Magnitude-sensitivity: rethinking decision-making

Angelo Pirrone[a,*], Andreagiovanni Reina[b], Tom Stafford[c], James A.R. Marshall[d], Fernand Gobet[a]

[a] *Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science, UK*
[b] *Institute for Interdisciplinary Studies on Artificial Intelligence (IRIDIA), Université Libre de Bruxelles, Belgium*
[c] *Department of Psychology, University of Sheffield, UK*
[d] *Department of Computer Science, University of Sheffield, UK*

Abstract

Magnitude-sensitivity refers to the result that performance in decision-making, across domains and organisms, is affected by the total value of the possible alternatives. This simple result offers a window onto fundamental issues in decision-making and has led to a reconsideration of ecological decision-making, prominent computational models of decision-making and optimal decision-making. Moreover, magnitude-sensitivity has inspired the design of new robotic systems that exploit natural solutions and apply optimal decision-making policies. In this article, we review the key theoretical and empirical results about magnitude-sensitivity and highlight the importance that this phenomenon has for the understanding of decision-making. Furthermore, we discuss open questions and ideas for future research.

Keywords: decision-making; magnitude-sensitivity; ecological decision-making; optimality; computational models

[*] Correspondence: a.pirrone@lse.ac.uk (Angelo Pirrone).

PREPRINT: please refer to the publisher's version if you intend to cite the work. There may be minor (stylistic) differences between this document and the published version.

Highlights

- Performance in decision-making is affected by the overall magnitude (i.e., summed values) of the alternatives; this result is known as 'magnitude-sensitivity'.
- Magnitude-sensitivity is observed for all types of decision-making (e.g., perceptual, value-based, economic, and collective) and for different organisms (e.g., unicellular organisms, honeybee swarms, monkeys, and humans).
- Magnitude-sensitivity allows organisms to optimise speed-value trade-offs for naturalistic decisions, in which a fast choice for high-magnitude stimuli is often preferred over a slow but correct one.
- Prominent accounts of theoretical, descriptive and normative decision-making need to be revisited in order to account for magnitude-sensitivity.
- Magnitude-sensitivity has inspired the design of decentralised algorithms for robot swarms that efficiently manage speed-value trade-offs.

Size matters

What do unicellular organisms making food choices [1], humans making perceptual decisions [2, 3], monkeys making economic decisions [3], honeybees choosing between nest sites [4] or robot swarms reaching an agreement [5, 6] have in common?

An answer to this question (among many other possibilities; [see 7]) is that in all those cases, decision-making is affected by the '**magnitude**' (see Glossary) of the alternatives. 'Magnitude' is defined as the summed value of the alternatives. If two food sources have arbitrary values of 3 and 5, the magnitude is simply $3 + 5 = 8$. In the case of a perceptual decision such as deciding which of two stimuli is brighter, the magnitude is the overall brightness of the alternatives, so that two stimuli having brightness of 30 cd/m^3 and 60 cd/m^3 have a magnitude of $30 + 60 = 90$ cd/m^3.

'**Magnitude-sensitivity**' (often also referred to as 'value-sensitivity') refers to the result that performance in decision-making is affected by the magnitude of the alternatives [1, 3, 8–21]. In particular, for higher magnitude conditions, decision-makers show responses that are faster and more likely to be random (i.e., more likely to be incorrect when there is a correct alternative) [3, 11, 22].

An interesting empirical result is that of magnitude-sensitivity with zero evidence [1, 3, 9, 17]: a choice between **equal alternatives** of high magnitude is made faster compared with a choice between equal alternatives of low magnitude. In this case, analyses are often based on reaction times alone

since choice over equal alternatives is generally uninformative. The focus on 'equal alternatives' is of particular interest, as it allows experimentalists to estimate the effect of magnitude-sensitivity in the absence of confounds due to psychophysical transformations.

Empirical evidence of magnitude-sensitivity with equal alternatives has been shown, for example, for human participants who are faster in deciding which of two equal stimuli is brighter when the magnitude of the alternatives increases [3], and in monkeys who are faster in choosing between two equally high-value rewards than between two equally low-value rewards [3]. A particularly evocative result is that even an unicellular, aneural organism exhibits magnitude-sensitivity [1] (Box 1).

A 'simple' behavioural finding such as magnitude-sensitivity has led to the reconsideration of key aspects of decision-making. This is because leading theories and prominent accounts of both descriptive and normative decision-making were *magnitude-insensitive*. That is, they did not take into account, or could not account for, the result of magnitude-sensitivity, and thus made incorrect predictions. In fact, major accounts of decision-making almost exclusively focus on the difference, or ratio in some cases, between alternatives [e.g., 23, 24]; in doing so, information about magnitude does not enter the decision process. As explained in detail in the next section, this is done in an attempt to study decision-making under the lens of statistical optimality [24], in which the difference between alternatives alone drives decision-making.

Here, we present in detail the key theoretical and empirical results about

magnitude-sensitivity and highlight the importance that this phenomenon has for ecological decision-making, computational models of decision-making, optimal accounts of decision-making, swarm cognition (Box 2), and the implementation of autonomous decision-making systems in robotics (Box 3). Given the relevance of magnitude-sensitivity to several disciplines, scientific progress in understanding the causes, consequences, mechanisms, and implications of such a phenomenon can be achieved through a combined interdisciplinary effort. Throughout the article, we discuss ideas for future interdisciplinary research.

