How to cope in a changing environment?

Adrian Radillo[1], Alan Veliz-Cuba[2], Krešimir Josić[1,◕], and Zachary Kilpatrick[3,◕]

1. U. of Houston ; 2. U. of Dayton ; 3. U. of Colorado Boulder. ◕. Equal contrib.
adrian@math.uh.edu ; https://math.uh.edu/~adrian/presentations.html

UNIVERSITY of HOUSTON

Summary

- In the *dynamic clicks task* (Piet, Hady & Brody, 2017) rats have shown the ability to discard old information in trials where the stimulus is **volatile**.
- We develop an **ideal-observer** model that maximizes the expected reward on this task and makes explicit the **evidence-discounting** phenomenon.
- A linear discounting model proves to be as efficient but less robust to parameter tuning.
- It remains unclear what model the rats use and how we may distinguish them based on data.

Debates and discussions are welcome!

The dynamic clicks task

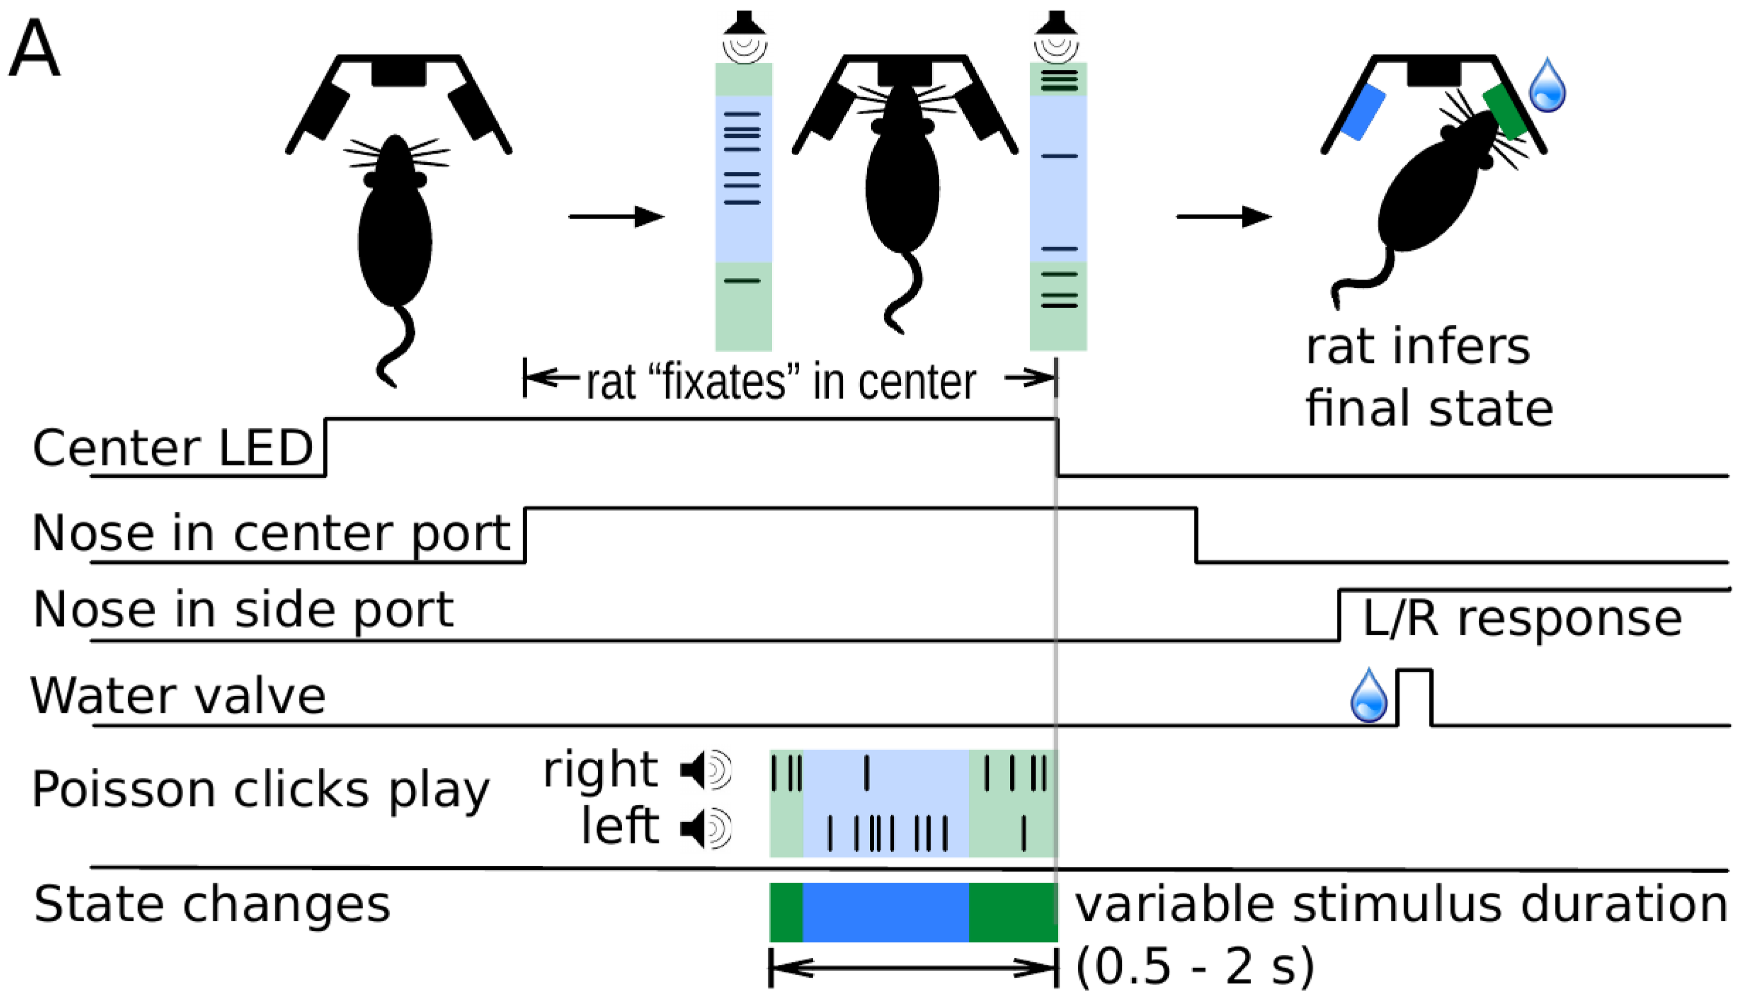

