

Input-dependent noise can explain magnitude-sensitivity in optimal value-based decision-making

Angelo Pirrone* Andreagiovanni Reina† Fernand Gobet‡

Abstract

Recent work has derived the optimal policy for two-alternative value-based decisions, in which decision-makers compare the subjective expected reward of two alternatives. Under specific task assumptions — such as linear utility, linear cost of time and constant processing noise — the optimal policy is implemented by a diffusion process in which parallel decision thresholds collapse over time as a function of prior knowledge about average reward across trials. This policy predicts that the decision dynamics of each trial are dominated by the difference in value between alternatives and are insensitive to the magnitude of the alternatives (i.e., their summed values). This prediction clashes with empirical evidence showing magnitude-sensitivity even in the case of equal alternatives, and with ecologically plausible accounts of decision making. Previous work has shown that relaxing assumptions about linear utility or linear time cost can give rise to optimal magnitude-sensitive policies. Here we question the assumption of constant processing noise, in favour of input-dependent noise. The neurally plausible assumption of input-dependent noise during evidence accumulation has received strong support from previous experimental and modelling work. We show that including input-dependent noise in the evidence accumulation process results in a magnitude-sensitive optimal policy for value-based decision-making, even in the case of a linear utility function and a linear cost of time, for both single (i.e., isolated) choices and sequences of choices in which decision-makers maximise reward rate. Compared to explanations that rely on non-linear utility functions and/or non-linear cost of time, our proposed account of magnitude-sensitive optimal decision-making provides a parsimonious explanation that bridges the gap between various task assumptions and between various types of decision making.

Keywords: value-based decision-making, optimality, noise, magnitude-sensitivity

*Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, London, UK. <https://orcid.org/0000-0001-5984-7853>. Email: a.pirrone@lse.ac.uk

†IRIDIA, Université Libre de Bruxelles, Belgium, and Department of Computer Science, University of Sheffield, Sheffield, UK. <https://orcid.org/0000-0003-4745-992X>. Email: andreagiovanni.reina@ulb.be

‡Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, London, UK. <https://orcid.org/0000-0002-9317-6886>. Email: f.gobet@lse.ac.uk

1 Introduction

In order to understand how decision making has evolved, it is crucial to understand what are the optimal policies (i.e., algorithms, behaviours) for decision making under different scenarios (Marshall, 2019; Pirrone et al., 2014; Bogacz et al., 2006). A common working hypothesis is that decision-making systems have evolved to approximate, through robust policies, optimal strategies for cost minimisation and reward maximisation across tasks and domains given the centrality of these factors for survival and reproduction (Pirrone et al., 2014; Tajima et al., 2016; Marshall, 2019; Bogacz et al., 2006).

Extensive work (Bogacz et al., 2006) has addressed the question of optimality with regard to accuracy-based choices — that is, choices for which there is a correct response. For decisions with two alternatives, and under specific constraints (for details see Bogacz et al., 2006; Moran, 2015), such choices are optimised by the well-known drift diffusion model in which agents integrate difference in evidence until a decision threshold for one of two alternatives is reached (Ratcliff & McKoon, 2008; Ratcliff et al., 2016).

Seminal work from Tajima et al. (2016) has focused instead on deriving the optimal policy for value-based choices. With value-based choices, participants are rewarded by the value of the chosen alternative, regardless of whether it is the best option available. The classical example for this type of choices is that of food choices — compared to accuracy-based scenarios, for food choices there is no ‘accurate’ choice. It is particularly important to study value-based choices because most naturalistic decisions are value-based (Pirrone et al., 2014). Even so-called ‘perceptual decisions’ are made in order to maximise reward or minimise loss such as, for example, avoiding an obstacle or detecting a prey.

Surprisingly, the optimal policy for value-based choices derived in Tajima et al. (2016) shows striking similarities to the optimal choice for accuracy-based choices (Bogacz et al., 2006; Tajima et al., 2016). Under specific task assumptions — such as linear utility, linear cost of time and constant processing noise — the optimal policy is implemented by a diffusion process in which parallel decision thresholds collapse over time as a function of prior knowledge about average reward across trials (Tajima et al., 2016). This mechanism ensures maximisation of the expected reward by having boundaries in highly rewarding environments collapsing faster than in low rewarding environments. ‘Parallel collapsing boundaries’ (see Figure 1 for an example) affect the amount of difference between alternatives that is needed to trigger a decision (Hawkins et al., 2015). In particular, the difference between alternatives that would trigger a decision decreases with time, so that less evidence

The authors declare that there is no conflict of interest regarding the publication of this article. The MATLAB code used for the simulations presented in this study is available at <https://github.com/joefresna/Optimal-policy-for-value-based-decision-making-with-value-sensitive-noise>. We thank James Marshall for helpful discussions. A.P. and F.G. acknowledge funding from the European Research Council (ERC-ADG-835002—GEMS). A.R. acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Chargé de Recherches.

Copyright: © 2021. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

that one alternative is superior to the other is needed to make a decision at late stages of evidence accumulation.

