The neural mechanisms of dynamic decisions

Paul Cisek

Groupe de recherche sur le système nerveux central GRSNC

University of Barcelona
Spring Seminar Series in Theoretical Neuroscience
April 25, 2023

Outline

- Quantitative models of decisions
  - The drift-diffusion model: a good foundation
  - How we can/should move beyond it
- What is the evidence to help us decide between models?
  - Behavioral studies
  - Neural recording studies
- Broader implications
  - Individual differences
  - Pathological conditions

Random dot motion discrimination task

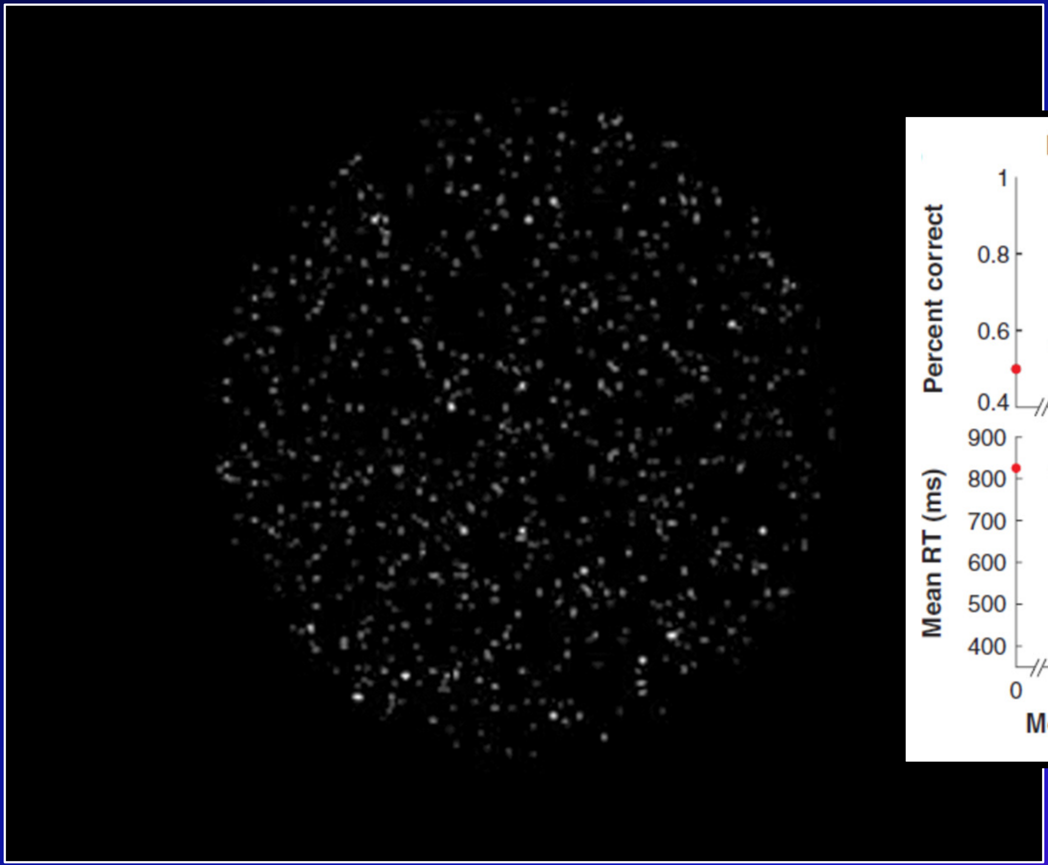

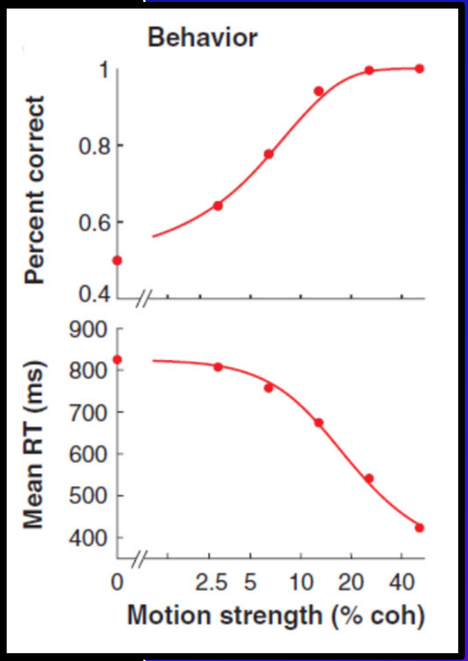

Gold & Shadlen (2007) *Annual Reviews in Neuroscience*

Which way are the dots moving?

How do you make decisions over time?

- ***The "Drift-diffusion model" (DDM)*** Ratcliff (1978) *Psychological Review*
  - Hypothesis: Deliberation is similar to a random walk to a threshold

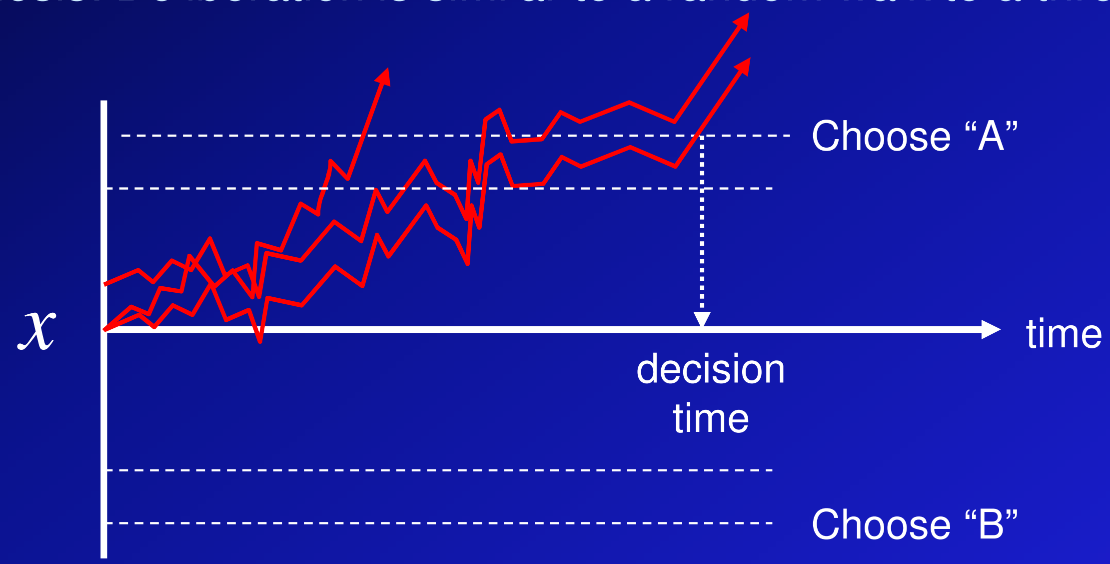

  - Some noisy neural variable (x) is changing over time, biased by sensory information, until it crosses a decision threshold
    - The strength of the evidence determines the rate of drift
    - Any prior information determines the starting point
    - The desired accuracy determines the threshold

The Drift-Diffusion Model (DDM)

- An integrator of evidence to a constant threshold

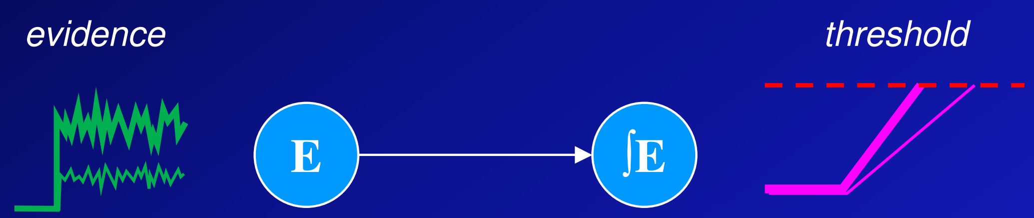

$$x_A(t) = Z_A + \int_0^t E_A(\tau)d\tau, \text{ until } x_A(t) = T$$

DDM ≈ sequential probability ratio test

- If you want to achieve a given level of accuracy α (e.g. 95%), how many samples should you take from the world before you commit to choice A or B?

$$x_A(n) = \log \frac{p(A|s_1 \ldots s_n)}{p(B|s_1 \ldots s_n)} < \log \frac{\alpha}{1-\alpha}$$

- From Bayes' Rule:

$$x_A(n) = \underbrace{\log \frac{p(A)}{p(B)}}_{\text{prior}} + \underbrace{\sum_{k=1}^{n} \log \frac{p(s_k|A)}{p(s_k|B)}}_{\substack{\text{accumulated evidence} \\ \text{(sum of log-likelihood ratios)}}} < \underbrace{\log \frac{\alpha}{1-\alpha}}_{\text{threshold}}$$

- Integration of evidence to a threshold implements the decision policy that optimizes time for any desired accuracy

$$x_A(t) = Z_A + \int_0^t E_A(\tau)\,d\tau < T$$

Behavioral evidence for the DDM

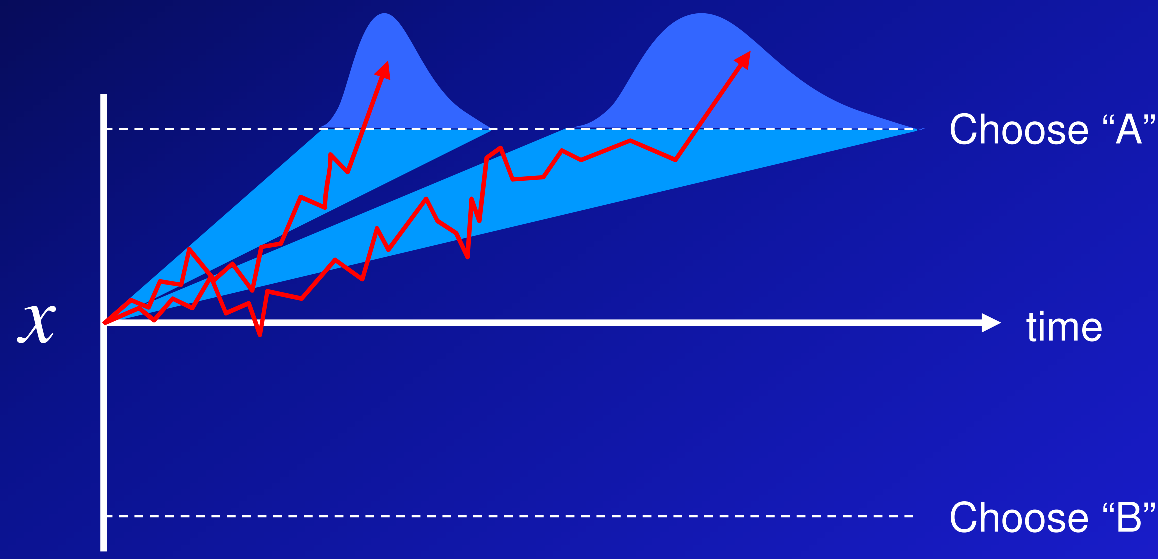

- Because of noise, the process will create RT distributions
  - Weaker sensory evidence, later decisions and broader distributions
- This is well-supported by behavioral data from many decision tasks

Neural evidence for the DDM

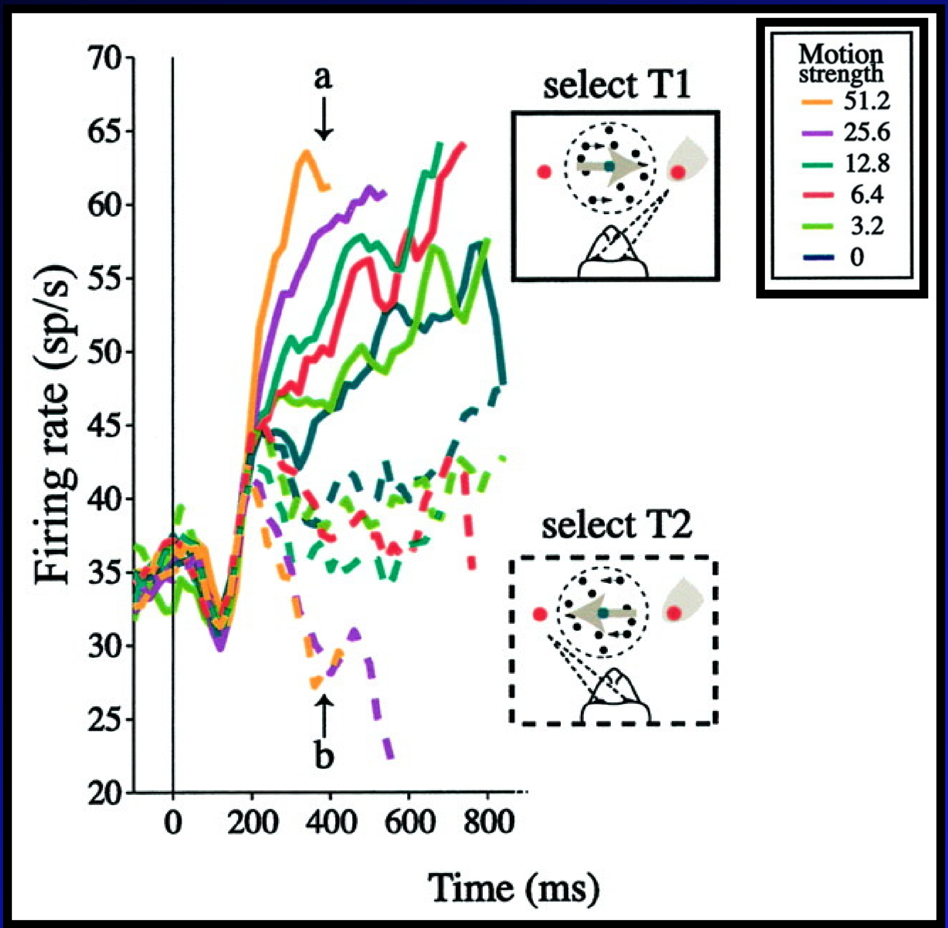

Roitman & Shadlen (2002) *J. Neurosci.*

- Neural activity in the lateral intraparietal area (LIP) grows at a rate related to the strength of evidence for a cell's preferred target
- Activity reaches a similar level just before the saccade
- Conclusion: LIP represents the integrated motion signal

The Drift-Diffusion Model (DDM)

- An integrator of evidence to a constant threshold

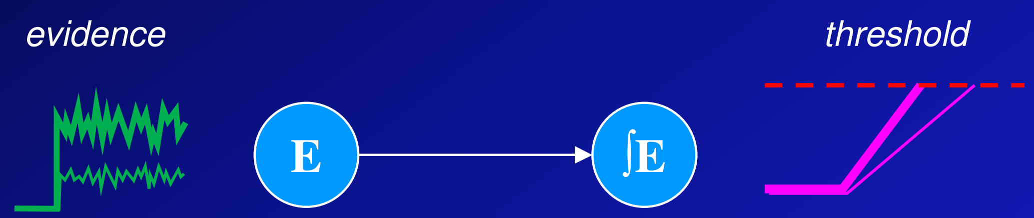

$$x_A(t) = Z_A + \int_0^t E_A(\tau)d\tau, \text{ until } x_A(t) = T$$

$$Z_A = \log\frac{p(A)}{p(B)},\ E_A(\tau) = \log\frac{p(s(\tau)|A)}{p(s(\tau)|B)},\ T = \log\frac{1-\alpha}{\alpha}$$

- The diffusion model is widely accepted as the explanation for decision-making during perceptual discrimination tasks
  - Makes good mathematical sense
  - Explains behavioral data on accuracy and RT distributions
  - Explains the widespread build-up of neural activity during decisions

However…

1. What if the world changes?
2. What do we want to integrate?
3. What do we want to optimize?

Q1: What if the world changes?

- During natural behavior, the world is always changing
- Any integrator will be sluggish in its response to such changes
- Need to be able to quickly respond to changes
  - Reset?
  - Sudden increase in gain?
  - Some other mechanism?

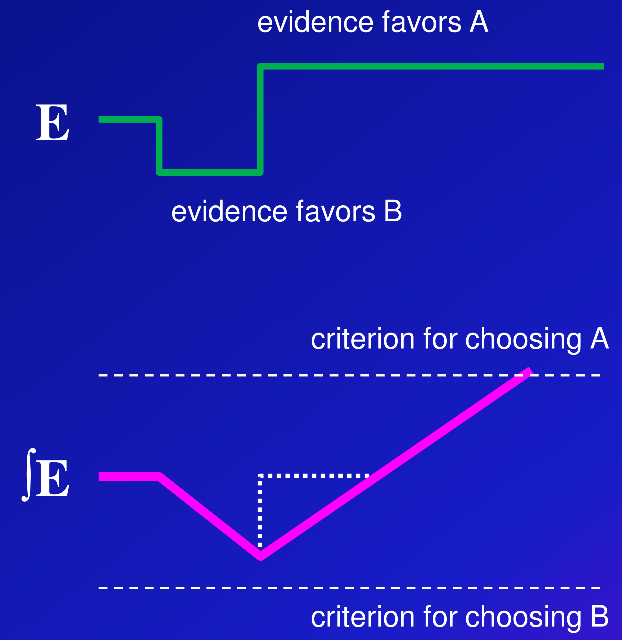

?

Q2. What ***should*** be integrated?

$$\log\frac{p(A)}{p(B)}+\sum_{k=1}^{n}\log\frac{p(s_k|A)}{p(s_k|B)}<\log\frac{1-\alpha}{\alpha}$$

- This assumed that sequential samples are ***statistically independent***
  - But that is not true in general. If you're looking at the same stimulus, later samples are partially or completely ***redundant***
  - To take this redundancy into account, we need to extend Bayes' Rule to $k{+}1$ variables, yielding:

$$\log\frac{p(A)}{p(B)}+\sum_{k=1}^{n}\log\frac{p(s_k|A,s_1,\ldots,s_{k-1})}{p(s_k|B,s_1,\ldots,s_{k-1})}<\log\frac{1-\alpha}{\alpha}$$

  - If sample s_k is completely independent of previous samples, then $p(s_k|X,s_1,\ldots,s_{k-1})=p(s_k|X)$, and this reduces to the equation on top i.e. **you *should* sum independent samples**
  - If sample s_k is completely predicted by previous samples, then $p(s_k|X,s_1,\ldots,s_{k-1})=1$, and $\log\frac{1}{1}=0$ i.e. **you should *ignore* redundant samples**

So what should we accumulate?

Should we accumulate ***all*** evidence?

$$\boxed{\log\frac{p(A)}{p(B)}+\sum_{k=1}^{n}\log\left(\frac{I_{k|A}}{I_{k|B}}\times\frac{p(s_k|A)}{p(s_k|B)}\right)}$$

$$I_{k|X}=\frac{p(s_1,\ldots,s_{k-1},s_k|X)}{p(s_1,\ldots,s_{k-1}|X)p(s_k|X)}$$

related to the mutual information between sample k and previous ones

- <u>Conclusion</u>: We should only accumulate evidence to the extent that it is ***novel***
- i.e. In order to properly implement the SPRT, we have to take statistical dependence between samples into account
- Ex: buying newspapers

What's optimal vs. possible vs. reasonable?

- Ideal mechanism:
  Compute the degree to which each sample contributes novel information

$$x_A(n) = \log\frac{p(A)}{p(B)} + \log\frac{p(s_1|A)}{p(s_1|B)} + \log\left(\frac{\frac{p(s_1,s_2|A)}{p(s_1|A)p(s_2|A)}}{\frac{p(s_1,s_2|B)}{p(s_1|B)p(s_2|B)}} \times \frac{p(s_2|A)}{p(s_2|B)}\right)$$
$$+ \log\left(\frac{\frac{p(s_1,s_2,s_3|A)}{p(s_1,s_2|A)p(s_3|A)}}{\frac{p(s_1,s_2,s_3|B)}{p(s_1,s_2|B)p(s_3|B)}} \times \frac{p(s_3|A)}{p(s_3|B)}\right) + \log\left(\frac{\frac{p(s_1,s_2,s_3,s_4|A)}{p(s_1,s_2,s_3|A)p(s_4|A)}}{\frac{p(s_1,s_2,s_3,s_4|B)}{p(s_1,s_2,s_3|B)p(s_4|B)}} \times \frac{p(s_4|A)}{p(s_4|B)}\right) + \cdots$$

- Another approach:
  Try predicting the next sample and accumulate that which you can't predict

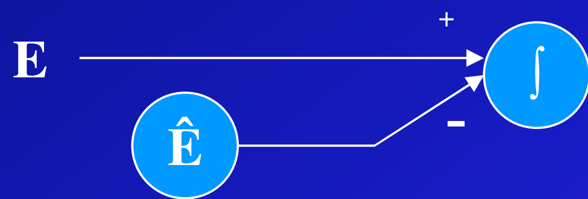

- Simple 1st-order approximation:
  - Novelty = change from previous sample
  - Assume fluctuations above a certain frequency are just noise
  - Low-pass filter

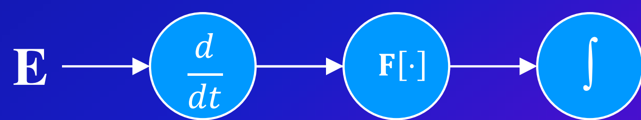

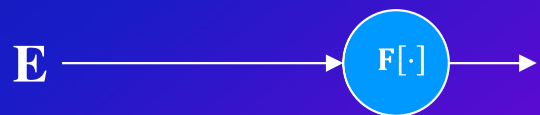

In a noisy task

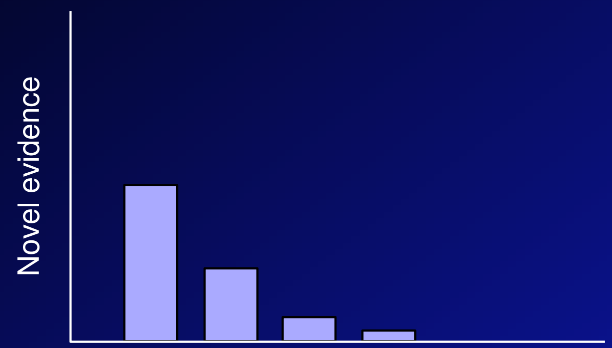

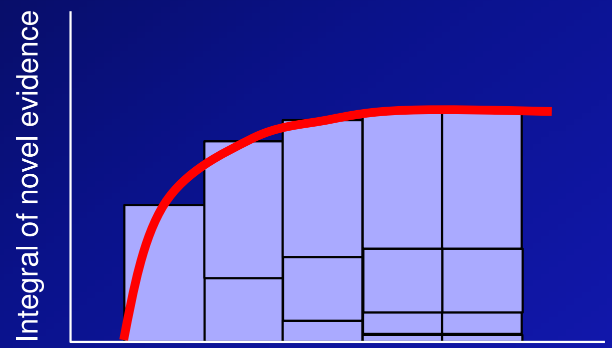

- First sample is independent
- Second gives some uncorrelated novel evidence
- Third sample gives less
- Etc.

- The integral…

  …looks like a low-pass filter

Leaky integration

- Accumulation of novelty will look like a low-pass filter
  - Leaky Competing Accumulator (Usher & McClelland, 2001)
- But then how do you get to the threshold?

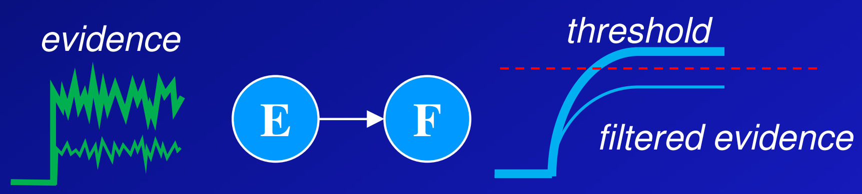

Q3: What kind of optimality do we want?

- If you're doing statistics then what you care about is ***accuracy***
  - The criterion of desired accuracy is determined by convention
    - $p<0.05$ = 95% confidence in your result
    - $p<0.01$ = 99% confidence in your result
  - The SPRT is optimal in that it minimizes the time required to reach that level of accuracy
    - If results not significant, get more data, even if it takes you another year to finish your thesis…
- But what if you're an animal in the wild?
  - Suppose you've reached 93% confidence
    - Do you wait to reach 95%?
      - What if it would take another thirty minutes?
      - What if it would take a year?
  - Time runs out, opportunities are lost, predators come and eat you!
  - You want to optimize ***reward rate***

Reward rate

- Reward Rate $$RR = \frac{p(t) \cdot U - C}{t + m + d}$$
- where
  - $p(t)$ is the probability of achieving a favorable outcome
  - t is the time spent deciding and planning
  - U is the subjective utility of that outcome
  - C is the subjective cost of trying
    (including "opportunity cost")
  - m is the time spent moving
  - d is the delay before you can try again
- "Time-discounted expected value"
- Similar form as "harvest intake" in foraging (Charnov, 1976)

Assumptions

$$RR = \frac{p(t) \cdot U - C}{t + m + d}$$

1. The probability of achieving a favorable outcome improves with time

2. But with diminishing returns

Therefore, the expected reward rate has a peak

– The peak of reward rate occurs when $RR'(t) = 0$ and $RR''(t) < 0$, or when

$$p(t) = p'(t)(t + m + d) + \frac{C}{U}$$

- This is the best time to commit

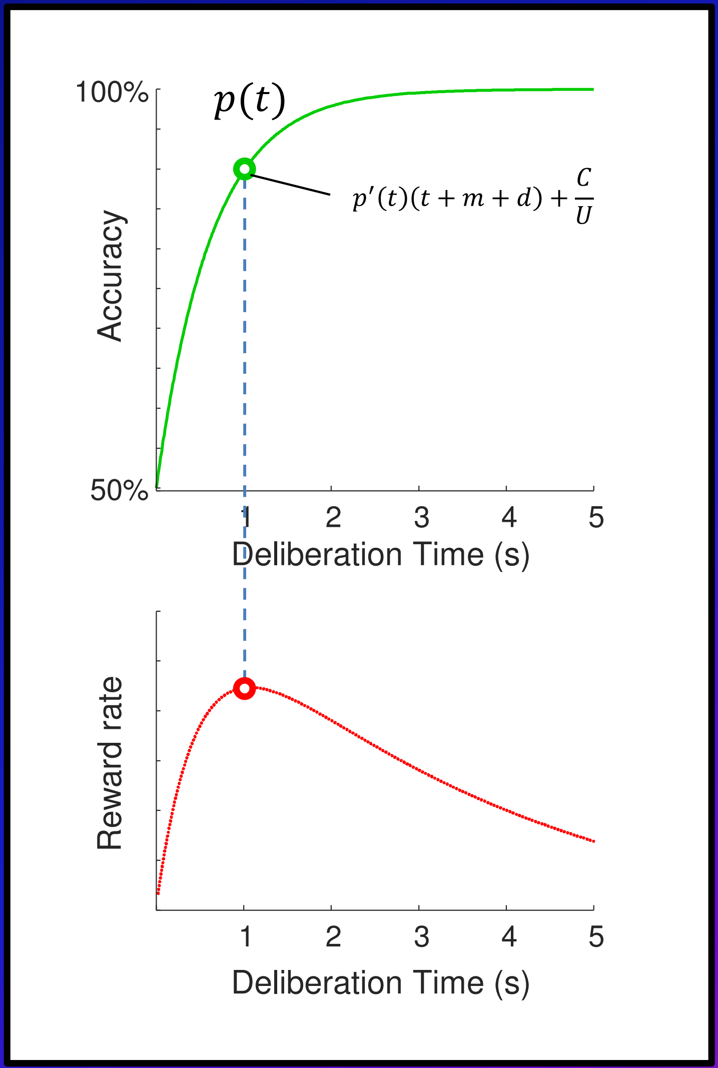

Geometric interpretation

$$p(t) = p'(t)(t + m + d) + \frac{C}{U}$$

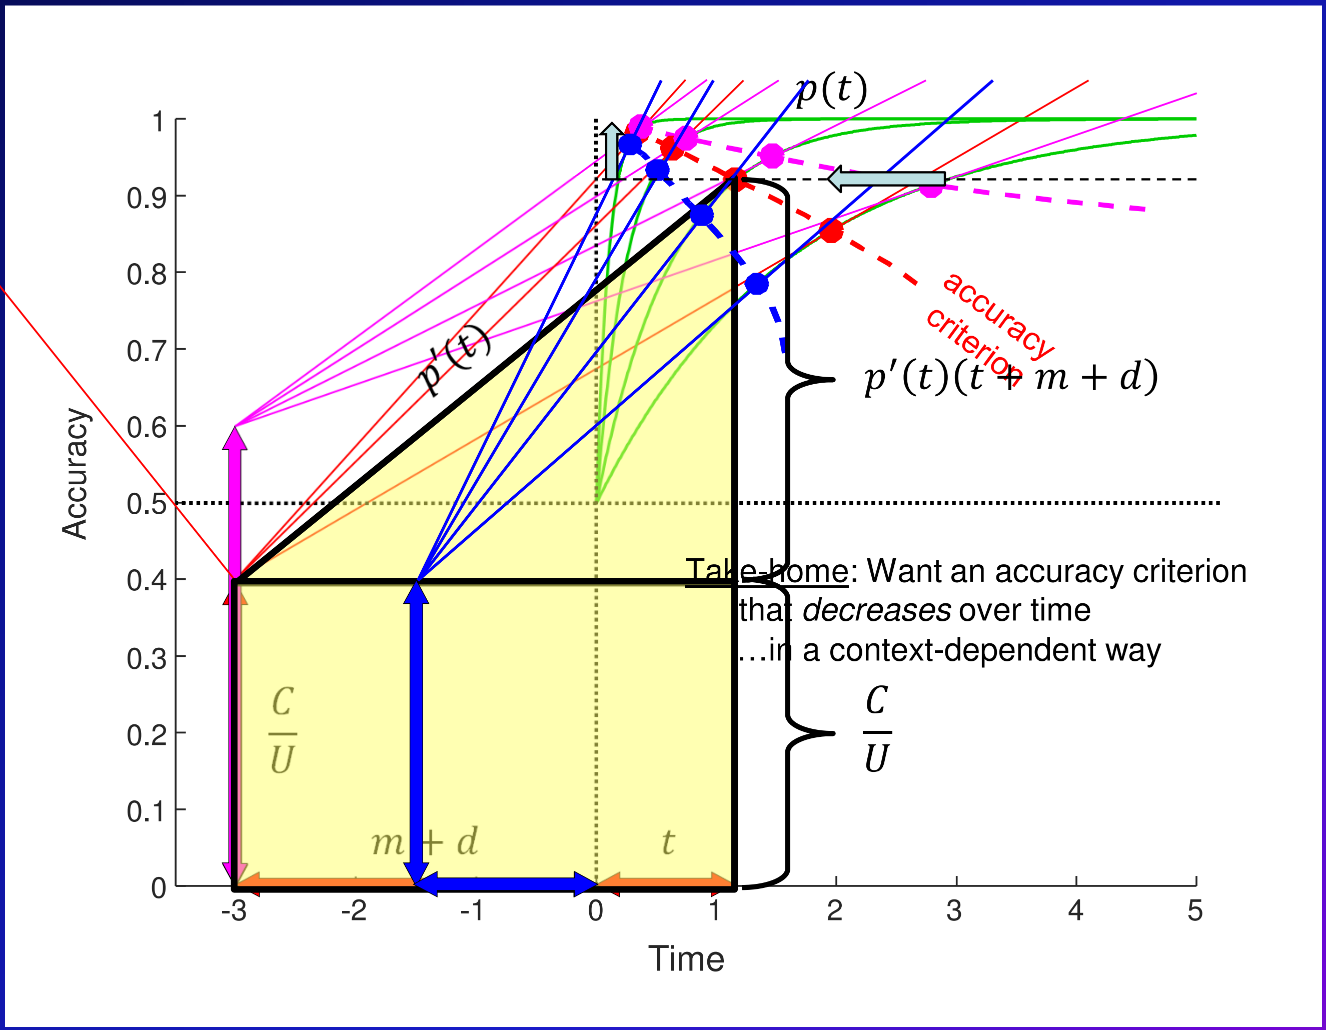

A dropping accuracy criterion

- The accuracy criterion drops, so you always reach it eventually
- But neural data suggests the threshold is constant
- Pr… thr… …terion
- Pr… …ditional sig… …me
- *"U…*
  - … …neural …old

