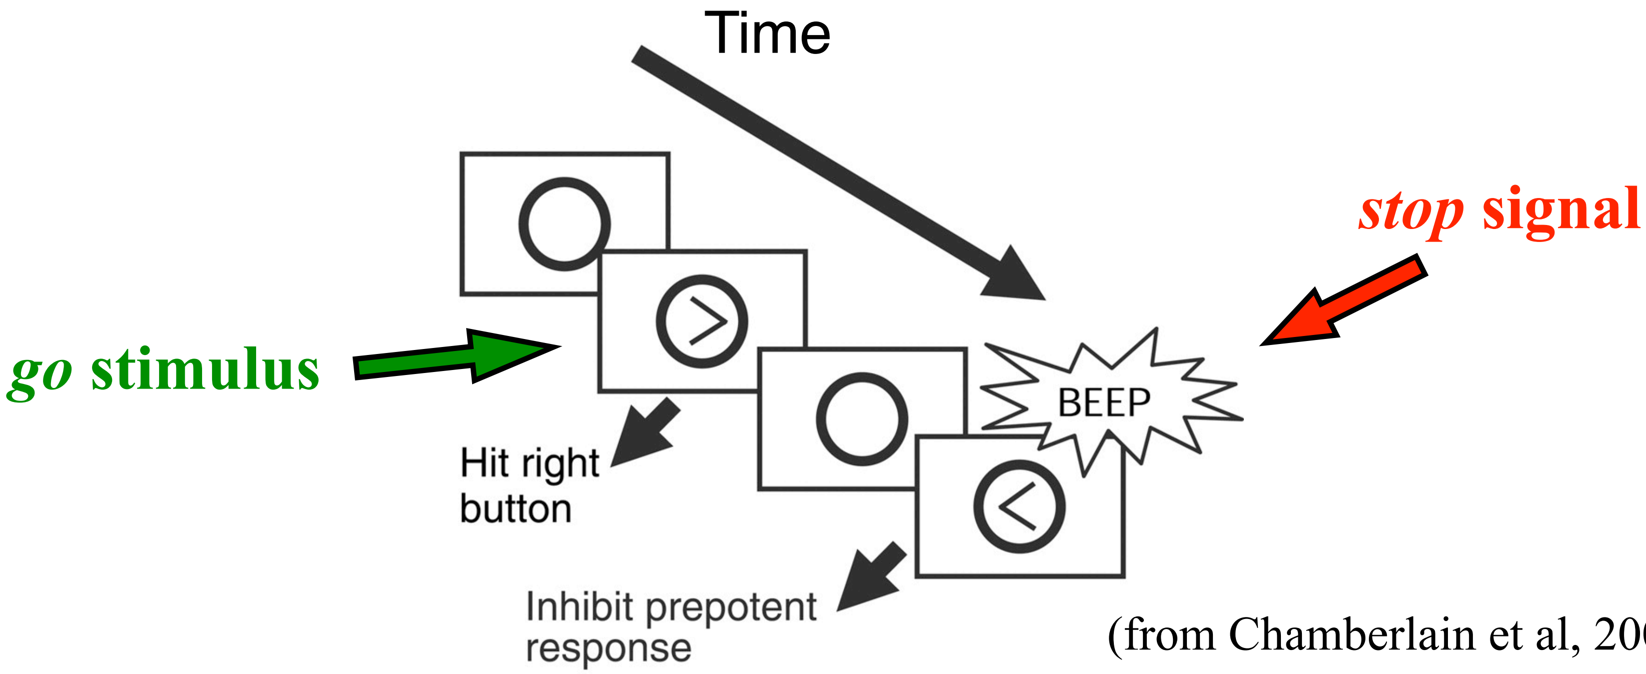

(from Chamberlain et al, 2006)

Go trial (no stop signal): correct

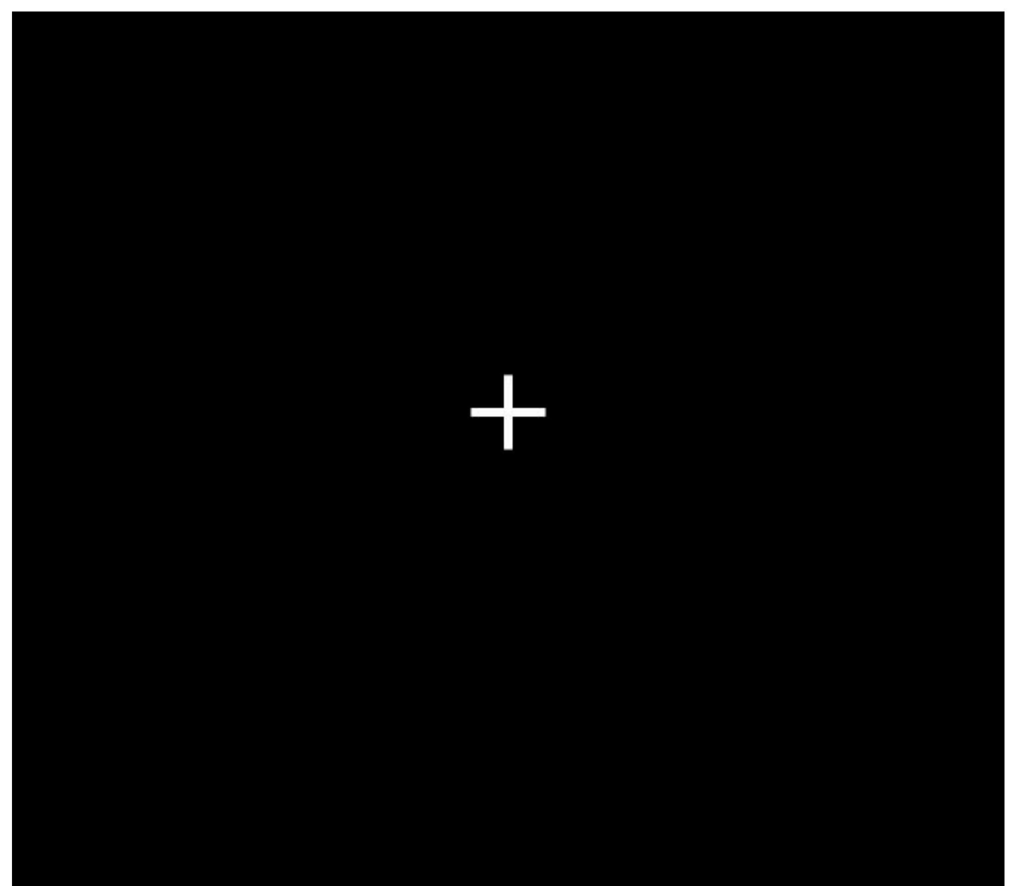

Experimental Paradigm: Stop Signal Task

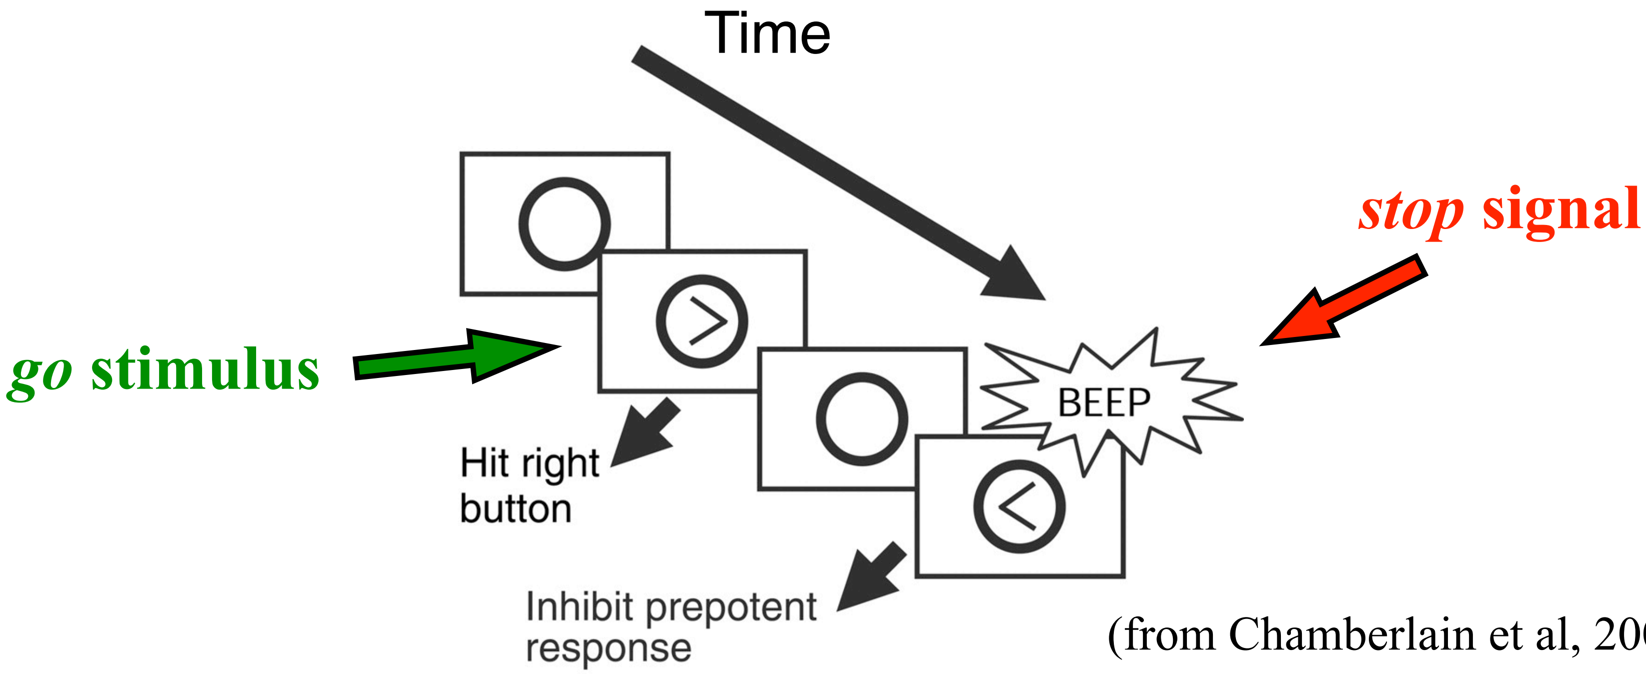

(from Chamberlain et al, 2006)

Go trial (no stop signal): error (misidentified)

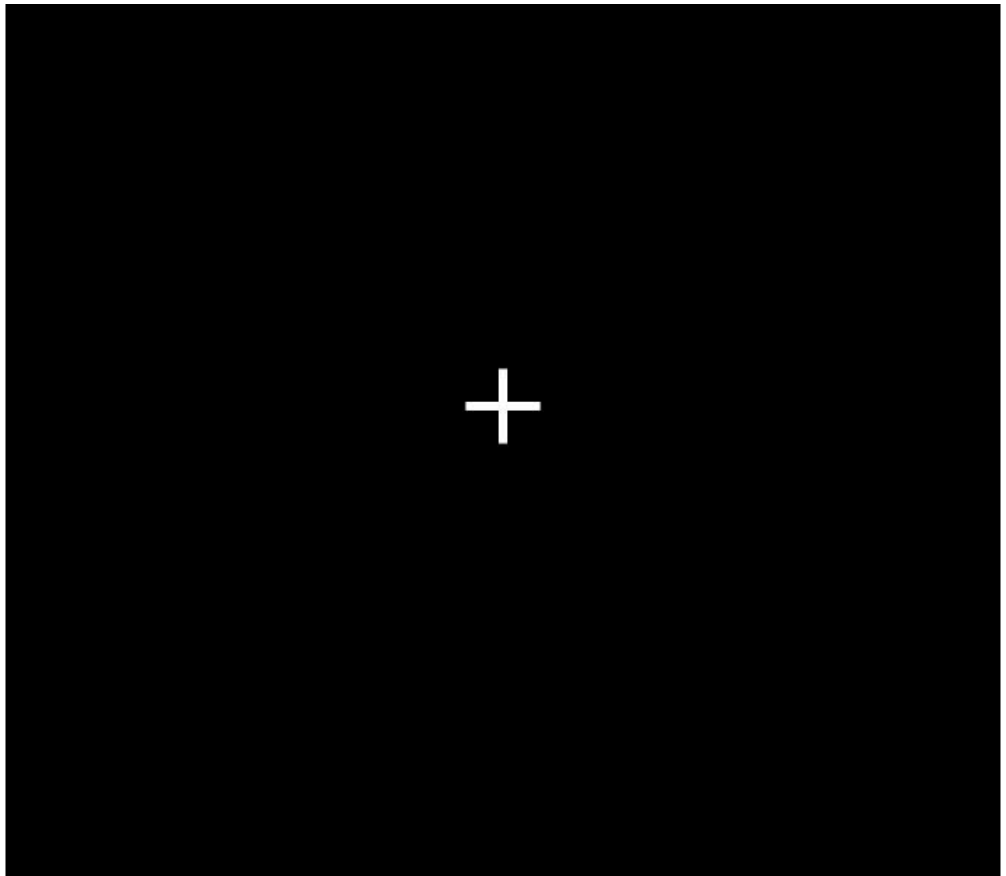

Outline

Outline

- **Model: brain implements rational (optimal) computations**
  - Sensory processing ⇔ Bayesian inference
  - Action selection ⇔ optimal stochastic control

Outline

- **Model: brain implements rational (optimal) computations**
  - Sensory processing ⇔ Bayesian inference
  - Action selection ⇔ optimal stochastic control
- **Model captures a range of behavioral results**
  - Classical results
  - Reward/motivation
  - Contextual effects, sequential effects

Outline

- **Model: brain implements rational (optimal) computations**
  - Sensory processing ⇔ Bayesian inference
  - Action selection ⇔ optimal stochastic control
- **Model captures a range of behavioral results**
  - Classical results
  - Reward/motivation
  - Contextual effects, sequential effects
- **Neural implementation**
  - Race (drift-diffusion) model as neurally plausible approximation

Inhibitory Control: An Example

Inhibitory Control: An Example

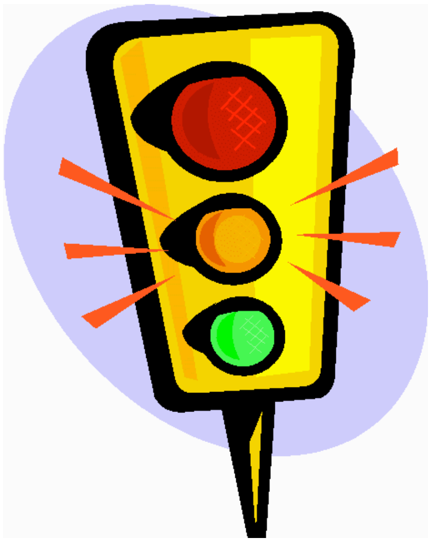

Inhibitory Control: An Example

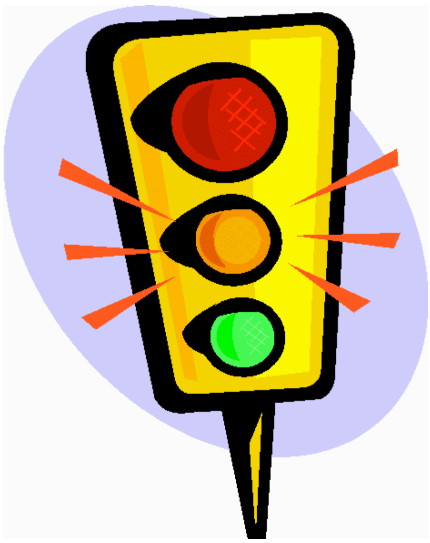

Possible actions

Inhibitory Control: An Example

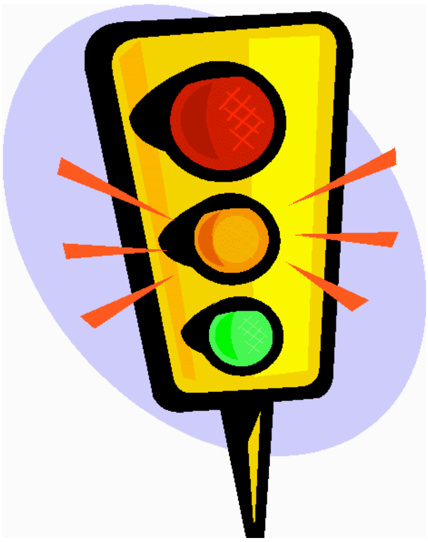

Possible actions

- stop

Inhibitory Control: An Example

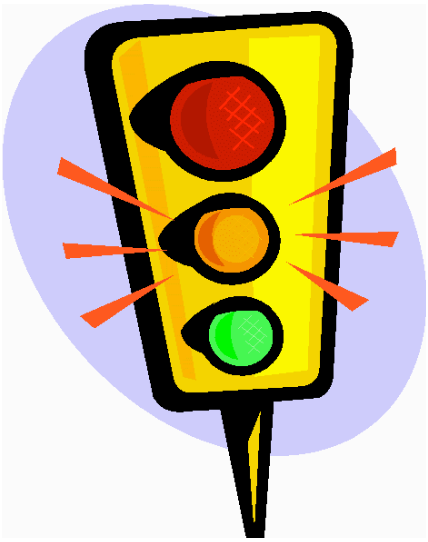

Possible actions

- stop
- go

Inhibitory Control: An Example

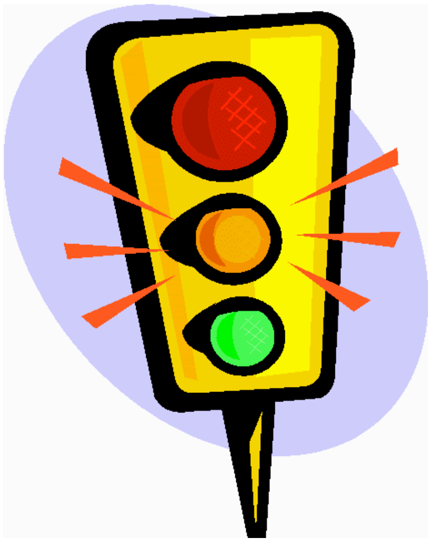

Possible actions

- stop
- go
- speed up!

Current & Future Work

Current & Future Work

- Neural implementation/approximation of optimal DM?
  - interactive race model (Boucher et al, 2007): fixation/movement neurons in FEF & SC
  - theory: other (neural) approximation of optimal DM?
  - experiments: fMRI (Li & Ide @ Yale), EEG (Makeig @ UCSD)

Current & Future Work

- Neural implementation/approximation of optimal DM?
  - interactive race model (Boucher et al, 2007): fixation/movement neurons in FEF & SC
  - theory: other (neural) approximation of optimal DM?
  - experiments: fMRI (Li & Ide @ Yale), EEG (Makeig @ UCSD)
- Effects of SSD distribution on stopping behavior
  - temporal expectancies $\Rightarrow$ stopping errors & SSRt

Current & Future Work

- Neural implementation/approximation of optimal DM?
  - interactive race model (Boucher et al, 2007): fixation/movement neurons in FEF & SC
  - theory: other (neural) approximation of optimal DM?
  - experiments: fMRI (Li & Ide @ Yale), EEG (Makeig @ UCSD)
- Effects of SSD distribution on stopping behavior
  - temporal expectancies ⇒ stopping errors & SSRt
- Population with impaired inhibitory control
  - depressives, stimulant users -- differentiate underlying cause

Results: Basic Effects

A. Model

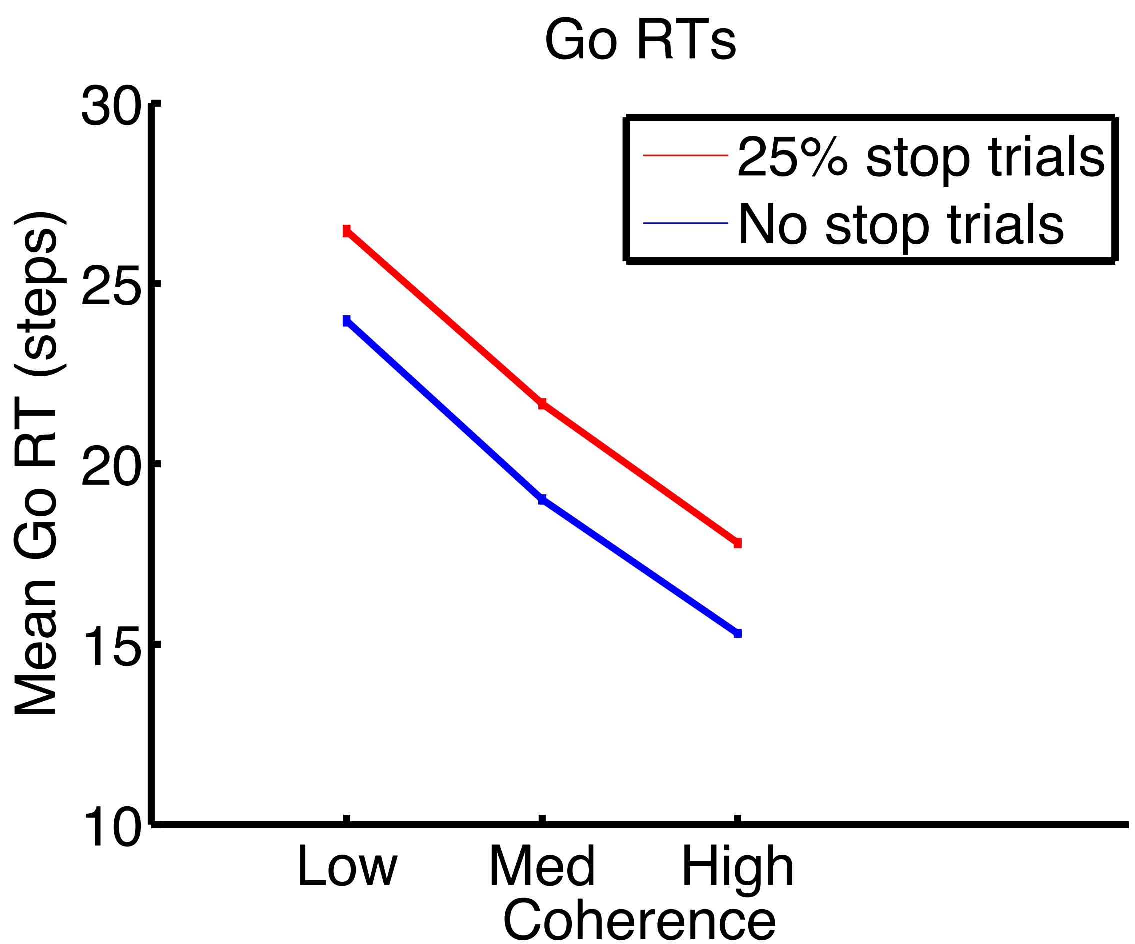

B. Data *n=5*

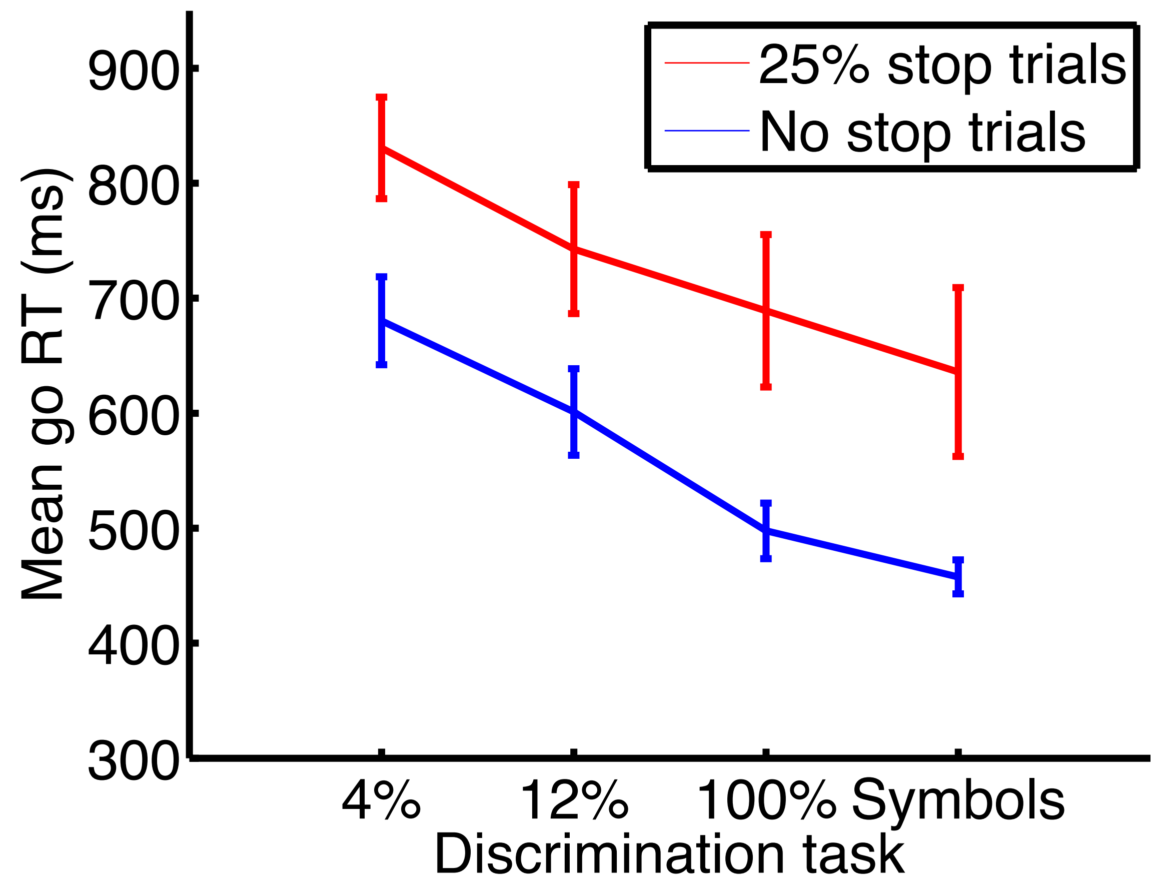

Results: Basic Effects

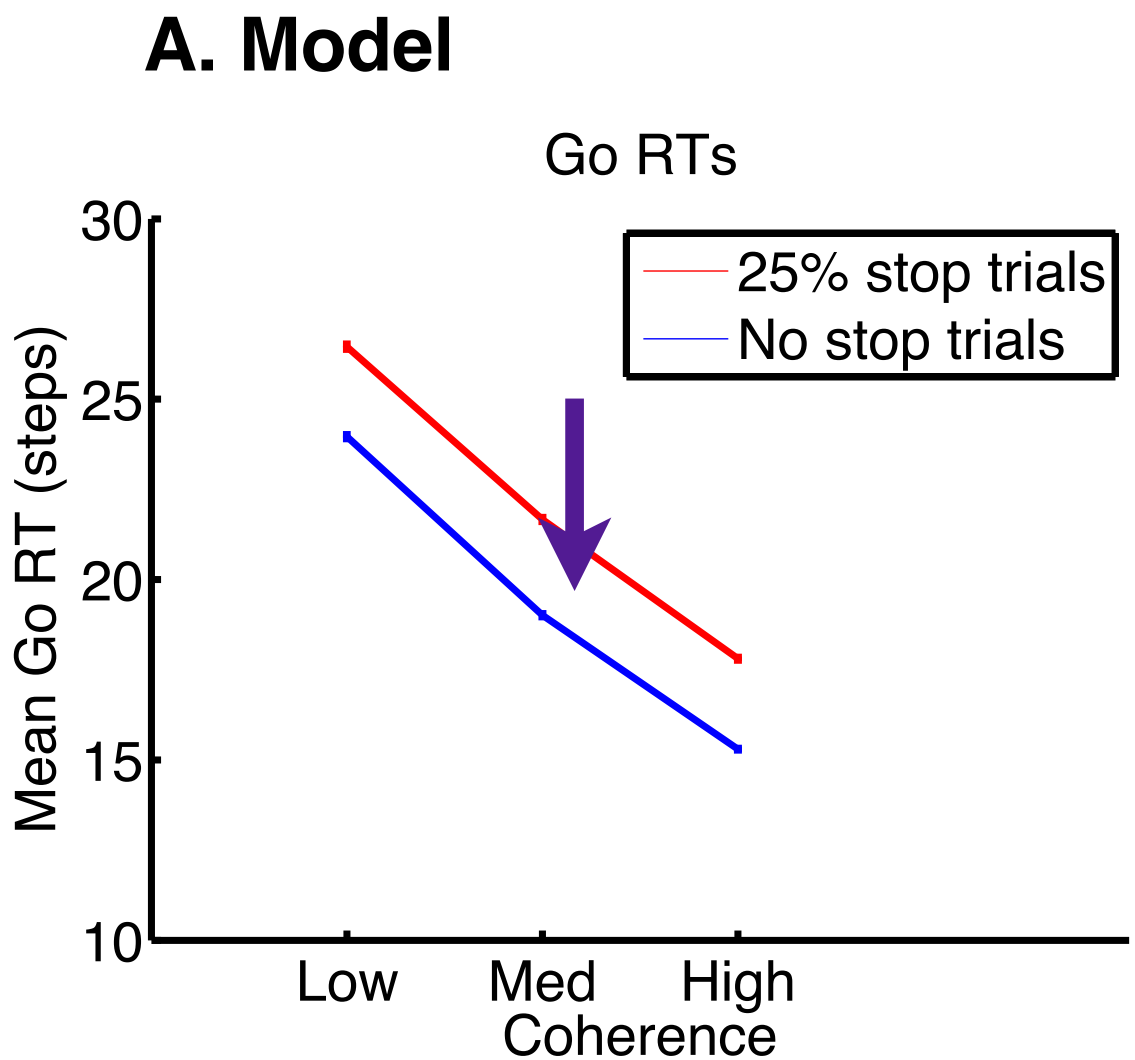

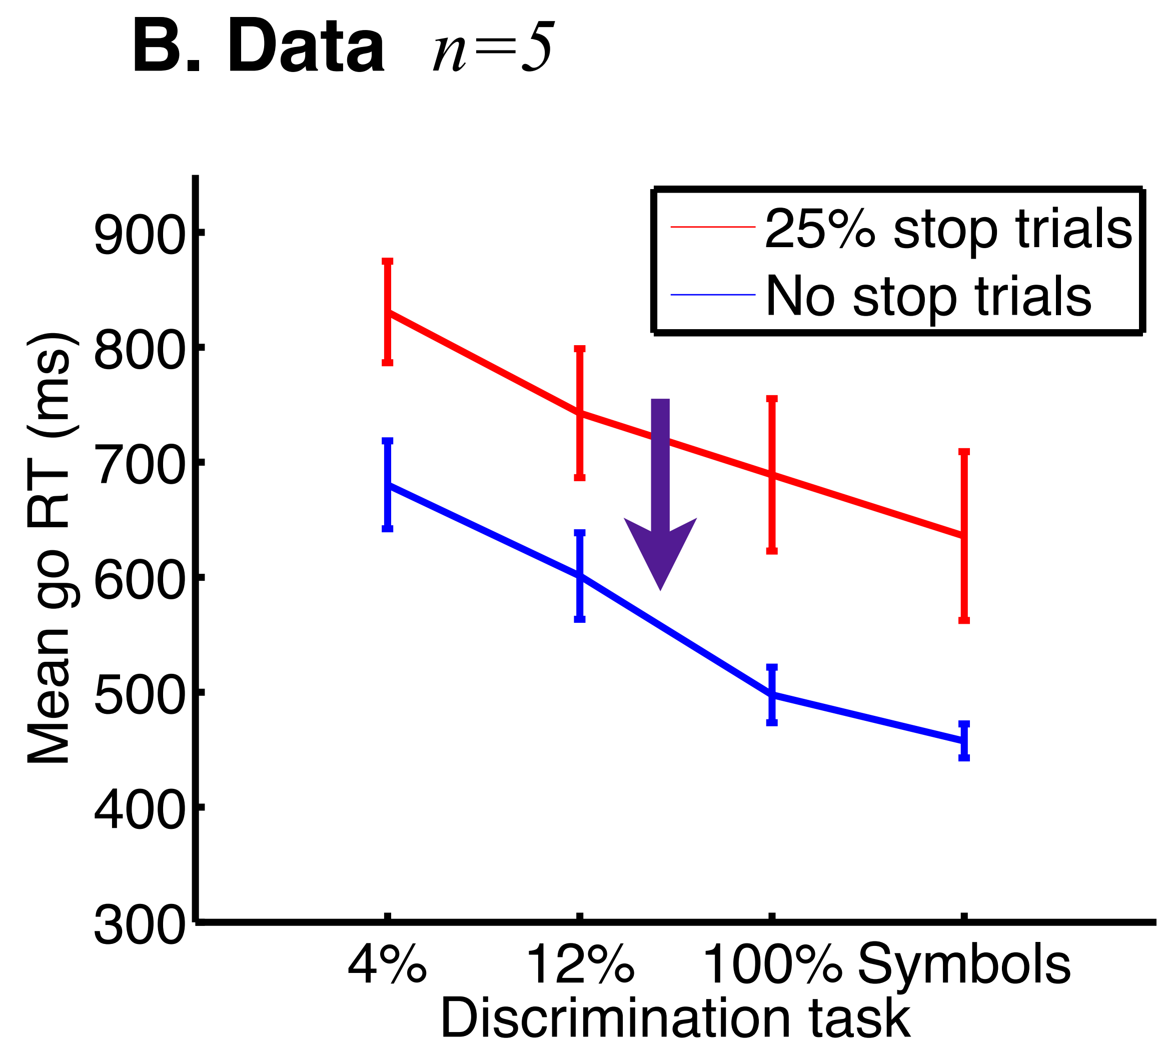

- *Stop* trials induce *slowing* of *Go* RT

Results: Basic Effects

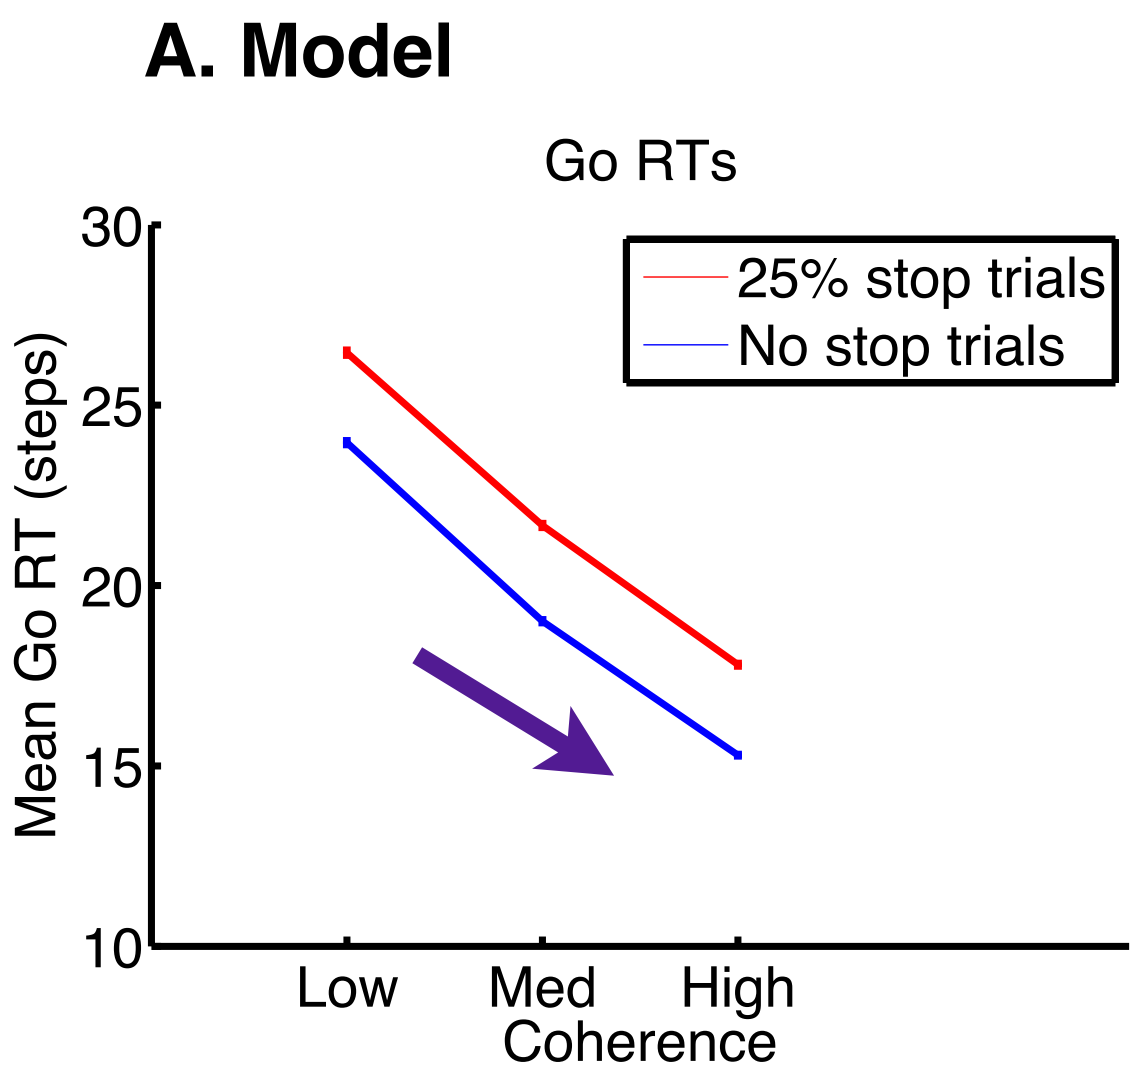

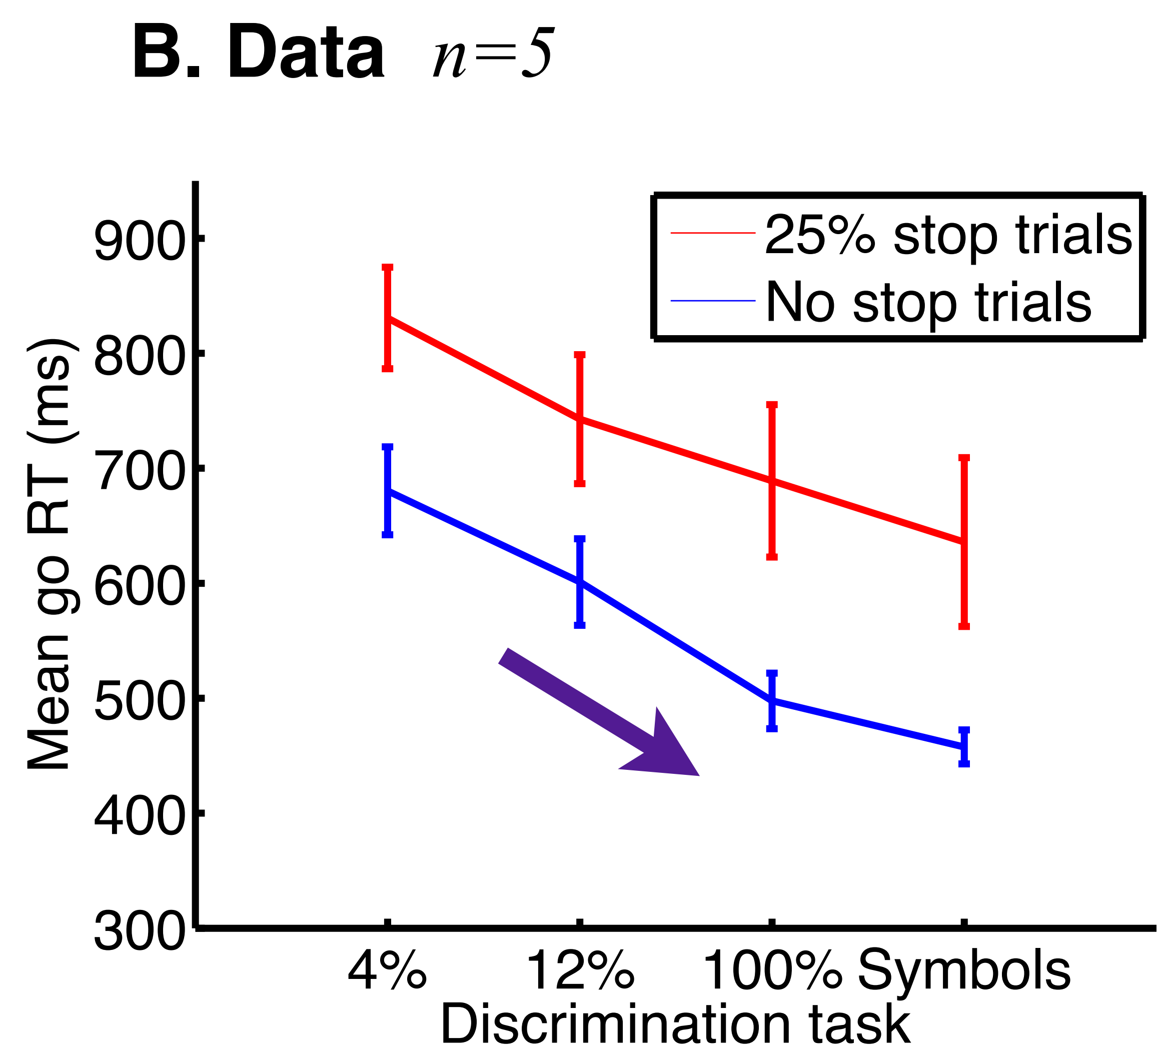

- ***Stop*** **trials induce** ***slowing*** **of** ***Go*** **RT**
- **Higher coherence (easier) induces faster** ***Go*** **response**