As discussed in detail in Marshall (2019), Pirrone et al. (2018b) and Steverson et al. (2019), one feature that characterises the optimal policy with linear subjective utility proposed by Tajima et al. (2016) is that single trial dynamics are magnitude-insensitive. The reason for this is straightforward: a purely relative decision process, in which difference between alternatives is integrated, cannot discriminate between conditions of different magnitude but with the same difference — even with the addition of parallel collapsing boundaries. This rationale is exemplified by the equal alternative case: an alternative pair of 2 vs 2 (low value) and an alternative pair of 8 vs 8 (high value) have both the same difference (null) and are indistinguishable for a purely relative model that processes only difference between alternatives. Even with collapsing boundaries, decisions among equal alternatives would, on average, be made in the same time.

Magnitude-sensitivity (Pirrone et al., 2014) refers to a value-maximising strategy in which small differences in accuracy between high-valued alternatives are disregarded in favour of a quick choice. This strategy has been deemed evolutionary advantageous in order to maximise speed-value trade-offs that characterise value-based decisions (Pirrone et al., 2014, 2018a).

Magnitude-sensitivity — faster choices as the magnitude of the alternatives increases — has been observed empirically in a number of studies and for different organisms, from unicellular organisms making food choices to humans and non-human primates involved in economic decision-making (Pais et al., 2013; Pirrone et al., 2018a,b; Bose et al., 2017; Teodorescu et al., 2016; Reina et al., 2017; Ratcliff et al., 2018; Dussutour et al., 2019; Steverson et al., 2019; Hunt et al., 2012; Kvam & Pleskac, 2016; Smith & Krajbich, 2019; Marshall et al., 2021). Magnitude-sensitivity has been observed even in the limit case of equal alternatives; compared to low but equally valued alternatives, agents show faster reaction times for high but equally valued alternatives (Pirrone et al., 2018a,b). For example, in choosing between rewards, monkeys are faster in choosing between two equally high rewards than two equally poor rewards (Pirrone et al., 2018a). Similarly, humans show faster reaction times as the value of equal alternatives increases in a typical value-based experiment in which participants have to choose between images of food that they had previously rated (Smith & Krajbich, 2019). Surprisingly, even unicellular organisms exhibit magnitude-sensitivity, being faster in reaching one of two equally-high than one of two equally-low valued food sources (Dussutour et al., 2019).

Tajima et al. (2016) have shown that if the assumption of linear subjective utility is relaxed in favour of non-linear subjective utility, the optimal policy for value-based decisions is implemented by non-parallel collapsing decision boundaries. In this case, choices for high-magnitude equal alternatives are made faster compared to choices for low-magnitude equal alternatives; that is, non-linear subjective utility can give rise to magnitude-sensitivity. However, given the widely documented result of magnitude-sensitivity, and the