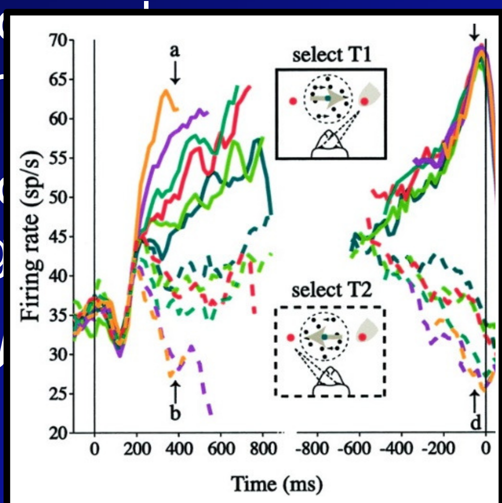

(Ditterich, 2006; Churchland et al. 2008)

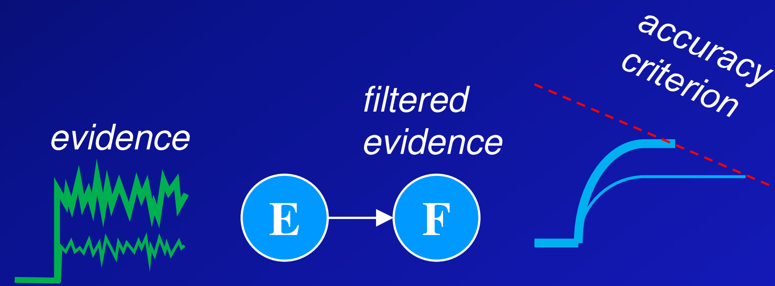

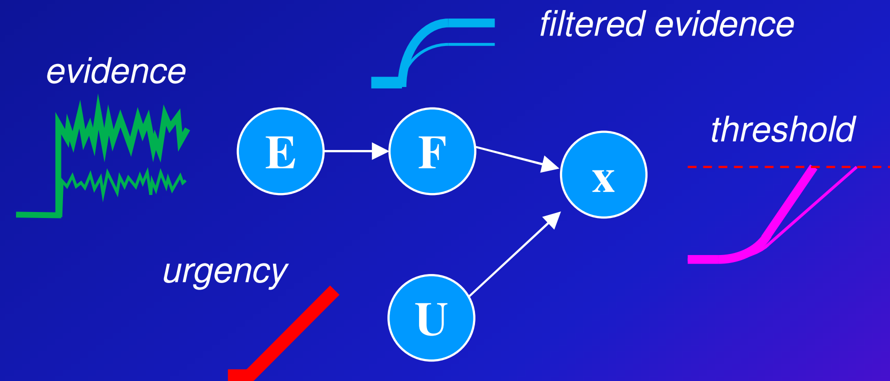

$$x_A(t) = (\underbrace{Z_A}_{\text{prior}} + \underbrace{\boldsymbol{F}[E_A(t)]}_{\text{filtered evidence}}) \cdot \underbrace{u(t)}_{\text{urgency}} < \underbrace{T}_{\text{threshold}}$$

Proposal: Urgency-Gating Model (UGM)

(Cisek et al. 2009; Thura et al. 2012)

Q1: What if the world changes?

- During natural behavior, the world is always changing
- Any integrator will be sluggish in its response to such changes
- Need to be able to quickly respond to changes
  - Reset?
  - Sudden increase in gain?
  - Some other mechanism?

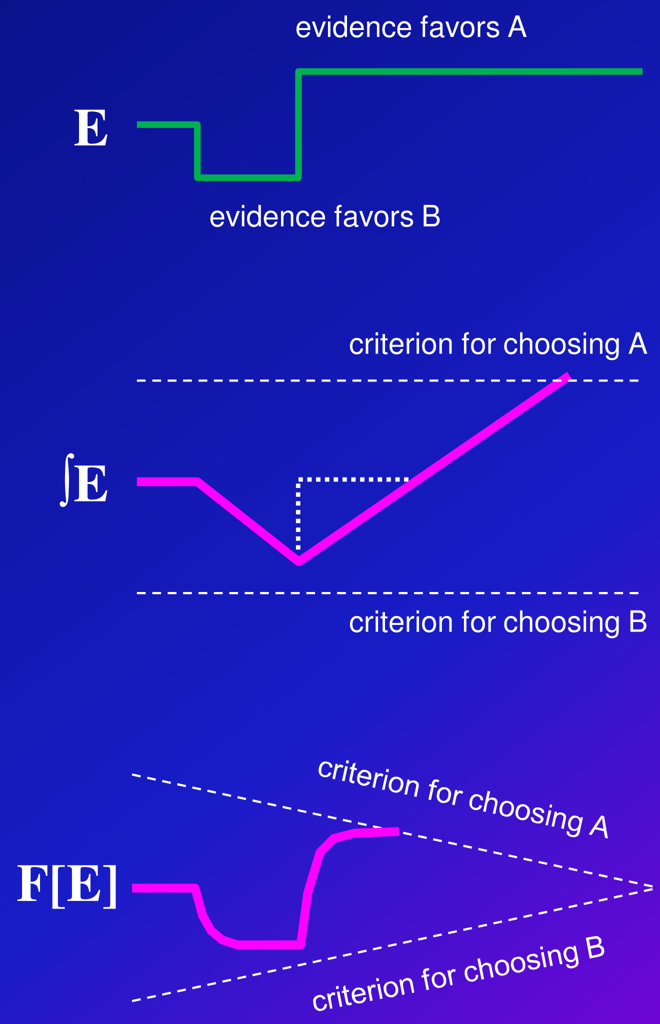

A low-pass filter has a short time constant – it responds to changes quickly

Two models

- Drift-diffusion model (DDM)
  - Integrate all samples
  - Stop at a constant criterion
- Urgency-gating model (UGM)
  - Integrate novel information (e.g. low pass filter)
  - Urgency pushes activity to threshold, implementing a dropping criterion
- Similar but different
  - Both “rise to threshold”
  - Different mechanisms responsible for build-up

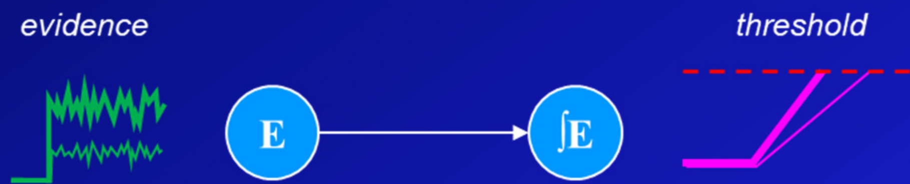

$$x_A(t) = Z_A + \int_0^t E_A(\tau)d\tau$$

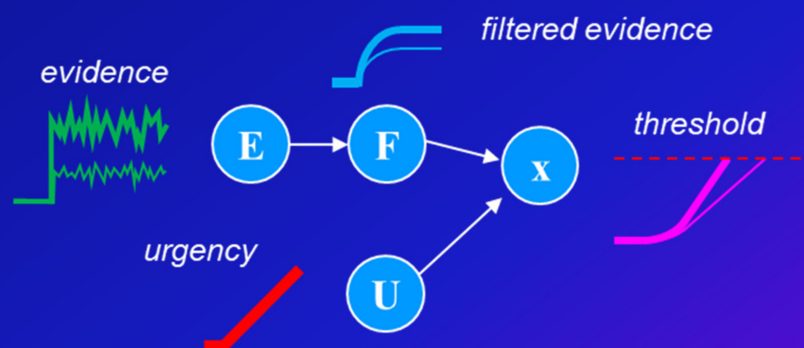

$$x_A(t) = (Z_A + \boldsymbol{F}[E_A(t)]) \cdot u(t)$$

A space of models

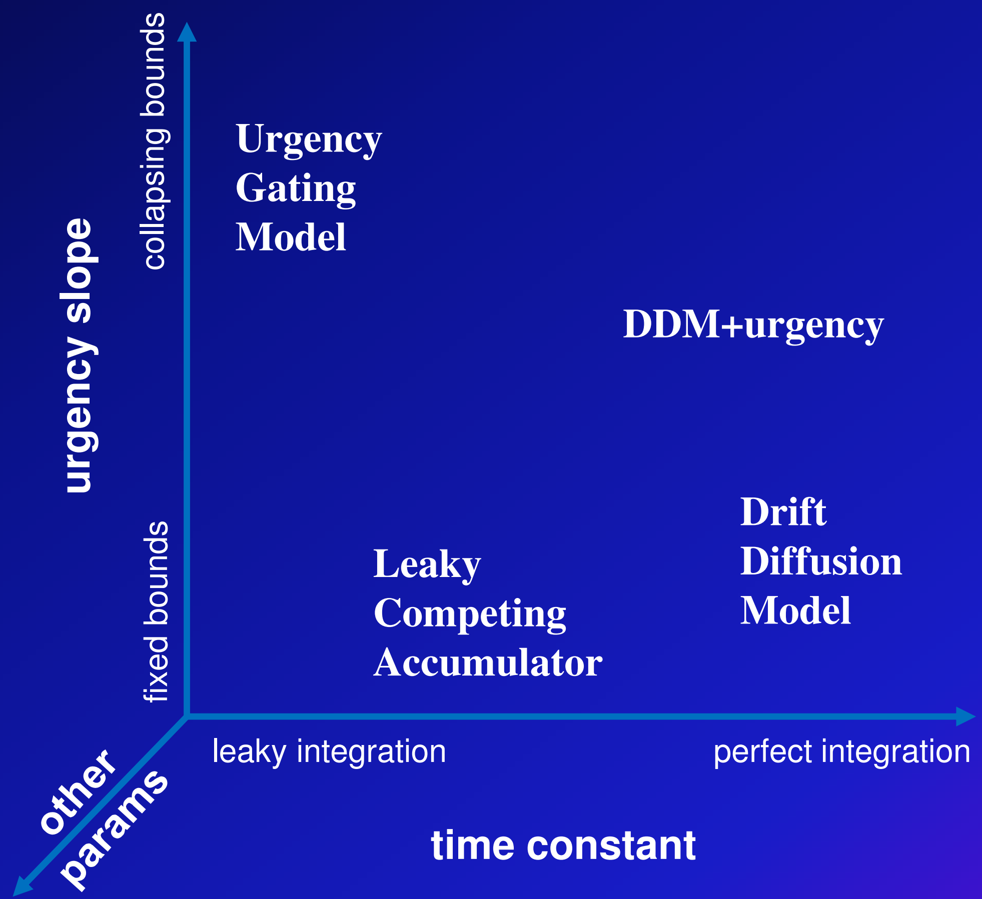

But what about all that data?

- Nearly all experiments have used constant-evidence $E_A(t) = E_A$
- In those conditions
  - Diffusion model:

$$x_A(t) = \int_0^t E_A(\tau)d\tau$$

$$x_A(t) = \int_0^t E_A d\tau = E_A \int_0^t d\tau$$

$$x_A(t) = E_A t$$

  - Urgency-gating:

$$x_A(t) = \mathrm{F}[E_A(t)] \cdot u(t)$$

$$x_A(t) = E_A \cdot u(t)$$

$$x_A(t) = E_A t$$

  - Both make similar predictions at the behavioral and neural level
  - To distinguish the models, need tasks with ***changing evidence***

A space of models

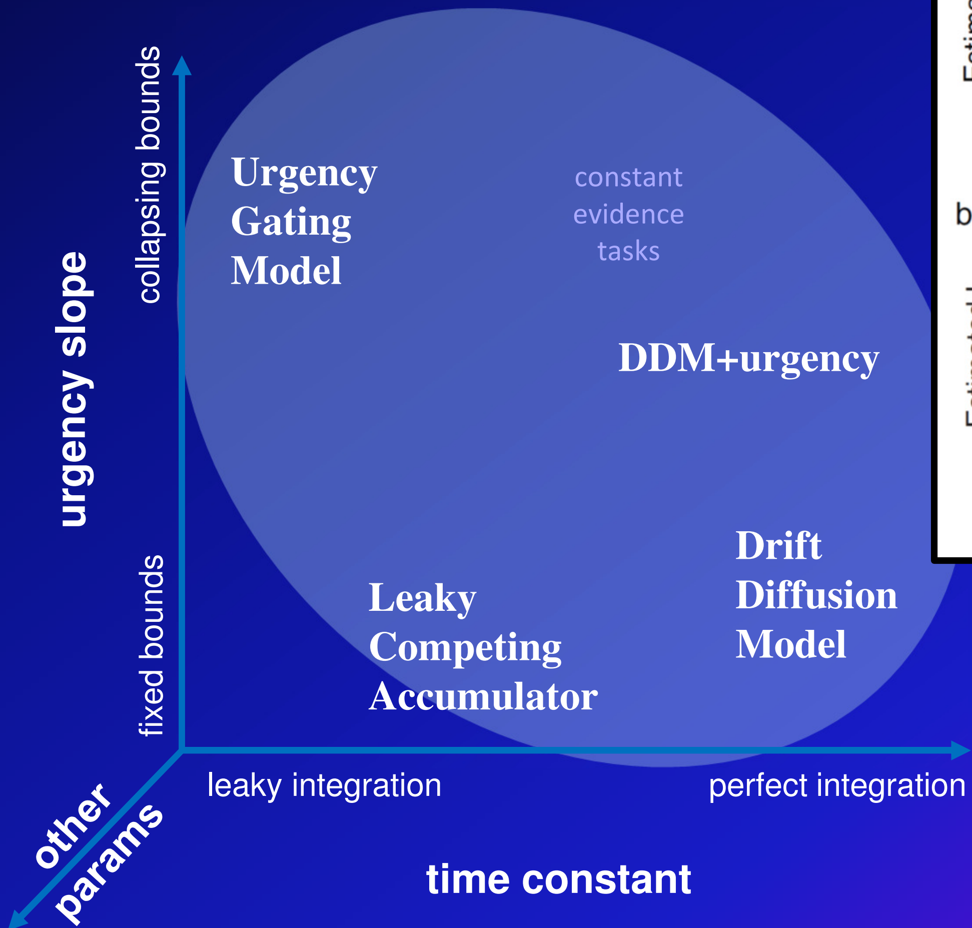

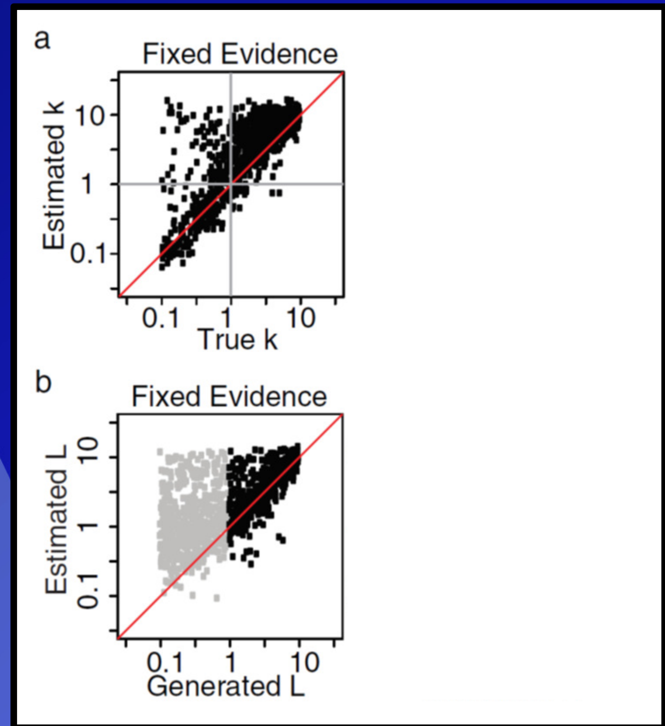

Trueblood et al. 2020

Noisy motion with pulses

Carland, Marcos, Thura, & Cisek (2016) *Journal of Neurophysiology*


Matt Carland


Encarni Marcos

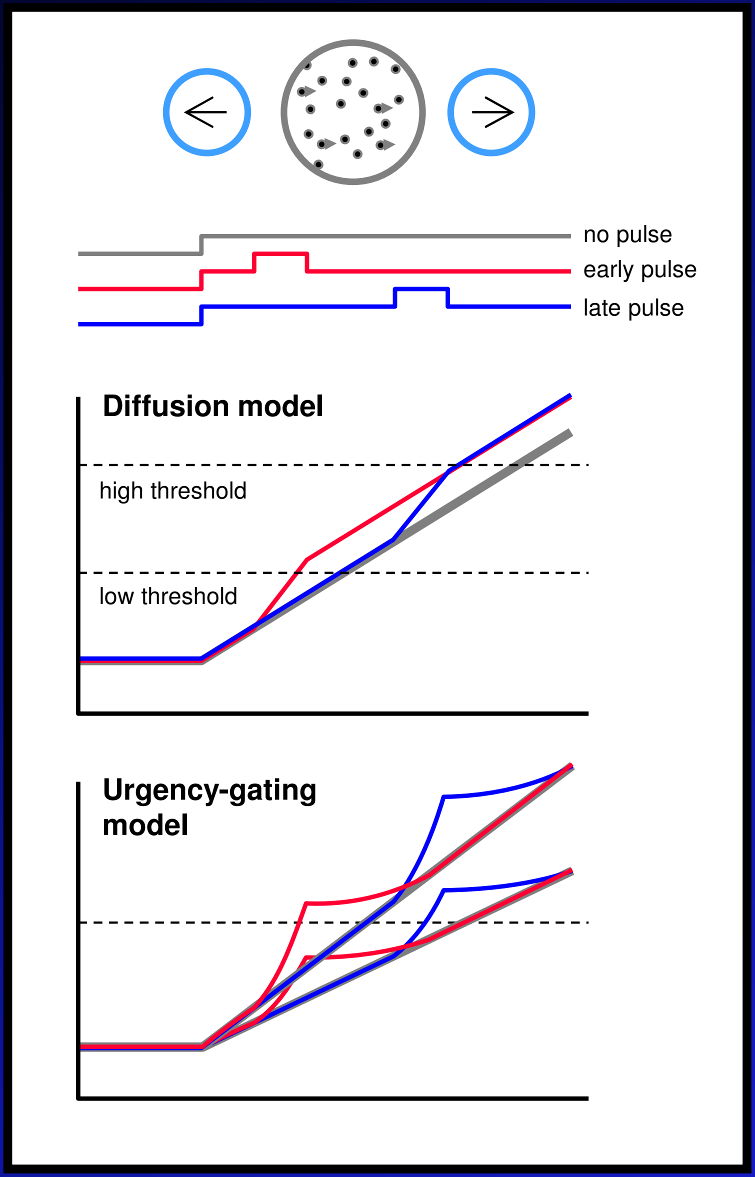

- Random dot motion discrimination task
  - Reaction time version
  - Low stimulus coherence (3%)
  - We add brief pulses of extra motion at different times during the trial
- Diffusion model predictions
  - Fast (low threshold): early pulses have an effect but late pulses are too late
  - Slow (high threshold): all pulses have an effect
- Urgency-gating model predictions
  - Fast (high urgency): early pulses have an effect but late pulses are too late
  - Slow (low urgency): late pulses have an effect but early pulses “leak out”
- Test subjects in fast and slow conditions

How to get subjects to slow down?

- Instruction is always the same:
  - "Indicate motion as soon as you see it"
- Two kinds of trials:
  - Constant Motion Discrimination (CMD) trials, with and without pulses
  - Variable Motion Discrimination (VMD) trials, with changing evidence (3% coherence step every 200ms)
  - In VMD trials, motion is sometimes stronger if you wait

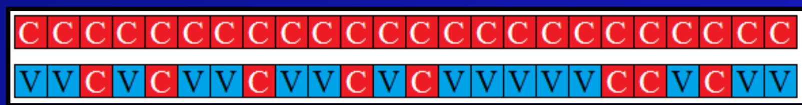

- Design:
  - **Blocked**: Just CMD trials
  - **Interleaved**: Identical CMD trials are interleaved with VMD trials
- <u>Question</u>: Do subjects slow down in CMD trials when they're interleaved with VMD?

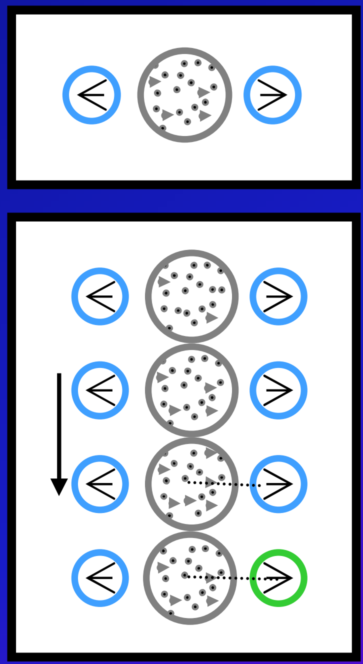

Do subjects slow down?

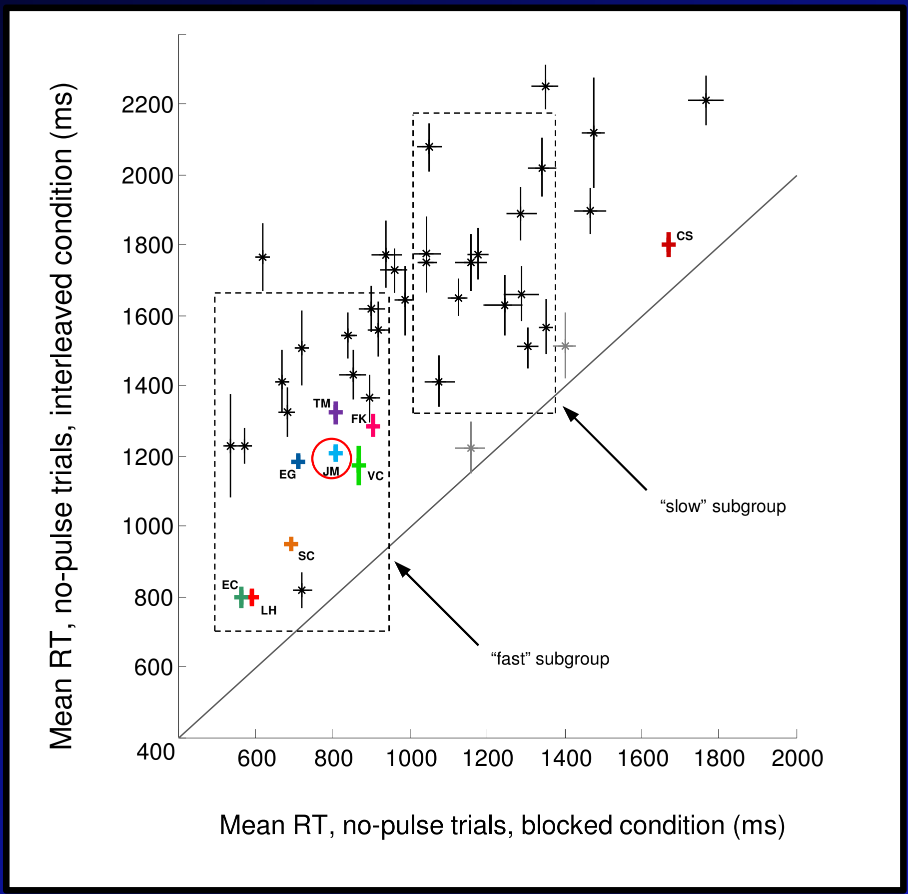

- Yes, 42/44 subjects are significantly slower on no-pulse trials during the interleaved sessions

- <u>Question</u>: Does the effect of pulses in CMD trials change between conditions?