Ecological decision-making and the speed-value trade-off

Psychology and neuroscience have often studied binary decision-making in the laboratory under the lens of the **speed-accuracy trade-off** [24–27], in which participants are required to make a 'correct' choice in the shortest time possible or, alternatively, are required to be as accurate as possible within a predefined time window. The statistically optimal strategy, which optimises the speed-accuracy trade-off, is known to be magnitude-insensitive, since it can be demonstrated that only information regarding the *difference* between alternatives is needed in order to make an optimal choice satisfying such a criterion [24]. In fact, to decide as quickly as possible which of two stimuli is the *best* available, information regarding the overall value is irrelevant. In the attempt to study decision-making under the working hypothesis of statistical optimality, previous research has disregarded magnitude-sensitivity *a priori*

[see for example 24].

It is certainly possible to design experiments for which the effect of magnitude may be negligible; for example, in value-based experiments in which the value of alternatives is near-identical across trials [28]. So, while not focusing on magnitude in specific studies could be perfectly reasonable, the problem arises when results from such studies are used to inform conclusions about the cognitive architecture for decision-making.

Interestingly, from a speed-accuracy trade-off perspective, not only is the optimal strategy magnitude-insensitive, but it also predicts, in the absence of ad hoc additions [29–31], decision deadlocks in the case of equal alternatives, regardless of their magnitude, as illustrated by the so-called Buridan's ass paradox [32], in which a decision-maker faced with equal alternatives starves, not being able to choose the best alternative.

Magnitude-sensitivity, while puzzling from a speed-accuracy trade-off perspective [24, 26], allows decision-makers to optimise another trade-off that was proposed for value-based decisions: the **speed-value trade-off** [3, 33]. This trade-off has been conceptualised in order to account both theoretically and empirically for magnitude-sensitivity. Such a decision trade-off is common in naturalistic settings in which agents are rewarded by the value of the alternative chosen, and not on making a correct choice nor by the overall accuracy of choices over many trials as commonly assumed in the laboratory.

In naturalistic environments [33–35] characterised by uncertainty of future alternatives due to factors such as competition, scarcity, and degradation of

alternatives, magnitude-sensitivity is advantageous since it allows the maximisation of long-term reward in a simple manner. This mechanism allows decision-makers to sacrifice small differences between high-valued alternatives and avoid time-costly decision deadlocks associated with ambivalence between similar high-valued options [3, 11, 33]. Additionally, it has been proposed that a fast choice is advantageous over a slow but correct one in the case of negative rewards of high intensity (e.g., the loud noise of a quickly approaching stimulus) that are to be avoided quickly [11]. In the case of low-valued alternatives, magnitude-sensitivity allows decision-makers not to make a decision, by maintaining a decision deadlock, in the hope that they will encounter better alternatives in the future [8, 33]. This rationale is at odds with classical experiments in **value-based decision-making**, characterised by the requirement of a 'forced' choice, according to which participants must make a choice even in the case of low-valued alternatives [e.g., 36]. Furthermore, this rationale is at odds with the kind of static stimuli used in laboratory settings; often, alternatives outside the laboratory are nonstatic, hence maintaining a deadlock with low-value alternatives may allow for better alternatives to be encountered.

Nonetheless, the fact that magnitude-sensitivity is also observed in other types of decision-making that are not characterised by speed-value trade-offs, such as is often the case in **perceptual decision-making** [3, 11, 37–44] [but see 45], reinforces a common belief [45, 46] that the same computational mechanism may be responsible for how organisms make decisions

across different domains, and that their decision-making architecture *is* indeed magnitude-sensitive. Given the importance of optimising speed-value trade-offs in naturalistic decisions, it is assumed that a signature of this evolutionarily plausible strategy is observed for all types of decision-making, even for those tasks in which it results in little or no benefit (e.g., in most perceptual choices and tasks with fixed or no reward).

In the next section we discuss the most prominent descriptive models of decision-making that can, or surprisingly cannot, account for magnitude-sensitivity, and present theoretical and empirical evidence in favour and against each account as a prime candidate for the explanation of magnitude-sensitivity. Our focus is on *dynamic* models of evidence accumulation able to account for both choice and reaction times; there is a large number of decision models besides the ones addressed here [e..g, 47, 48], which however are not 'computational' [49] and generally only focus on the choice made but not on reaction times.

Magnitude-sensitivity in computational models of decision-making

A common assumption [52–54] when modelling decision-making is that evidence is accumulated over time until a threshold for a decision is reached and a decision is made in favour of that alternative, as depicted in Figure 1.

Despite their fundamental differences, computational models of decision-making show a high degree of model mimicry when fitted to decision-making data [10, 24, 55, 56] but, interestingly, magnitude-sensitivity has been shown

