

## **Commentary on Rahnev & Denison**

Word Counts: Abstract: 59; Main text: 999; References: 231; Entire 1359:

### **The Standard Bayesian Model is Normatively Invalid for Biological Brains**

Rani Moran

Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University

College London, 10-12 Russell Square, London WC1B 5EH, UK;

Wellcome Centre for Human Neuroimaging, University College London, London WC1N  
3BG, United Kingdom.

Phone: +44 7597700692;

email: [rani.moran@gmail.com](mailto:rani.moran@gmail.com)

URL: <https://www.mps-ucl-centre.mpg.de/en/people/rani-moran>

Konstantinos Tsetsos:

Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-  
Eppendorf, N43, Martinstraße 52, 20246 Hamburg, Germany.

Phone: +49 40741054305;

email: [k.tsetsos62@gmail.com](mailto:k.tsetsos62@gmail.com)

URL: <https://sites.google.com/site/konstantinostsetsos/>

**Abstract:** We show that the benchmark Bayesian framework that R&D used to assess optimality is actually suboptimal under realistic assumptions about how noise corrupts decision-making in biological brains. This model is therefore invalid qua normative standard. We advise against generally forsaking optimality and argue that a biologically constrained definition of optimality could serve as an important driver for scientific progress.

**Main Text:** R&D's extensive review of the perceptual decision-making literature points to the fact that human behavior substantially deviates from optimality. Notably, R&D define optimality according to a standard Bayesian framework (henceforth, the "benchmark model"). Here, we would like to support and elaborate on the claim that this benchmark model has limited validity in describing perceptual decisions. However, rather than addressing the *descriptive* (in)ability of the model to account for behavioral data, we would like to go a step further and postulate that this model has limited *normative* validity once constraints of information-processing in biological brains are taken into account.

Specifically, we argue that the assumptions of the benchmark model are over-simplistic, considering how perceptual processing is implemented in biological brains, and that under more realistic assumptions the benchmark model ceases to be optimal. A key assumption in the benchmark model is that encoding noise is the sole corrupting element of perceptual decisions (Sec. 2.2). Alas, decision-making operations are never limited to stimulus encoding. Rather, they involve a sequential cascade of processes past the encoding stage, such as integration of information across different sources, mental inference, decision-formation, response selection and motor execution. In biological brains, these processes are performed across a hierarchy of cortical layers, which are prone to different sources of noise (Servan-

Schreiber et al., 1990). Thus, perceptual choices are unavoidably corrupted by “late”, post-encoding, noise. This notion was corroborated in a recent study that dissociated the contributions of three noise-sources to decision sub-optimality: sensory-encoding, response selection and mental inference (Drugowitsch et al., 2016). Strikingly, noise in mental-inference, rather than encoding, was found to be the main contributor to sub-optimality.

Here, we argue that in the presence of post-encoding sources of noise, the benchmark model ceases to be optimal. As we next show, higher accuracy can be obtained by down-weighting some aspects of the available information via a “selective integration” process (Tsetsos et al., 2016a; Tsetsos et al., 2016b; See Osmmy et al., 2013 for an illustration of how a different form of information down-weighting can facilitate signal detection performance).

To illustrate the superiority of the selective integration model, we consider a simple binary choice scenario, as in R&D’s Fig. 1. We assumed that the measurement distributions of these two stimuli are  $N(\pm 1, 0.7^2)$ , that both stimuli are presented with equal prior probabilities and that observers are awarded or penalized one point for each correct or erroneous choice, respectively. Critically, unlike R&D, we assumed that predicted action costs calculated by observers are prone to an additional source of late Gaussian noise denoted by  $N(0, \xi^2)$ . Note that this late noise affects only the costs that observers predict but not the actual rewards and penalties they receive ( $\pm 1$ ). We assumed that late noise is independent across actions and independent of the encoding noise. Using R&D notations, the predicted cost of each action is:

$$cost_{predicted}(a) = \sum_{s \in S} [l(s|x, \theta)\pi(s)\mathcal{L}(s, a)] + N(0, \xi^2)$$

According to the benchmark model, the chosen action on each trial minimizes predicted cost:

$$\delta(x) = \operatorname{argmin}_{a \in A} \{cost_{predicted}(a)\}$$

Unsurprisingly, the actually realized cost of the model-selected action increases as a function of late noise,  $\xi$  (Fig. 1). Focal to our interest, however, is the comparison between the benchmark and the selective integration models. In the selective integration model, cost predictions are based on ‘pseudo-likelihoods’  $l'$  defined as follows:

$$l'(x|s_i, \theta) = \begin{cases} l(x|s_i, \theta), & \text{if } l(x|s_i, \theta) \geq l(x|s_j, \theta) \\ l(x|s_i, \theta)w, & \text{if } l(x|s_i, \theta) < l(x|s_j, \theta) \end{cases}$$