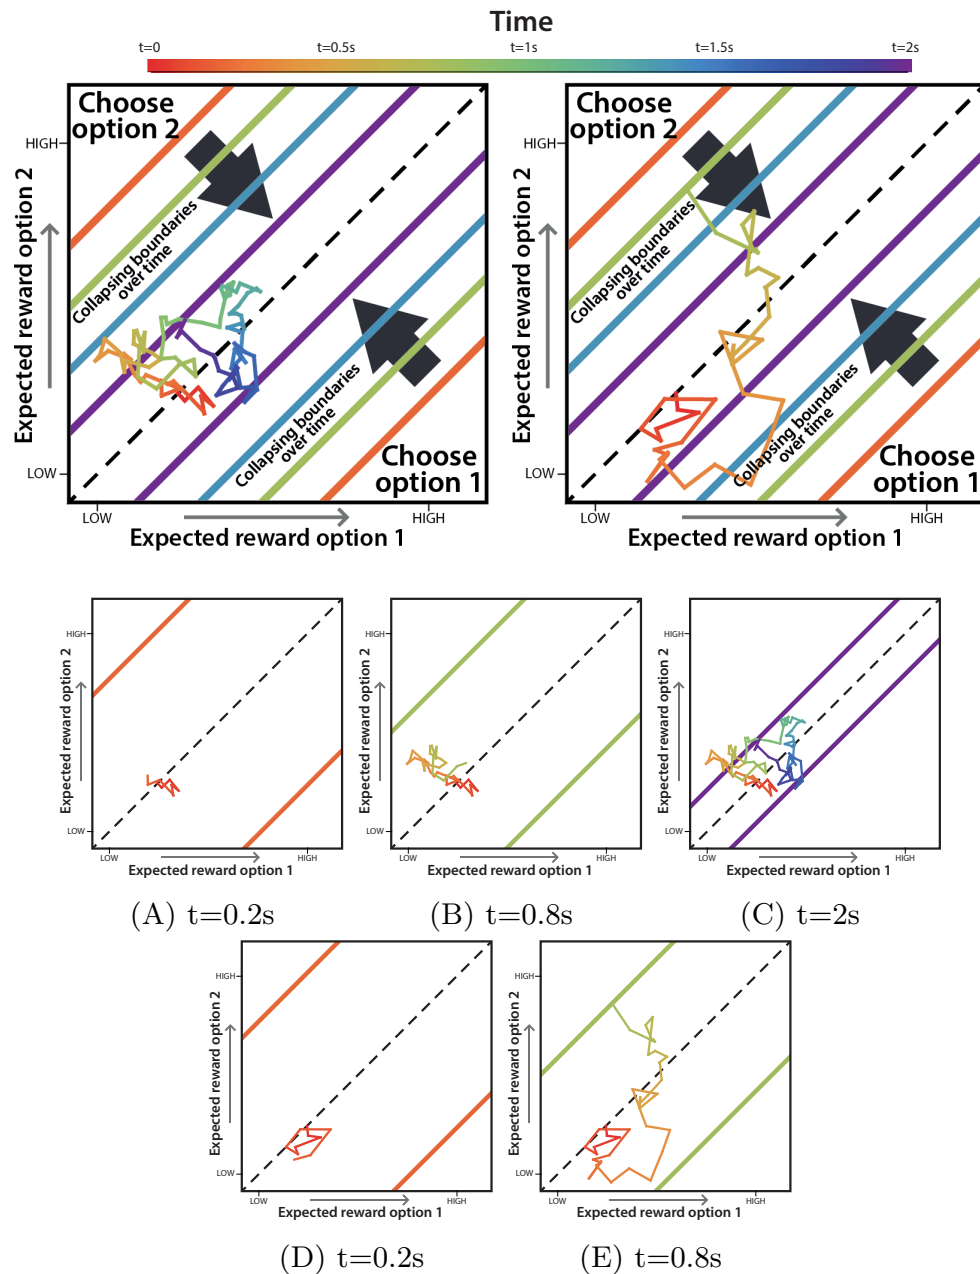


FIGURE 1: Optimal policy for binary value-based decision-making with input-dependent noise. The policy determines when an optimal decision-maker should choose an option: decision-makers continue to accumulate evidence until a decision boundary is reached and a decision is made. In the top row, the two panels show two representative sampling trajectories for equal alternatives with low (left) and high (right) magnitude conditions. The panels below show the time course for the low magnitude condition, in (A) to (C), and for the high magnitude condition, in (D) and (E). Both trajectories and collapsing boundaries are colour-coded, representing time (top legend). With input-dependent noise, the size of the random fluctuations varies with the input magnitude, therefore the high-magnitude conditions have on average larger fluctuations that hit a decision boundary faster compared to the low-magnitude conditions (0.8 s, compared to 2 s). In the absence of input-dependent noise, low and high-magnitude conditions would be indistinguishable and reach a boundary in the same time, exhibiting magnitude-insensitivity.

theoretical arguments supporting why it is expected for *optimal* decision-making, Tajima et al.'s (2016) model has been modified in order to account for magnitude-sensitivity in the linear utility case. One line of research has questioned Tajima et al.'s (2016) assumption of linear cost of time in favour of an ecologically plausible non-linear cost of time of future rewards (Steverson et al., 2019; Marshall, 2019; Marshall et al., 2021); in this case, magnitude-sensitivity is observed even with linear utility functions. However, it remains to be understood if and how a non-linear cost of time could explain magnitude-sensitivity in tasks in which reward is either fixed, non-delayed or even absent (Pirrone et al., 2018a,b; Teodorescu et al., 2016; Smith & Krajbich, 2019; Ratcliff et al., 2018).

Here, building on previous strong empirical and theoretical evidence (Brunton et al., 2013; Teodorescu et al., 2016; Ratcliff et al., 2018; Lu & Doshier, 2008; Louie et al., 2013; Geisler, 1989), we investigate whether magnitude-sensitive noise in the accumulation of evidence could give rise to magnitude-sensitive optimal decision-making. In other words, we question the assumption of constant processing noise made by Tajima et al. (2016).

Extensive work supports the hypothesis that input-dependent noise is neurally-plausible (Albrecht & Geisler, 1991; Albrecht & Hamilton, 1982; Bonds, 1991; Derrington & Lennie, 1984; Heeger, 1993; Kaplan & Shapley, 1982; Ohzawa et al., 1982; Sclar et al., 1990), and there is evidence that during evidence accumulation in both humans and rats, input-dependent noise plays a dominant role, while constant processing noise is null (Brunton et al., 2013). Hence, we want to stress that input-dependent noise is not *just* a technical ad-hoc assumption made in order to accommodate magnitude-sensitivity, but it is instead a principled account of evidence accumulation that warrants further investigation. Here, we report theoretical evidence that input-dependent noise is one of the key candidate explanations for magnitude-sensitivity, as previously suggested by Teodorescu et al. (2016), Ratcliff et al. (2018) and Bose et al. (2020). Our approach is in contrast with how noise is parametrised in computational models of choice (Ratcliff & McKoon, 2008; Usher & McClelland, 2001; Bogacz et al., 2006; Brown & Heathcote, 2008), where input-dependent noise is absent and only constant processing noise affects the decision-making process. Our approach is instead in line with influential work by Lu & Doshier (2008), who have shown that including input-dependent noise in models of human perception is necessary in order to satisfactorily explain empirical data. Including input-dependent noise in the accumulation of evidence does not necessarily predict that the optimal policy should be magnitude-sensitive; this needs to be investigated with mathematical simulations and cannot be claimed a priori as there is not a simple, direct correspondence between evidence accumulation dynamics and the optimal policy.