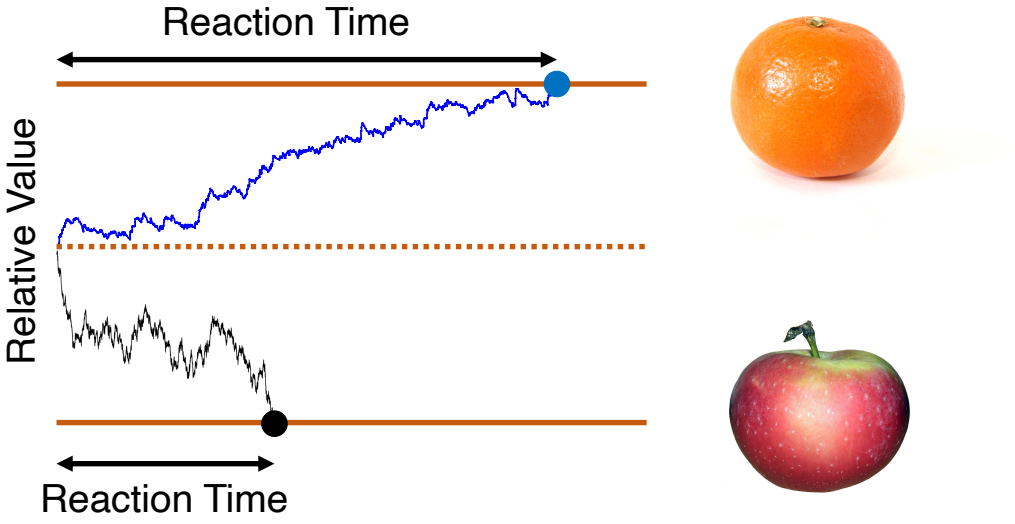

Figure 1: **Accumulation of evidence in value-based decision-making**. According to evidence accumulation models [e.g., 23, 24, 50], decision-makers accumulate relative value over time (to choose the orange versus to choose the apple). Decision-makers, if unbiased, start to accumulate value at the same distance between the two alternatives (dashed line). The value accumulation is subject to noise and a decision is made when a boundary threshold is reached. The time taken to reach a threshold is the reaction time. The blue line shows a case in which a decision-maker has a slight preference for the orange (small evidence drift); hence, the decision is slow because information needs to be sampled and integrated over many steps in order to reach a boundary. The black line shows a case in which a decision-maker has instead a strong preference for the apple (large evidence drift); in this case, the boundary is reached quickly. In the case where alternatives are indistinguishable (both have the same value, regardless of whether it is a high value or a low value), a decision deadlock is maintained: noisy fluctuations around the dashed line of indifference are expected [3, 23, 24, 32, 33]. However, a boundary could eventually be reached because of noise in the accumulation process [51]. A magnitude-sensitive mechanism ensures that if both alternatives are equally poor no choice is made, while if both alternatives are equally good a fast random choice is made [3, 11, 33]. In the case in which alternatives have different values, magnitude-sensitivity predicts that, on average, the most valuable alternative is selected but with lower consistency (i.e., lower accuracy) compared with when decision-makers are optimising speed-accuracy trade-offs [11].

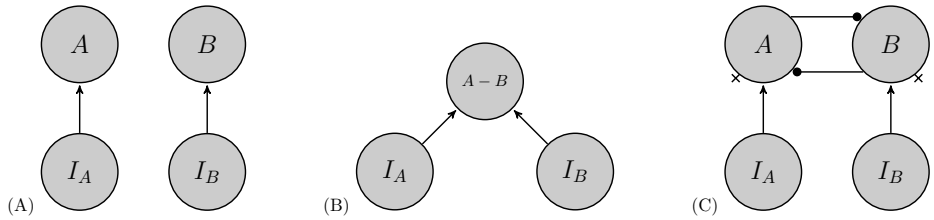

Figure 2: **Schematics of accumulator models**. A and B are evidence accumulators integrating evidence over time, I_A and I_B are the respective sensory inputs. (A) In the race model [59, 60], two accumulators integrate evidence for each alternative separately. (B) In the drift diffusion model [61, 62], a single accumulator integrates difference in evidence between the alternatives. (C) In the leaky competing accumulator [50], two accumulators integrate evidence for each alternative separately; the two accumulators inhibit each other (proportionally to evidence accumulated), and information is lost (i.e., leaked) over time. Solid arrows denote excitatory connections, arrows ending with a dot denote inhibitory connections, crosses on the evidence accumulators denote leak of evidence. Figure adapted from [24, 63].

to be a feature of decision-making that allows researchers to discriminate between models of choice, either qualitatively (i.e., magnitude-sensitive versus insensitive models) or quantitatively [9–12, 57, 58].

While it is beyond the scope of our review to present all the decision-making models that have been proposed, we introduce the reader to the most relevant ones [for in depth discussion, see 24] for the topic of our paper. Such models differ in substantial ways, including: the mechanism that implements magnitude-sensitivity (or insensitivity); their goodness of fit to empirical data; whether they can explain magnitude-sensitivity across domains and tasks; their ecological validity and/or biological plausibility; and the importance that they assign to additional, unique, factors such as visual fixations.

Race models

One of the oldest proposed accumulator models is the race model [59, 64], Figure 2(A). In this model, two accumulators integrate evidence separately for the two respective inputs and once one of the two accumulators reaches a threshold, a decision is made. This is a prototype of a magnitude-sensitive model, since each accumulator integrates evidence separately and conditions with higher magnitude reach a decision threshold faster. However, it has been shown in studies of magnitude-sensitivity that race models cannot account for data quantitatively (i.e., obvious mismatches in choice and reaction times) when, compared with a baseline condition, the magnitude of the alternatives is increased by keeping either difference or ratio between alternatives constant [11]; similarly, in other studies that compared models extensively, race models performed poorly when fitted to data [58]. For these reasons [2, 11, 58], this account is not generally considered as a plausible candidate in investigations about magnitude-sensitivity.

The drift diffusion model

A prominent model of binary decision-making is the drift diffusion model [DDM; 23, 61, 62, 65, 66] – arguably the most celebrated computational model of binary decision-making [for a review of its applications see 62], which under specific parametrisations [24, 25, 67] optimises speed-accuracy trade-offs. In the DDM, a single accumulator integrates noisy *difference* in evidence between the alternatives. This model cannot explain magnitude-

sensitivity since it is purely *relative* (i.e., only information about the *difference* between alternatives is processed, so magnitude information is lost) and, as a consequence, it is magnitude-insensitive.