,where  $w < 1$  (in Fig. 1,  $w = 0.1$ ). Note that the pseudo-likelihood equals the standard likelihood for the more likely of the two stimuli. However, the likelihood of the less likely stimulus is selectively down-weighted. Action-selection is determined by minimization of the predicted pseudo-costs:

$$cost'_{predicted}(a) = \sum_{s \in S} [l'(s|x, \theta)\pi(s)\mathcal{L}(s, a)] + N(0, \xi^2)$$

$$\delta'(x) = \operatorname{argmin}_{a \in A} \{cost'_{predicted}(a)\}$$

By down-weighting the likelihood of the less likely stimulus, observers attenuate the harmful influences that late noise exerts on performance (see Tsetsos et al., 2016a for a detailed discussion). Indeed, the selective integration model strikingly outperforms the benchmark model by achieving lower actual costs (Fig. 1; negative costs correspond to positive gains). Therefore, the benchmark model is suboptimal when post-encoding noise is present and is thus an inadequate standard for assessing optimality in behavior. We recently showed that an intriguing violation of rational choice theory, intransitivity of choices, was a byproduct of adaptive selective integration processes (Tsetsos et al., 2016a). Future research should investigate whether and which of the sundry ‘‘suboptimal’’ behaviors reviewed by R&D actually reflect biologically constrained adaptive processes.




Figure 1. Comparison between the Bayesian “benchmark” and the selective integration models. Average actual cost is displayed as a function of late noise ( $\xi$ ). See [https://osf.io/gexrd/?view\\_only=aa02df150be94beebf310f1e56cec16f](https://osf.io/gexrd/?view_only=aa02df150be94beebf310f1e56cec16f) for simulation code.

Next, we wish to qualify R&D’s arguments about the overall utility of assessing optimality (Sec. 4). While we agree that addressing optimality should not be in itself the ultimate *goal* of the study of perception, we still think that the notion of optimality can serve as an important driver to scientific progress. From an evolutionary perspective, perceptual processing reflects an extended adaptation process and as such, is *ex ante* expected to be optimal (Moran, 2015). Therefore, when theories of optimality are defined appropriately, taking into account the biological constraints of human information processing, they can provide an invaluable benchmark, guiding both theory development and behavioral assessments. On the one hand, when organisms are found to behave optimally, it raises questions about how optimality is achieved at the algorithmic and implementation levels. When, on the other hand, suboptimal

behaviors are found, it raises questions pertaining to why these alleged deviations from optimality occur. Such questions can lead to a better understating of the constraints and limitations of human information processing (Tsetsos et al., 2016a), of the cost function that neural systems strive to minimize (Soltani et al., 2012), and of the statistical structure of the environment in which cognitive processes evolved (Fawcett et al., 2014). This understating often leads to a subtler definition of optimality, which supports novel behavioral hypotheses and assessments. In sum, we conceive of scientific progress as an iterative process, in which the notion of optimality, rather than being relinquished, continually evolves and undergoes refinement. It is challenging to define optimality but we nevertheless think that such attempts are instrumental in scientific progress.

## References

- Drugowitsch, J., Wyart, V., Devauchelle, A. D. & Koehlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, 92, 1398-1411.
- Fawcett, T. W., Fallenstein, B., Higginson, A. D., Houston, A. I., Mallpress, D. E., Trimmer, P. C., & McNamara, J. M. (2014). The evolution of decision rules in complex environments. *Trends in cognitive sciences*, 18(3), 153-161.
- Moran, R. (2015). Optimal decision making in heterogeneous and biased environment. *Psychonomic Bulletin & Review*, 22, 38-53.
- Osmmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., & Donner, T. (2013). The time scale of perceptual evidence integration can be adapted to the environment. *Current Biology*, 23, 981 – 986.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science*, 249(4971), 892-895.
- Soltani, A., De Martino, B., & Camerer, C. (2012). A range-normalization model of context-dependent choice: a new model and evidence. *PLoS computational biology*, 8(7), e1002607.
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield C. (2016a). Economic irrationality is optimal during noisy decision-making. *Proceedings of the National Academy of Sciences*, 113, 3102–3107.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield C. (2016b). Reply to Davis-Stober et al.: Violations of rationality in Tsetsos et al. (2016) are not aggregation artifacts. *Proceedings of the National Academy of Sciences* 113(33), E4764–4766.