Piet, Hady, & Brody, (2017). Rats optimally accumulate and discount evidence in a dynamic environment. bioRxiv. https://doi.org/10.1101/204248

Key features of the stimulus

- Two coupled inhomogeneous Poisson processes (one per ear).
- Each intensity $\lambda^R(t), \lambda^L(t) \in \{\lambda_{\text{low}}, \lambda_{\text{high}}\}$ alternates stochastically within a trial according to the **hazard rate** h.

Ideal-observer model

$$y_t := \log \frac{P(\lambda_{\text{high}}(t) \text{ on the right} \mid \text{observed streams})}{P(\lambda_{\text{high}}(t) \text{ on the left} \mid \text{observed streams})}$$

$$\frac{dy}{dt} = -2h\sinh y + \log\frac{\lambda_{\text{high}}}{\lambda_{\text{low}}} \sum_{t_{\text{clicks}}} \delta_i^R - \delta_j^L; \quad y_0 = 0; \quad h = \text{hazard rate}$$

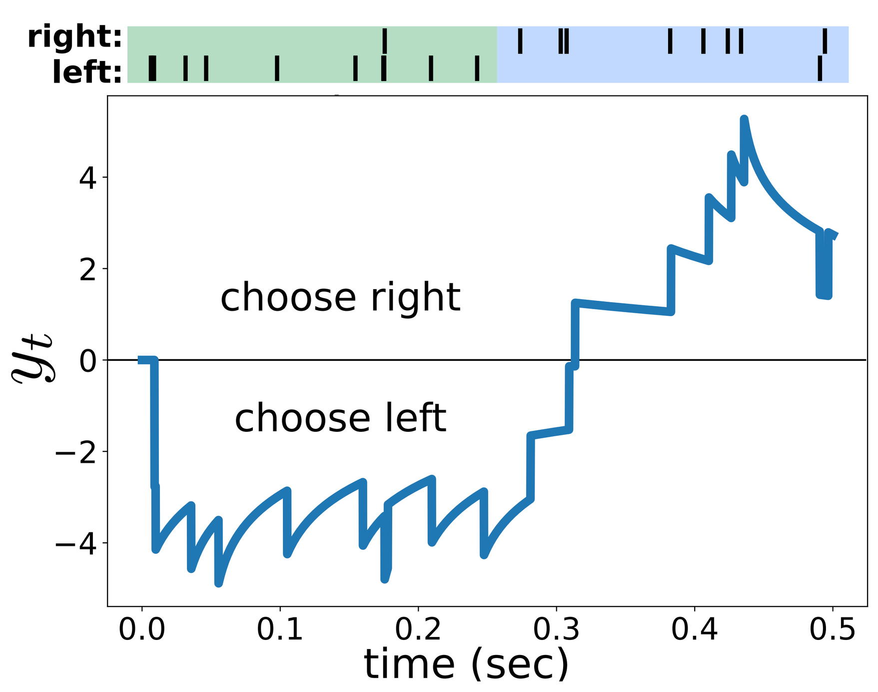

- Discontinuous jumps at clicks, nonlinear decay in-between
- Interrogate at time T and decide based on sign of y_T
- This model maximizes the expected reward on every trial