Results: Basic Effects

A. Model

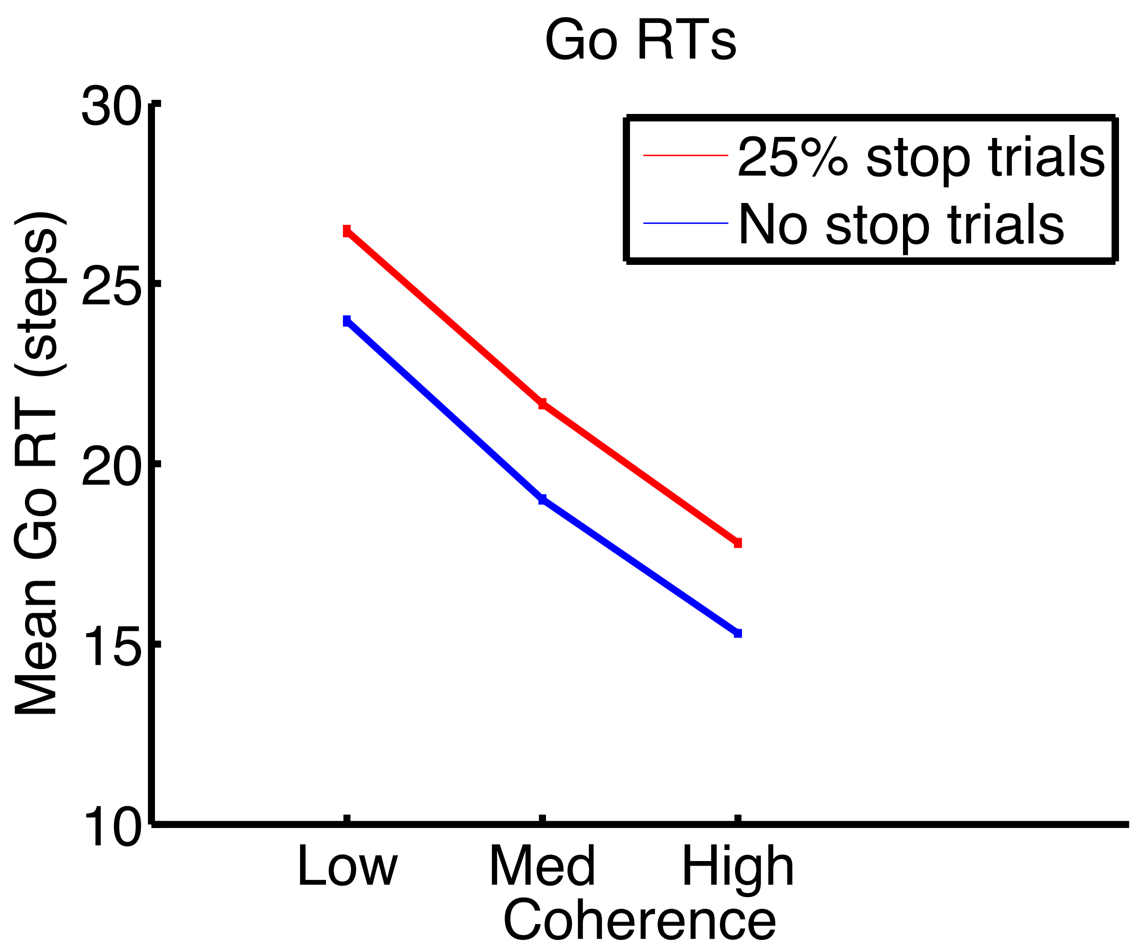

B. Data *n=5*

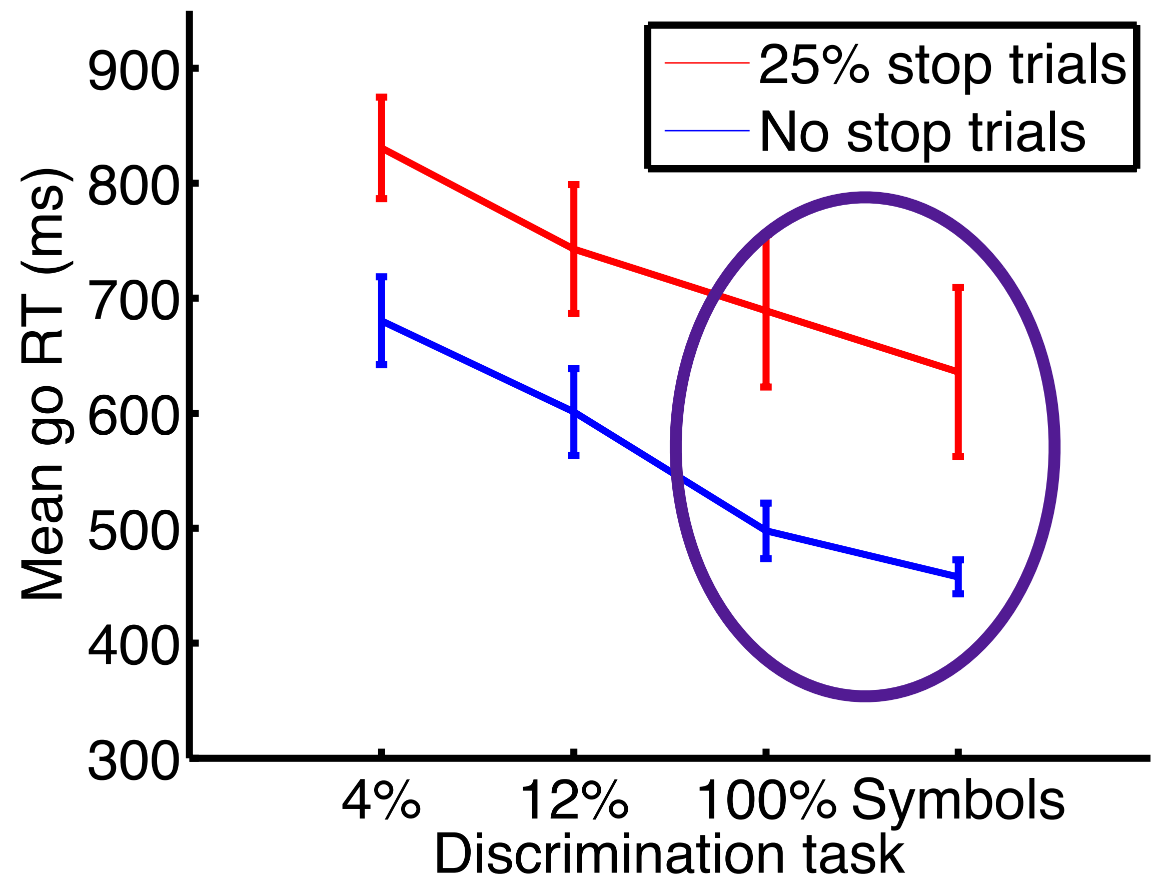

- ***Stop*** **trials induce** ***slowing*** **of** ***Go*** **RT**
- **Higher coherence (easier) induces faster** ***Go*** **response**
- **100% coherence similar to standard symbol discrimination**

Results: Harder *Go* Task Reduces *Stop* Errors

A. Model

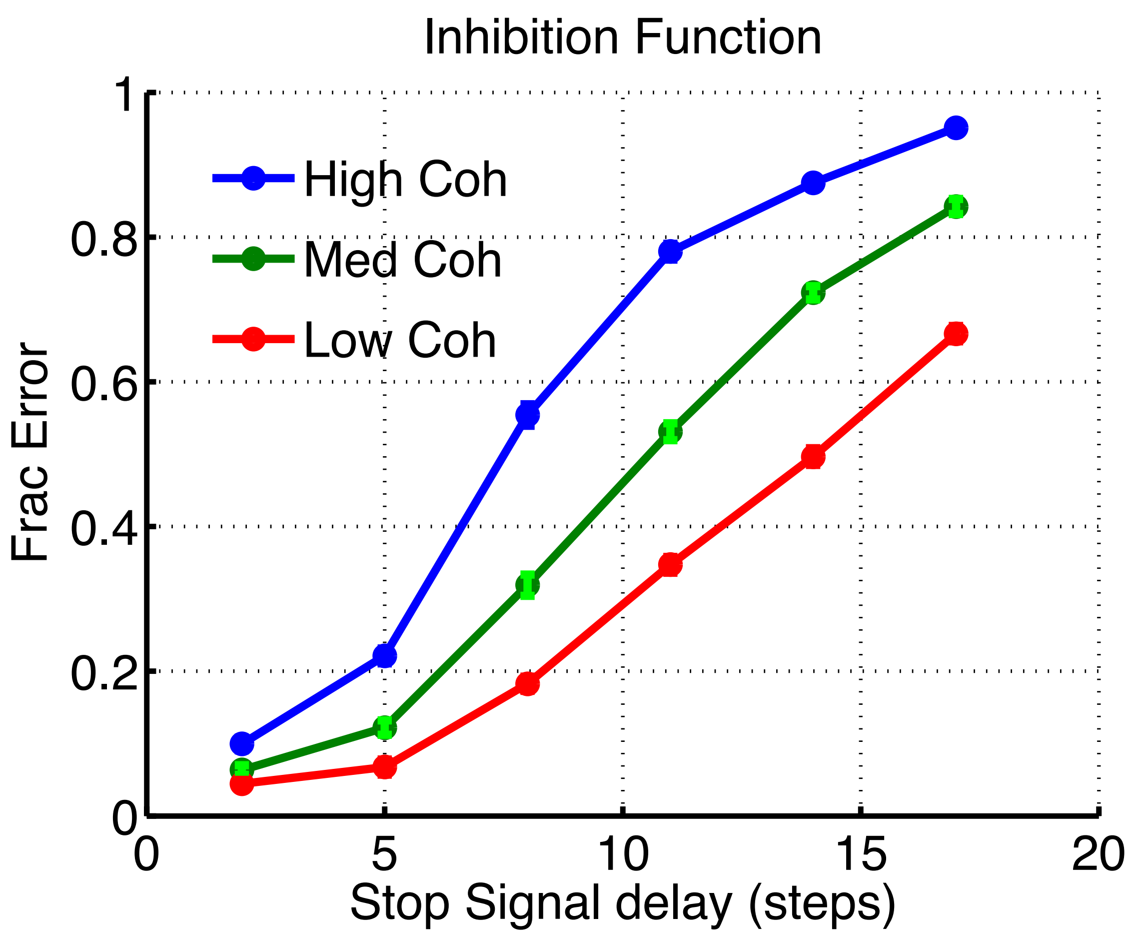

B. Data *n=5*

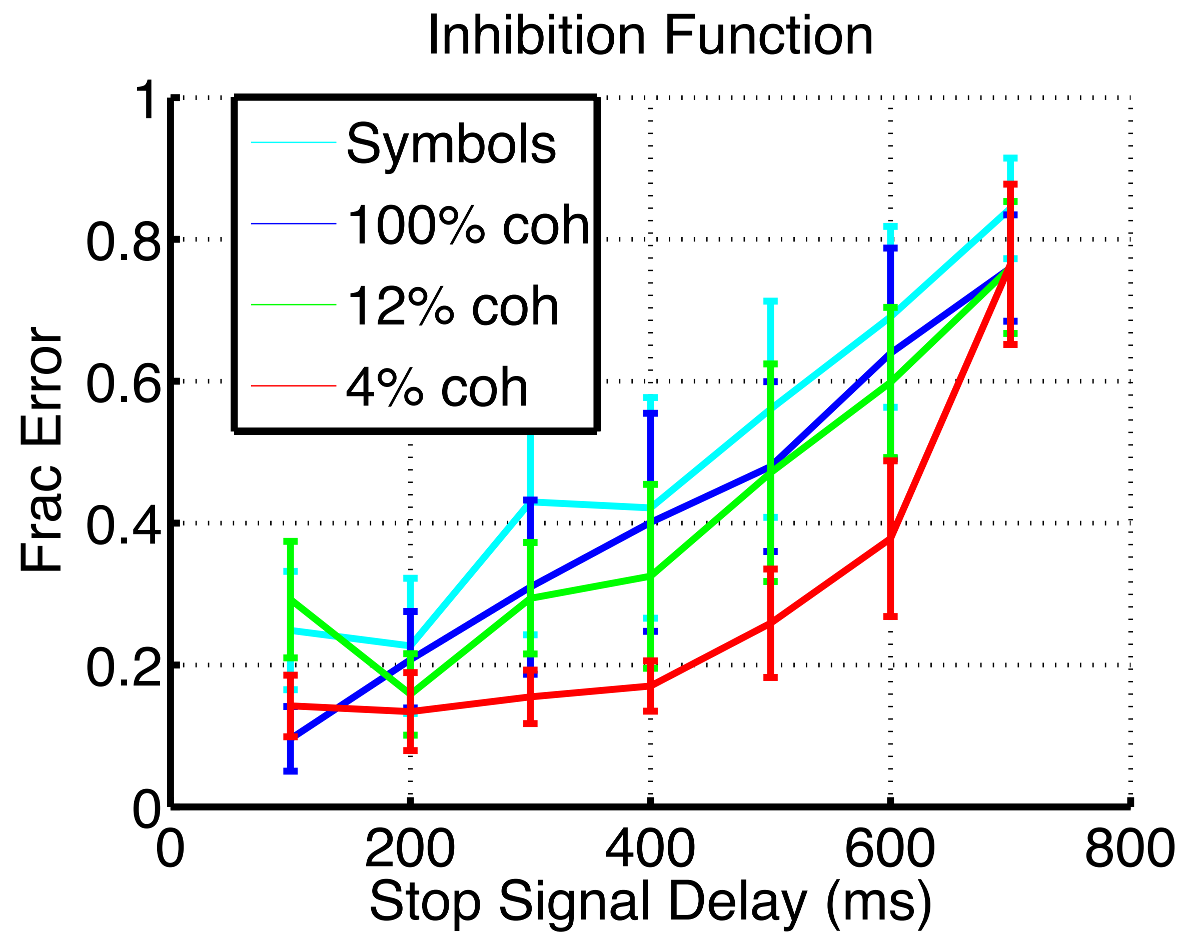

Results: Harder *Go* Task Reduces *Stop* Errors

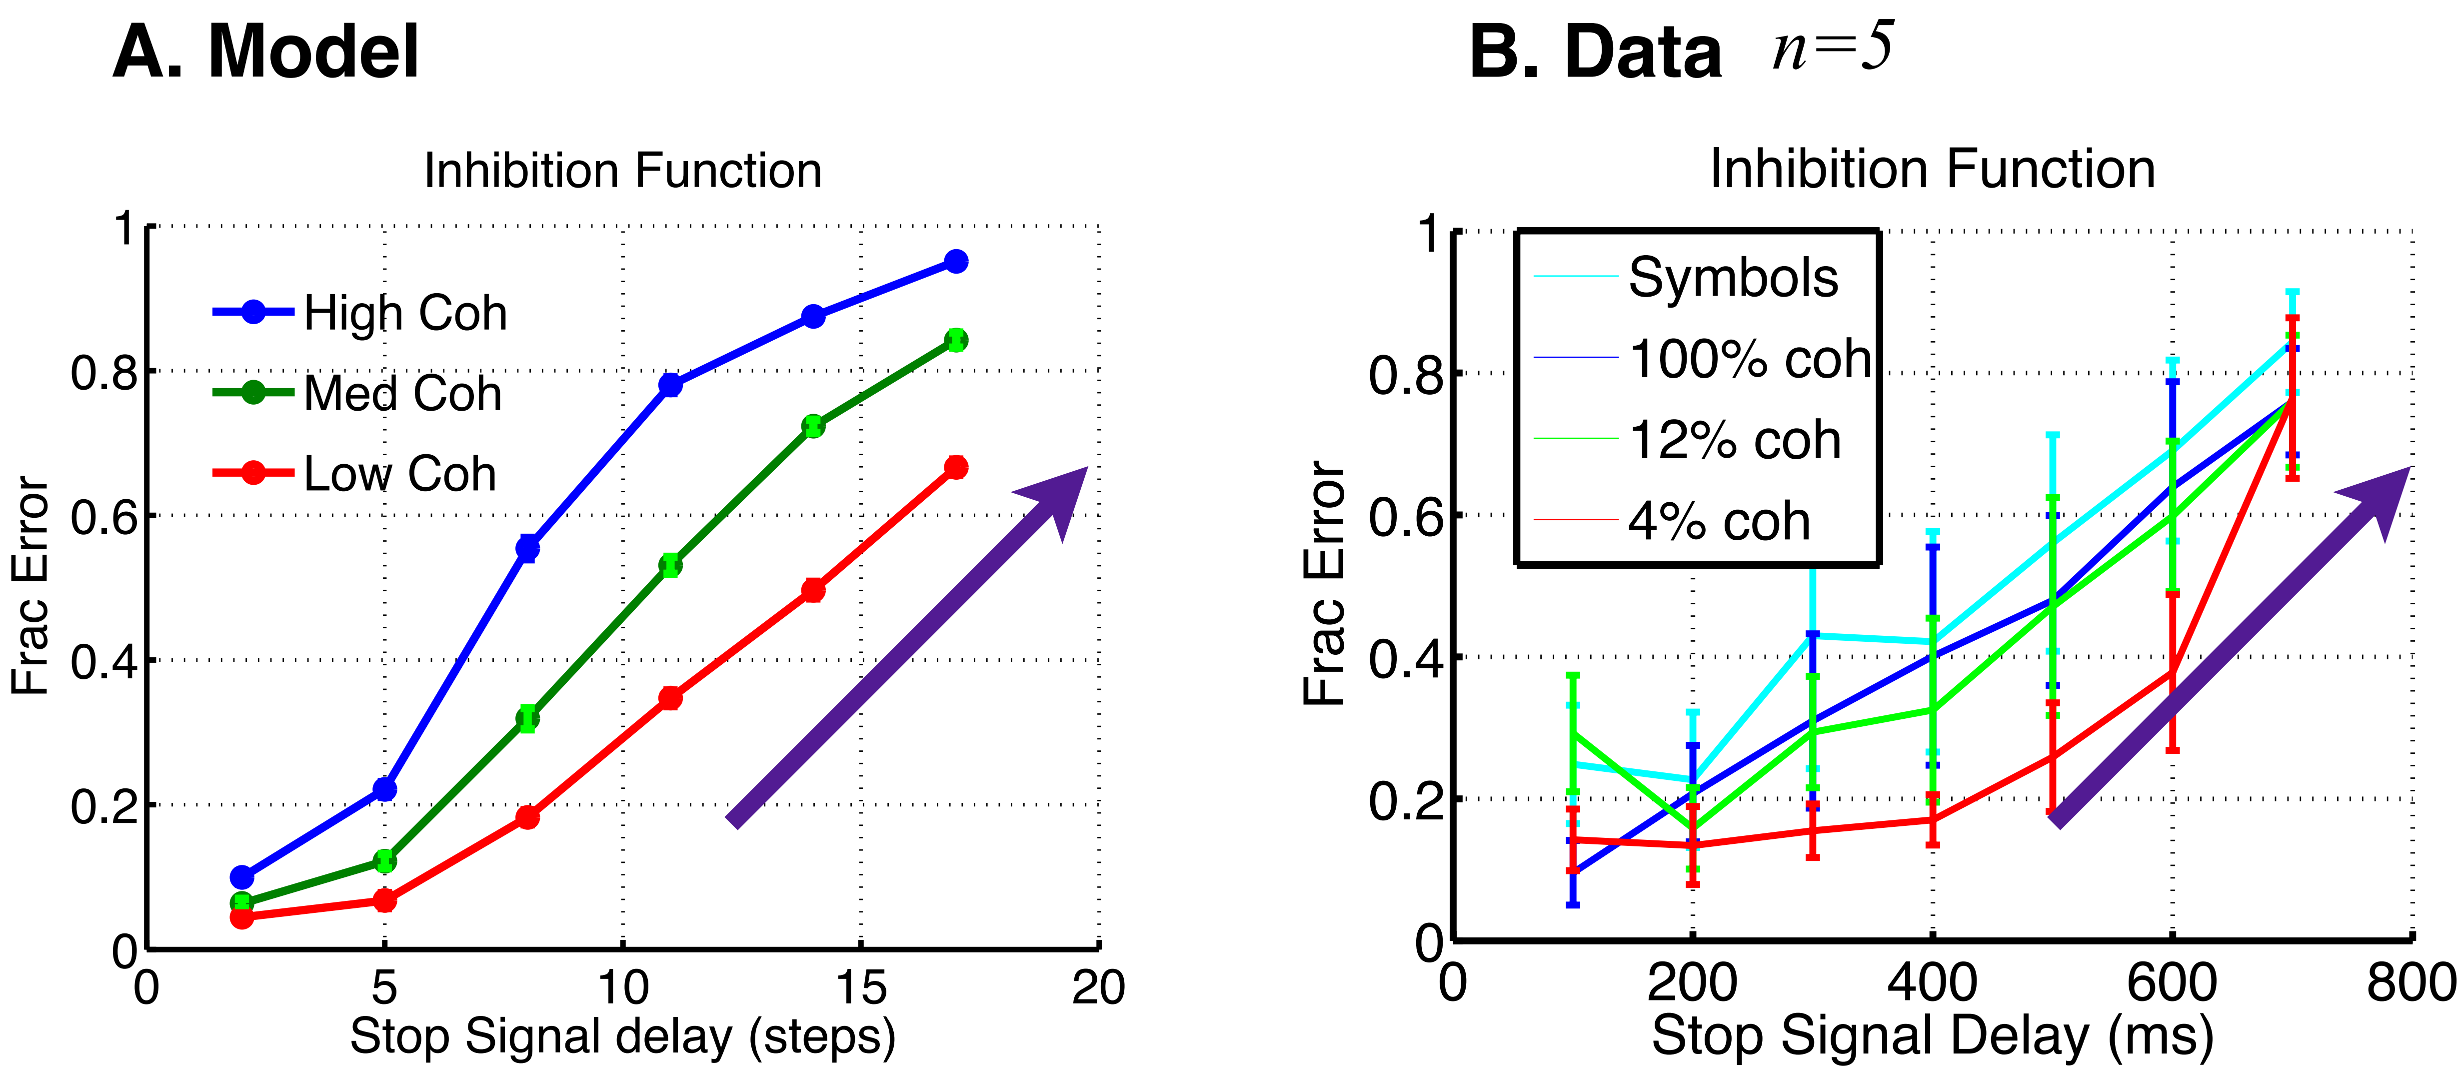

- Later stop signal ⇒ more stop errors

Results: Harder *Go* Task Reduces *Stop* Errors

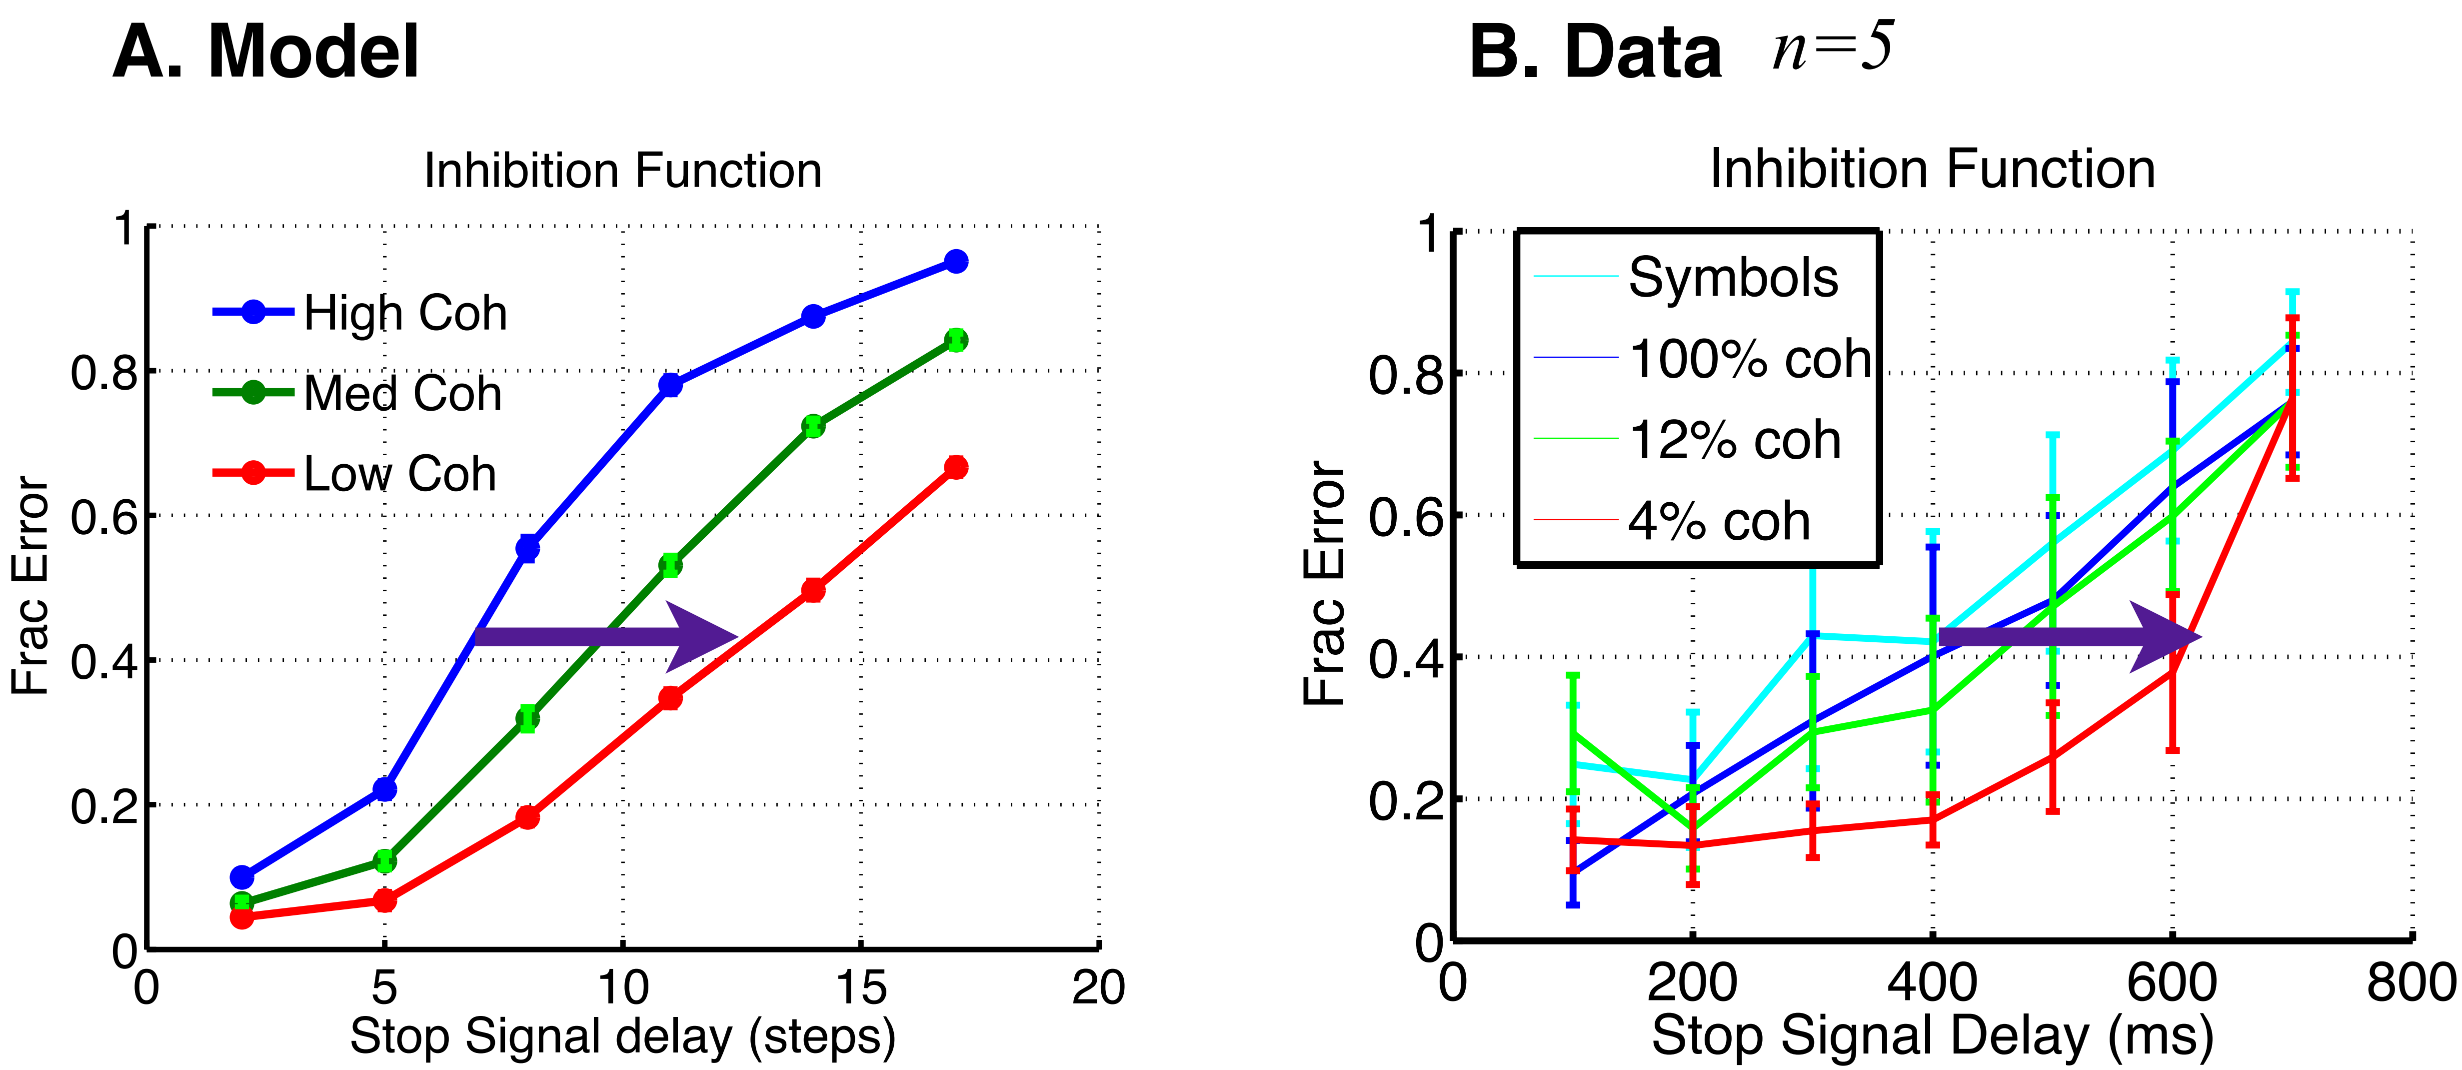

- Later stop signal ⇒ more stop errors
- More difficult ***go*** task ⇒ fewer stop errors

Results: Harder *Go* Task Reduces *Stop* Errors

A. Model

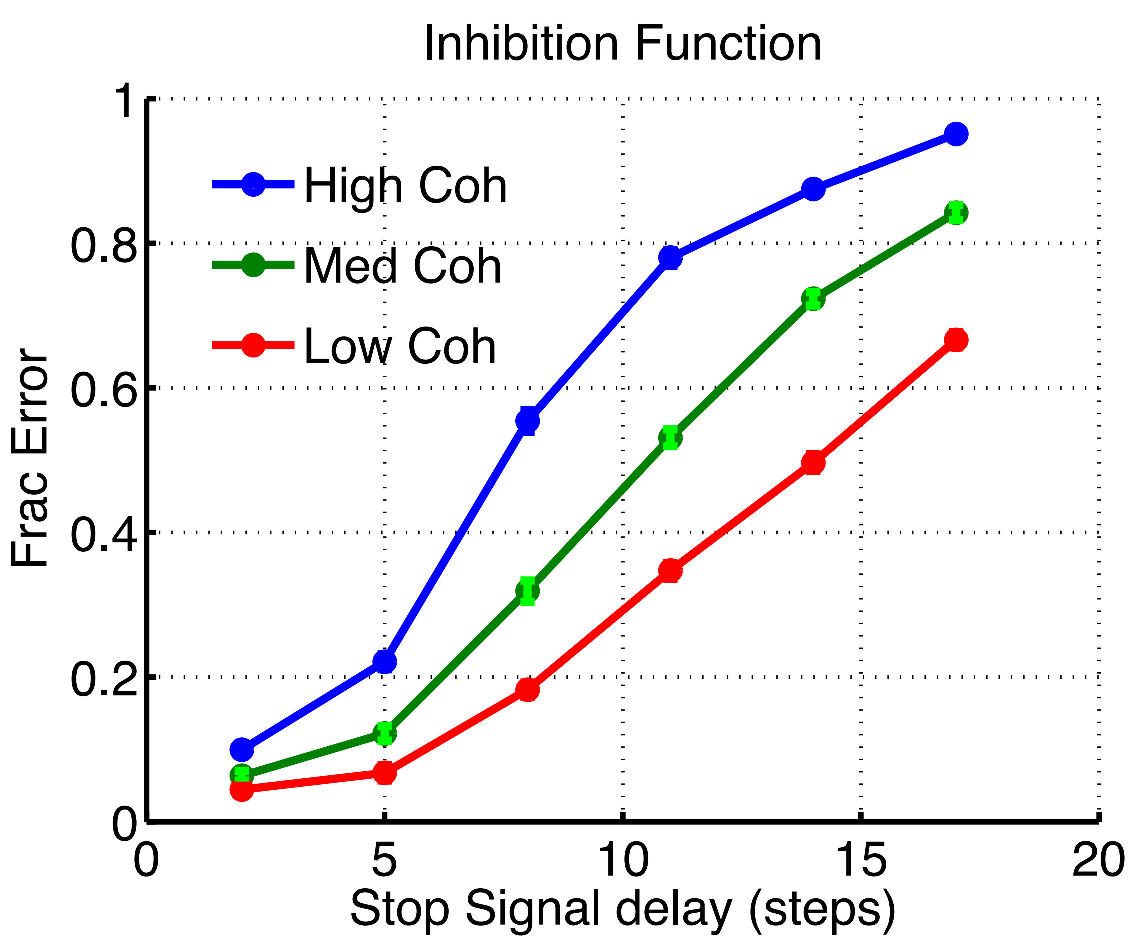

B. Data *n=5*

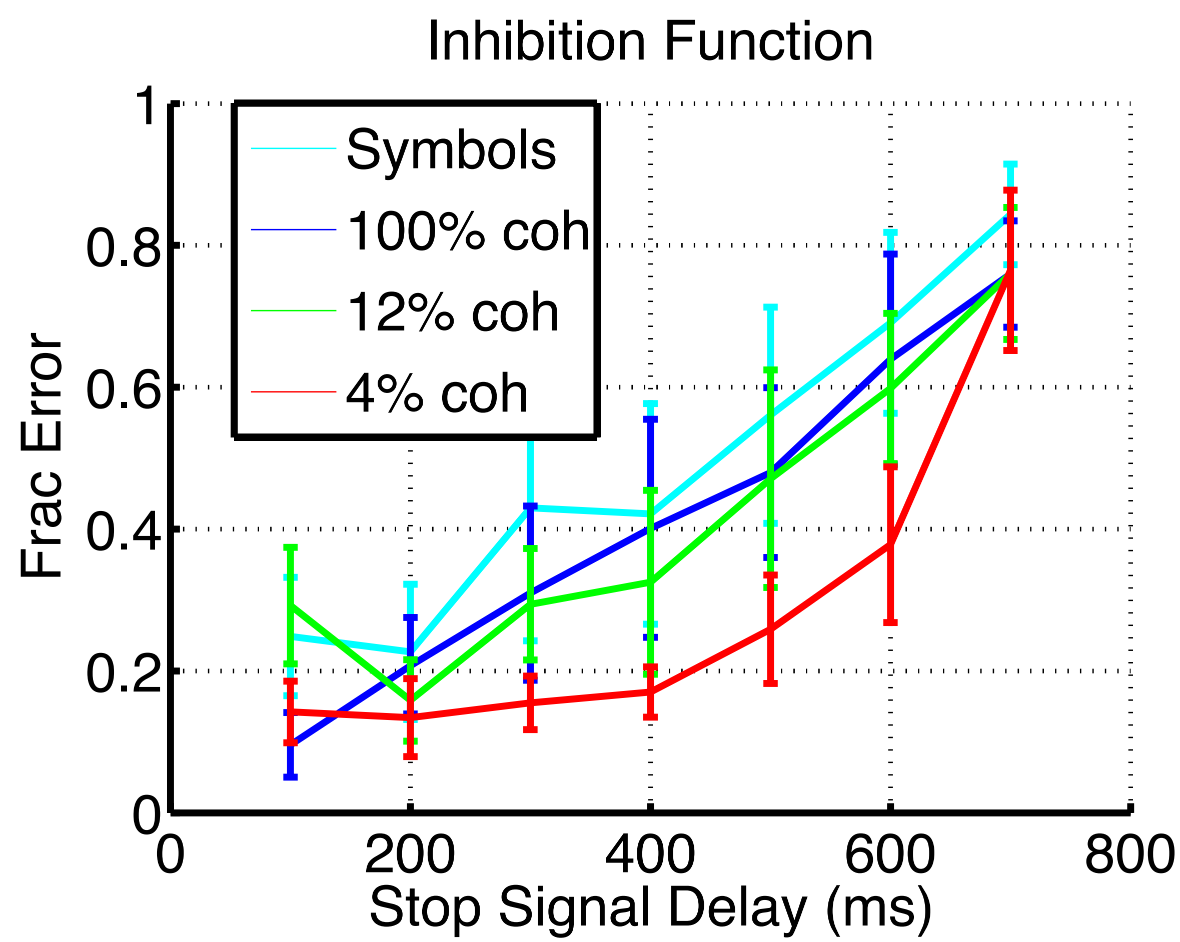

- Later stop signal ⇒ more stop errors
- More difficult ***go*** task ⇒ fewer stop errors
- Consequence of slower ***go*** RT