The attentional drift diffusion model

A modification of the DDM, particularly popular in studies of value-based decision-making, is the attentional DDM [36] in which the accumulation of difference in evidence between the alternatives is biased by visual fixation patterns (which are de facto equated to attention) and duration (see [68, 69] for a critique of this key assumption). When one of the alternatives is fixated, the value of the nonfixated alternative is discounted by a factor $0 \leq \theta \leq 1$.

Fitting this model to previous datasets, it has been shown that this model can accommodate magnitude effects [16]. This is because the biasing effect of visual fixation ensures that net evidence is higher for high-magnitude alternatives compared to low-magnitude alternatives, even in the case of equal alternatives. Consider a scenario of alternatives having values of 5 versus 5 (high magnitude) and 1 versus 1 (low magnitude); if the discount factor θ has a value of .3, net evidence for the first case is 5 - (5×.3) = 3.5, while for the second case it is 1 - (1×.3) = .7; thus, the first case is more likely to reach a boundary faster, exhibiting magnitude-sensitivity.

The visual fixation model was not compared with competing accounts of magnitude-sensitivity, but only to magnitude-insensitive models [16, 70]. Although the goodness of fit from visual-fixation accounts is adequate (but

see [71]), and this model has also been extended to the multi-alternative case [72], we believe that an account of magnitude-sensitivity based on visual fixations alone may be unsatisfactory on theoretical grounds.

In fact, a causal effect of visual fixations is not necessary for magnitude-sensitivity. In most experiments investigating magnitude-sensitivity [e.g., 3, 11, 12, 17], participants were required to fixate the centre of the screen for the whole duration of the experiment, minimising in this way any effect due to visual fixations, which in some cases were even controlled [17]. In all these cases, magnitude-sensitivity was observed, meaning that, while a visual fixation bias could provide an additional source of magnitude-sensitivity, the core mechanism explaining magnitude-sensitivity may be unrelated to visual fixations alone. Noticeably, one could speculate that the mechanisms of the attentional DDM may extend to covert attention (i.e., directing attention independent of fixations). In tasks where they are required to fixate at the centre of the screen, participants could use covert attention which discounts the value of nonattended items [73]; in this case, the model could account for the cases mentioned earlier, which are currently problematic. This hypothesis, and a theoretical reconsideration of the model, would need to be tested by future research.

Furthermore, while experiments in laboratory settings include the simultaneous presence of two (or multiple) alternatives, in numerous cases alternatives are likely to be encountered sequentially [e.g., 74] – that is, one at a time. In the case of a single isolated stimulus, in which agents have to decide

whether to obtain the reward or not, it remains unclear how visual fixations could bias the decision in order to produce magnitude-sensitivity.

The input-dependent noise DDM

Another magnitude-sensitive modification of the canonical DDM is the input-dependent noise DDM [11, 12]. In this model, there is an additional noise component, on top of constant processing noise common to all DDMs. This additional noise scales with the magnitude of the alternatives.

Input-dependent noise is an attractive account since it links biologically and neurally plausible principles [75, 76] that have been shown to provide excellent fits to data [75–77] and, in particular, to magnitude-sensitive data [10–12, 78]. Experimentally, there is evidence [75, 76] for a dominant role of input-dependent noise during evidence accumulation. Even in the case of equal alternatives, input-dependent noise ensures that decisions between high-magnitude alternatives are made faster than decisions between low-magnitude alternatives; higher noise, associated with higher magnitude, leads to stronger fluctuations in the accumulated evidence which, in turn, is more likely to reach a decision threshold faster.

The input-dependent noise account suggests that magnitude-sensitivity is an artefact of the limitations of information processing [11]. Magnitude-sensitivity has not been selected against by natural selection because of the advantages it provides in various scenarios, for example those characterised by speed-value trade-offs.

To the best of our knowledge, an account that combines input-dependent noise with integration of differences in evidence between the alternatives has been applied only to the two-alternative case (however, see [17] for an input-dependent noise account with separate accumulators for each of three possible alternatives). Extending input-dependent models to more than two alternatives would require ad hoc, nontrivial theoretical and mathematical adaptations [e.g., 2, 11].

Lateral inhibition

Finally, the leaky competing accumulator model [LCA; 50], Figure 2(C), is a biologically inspired model of choice [see 50, 79], in which two separate accumulators leak evidence over time (i.e., evidence is lost/forgotten) and there is lateral inhibition between evidence accumulators. In this case, magnitude-sensitivity arises as a function of lateral inhibition between evidence accumulators [10, 11, 50]. Inhibition between accumulators is proportional to overall activation, and it increases with time [see 11, 50, 79]. Consequently, early decision dynamics are similar to those of race models (due to low inhibition) and are dominated by the magnitude of alternatives; as time progresses, inhibition increases (which in turn decreases the effect of magnitude) and this leads to late dynamics to be dominated by the alternatives' difference, in a DDM-like fashion [11, 50, 79]. In short, the resulting temporal dynamics of lateral inhibition between evidence accumulators [24, 54] ensure that decisions for high-magnitude conditions are made faster and more inaccurately,

due to ‘early’ magnitude-sensitivity.

