Lecture 7: Perceptual decision making

Samuel Gershman

Harvard University

Roadmap

- This lecture analyzes decisions about discrete hidden states based on perceptual evidence.

Roadmap

- This lecture analyzes decisions about discrete hidden states based on perceptual evidence.
- We first formalize the computational problem, and then derive algorithms for solving it.

Roadmap

- This lecture analyzes decisions about discrete hidden states based on perceptual evidence.
- We first formalize the computational problem, and then derive algorithms for solving it.
- The basic algorithmic motif is evidence accumulation to a decision threshold. We look at how this motif can be implemented in neural circuitry, and how it can adaptively optimize reward rate.

Perceptual discrimination problems

- Hidden state corresponds to one of several discrete states: $s \in \{1, \ldots, S\}$, generating a perceptual signal (x), which an agent uses to discriminate the hidden state.

Perceptual discrimination problems

- Hidden state corresponds to one of several discrete states: $s \in \{1, \ldots, S\}$, generating a perceptual signal (x), which an agent uses to discriminate the hidden state.
- Because the evidence provided by the neural encoding of sensory signals is noisy, a single snapshot will in general not suffice to make a good decision. Rather, the optimal algorithm is to accumulate evidence across time until it crosses a decision threshold.

Perceptual discrimination problems

- Simplest case: two hidden states, $s \in \{A, B\}$, sampled with probability $p(s)$ on each trial.

Perceptual discrimination problems

- Simplest case: two hidden states, $s \in \{A, B\}$, sampled with probability $p(s)$ on each trial.
- Agent observes a time series of signals $x(t)$, chooses an action $a \in \{A, B\}$, and receives reward $r = R_s$ if $a = s$ (0 otherwise).

Perceptual discrimination problems

- Simplest case: two hidden states, $s \in \{A, B\}$, sampled with probability $p(s)$ on each trial.
- Agent observes a time series of signals $x(t)$, chooses an action $a \in \{A, B\}$, and receives reward $r = R_s$ if $a = s$ (0 otherwise).
- Two versions of the discrimination problem: (1) *interrogation paradigm*, where a fixed response deadline is imposed on the agent; (2) *free response paradigm*, where the agent chooses when to respond.

Perceptual discrimination problems

- Posterior conditional on history $X(t)$ in log-odds form:

$$\log \frac{p(s = A|X(t))}{p(s = B|X(t))} = \log \frac{p(X(t)|s = A)}{p(X(t)|s = B)} + \log \frac{p(s = A)}{p(s = B)}.$$

Perceptual discrimination problems

- Posterior conditional on history $X(t)$ in log-odds form:

$$\log \frac{p(s = A|X(t))}{p(s = B|X(t))} = \log \frac{p(X(t)|s = A)}{p(X(t)|s = B)} + \log \frac{p(s = A)}{p(s = B)}.$$

- The first term on the right-hand side is the log likelihood ratio, and the second term is the log prior ratio.

Evidence accumulation

- If the sensory signals are independent, the log posterior odds can be written recursively.

Evidence accumulation

- If the sensory signals are independent, the log posterior odds can be written recursively.
- Let $L(t)$ denote the log posterior odds at time t. Initializing to $L(0) = \log \frac{p(s=A)}{p(s=B)}$, the log odds is incremented over a short interval δ according to the momentary evidence supplied by sensory signals:

$$L(t) - L(t - \delta) = \log \frac{p(x(t)|s = A)}{p(x(t)|s = B)}$$

Utility

- Assuming utility is reward, $u(r) = r$, the expected utility of choosing action a at time t is:

$$\mathbb{E}[u(r)|s, a, L(t)] = \begin{cases} \sigma(L(t))R_s, & \text{if } a = A \\ \sigma(-L(t))R_s, & \text{if } a = B \end{cases}$$

where $\sigma(L) = 1/(1 + e^{-L})$ is the logistic sigmoid.

Utility

- Assuming utility is reward, $u(r) = r$, the expected utility of choosing action a at time t is:

$$\mathbb{E}[u(r)|s, a, L(t)] = \begin{cases} \sigma(L(t))R_s, & \text{if } a = A \\ \sigma(-L(t))R_s, & \text{if } a = B \end{cases}$$

  where $\sigma(L) = 1/(1 + e^{-L})$ is the logistic sigmoid.