Results: Harder *Go* Task Increases *SSRT*

A. Model

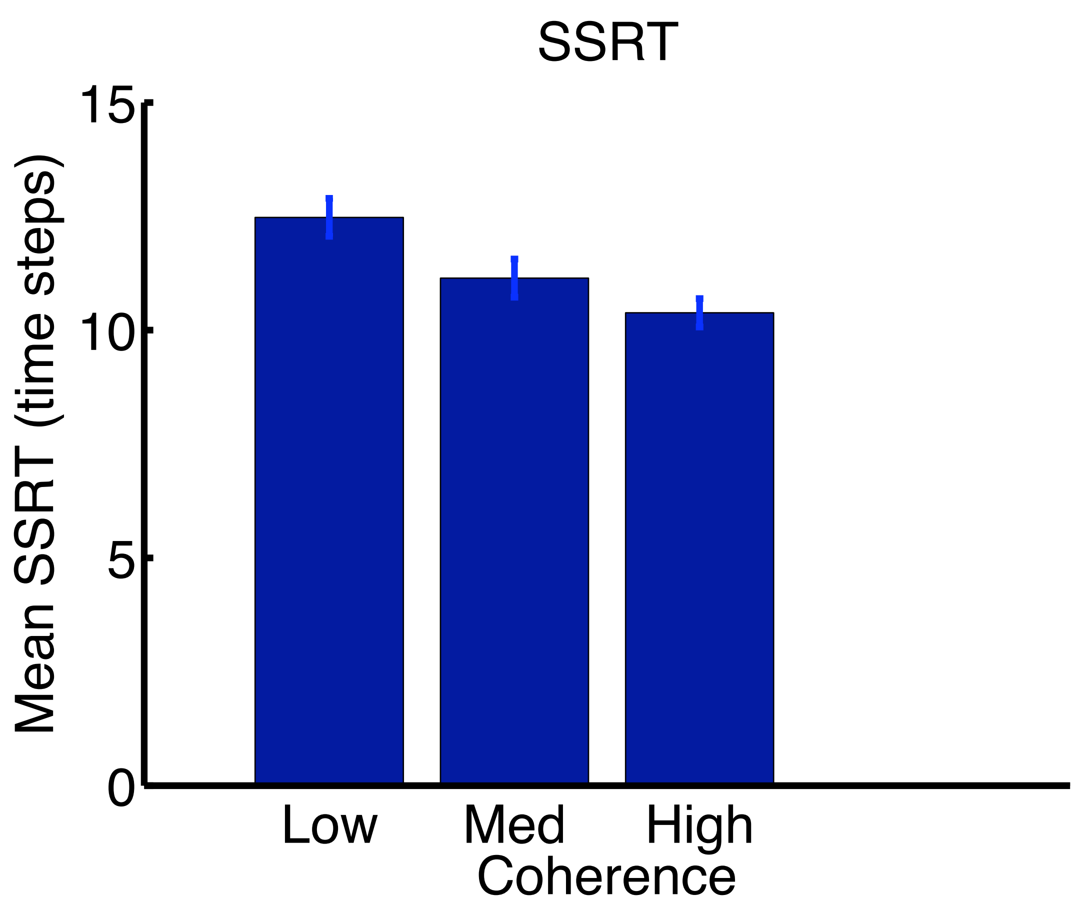

Results: Harder *Go* Task Increases *SSRT*

A. Model

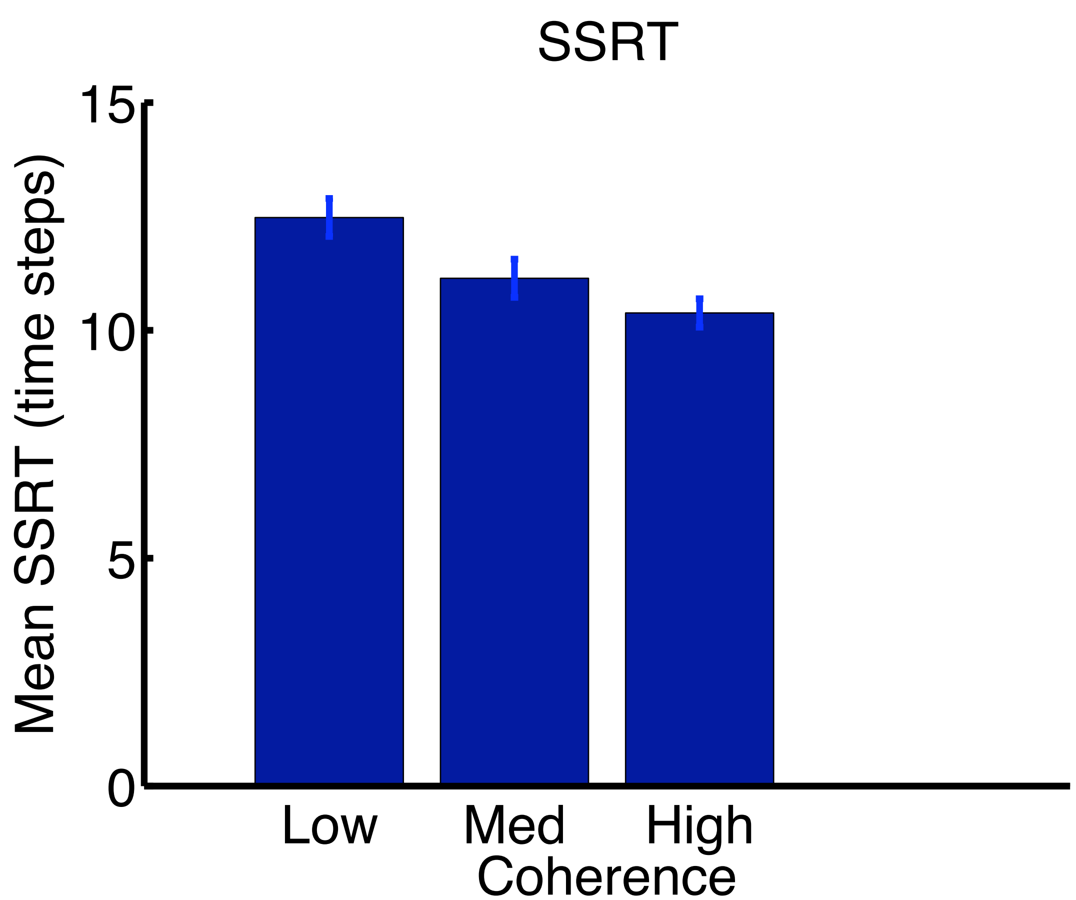

- Optimal model: difficult ***go*** processing lengthens ***SSRT***

Results: Harder *Go* Task Increases *SSRT*

A. Model

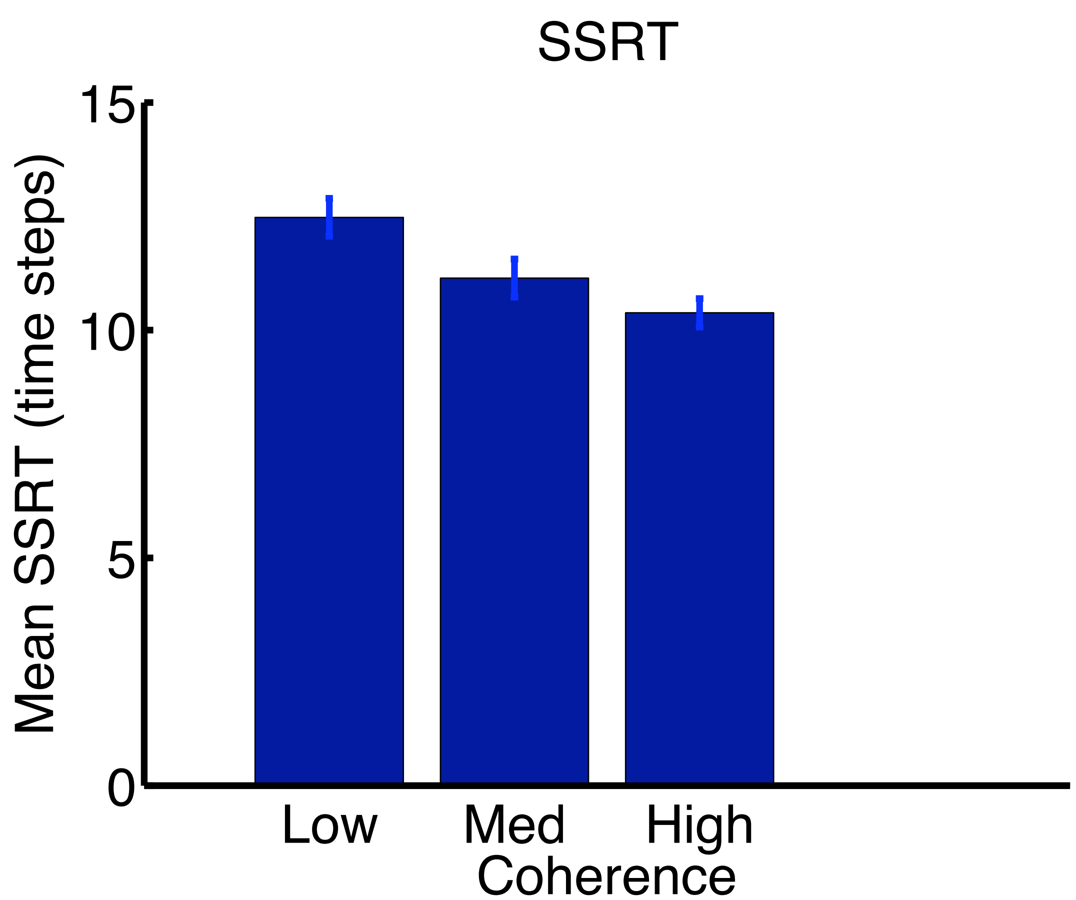

- Optimal model: difficult ***go*** processing lengthens ***SSRT***
- Race model: ***go*** & ***stop*** independent processes

Results: Harder *Go* Task Increases *SSRT*

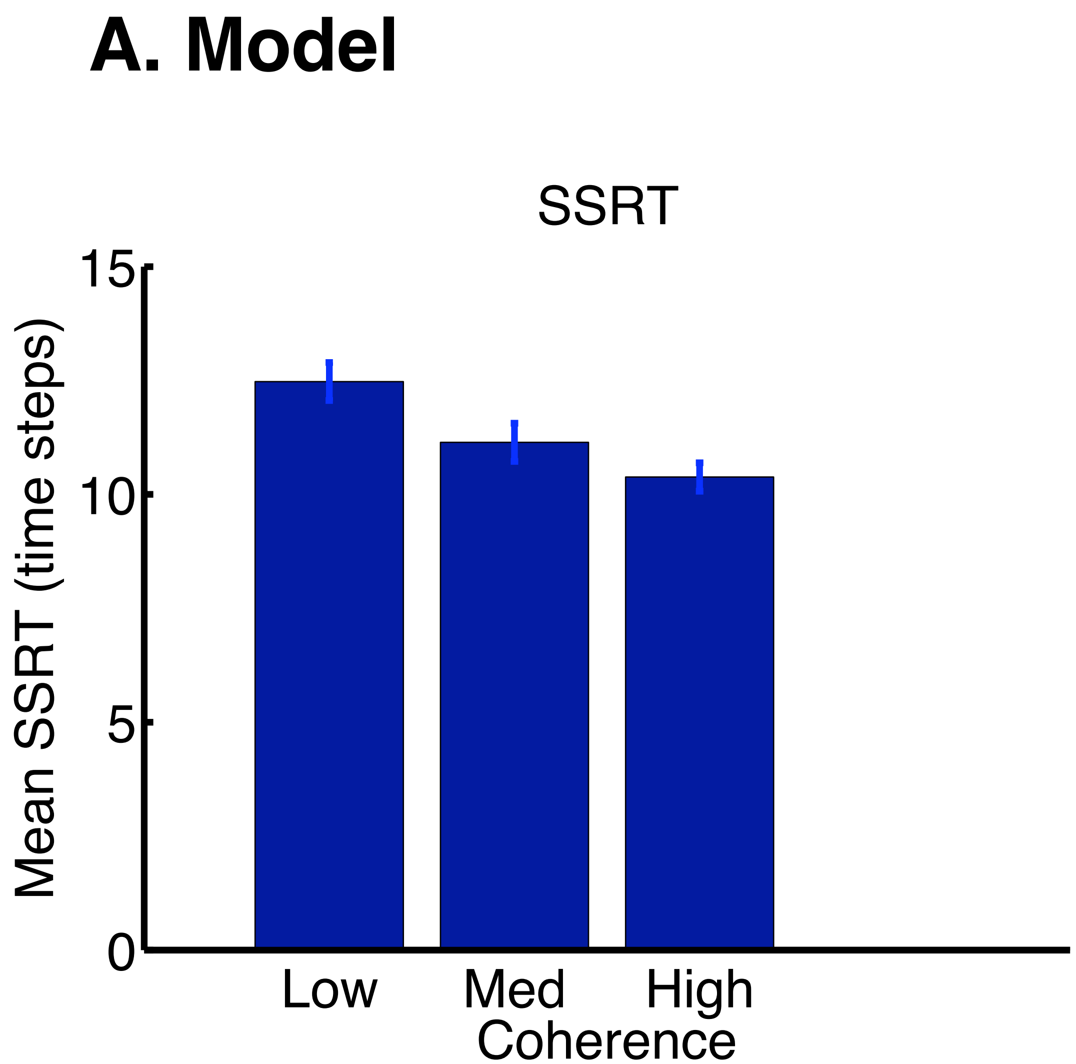

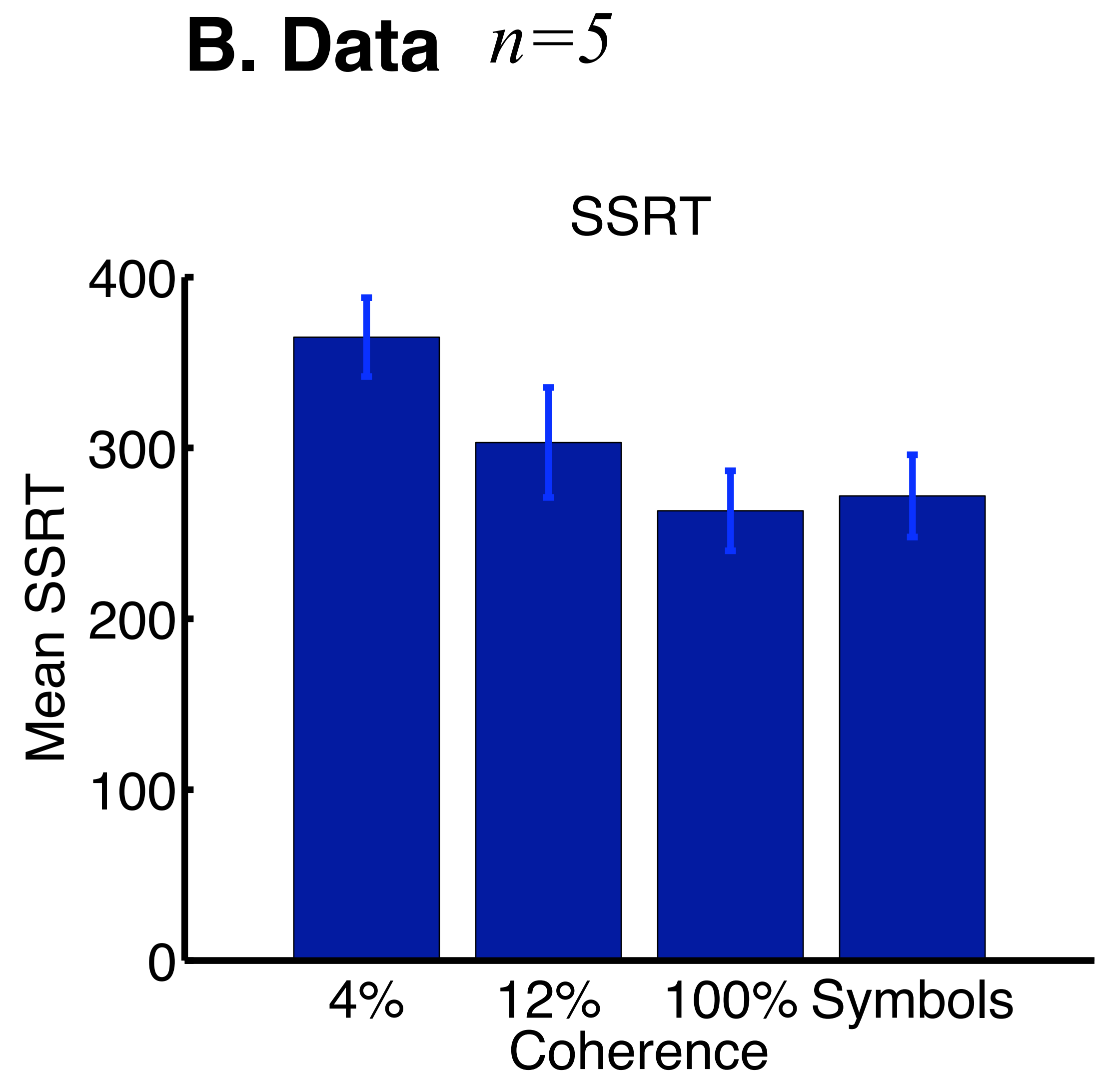

- Optimal model: difficult ***go*** processing lengthens ***SSRT***
- Race model: ***go*** & ***stop*** independent processes
- ***SSRT*** increase ⇒ ***go*** & ***stop*** processes fundamentally inter-related

Results: Harder *Go* Task Increases *SSRT*

A. Model

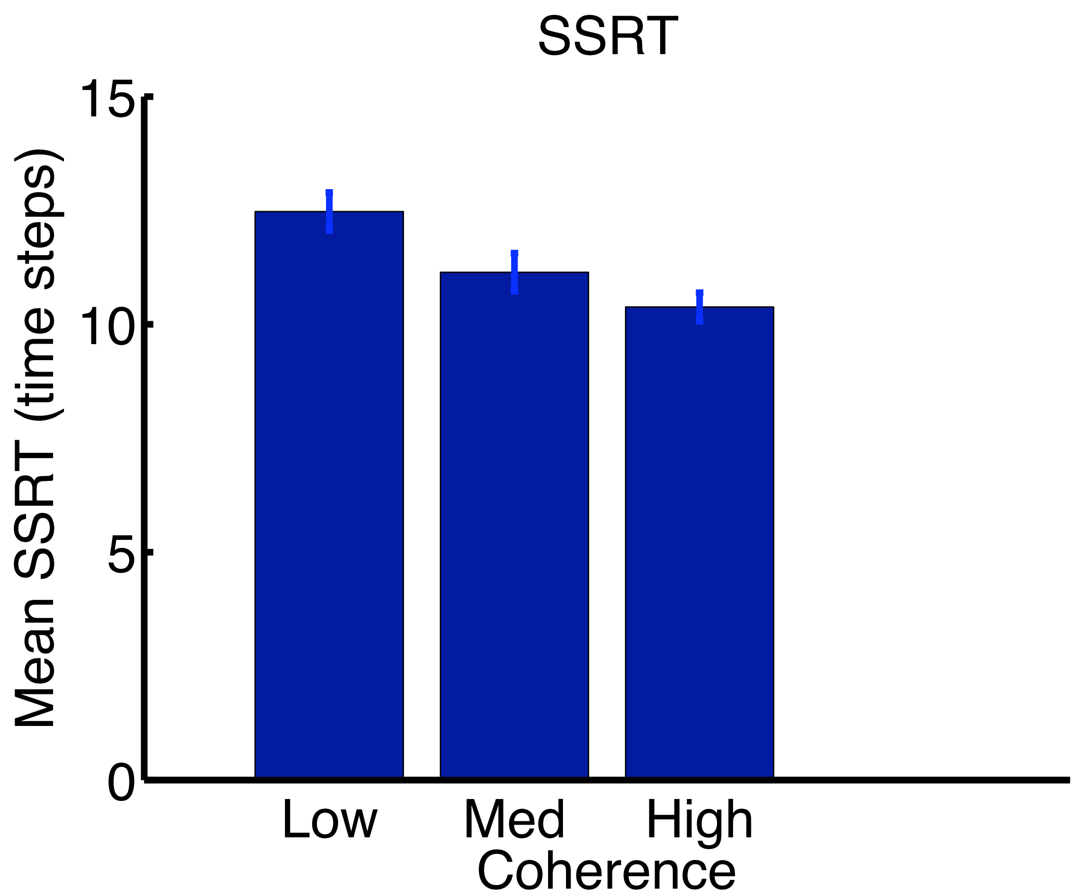

B. Data *n=5*

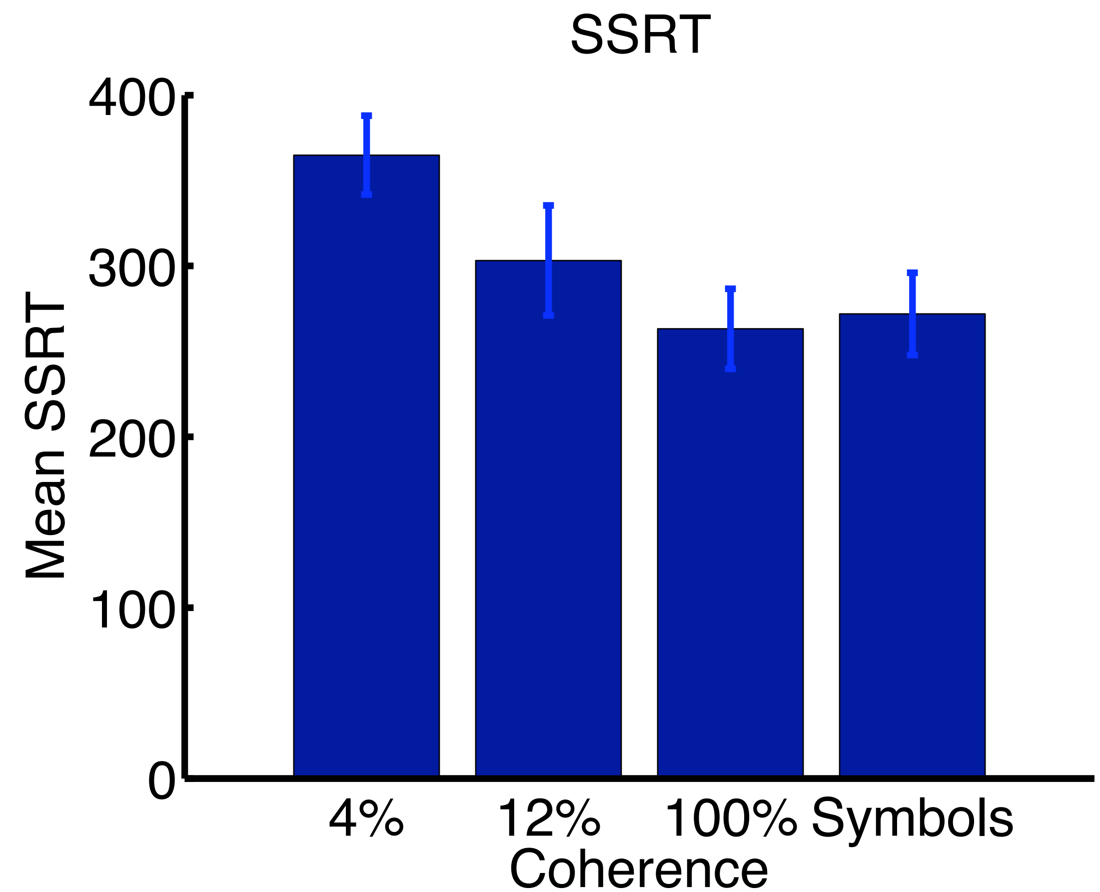

- Optimal model: difficult ***go*** processing lengthens ***SSRT***
- Race model: ***go*** & ***stop*** independent processes
- ***SSRT*** increase ⇒ ***go*** & ***stop*** processes fundamentally inter-related

Race Model

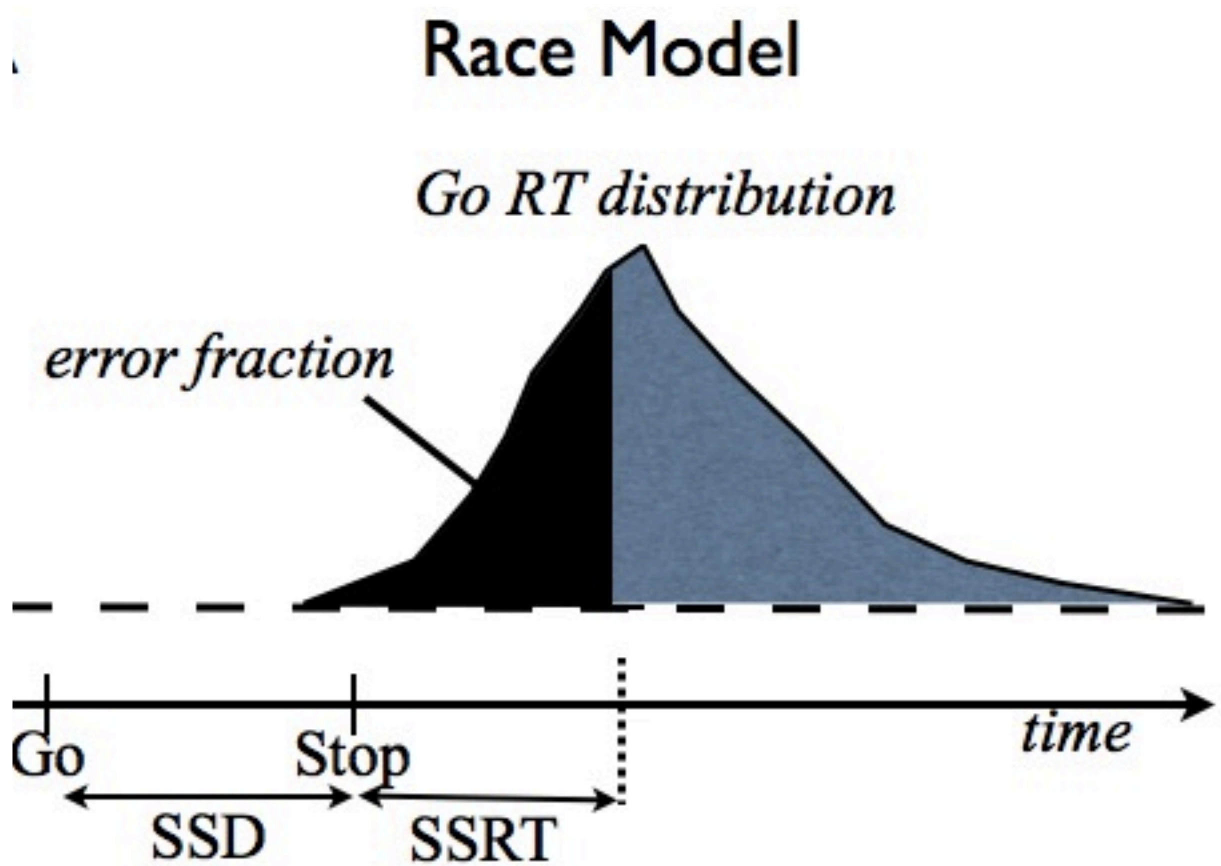

(Logan & Cowan, 1984)

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

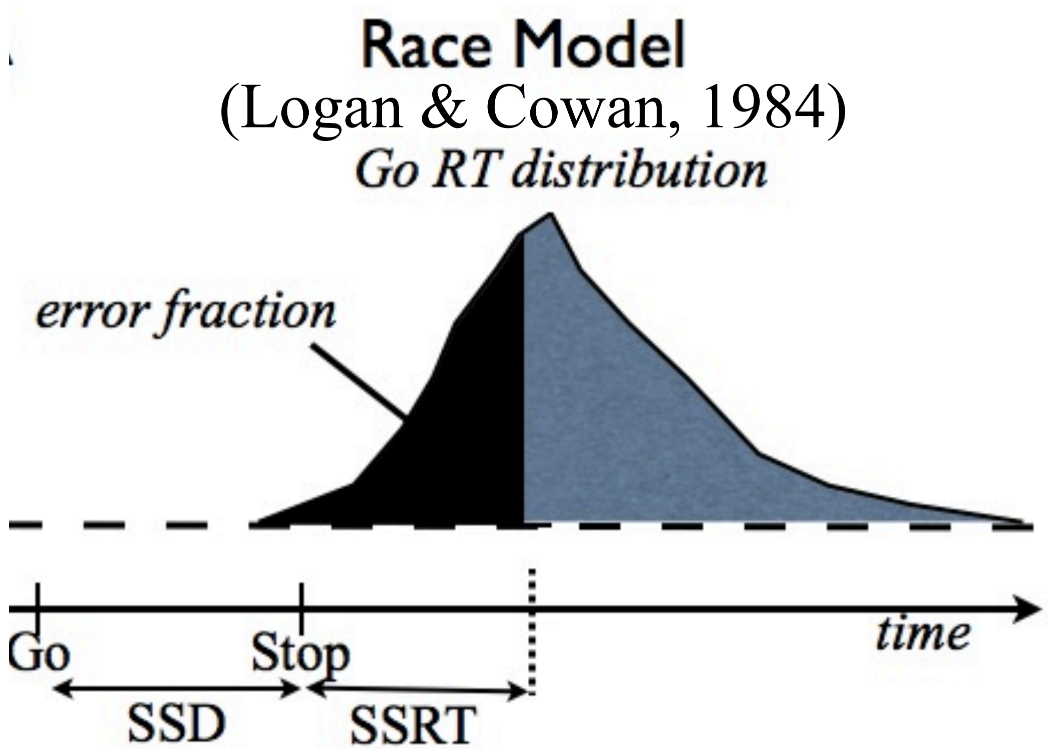

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

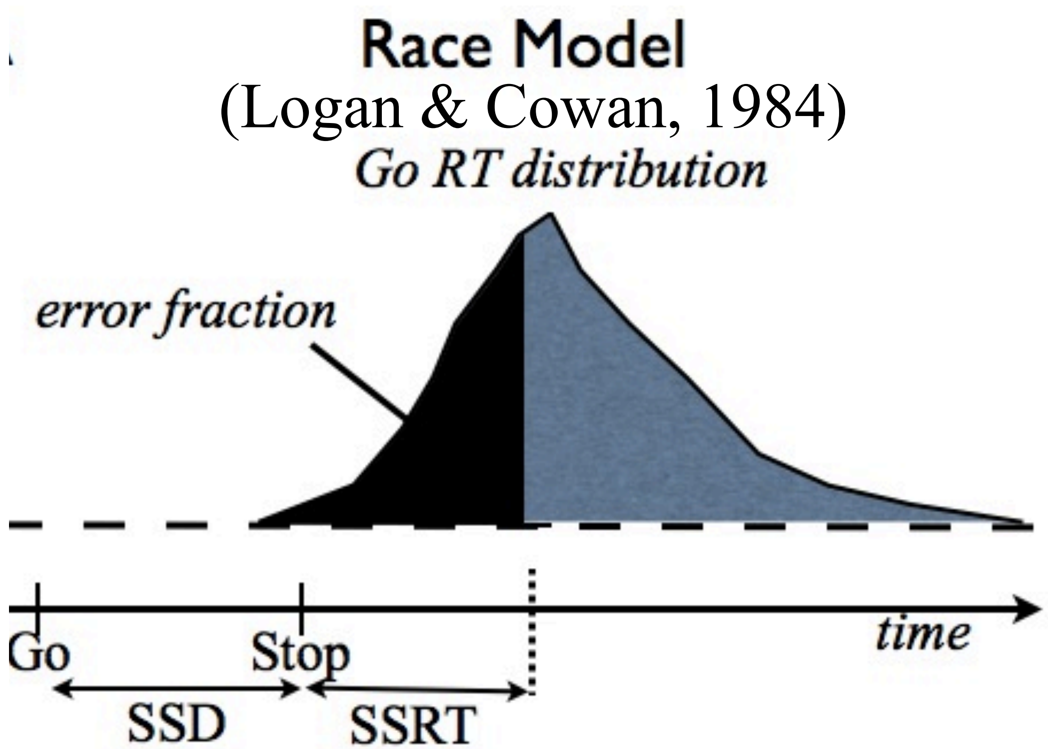

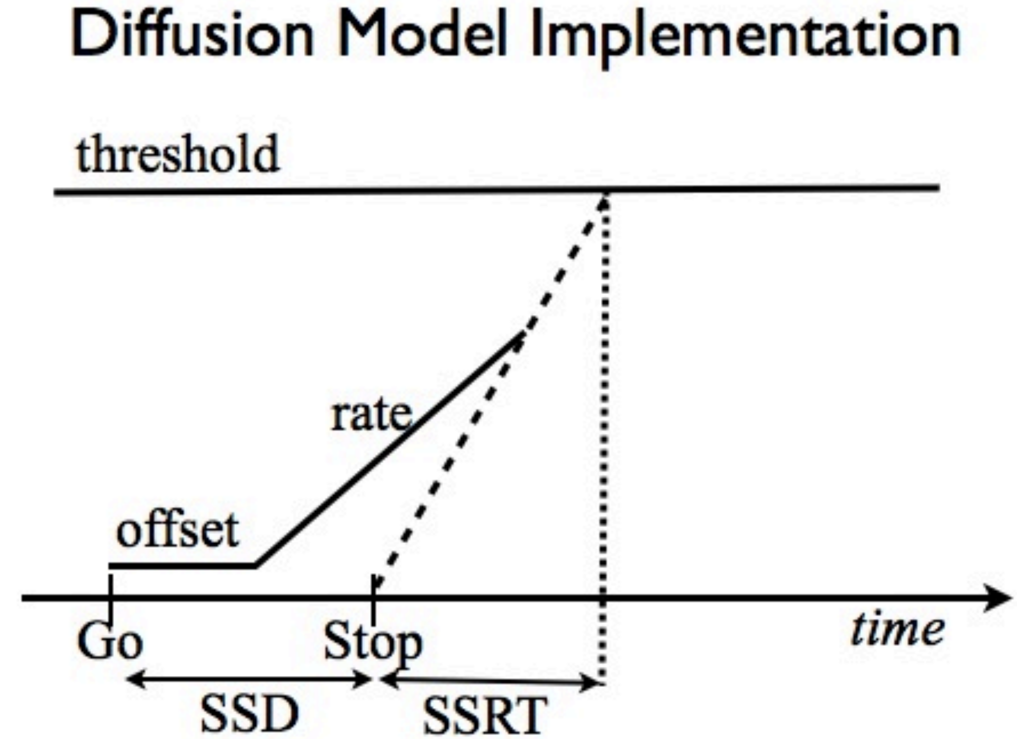

Contextual Effects ⇒ Stopping Behavior

Race model Approximation to Optimal DM

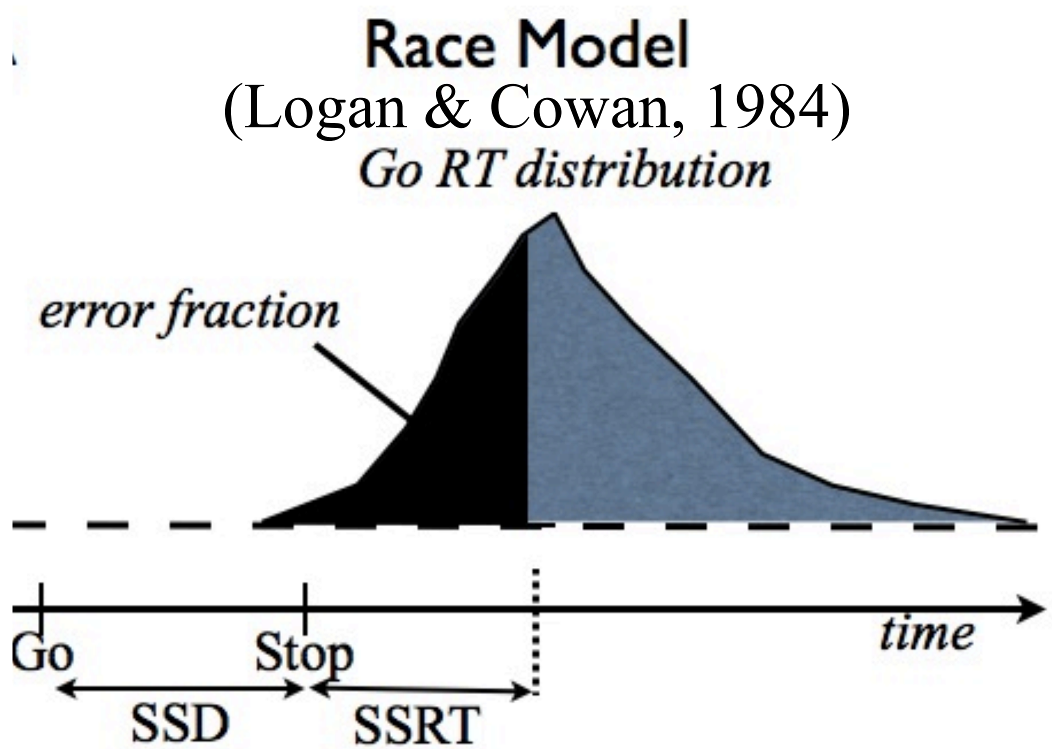

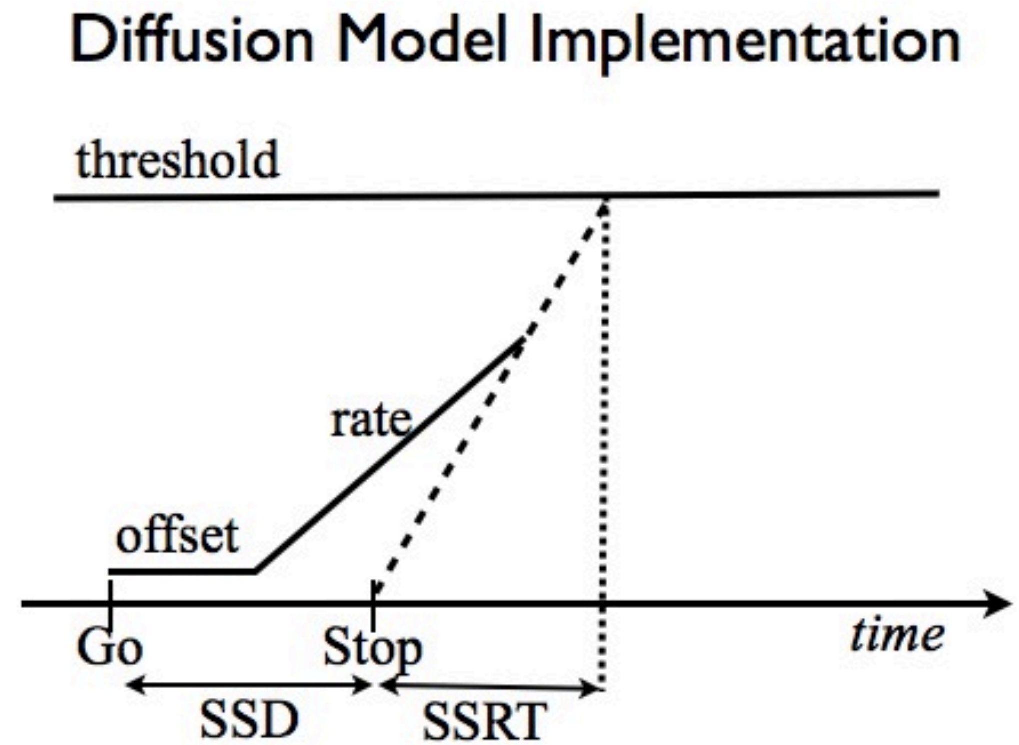

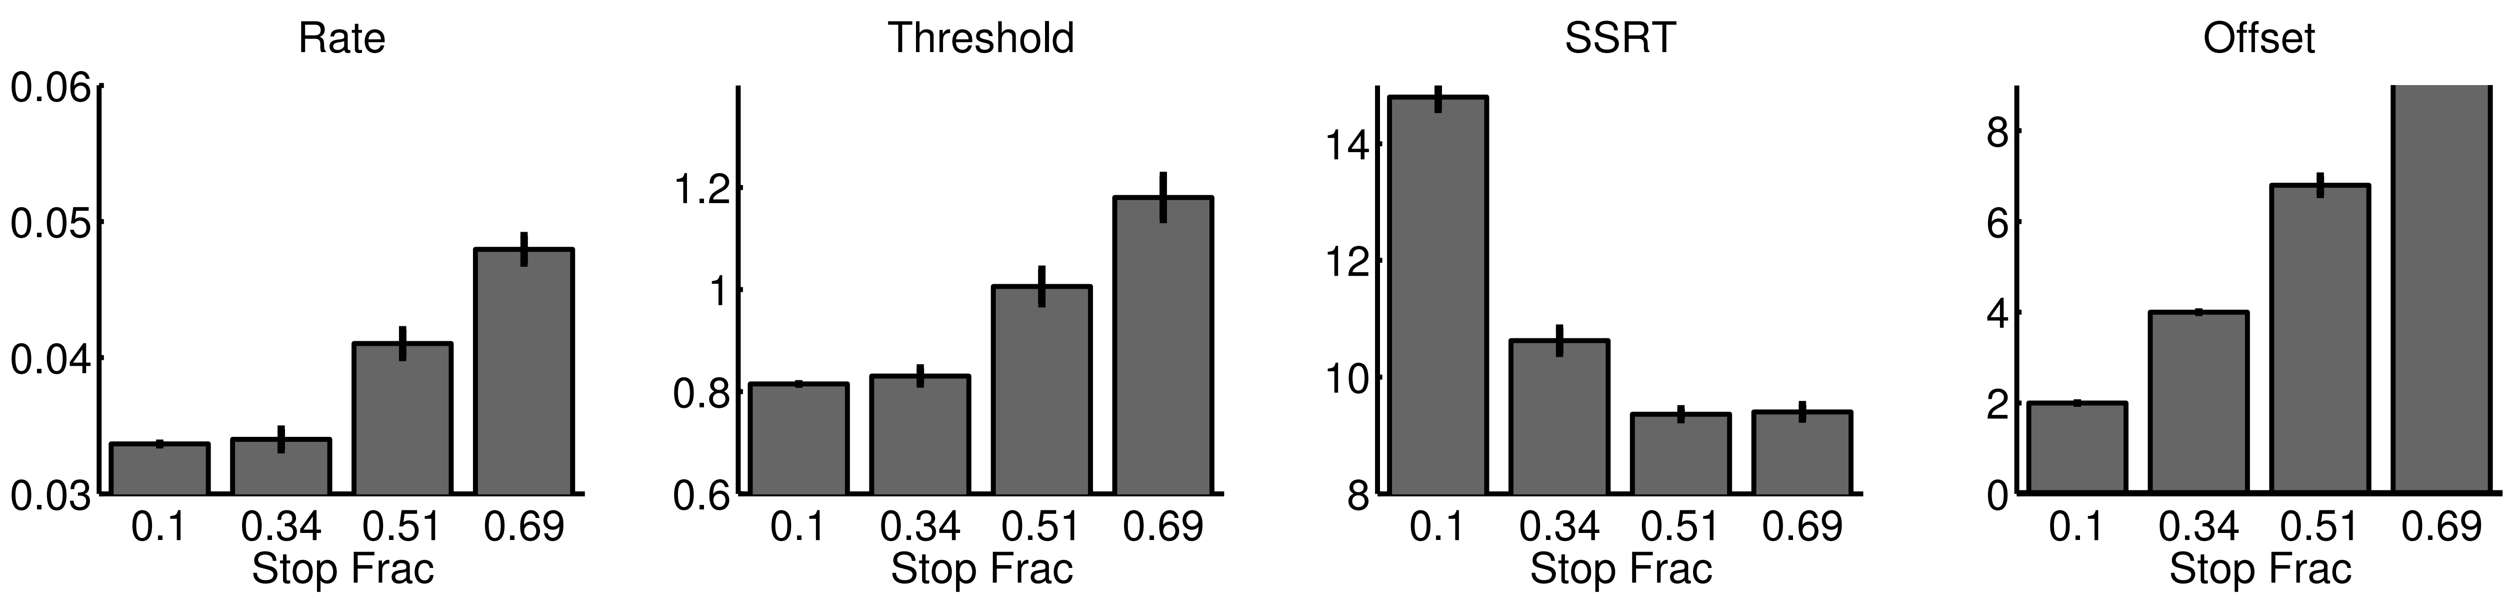

SSRT As a Measure of Inhibitory Control

SSRT As a Measure of Inhibitory Control

- SSRT slower in populations with inhibitory deficits
  - e.g., ADHD, substance abuse, OCD: longer SSRT

SSRT As a Measure of Inhibitory Control

- SSRT slower in populations with inhibitory deficits
  - e.g., ADHD, substance abuse, OCD: longer SSRT
- Behavioral SSRT linked closely to neural activity
  - Neural response in frontal eye field, superior colliculus (Hanes et al., 1996, Pare & Hanes, 2003)
  - Suggests a *neural mechanism* underlying stopping behavior

Race Model: Wherefore the Race?