This account links a biologically and neurally plausible mechanism [80–85] of decision-making to magnitude-sensitivity; it has also been shown to fit magnitude-sensitive data well [10, 11]. It is interesting to note that several models characterised by nonlinear decision dynamics [8, 19, 86–89] under specific parametrisations can approximate a lateral inhibition model [10, 11, 50] and show a very high degree of model mimicry when simulated or fitted to data [10, 24] – see also Box 2 and Box 3. Moreover, this model is naturally extendable to more than two alternatives [2, 79] in a biologically plausible way.

A strong model mimicry between accounts based on input-dependent noise and those based on lateral inhibition has also been documented [10, 11, 24] and it has been debated which of the two provides a better fit to the data. While a first study [11] showed that both models provide equally good numerical fit to data, subsequent studies comparing the two models in more detail have suggested that the fit of the input-dependent noise account is superior [10, 12]. However, a recent study found lateral inhibition to be a better quantitative account of magnitude-sensitivity [58] when compared with input-dependent noise and other accounts. Contrasting modelling results highlight the importance of designing tasks that allow unique qualitative predictions of different models to be tested experimentally; specific parametrisations and fitting methods are likely to drive contrasting results across modelling studies. Alternatively, a modelling approach would have to

rely on large-scale experiments and datasets [90] to shed light on the best descriptors of magnitude-sensitivity.

From an evolutionary perspective, compared with the input-dependent noise account, the LCA account proposes that magnitude-sensitivity has evolved as an adaptive feature that allows organisms to directly optimise speed-value trade-offs [3, 11, 13, 33]. In this case, magnitude-sensitivity is implemented as part of the decision architecture rather than as an artefact of the limitations of information processing.

Honourable mentions

Another theoretical modification of the DDM proposes that magnitude-sensitivity arises as a consequence of input-dependent *across-trial variability* in evidence accumulation [12] – which in all other models is generally kept constant. However, we note that while a proposal based on input-dependent across-trial variability may account for magnitude-sensitivity in laboratory experiments, it lacks ecological plausibility. It is unclear how input-dependent *across-trial variability* could give rise to magnitude-sensitivity for isolated decisions [8, 9, 74, 91, 92], as is the case in naturalistic decision-making in which decisions are not embedded in a sequence of repeated encounters, but are met as isolated decision problems.

Lastly, we present a model that differs significantly from those presented up to this point. Compared with all models discussed earlier, the urgency-gating model [93, 94] assumes that evidence is not accumulated at all; instead,

samples for each alternative are only low-pass filtered. Evidence signals are then combined with a time-dependent 'urgency signal' which brings, over time, activity nearer to the decision threshold; once a decision threshold is reached, a decision is made. This model avoids long reaction times for low evidence trials and, in theory, it could be adopted to explain magnitude-sensitivity. However, to the best of our knowledge, there is no research on how the urgency-gating model accounts for magnitude-sensitivity theoretically and empirically. Future research should assess the ability of this account to explain magnitude-sensitivity, especially in view of recent studies [95] showing a mixture of evidence accumulation and time-varying urgency in decision-making (but see [96]).

Magnitude-sensitive optimal decision-making

So far, we have reviewed descriptive accounts of decision-making; that is, accounts of *how decisions are made*. Normative accounts of decision-making focus instead on *how decisions should be made*. The study of optimal decision-making makes it possible to propose models that define the ideal decision-making strategy ('policy') given an objective criterion to be optimised and a number of underlying assumptions regarding the environment. Under the assumption that evolutionary pressure selects optimal strategies for decision-making [99], normative models can then be used to predict behaviour and/or explain (ir)regularities in decision-making. There is obviously a direct link between the descriptive and normative accounts; in some cases,

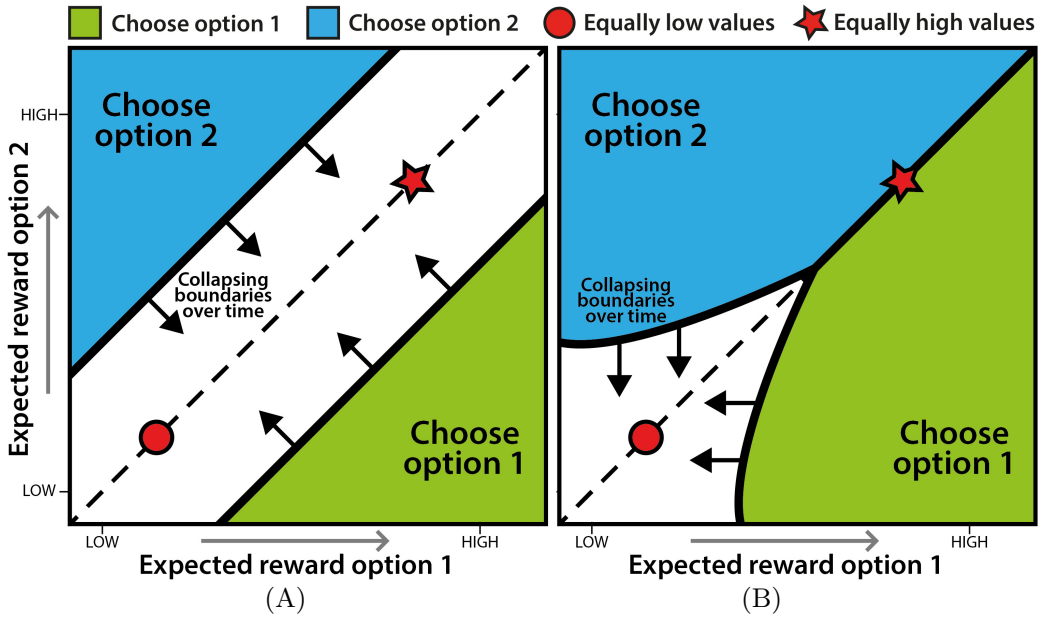

Figure 3: **Optimal policy for value based decision-making**. (A) For binary decisions with linear subjective utility (and linear cost of time and constant processing noise) [97], the optimal policy is magnitude-insensitive. (B) For binary decisions with nonlinear subjective utility [97], the optimal policy is magnitude-sensitive. Nonparallel collapsing boundaries implementing magnitude-sensitivity can also be observed for linear subjective utility when time has a multiplicative cost [13, 34]. In the presence of input-dependent noise [98], boundaries remain parallel to each other, however high noise (associated with high magnitude) results in stronger fluctuations in accumulated evidence which hit a decision boundary faster compared with low-magnitude conditions.