Results: Subject "JM"

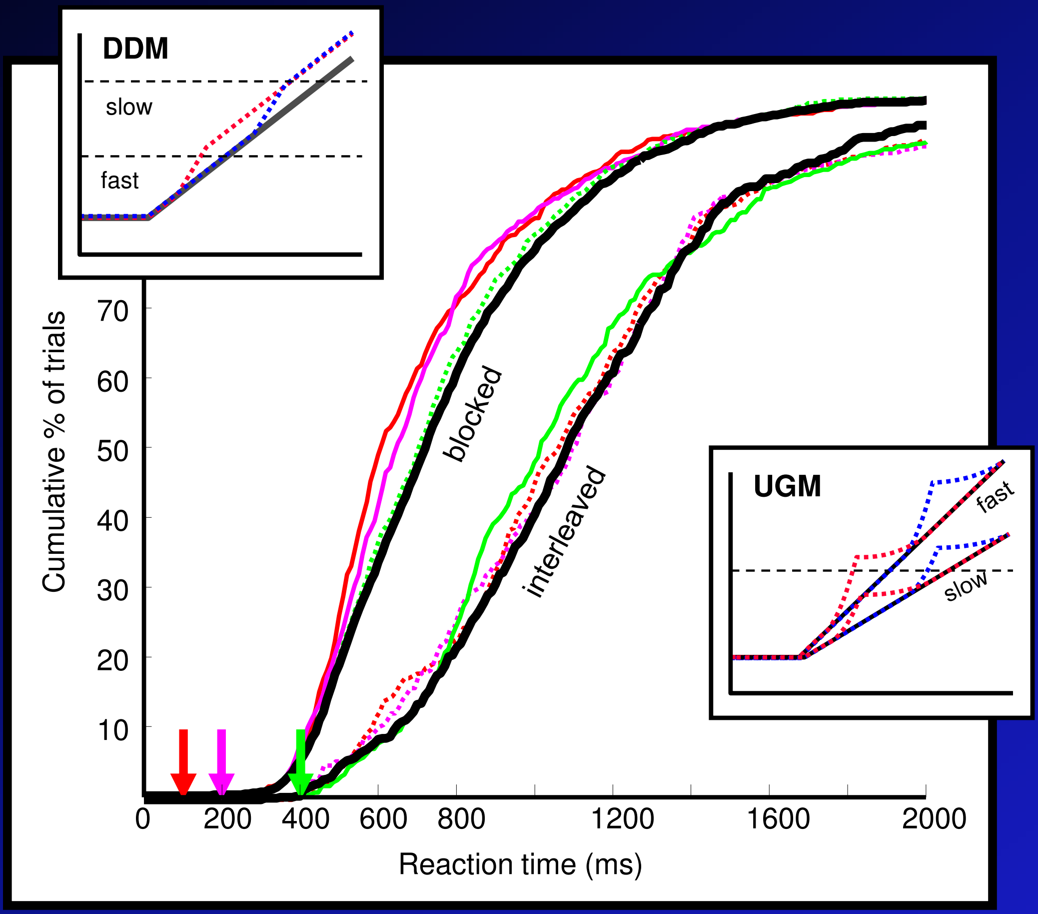

- Blocked (fast) condition
  - Pulses at 100 and 200ms have a significant effect
  - Pulse at 400ms does not
  - Consistent with both models
  - Consistent with previous studies of motion pulses (Huk & Shadlen 2005; Kiani et al. 2008)
- Interleaved (slow) condition
  - Pulse at 400ms has a significant effect
  - Pulses at 100 and 200ms ***do not*** have an effect
  - Suggests strong leak Time constant of 100-250ms

Results: Fast subgroup

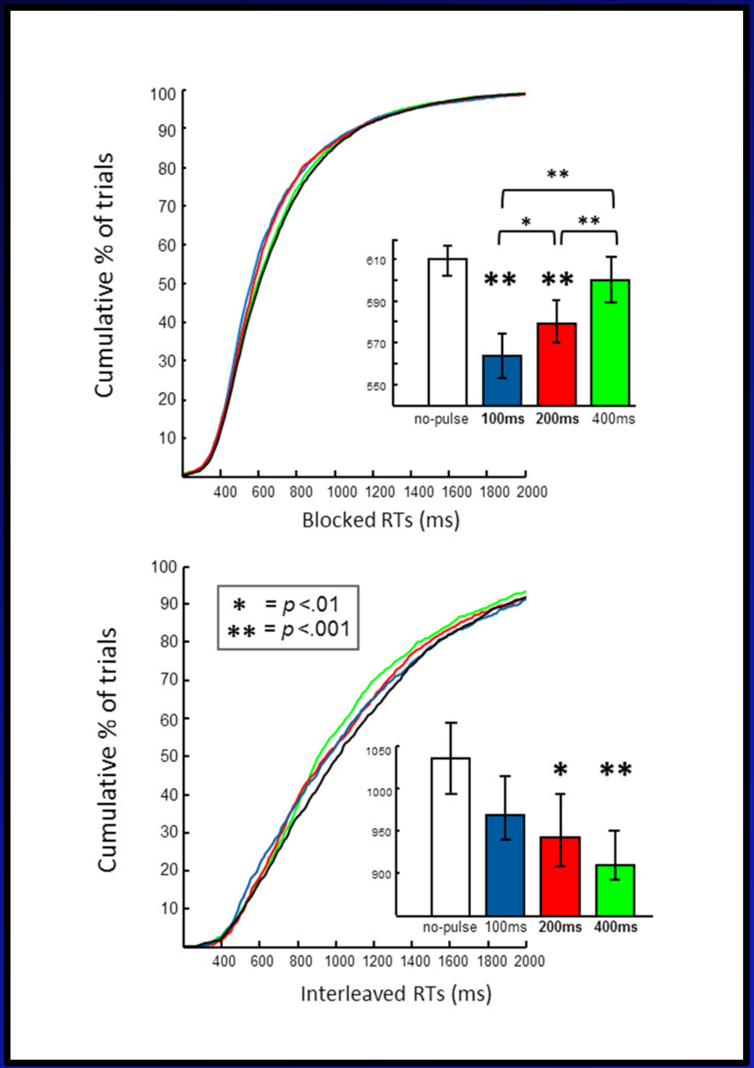

- Blocked (fast) condition
  - Pulses at 100 and 200ms have a significant effect
  - Pulse at 400ms does not
- Interleaved (slow) condition
  - Pulses at 200ms and 400ms have a significant effect
  - Pulse at 100 ***does not*** have a significant effect

Results: Slow subgroup

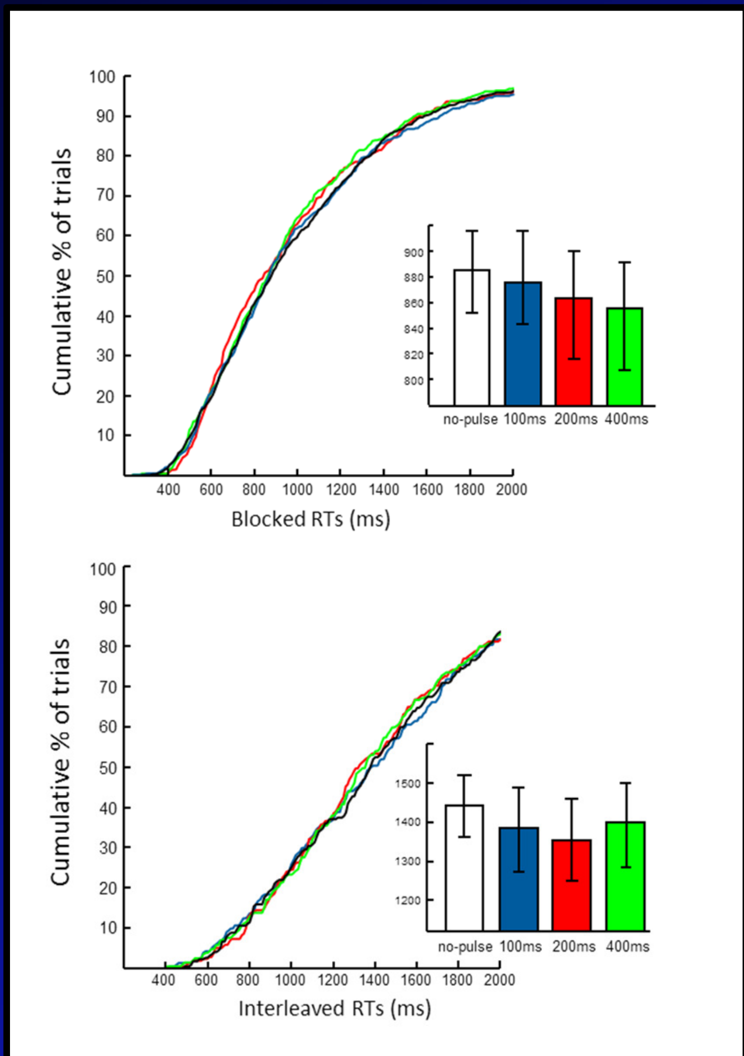

- Blocked (fast) condition
  - No effect of pulses
- Interleaved (slow) condition
  - No effect of pulses

Results: Subject "JM"

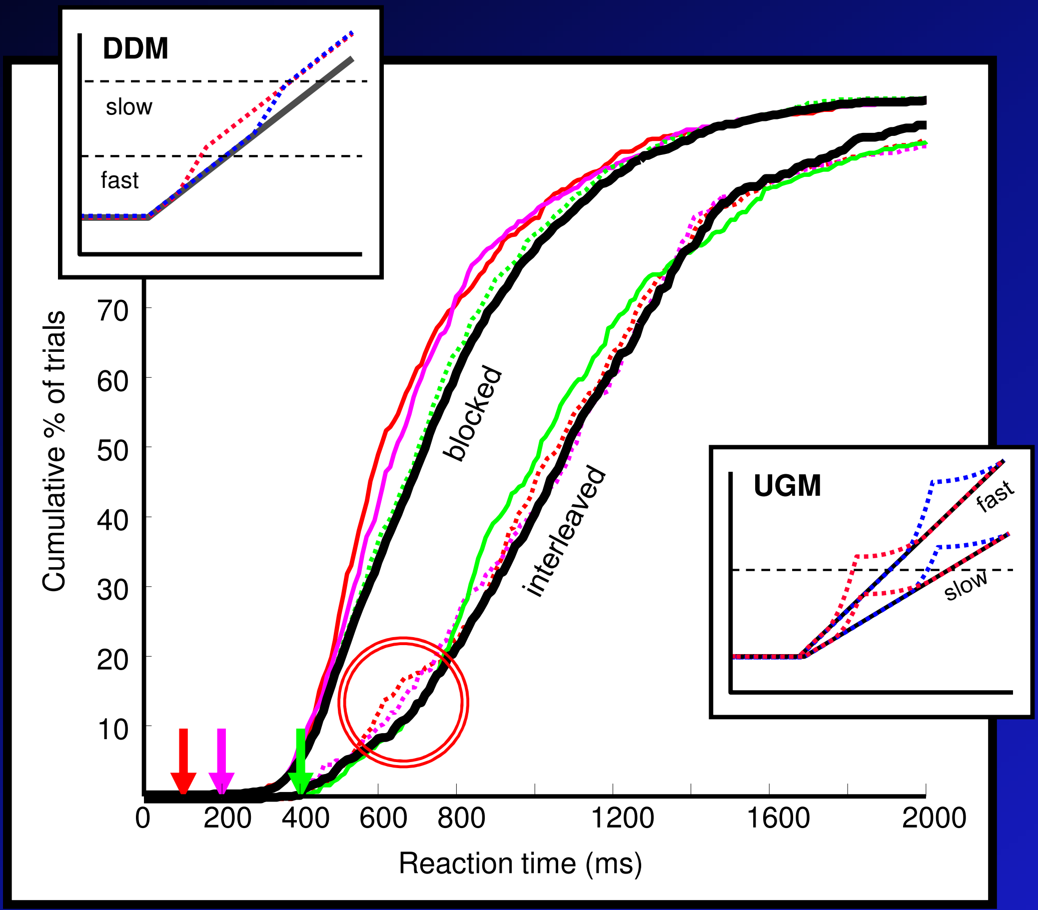

- Blocked (fast) condition
  - Pulses at 100 and 200ms have a significant effect
  - Pulse at 400ms does not
  - Consistent with both models
  - Consistent with previous studies of motion pulses (Huk & Shadlen 2005; Kiani et al. 2008)
- Interleaved (slow) condition
  - Pulse at 400ms has a significant effect
  - Pulses at 100 and 200ms ***do not*** have an effect
  - Suggests strong leak
    Time constant of 100-250ms

Individual RT-dependent trends

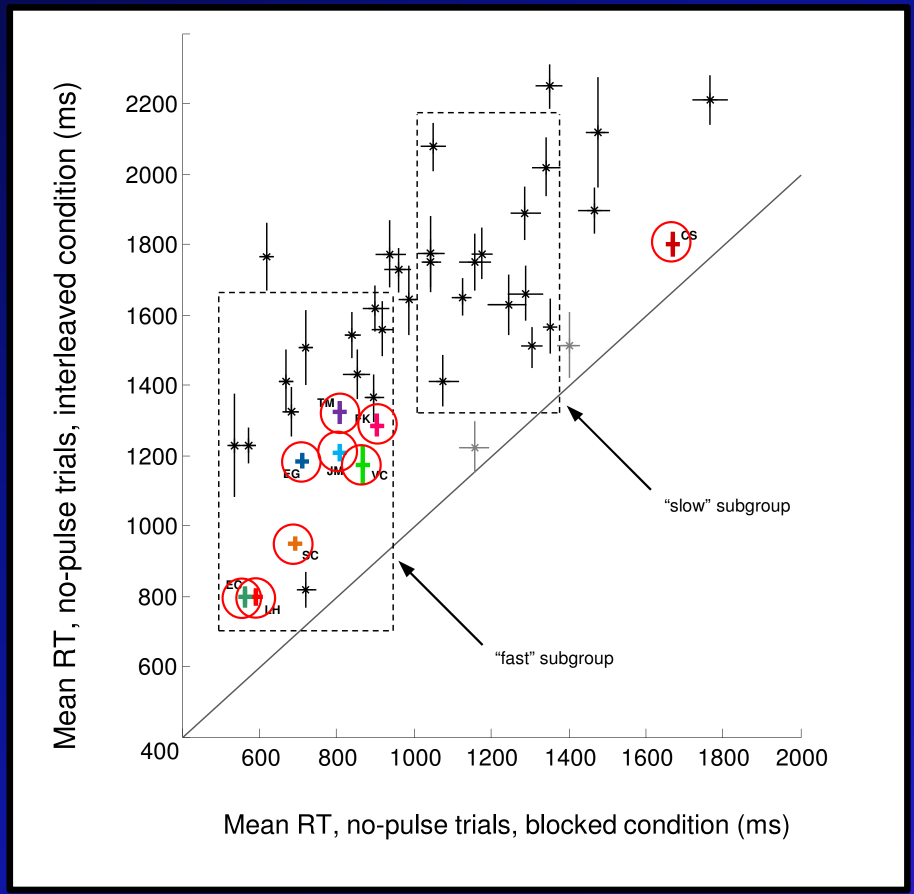

Predict: The time window in which a pulse is effective should depend on an individual's reaction time

Individual RT-dependent trends

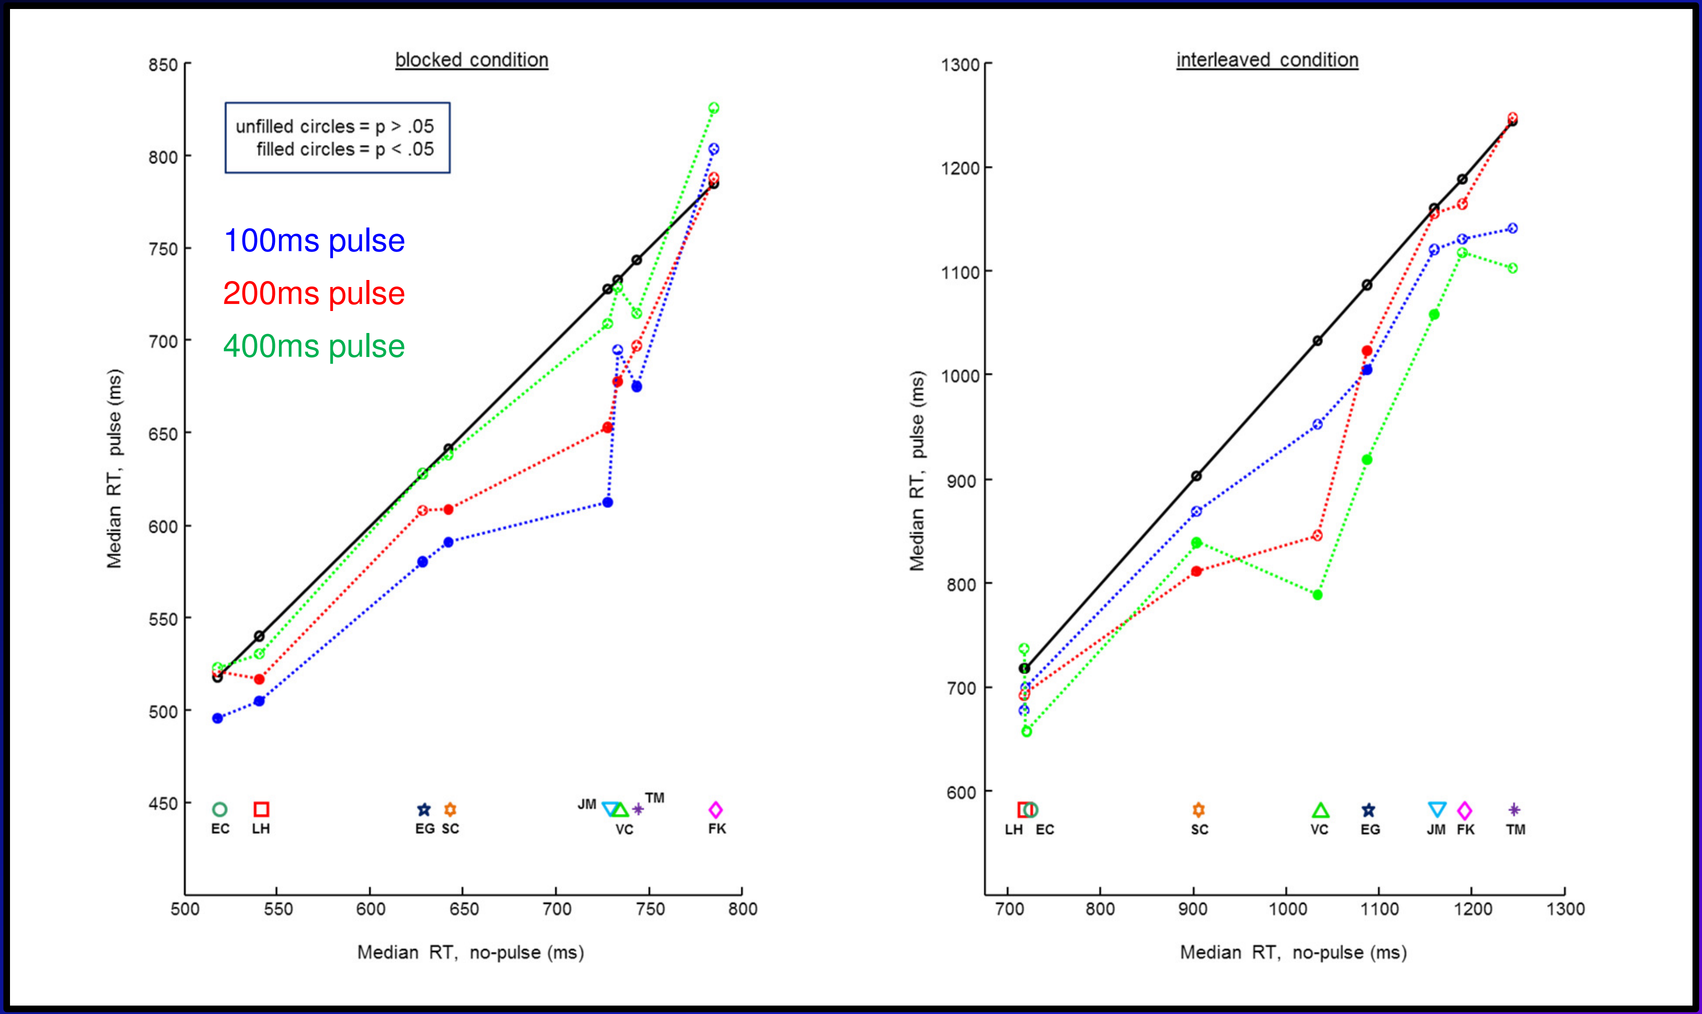

Modeling with UGM & DDM

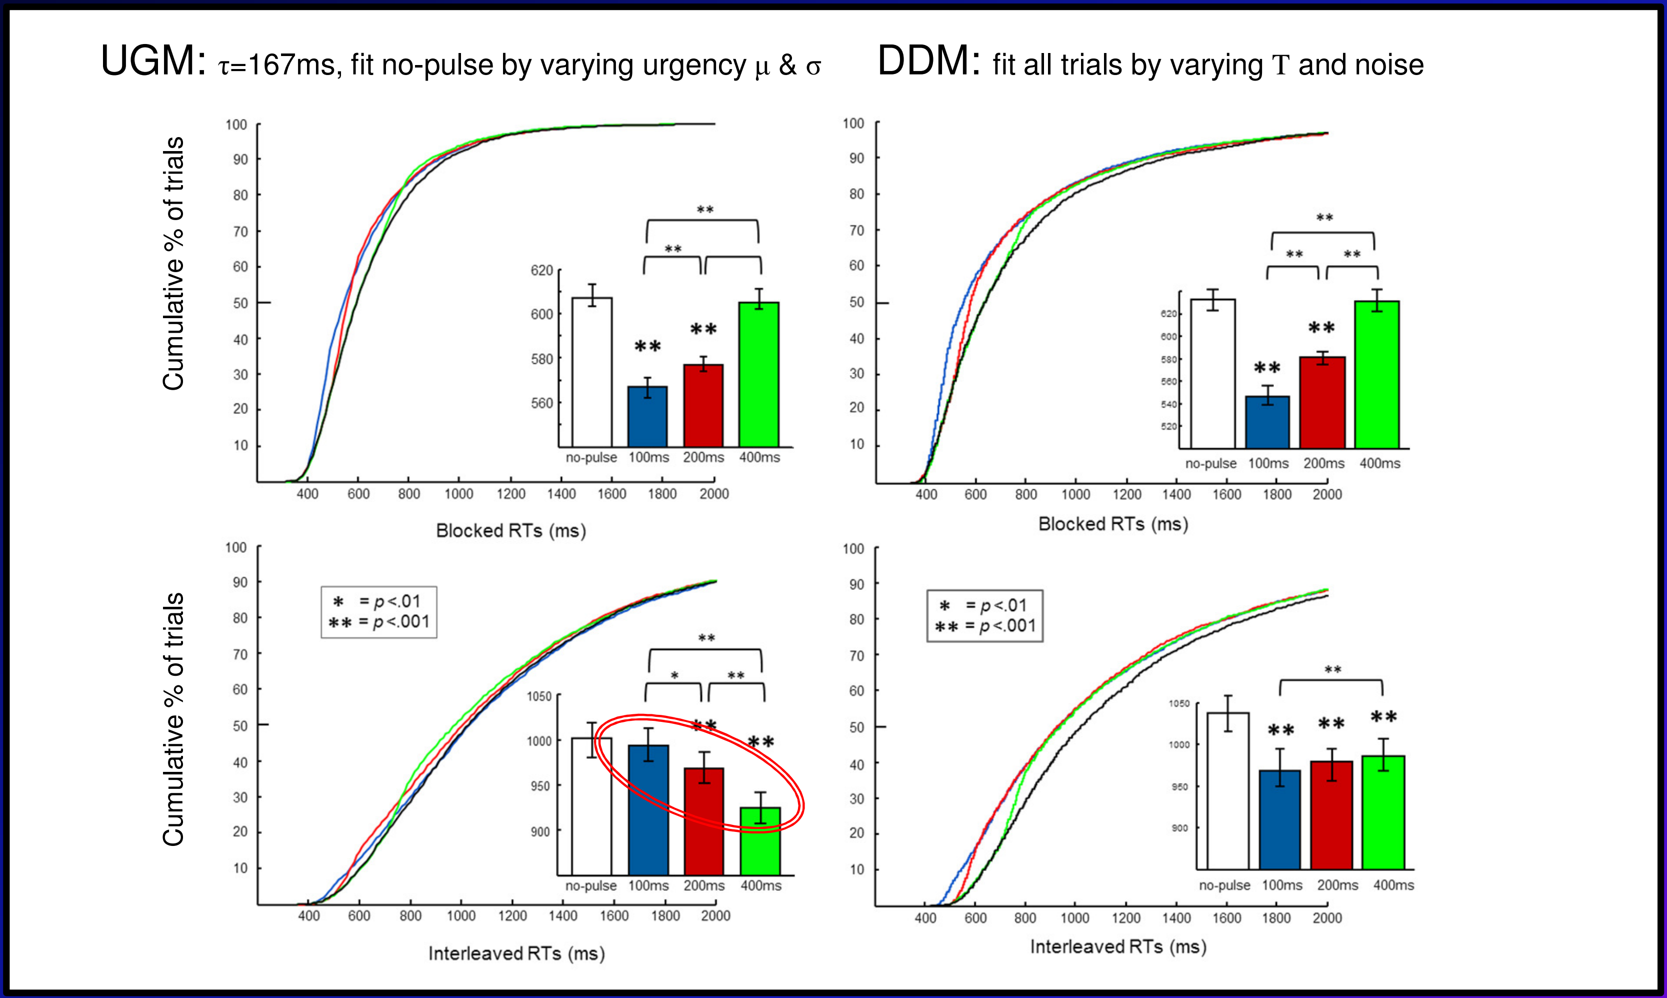

Claim: This result cannot be explained with the DDM, regardless of parameter settings

Task-dependence?

- "Two-model solution"
  - Use DDM when doing a "constant evidence" task
  - Use UGM when doing a "changing evidence" task
- Lacks parsimony
  - Use model X to explain data set A, and model Y for data set B…
  - …but model Y already explains *both* data sets A and B
- Goodness of fit?
  - Hawkins et al. 2015, 2016 suggest DDM fits better than UGM in constant-evidence tasks (but don't test changing-evidence tasks)
  - A two-model solution includes parameters of both models, plus a switching mechanism
  - The fit to *any* given data set must be penalized for all of these parameters, including those of the "unused model"

A space of models

Cisek, Puskas, El-Murr (2009)
Thura, Beauregard-Racine, Fradet, Cisek (2012)
Gluth, Reiskamp, Buchtel (2012)
Thura & Cisek (2014); Thura, Cos, Trung, Cisek (2014)
Winkel et al. (2014) → Carland, Thura, Cisek (2015)
Murphy, Boonstra, Niewenhuis (2016)
Carland, Marcos, Thura, Cisek (2016)
Thura & Cisek (2016); Thura & Cisek (2017)
Evans, Hawkins, Boehm, Wagenmakers, Brown (2017)
Malhotra, Leslie, Ludwig, Bogacz (2017)
Palestro, Weichart, Sederberg, Turner (2018)
Evans, Trueblood, Holmes (2020)
Yau, Dadar, Taylor, Zeighami, Fellows, Cisek, Dagher (2020)
Trueblood, Heathcote, Evans, Holmes (2020)
…

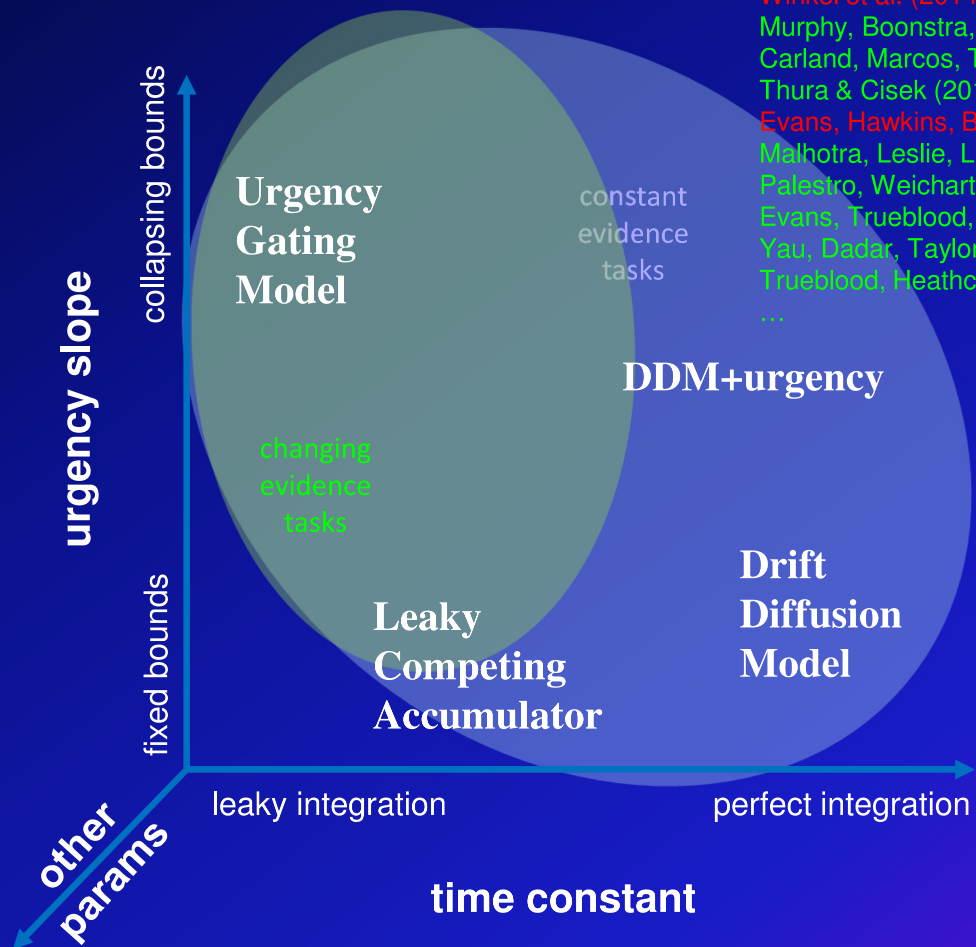

Conclusion, so far

- The "urgency-gating model" offers a better explanation of decision-making than the "drift-diffusion model"
  - Theoretically (considers redundancy, maximizes reward rate)
  - Empirically (fits a larger class of experiments)

  (some confessions)

- Where do we go from here?
  - Neural mechanisms
  - Broader phenomena

Neural mechanisms

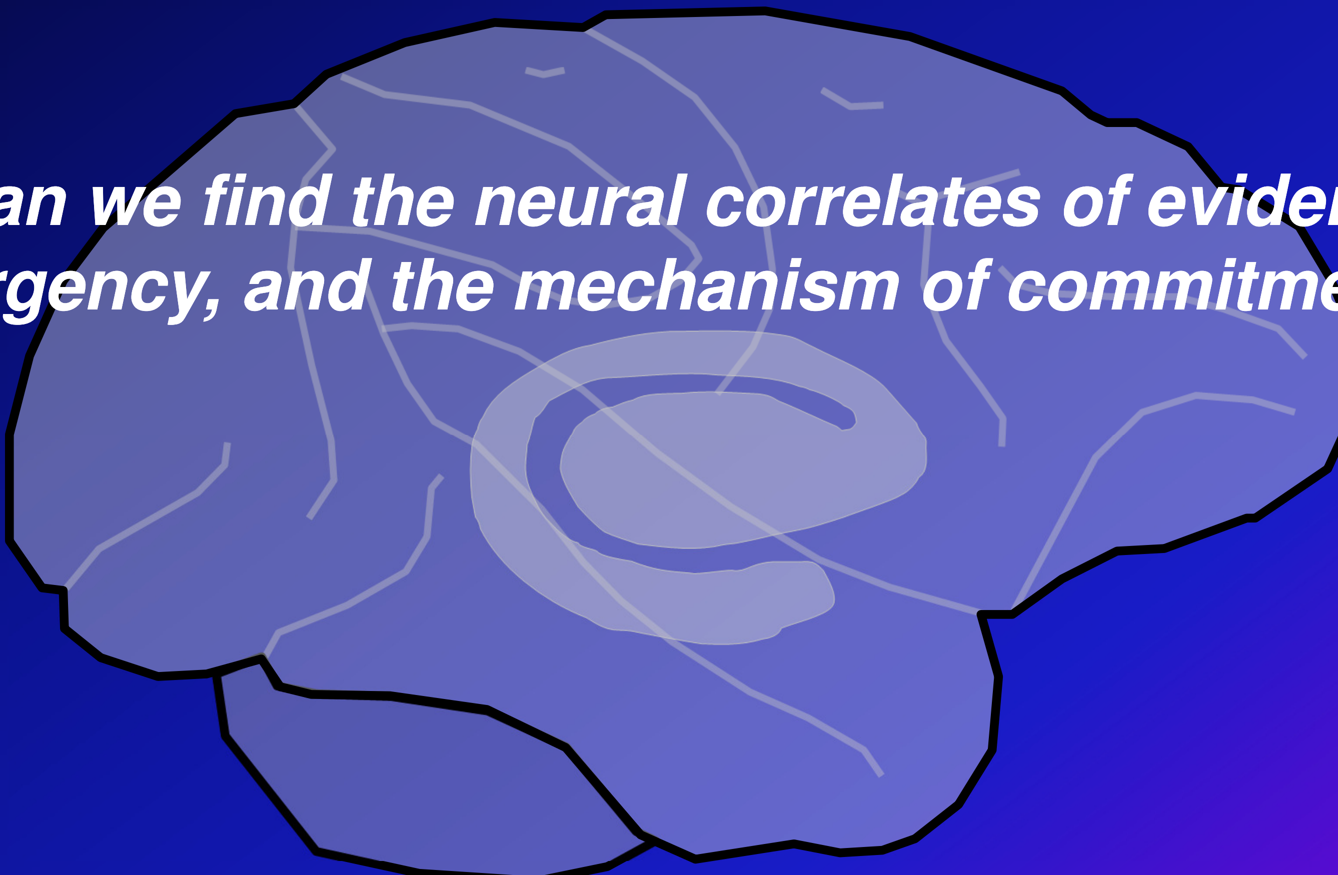

Can we find the neural correlates of evidence, urgency, and the mechanism of commitment?