Race Model: Wherefore the Race?

- Silent on how SSRT arises from or recruits different facets of cognitive processing
  - sensory processing, attention, decision-making, learning

Race Model: Wherefore the Race?

- Silent on how SSRT arises from or recruits different facets of cognitive processing
  - sensory processing, attention, decision-making, learning
- Cannot explain *how* and *why* SSRT changes with experimental parameters
  - motivational factors (Leotti & Wager, 2009)
  - fraction of stop trials, immediate history (Emeric et al, 2007)
  - go response difficulty (Logan et al, JEP:HPP, 1984)

Race Model: Wherefore the Race?

- Silent on how SSRT arises from or recruits different facets of cognitive processing
  - sensory processing, attention, decision-making, learning
- Cannot explain *how* and *why* SSRT changes with experimental parameters
  - motivational factors (Leotti & Wager, 2009)
  - fraction of stop trials, immediate history (Emeric et al, 2007)
  - go response difficulty (Logan et al, JEP:HPP, 1984)
- Single parameter (SSRT) cannot explain full range of data

Experiment: Vary *Go* Discrimination Difficulty

Experiment: Vary *Go* Discrimination Difficulty

- Race model: ***go*** & ***stop*** independent processes

Experiment: Vary *Go* Discrimination Difficulty

- Race model: ***go*** & ***stop*** independent processes
- Optimality model: interaction b/t ***go*** & ***stop*** processing

  $\Rightarrow$ difficulty of ***go*** processing improves stopping

Experiment: Vary *Go* Discrimination Difficulty

- Race model: ***go*** & ***stop*** independent processes
- Optimality model: interaction b/t ***go*** & ***stop*** processing
  $\Rightarrow$ difficulty of ***go*** processing improves stopping
- Design: *random-dot motion task* -- coherence controls difficulty

Experiment: Vary *Go* Discrimination Difficulty

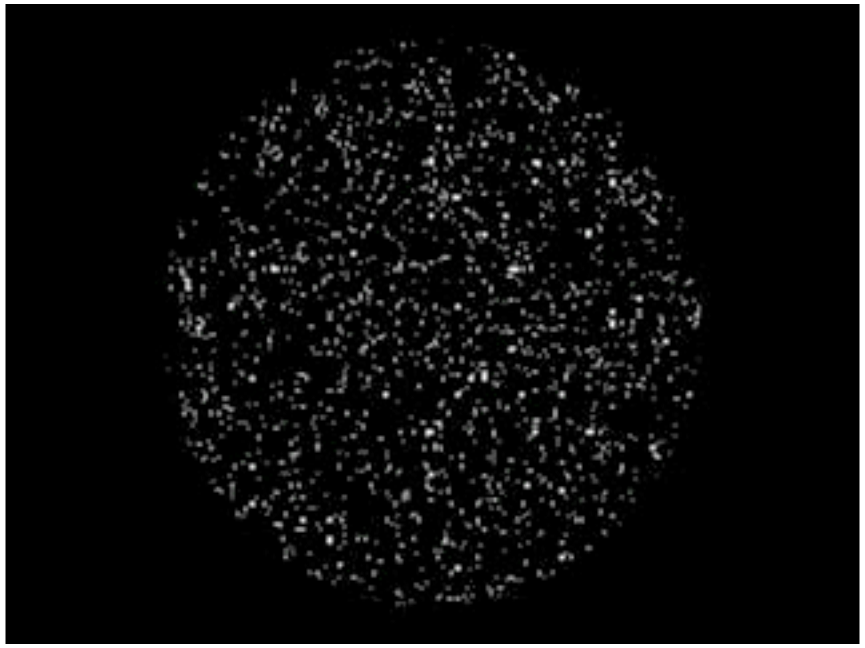

30% coh

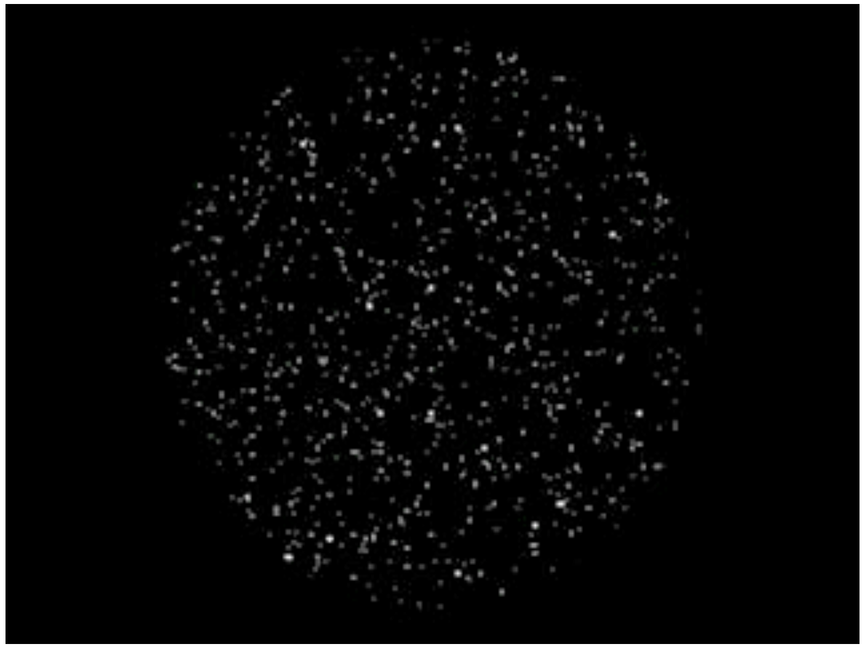

5% coh

- Race model: ***go*** & ***stop*** independent processes
- Optimality model: interaction b/t ***go*** & ***stop*** processing
  $\Rightarrow$ difficulty of ***go*** processing improves stopping
- Design: *random-dot motion task* -- coherence controls difficulty

Race -- Diffusion Model

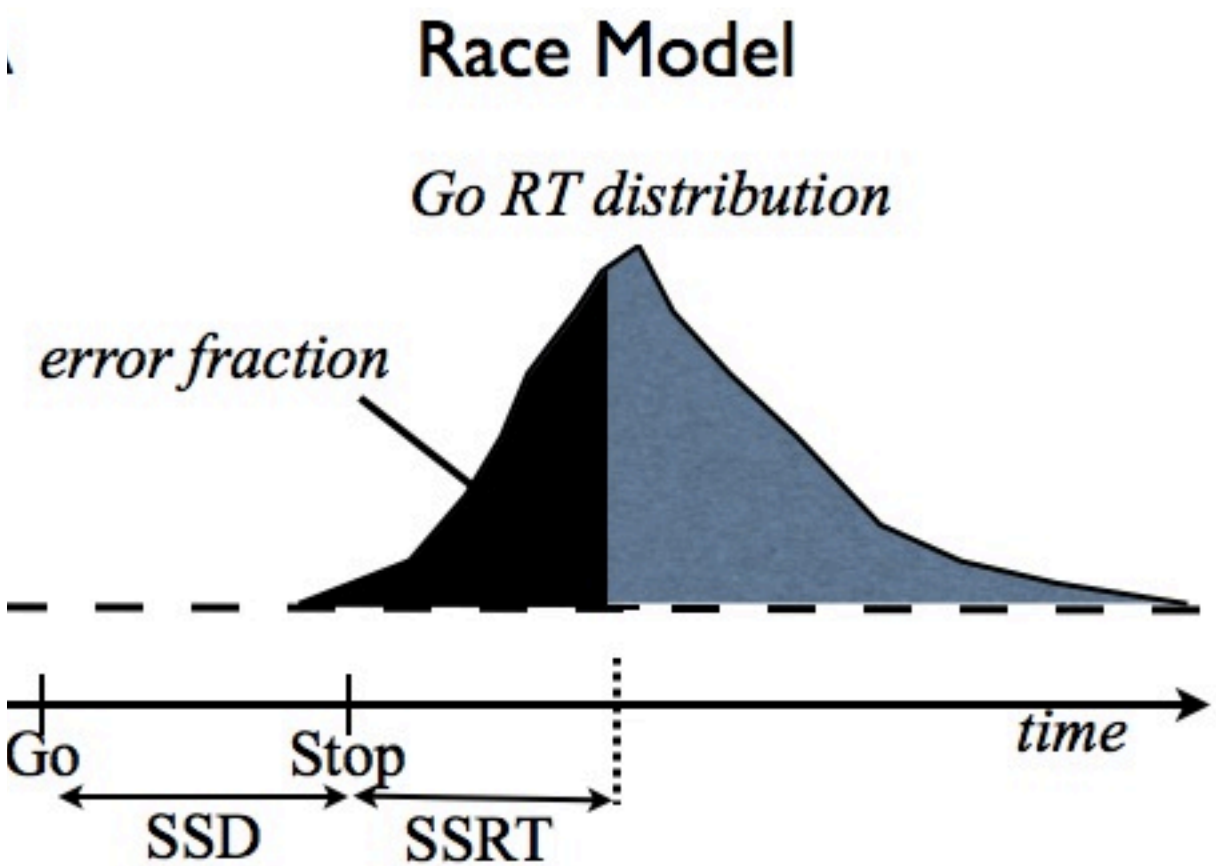

(Logan & Cowan, 1984)

Race -- Diffusion Model

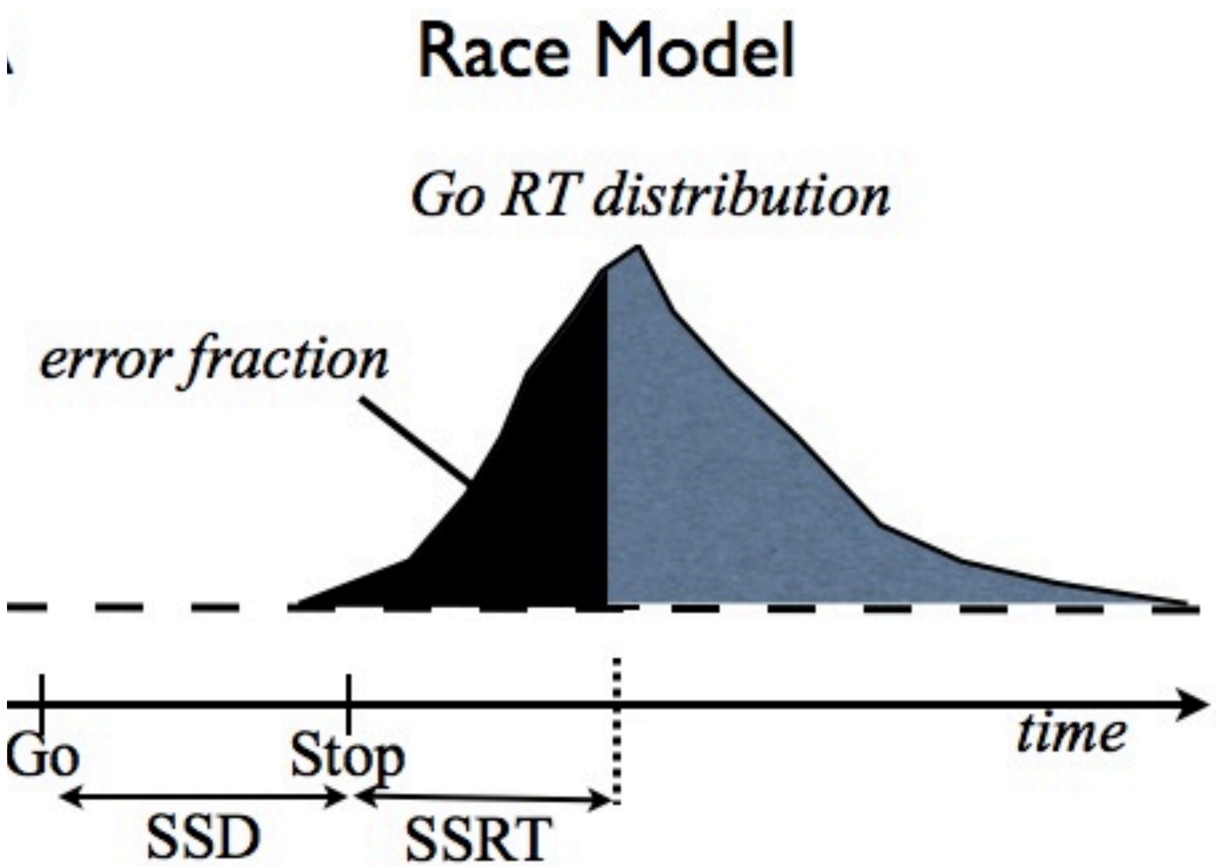

(Logan & Cowan, 1984)

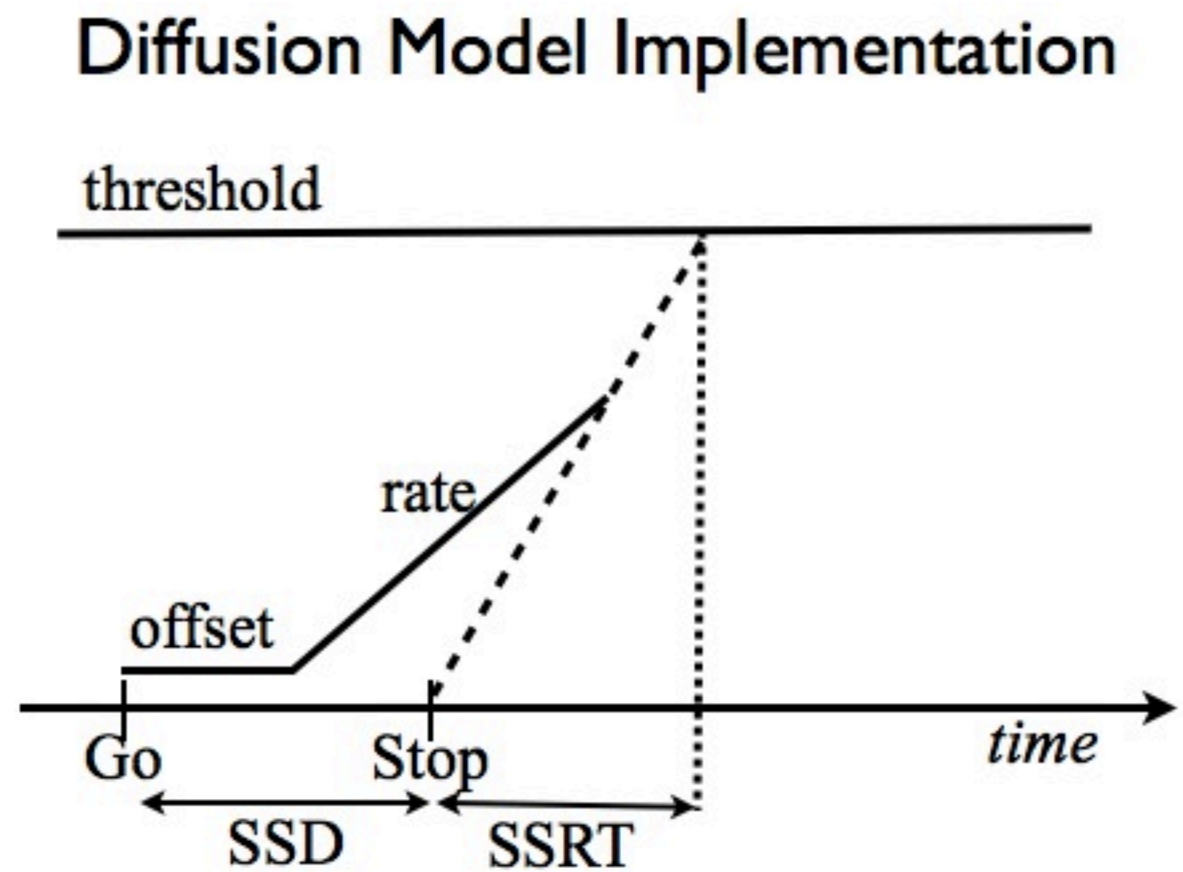

(e.g. Verbruggen & Logan, 2007)

Rational Framework for Stop Signal Task

Fundamental decision: when (whether) to go?

Reasons to go **fast**	Reasons to go **slow**
sensory	**sensory**
• *go* discrimination easy	• *go* discrimination difficult
prior knowledge	**prior knowledge**
• stop signal rare	• stop signal frequent
costs	**costs**
• time cost high	• time cost low
• deadline penalty large	• deadline penalty small
• stop error penalty small	• stop error penalty large

Decision Policy: *Go* & *Wait* Regions of Belief State

Simulation Results

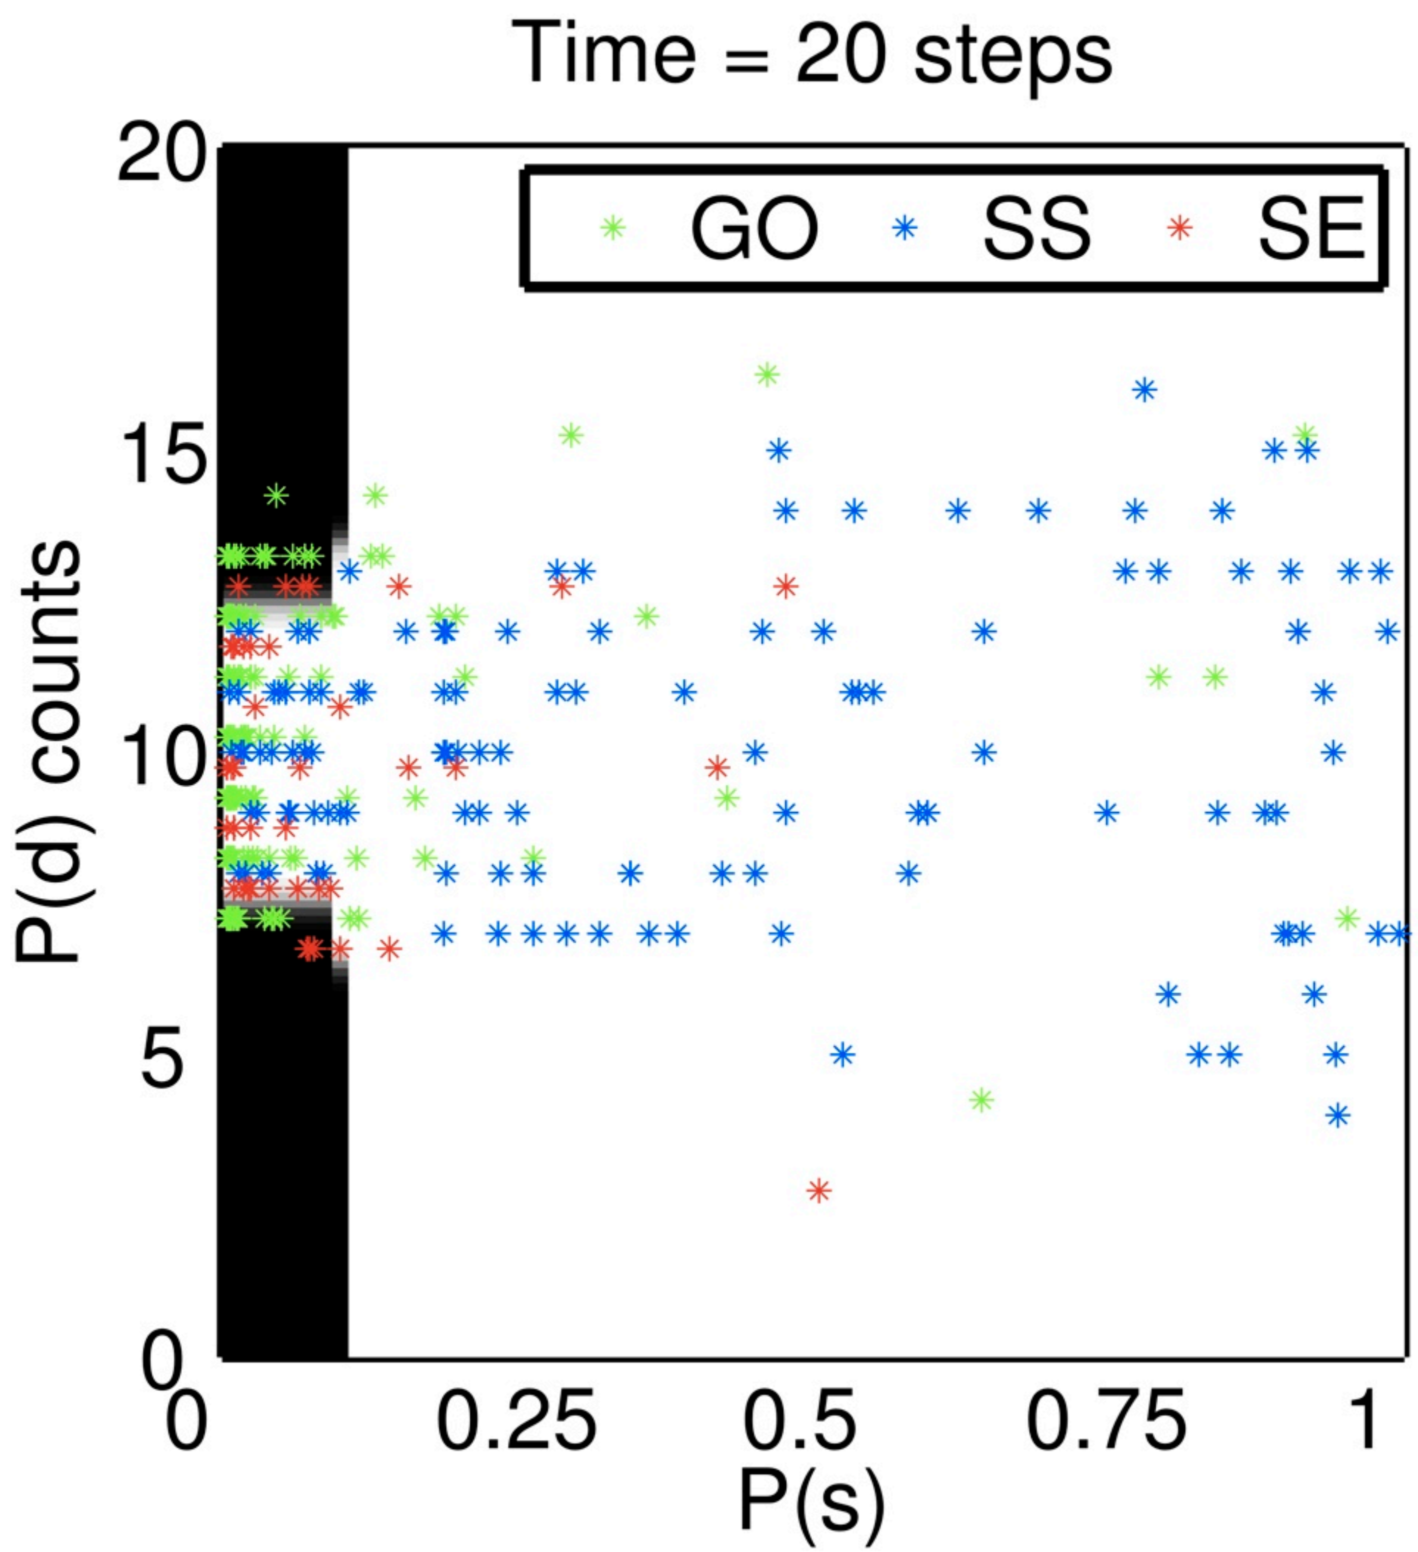

Neural Coding (FEF) of Instantaneous Action Value?

Data

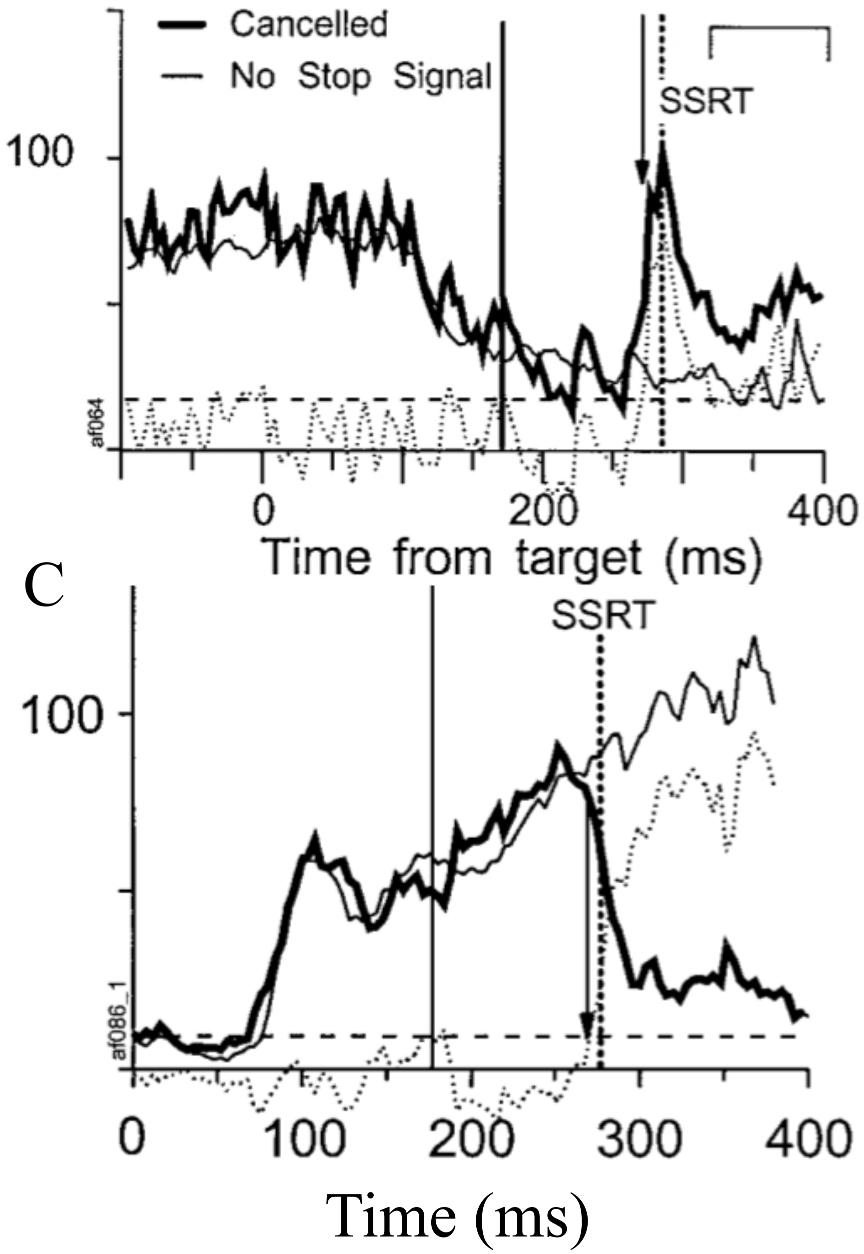

(Hanes, Patterson & Schall, 1998)

Model

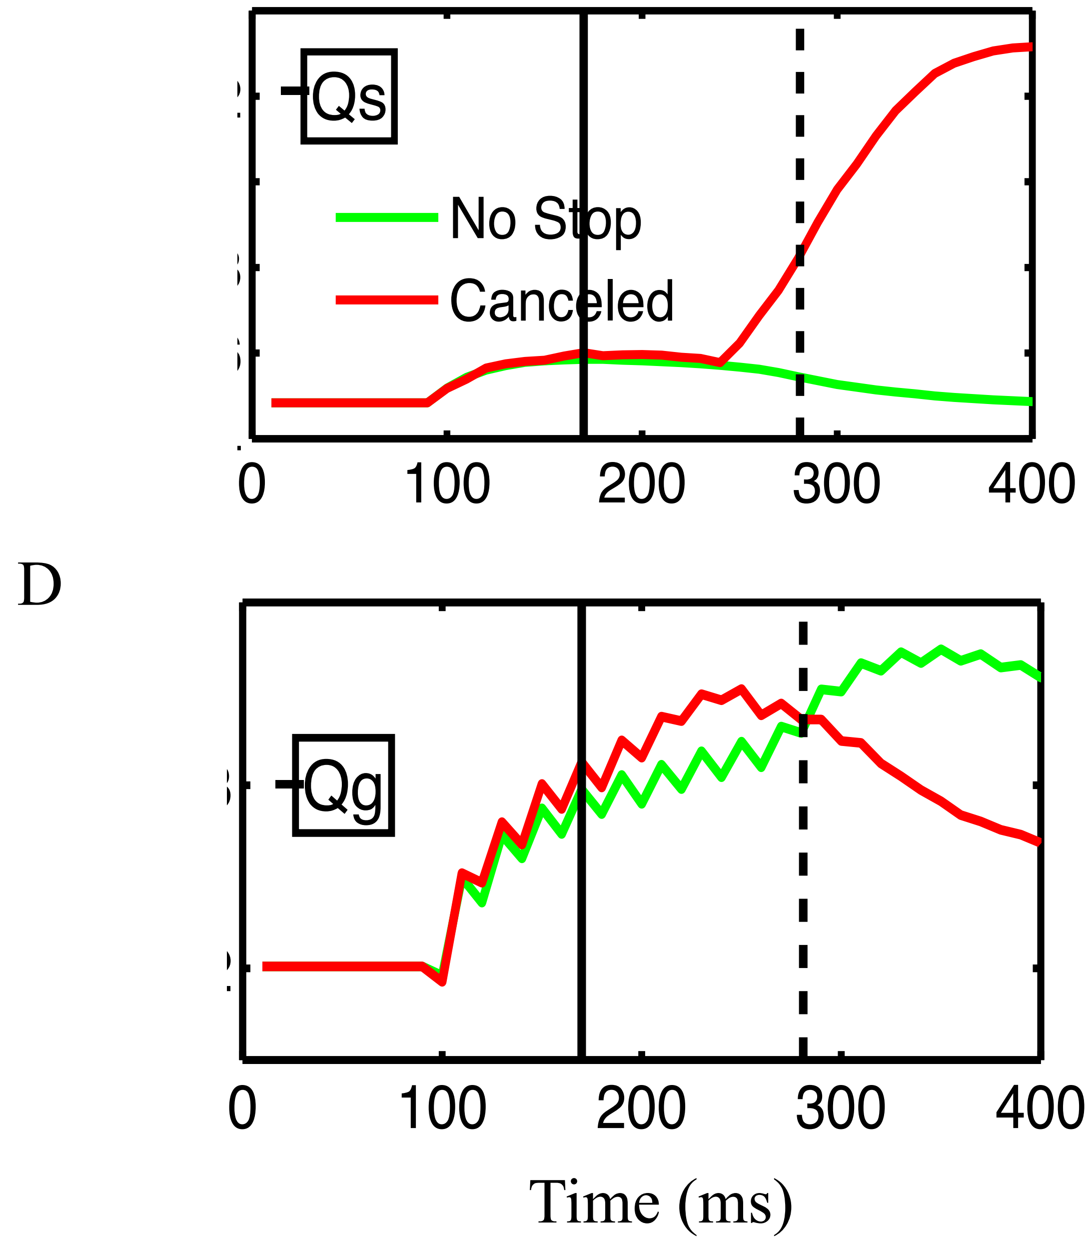

Neural Coding (FEF) of Instantaneous Action Value?

Data

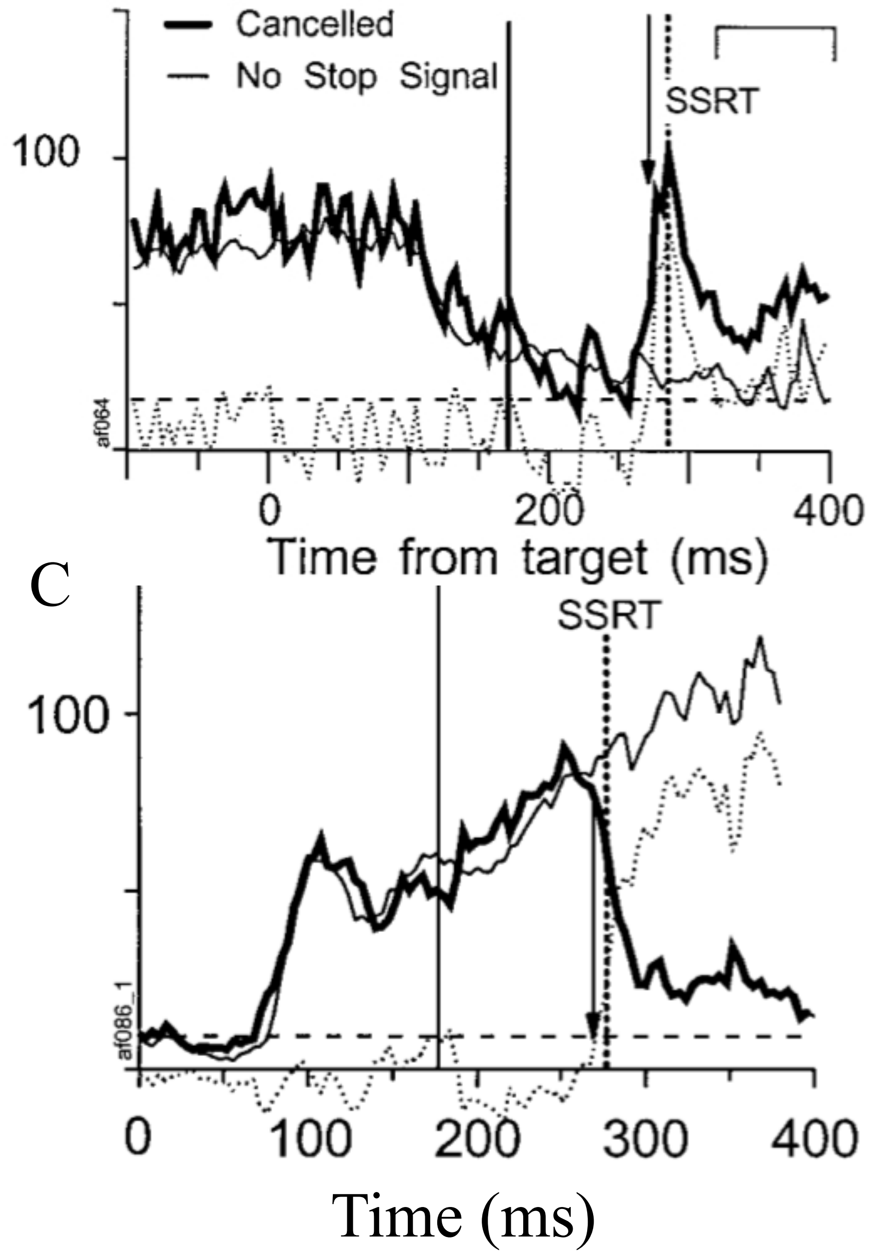

(Hanes, Patterson & Schall, 1998)

Model

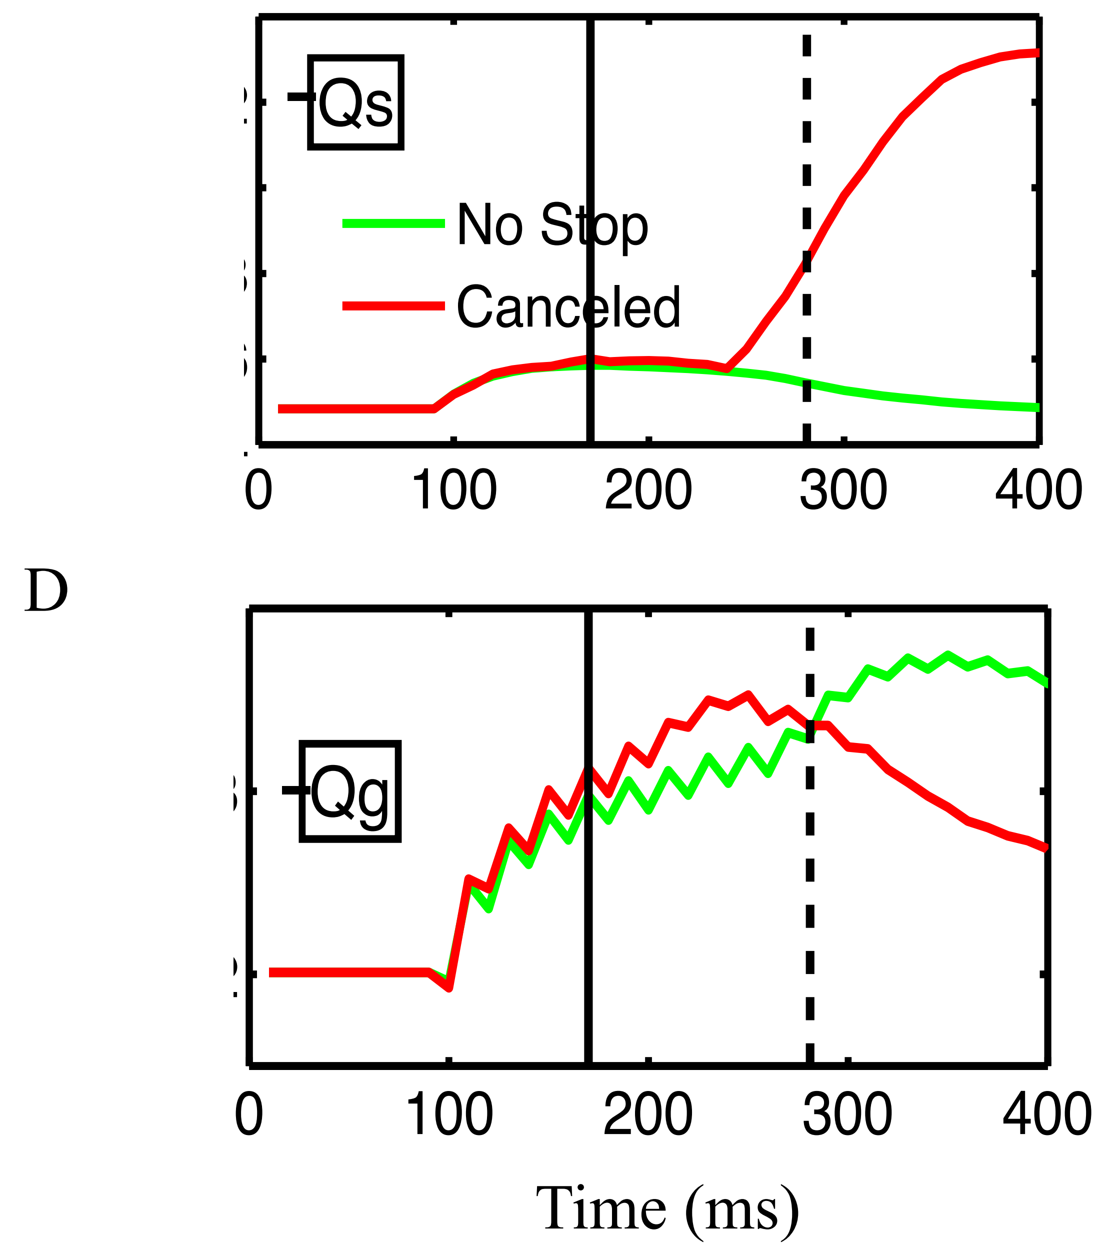

- FEF: fixation neurons (A) and movement neurons (C) diverge around SSRT between go and successfully stopped trials

Neural Coding (FEF) of Instantaneous Action Value?