Investigating the consequences for optimal decision-making when input-dependent noise is added to decision process was done by modifying the code made available by Tajima et al. (2016) from their pioneering study. In the next section we report the technical details of our simulation, and in the final section we discuss the implications of our results for decision making research.

2 Methods and Results

Through numerical simulations, we investigate the effect of magnitude-sensitive noise on binary decision-making. We follow the same assumptions of the value-based decision-making framework described by Tajima et al. (2016). The decision-maker must choose between two alternatives with potentially different rewards, r_1 and r_2 (e.g., nutritional or monetary value). The rewards are unknown to the decision-maker, who acquires through observation some momentary evidence $dr_{i,t} \sim \mathcal{N}(r_i dt, \Gamma(r_1, r_2) dt)$ for both options $i \in \{1, 2\}$ simultaneously, in repeated small time steps of duration $dt \ll 1$. Momentary evidence is sampled from a normal distribution with mean proportional to the true reward value and its variance representing ambiguity, due to both exogenous and endogenous noise, that in line with previous work (Teodorescu et al., 2016; Bose et al., 2020), we model as an input-dependent function, which reads as

$$\Gamma(r_1, r_2) = \sigma^2 + \Phi(r_1^2 + r_2^2), \quad (1)$$

where the parameters σ and Φ are the strength of input-independent and input-dependent noise, respectively (Teodorescu et al., 2016; Bose et al., 2020). Therefore, for $\Phi = 0$, evidence integration has constant noise only, while for $\Phi > 0$, we can observe the effect of magnitude-sensitive noise.

Our decision-maker, at the beginning of a trial, has equal prior expectations for both alternatives, that we model as normally distributed prior beliefs $\mathcal{N}(\mu_\pi, \sigma_\pi^2)$. We assume that prior expectation is the same for both options. According to Bayesian theory, after time t , the posterior mean, or expected reward $\hat{r}_1(t)$, is:

$$\hat{r}_i(t) = \frac{\mu_\pi \sigma^2 + \sigma_\pi^2 \sum_{\tau \in t} dx_{i,\tau}}{\sigma^2 + t \sigma_\pi^2}, \quad (2)$$

where $\sum_{\tau \in t} dx_{i,\tau}$ is the sum of evidence for option i , with $i \in \{1, 2\}$, at time $\tau \in \{dt, 2dt, \dots, t\}$. The decision-makers also incurs a decision cost $c = 0.1$ per temporal unit taken to make the decision. Therefore, when making a decision for option i at time t , the decision-maker receives the reward r_i reduced by the temporal cost ct (for example, the energy or cognitive cost invested in integrating evidence).

In order to maximise reward and minimise cost, the decision-maker updates over time the expected rewards, $\hat{r}_1(t)$ and $\hat{r}_2(t)$, until the integrated evidence has reduced ambiguity sufficiently enough to determine reliably which option has the higher expected reward.

We test both the case of single decisions and of sequential decisions. In the latter, we assume a constant waiting time between decisions $t_w = 1$, thus the total temporal cost is $ct + t_w$, and the decision-maker aims to maximise the reward rate.

Tajima et al. (2016) showed that, in both single and sequential decision-making, through dynamic programming and the Bellman's equation it is possible to compute the optimal policy, which consists in sampling new information until the difference of the expected

rewards, $x(t) = \hat{r}_1(t) - \hat{r}_2(t)$, is larger than a threshold $z(t)$ that decreases over time (collapsing boundaries), i.e. $x(t) \geq z(t)$ or $x(t) \leq -z(t)$. Note that in Tajima et al. (2016), and in our current work, the collapsing boundaries are not a preexisting assumption; the collapsing boundaries are *derived* (i.e., found) as part of the optimal policy. Once the threshold is reached, the decision-maker chooses the alternative with the highest expected reward: $\max(\hat{r}_1(t), \hat{r}_2(t))$. This optimal policy can be implemented by the drift diffusion model (Ratcliff & McKoon, 2008; Ratcliff et al., 2016) with collapsing boundaries. The drift diffusion model is composed by two terms that describe the momentary change of $x(t)$ as

$$dx = (r_1 - r_2)dt + \sqrt{\Gamma(r_1, r_2)} dW(t), \quad (3)$$

where dW is the increment of a normally distributed Wiener process, $dW \sim \mathcal{N}(0, 1)$.