- Decision variable (log expected utility ratio):

$$U(t) = \log \frac{\mathbb{E}[u(r)|s, a = A, L(t)]}{\mathbb{E}[u(r)|s, a = B, L(t)]} = 2L(t) + \log \frac{R_A}{R_B}$$

Expected utility maximization for the interrogation paradigm

- Expected utility is maximized by choosing

$$a^* = \begin{cases} A, & \text{if } U(t) > 0 \\ B, & \text{if } U(t) < 0 \end{cases}$$

Expected utility maximization for the interrogation paradigm

- Expected utility is maximized by choosing

$$a^* = \begin{cases} A, & \text{if } U(t) > 0 \\ B, & \text{if } U(t) < 0 \end{cases}$$

- This is the Bayes-optimal policy for the interrogation paradigm, where the agent is obligated to make a choice at a particular time.

Expected utility maximization for the free response paradigm

- Agent can potentially increase their reward rate by responding more quickly. However, there is a speed-accuracy trade-off: faster responses will tend to yield less reward because the quality of evidence is lower.

Expected utility maximization for the free response paradigm

- Agent can potentially increase their reward rate by responding more quickly. However, there is a speed-accuracy trade-off: faster responses will tend to yield less reward because the quality of evidence is lower.
- Agent can set decision thresholds U_A and U_B, such that $a = A$ if $U(t) > U_A$ and $a = B$ if $U(t) < U_B$. If no threshold has been crossed, the agent withholds a response.

Expected utility maximization for the free response paradigm

- Agent can potentially increase their reward rate by responding more quickly. However, there is a speed-accuracy trade-off: faster responses will tend to yield less reward because the quality of evidence is lower.
- Agent can set decision thresholds U_A and U_B, such that $a = A$ if $U(t) > U_A$ and $a = B$ if $U(t) < U_B$. If no threshold has been crossed, the agent withholds a response.
- Thresholds determine preference for speed vs. accuracy. When the threshold separation $\Delta U = U_A - U_B$ is larger, the agent will take longer to respond and the responses will be more accurate on average.

The drift-diffusion model

- Special case of Gaussian signals:

$$x(t) \sim \mathcal{N}(\delta\mu_s, \delta\nu)$$

where μ_s is the signal strength and ν is the signal variance.

The drift-diffusion model

- Special case of Gaussian signals:

$$x(t) \sim \mathcal{N}(\delta\mu_s, \delta\nu)$$

  where μ_s is the signal strength and ν is the signal variance.
- Momentary evidence accumulated over δ is also Gaussian distributed.

The drift-diffusion model

- Special case of Gaussian signals:

$$x(t) \sim \mathcal{N}(\delta\mu_s, \delta\nu)$$

  where μ_s is the signal strength and ν is the signal variance.
- Momentary evidence accumulated over δ is also Gaussian distributed.
- In the limit $\delta \to 0$, this becomes the *drift-diffusion model* (DDM), which can be written as a stochastic differential equation:

$$dL = \theta_s dt + \kappa dW,$$

  where dW is the differential of the Wiener process (Brownian motion).

The drift-diffusion model

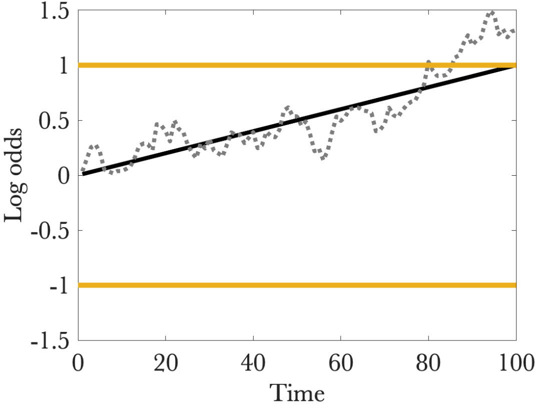

Black line: drift component. Dotted gray line: trajectory of the log odds over time (the sum of the drift and diffusion components). Horizontal lines: decision thresholds.

LIP and MT activity during evidence accumulation and termination

MT: sensory drive (motion energy). LIP: evidence accumulator.