Natural parameters

- We found numerical evidence that accuracy at time T remains constant as long as $SNR(T)$ and h/S^2 are kept constant:

$$SNR(T) := \sqrt{T} \cdot S$$
$$S := \frac{\lambda_{\text{high}} - \lambda_{\text{low}}}{\sqrt{\lambda_{\text{high}} + \lambda_{\text{low}}}}$$

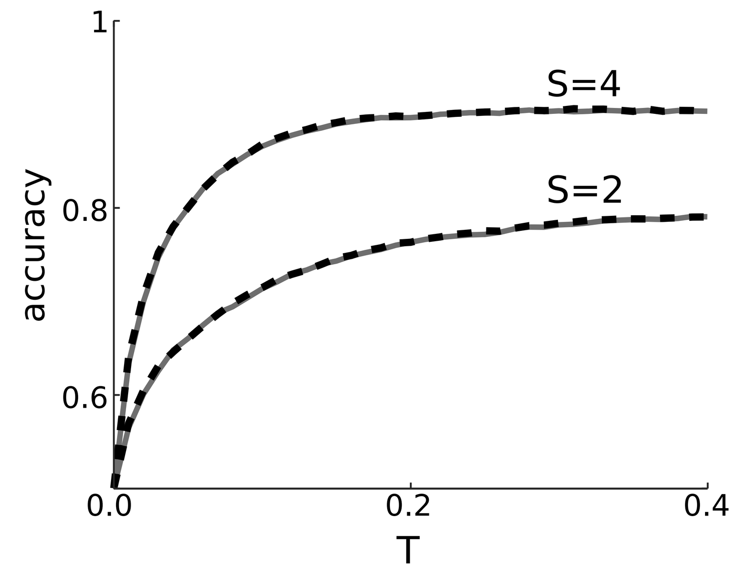

- So our system really is 2-dimensional
- If $hT =$ const. the system becomes 1-dimensional!
- On the right, accuracy varies as a function of S. Black lines are the level curves.

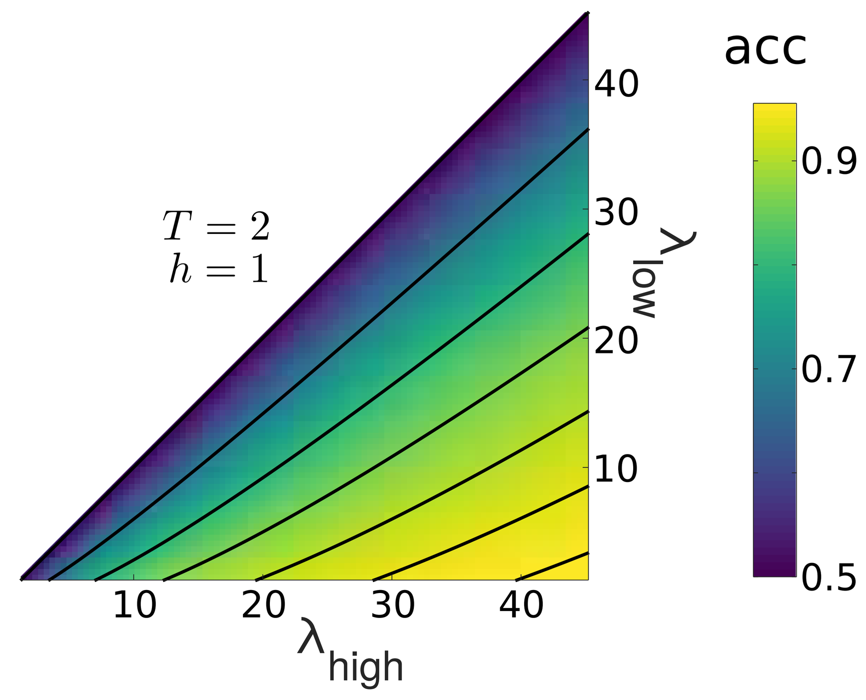

Linear model is near-optimal...

$$\frac{dy}{dt} = -2\gamma \cdot y + \log\frac{\lambda_{\text{high}}}{\lambda_{\text{low}}} \sum_{t_{\text{click}}} \delta_i^R - \delta_j^L; \quad y_0 = 0; \quad \gamma = \text{discounting rate}$$