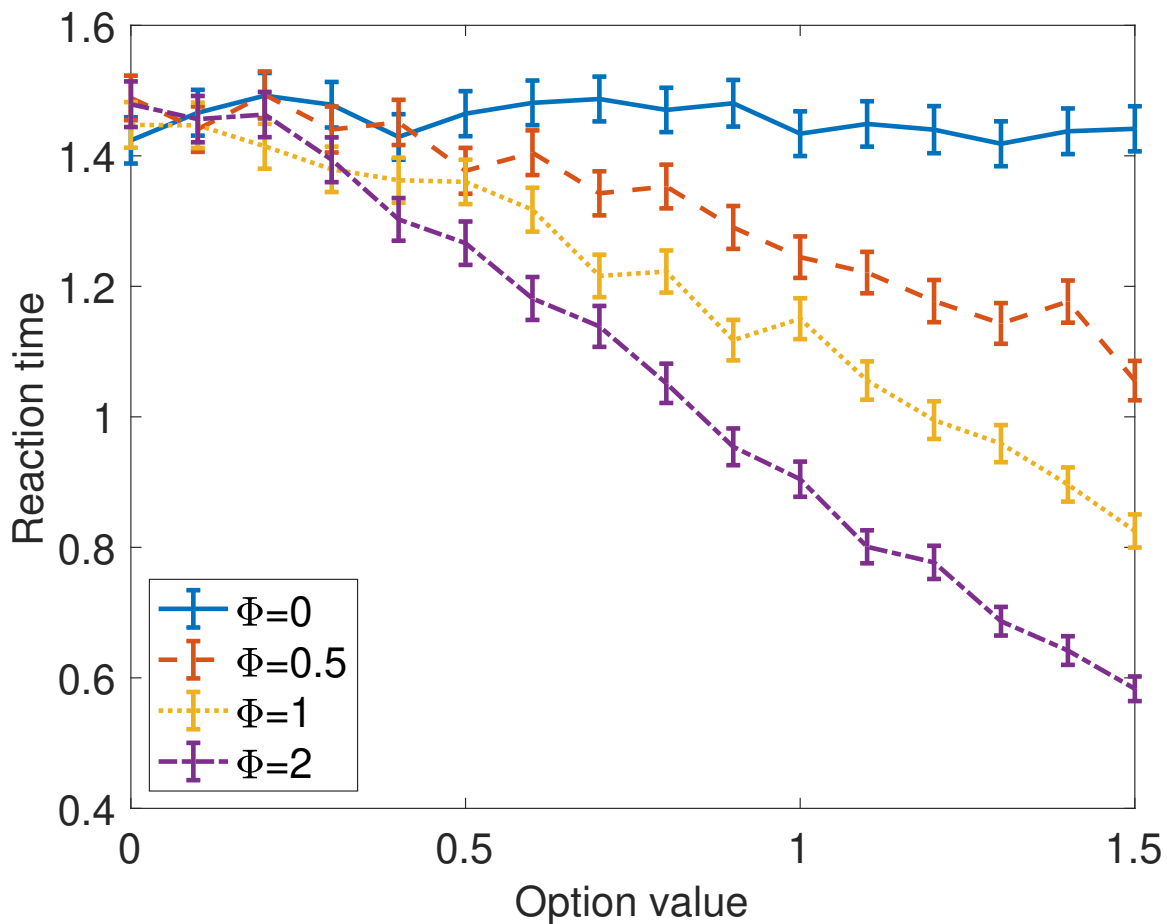


FIGURE 2: Results from stochastic simulations for a single choice: input-dependent noise can explain magnitude-sensitive optimal policies. Φ quantifies the strength of the input-dependent noise. The figure shows mean reaction time as a function of the magnitude of equal alternatives (the bars are 95% confidence intervals). When $\Phi=0$, the magnitude-insensitive optimal policy is derived (Tajima et al., 2016). This figure shows magnitude-sensitive optimal reaction times for a single choice (i.e., expected reward for each individual choice is maximised) as a function of input-dependent noise and magnitude of the stimuli.