The "tokens task"

David Thura

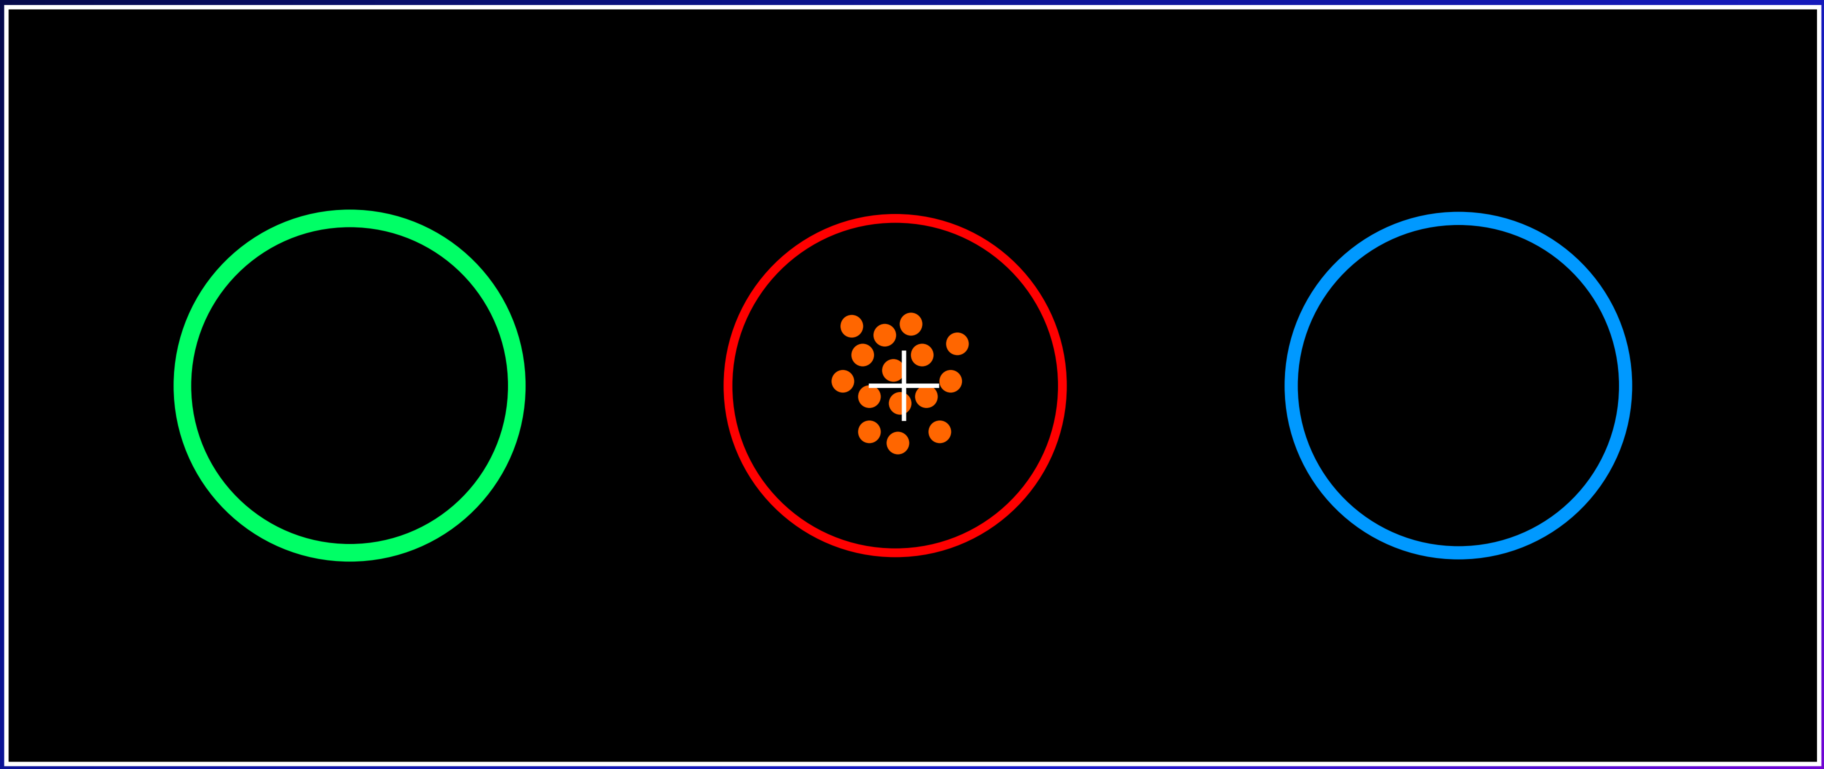

The "tokens task"

David Thura

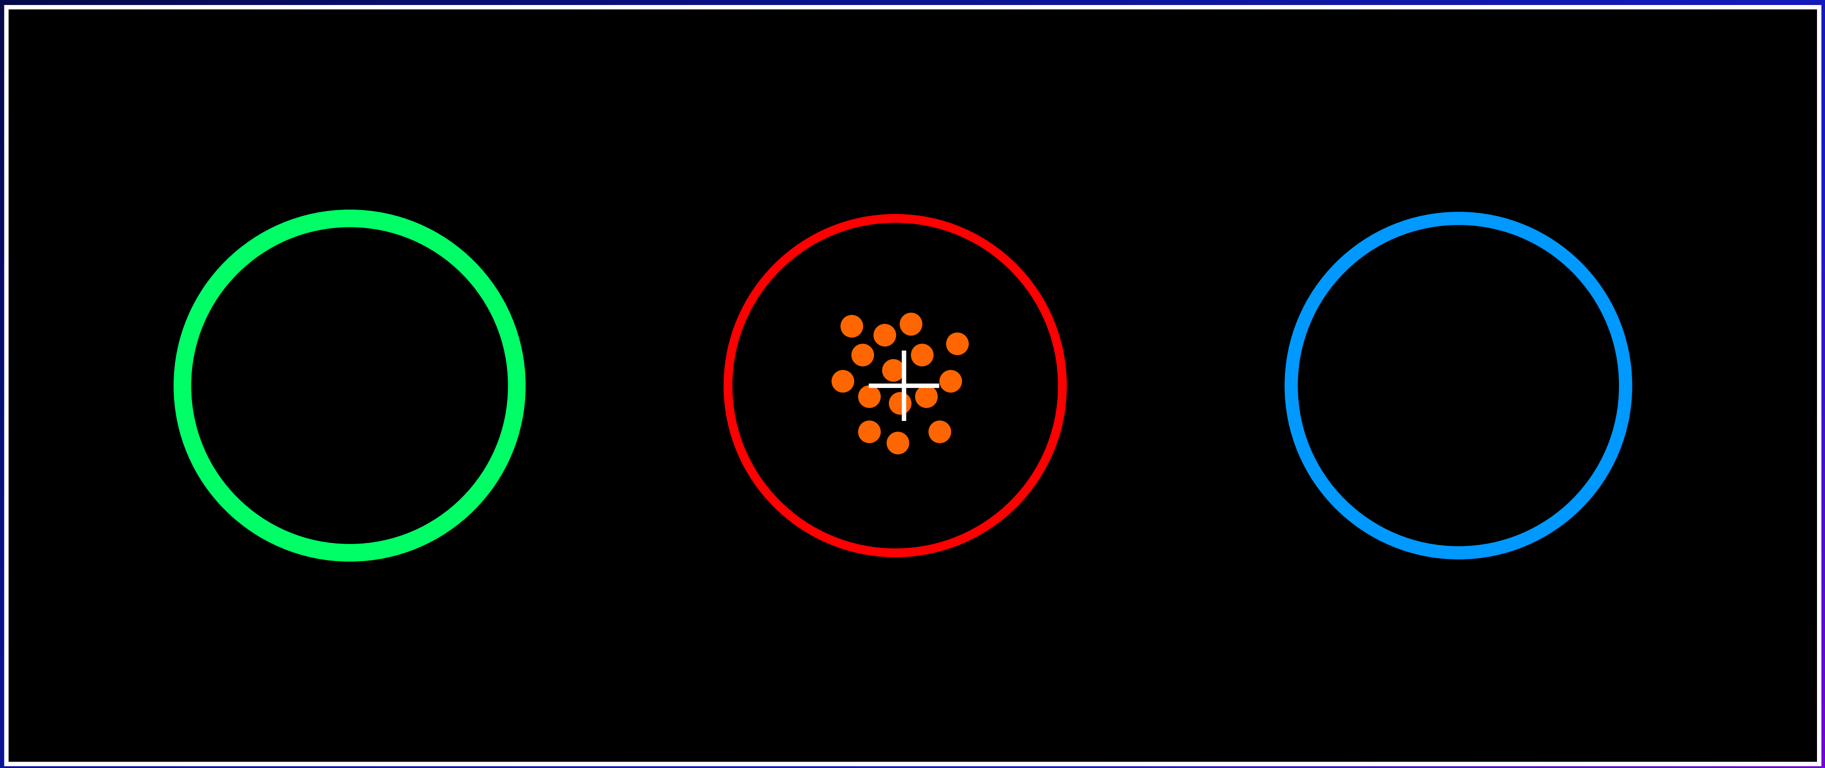

David Thura

The "tokens task"

$$p(R \mid N_R, N_L, N_C) = \frac{N_C!}{2^{N_C}} \sum_{k=0}^{\min(N_C, 7-N_L)} \frac{1}{k!(N_C-k)!}$$

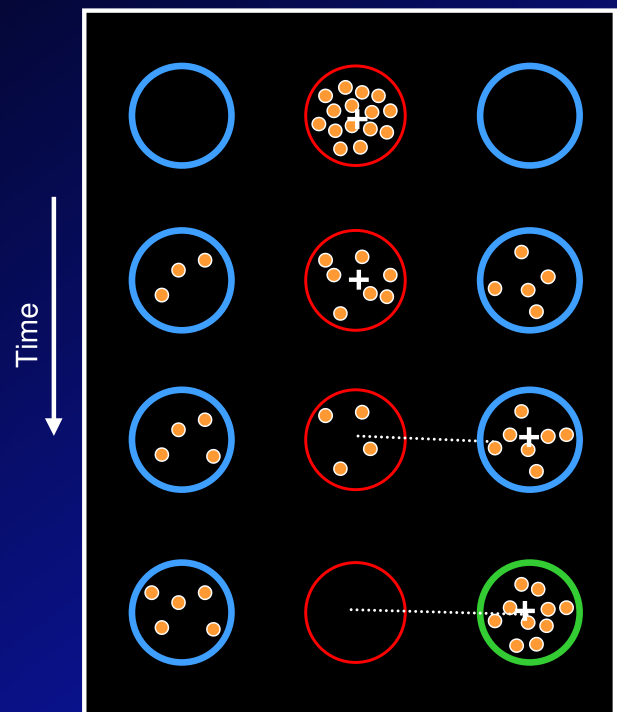

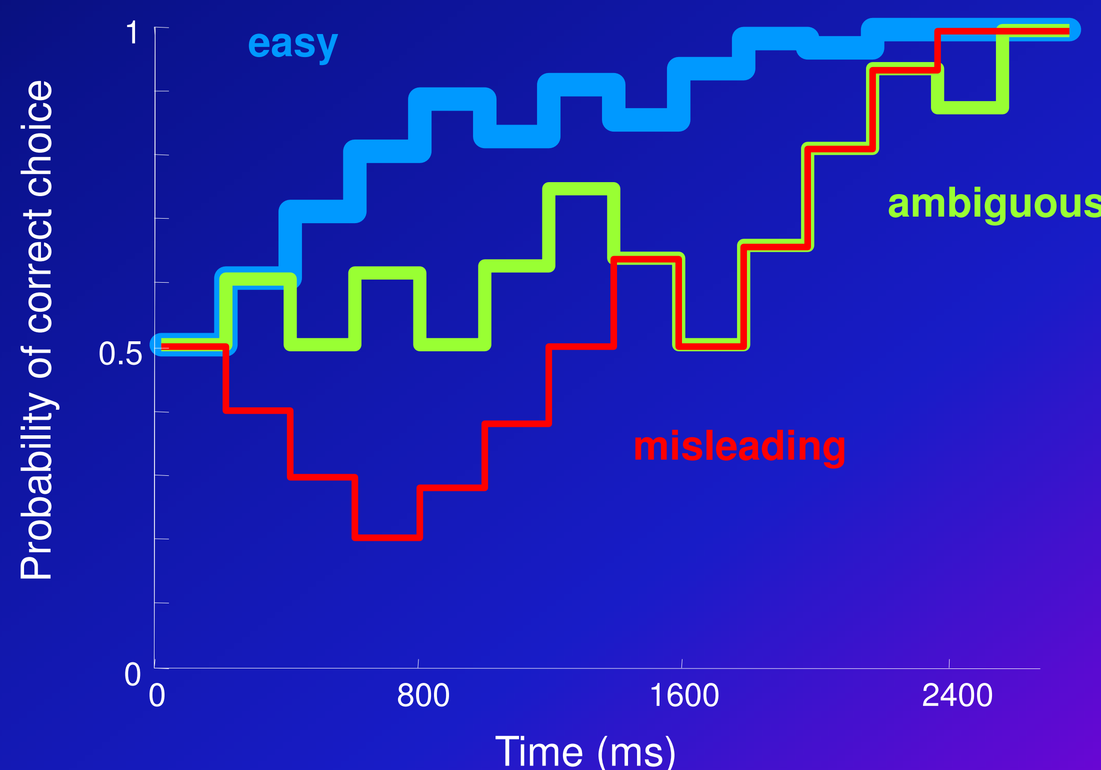

The "tokens task"

David Thura

$$p(R \mid N_R, N_L, N_C) = \frac{N_C!}{2^{N_C}} \sum_{k=0}^{\min(N_C, 7-N_L)} \frac{1}{k!(N_C-k)!}$$

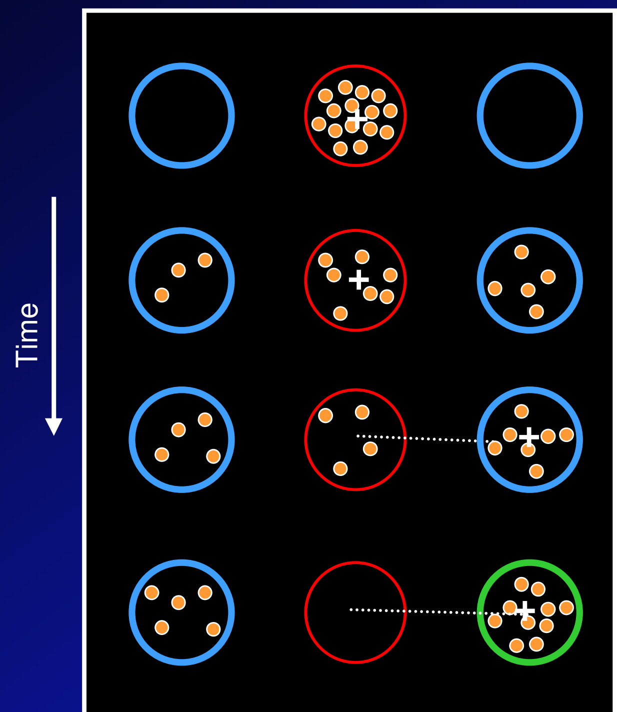

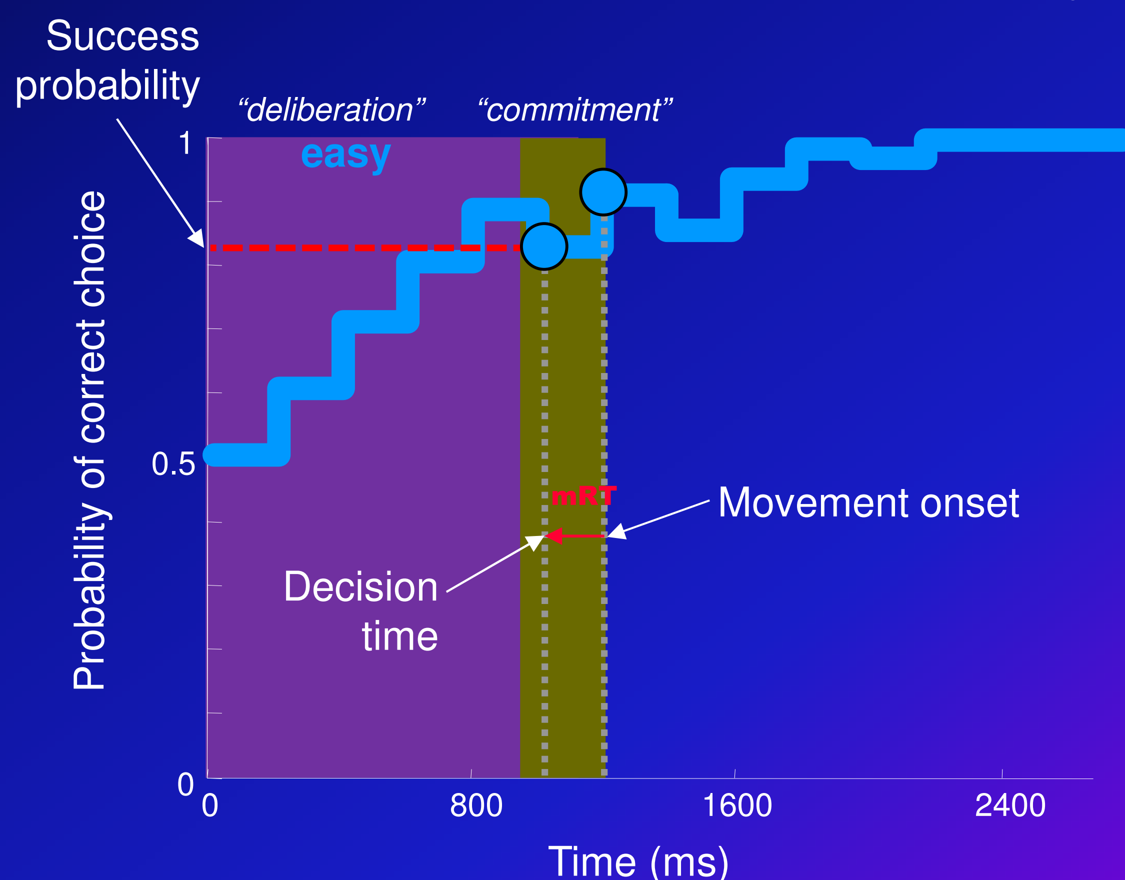

Behavioral data

David Thura

Ignasi Cos

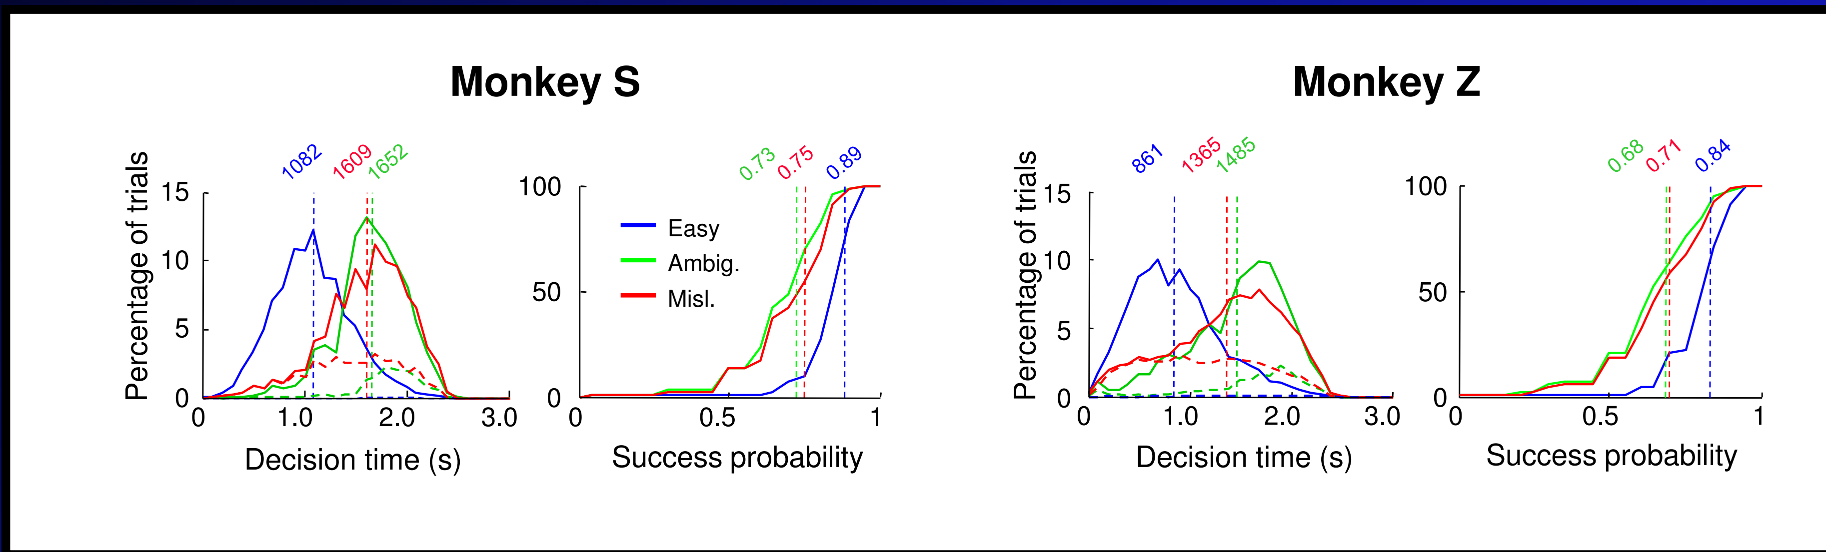

Thura, Cos, Trung & Cisek (2014) *J. Neurosci.*

- In easy trials, monkeys respond more quickly than in ambiguous or misleading trials
- In longer trials, decisions are made at a lower level of probability
  - Monkeys drop their accuracy criterion over time

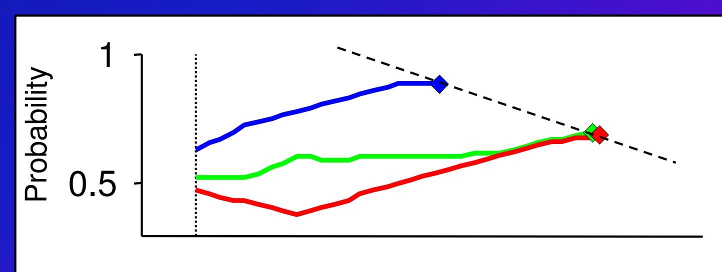

Deriving the urgency signal

David Thura

Ignasi Cos

- For each trial, calculate the evidence available at the time the monkey made his decision
- Group trials according to decision time, and calculate the mean
- Dropping accuracy criterion = urgency signal
- Model with UGM

$$x_i = \underbrace{[\text{illegible}]}_{\text{evidence}}\ \underbrace{[\text{illegible}]}_{\text{urgency}}\ \underbrace{[\text{illegible}]}_{\text{variability}} < T$$

- Find m and b that provide the best fit

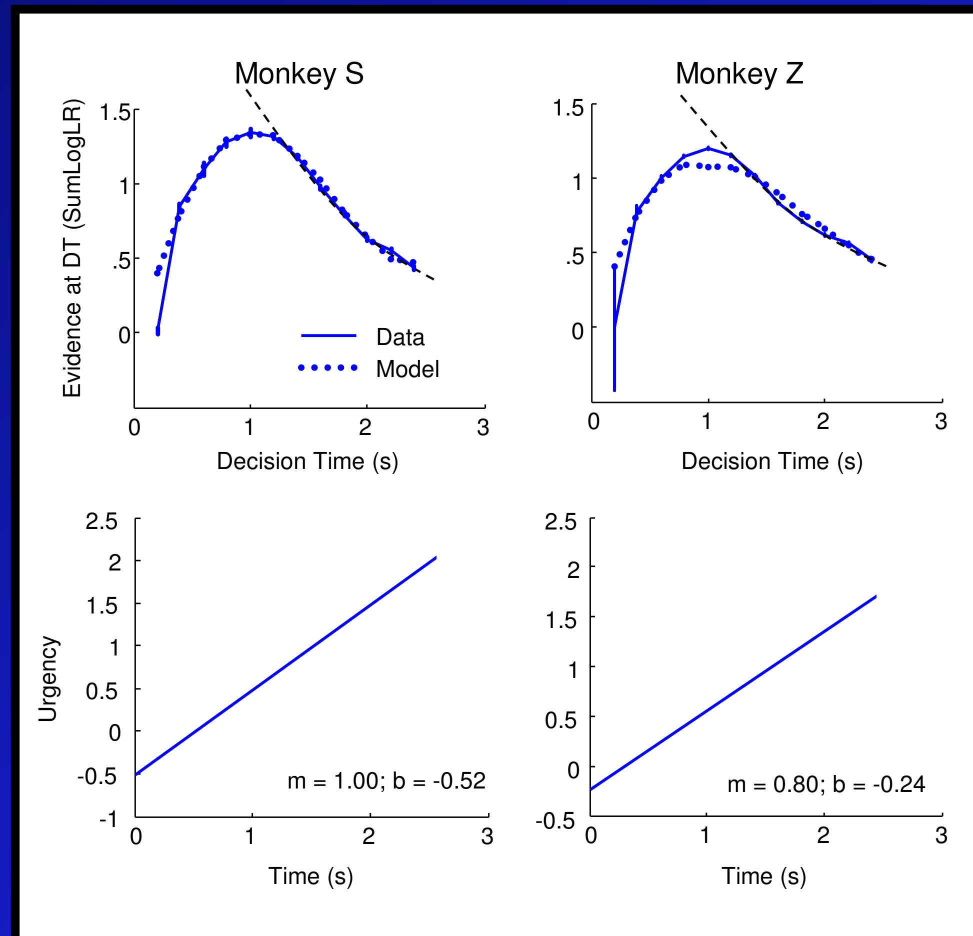

Thura, Cos, Trung & Cisek (2014) *J. Neurosci.*

The “tokens task”

David Thura

$$p(R \mid N_R, N_L, N_C) = \frac{N_C!}{2^{N_C}} \sum_{k=0}^{\min(N_C, 7-N_L)} \frac{1}{k!(N_C-k)!}$$

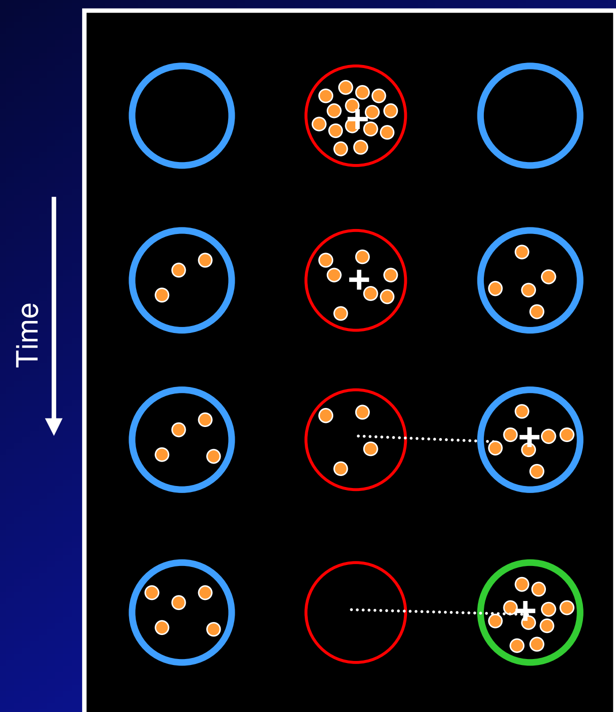

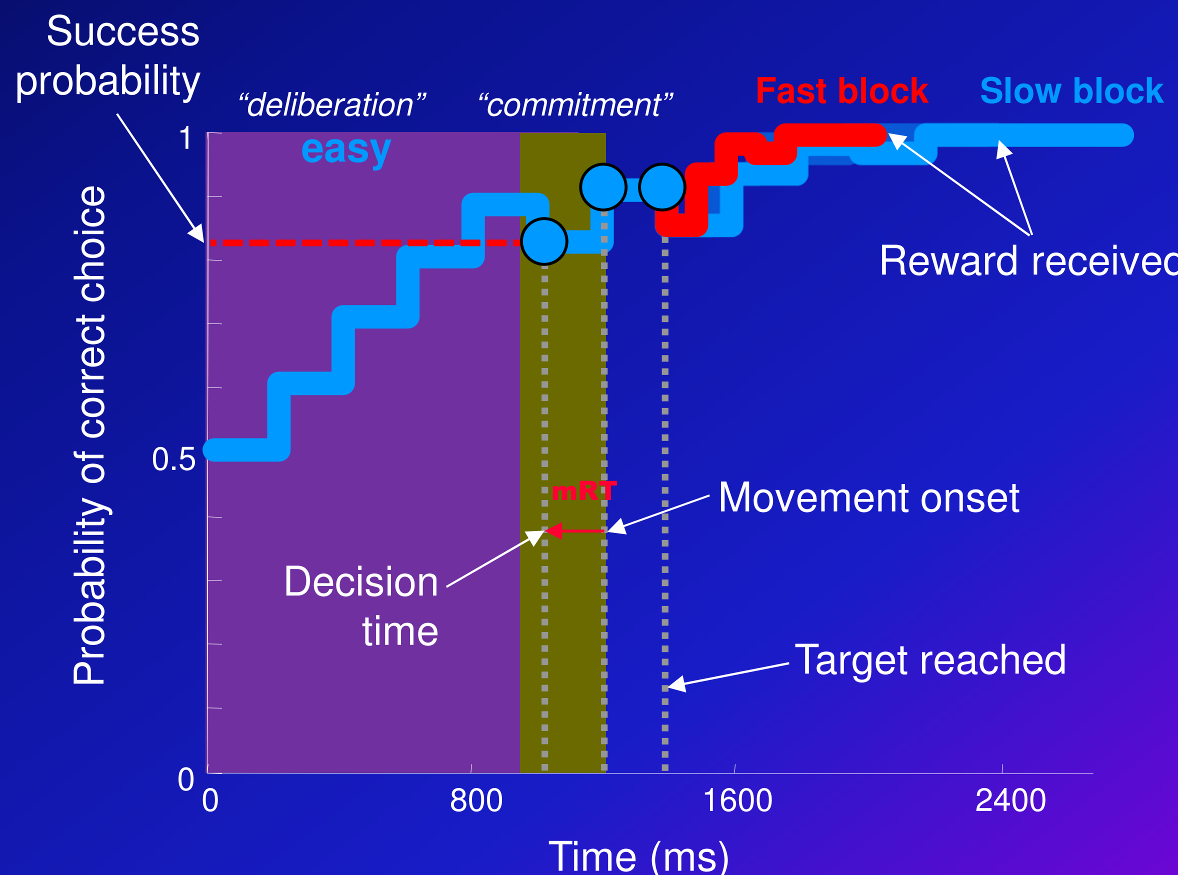

Do monkeys adjust their behavior?