researchers use descriptive accounts and investigate under what parametrisations and in which scenarios such models implement optimal strategies [24], while in other cases researchers derive optimal models of decision-making [57, 97, 100] for specific scenarios and then compare the optimal models to the descriptive accounts in order to highlight similarities and differences. A word of caution with regard to the concept of 'optimality': any model/rule implements optimality *only under specific, limited, assumptions*. Therefore, these assumptions are often the matter of debate in research about optimal decision-making [25, 33, 34, 57].

Recent research has focused on deriving the optimal policy for value-based decision-making [13, 34, 97, 98, 101, 102]. The general question that this research has addressed is: how much evidence should be accumulated before committing to a choice, in the case in which the rewards associated with the stimuli are uncertain?

For the binary case, the optimal policy for value-based decision-making [97] shows, under certain assumptions, striking similarities to the optimal strategy for speed-accuracy trade-offs [24]. As shown in Figure 3(A), for value-based decisions, the optimal strategy is implemented by a purely relative model (i.e., decision dynamics are dominated by the *difference* between alternatives), resembling a classical DDM, in which, unlike the canonical DDM, parallel boundaries collapse over time. Figure 3(A) shows that the boundaries are both parallel to each other in the reward space of the two options, and parallel to the identity line on which *all* equal alternatives lie.

As a result, the optimal policy is magnitude-insensitive since the **collapsing boundaries** would hit a choice for equally low or equally high-value alternatives in the same time. The speed at which boundaries collapse is a function of overall reward received; this ensures that for highly rewarding environments thresholds collapse faster; nonetheless, the dynamics of each trial remain magnitude-insensitive, a result that clashes with empirical evidence [3, 9, 11].

Seminal work has shown that the optimal policy for two-alternative value-based decision-making could predict magnitude-sensitive reaction times when utility grows nonlinearly in the reward (under specific functional forms of the utility) [97]. In this case, collapsing boundaries are not parallel to the identity line in the reward space, ensuring faster decision with high-magnitude equal alternatives, as illustrated in Figure 3(B).

Building on this work, two studies [13, 34] have shown that once the ecologically plausible assumption of **multiplicative cost of time** [103] is added to the decision-making process, the optimal policy for two-alternative value-based decision-making predicts magnitude-sensitive reaction times regardless of the specific details of the subjective utility function (i.e., linear or nonlinear). In these accounts, there is a discount of future rewards that is nonlinearly proportional to the time it takes to realise such rewards, a common assumption [34, 104] made in behavioural ecology [but see 105]. Multiplicative cost of time can account for the inherent uncertainty of future rewards in some naturalistic environments and for uncertainty in value

representations of future events [106].

Recent work [100] has shown that, also in the multi-alternative case, magnitude-sensitivity is observed, but with surprising differences compared to the case of binary decisions [13, 57]. Theory shows that, for multi-alternative decision-making, nonlinear subjective utility combined with linear time discounting induces only very weak magnitude-sensitivity; multiplicative time discounting instead predicts strong magnitude-sensitivity for both linear and nonlinear utility, in line with empirical evidence showing strongly magnitude-sensitive reaction times with three alternatives [57]. Overall, the multiplicative cost of time account assumes that magnitude-sensitivity is the signature of a common mechanism for decision-making that has evolved to directly optimise speed-value trade-offs and that, as a consequence, is observed even when there is no reward associated with a choice [e.g., in perceptual tasks, see 3, 9, 11, 12] or in experiments for which the cost of time is controlled/minimised (e.g., when the experiment duration and number of trials are fixed). That is, even though decision-makers perform a purely perceptual task, they perform it *as if* optimising a speed-value trade-off for value-based choices [34, 57].

Lastly, a recent study on normative magnitude-sensitivity has questioned the assumption of constant processing noise made in all models of optimal binary decision-making [13, 34, 97, 100]. This study has investigated the optimal policy for value-based decision-making with input-dependent noise [98]. Interestingly, input-dependent noise in the evidence accumulation pre-

dicts magnitude-sensitive optimal decision-making, without requiring further assumptions regarding the utility function or cost of time [98]. As discussed by the authors [98], this account could provide a parsimonious explanation that bridges the gap between various task assumptions (i.e., there is always noise, endogenous or exogenous, associated with the stimuli regardless of the task) and between various types of decision-making (i.e., value-based and perceptual).

Overall, the normative approach seems to suggest a *hybrid* account in which a relative model, resembling a classical DDM, is coupled with either input-dependent noise, nonlinear time cost or nonlinear utility functions in order to give rise to magnitude-sensitive optimal decisions. Interestingly, none of these accounts are mutually exclusive; indeed, these accounts may work in conjunction, an interesting avenue for future research.