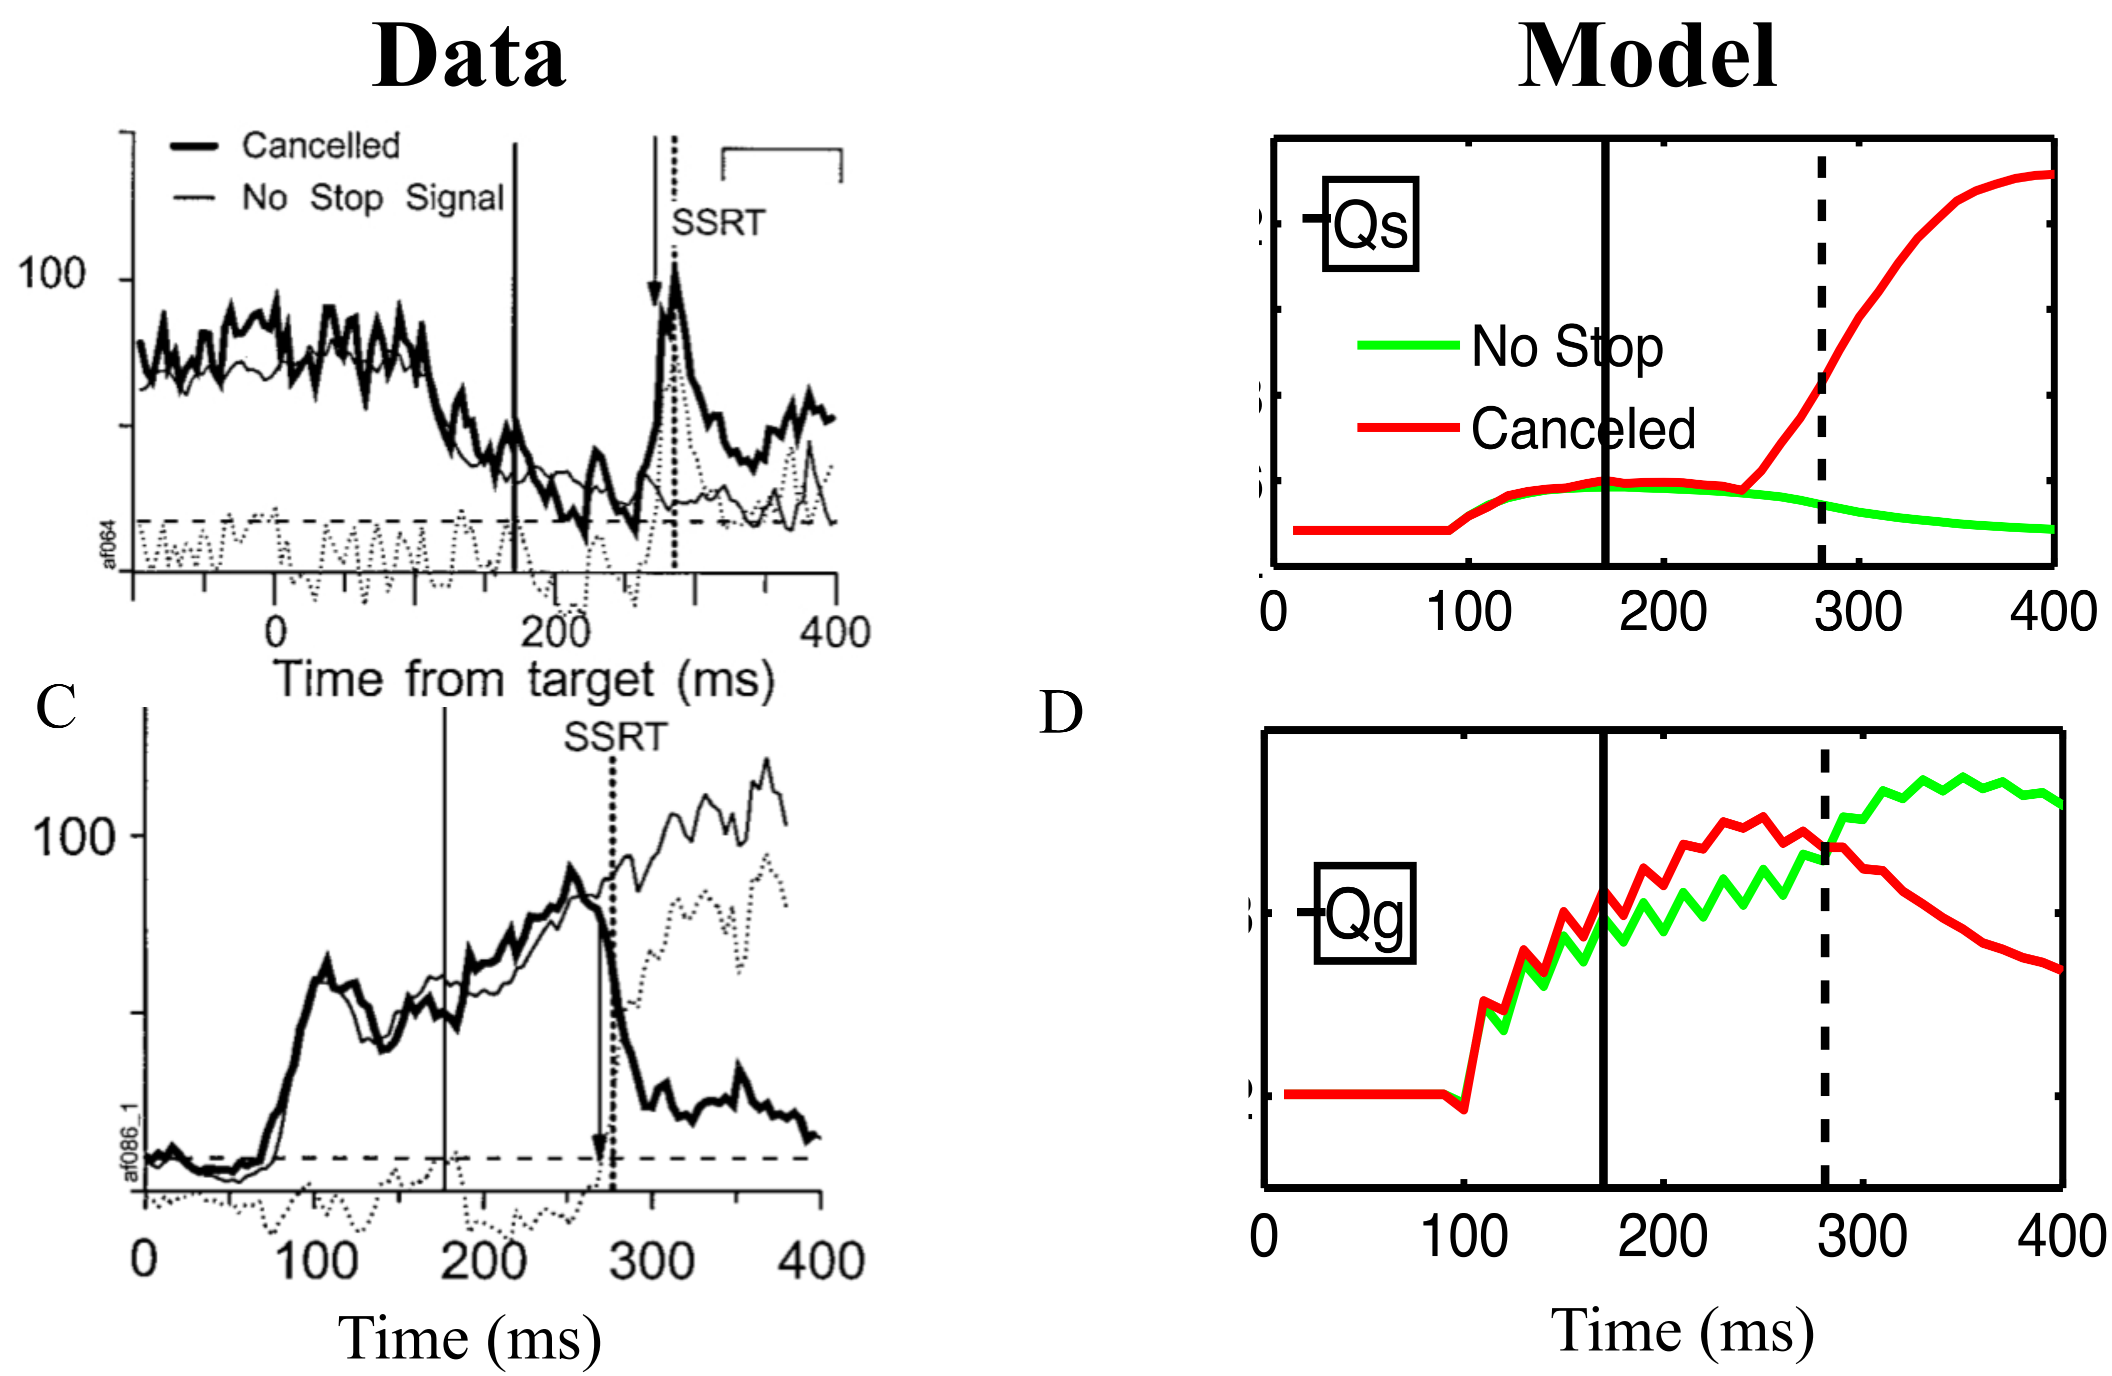

(Hanes, Patterson & Schall, 1998)

- FEF: fixation neurons (A) and movement neurons (C) diverge around SSRT between go and successfully stopped trials
- Model: trajectories of stop/go action values mimic neural activity (B,D)

The Race Model of Stopping

- A race between independent ***go*** and ***stop*** processes (Logan & Cowan, 1984)
- Winner determines trial outcome
- Stopping latency (SSRT) not directly observable
- SSRT estimated from go RT and stopping errors

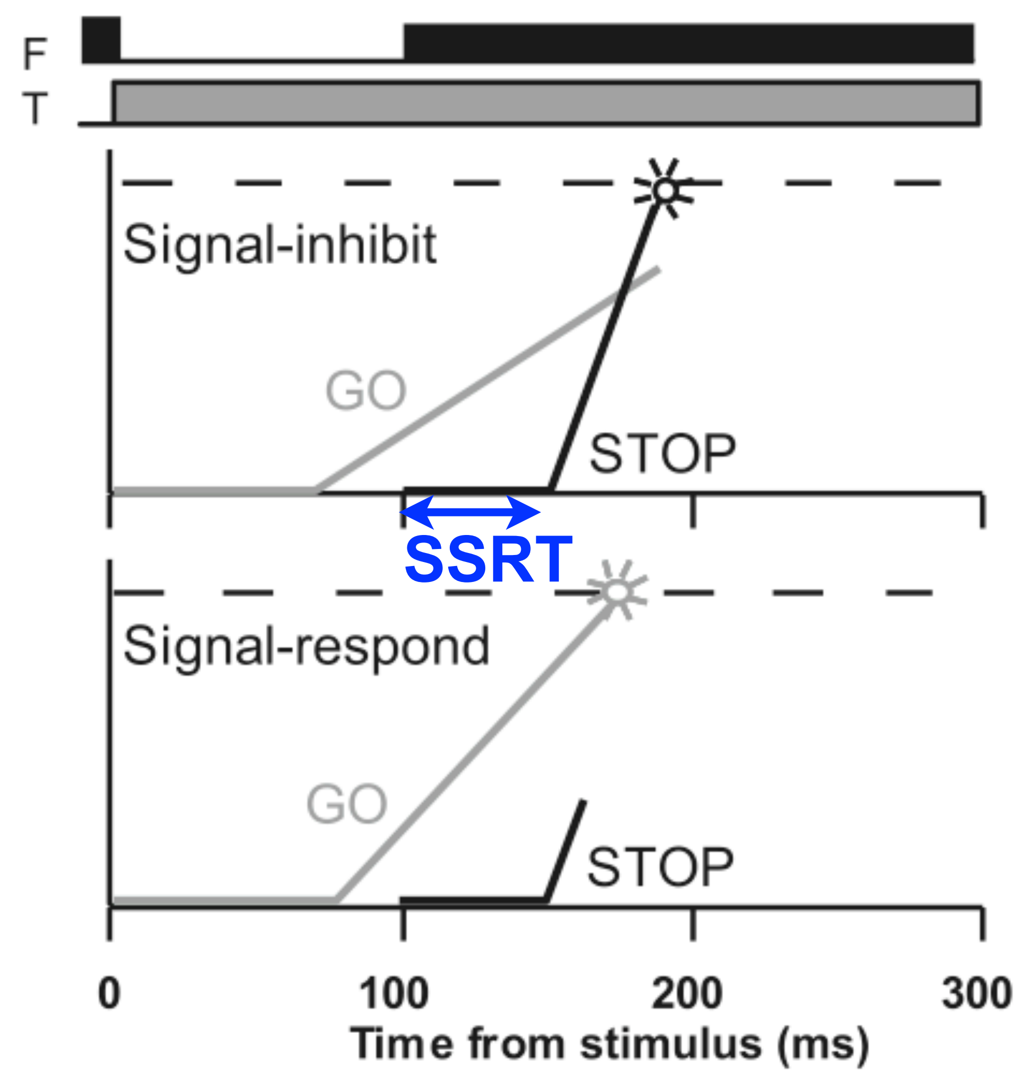

Rational Agent: Perception

Ex 1: visual illusions & ideal observer

(Adelson, 1995)

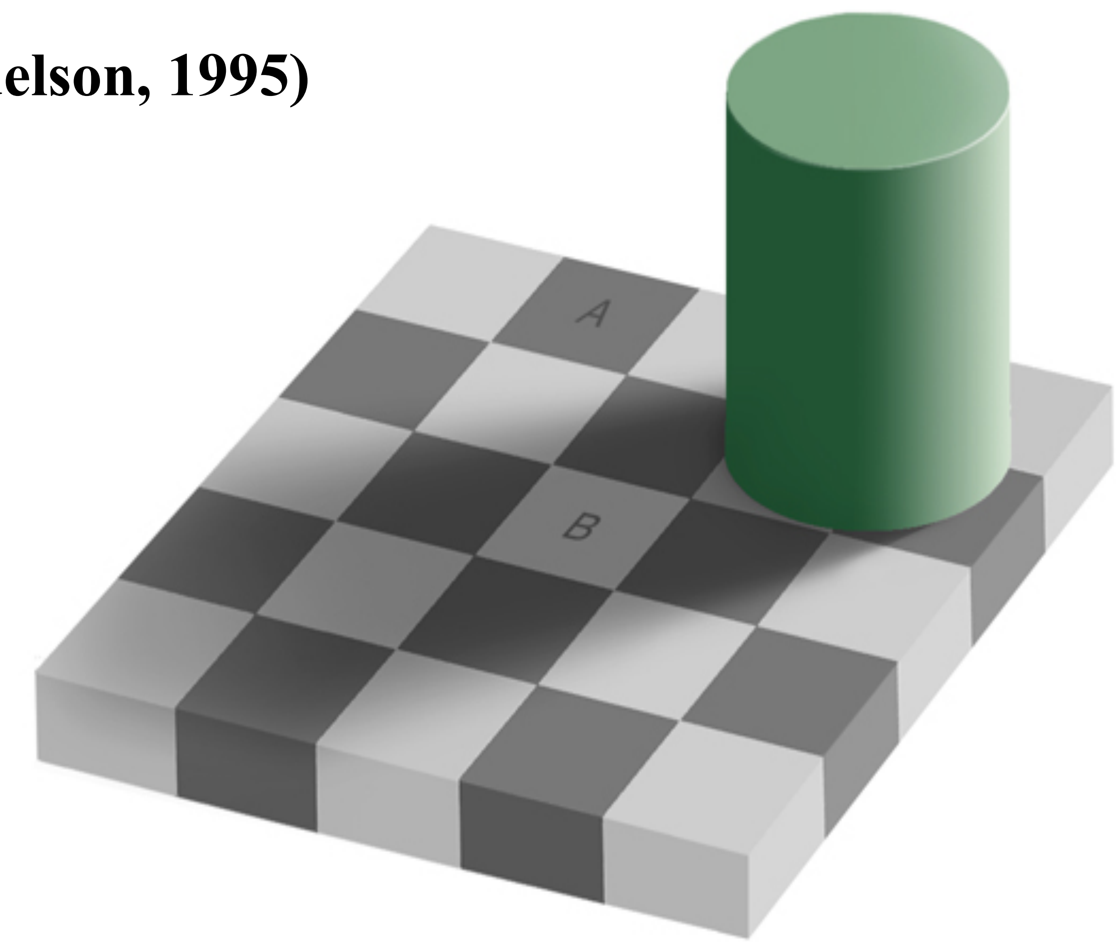

Rational Agent: Perception

Ex 1: visual illusions & ideal observer

(Adelson, 1995)

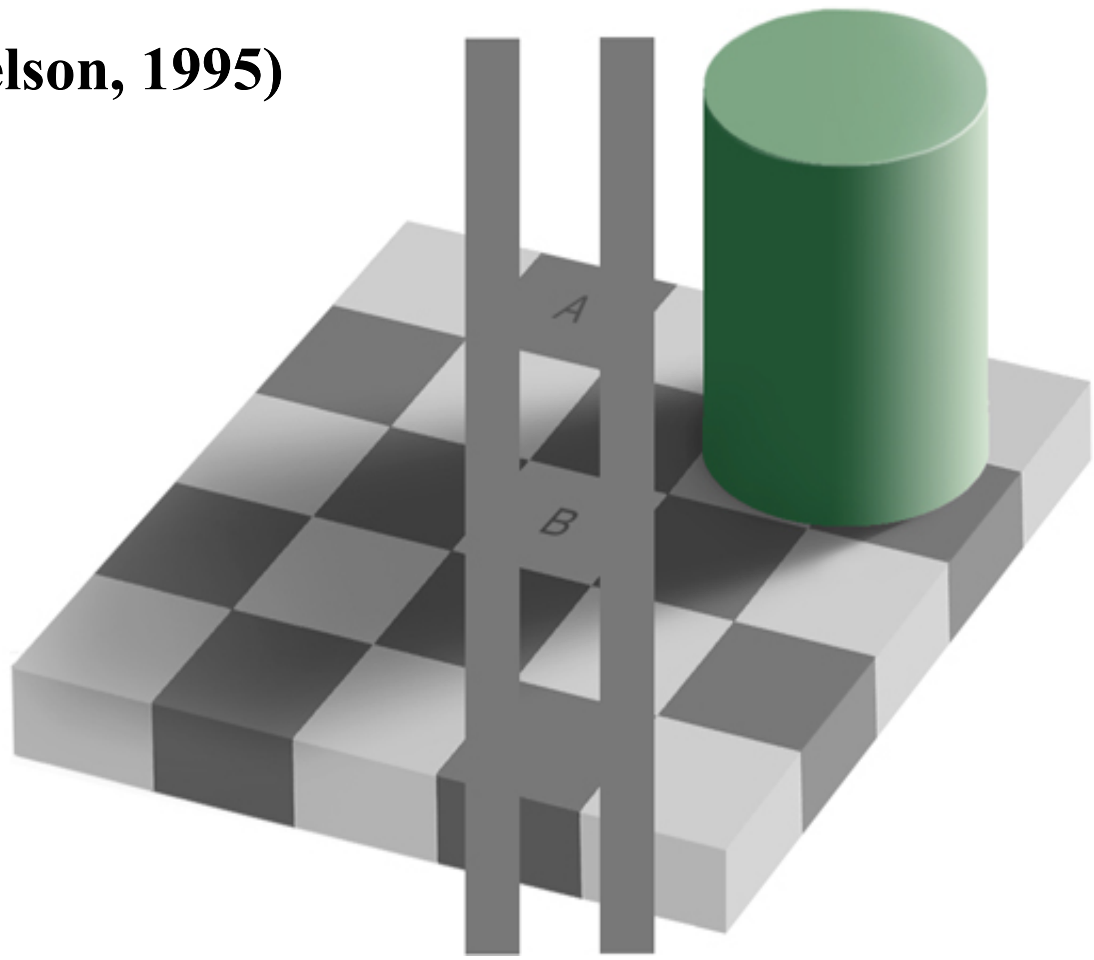

Rational Agent: Perception

Ex 1: visual illusions & ideal observer

(Adelson, 1995)

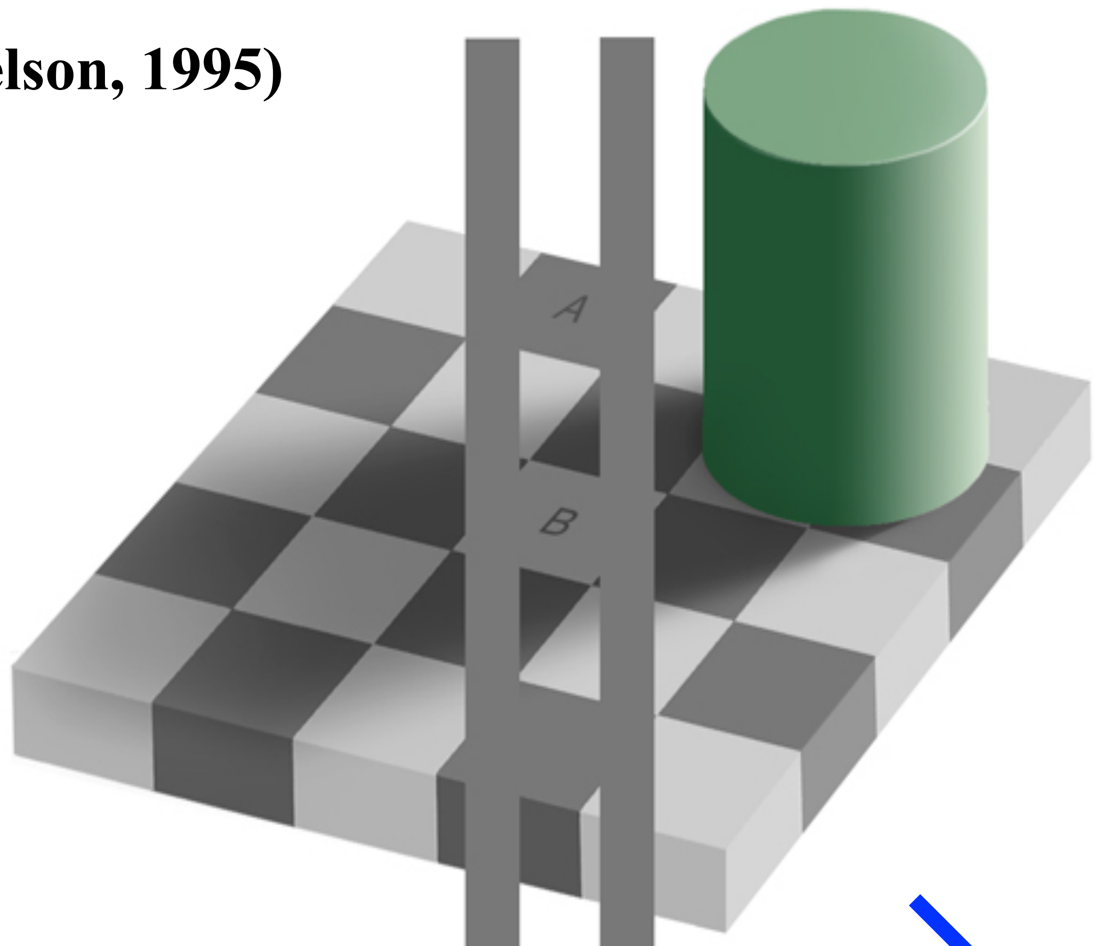

depth
lighting
shadow
spatial regularity
....

sensory input

visual percept

Race Model -- Diffusion Model

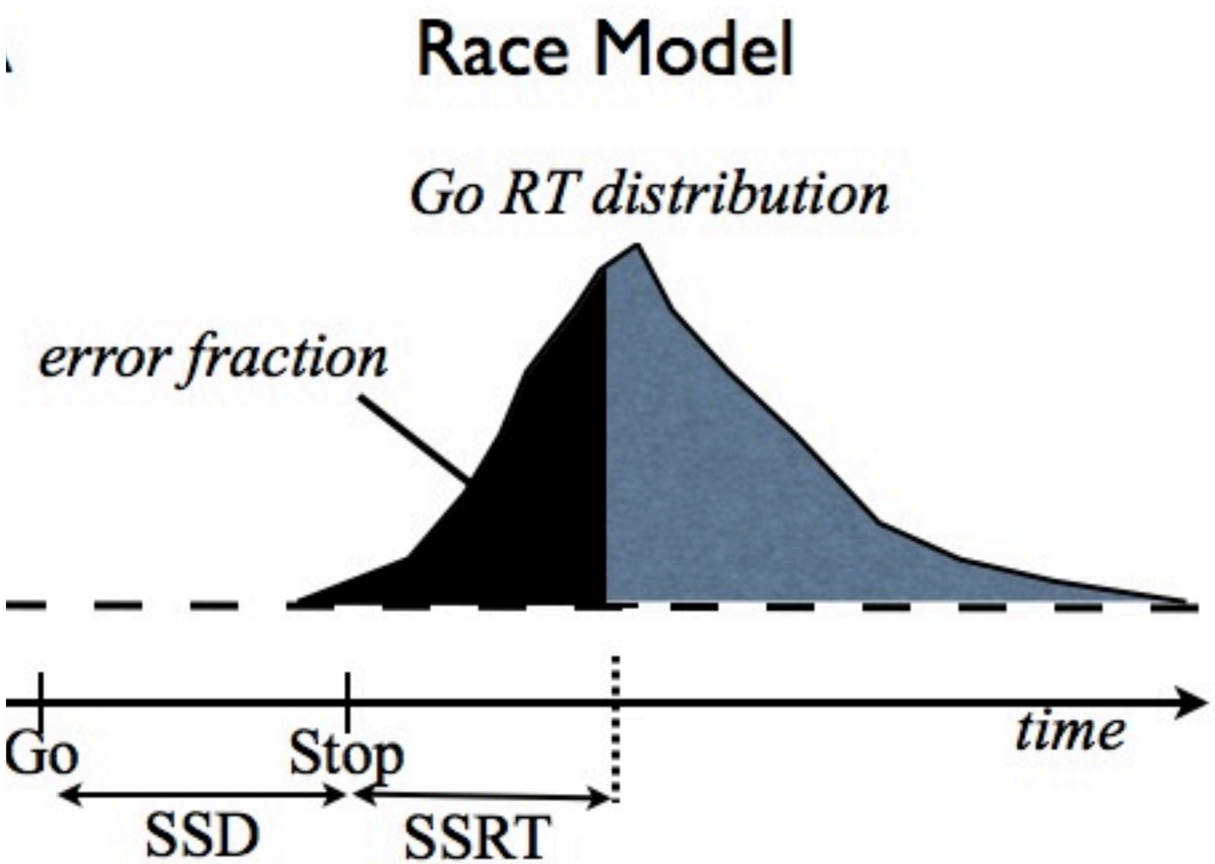

Race Model -- Diffusion Model

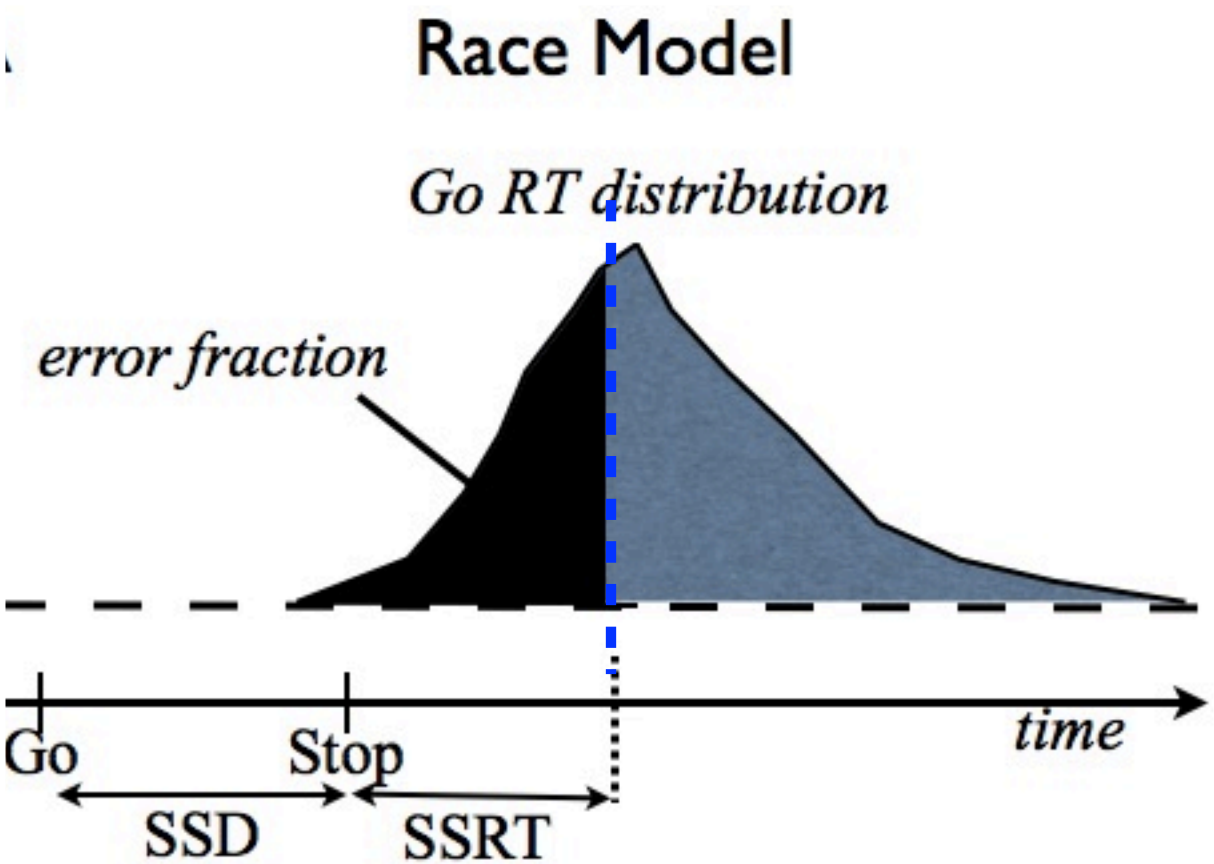

Race Model -- Diffusion Model

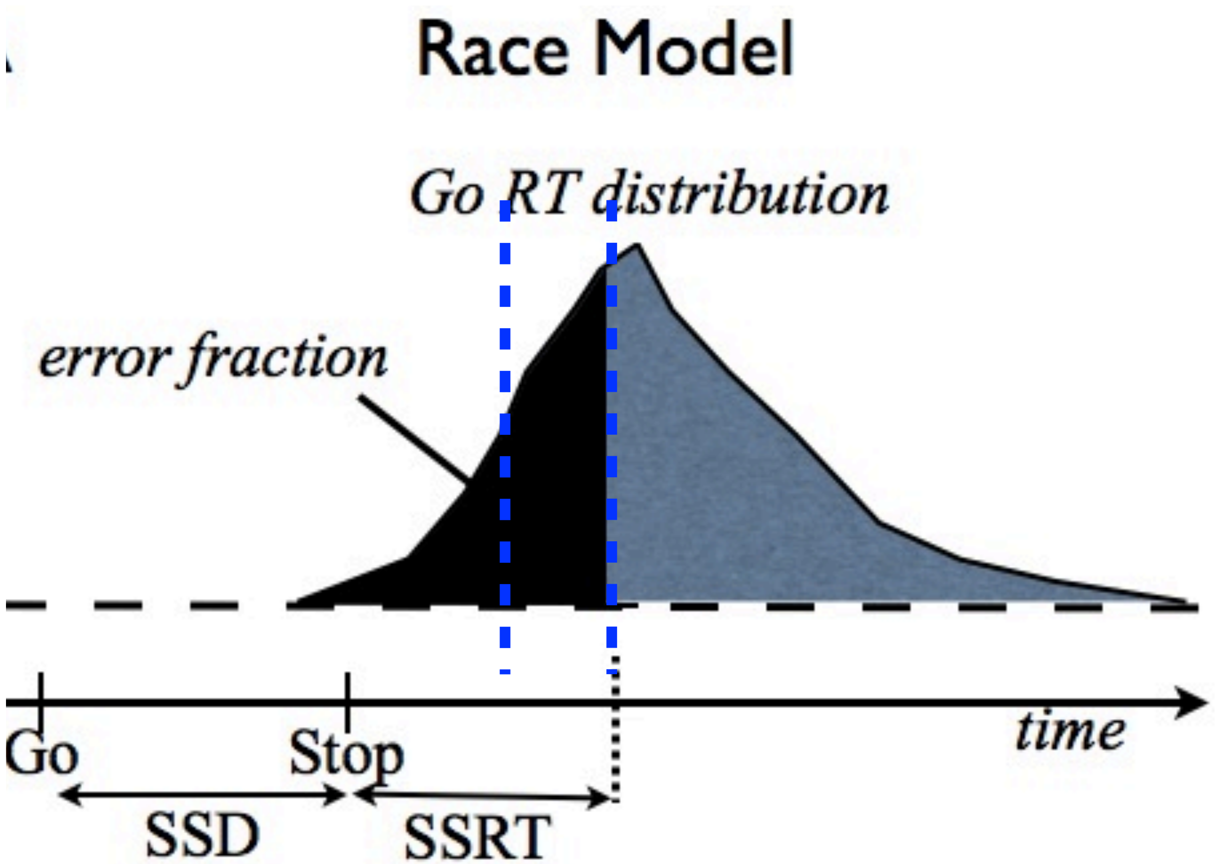

Race Model -- Diffusion Model

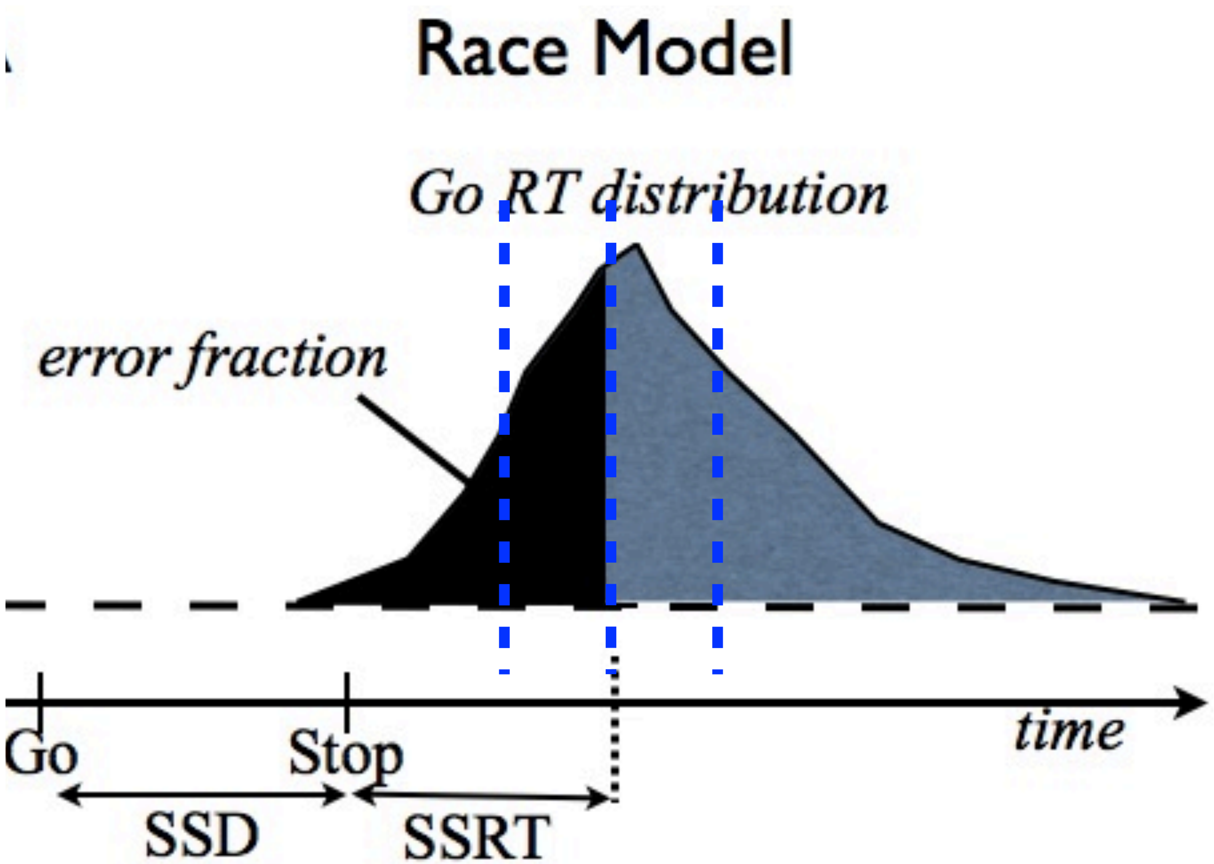

Race Model -- Diffusion Model

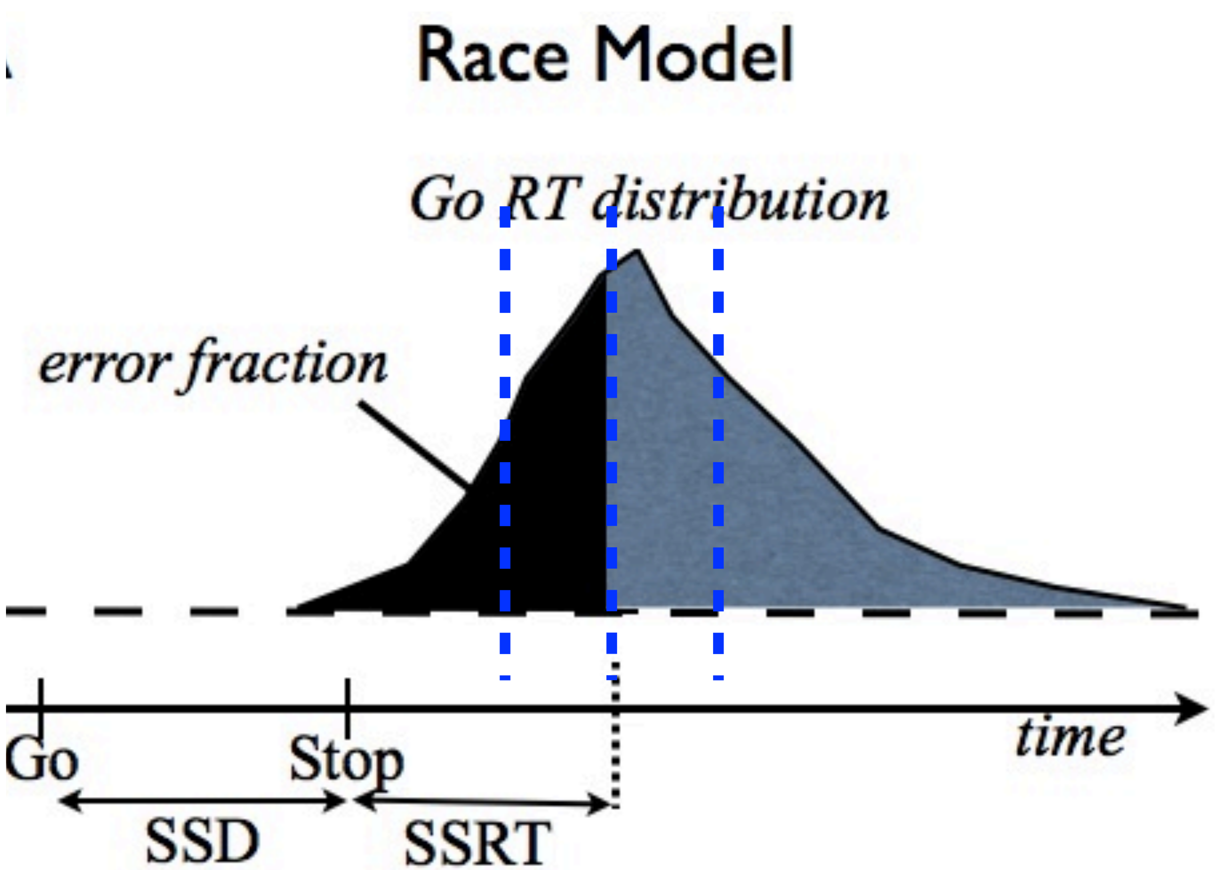

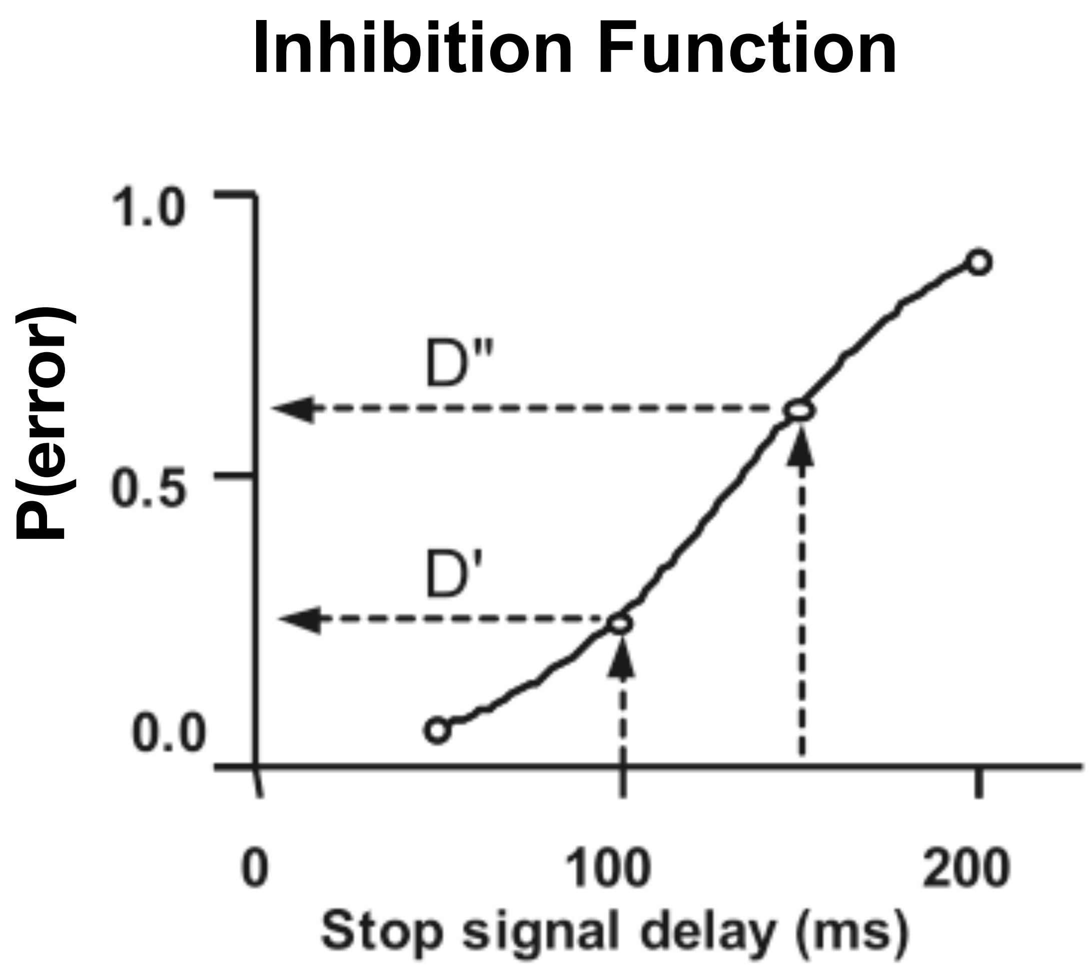

Sensory Processing = Bayesian Inference

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto \left(p_z^{t-1} + (1 - p_z^{t-1}) h(t)\right) f_z(y^t)$$