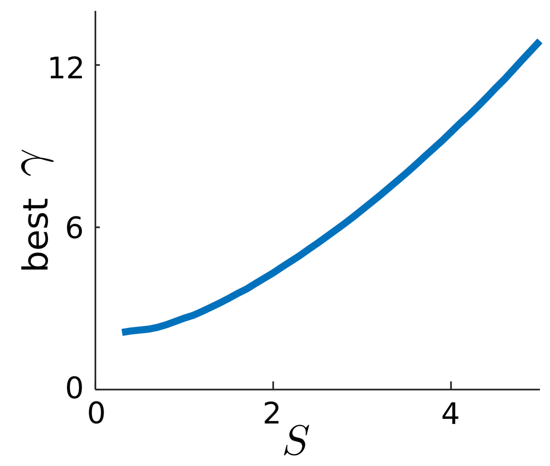

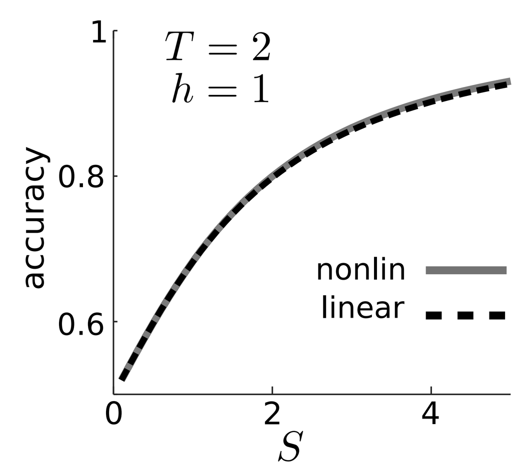

... but more sensitive to errors

- When discounting parameter is misspecified, performance drops faster in the linear model

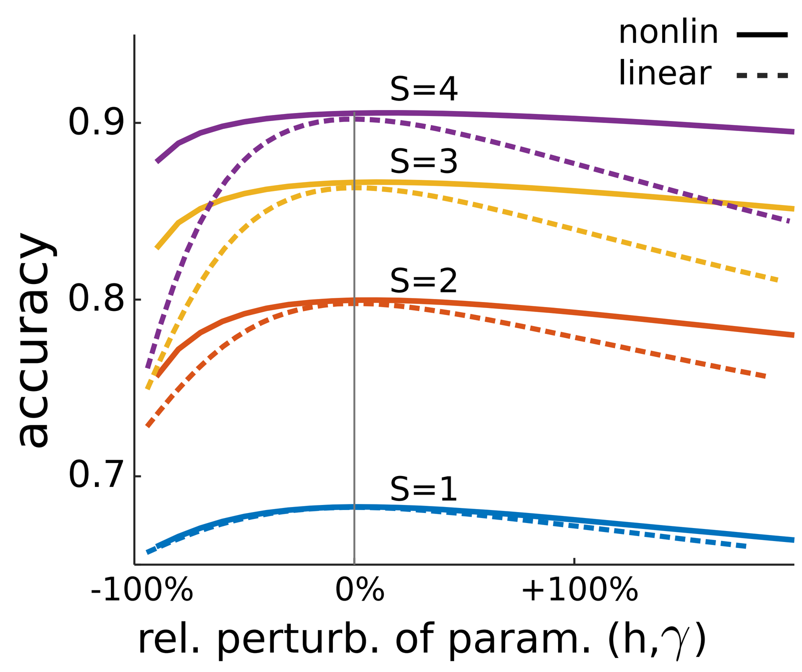

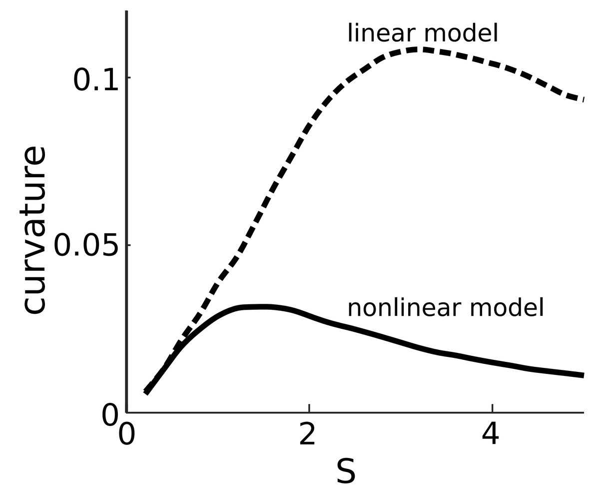

The inverse problem

- Given the stimulus and decision data from our models, how well can we:
  - identify the model?
  - recover its parameters?