Concluding Remarks

The research synthesised here shows how an entire field can be significantly reshaped in view of a 'simple' experimental result that challenges untested assumptions about key aspects of decision-making. Specifically, the result that decisions are made faster and are more likely to be random with high-magnitude alternatives [3, 11, 33] has led to reconsideration of current descriptive, theoretical, and normative accounts of decision-making. In particular, purely relative models of descriptive and normative decision-making that have dominated the literature for decades had to be revisited.

This is a dramatic change imposed by results of magnitude-sensitivity. Similarly, the dominant view of equating optimal decision-making with statistical optimality seems now obsolete, in view of theoretical arguments driven by magnitude-sensitivity and its associated speed-value trade-off.

An overview of the literature on magnitude-sensitivity shows that while it is established that decision-making *is* magnitude-sensitive and that this is a problematic result for several dominant views on decision-making, the computational, algorithmic, and implementational levels of analysis [107] concerning magnitude-sensitivity are far from being fully understood. Crucially, it remains to be understood whether embracing existing models, or patched up versions of them, is the best solution, or whether magnitude-sensitivity calls for bigger scale changes in how we model decisions. We believe that one of the most promising avenues for future research consists of considering optimal strategies from first principles [108], by understanding the context in which decisions are made [109] rather than tweaking existing models with ad hoc fixes.

Given the number of timely questions that remain unanswered (see Outstanding Questions), we believe that magnitude-sensitivity, being a decisive feature in view of which decision-making should be understood, will continue to motivate further research into the fundamental aspects of decision-making.

Box 1 - Magnitude-sensitive decision-making without a brain

The slime mould *Physarum polycephalum* is a macroscopic, unicellular organism, abundant in nature. *P. polycephalum* explores the environment, searching for nutrients by extending tubular structures called pseudopods. Despite its lack of a nervous system, *P. polycephalum* is capable of complex behaviours, such as, for example, sophisticated navigational skills [110]. Furthermore, a decade of research has shown that fundamental features of cognition and information processing observed in humans and animals, such as the speed-accuracy trade-off [111], habituation [112], learning [113–115] and memory (externalised spatial "memory") [116] are also observed in slime moulds. Surprisingly, even violations of rationality in *P. polycephalum* are in line with those observed in humans and other animals [117]. While obviously the architecture of information processing differs between different biological taxa, there are striking similarities in observable behaviours [1], postulated cognitive mechanisms [118] and common properties for information processing [119].

Of particular interest for our review is the evidence of magnitude-sensitivity in slime moulds [1]. In their study [1], the authors used three different strains of *P. polycephalum*. The experimental arena consisted of a Petri dish filled with a plain agar. Three equidistant holes were punched in the agar (in a triangular arrangement); on one

the slime mould was introduced, while two equal sources of food were introduced in the remaining two holes. The nutritional value of equal food sources could be poor, medium, or rich. Each decision problem was repeated 60 times for each of the three different strains, and the time it took to the slime mould to reach any of the two stimuli was recorded. Results showed that, for all three strains, time decreased significantly as a function of the food quality, in line with predictions from the speed-value trade-off. Moreover, compared with decisions with a single food source, slime moulds exhibited faster decision times in the equal alternatives case. Magnitude-sensitivity was also observed in a similar recent study investigating multi-alternative decisions in slime moulds [57].

These experiments show that even a biologically simple organism can maximise value over accuracy of decisions, in order to avoid costly decision deadlocks with high-value alternatives, and that the signature of a magnitude-sensitive decision-making architecture is observed in a wide range of taxa. These results reinforce the argument that magnitude-sensitivity is a fundamental mechanism of decision-making that allows resolution of decision deadlocks adaptively.

Box 2 - Value-sensitive decision-making in animal swarms

Predictions and empirical results from collective behaviour [8, 92] have inspired research about magnitude-sensitivity. In particular, the study of decision-making in social insect colonies has led to the questioning of the dominant view of speed-accuracy trade-offs for human decision-making [33] and has inspired the first magnitude-sensitivity tests in humans [3, 11].

Collective behaviour can generate similar patterns to those observed in diverse species at the individual level and it is therefore interesting to study it through the lens of cognitive science [120–123]. A fundamental cognitive process is collective decision-making, through which groups of individuals either reach a consensus to select one of several available options, or distribute themselves over different tasks [124, 125]. When natural selection acts at a super-organismal level, the mechanisms of information processing by groups can be similar to the brain mechanisms for optimal decision-making [126]. In fact, some of the accumulator models of Figure 2 that describe decision-making as excitation and inhibition of neuron populations can also describe the house-hunting process of social insects [54]. Analogously to neuronal mechanisms, the house-hunting insects accumulate noisy evidence of the environment and interact through simple excitatory and inhibitory signals. Through numerous interactions, the colony collectively pro-

cesses the collected information, until a decision quorum is reached [127]. The house-hunting process can be seen as a collective form of the two-alternatives forced-choice tasks in which the decision value is the colony reward in terms of the dwelling place quality. The analysis of an empirically motivated model of European house-hunting honeybees has predicted that the emergent group-level response is magnitude-sensitive with regard to the quality of the nest-site location [4, 8, 92]. Individual bees do not have a magnitude-sensitive response to environmental stimuli; however, the collective response shows such a response pattern [7]. Interestingly, the analysis predicts that such emergent dynamics also generalise to other psychophysical laws normally observed in human individuals and other animals, such as Hick-Hyman's and Weber's laws [7].