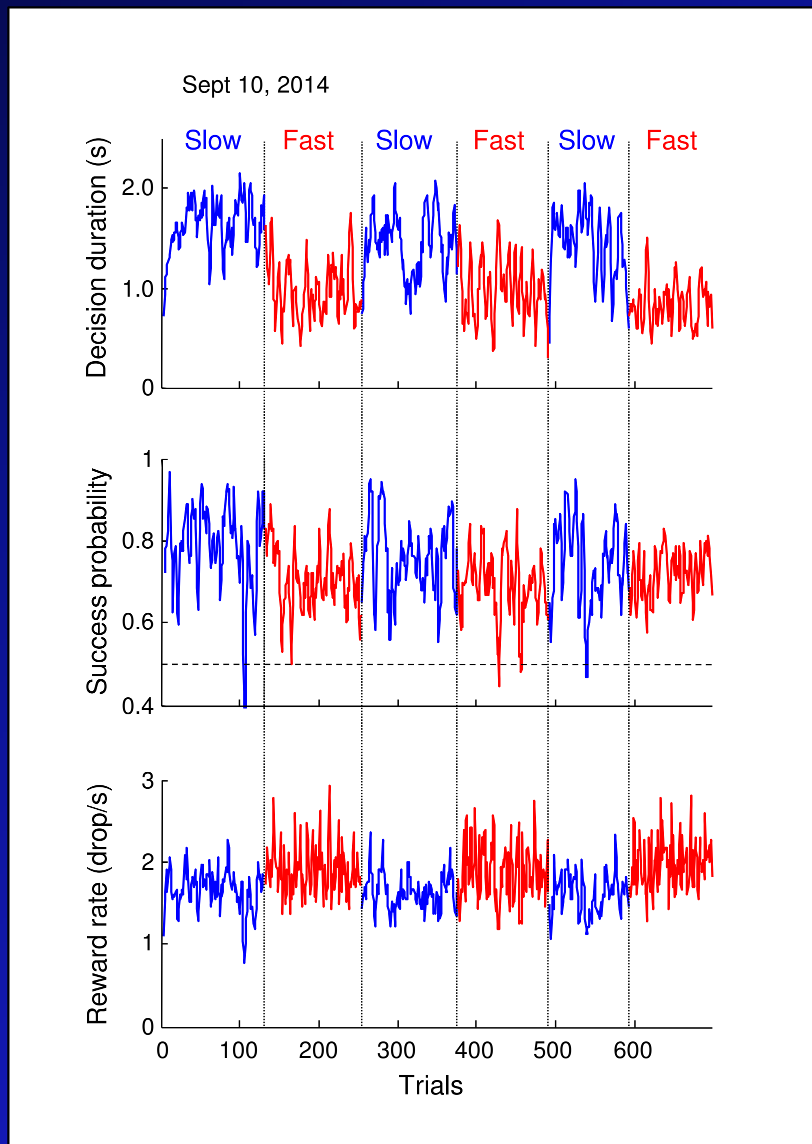

Do monkeys adjust their behavior?

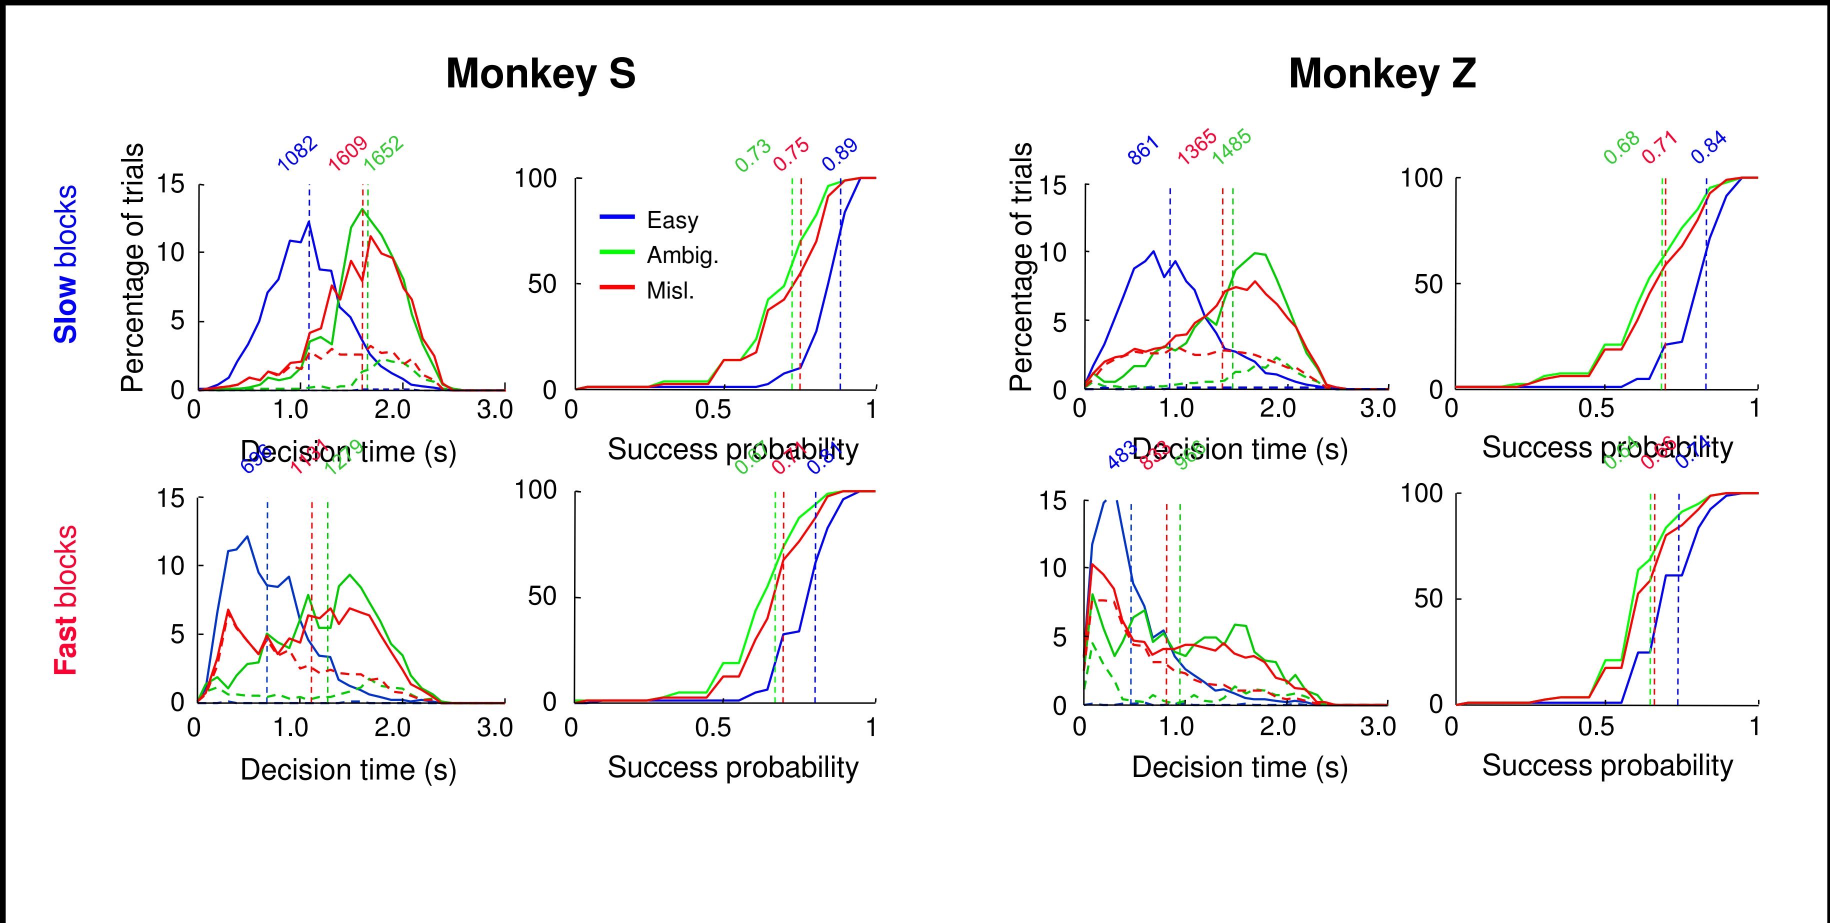

Thura, Cos, Trung & Cisek (2014) *J. Neurosci.*

Urgency is context-dependent

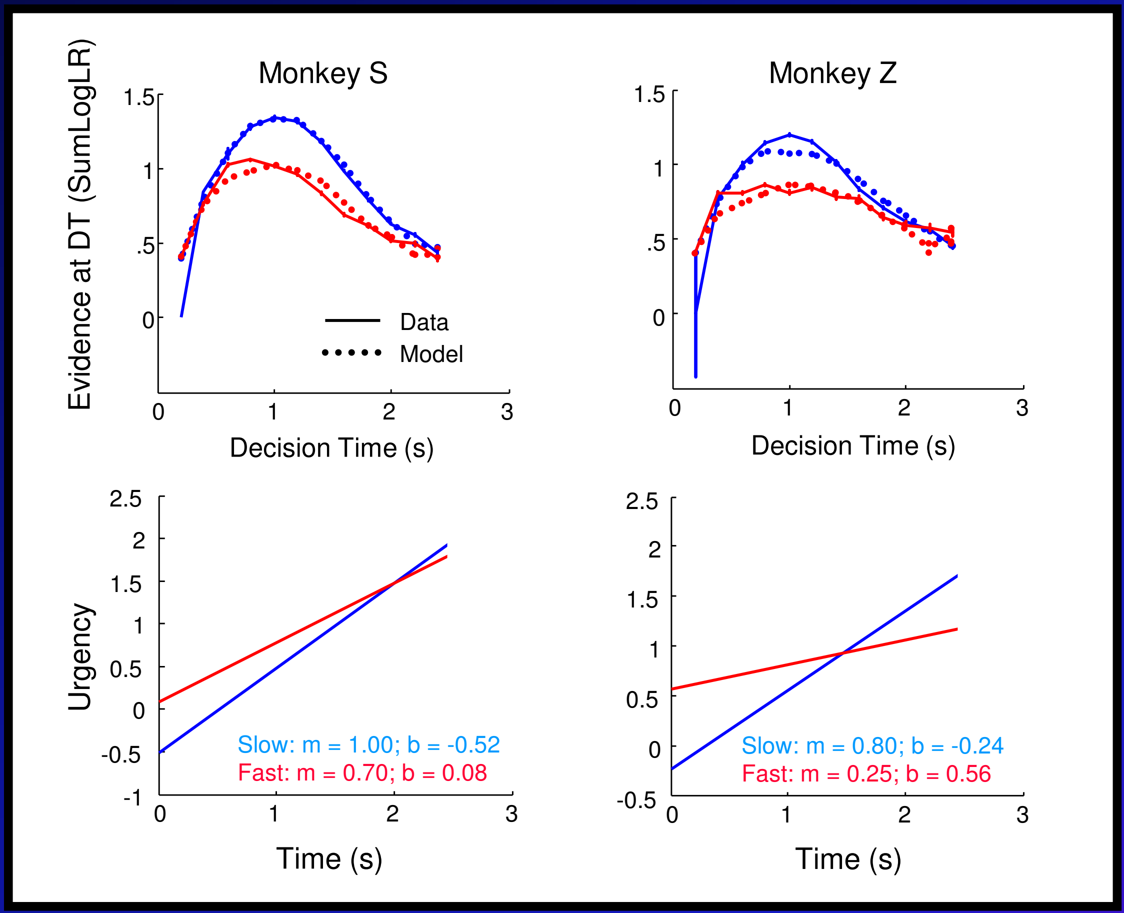

- We derive distinct urgency functions for the two blocks
- Fast starts higher than Slow, and they converge over time

Thura, Cos, Trung & Cisek (2014) *J. Neurosci.*

Neural correlates

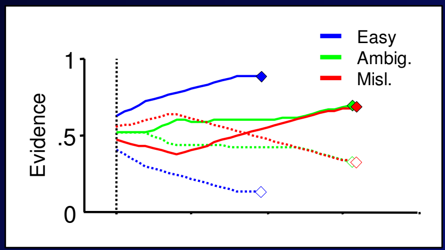

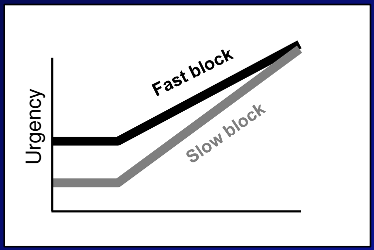

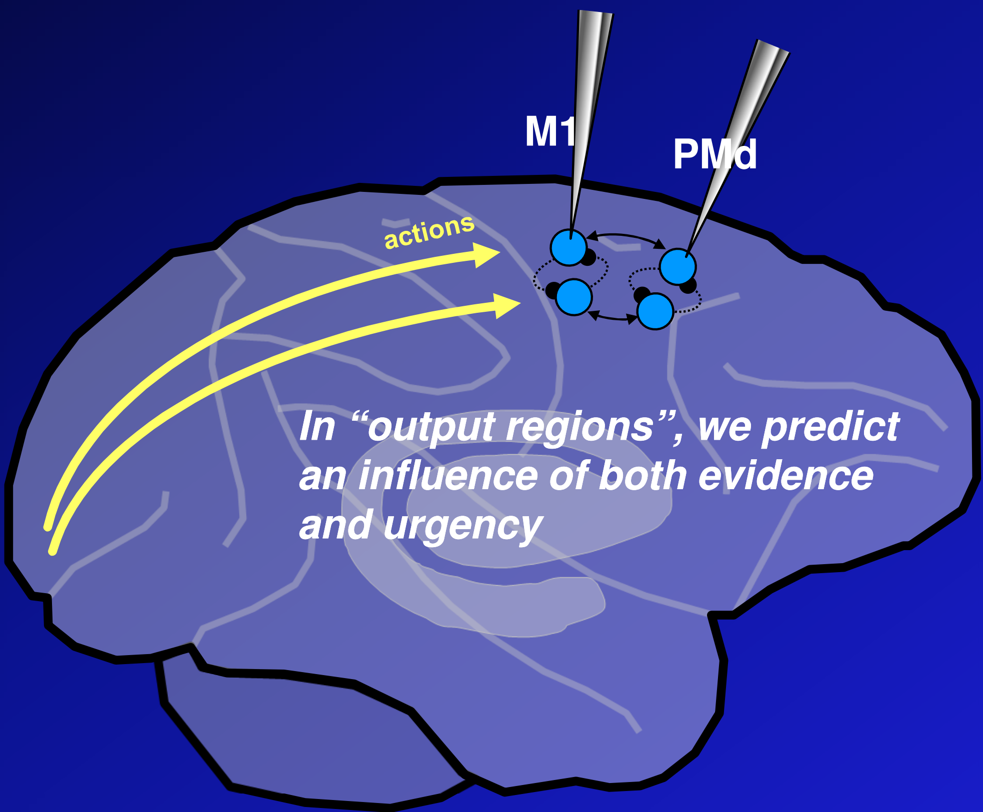

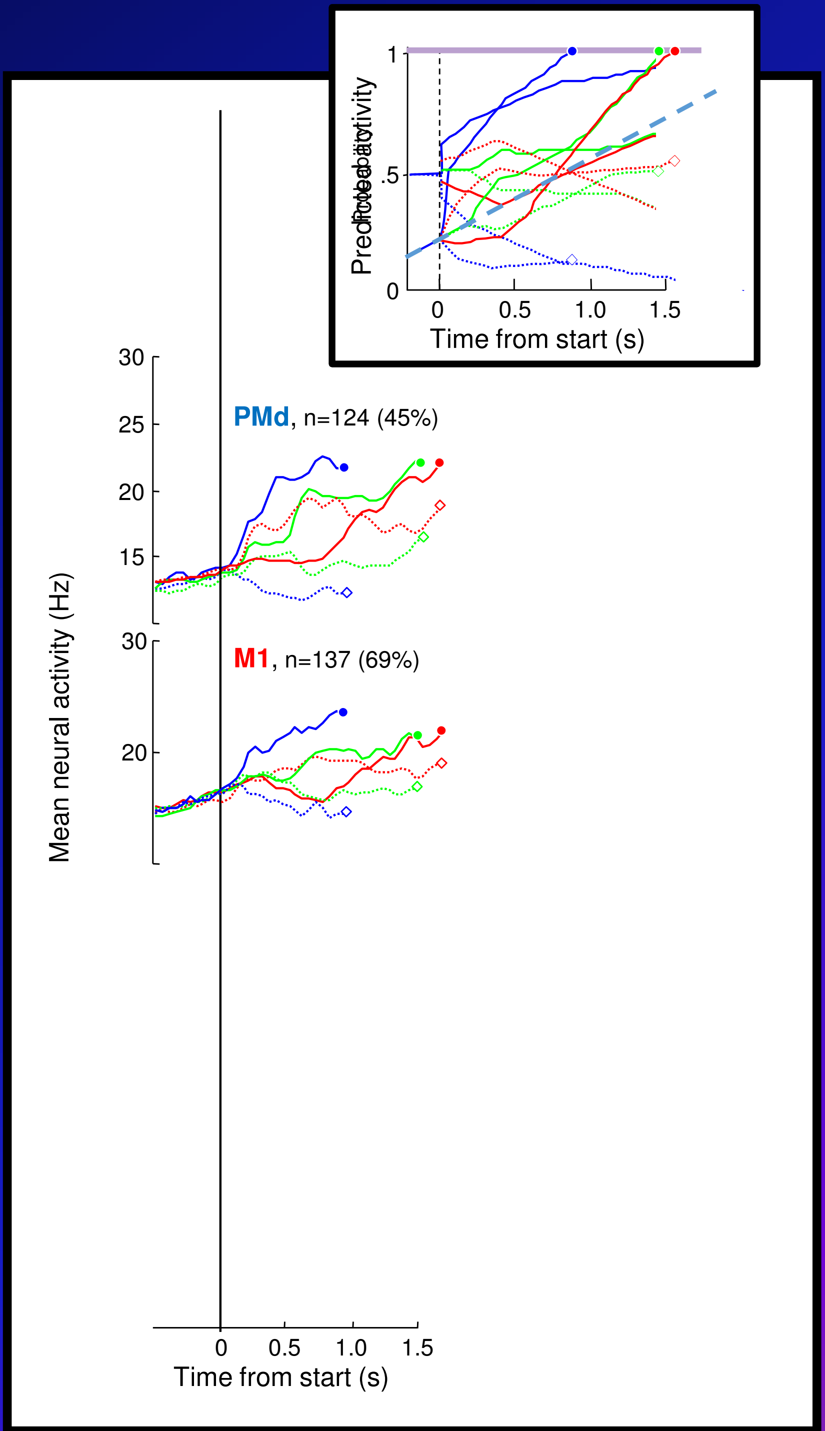

Thura & Cisek (2014) *Neuron*

The influence of "urgency"