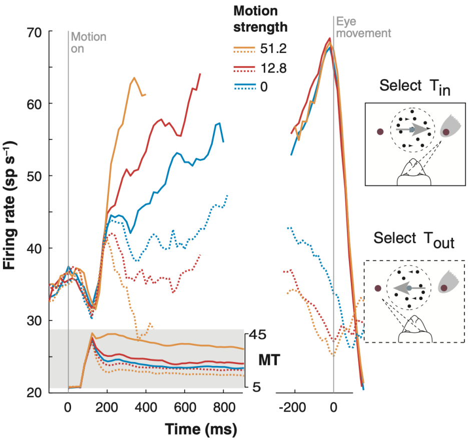

[Britten et al 1992; Roitman et al 2002]

How does LIP know when to stop accumulating evidence?

- Requires feedback from a downstream area registering the eye movement decision.

How does LIP know when to stop accumulating evidence?

- Requires feedback from a downstream area registering the eye movement decision.
- Superior colliculus is a good candidate for providing the feedback signal; sends projections to cortical areas, including LIP.

How does LIP know when to stop accumulating evidence?

- Requires feedback from a downstream area registering the eye movement decision.
- Superior colliculus is a good candidate for providing the feedback signal; sends projections to cortical areas, including LIP.
- Burst neurons fire transiently at high rates immediately before an eye movement [Sparks 1978].

How does LIP know when to stop accumulating evidence?

- Requires feedback from a downstream area registering the eye movement decision.
- Superior colliculus is a good candidate for providing the feedback signal; sends projections to cortical areas, including LIP.
- Burst neurons fire transiently at high rates immediately before an eye movement [Sparks 1978].
- These cells might report threshold-crossing events in LIP (and possibly other cortical areas) and then inhibit the cortical inputs, terminating evidence accumulation [Lo et al 2006].

Superior colliculus can terminate evidence accumulation

- Bursts in superior colliculus are triggered by high firing rates in LIP, and inactivation of superior colliculus causes a delay in decision termination [Stine et al 2023].

Superior colliculus can terminate evidence accumulation

- Bursts in superior colliculus are triggered by high firing rates in LIP, and inactivation of superior colliculus causes a delay in decision termination [Stine et al 2023].
- The inactivation led to paradoxically better performance, because more evidence was accumulated prior to a decision.

Resource constraints on evidence accumulation

- In the interrogation paradigm, it is commonly assumed that subjective decision thresholds don't matter; decision time is determined exogenously by the experimenter.

Resource constraints on evidence accumulation

- In the interrogation paradigm, it is commonly assumed that subjective decision thresholds don't matter; decision time is determined exogenously by the experimenter.
- Implication: performance should approach perfect accuracy as the decision time gets longer, since eventually the accumulated evidence in favor of the correct decision will overpower the diffusion noise.

Resource constraints on evidence accumulation

- In the interrogation paradigm, it is commonly assumed that subjective decision thresholds don't matter; decision time is determined exogenously by the experimenter.
- Implication: performance should approach perfect accuracy as the decision time gets longer, since eventually the accumulated evidence in favor of the correct decision will overpower the diffusion noise.
- In fact, human performance falls short of perfection—why?

Resource constraints on evidence accumulation

- Hypothesis: evidence cannot be accumulated without bound by the brain, because neurons cannot produce arbitrarily high firing rates.

Resource constraints on evidence accumulation

- Hypothesis: evidence cannot be accumulated without bound by the brain, because neurons cannot produce arbitrarily high firing rates.
- Ratcliff [2006] showed that positing "implicit" thresholds provides a good account of human accuracy in the interrogation paradigm.

Resource constraints on evidence accumulation

- Hypothesis: evidence cannot be accumulated without bound by the brain, because neurons cannot produce arbitrarily high firing rates.
- Ratcliff [2006] showed that positing "implicit" thresholds provides a good account of human accuracy in the interrogation paradigm.
- Kiani et al [2008] showed that LIP neurons do not continue to increase their activity until the response signal, but instead saturate at a fixed level regardless of the decision time.