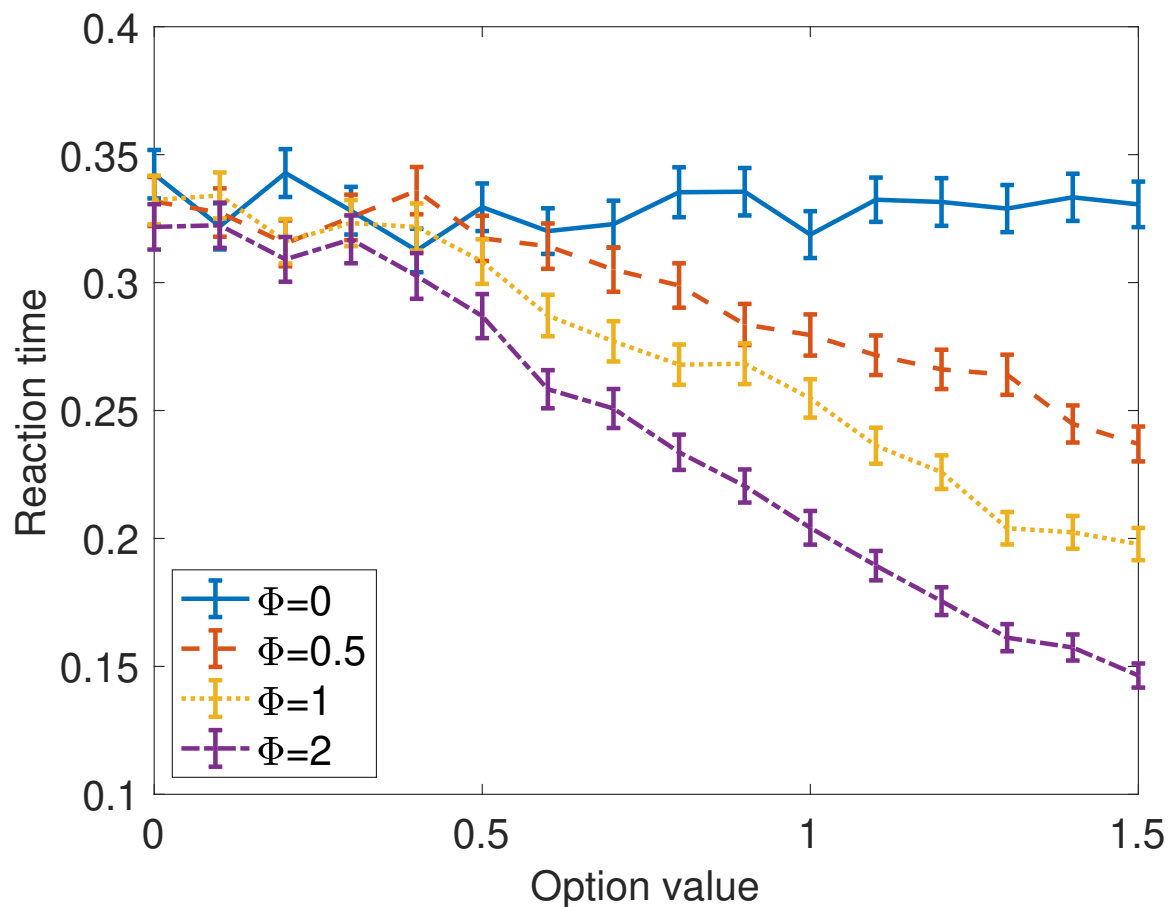


FIGURE 3: Results from stochastic simulations for a sequence of choices: input-dependent noise can explain magnitude-sensitive optimal policies. This figure shows magnitude-sensitive optimal reaction times for a sequence of choices (i.e., total expected reward within a fixed time period is maximised) as a function of input-dependent noise and magnitude of the stimuli.

Figure 1 shows how the threshold $\pm z(t)$ moves over time in the bidimensional space of the two expected rewards, $\hat{r}_1(t)$ and $\hat{r}_2(t)$. In graphical representations of the drift diffusion model, the x-axis generally represents time and the y-axis represents difference in evidence (or value) between the alternatives. In this case the collapsing boundaries are parallel to the x-axis and orthogonal to the y-axis. However, in the case of the optimal strategy for value-based decisions, it is easier to communicate interesting decision dynamics in terms of a rotated space in which the two axes represent the value of each alternative and the boundaries are parallel to the diagonal with unitary slope in the 2-dimensional reward space, as in Figure 1. The rotation of axes does not change the interpretation of decision dynamics in any way; it only simplifies the graphical representation of the optimal decision policy.

The two boundaries are parallel to each other with unity slopes, separating the space into three regions. When the expected difference between the rewards, $x(t)$, exceeds the threshold $\pm z(t)$ (top-left and bottom-right regions of the plots of Figure 1), the decision

is made in favour of the highest expected reward; instead, when the difference is not large enough (central region), the decision-maker chooses to accumulate further evidence. As the policy depends only on the difference between rewards, it is insensitive to the overall magnitude of the alternatives ($r_1 + r_2$), therefore choices for equal alternatives with low and high magnitude have the same decision time (see also Steverson et al., 2019; Marshall, 2019).

We simulated decisions for equal quality alternatives (i.e., $r_1 = r_2 = r$) where we varied only their magnitude $r \in \{0, 0.1, 0.2, \dots, 1.5\}$. We computed the optimal thresholds $\pm z(t)$ using the code that Satoshi Tajima shared with us (code that was further modified by James A.R. Marshall, and is available on GitHub¹), from his 2016 paper. Figure 2 shows the average reaction time for 10^3 simulations in each condition with time step length $dt = 0.01$, prior mean $\mu_\pi = 0$, and prior variance $\sigma_\pi^2 = 5$. We can see that when $\Phi = 0$, the noise is input-independent, constant to a fixed value $\sigma^2 = 2$, and in turn the reaction time is also constant. This result is in agreement with previous analyses (Tajima et al., 2016; Steverson et al., 2019; Marshall, 2019). Instead, when $\Phi > 0$, we can appreciate a decrease in the reaction time with increasing magnitude. As Φ increases, value-sensitivity is more evident. This effect is qualitatively similar for both single and sequential decisions, as results show in Figures 2 and Figure 3, respectively.

Note that input-dependent noise predicts faster and less ‘accurate’ responses, meaning that accuracy over near-equal high-magnitude alternatives is sacrificed in favour of a fast response. This pattern was observed empirically (Teodorescu et al., 2016; Ratcliff et al., 2018) and in simulation-based studies (Bose et al., 2020). Overall, this is a key prediction of any magnitude-sensitive mechanism (Pirrone et al., 2014, 2018a,b; Teodorescu et al., 2016; Kirkpatrick et al., 2021; Steverson et al., 2019; Marshall et al., 2021; Marshall, 2019). However, in our study, in line with previous investigation of magnitude sensitivity (Pirrone et al., 2014, 2018a; Dussutour et al., 2019), we focus exclusively on equal alternatives; that is, in each trial, the two alternatives are identical. Equal alternatives allow to appreciate magnitude effects in the absence of confounds introduced by maintaining differences between unequal alternatives constant while increasing their magnitude (Teodorescu et al., 2016; Ratcliff et al., 2018; Smith & Krajbich, 2019). As such, our simulations and results are based on reaction times alone since it is not possible to define accuracy in a choice between equal alternatives.