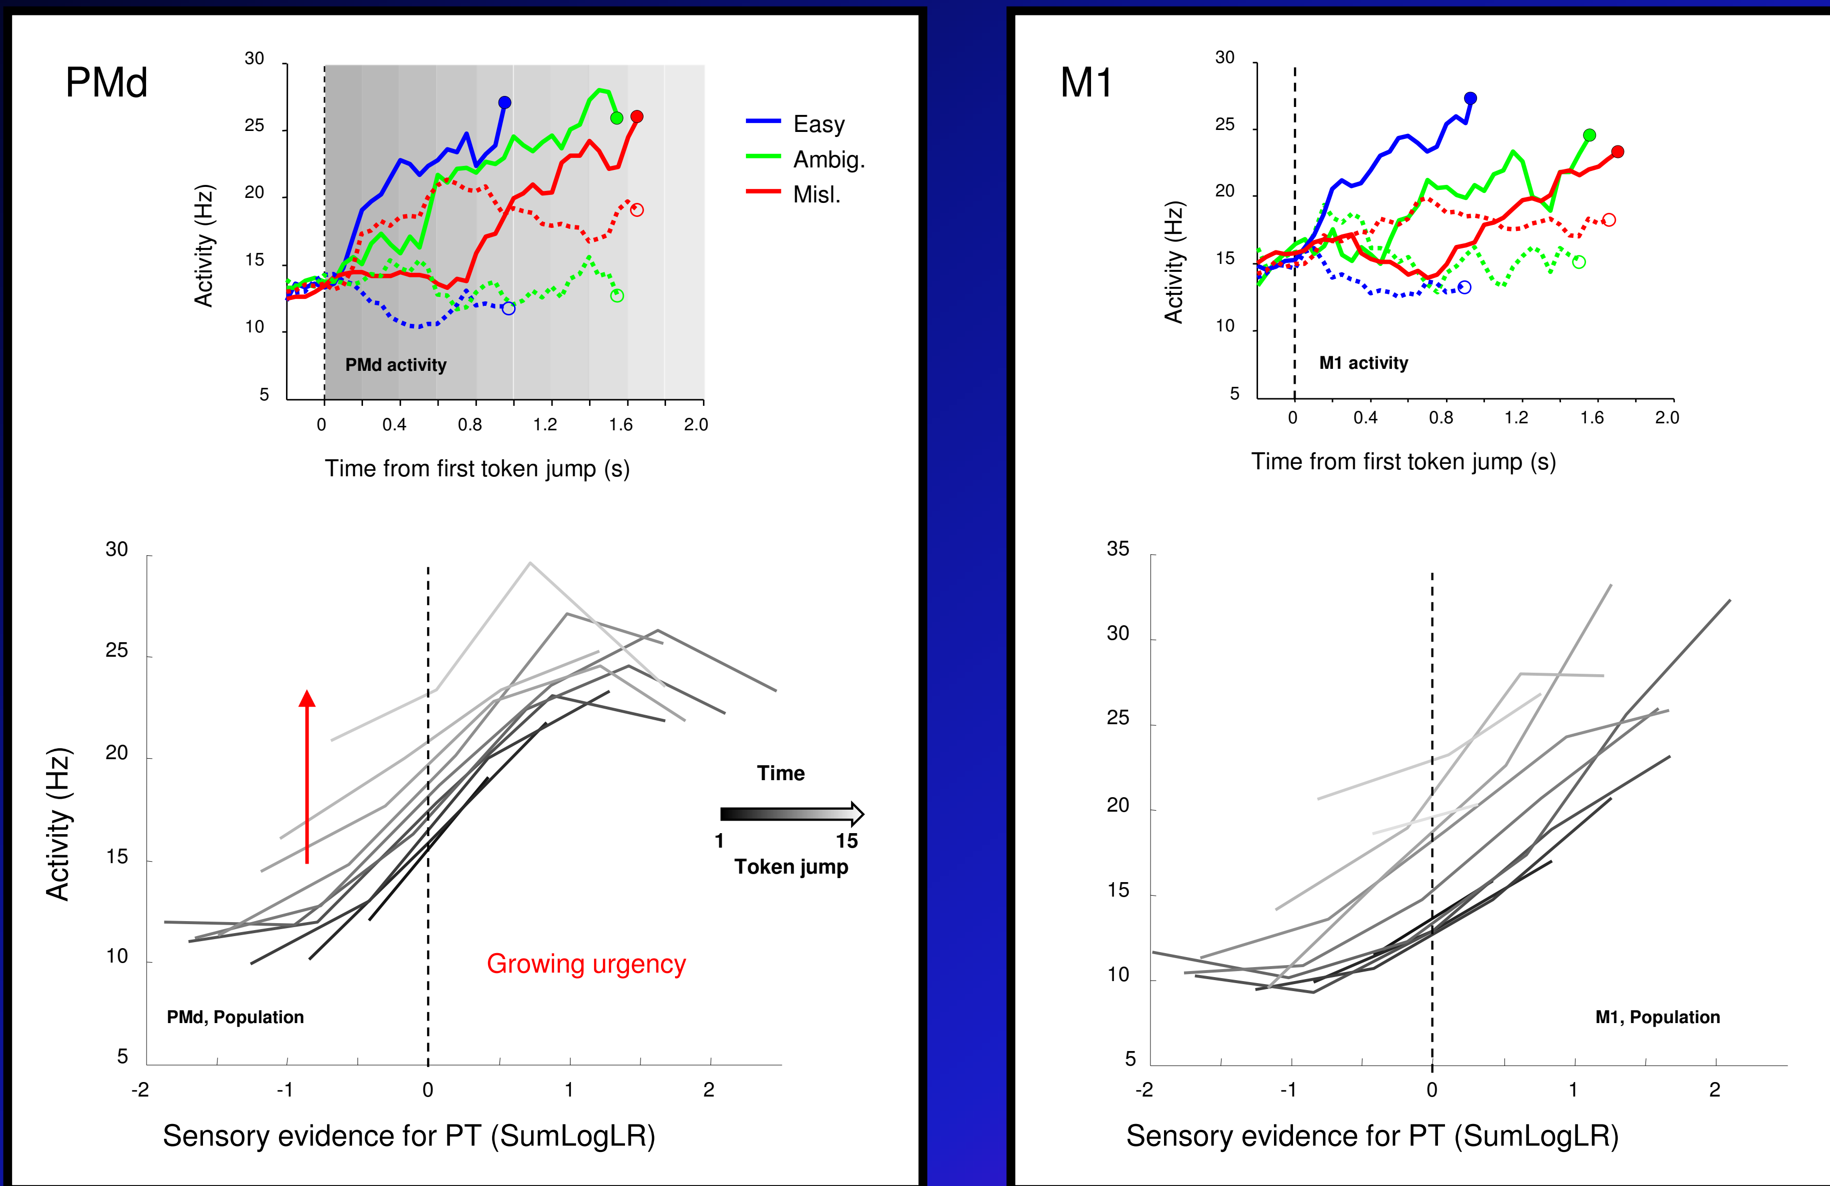

Deriving the urgency signal from neural data

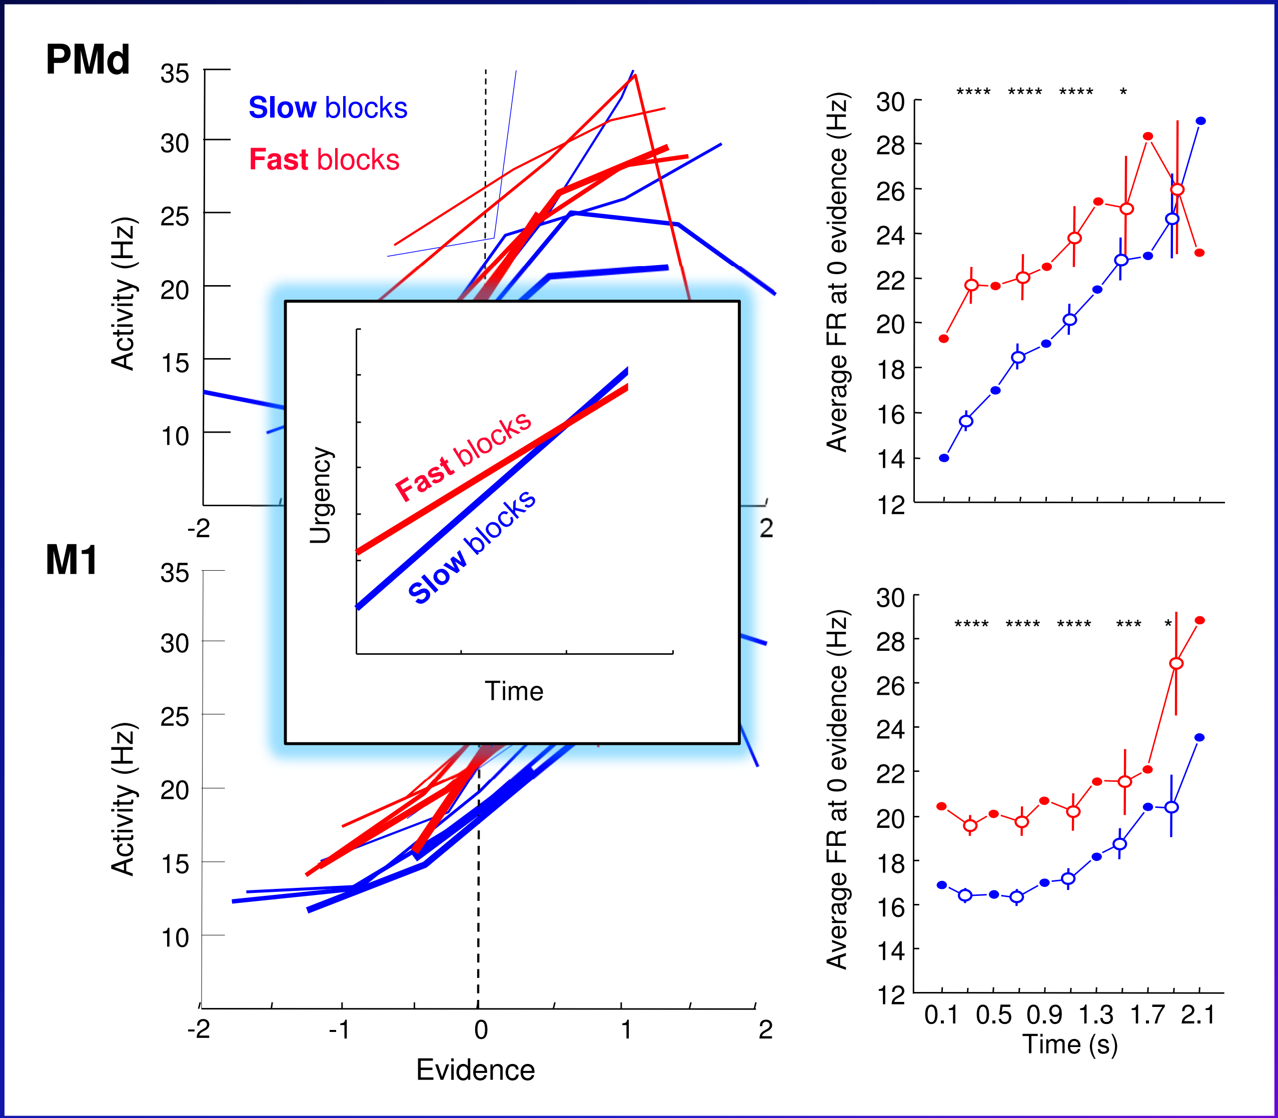

Thura & Cisek (2016) *J. Neurosci.*

Neural correlates

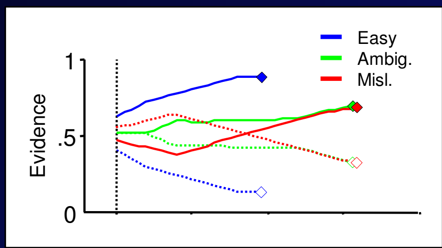

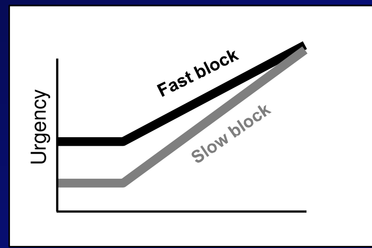

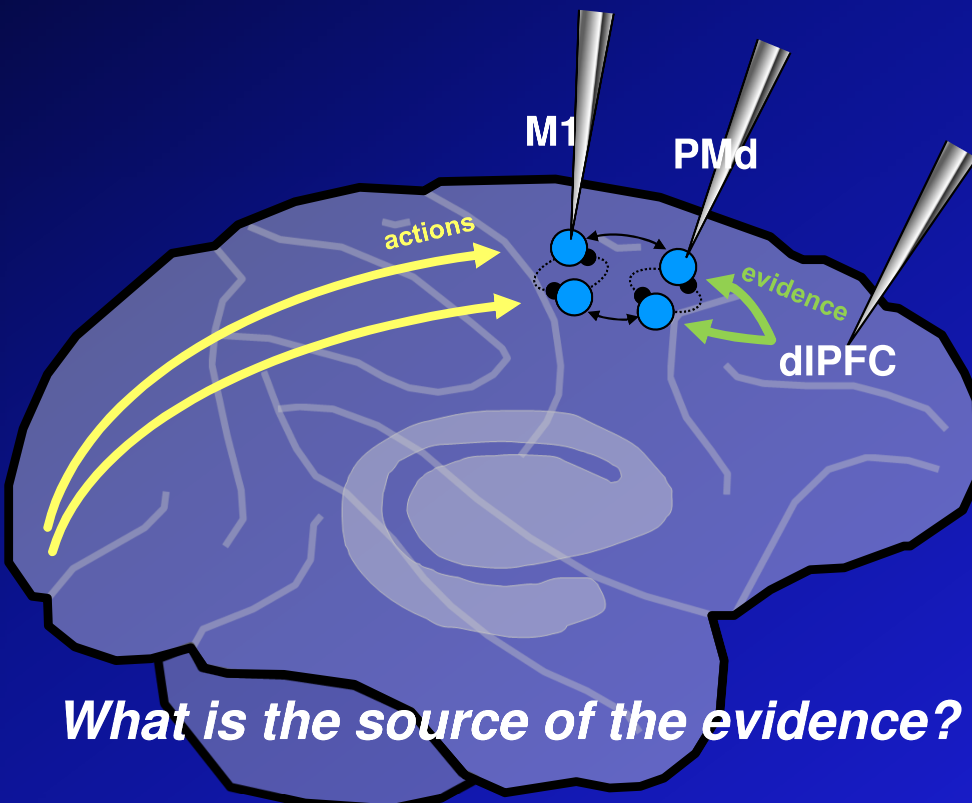

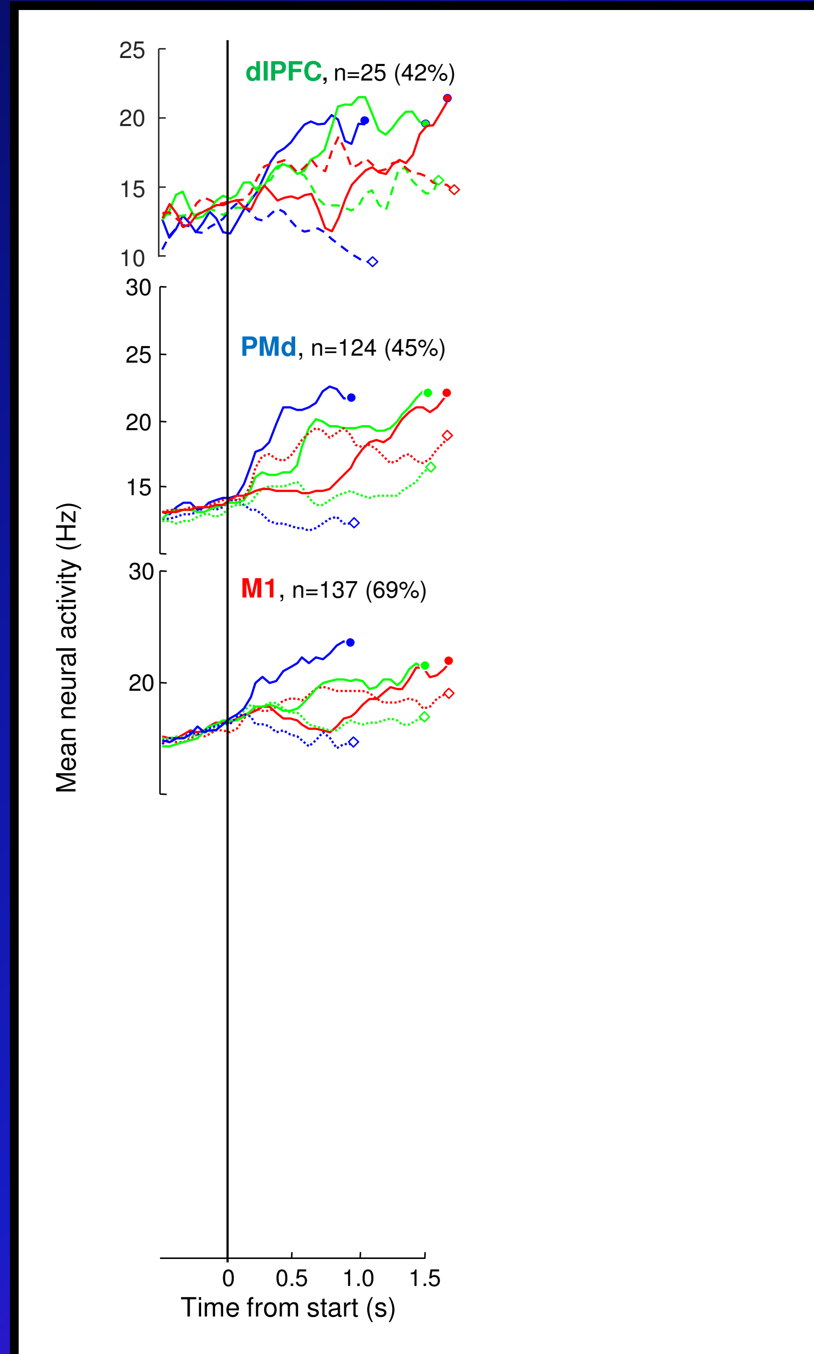

What is the source of the evidence?

What about urgency?

Thura & Cisek (2014) *Neuron*; Thura & Cisek (in prep)

Brief detour

- Reward rate

$$RR = \frac{p(t) \cdot U - C}{t + m + d}$$

- It also depends on movement time m
  - If you take more time to decide, you can save time by moving faster
  - If you took a guess, less need for accuracy
- Effort is a function of speed
  - Predict correlations between urgency and vigor

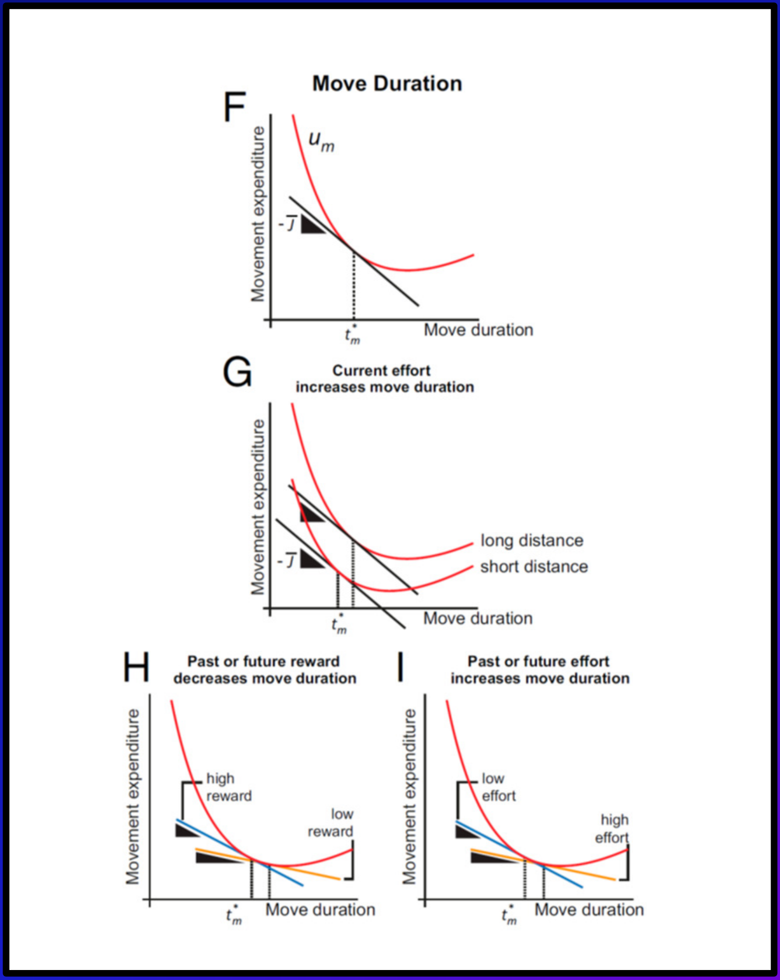

Yoon, Geary, Ahmed, Shadmehr (2018) *PNAS*

Correlation between urgency and vigor

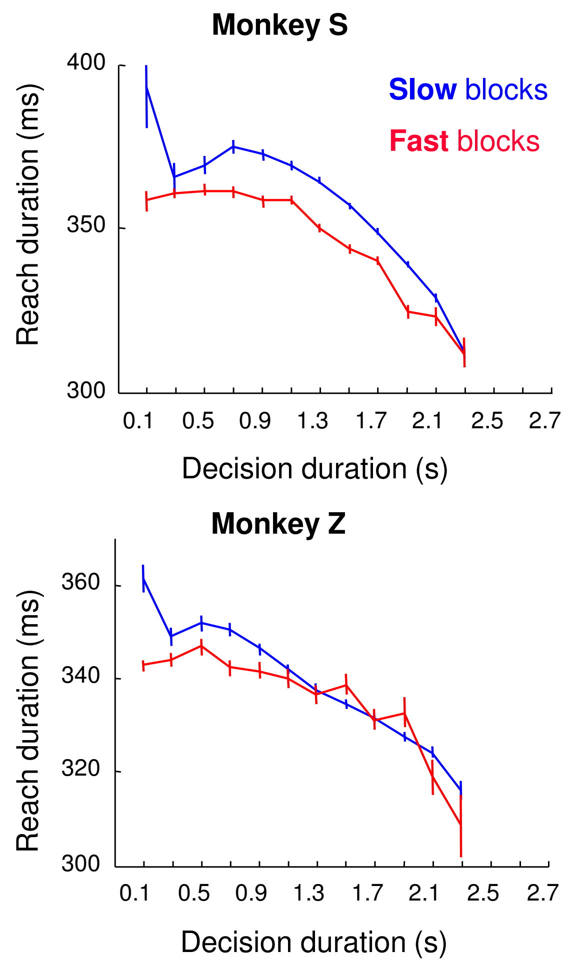

Thura, Cos, Trung & Cisek (2014) *J. Neurosci.*

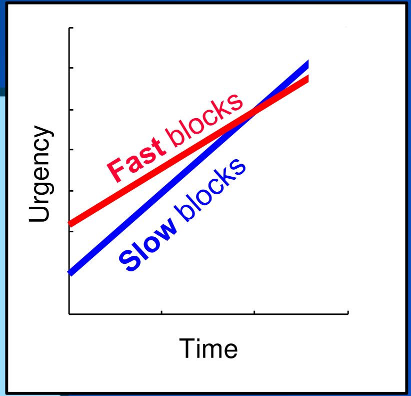

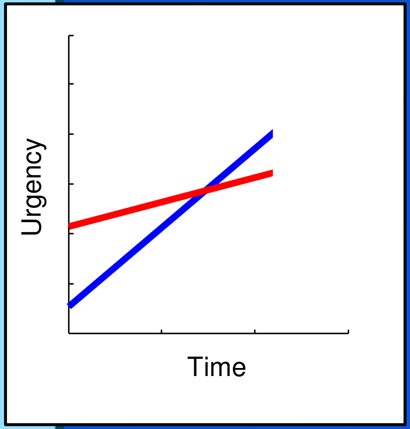

Proposal: The level of urgency at time of decision influences the vigor of action

Basal Ganglia?

Neural correlates

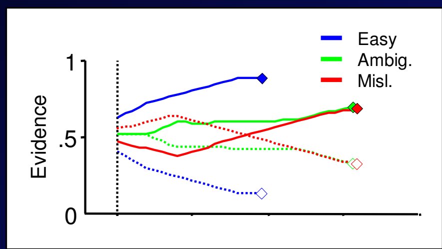

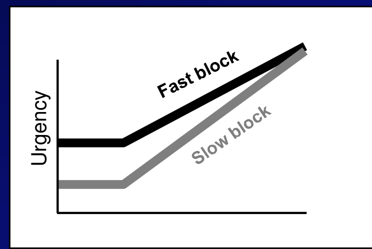

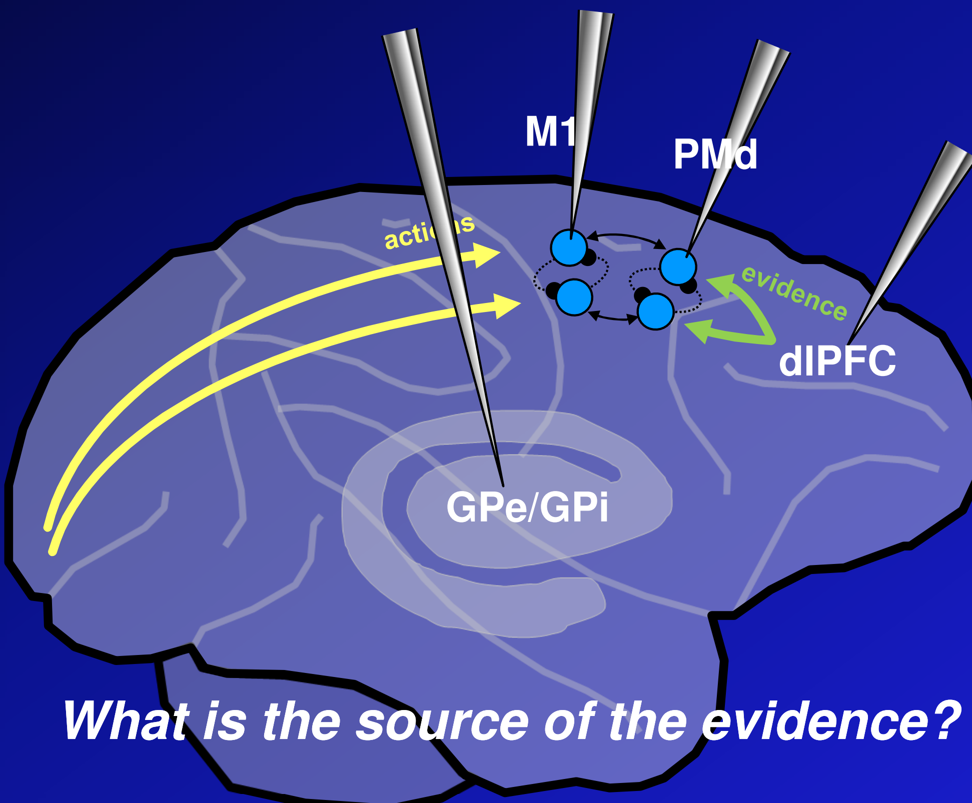

What is the source of the evidence?

What about urgency?

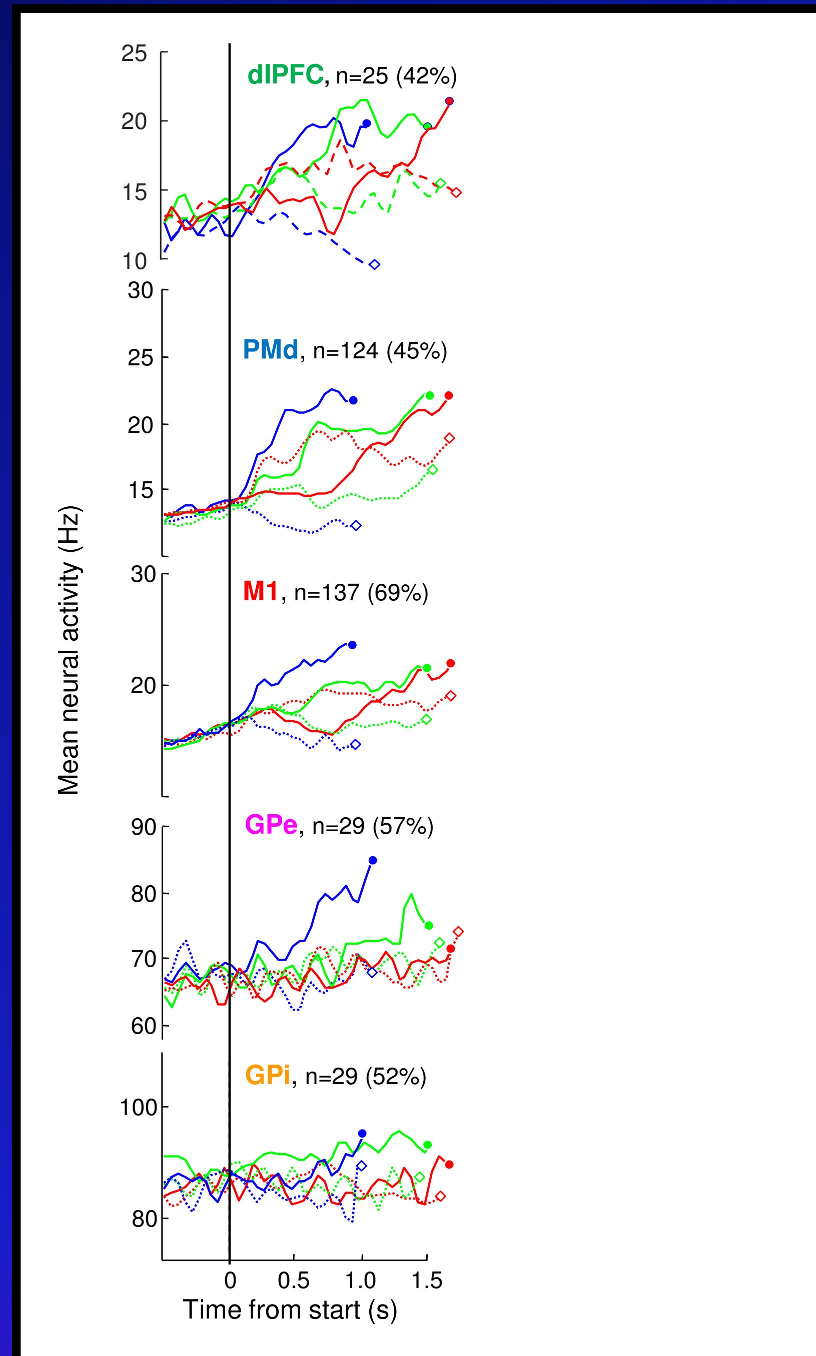

Thura & Cisek (2014) *Neuron*; Thura & Cisek (2016) *J. Neurosci.*; Thura & Cisek (2017) *Neuron*

Neural correlates

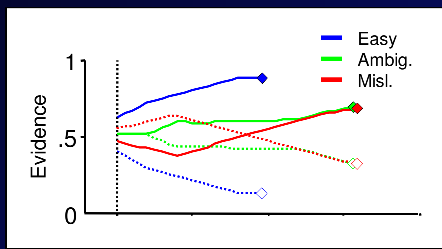

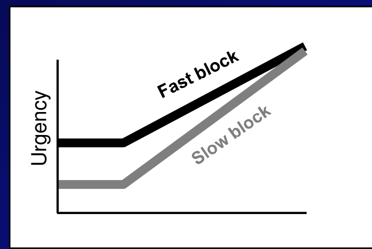

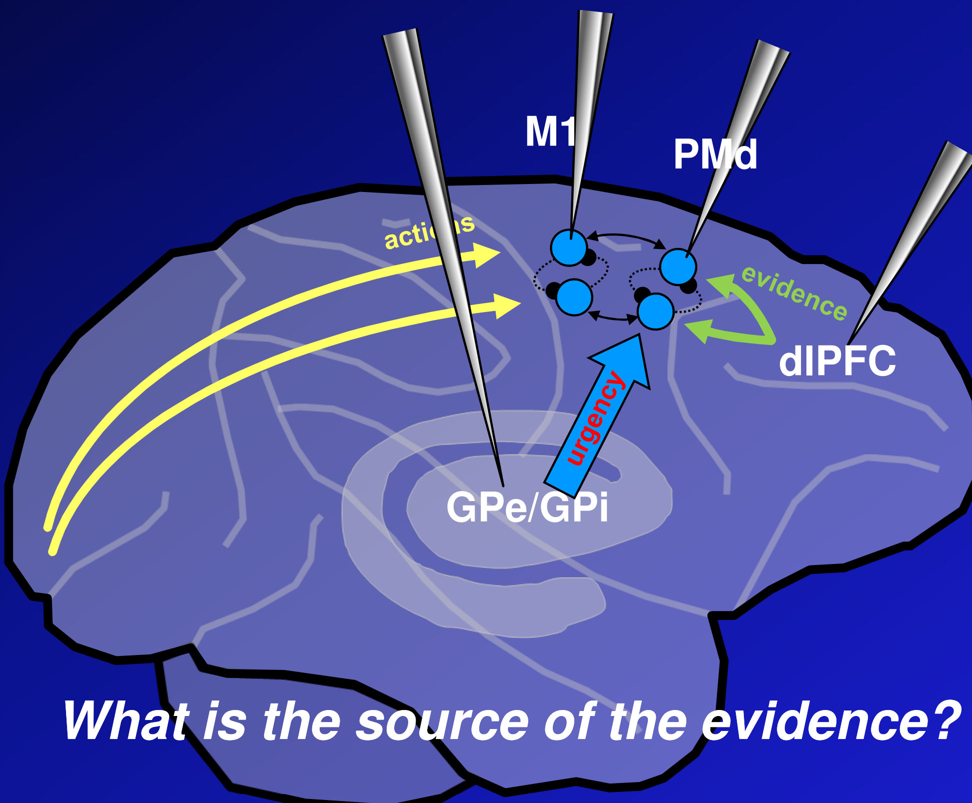

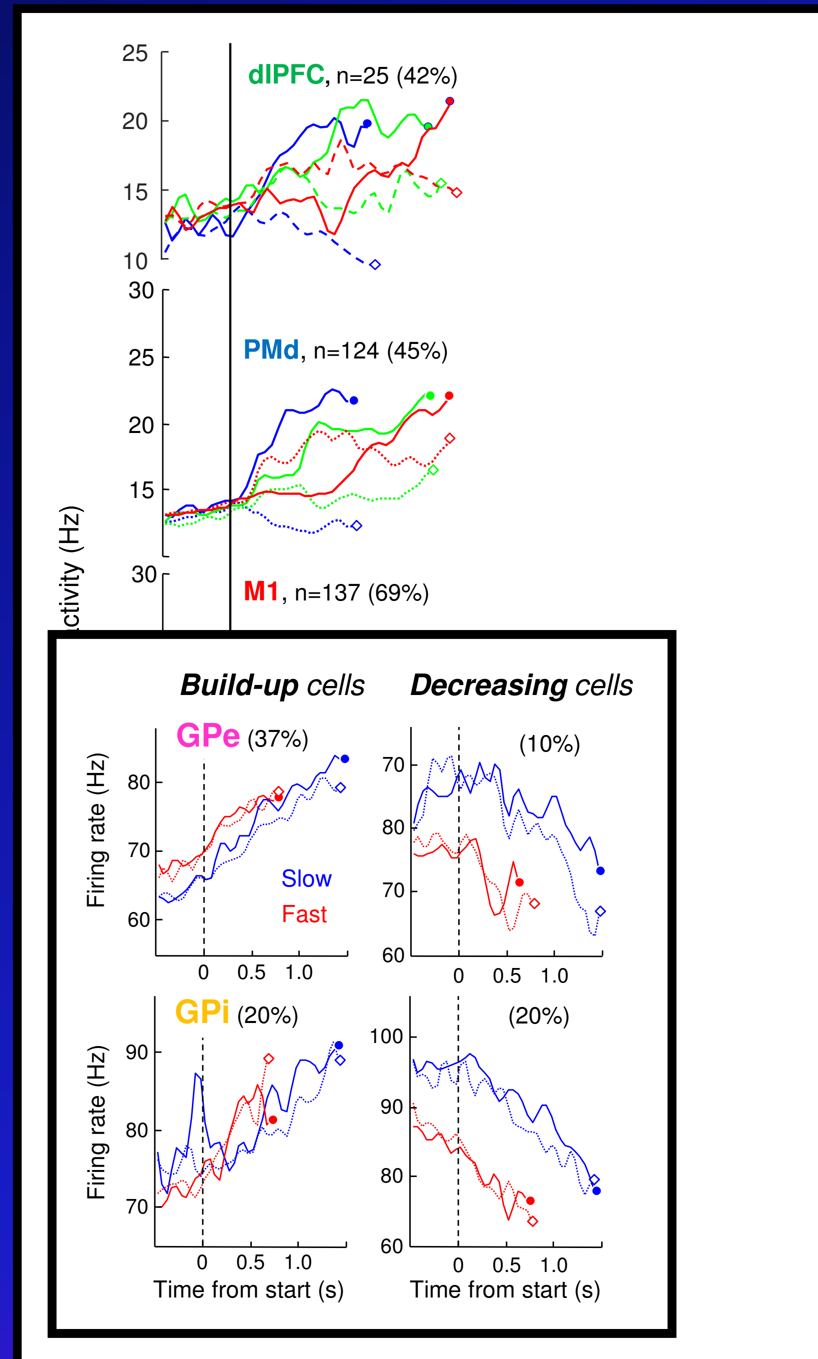

What is the source of the evidence?