Monkey performance on motion direction discrimination in the interrogation paradigm

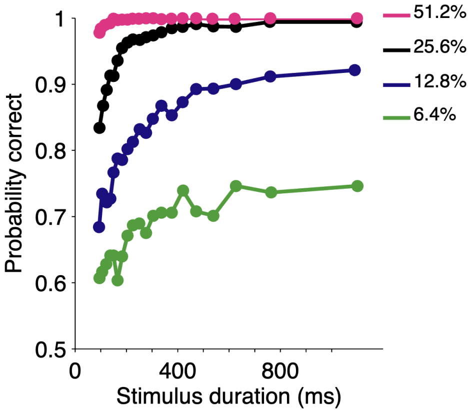

[Kiani et al 2008]

Leaky, competitive dynamics

- The DDM assumes perfect integration—it forgets nothing. Real neurons always have leakage, because membranes are not perfect insulators.

Leaky, competitive dynamics

- The DDM assumes perfect integration—it forgets nothing. Real neurons always have leakage, because membranes are not perfect insulators.
- Decision circuits in the brain have recurrent connectivity: increasing the activity of neurons favoring one decision inhibit neurons favoring a different decision.

Leaky, competitive dynamics

- The DDM assumes perfect integration—it forgets nothing. Real neurons always have leakage, because membranes are not perfect insulators.
- Decision circuits in the brain have recurrent connectivity: increasing the activity of neurons favoring one decision inhibit neurons favoring a different decision.
- Both of these features are incorporated into the leaky competing accumulator (LCA) model [Usher & McClelland 2001].

Leaky, competitive dynamics

- The LCA models the dynamics of two units (populations of neurons), y_1 and y_2, according to:

$$dy_1 = [-\gamma y_1 - \beta g(y_2) + \mu_1]dt + \kappa dW_1$$

$$dy_2 = [-\gamma y_2 - \beta g(y_1) + \mu_2]dt + \kappa dW_2$$

  where γ is the leak rate, β controls the mutual inhibition, $g(\cdot)$ is a static nonlinearity, μ_i is the excitatory drive for population i, κ is the noise standard deviation, and $W_i(t)$ is a Wiener process.

Leaky, competitive dynamics

- The LCA models the dynamics of two units (populations of neurons), y_1 and y_2, according to:

$$dy_1 = [-\gamma y_1 - \beta g(y_2) + \mu_1]dt + \kappa dW_1$$

$$dy_2 = [-\gamma y_2 - \beta g(y_1) + \mu_2]dt + \kappa dW_2$$

  where γ is the leak rate, β controls the mutual inhibition, $g(\cdot)$ is a static nonlinearity, μ_i is the excitatory drive for population i, κ is the noise standard deviation, and $W_i(t)$ is a Wiener process.
- When leak and inhibition are balanced ($\beta = \gamma$), the model is equivalent to the DDM.

What happens when leak and inhibition are not balanced?

- Corresponds to a generalization of the DDM (O-U process).

What happens when leak and inhibition are not balanced?

- Corresponds to a generalization of the DDM (O-U process).
- If leak is stronger than inhibition ($\gamma > \beta$), the O-U process converges to a Gaussian equilibrium distribution.

What happens when leak and inhibition are not balanced?

- Corresponds to a generalization of the DDM (O-U process).
- If leak is stronger than inhibition ($\gamma > \beta$), the O-U process converges to a Gaussian equilibrium distribution.
- If inhibition is stronger than leak ($\beta > \gamma$), there is no stable equilibrium; the mean and variance of the process grow exponentially (but the process still terminates when the threshold is crossed).

What happens when leak and inhibition are not balanced?

- Corresponds to a generalization of the DDM (O-U process).
- If leak is stronger than inhibition ($\gamma > \beta$), the O-U process converges to a Gaussian equilibrium distribution.
- If inhibition is stronger than leak ($\beta > \gamma$), there is no stable equilibrium; the mean and variance of the process grow exponentially (but the process still terminates when the threshold is crossed).
- The leak-dominant regime offers an alternative explanation for suboptimal performance in the interrogation paradigm: if the equilibrium mean is below the decision threshold, then the process relies on noise to cross the threshold. This means that performance will not be perfect even for long decision times.