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1-p_z^{t-1})h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1 - p_z^{t-1}) h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

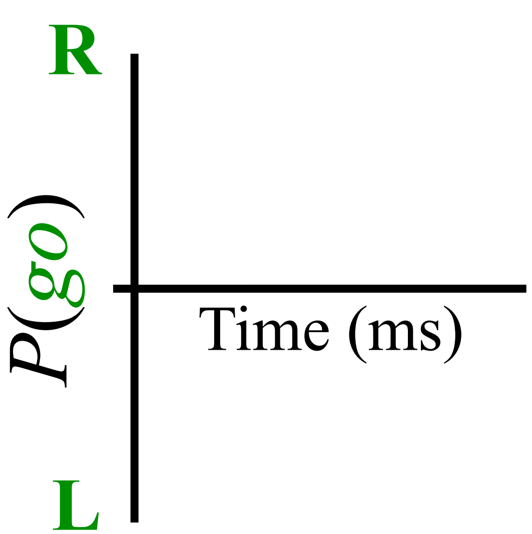

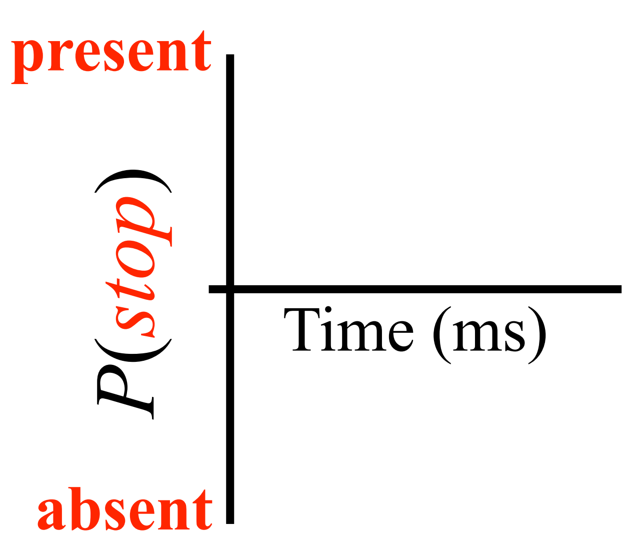

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1-p_z^{t-1})h(t))f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

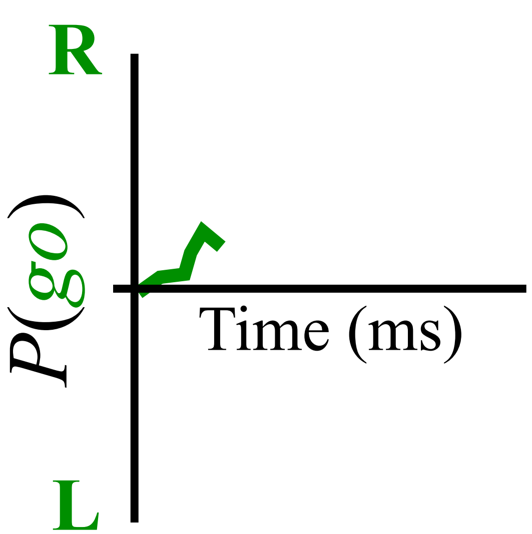

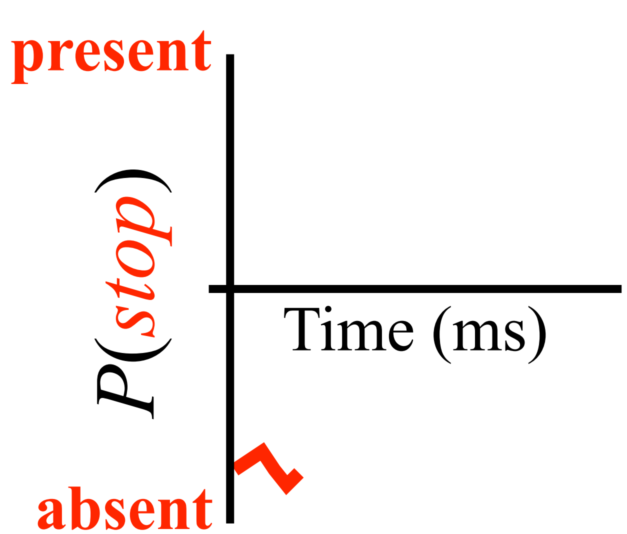

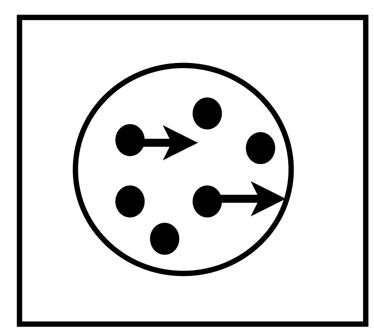

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1-p_z^{t-1})h(t))\, f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

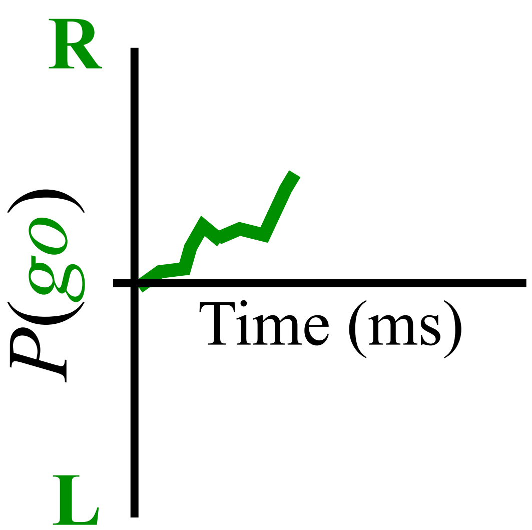

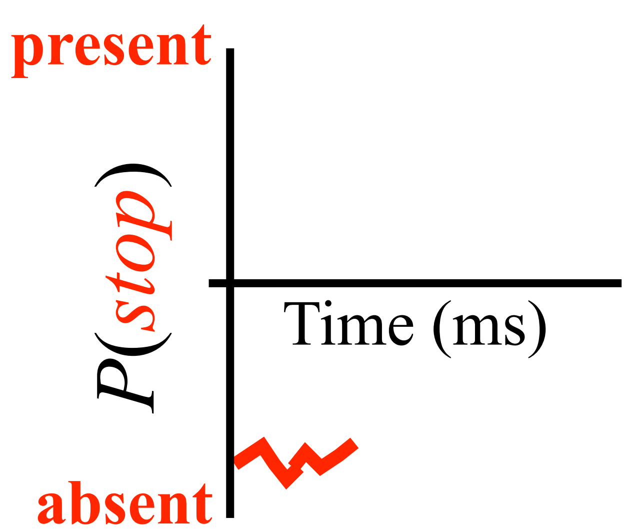

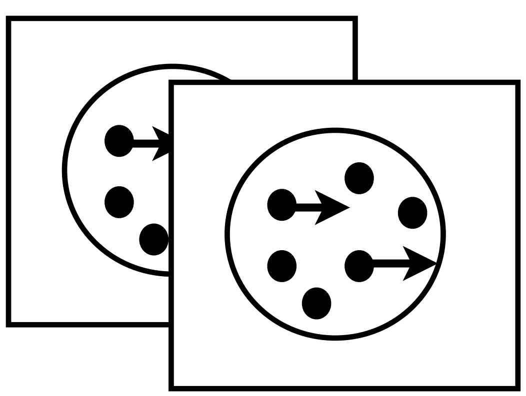

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1 - p_z^{t-1}) h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

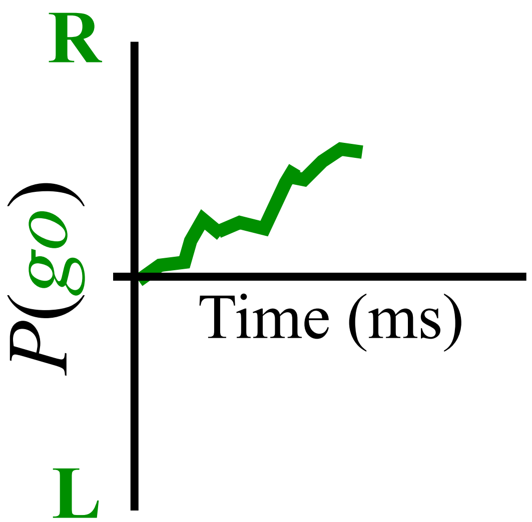

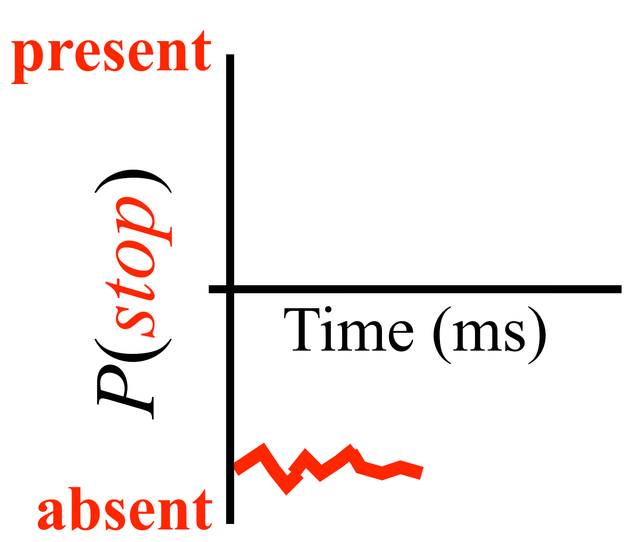

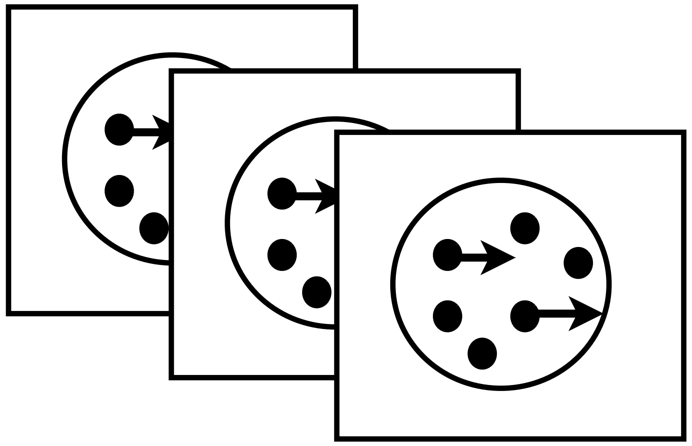

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1-p_z^{t-1})h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

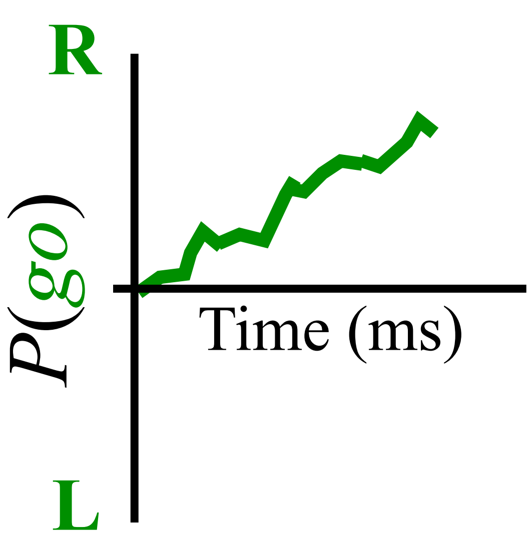

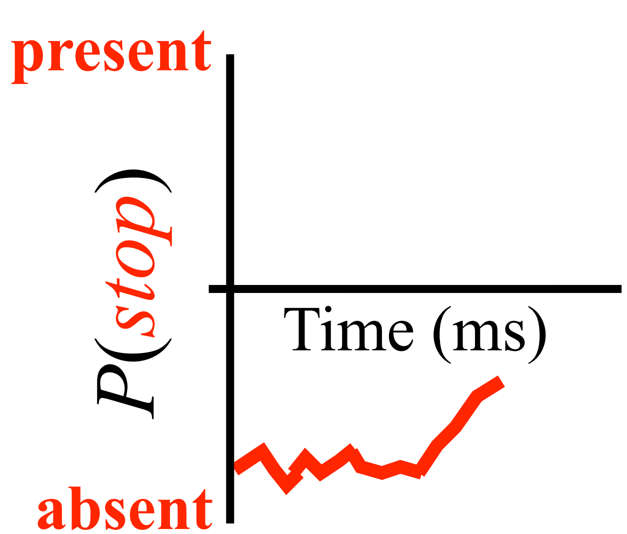

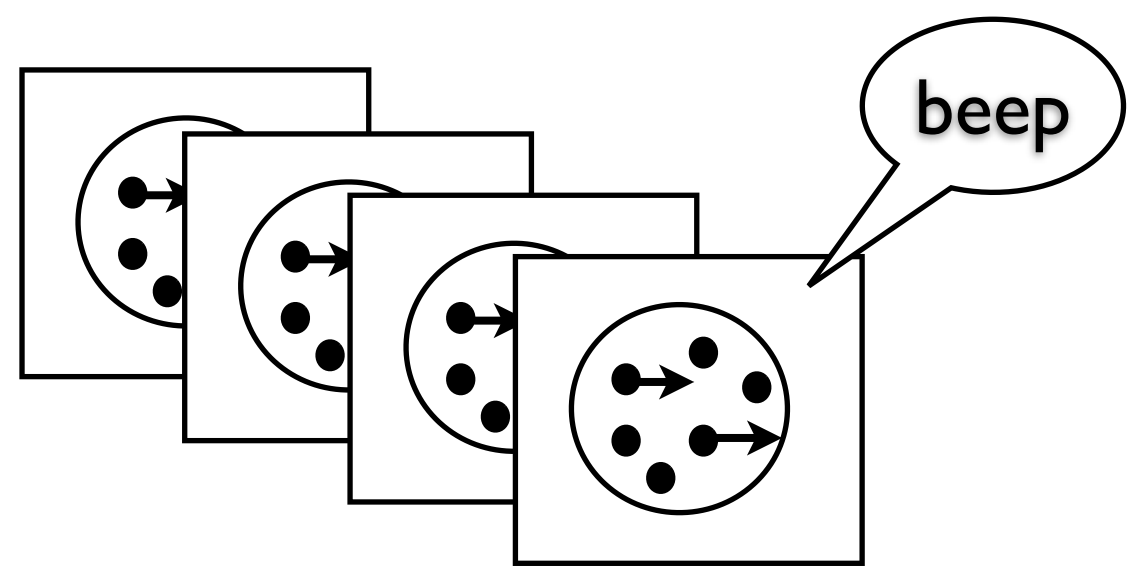

Sensory Processing = Bayesian Inference

Target = L/R?
(Bayes' Rule)

$$p_d^t \propto p_d^{t-1} f_d(x^t)$$

Stop signal present?
(Bayes' Rule)

$$p_z^t \propto (p_z^{t-1} + (1 - p_z^{t-1}) h(t)) f_z(y^t)$$

Stop trial?

$$p_s^t = p_z^t + P\{\text{stop signal in future}\}$$

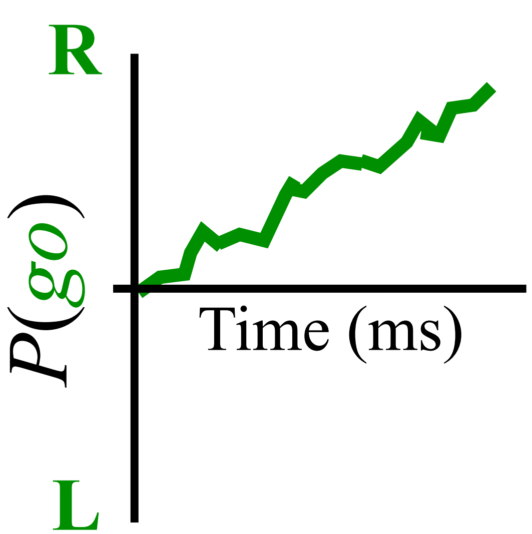

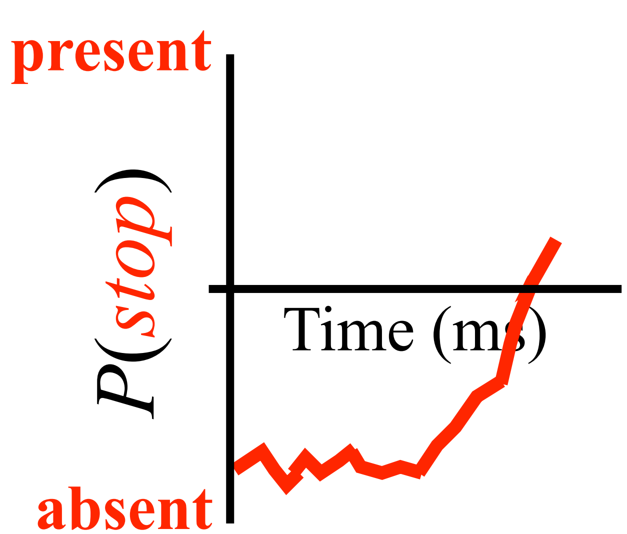

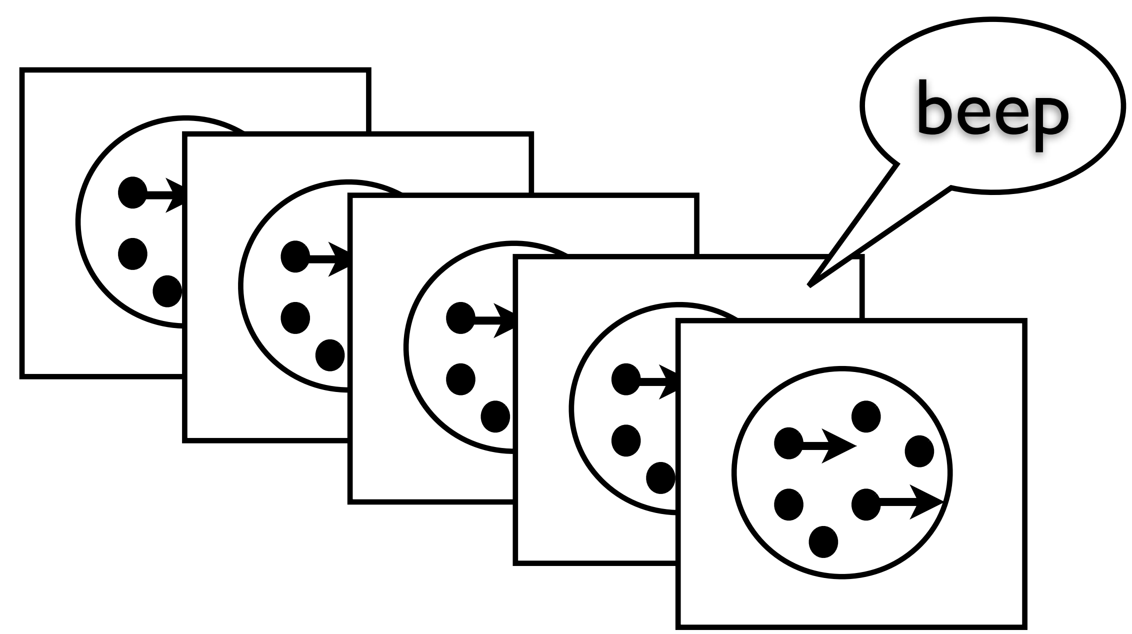

Rational Agent: Perceptual Discrimination

Ex: 2AFC motion discrimination

(from Roitman & Shadlen, 2002)

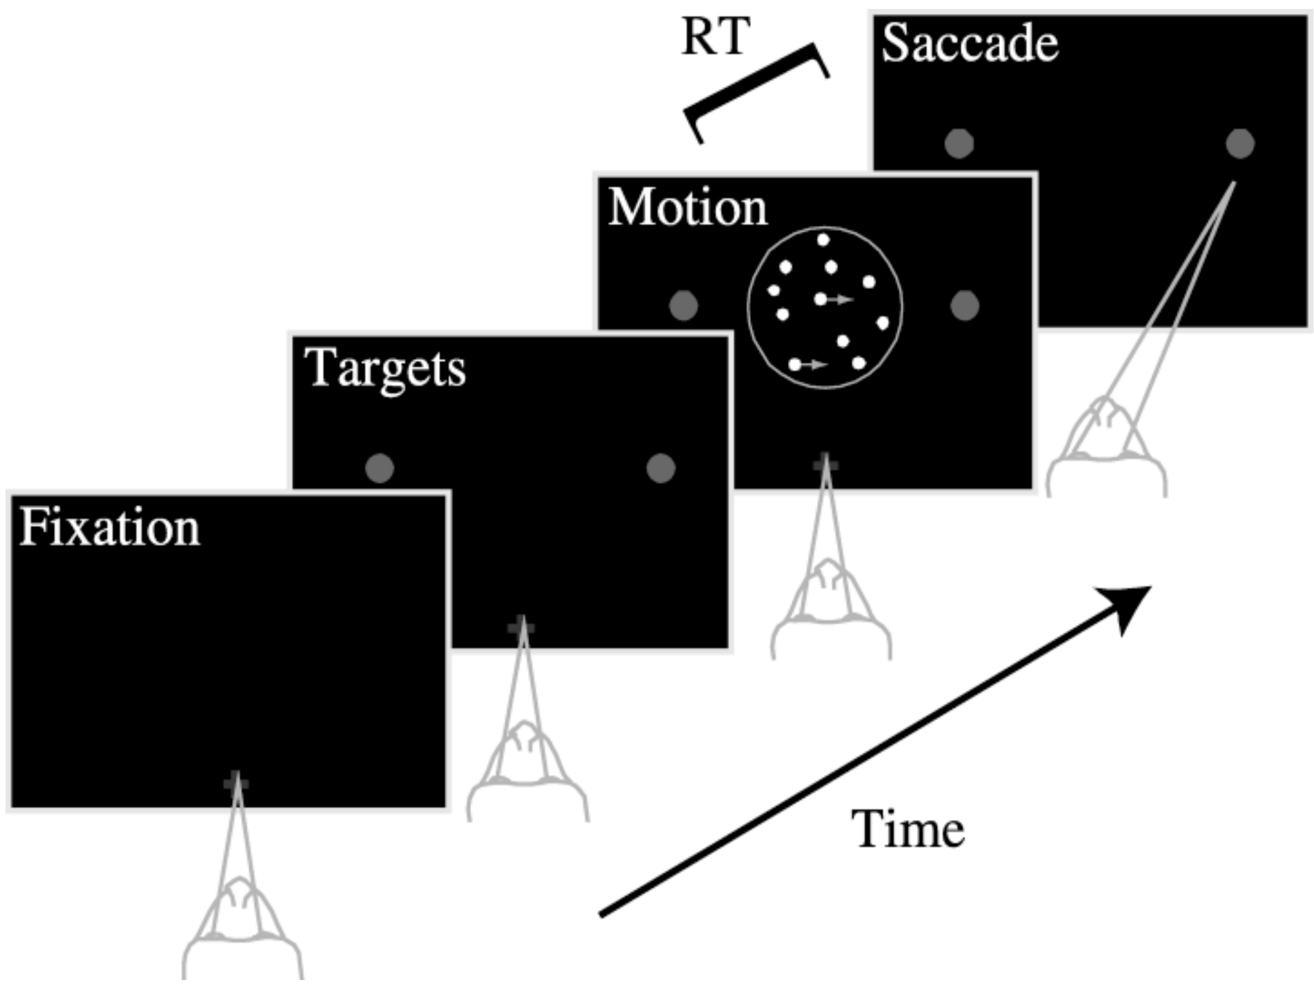

Rational Agent: Perceptual Discrimination

Ex: 2AFC motion discrimination

(from Roitman & Shadlen, 2002)

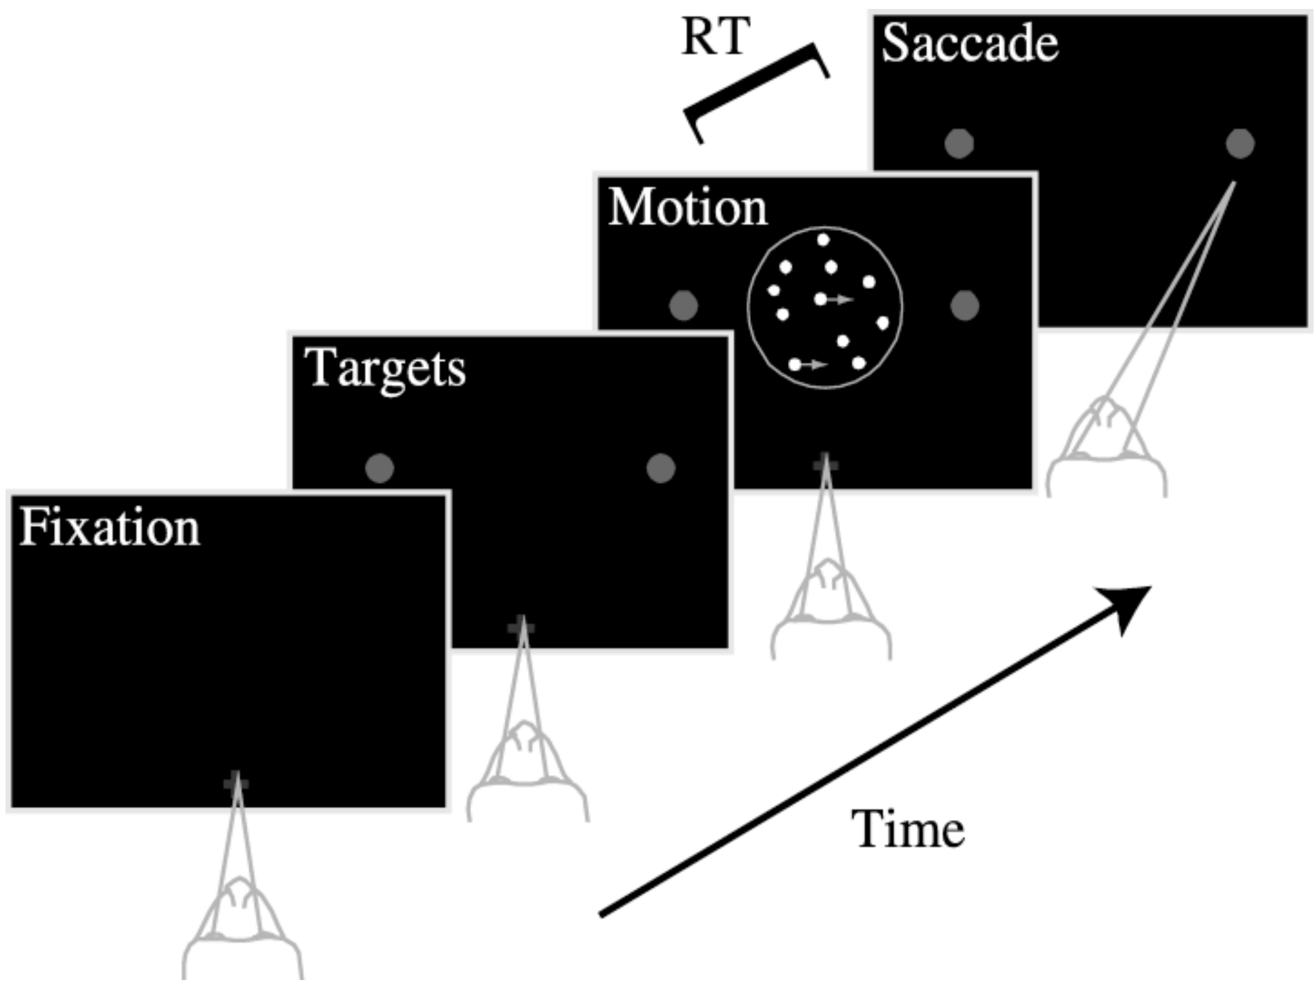

30% coh

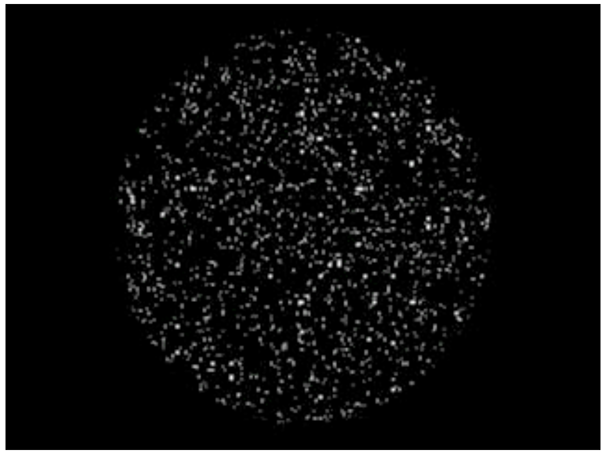

Rational Agent: Perceptual Discrimination

Ex: 2AFC motion discrimination

(from Roitman & Shadlen, 2002)

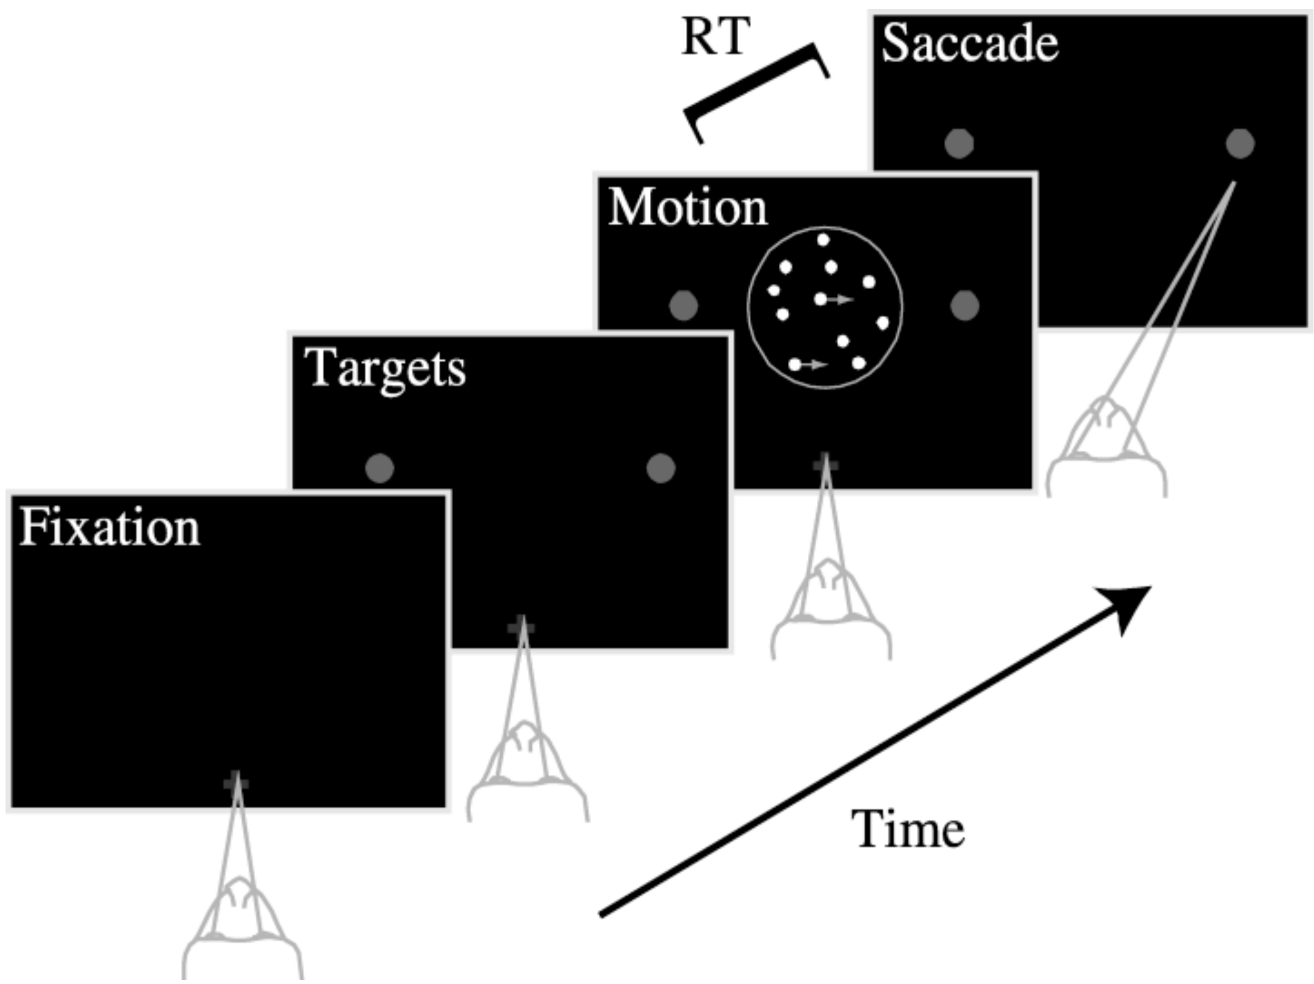

30% coh

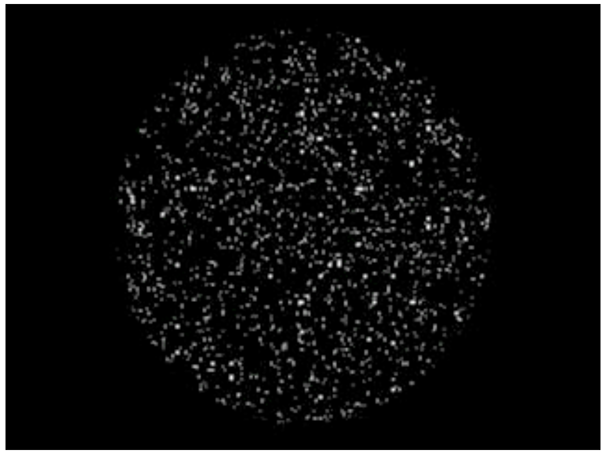

5% coh

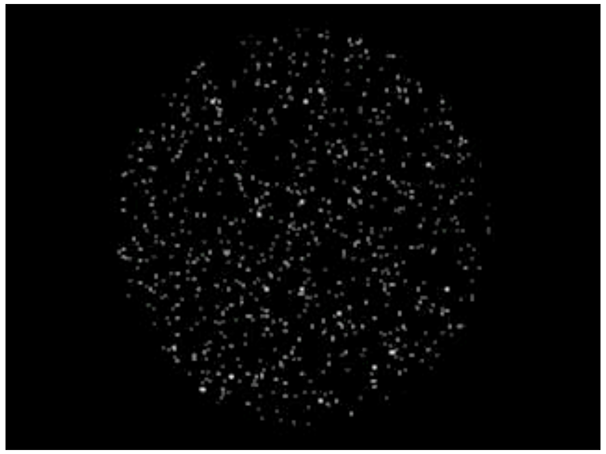

Rational Agent: Perceptual Discrimination

Ex: 2AFC motion discrimination

(from Roitman & Shadlen, 2002)

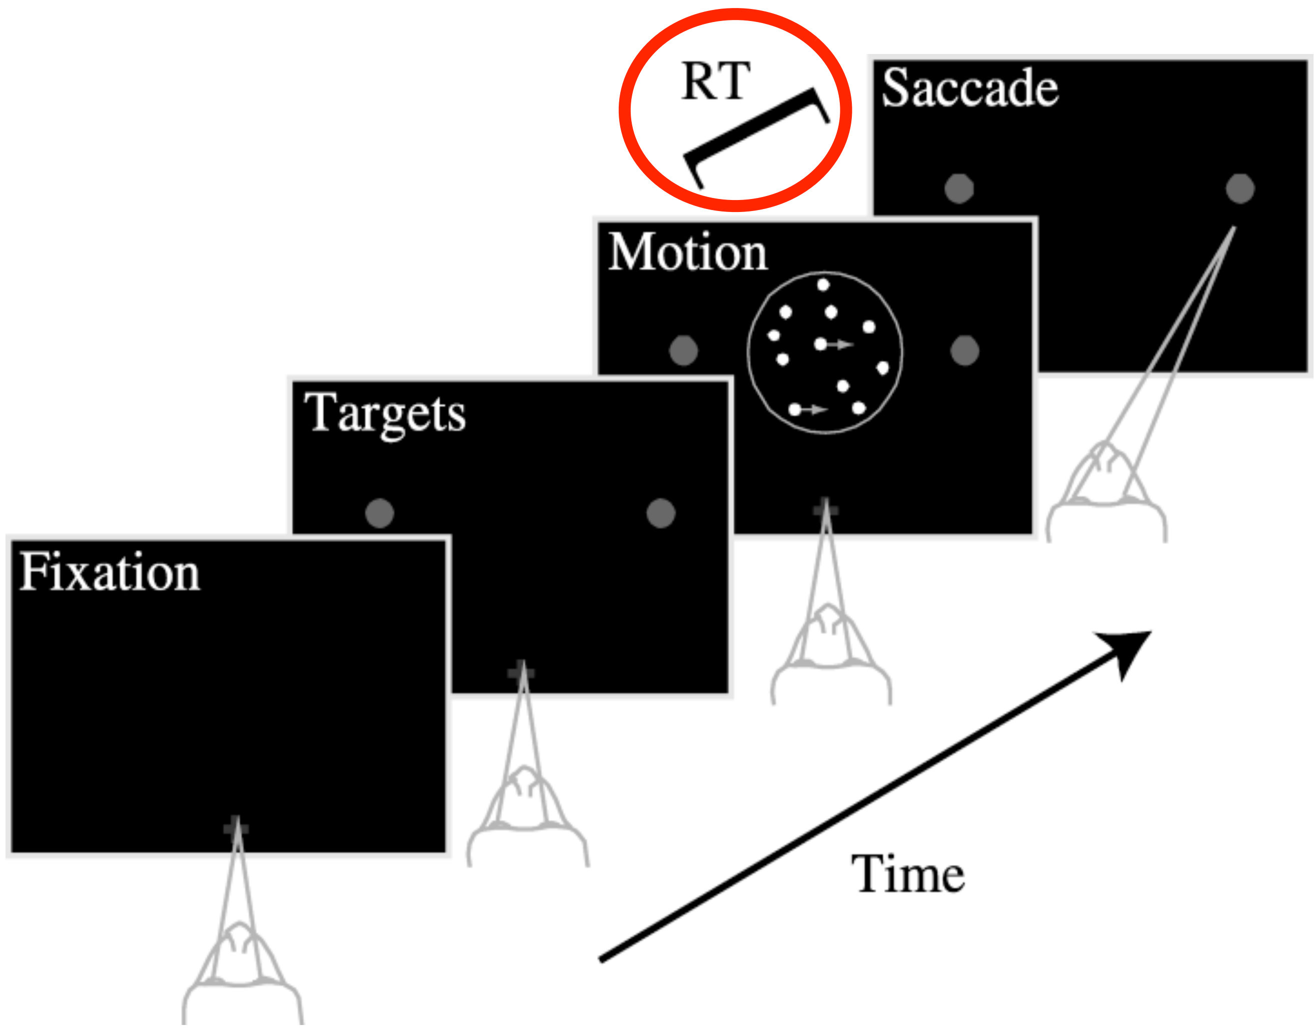

30% coh

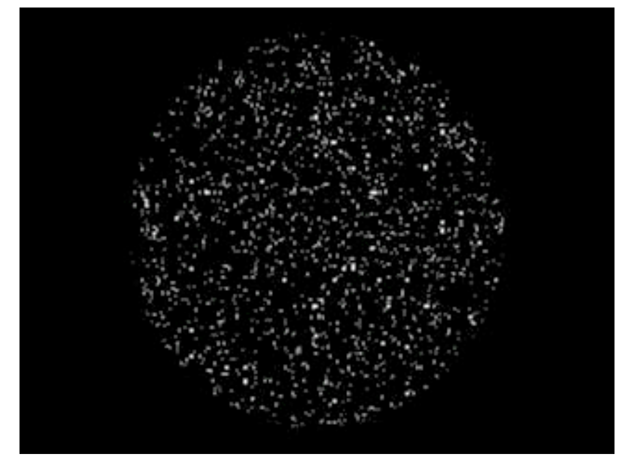

5% coh

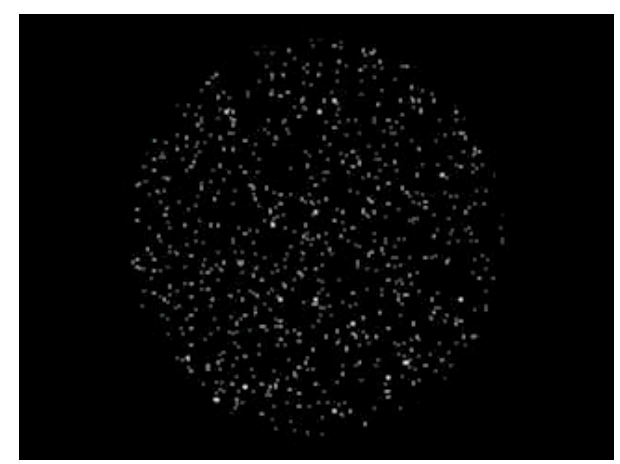

Rational Agent: Perceptual Discrimination

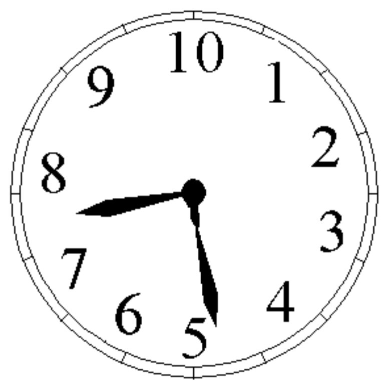

Speed **vs.** **Accuracy**

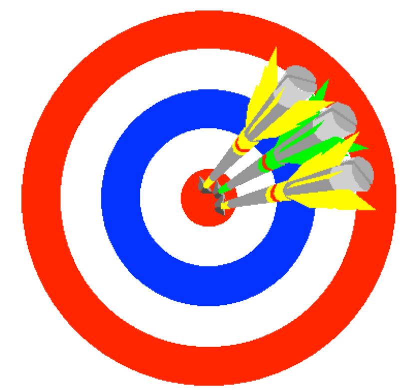

Rational Agent: Perceptual Discrimination

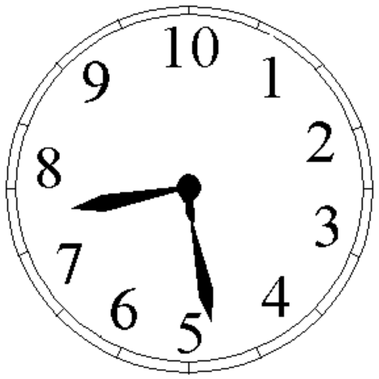

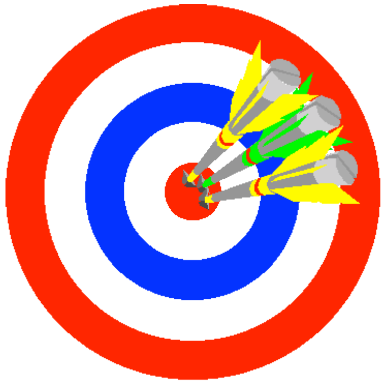

Speed vs. **Accuracy**

- Slow response ⇒ fewer errors, higher opportunity cost

Rational Agent: Perceptual Discrimination

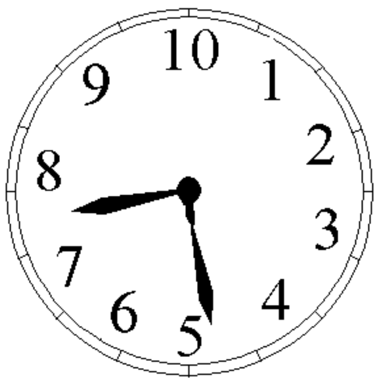

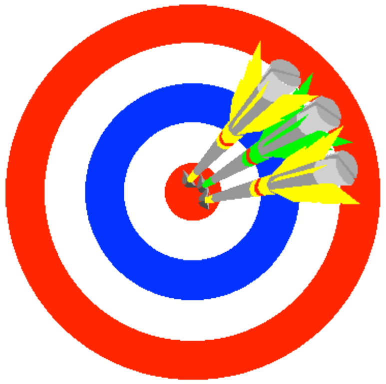

Speed **vs.** **Accuracy**

- Slow response $\Rightarrow$ fewer errors, higher opportunity cost
- What is optimal tradeoff? What computations involved?