What about urgency?

Thura & Cisek (2014) *Neuron*; Thura & Cisek (2016) *J. Neurosci.*; Thura & Cisek (2017) *Neuron*

Neural correlates

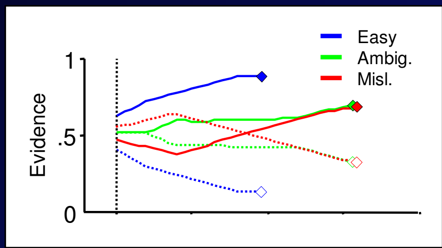

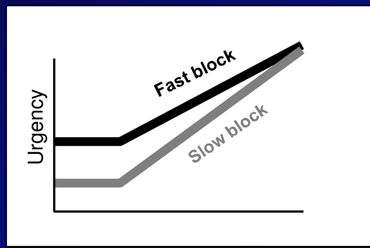

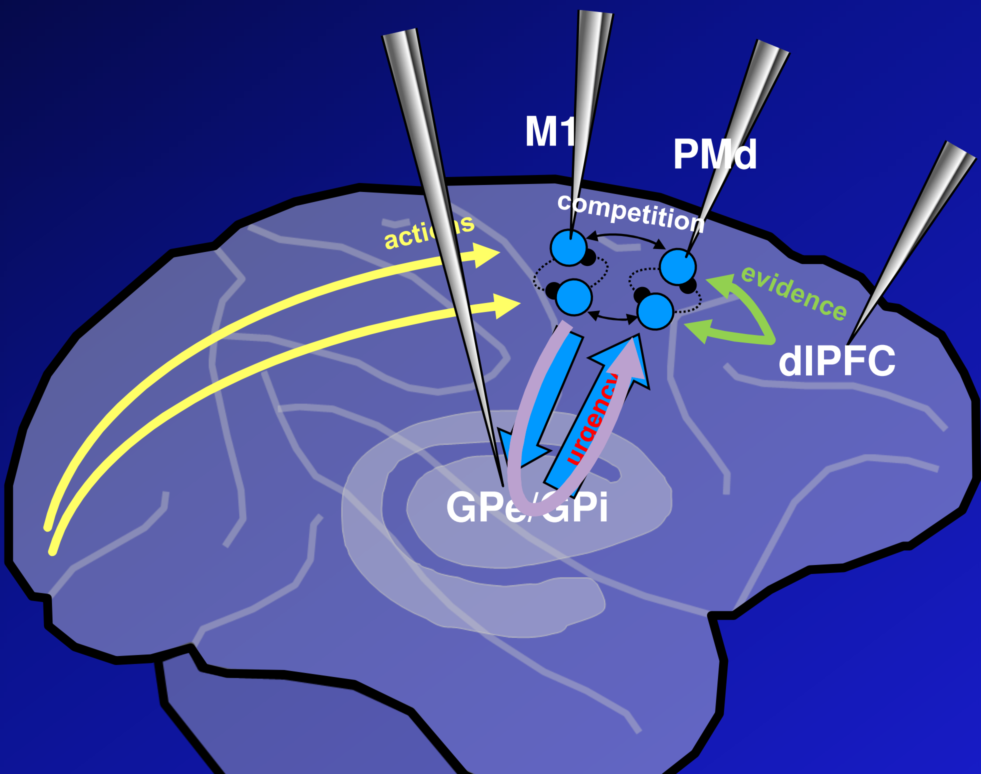

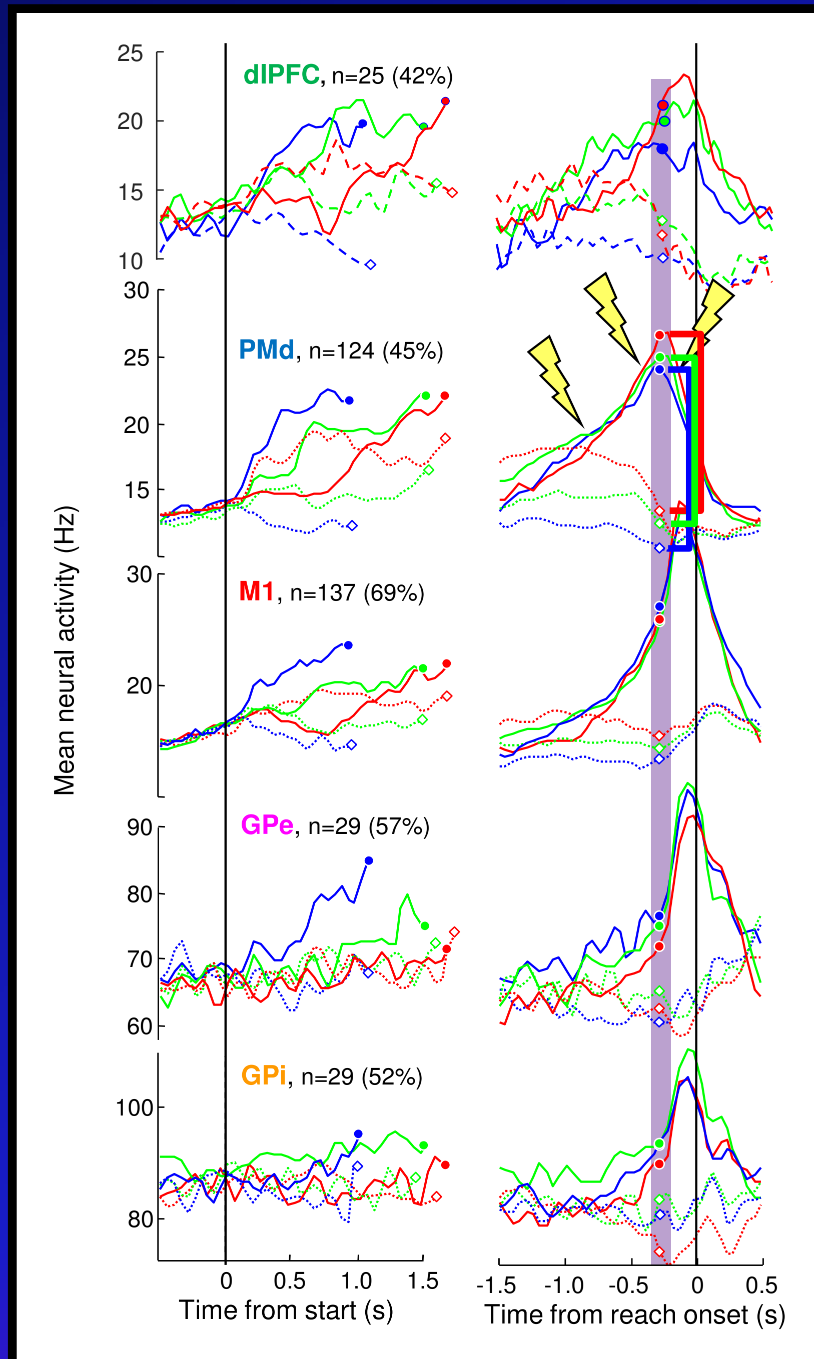

What about commitment?

Thura & Cisek (2014) *Neuron*; Thura & Cisek (2016) *J. Neurosci.*; Thura & Cisek (2017) *Neuron*
Thura & Cisek (2020) *J. Neurophys.*

Summary so far

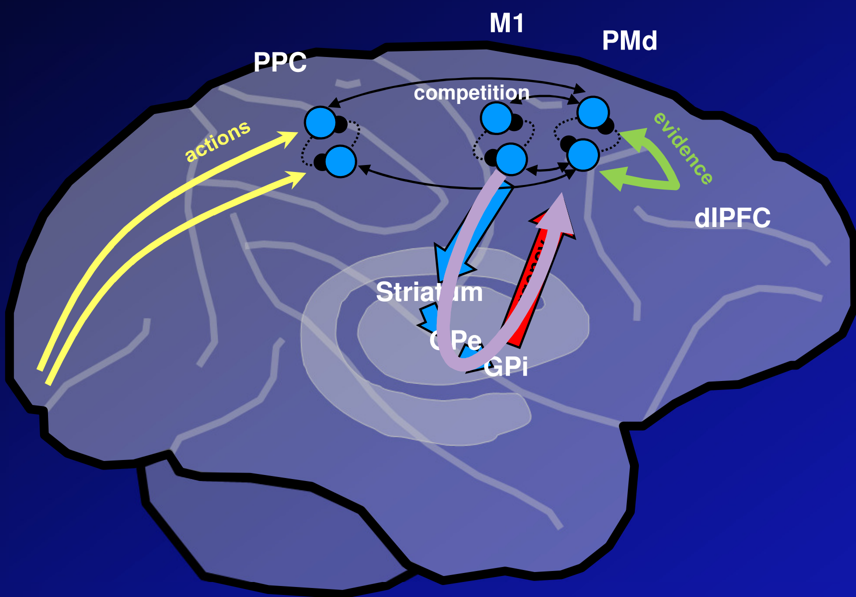

Proposal:
Commitment occurs when the recurrent dynamical system (PMd/M1 ↔ BG) falls into an "attractor"

- Deliberation as competition
  - Potential actions from dorsal stream
  - Competition in PPC/PMd/M1
  - Evidence from dlPFC
  - Urgency from BG
- Cortex develops a contrast
  - Starts to spill into GPe
  - Reaches critical point and spills into GPi
  - Positive feedback
  - Volitional commitment to a *reaching* action

Recurrent attractor model

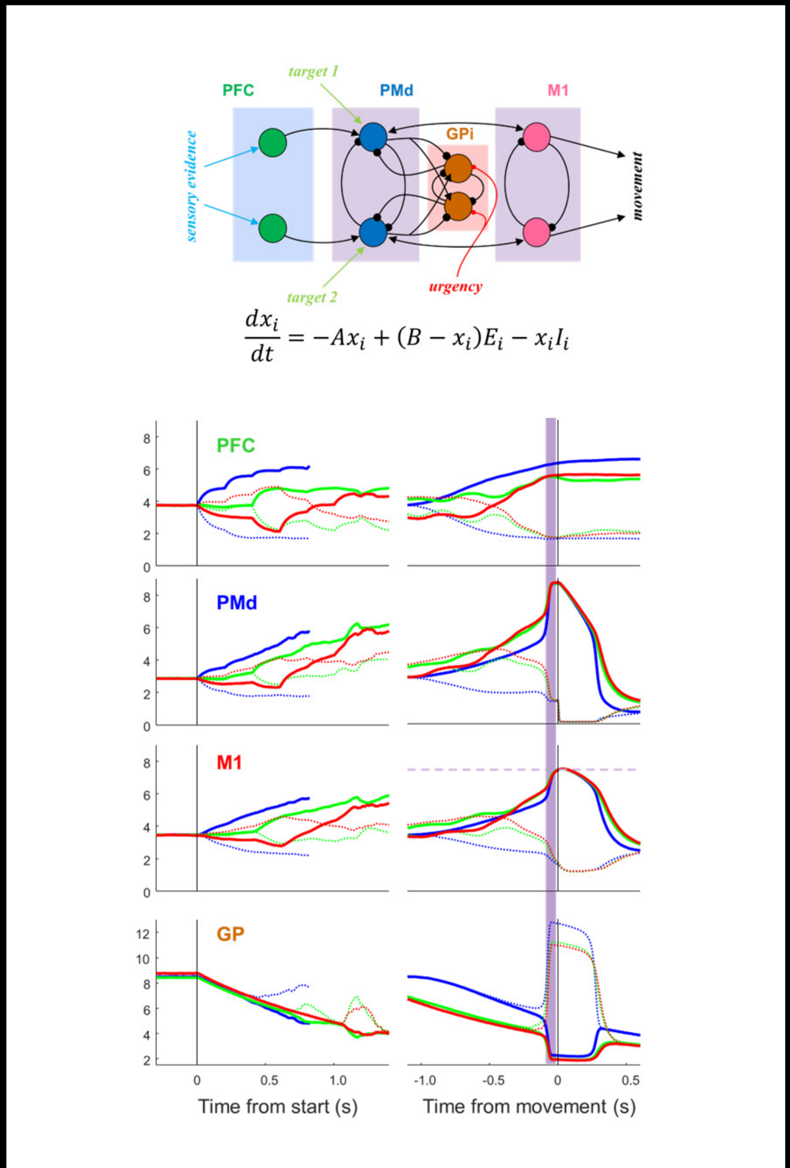

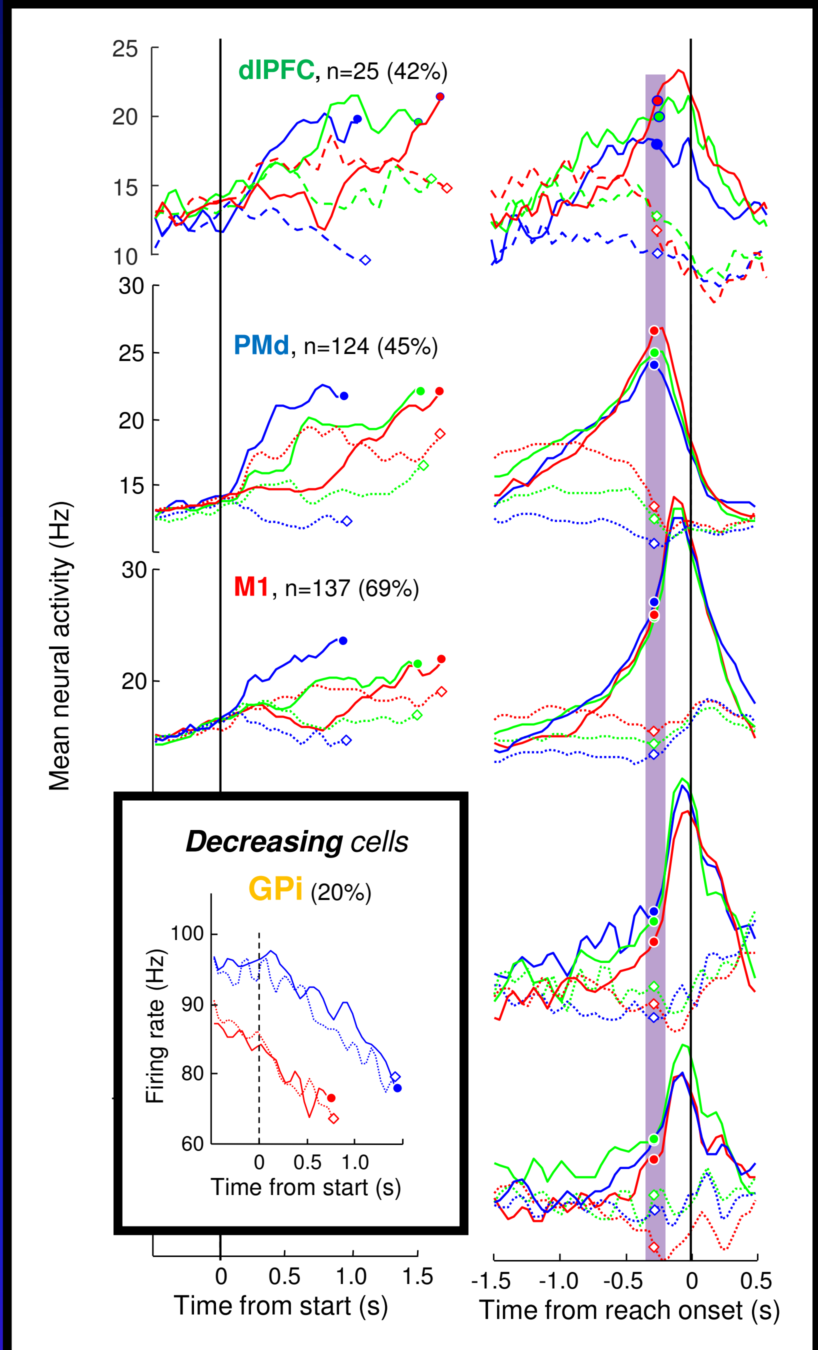

Thura & Cisek (2014) *Neuron*; Thura & Cisek (2016) *J. Neurosci.*; Thura & Cisek (2017) *Neuron*

Visualizing the dynamics

- Compute neural state in the high-dimensional space of recorded cells
- Extract the principal components (PCs) that capture the variance

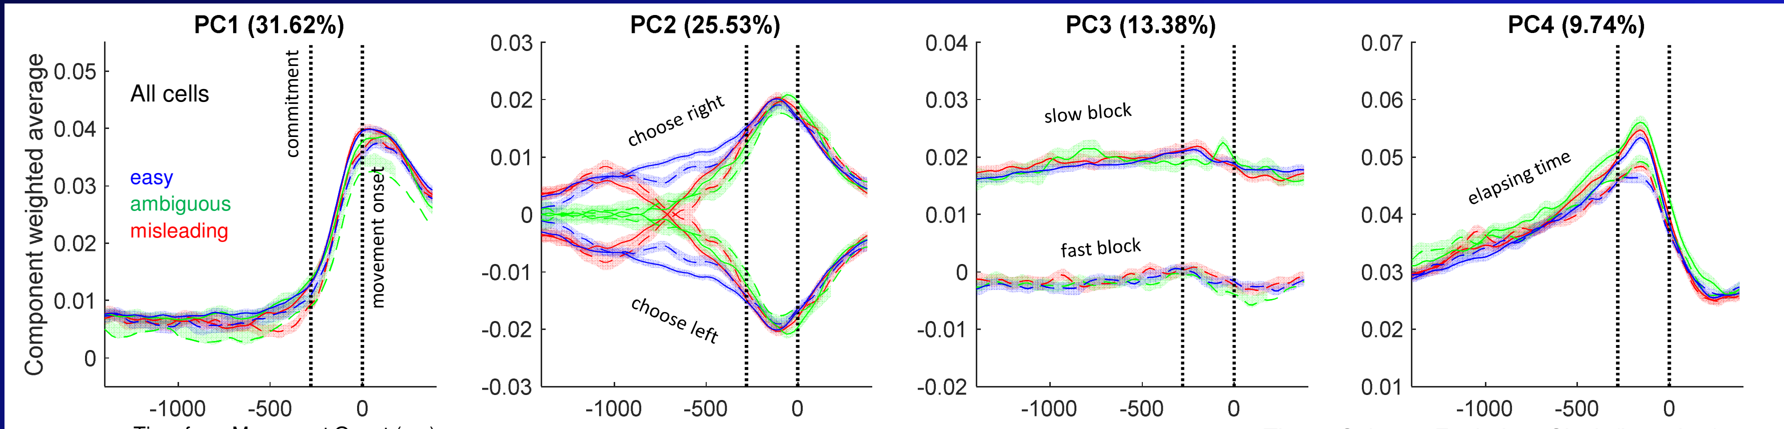

Thura, Cabana, Feghaly & Cisek (in review)

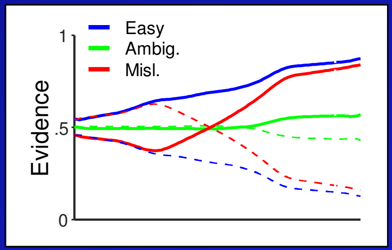

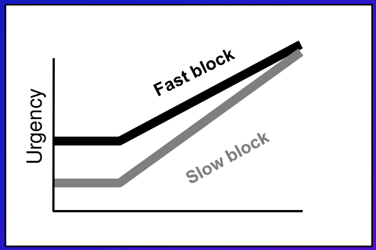

- The top four components reflect the key elements of the urgency gating model: ***commitment***, ***evidence***, baseline and slope of ***urgency***

Visualizing the dynamics

- Trajectories in PC space
  - States during deliberation
    - “Decision manifold”

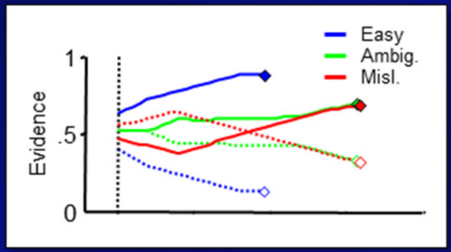

  - Moment of commitment

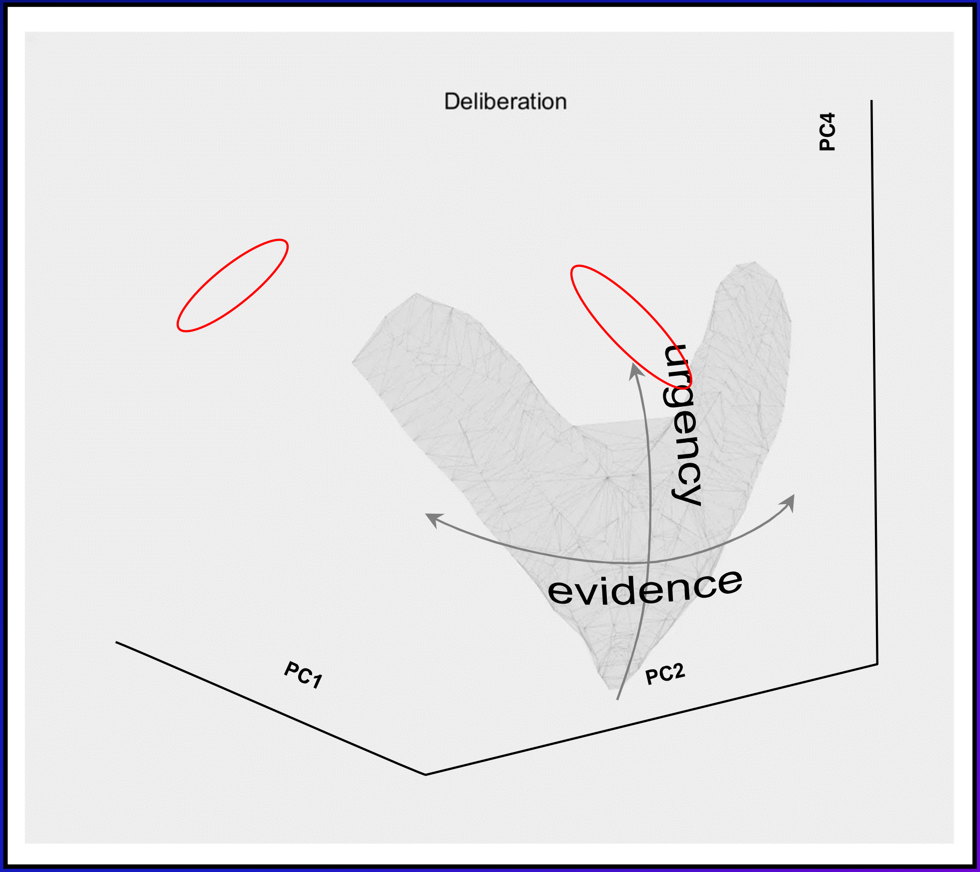

Visualizing the dynamics

- Trajectories in PC space
  - States during deliberation
    - “Decision manifold”

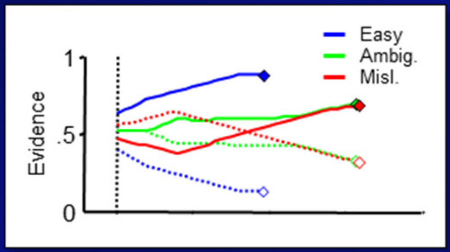

  - Moment of commitment
  - Choice-specific attractors

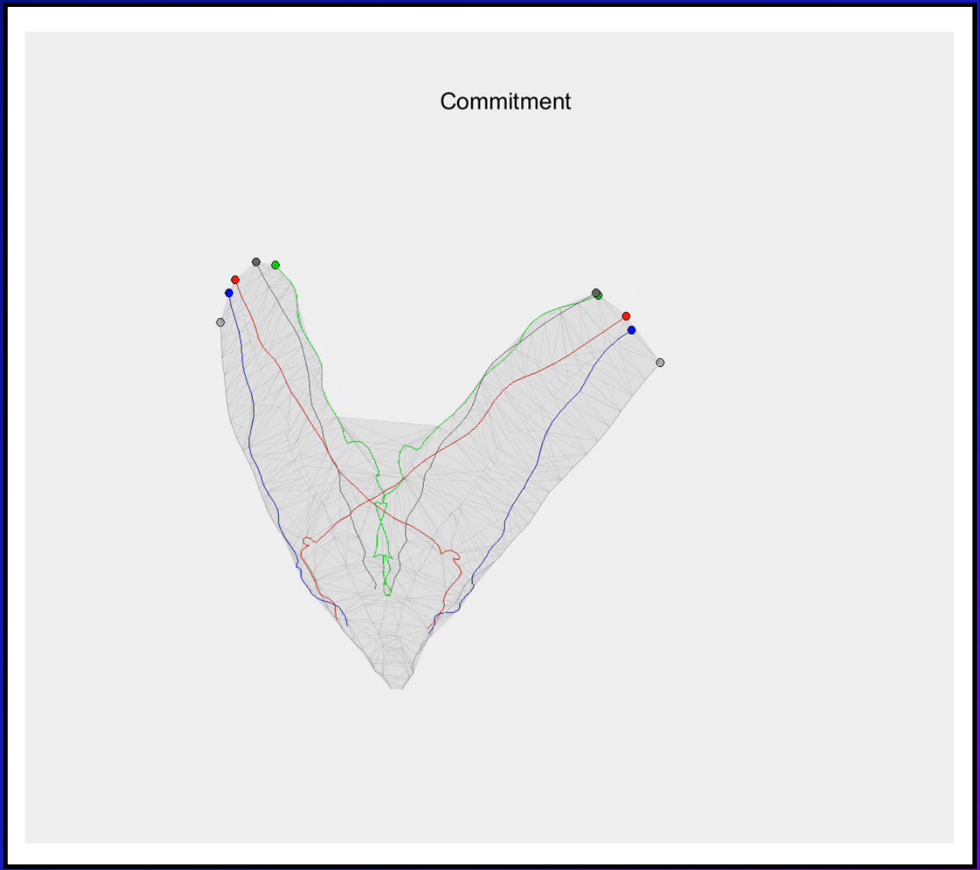

- Transition from deliberation to commitment is a state transition within a unified dynamical system

Unified system, but with specialization

- Different dynamics in different regions
  - Dorsal premotor cortex
    - Curved decision manifold
  - Primary motor cortex
    - Planar decision manifold

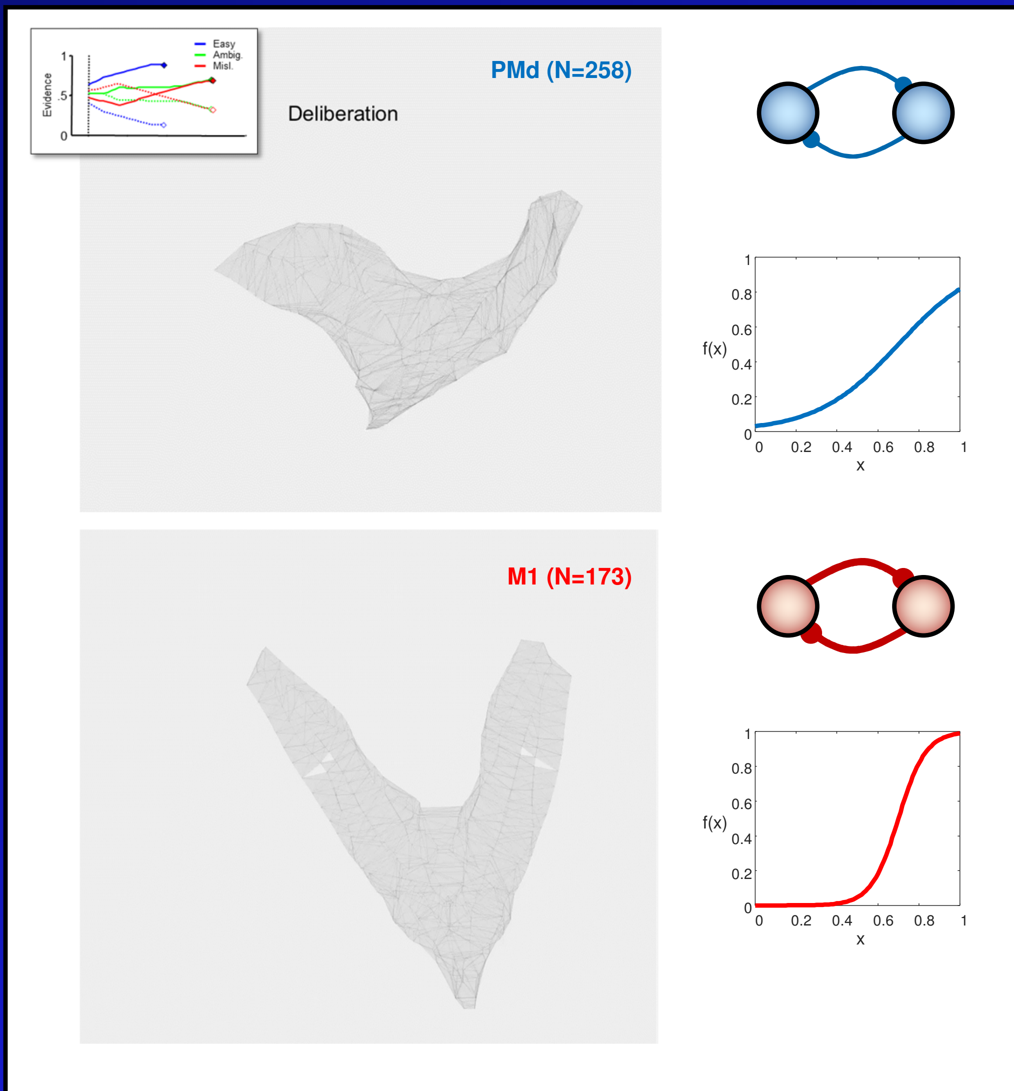

Unified system, but with specialization

- Different dynamics in different regions
  - Dorsal premotor cortex
    - Curved decision manifold
  - Primary motor cortex
    - Planar decision manifold
  - Dorsolateral prefrontal cortex
    - Extended mostly along PC2 (evidence)
  - Globus pallidus
    - Extended mostly along PC4 (urgency)
- PFC provides evidence, basal ganglia provide urgency, PMd/M1 put them together