Leak-dominant vs. inhibition-dominant regimes

- Leak-dominant regime produces "conservative" behavior: the dynamics slow down as the activity approaches the fixed point, which lies close to the correct threshold.

Leak-dominant vs. inhibition-dominant regimes

- Leak-dominant regime produces "conservative" behavior: the dynamics slow down as the activity approaches the fixed point, which lies close to the correct threshold.
- Inhibition-dominant regime produces "risky" behavior: the dynamics speed up as the activity approaches both the correct and incorrect thresholds. This produces faster responses, but also more errors.

Leak-dominant vs. inhibition-dominant regimes

- Leak-dominant regime produces "conservative" behavior: the dynamics slow down as the activity approaches the fixed point, which lies close to the correct threshold.
- Inhibition-dominant regime produces "risky" behavior: the dynamics speed up as the activity approaches both the correct and incorrect thresholds. This produces faster responses, but also more errors.
- Busemeyer & Townsend [1993] suggested that rewards push the imbalance towards inhibition-dominance, whereas punishments push the imbalance towards leak-dominance.

Study question

Compare how the DDM and leaky competing accumulator model account for the speed-accuracy trade-off. What are the distinctive computational and neural mechanisms each emphasizes?

Why is evidence accumulation leaky?

- If the non-leaky DDM achieves optimal performance, then why does evidence accumulation in the brain appear to be (at least sometimes) leaky?

Why is evidence accumulation leaky?

- If the non-leaky DDM achieves optimal performance, then why does evidence accumulation in the brain appear to be (at least sometimes) leaky?
- One answer: this is just an unavoidable property of biology.

Why is evidence accumulation leaky?

- If the non-leaky DDM achieves optimal performance, then why does evidence accumulation in the brain appear to be (at least sometimes) leaky?
- One answer: this is just an unavoidable property of biology.
- BUT: networks of neurons can resist leak through recurrent connectivity, producing "reverberating" activity that outlives the time constant of any individual neurons.

Why is evidence accumulation leaky?

- Another answer: leak is actually useful.

Why is evidence accumulation leaky?

- Another answer: leak is actually useful.
- Consider what happens when the hidden state is not fixed over time but instead switches intermittently. In this case, an optimal decision circuit should "forget" the evidence it accumulated prior to a switch.

Why is evidence accumulation leaky?

- Another answer: leak is actually useful.
- Consider what happens when the hidden state is not fixed over time but instead switches intermittently. In this case, an optimal decision circuit should "forget" the evidence it accumulated prior to a switch.
- In the case where the circuit doesn't have direct access to the switch times, it should forget gradually—precisely what is accomplished by a leaky accumulator.

Why is evidence accumulation leaky?

- Another answer: leak is actually useful.
- Consider what happens when the hidden state is not fixed over time but instead switches intermittently. In this case, an optimal decision circuit should "forget" the evidence it accumulated prior to a switch.
- In the case where the circuit doesn't have direct access to the switch times, it should forget gradually—precisely what is accomplished by a leaky accumulator.
- In support of this interpretation, humans [Glaze et al 2015] and rats [Piet et al 2018] adopt leakier accumulation when the switch rate is higher.

Thresholds that optimize reward rate

- Natural optimization objective for the free response paradigm is the reward rate, defined as the ratio of expected reward to expected trial duration:

$$\bar{r} = \frac{\mathbb{E}[r]}{\mathbb{E}[T] + T_0 + \mathrm{ITI}}$$

where ITI is the intertrial interval (the time between the response and the next trial) and T_0 is non-decision time.

Thresholds that optimize reward rate

- Natural optimization objective for the free response paradigm is the reward rate, defined as the ratio of expected reward to expected trial duration:

  $$\bar{r} = \frac{\mathbb{E}[r]}{\mathbb{E}[T] + T_0 + \mathrm{ITI}}$$

  where ITI is the intertrial interval (the time between the response and the next trial) and T_0 is non-decision time.
- The problem facing the agent is how to set the thresholds to optimize reward rate.