Rational Agent: Perceptual Discrimination

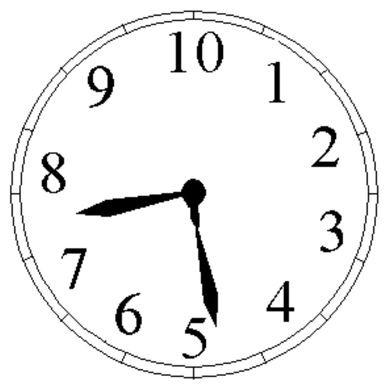

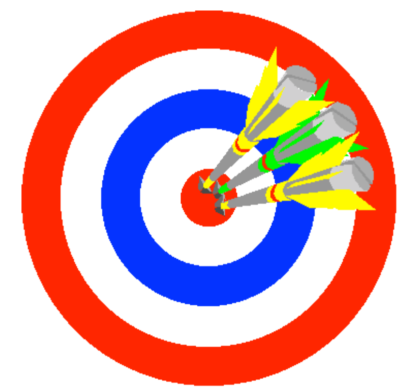

Speed vs. **Accuracy**

- Slow response $\Rightarrow$ fewer errors, higher opportunity cost
- What is optimal tradeoff? What computations involved?
- Are humans/animals optimal?

Rational Agent: Perceptual Discrimination

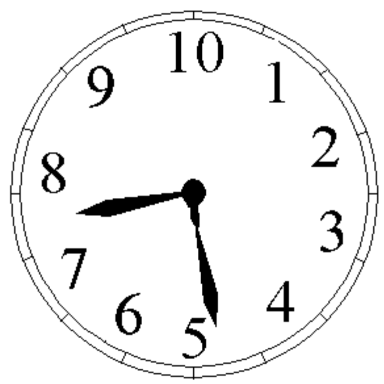

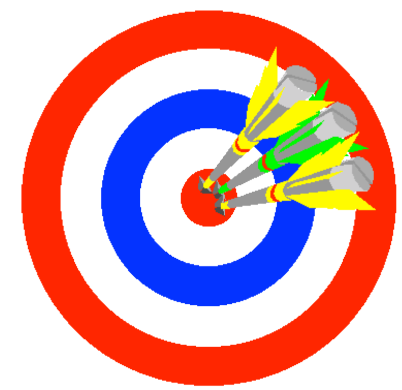

Speed **vs.** **Accuracy**

- Slow response $\Rightarrow$ fewer errors, higher opportunity cost
- What is optimal tradeoff? What computations involved?
- Are humans/animals optimal?
- Neural implementation?

Reward/Motivation ⇒ Stopping Behavior

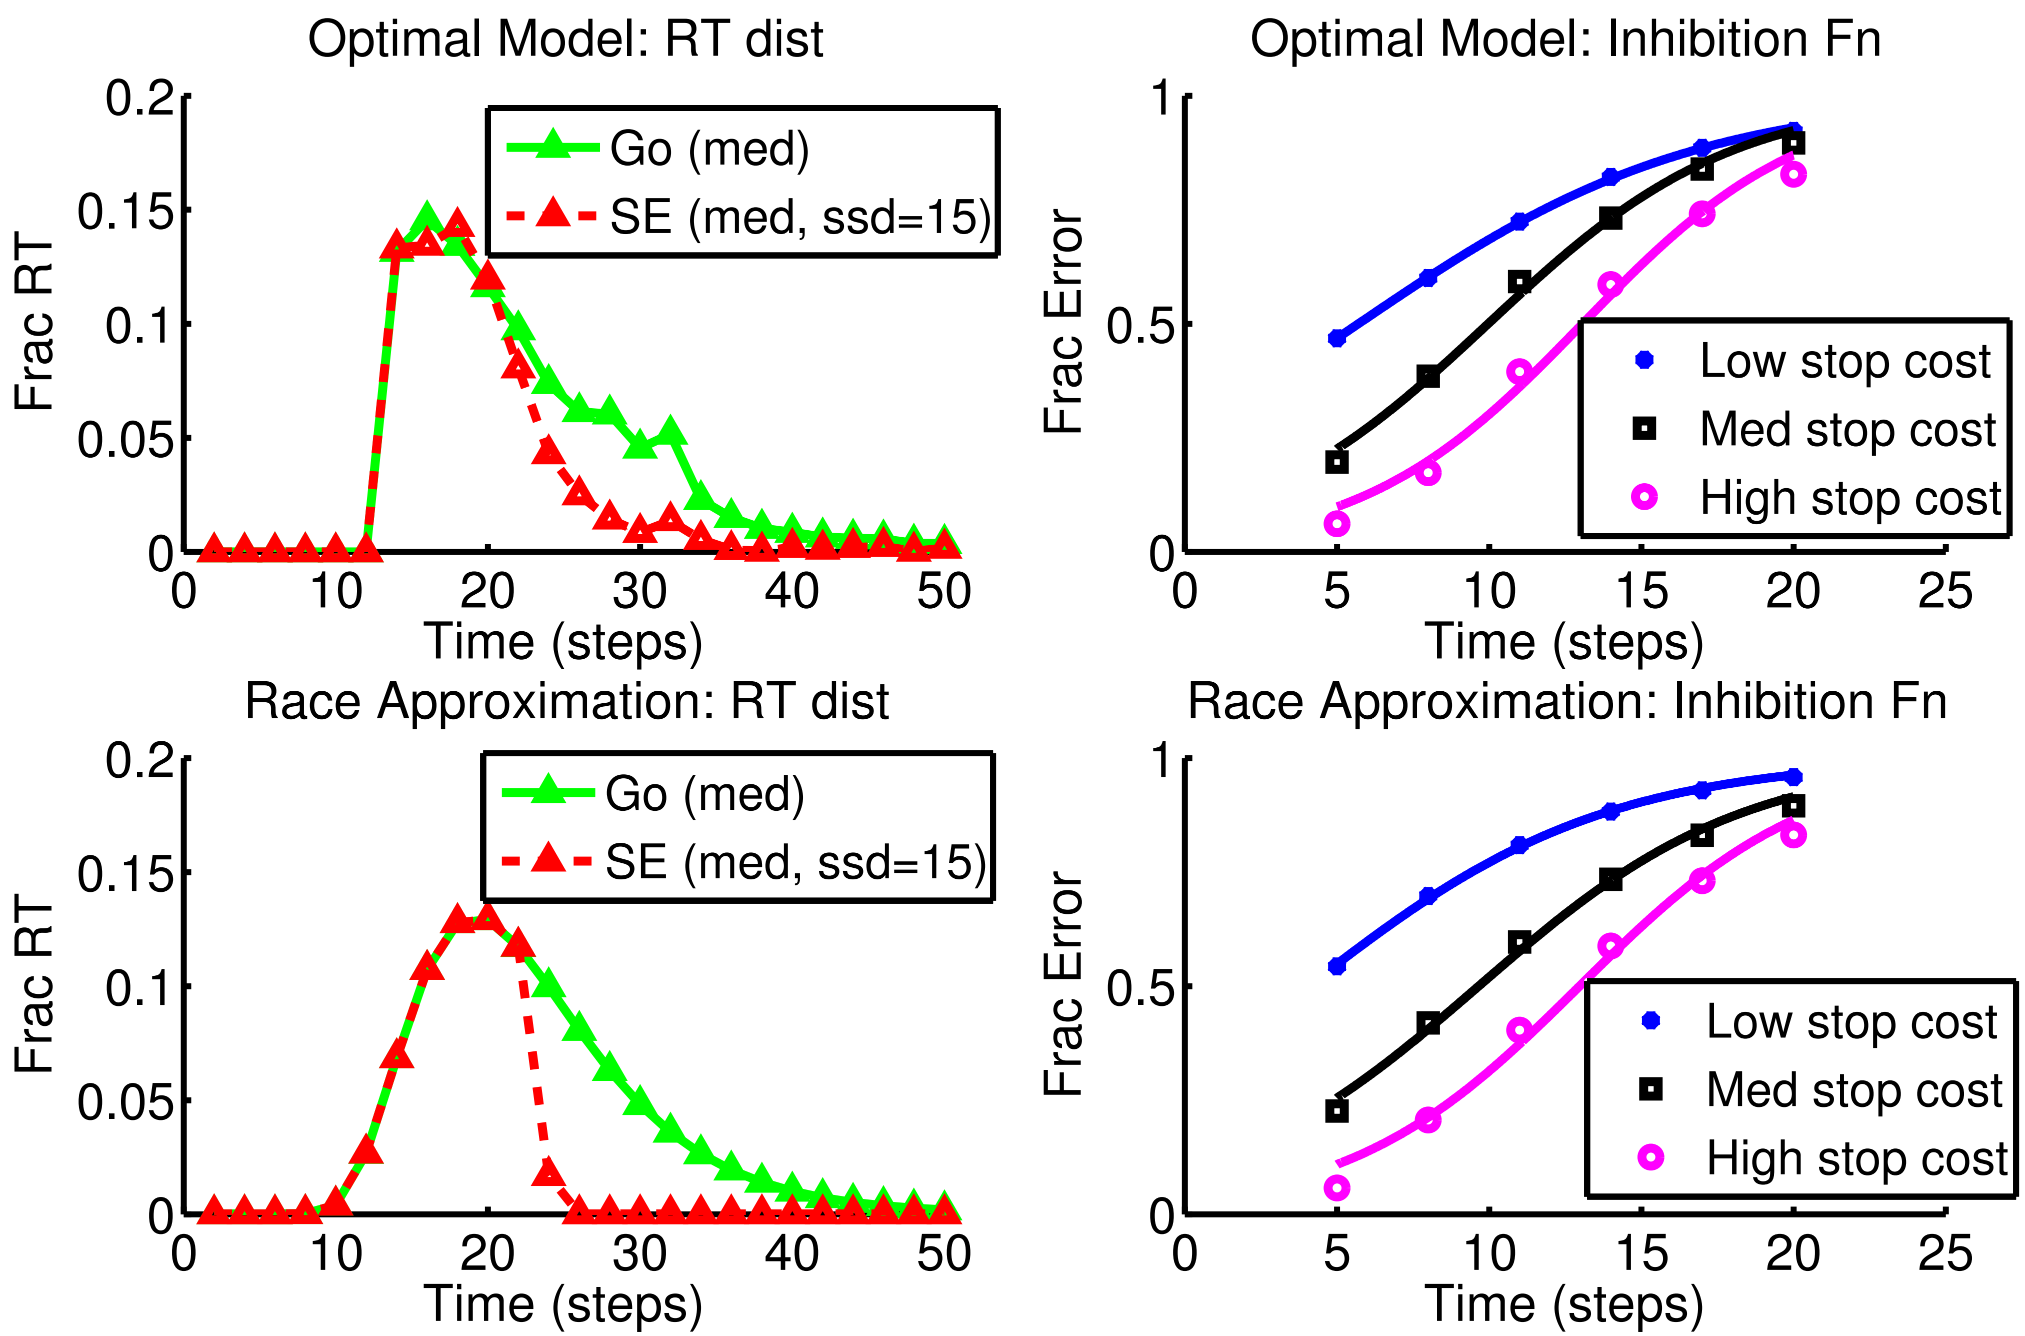

Reward/Motivation ⇒ Stopping Behavior

Race Model *Approximation* to Rational Decision-Making

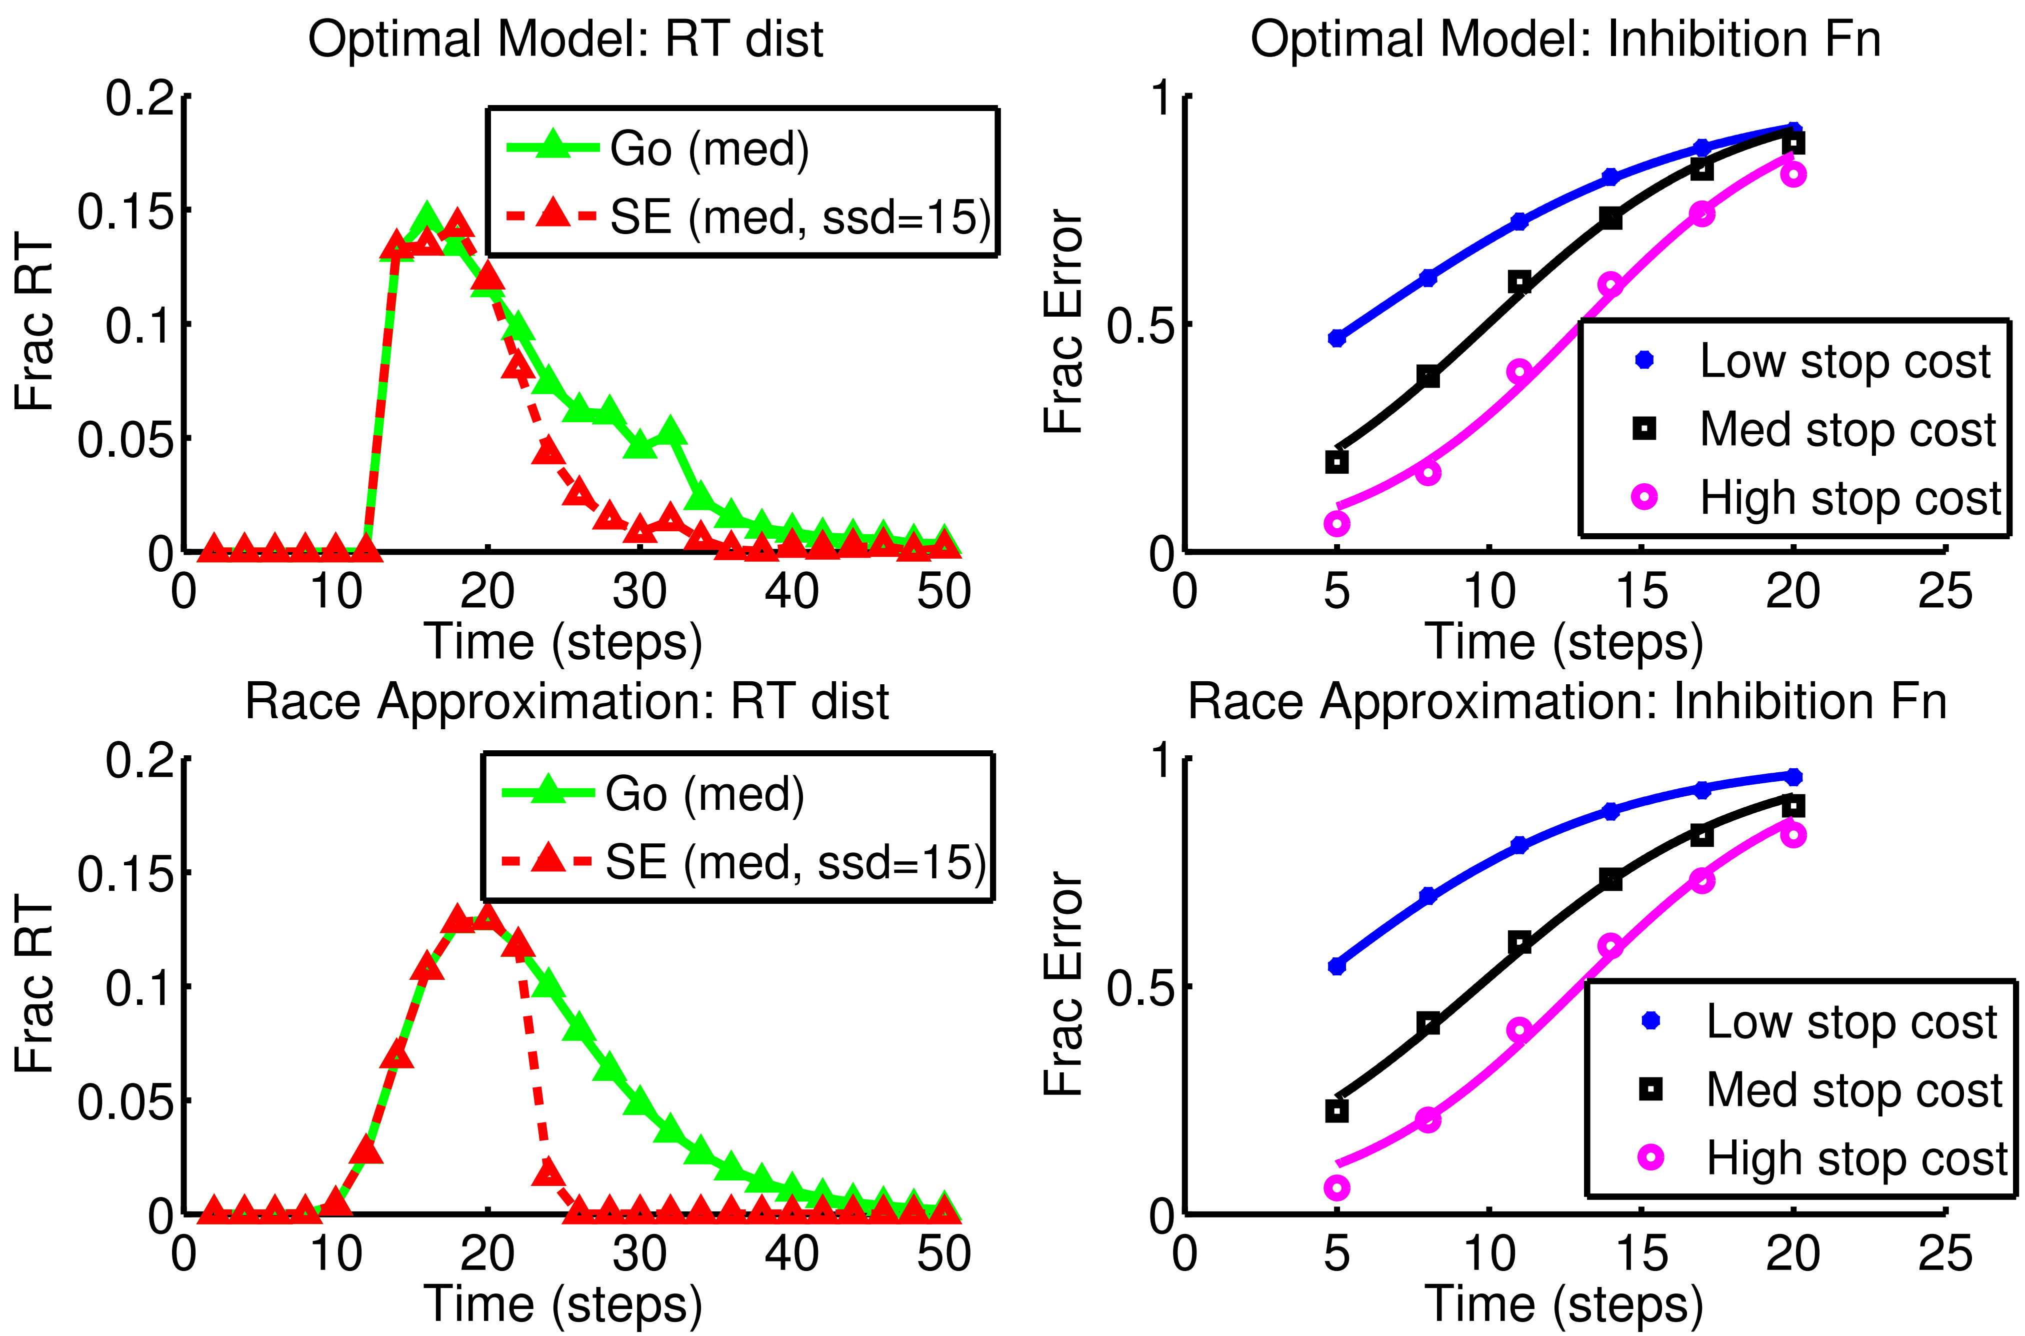

Rational Agent: Perceptual Discrimination

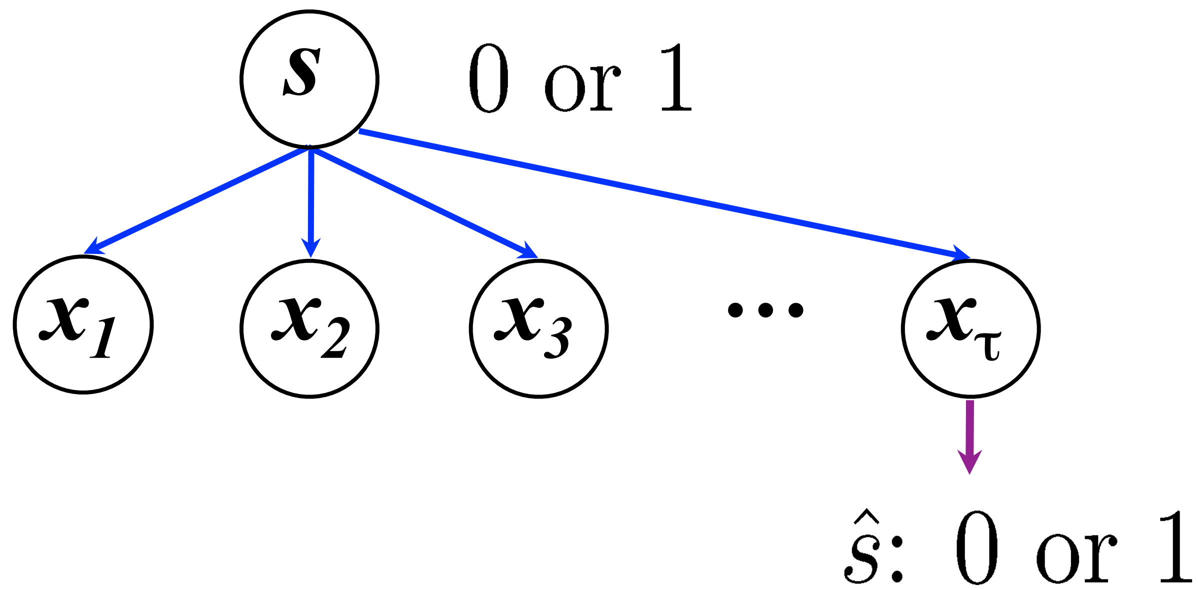

Rational Agent: Perceptual Discrimination

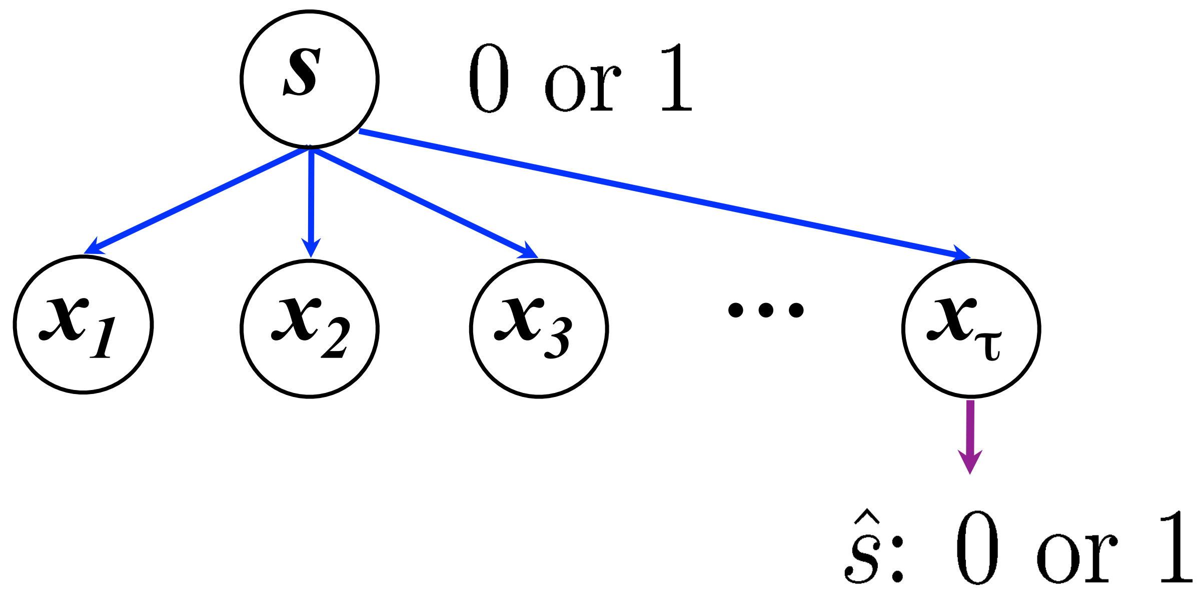

Monitoring Process

Incorporate evidence iteratively (Bayes' Rule):

$$q_t \triangleq P(s=1|x_1,\ldots,x_t) = \frac{p(x_t|s=1)q_{t-1}}{p(x_t|x_1,\ldots,x_{t-1})}$$

Rational Agent: Perceptual Discrimination

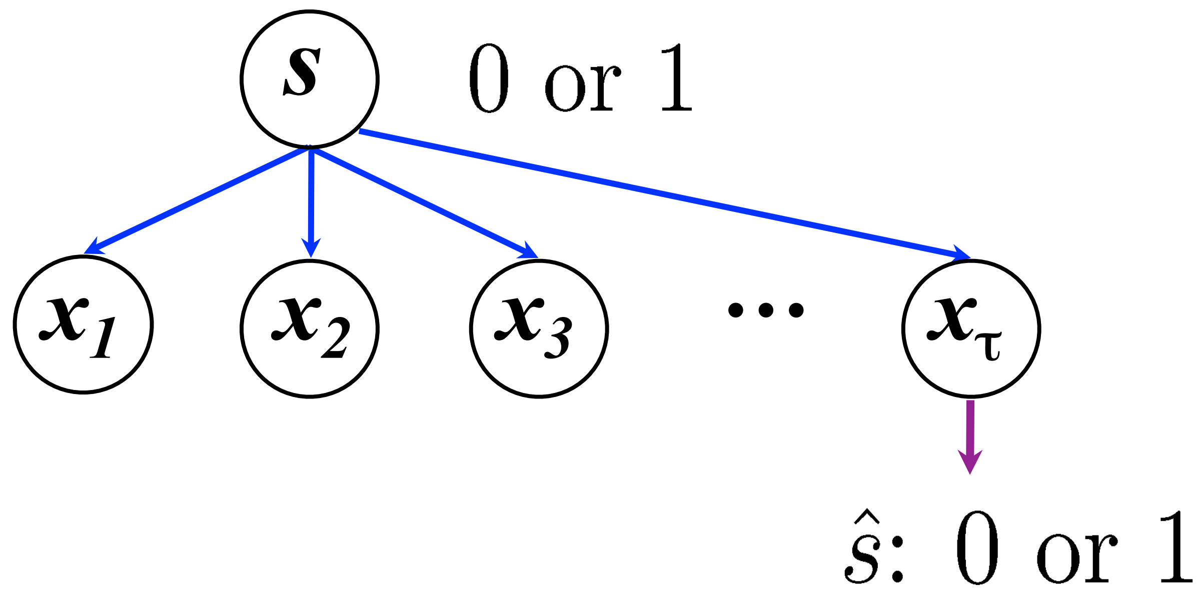

Monitoring Process

Incorporate evidence iteratively (Bayes' Rule):

$$q_t \triangleq P(s=1|x_1,\ldots,x_t) = \frac{p(x_t|s=1)q_{t-1}}{p(x_t|x_1,\ldots,x_{t-1})}$$

Decision Process

$$\pi(x_1,\ldots,x_t) \rightarrow \{0, 1, \text{cont}\}$$

Rational Agent: Perceptual Discrimination

Optimal decision policy (SPRT; Wald & Wolfowitz, 1948)

Rational Agent: Perceptual Discrimination

Optimal decision policy (SPRT; Wald & Wolfowitz, 1948)

- at time t, ***wait*** if $b < q_t < a$
- ***go*** & choose $\hat{s} = 1$ if $q_t > b$, choose $\hat{s} = 0$ if $q_t < a$

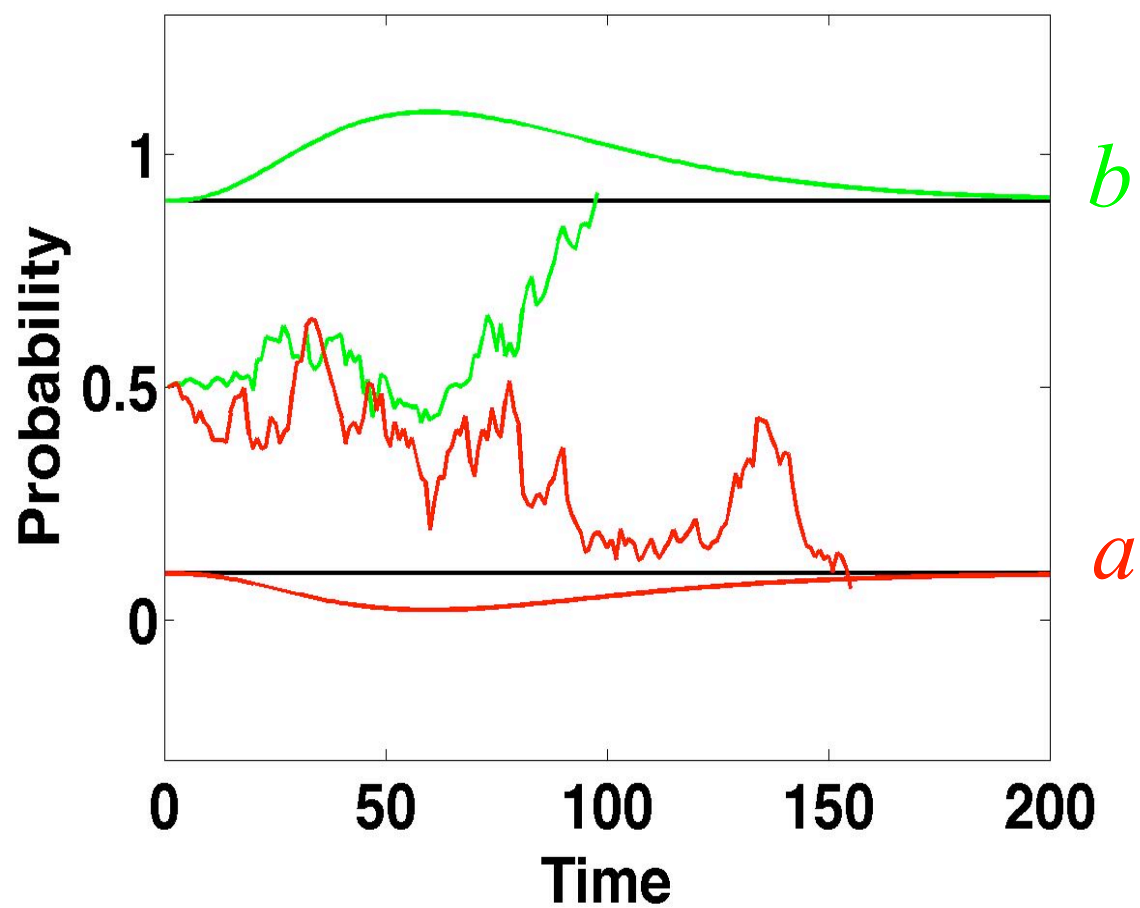

Rational Agent: Perceptual Discrimination

Optimal decision policy (SPRT; Wald & Wolfowitz, 1948)

- at time t, ***wait*** if $b < q_t < a$
- ***go*** & choose $\hat{s} = 1$ if $q_t > b$, choose $\hat{s} = 0$ if $q_t < a$

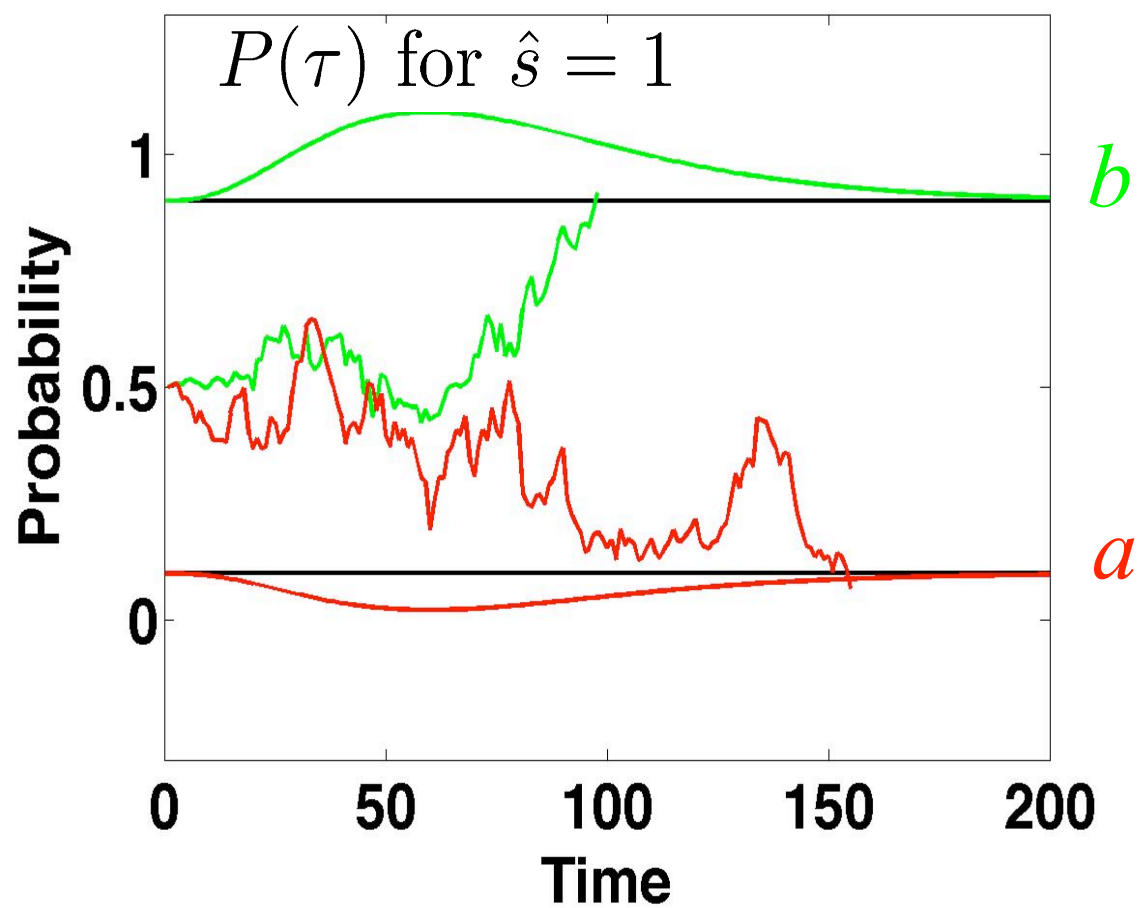

Rational Agent: Perceptual Discrimination

Optimal decision policy (SPRT; Wald & Wolfowitz, 1948)

- at time t, ***wait*** if $b < q_t < a$
- ***go*** & choose $\hat{s} = 1$ if $q_t > b$, choose $\hat{s} = 0$ if $q_t < a$

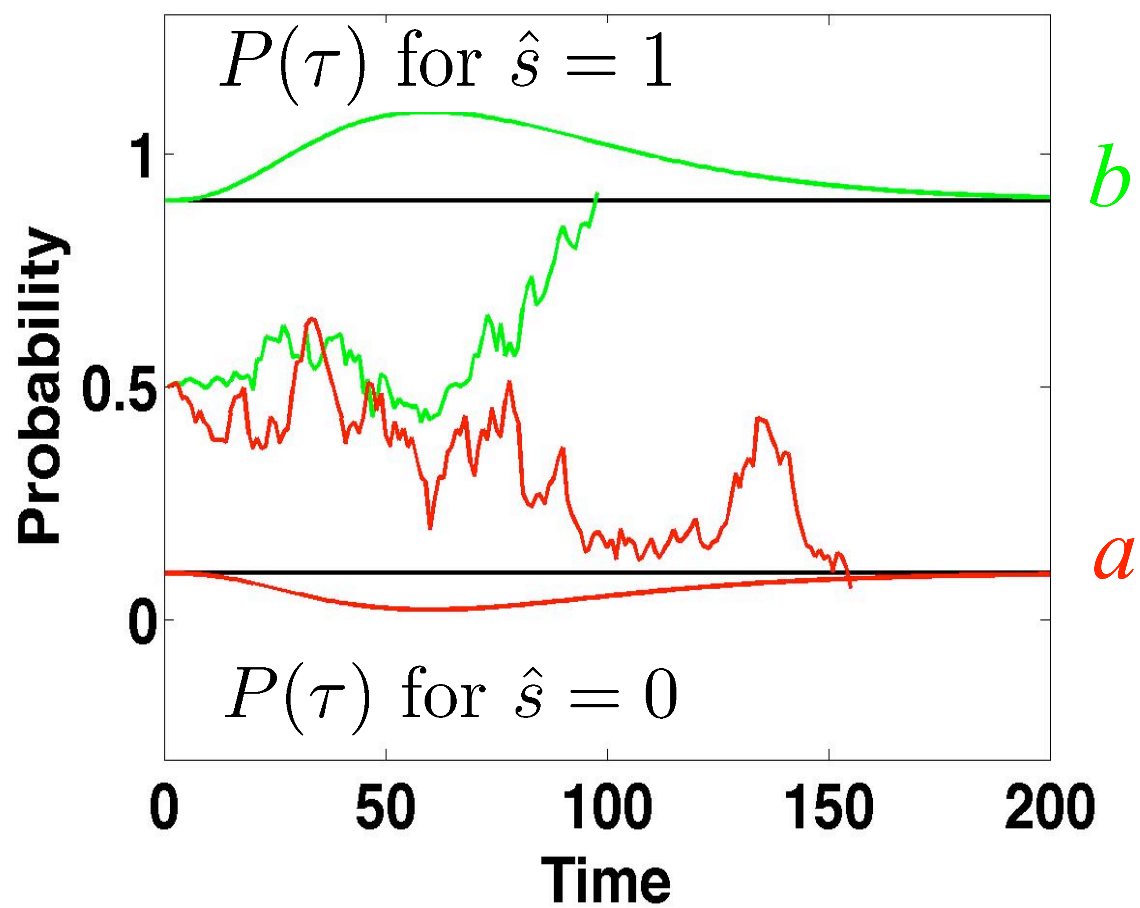

Rational Agent: Perceptual Discrimination

Optimal decision policy (SPRT; Wald & Wolfowitz, 1948)