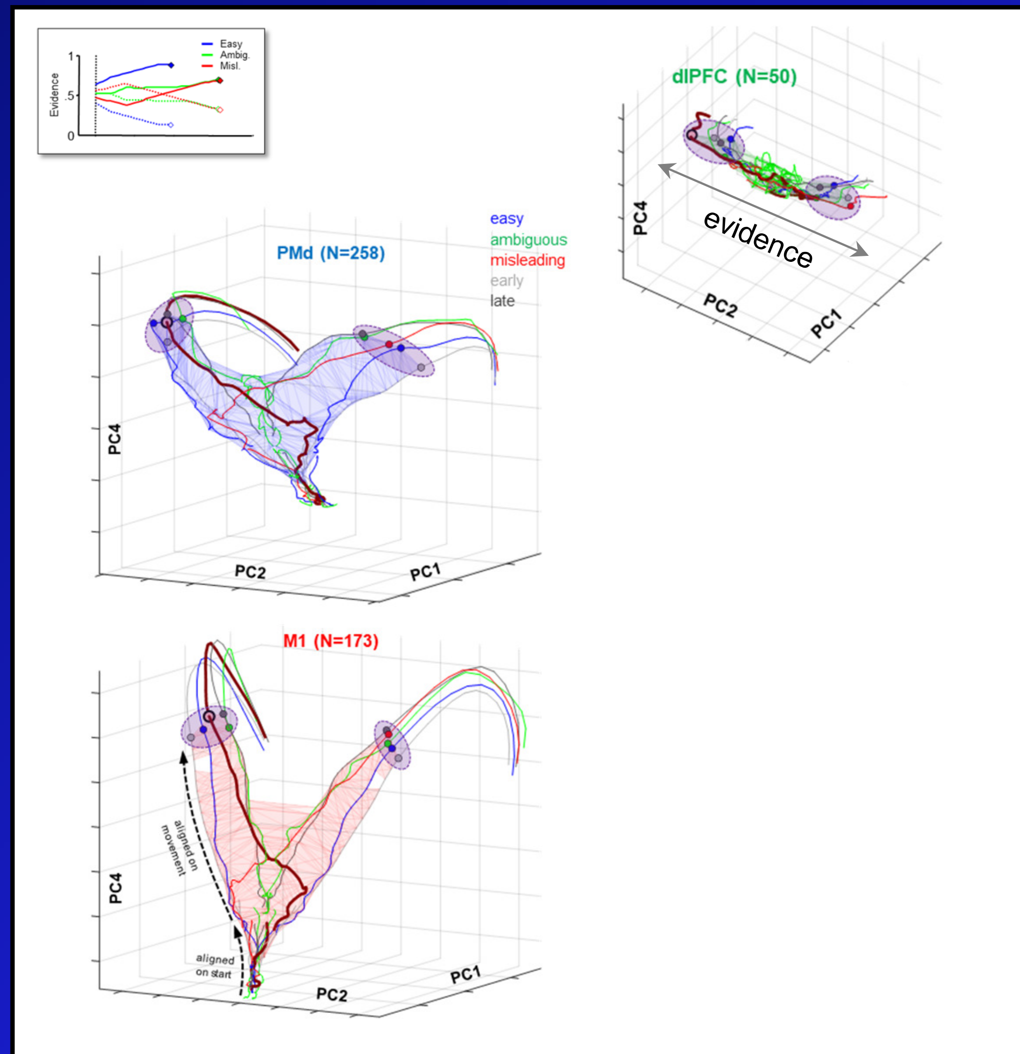

Clinical implications?

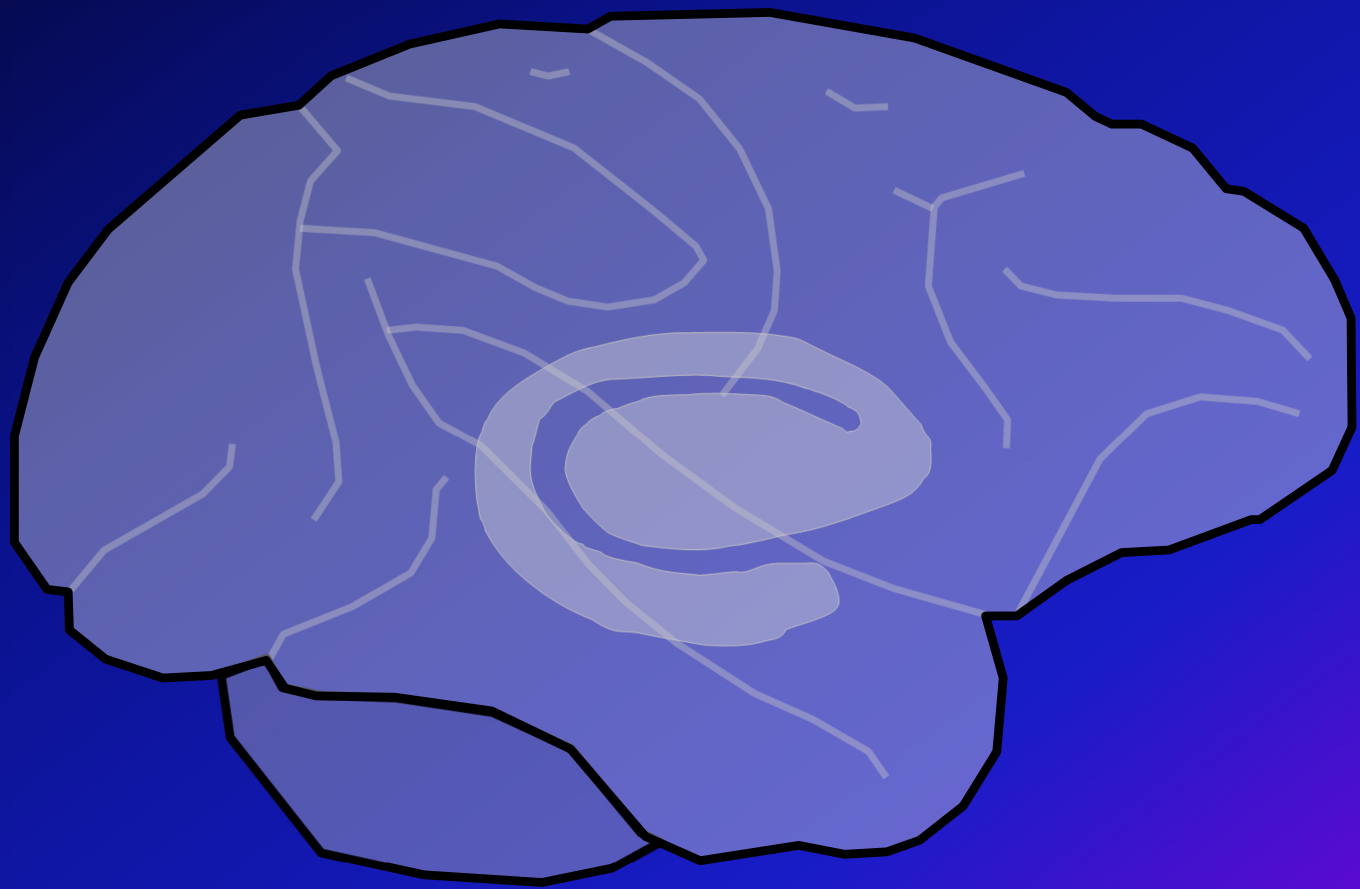

Urgency as an individual trait

Matt Carland

- Subjective differences
  - Feeling of effort
  - Importance of time
  - …
- Different urgency signals

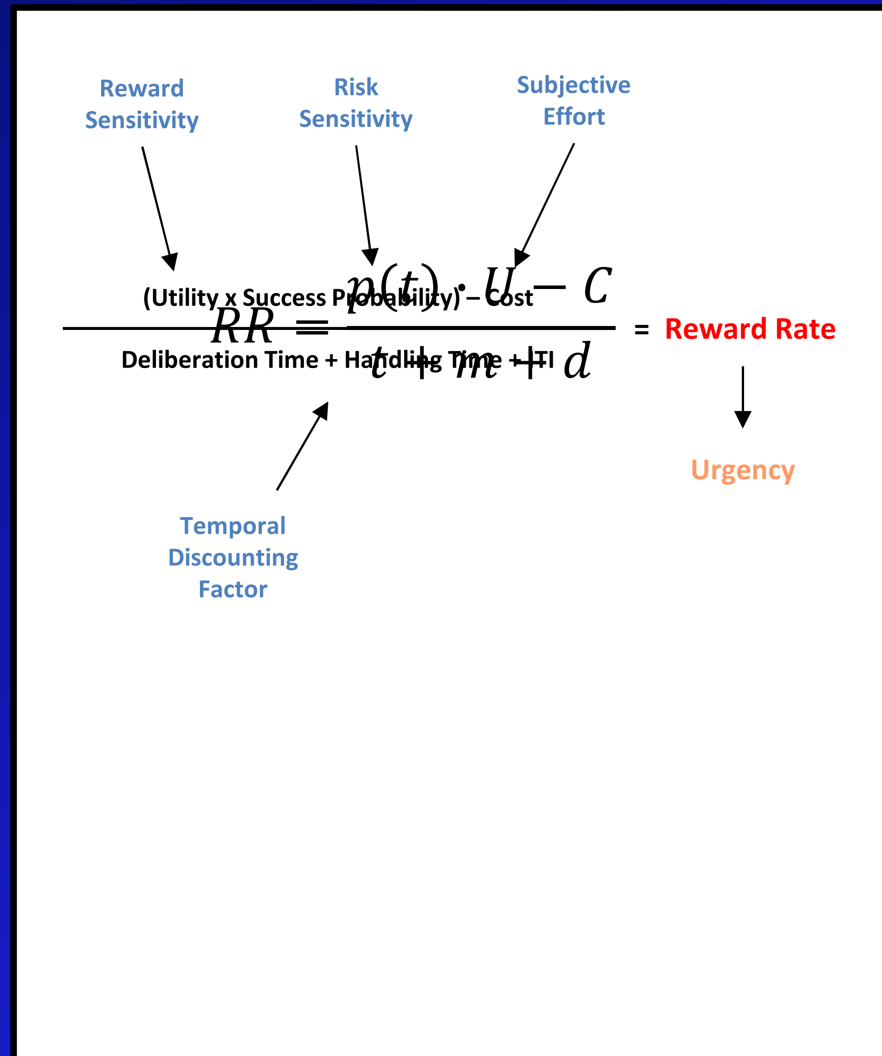

Urgency as an individual trait

- Subjective differences
  - Feeling of effort
  - Importance of time
  - …
- Different urgency signals
  - Low-urgency individual
  - High-urgency individual

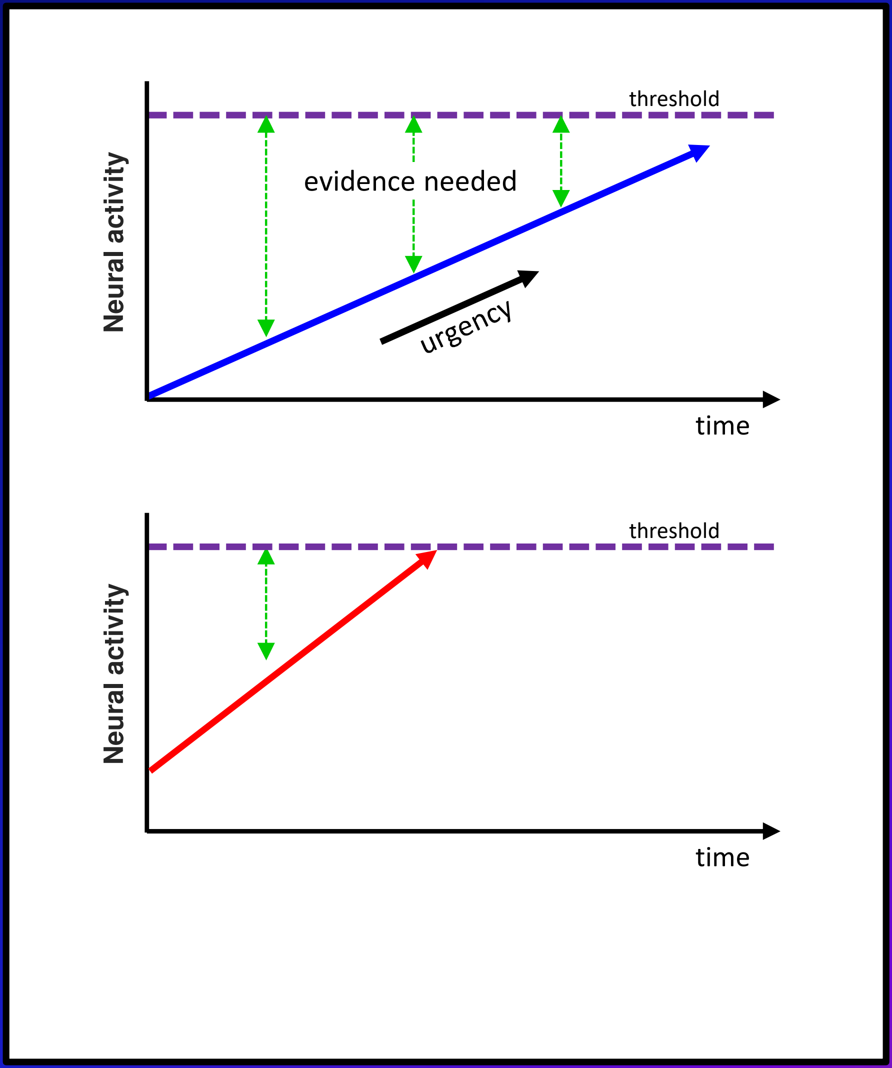

Urgency as an individual trait

- Subjective differences
  - Feeling of effort
  - Importance of time
  - …
- Different urgency signals
- Free response tasks
  - High-urgency → short RT
  - Low-urgency → long RT
- Fixed duration tasks
  - High-urgency → fast mvmt
  - Low-urgency → slow mvmt

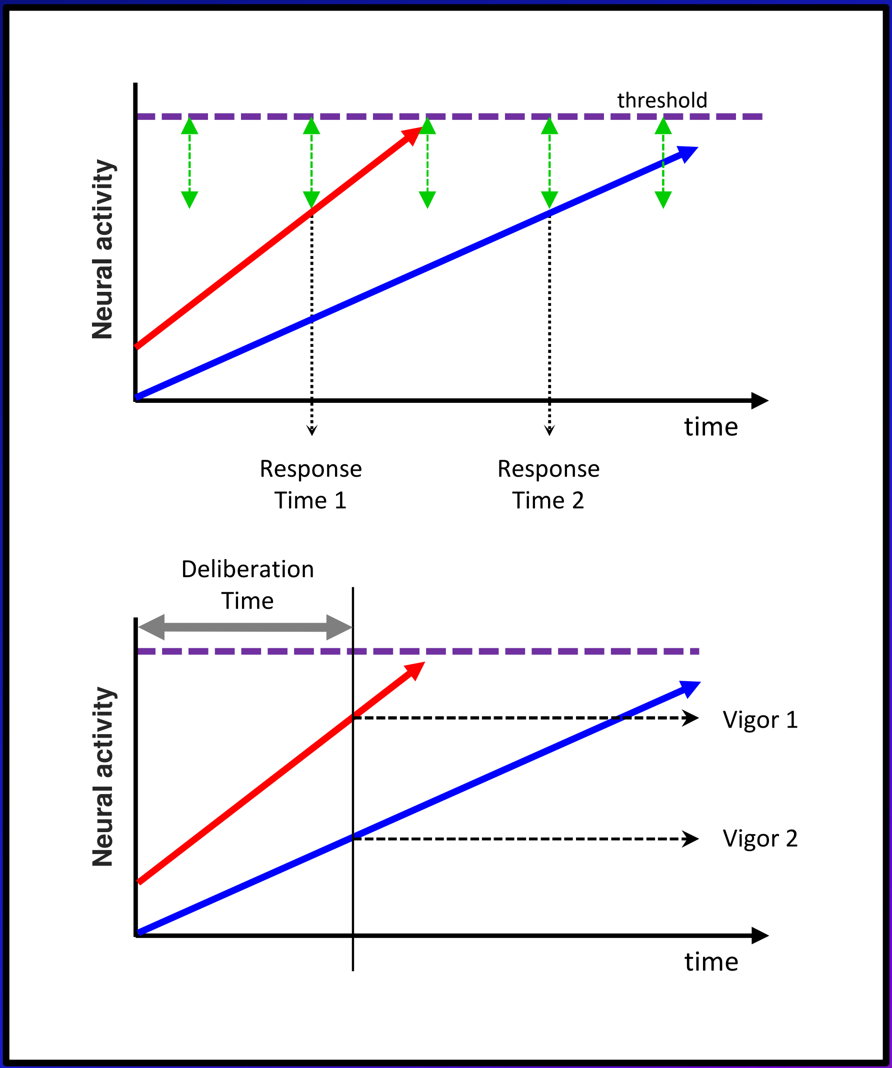

Individual differences

- Those who move with more vigor tend to respond more quickly
  - Jaśkowski et al. 2000; Reppert et al. 2018
- These traits are stable, and can be used to identify individuals
  - Choi et al. 2014; Reppert et al. 2015; Rigas et al. 2016; Bargary et al. 2017; Friedman et al. 2017; Berret et al. 2018
- “Impulsive” individuals
  - Make decisions quickly (Burnett-Heyes et al. 2012; Voon et al. 2014)
  - Have trouble withholding responses (Aichert et al. 2012; Choi et al. 2014; Speiser et al. 2017)
  - Sensation-seeking, prone to boredom (Watt & Vodanovich 1992; Whiteside & Lynam 2001; Berret et al. 2018)
  - Steeper temporal discounting (Shadmehr et al. 2010; Haith et al. 2012; Dalley & Robbins 2017; Summerside et al. 2018)
- “Conservative” individuals

Parkinson's Disease (PD)

- Bradykinesia
  - Slowing, but not an inability to move (Mazzoni et al. 2007)
  - Insufficient modulation of motor system by reward (Pekny et al. 2015)
  - Proposal: Pathologically ***diminished*** urgency
- Therapy (dopamine replacement, DBS)
  - Significantly exaggerated temporal discounting (Housden et al. 2010)
  - Symptoms of trait impulsivity (Djamishidian et al. 2014; Kojovic et al. 2016; Frank et al. 2007)
  - Prolonged treatment can lead to mania and impulse control disorders (Maier et al. 2014; Lopez et al. 2017; Molina et al. 2000; Seedat et al. 2000)
  - Proposal: Pathologically ***elevated*** urgency

Major Depressive Disorder (MDD)

- Affective symptoms
- Non-affective symptoms
  - "Activational" aspects of motivation
    - Low energy, apathy, and fatigue
    - Psychomotor slowing, neurocognitive retardation
    - Less sensitivity to rewards, less willing to exert effort for rewards, experience greater subjective difficulty in producing force
    - High comorbidity with Parkinson's (Koerts et al. 2007; Rana et al. 2015)
  - Not alleviated by SSRI treatments of the affective symptoms of MDD (Stahl 2002; Treadway & Zald 2011; Fava et al. 2014; Gorwood et al. 2014)
  - Instead, they are more responsive to *noradrenaline* and *dopamine* reuptake inhibitors (NDRIs) (Pae et al. 2007; Stahl 2002; Treadway & Zald 2011; Zisook et al. 2006; Demyttenaere et al. 2005; Stahl et al. 2003)
  - Proposal: Pathologically ***diminished*** urgency
    - Absence of these symptoms (intact urgency?) is often a predictor of response to SSRI treatment

A “dimensional” view of urgency

Carland, Thura & Cisek (2019) *The Neuroscientist*

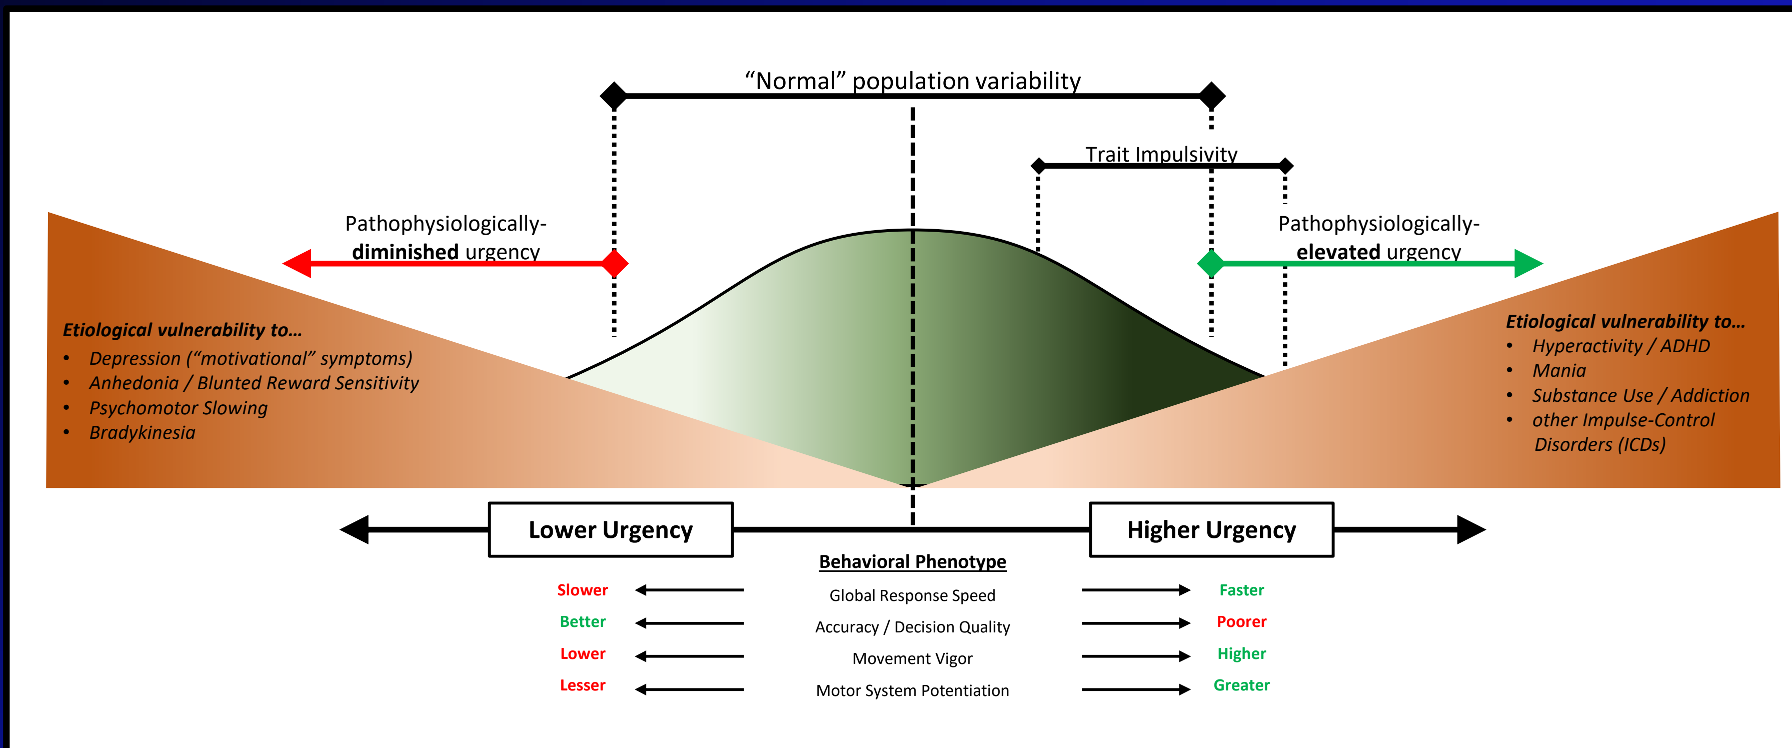

- A continuum of variation in urgency/vigor across individuals
- The extrema correspond to pathologies
- Reward rate → neural mechanisms → psychological phenomena

Conclusions

- Need to move beyond the drift-diffusion model
  - Accumulation of ***novel*** evidence, resembling a low-pass filter with a short time constant
  - ***Urgency*** signal pushes activity to commitment, implementing a decreasing accuracy criterion
- Urgency is an individual trait
  - Too high: impulse control disorders, ADHD, mania
  - Too low: bradykinesia, psychomotor slowing, motivational aspects of depression
- Dynamical system for dynamic decisions
  - Cerebral cortex implements the evolving deliberation
  - Basal ganglia energize the decision, confirm commitment, and invigorate the action
  - Decisions are made through recurrent attractor dynamics

Thank you

"The great end of life is not knowledge but action" T. H. Huxley (1825-1895)

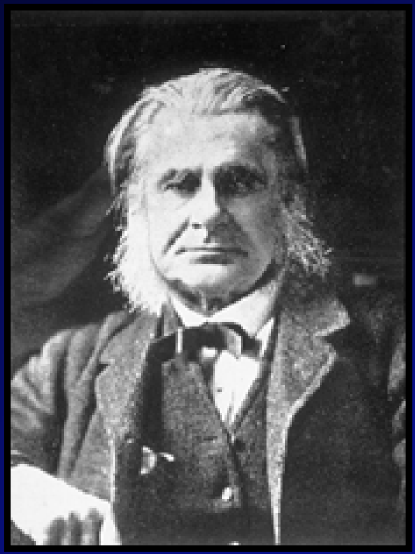

"Your head is there to move you around" R.E.M. (1980-2011)

Current lab members

- Marie-Claude Labonté (technician)
- **Matthew Carland** (PhD student)
- Ayuno Nakahashi (PhD student)
- **Tyler Peel** (postdoc)
- Cesar Canaveral (PhD student)
- **Poune Mirzazadeh** (MSc student)

Alumni

- Jean-Philippe Thivierge (Ottawa)
- Thomas Michelet (Bordeaux)
- Valeriya Gritsenko (UWV)
- **Ignasi Cos** (Barcelona)
- Alex Pastor-Bernier (McGill)
- **David Thura** (Lyon)
- Timothy Meehan
- Julien Michalski (MSc)
- Thomas Lusignan (MSc)

Collaborators

- **John Kalaska** (U Montréal)
- Andrea Green (U Montréal)
- **Karim Jerbi** (U Montréal)
- **Guillaume Lajoie** (U Montréal)
- **Max Puelma-Touzel** (U Montréal)
- **Alain Dagher** (McGill)
- **Julie Duque** (Brussels)
- **Giovanni Pezzulo** (Rome)
- Alexandra Battaglia-Mayer (Rome)
- Aaron Batista (U Pittsburgh)
- Dang Nguyen (U Montréal)

Visiting students / interns

- **Genevieve Aude Puskas**
- Elisabeth Rounis (London)
- **Stephany El-Murr**
- Nicolas Belanger
- **Julie Beauregard-Racine**
- **Charles-William Fradet**
- Farid Medleg
- Elsa Tremblay
- **Encarni Marcos** (Barcelona)
- **Jessica Trung**
- **Jean-François Cabana**
- **Albert Feghaly**
- **Gerard Derosiere** (Brussels)
- **Guido Guberman**
- Philippe Castonguay
- Thomas Lusignan
- Jia Dong Wang
- Sandra Ferland
- Léo Demange-Hamel
- Ikrame Housni
- William Lata
- Tianhui Deng
- Jakob Boulanger

Groupe de recherche sur le système nerveux central GRSNC
CIHR IRSC Canadian Institutes of Health Research Instituts de recherche en santé du Canada
NSERC CRSNG
Fonds de recherche Nature et technologies Québec
Fonds de recherche Santé Québec
Canada Foundation for Innovation Fondation canadienne pour l'innovation
THE EJLB FOUNDATION