Thresholds that optimize reward rate

- Natural optimization objective for the free response paradigm is the reward rate, defined as the ratio of expected reward to expected trial duration:

$$\bar{r} = \frac{\mathbb{E}[r]}{\mathbb{E}[T] + T_0 + \mathrm{ITI}}$$

  where ITI is the intertrial interval (the time between the response and the next trial) and T_0 is non-decision time.
- The problem facing the agent is how to set the thresholds to optimize reward rate.
- Generally speaking, reward rate is optimized by choosing a higher threshold when the evidence is weaker. Intuitively, this is necessary to avoid making erroneous decisions based on noise.

Thresholds that optimize reward rate

Effect of threshold on reward, decision time, and reward rate

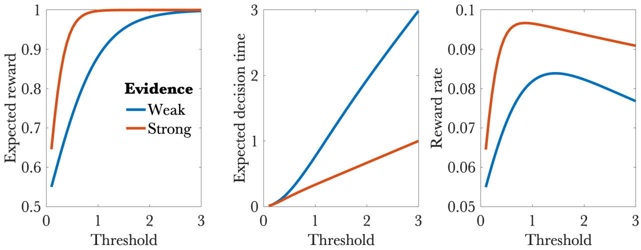

Do humans optimize reward rate?

- To answer this question, it would be nice to have a performance measure that doesn't depend on the choice of parameters (drift rate, noise variance).

Do humans optimize reward rate?

- To answer this question, it would be nice to have a performance measure that doesn't depend on the choice of parameters (drift rate, noise variance).
- Optimal performance curve [Bogacz et al 2006]:

$$\frac{\mathbb{E}[T]}{T_0 + \mathrm{ITI}} = \left[\frac{1}{P_e \log \frac{1-P_e}{P_e}} + \frac{1}{1-2P_e} \right]^{-1}$$

where P_e is the error rate (the proportion of trials on which the wrong decision is made). The quantity on the left-hand-side is the *normalized decision time*.

Do humans optimize reward rate?

- The optimal performance curve predicts that decision times will be fastest when the error rate is very large and when it is very small.

Do humans optimize reward rate?

- The optimal performance curve predicts that decision times will be fastest when the error rate is very large and when it is very small.
- When the evidence is very weak relative to the noise level, the threshold should be set low because the stimulus carries little information about the optimal decision, making it more advantageous to move as quickly as possible to the next trial.

Do humans optimize reward rate?

- The optimal performance curve predicts that decision times will be fastest when the error rate is very large and when it is very small.
- When the evidence is very weak relative to the noise level, the threshold should be set low because the stimulus carries little information about the optimal decision, making it more advantageous to move as quickly as possible to the next trial.
- When the evidence is very strong relative to the noise level, the threshold should again be set low, in this case because the stimulus carries substantial information about the optimal decision—the decision is easy.

Do humans optimize reward rate?

Optimal performance curve compared to human performance.

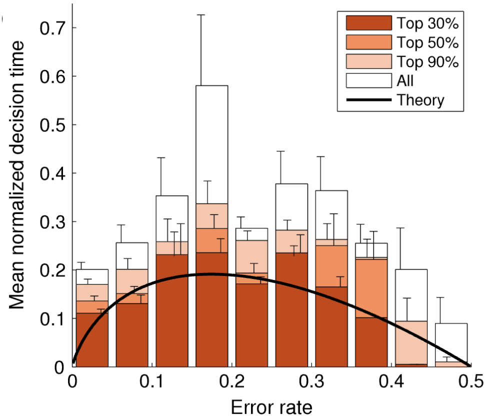

[Bogacz et al 2010]

Why are so many people suboptimal?

- One possibility: optimization process is limited by timing uncertainty.

Why are so many people suboptimal?

- One possibility: optimization process is limited by timing uncertainty.
- Perception of elapsed time is subject to irreducible noise. In the context of reward rate optimization, timing uncertainty means that humans may not have direct access to temporal variables such as the ITI, the non-decision time, and decision time.

Why are so many people suboptimal?

- One possibility: optimization process is limited by timing uncertainty.
- Perception of elapsed time is subject to irreducible noise. In the context of reward rate optimization, timing uncertainty means that humans may not have direct access to temporal variables such as the ITI, the non-decision time, and decision time.
- Zacksenhouse et al [2010] proposed that some subjects adopt a *maximin* strategy, attempting to maximize the worst reward rate given their level of timing uncertainty. This yields a rescaled optimal performance curve: longer decision times compared to reward rate optimization without timing uncertainty.