3 Discussion

Our work investigates the repercussions for optimal value-based decision-making if an input-dependent noise component is added to the decision making process. Input-dependent noise has received ample support (Brunton et al., 2013; Teodorescu et al., 2016; Ratcliff et al., 2018; Lu & Doshier, 2008; Louie et al., 2013), with experimental and modelling

¹<https://github.com/joefresna/Optimal-policy-for-value-based-decision-making-with-value-sensitive-noise>

work showing that, during evidence accumulation, *the dominant* source of noise is input-dependent. This contrasts with classical drift diffusion models (Ratcliff & McKoon, 2008), in which only a source of constant processing noise is *assumed*. It is important to highlight that input-dependent noise per se does not assume or predict *optimal magnitude-sensitivity* — there is no a priori relationship between the two. In this paper, we have established through numerical simulations that the optimal policy for value-based decision-making, which was derived with input-dependent noise, gives rise to magnitude-sensitivity. In the optimal policy, boundaries are still parallel; however, the noise makes the signal fluctuate more for high-magnitude conditions compared to low-magnitude conditions. In the case of equal alternatives, the boundaries are hit only through noise, and therefore higher noise makes the accumulated evidence (which is on average null) fluctuate more and hit a random boundary quicker than when lower noise is applied. Interestingly, while input-dependent noise accounts for magnitude-sensitivity with parallel boundaries, all other magnitude-sensitive optimal accounts (i.e., non-linear utility, non-linear cost of time) predict instead that magnitude-sensitivity arises as a function of non-parallel collapsing boundaries (Tajima et al., 2016; Marshall, 2019; Steverson et al., 2019). While there is evidence that in some cases decisions are best described by parallel collapsing boundaries (Milosavljevic et al., 2010; Palestro et al., 2018; Hawkins et al., 2015), there is no empirical evidence for non-parallel collapsing boundaries in decision making, as predicted by the non-linear utility and cost of time accounts.

Input-dependent noise enriches the modelling account of decision making by including a neurally plausible assumption (Brunton et al., 2013; Lu & Doshier, 2008; Teodorescu et al., 2016). Furthermore, previous studies have demonstrated that input-dependent noise increases goodness of fit (Teodorescu et al., 2016; Bose et al., 2020; Ratcliff et al., 2018) compared to some competing accounts (e.g., the leaky competing accumulator model, race models, the canonical drift diffusion model; see Teodorescu et al. 2016; Bose et al. 2020; Ratcliff et al. 2018; but also see Kirkpatrick et al. 2021). Moreover, input-dependent noise is a feature that could allow magnitude-sensitivity, and hence the maximisation of reward, across various types of decision making and tasks. This latter aspect — magnitude-sensitivity *across* tasks and domains — makes input-dependent noise a particularly attractive account for magnitude-sensitivity: while explanations of magnitude-sensitive reaction times based on non-linear utility and/or cost of time could be applied ad-hoc to a number of cases, there are numerous scenarios in which the decision-making problem faced by agents may be better described by linear utility and linear cost of time – for example in tasks in which reward is fixed and there is no penalty for a wrong response. Theoretically, we believe that the assumption of linear cost of time and linear subject utility are a reasonable first hypothesis to be explored before considering non-linear functions.

The hypothesis of input-dependent noise addresses all problems discussed above: input-dependent noise is based on strong empirical data and applies to *any* task, regardless of the nature of the stimuli, the number of alternatives, the specific loss function, the utility

function and/or the subject's utility. In fact, regardless of whether it is endogenous or exogenous, noise characterises virtually all decision-making problems, regardless of their specific details. Hence, we believe that input-dependent noise could provide a theoretically parsimonious explanation of descriptive and optimal magnitude-sensitive decision-making.

Interestingly, we show that both single choices and sequence of choices (i.e., the policy maximising reward of a sequence of trials) are magnitude-sensitive with input-dependent noise. This result is in line with the observed results of magnitude-sensitivity that characterises decision-making from unicellular organisms (Dussutour et al., 2019) to monkeys (Pirrone et al., 2018a) and humans across a variety of tasks — both in perceptual and value-based choices (Pais et al., 2013; Pirrone et al., 2018b; Bose et al., 2017; Teodorescu et al., 2016; Ratcliff et al., 2018; Steverson et al., 2019; Hunt et al., 2012; Kvam & Pleskac, 2016; Smith & Krajbich, 2019; Kirkpatrick et al., 2021) and for both single trials and sequence of choices.

However, it is important to mention that the quantitative predictions of optimal decision-making with input-dependent noise have not yet been compared to those of non-linear utility and non-linear cost of time accounts, and this is a timely question for future research that should aim at selecting the best candidate. Furthermore, future empirical studies should investigate the extent to which participants are able to adjust decision boundaries in order to approach optimality as predicted by numerical simulations.