- at time t, ***wait*** if $b < q_t < a$
- ***go*** & choose $\hat{s} = 1$ if $q_t > b$, choose $\hat{s} = 0$ if $q_t < a$
- model of both accuracy and RT

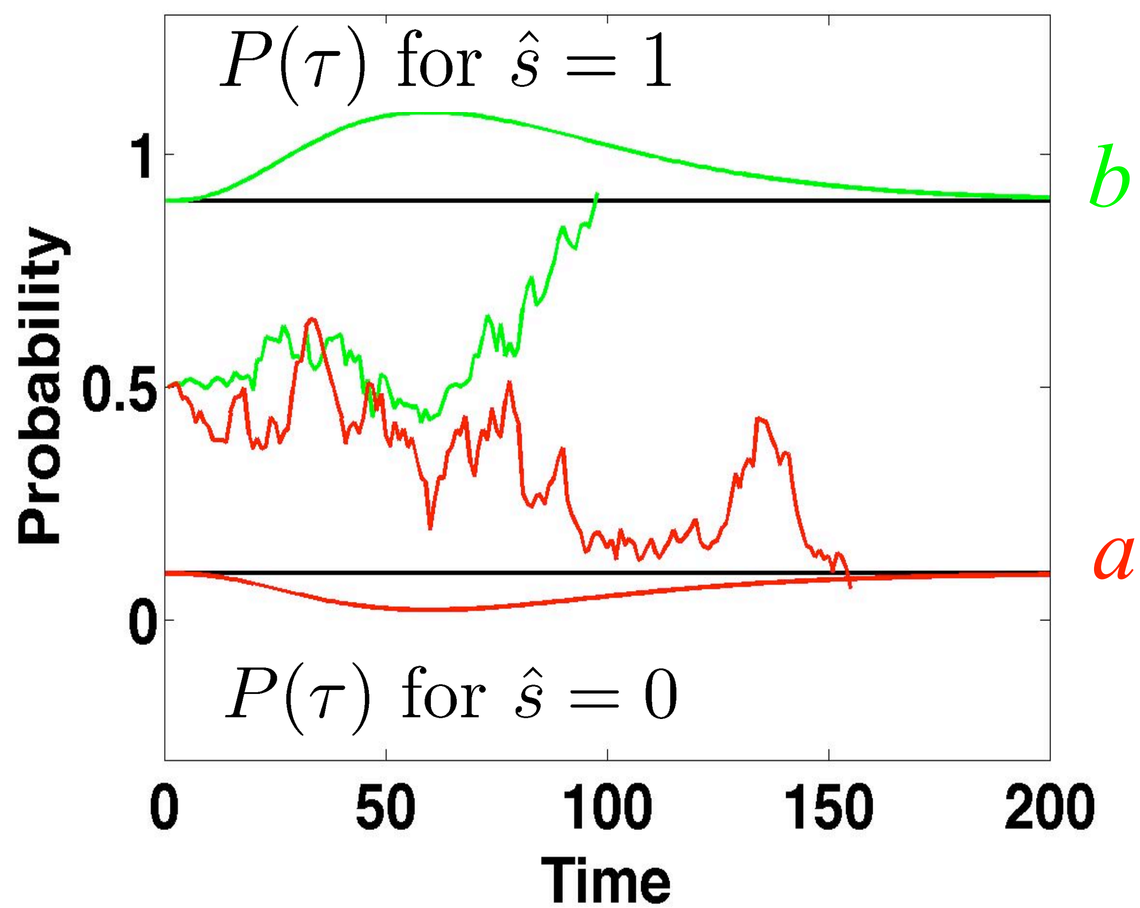

Rational DM Explains Behavioral Data

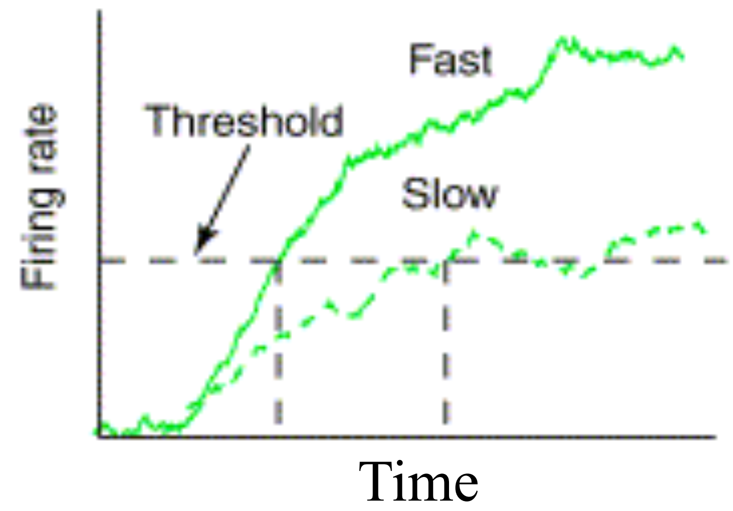

(from Smith & Ratcliff, 2004)

Accuracy vs. Coherence

RT vs. Coherence

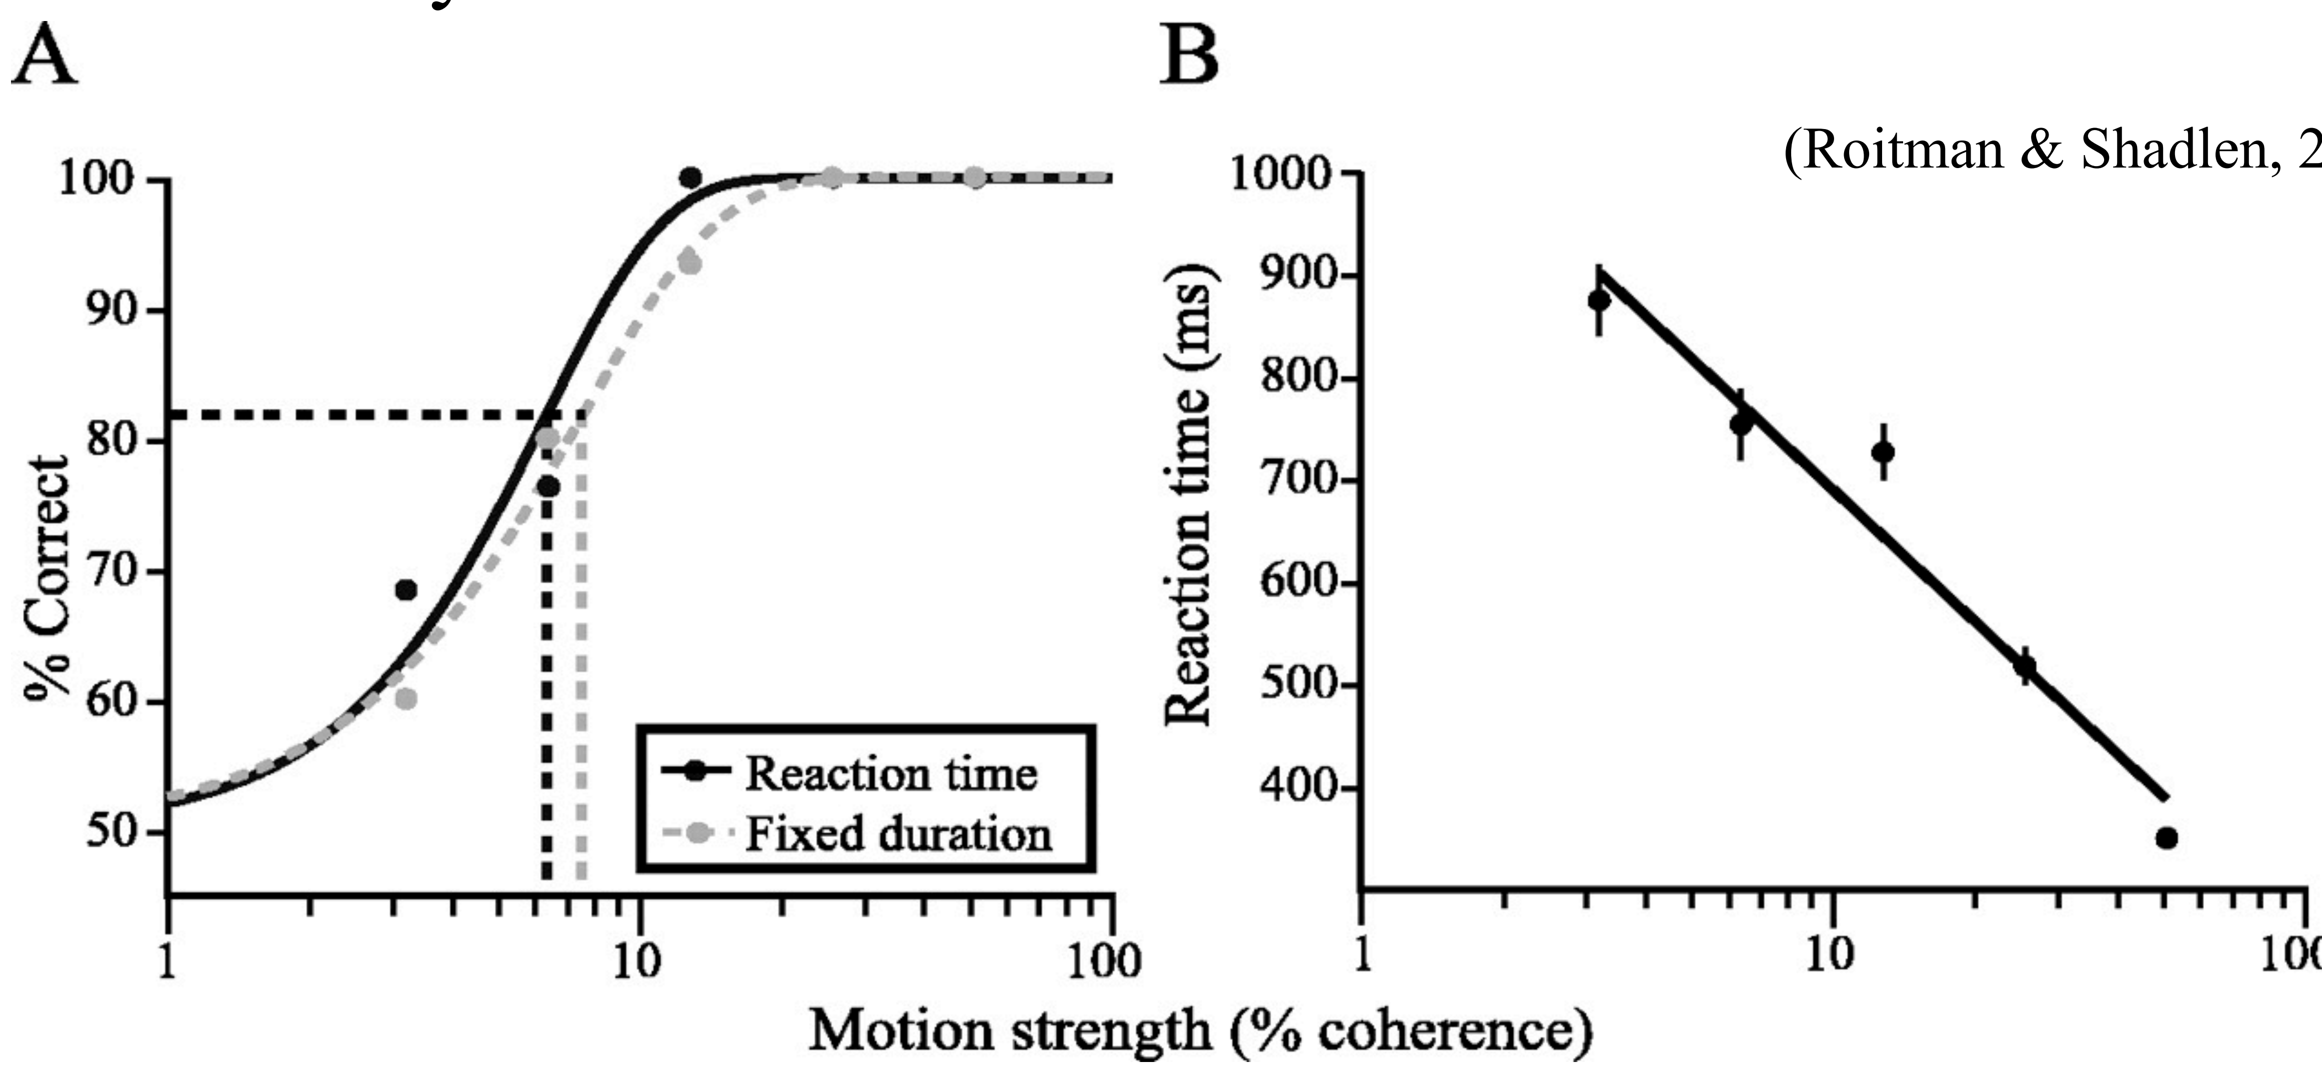

(Roitman & Shadlen, 2002)

Rational DM Explains Neural Data

Saccade generation

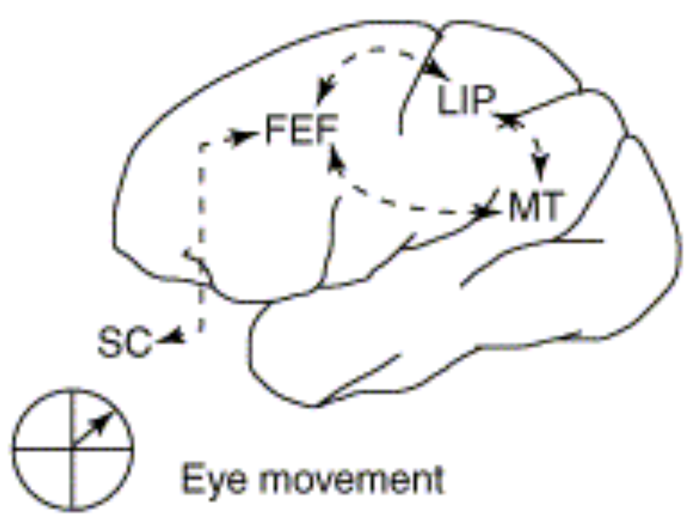

(from Smith & Ratcliff, 2004)

Rational DM Explains Neural Data

LIP = neural SPRT integrator?

(Roitman & Shalden, 2002; Gold & Shadlen, 2004)

Saccade generation

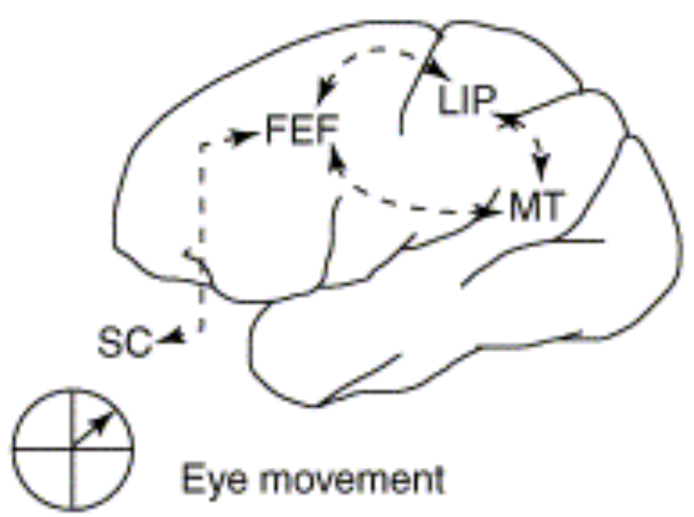

(from Smith & Ratcliff, 2004)

LIP Response & Coherence

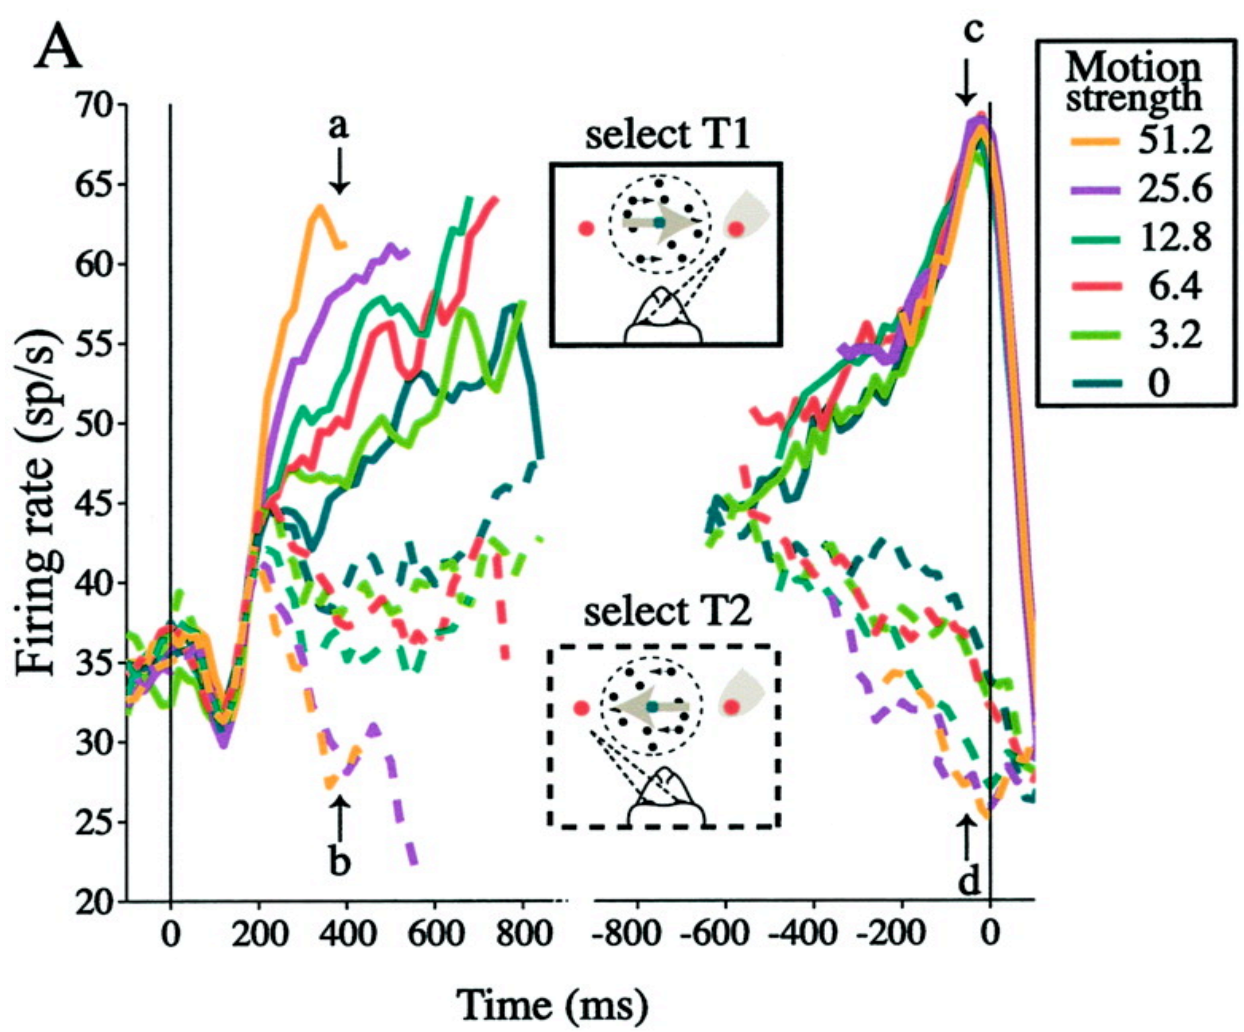

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

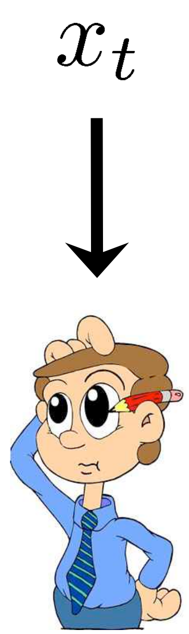

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

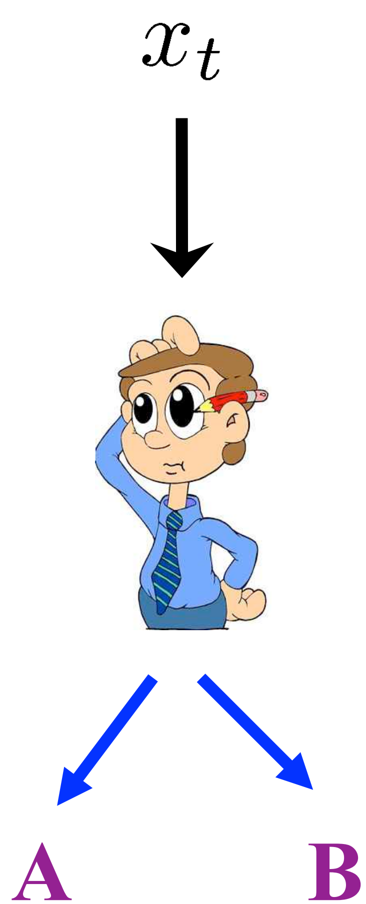

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

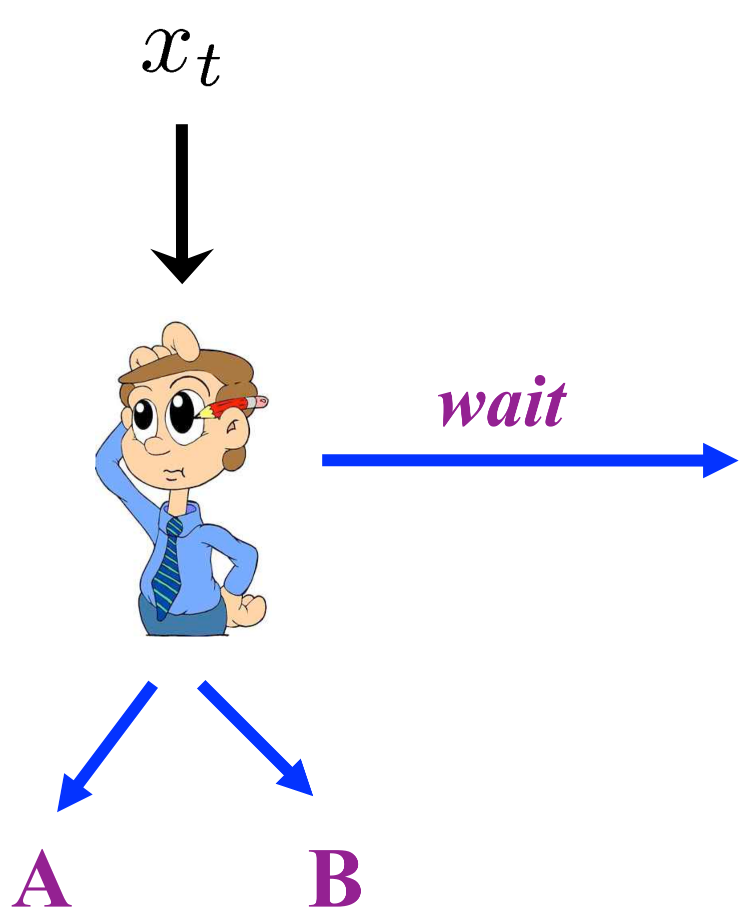

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

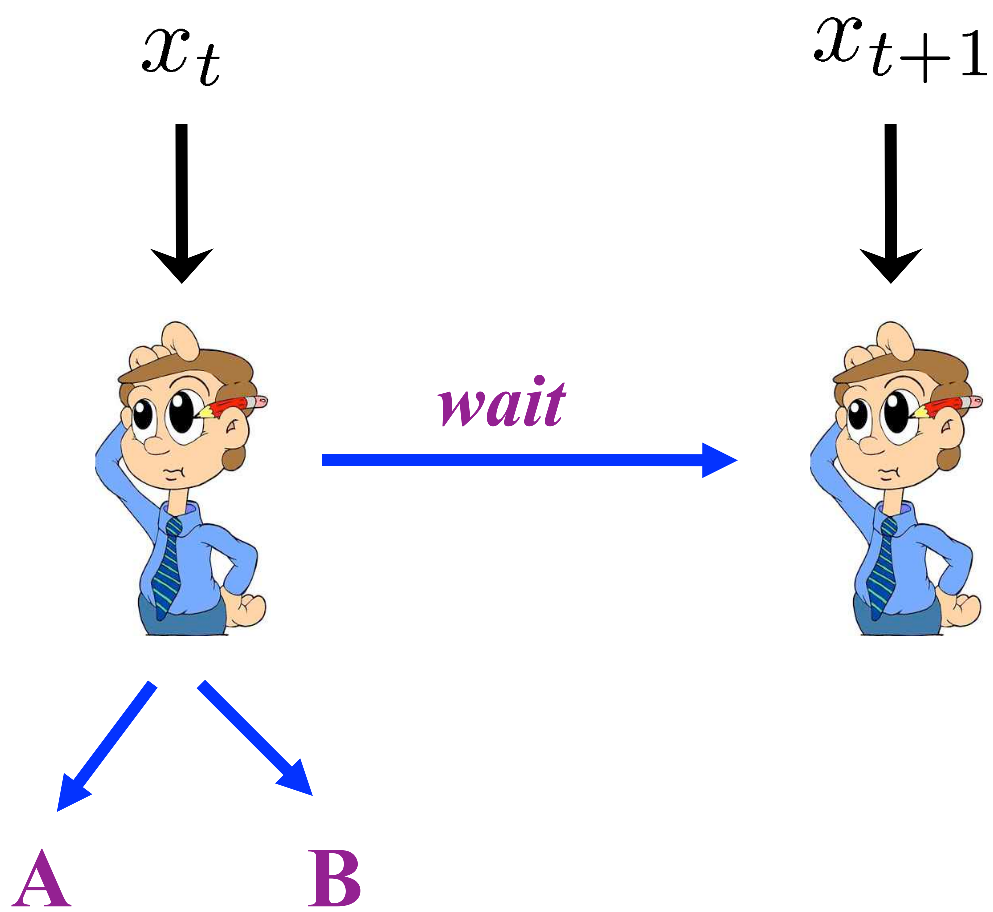

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

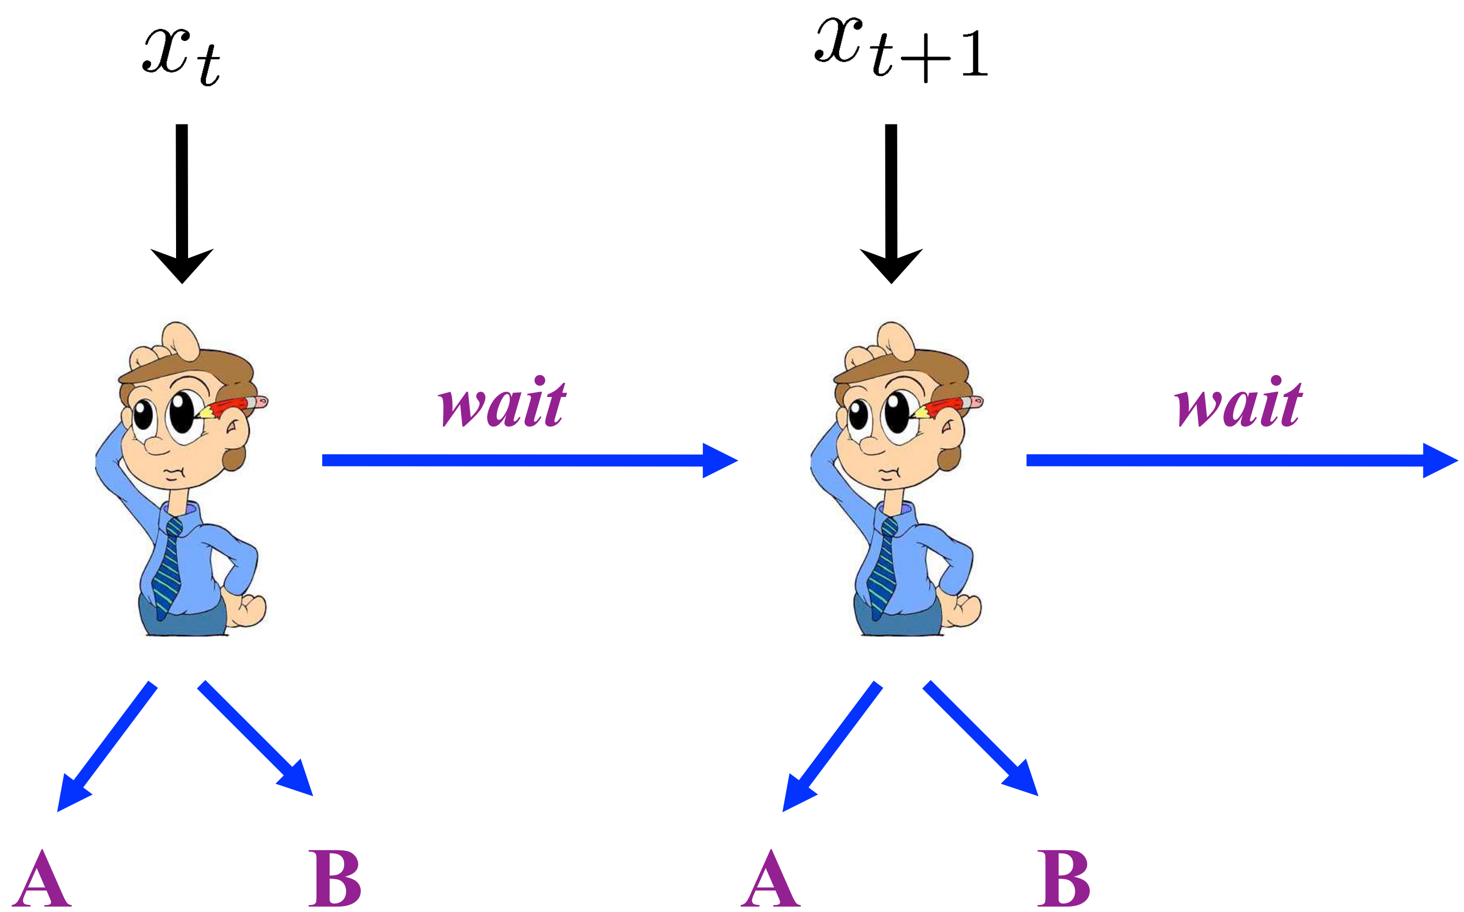

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

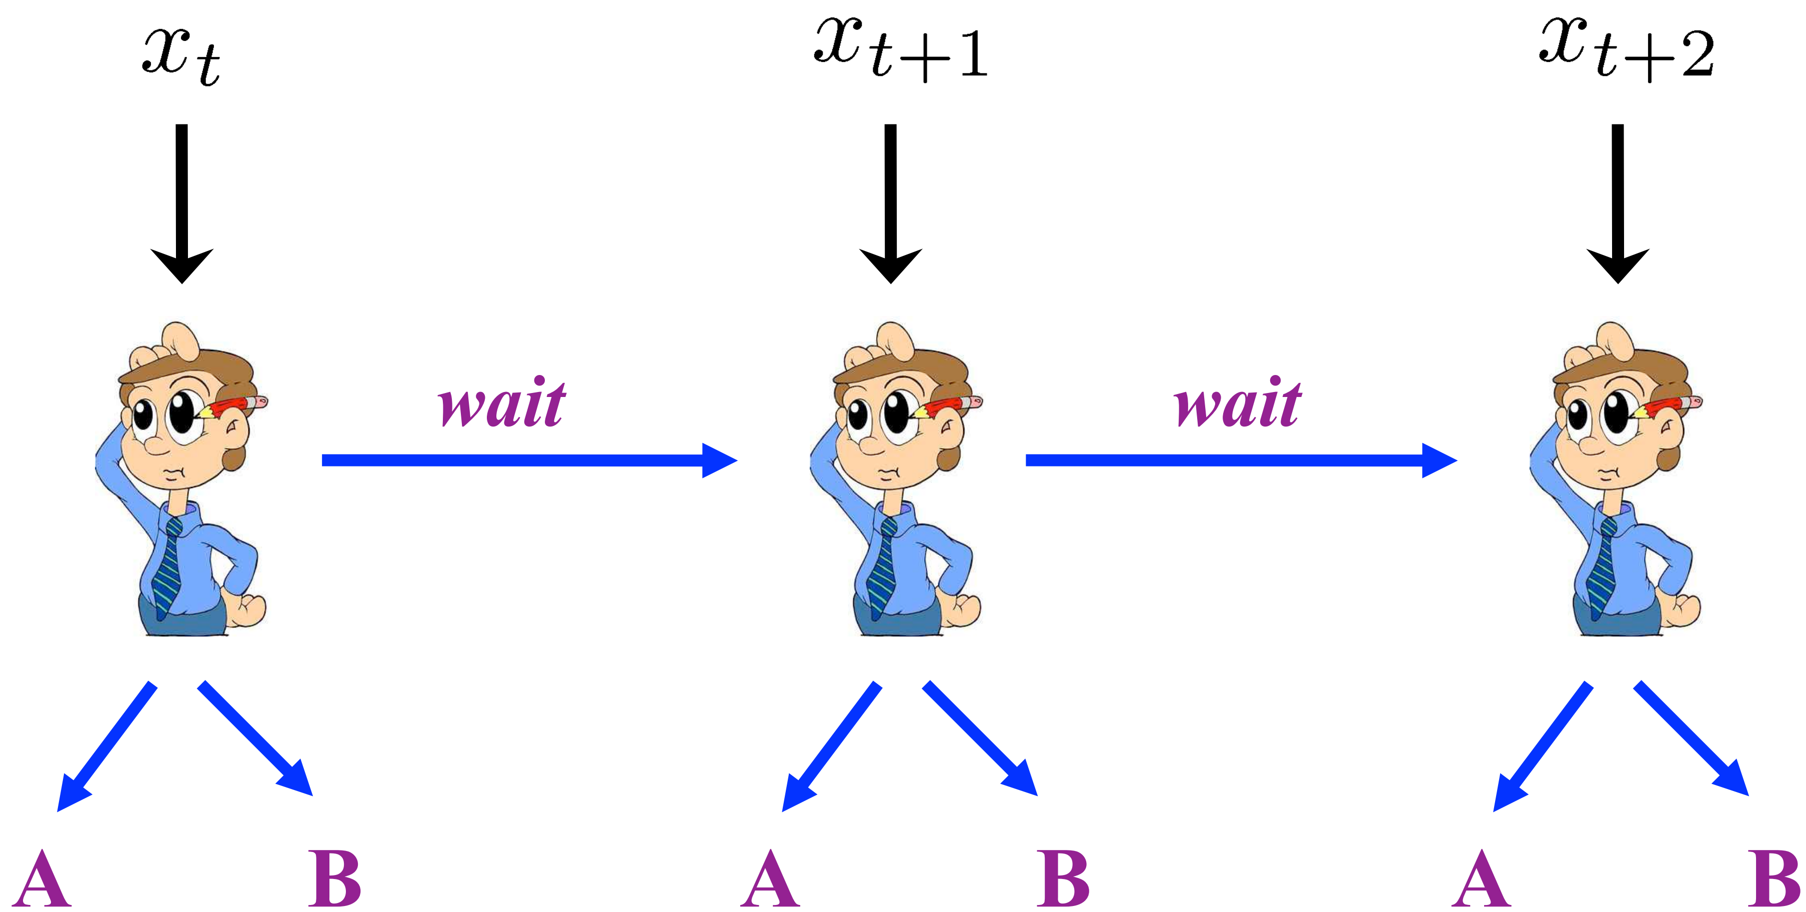

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making

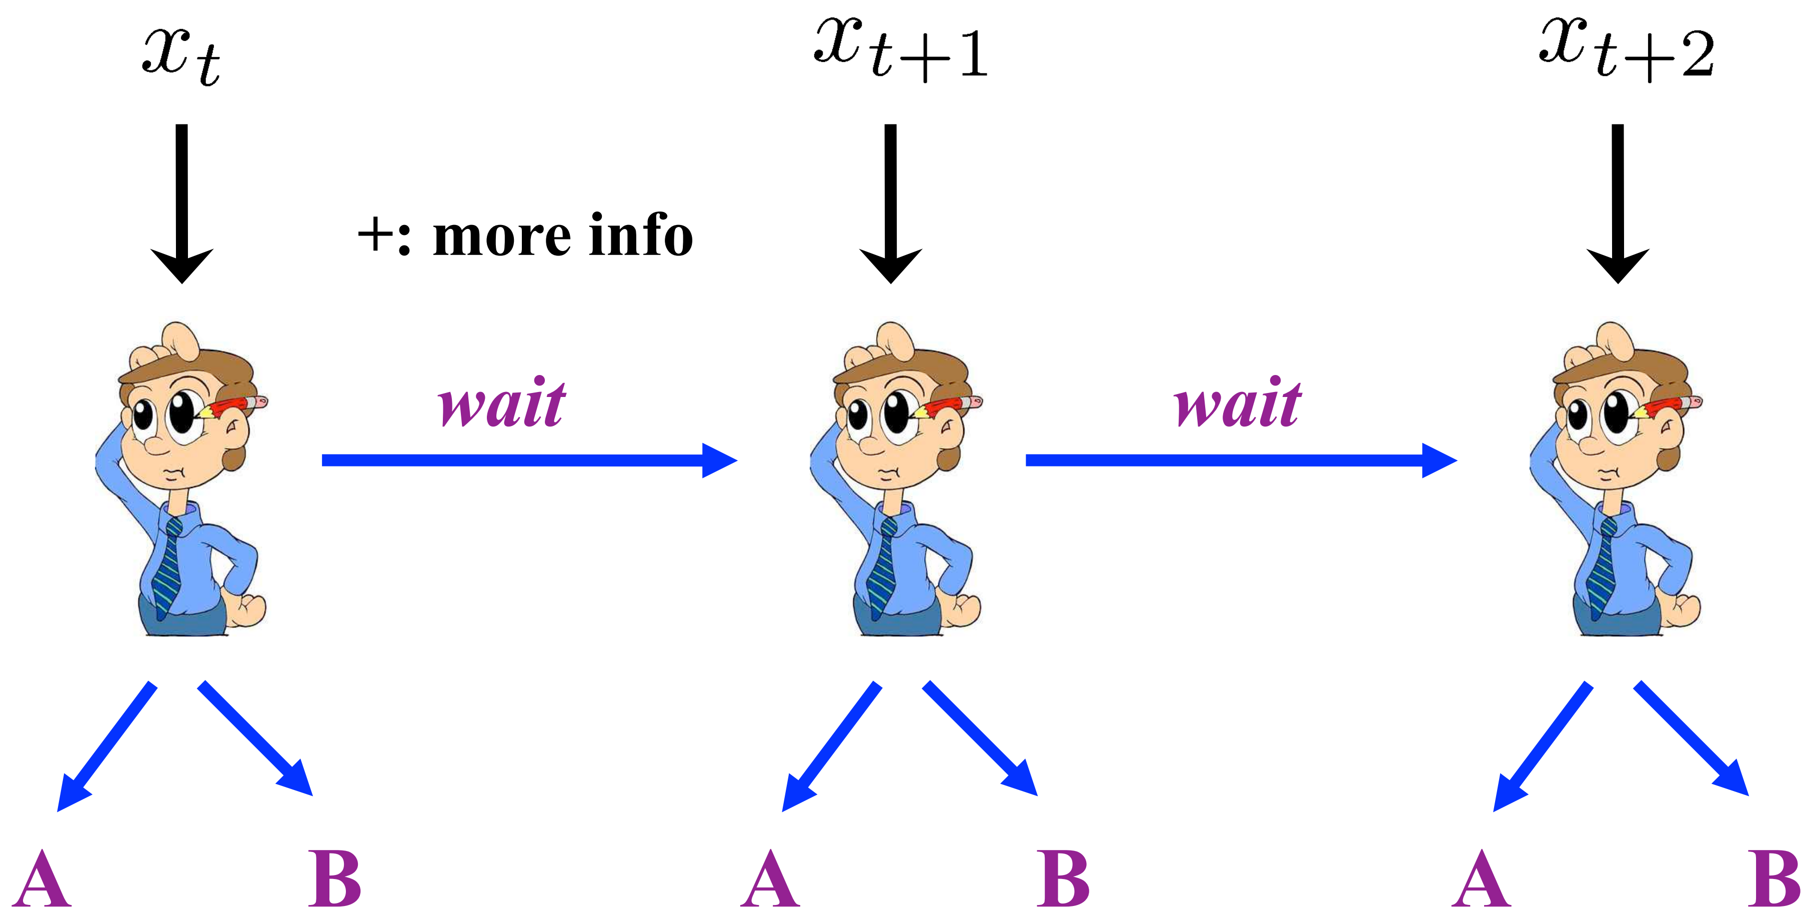

Rational Agent: Perceptual Discrimination

Problem: sequential decision-making