Why are so many people suboptimal?

- One possibility: optimization process is limited by timing uncertainty.
- Perception of elapsed time is subject to irreducible noise. In the context of reward rate optimization, timing uncertainty means that humans may not have direct access to temporal variables such as the ITI, the non-decision time, and decision time.
- Zacksenhouse et al [2010] proposed that some subjects adopt a *maximin* strategy, attempting to maximize the worst reward rate given their level of timing uncertainty. This yields a rescaled optimal performance curve: longer decision times compared to reward rate optimization without timing uncertainty.
- The maximin strategy does a good job explaining why most subjects tended to have surprisingly long decision times.

Timing uncertainty

- Consistent with the timing uncertainty theory, Balci et al [2011] showed that that individual differences in timing ability (assessed by an independent test) are correlated with deviations from the optimal performance curve.

Timing uncertainty

- Consistent with the timing uncertainty theory, Balci et al [2011] showed that that individual differences in timing ability (assessed by an independent test) are correlated with deviations from the optimal performance curve.
- This suggests that individuals take into account their abilities when setting thresholds.

Threshold optimization

- Threshold optimization could be implemented through an iterative adjustment process, where errors induce increments in the threshold and correct decisions induce decrements in the threshold.

Threshold optimization

- Threshold optimization could be implemented through an iterative adjustment process, where errors induce increments in the threshold and correct decisions induce decrements in the threshold.
- People become increasingly close to the optimal performance curve with practice [Balci et al 2011].

Threshold optimization

- Threshold optimization could be implemented through an iterative adjustment process, where errors induce increments in the threshold and correct decisions induce decrements in the threshold.
- People become increasingly close to the optimal performance curve with practice [Balci et al 2011].
- After an error, people slow down on the next trial, making them more accurate [Rabbitt 1966].

Collapsing thresholds

- So far, we have assumed a fixed threshold. Under certain assumptions, optimal thresholds should decline within a trial.

Collapsing thresholds

- So far, we have assumed a fixed threshold. Under certain assumptions, optimal thresholds should decline within a trial.
- Fudenberg et al [2018] analyze the case of an agent maximizing expected reward penalized by a linear time cost:

$$\pi^* = \underset{\pi}{\text{argmax}}\, \mathbb{E}[r - \alpha T | \pi]$$

where $\alpha > 0$ is a time cost coefficient.

Collapsing thresholds

- So far, we have assumed a fixed threshold. Under certain assumptions, optimal thresholds should decline within a trial.
- Fudenberg et al [2018] analyze the case of an agent maximizing expected reward penalized by a linear time cost:

$$\pi^* = \underset{\pi}{\operatorname{argmax}}\, \mathbb{E}[r - \alpha T | \pi]$$

  where $\alpha > 0$ is a time cost coefficient.
- They derive an asymptotic approximation of the optimal threshold for large t and κ:

$$U^*(t) \approx \frac{1}{2\alpha(\chi^{-1} + \kappa^{-2} t)}$$

Collapsing thresholds

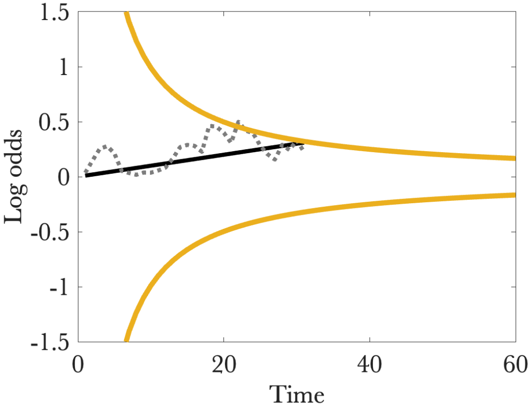

Collapsing thresholds

$$U^*(t) \approx \frac{1}{2\alpha(\chi^{-1} + \kappa^{-2}t)}$$

- The threshold is a hyperbolic function of time (i.e., declining with rate $1/t$), collapsing more quickly when signals are noisier (large κ).