Overall, our contribution enriches Tajima et al.'s (2016) work; we believe that future research could benefit from a similar approach in which, building on Tajima et al.'s (2016) work (and code), assumptions are relaxed in order to account for ecological and naturalistic decision-making.

References

- Albrecht, D. G. & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience*, 7(6), 531–546.
- Albrecht, D. G. & Hamilton, D. B. (1982). Striate cortex of monkey and cat: contrast response function. *Journal of Neurophysiology*, 48(1), 217–237.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700–765.
- Bonds, A. (1991). Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Visual Neuroscience*, 6(3), 239–255.
- Bose, T., Pirrone, A., Reina, A., & Marshall, J. A. R. (2020). Comparison of magnitude-sensitive sequential sampling models in a simulation-based study. *Journal of Mathematical Psychology*, 94, 102298.
- Bose, T., Reina, A., & Marshall, J. A. R. (2017). Collective Decision-Making. *Current Opinion in Behavioral Sciences*, 6, 30–34.

- Brown, S. D. & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3), 153–178.
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, 340(6128), 95–98.
- Derrington, A. & Lennie, P. (1984). Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *The Journal of Physiology*, 357(1), 219–240.
- Dussutour, A., Ma, Q., & Sumpter, D. (2019). Phenotypic variability predicts decision accuracy in unicellular organisms. *Proceedings of the Royal Society B*, 286(1896), 20182825.
- Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, 96(2), 267–314.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6), 2476–2484.
- Heeger, D. J. (1993). Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. *Journal of Neurophysiology*, 70(5), 1885–1898.
- Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F., & Behrens, T. E. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, 15(3), 470–476.
- Kaplan, E. & Shapley, R. (1982). X and Y cells in the lateral geniculate nucleus of macaque monkeys. *The Journal of Physiology*, 330(1), 125–143.
- Kirkpatrick, R. P., Turner, B. M., & Sederberg, P. B. (2021). Equal evidence perceptual tasks suggest a key role for interactive competition in decision-making. *Psychological Review*, <https://doi.org/10.1037/rev0000284>.
- Kvam, P. D. & Pleskac, T. J. (2016). Strength and weight: The determinants of choice and confidence. *Cognition*, 152, 170–180.
- Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, 110(15), 6139–6144.
- Lu, Z.-L. & Doshier, B. A. (2008). Characterizing observers using external noise and observer models: assessing internal representations with external noise. *Psychological Review*, 115(1), 44–82.
- Marshall, J. A. R. (2019). Comment on ‘Optimal Policy for Multi-Alternative Decisions’. *bioRxiv*, <https://doi.org/10.1101/2019.12.18.880872>.
- Marshall, J. A. R., Reina, A., & Pirrone, A. (2021). Magnitude-sensitive reaction times reveal non-linear time costs in multi-alternative decision-making. *bioRxiv*, <https://doi.org/10.1101/2021.05.05.442775>.
- Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, 5(6), 437–449.

- Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, 22(1), 38–53.
- Ohzawa, I., Sclar, G., & Freeman, R. (1982). Contrast gain control in the cat visual cortex. *Nature*, 298(5871), 266–268.
- Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., & Marshall, J. A. R. (2013). A mechanism for value-sensitive decision-making. *PloS One*, 8(9).
- Palestro, J. J., Weichart, E., Sederberg, P. B., & Turner, B. M. (2018). Some task demands induce collapsing bounds: Evidence from a behavioral analysis. *Psychonomic Bulletin & Review*, 25(4), 1225–1248.
- Pirrone, A., Azab, H., Hayden, B. Y., Stafford, T., & Marshall, J. A. R. (2018a). Evidence for the speed–value trade-off: Human and monkey decision making is magnitude sensitive. *Decision*, 5(2), 129–142.
- Pirrone, A., Stafford, T., & Marshall, J. A. R. (2014). When natural selection should optimize speed-accuracy trade-offs. *Frontiers in Neuroscience*, 8, 73.
- Pirrone, A., Wen, W., & Li, S. (2018b). Single-trial dynamics explain magnitude sensitive decision making. *BMC Neuroscience*, 19(1), 1–10.
- Ratcliff, R. & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4), 260–281.
- Ratcliff, R., Voskuilen, C., & Teodorescu, A. (2018). Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. *Cognitive Psychology*, 103, 1–22.
- Reina, A., Marshall, J. A. R., Trianni, V., & Bose, T. (2017). Model of the best-of-N nest-site selection process in honeybees. *Physical Review E*, 95(5), 052411.
- Sclar, G., Maunsell, J. H., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, 30(1), 1–10.
- Smith, S. M. & Krajbich, I. (2019). Gaze amplifies value in decision making. *Psychological Science*, 30(1), 116–128.
- Steverson, K., Chung, H.-K., Zimmermann, J., Louie, K., & Glimcher, P. (2019). Sensitivity of reaction time to the magnitude of rewards reveals the cost-structure of time. *Scientific Reports*, 9(1), 1–14.
- Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature Communications*, 7(1), 1–12.
- Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute: violations of value invariance in human decision making. *Psychonomic Bulletin & Review*, 23(1), 22–38.
- Usher, M. & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592.