Collapsing thresholds

$$U^*(t) \approx \frac{1}{2\alpha(\chi^{-1} + \kappa^{-2} t)}$$

- The threshold is a hyperbolic function of time (i.e., declining with rate $1/t$), collapsing more quickly when signals are noisier (large κ).
- This fits with the intuition that more difficult trials should induce a lower threshold.

Collapsing thresholds

$$U^*(t) \approx \frac{1}{2\alpha(\chi^{-1} + \kappa^{-2}t)}$$

- The threshold is a hyperbolic function of time (i.e., declining with rate $1/t$), collapsing more quickly when signals are noisier (large κ).
- This fits with the intuition that more difficult trials should induce a lower threshold.
- Thresholds are lower when the time cost is greater (large α), capturing an agent's disinclination to deliberate.

Collapsing thresholds

$$U^*(t) \approx \frac{1}{2\alpha(\chi^{-1} + \kappa^{-2}t)}$$

- The threshold is a hyperbolic function of time (i.e., declining with rate $1/t$), collapsing more quickly when signals are noisier (large κ).
- This fits with the intuition that more difficult trials should induce a lower threshold.
- Thresholds are lower when the time cost is greater (large α), capturing an agent's disinclination to deliberate.
- Thresholds are lower when the agent is more certain *a priori* that the drift rate is close to the prior mean of 0 (small χ), capturing the intuition that more difficult decisions should induce a lower threshold.

Collapsing thresholds

- The empirical data supporting collapsing thresholds is mixed.

Collapsing thresholds

- The empirical data supporting collapsing thresholds is mixed.
- One reason for these mixed results might be that the optimality analysis is more subtle than it first appears. Collapsing thresholds are only optimal when some trials are very difficult; if these trials are absent, the optimal thresholds are constant or can even increase [Malhotra et al 2018].

Collapsing thresholds

- The empirical data supporting collapsing thresholds is mixed.
- One reason for these mixed results might be that the optimality analysis is more subtle than it first appears. Collapsing thresholds are only optimal when some trials are very difficult; if these trials are absent, the optimal thresholds are constant or can even increase [Malhotra et al 2018].
- Even when collapsing thresholds are theoretically optimal, there are many decision problems in which a fixed threshold can achieve a comparable reward rate [Boehm et al 2020].

Challenges

- The study of two-alternative perceptual decisions has fostered the development of highly accurate quantitative models. Nonetheless, it is important to consider some limitations and challenges of this setup.

Challenges

- The study of two-alternative perceptual decisions has fostered the development of highly accurate quantitative models. Nonetheless, it is important to consider some limitations and challenges of this setup.
- Many decision problems involve additional complexities: time-varying hidden states, larger action spaces, state-dependent rewards, and so on.

Challenges

- The study of two-alternative perceptual decisions has fostered the development of highly accurate quantitative models. Nonetheless, it is important to consider some limitations and challenges of this setup.
- Many decision problems involve additional complexities: time-varying hidden states, larger action spaces, state-dependent rewards, and so on.
- Another challenge is evidence against the role of LIP as an evidence accumulator. For example, one study [Katz et al 2016] found that inactivation of LIP surprisingly had no effect on decision making performance. Another study [Latimer et al 2015] argued that many LIP neurons do not ramp up gradually with evidence, but instead exhibit discrete jumps.

Steps vs. ramps

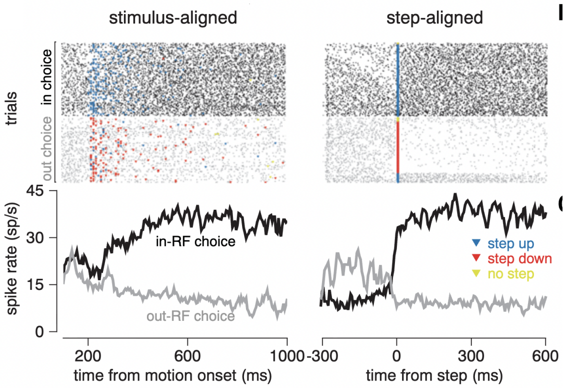

[Latimer et al 2015]

Study question

Some studies found that inactivation of LIP or discrete jumps in neural activity contradict the ramping evidence accumulator view of LIP. How might computational neuroscientists reconcile these discrepancies?