Box 3 - Magnitude-sensitive decision-making in robot swarms

Inspired by the honeybee house-hunting models [4, 8, 92], a series of studies implemented collective site-selection behaviours in swarms of robots (Figure I) [5, 6, 128–130]. The robots explore the environment comprising n available sites and exchange excitatory and in-

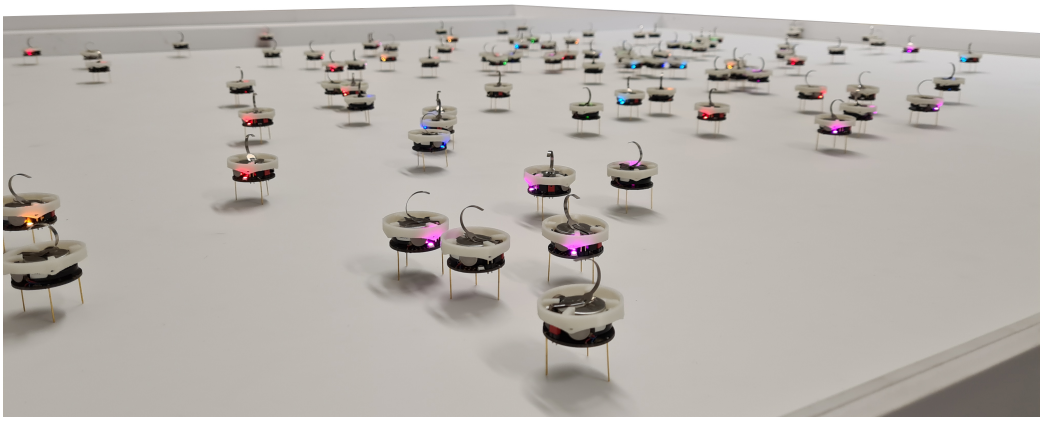

Figure I: **Robot swarms can efficiently manage the speed-value trade-off.** Efficient solutions from nature have always been the source of inspiration for the design of machines and their algorithms. Swarm robotics looks at collective animal behaviour to design large groups of robots able to self-organise. In the depicted experiment, a swarm of Kilobot robots exchanges excitatory and inhibitory messages to reach consensus on the best site. Each robot shows its opinion through a coloured LED. The interaction mechanism of the robots is similar to the ones of house-hunting honeybees and neuronal accumulator models. The results show that the swarm manages the speed-value trade-off efficiently and exhibits magnitude-sensitive behaviour [5, 6].

hibitory messages with communication patterns similar to the accumulator model of Figure 2(C). Each robot forms its opinion on noisy estimates and, similarly to individual bees, exchanges its opinion for a quantity of time proportional to the estimated value of the site's quality. Through this quality-dependent behaviour, the swarm solves the so-called best-of-n problem by reaching a consensus in favour of the best of the n sites. Swarm robotics experiments have found two qualitatively different responses when the magnitude of the stimuli is changed [5]. In a form of *'good-enough rule'*, the 150 robot swarm

formed a decision deadlock and refrained from making a collective decision when the quality of the two equal-quality sites was low. Instead, when the environment had two high quality sites, the swarm broke the symmetry and selected either of the two. The minimum quality to trigger the decision (the good-enough threshold) can be controlled by changing the frequency of the social interactions among robots. Thus, robot swarms display magnitude-sensitive behaviour.

Biological observations and robotics experiments confirm that when the values of the option qualities are similar and obtaining high accuracy is difficult, the collective system can neglect accuracy in favour of better performance in terms of the speed-value trade-off. An open question for future research consists of understanding the mechanistic causal relationship between individual rules that are magnitude-insensitive and the emergent magnitude-sensitive collective response.

Algorithms for decentralised artificial systems, such as robot swarms, have been largely inspired by models of animal behaviour and brains [131, 132]. In turn, implementation of the theoretical models in physical robots that are limited in sensory and communication abilities can further help in the understanding of the link between individual rules and collective behaviour [133, 134]. Rather than simulating the behaviour, implementing it on physical devices forces researchers to think about social behaviours in relation to the perceptual, cognitive,

and motor abilities of the individual agents [135]. Identifying how the information flows between decision agents, whether neurons, bees, or robots, can uncover similarities that can be understood through interdisciplinary research.

Outstanding Questions

- How many mechanisms contribute to magnitude-sensitivity? While the assumption of a single mechanism for magnitude-sensitivity is attractive, future investigations could focus on elucidating the unique contribution of the different mechanisms that, depending on the nature of the task, explain magnitude-sensitivity. Overall, depending on the specific domain and environmental constraints, it is possible that *multiple* sources contribute to the observed result of faster reaction times for higher magnitude conditions. Interestingly, the mechanisms reviewed in this paper (e.g., overt/covert attention, input-dependent noise, lateral inhibition, nonlinear utility, nonlinear time cost) are not mutually exclusive.

- What are the temporal dynamics of magnitude-sensitivity? That

is: do both early and late stages of evidence accumulation show magnitude-sensitivity, or exclusively one of the two? In the leaky competing accumulator (under certain assumptions), early stages of evidence accumulation are magnitude-sensitive and late stages are dominated by the difference between alternatives, while the input-dependent noise account predicts a constant effect of magnitude throughout a trial.

- How to extend the input-dependent noise explanation to the multi-alternative case?
- Does magnitude-sensitivity extend to all sensory modalities?
- What is the link between descriptive/normative accounts of magnitude-sensitivity and their neural implementation?
- Can models and accounts based on other mechanisms than accumulation (e.g., urgency-gating) explain magnitude-sensitivity?
- Moving beyond reaction times and accuracy alone, what other features of decision-making (e.g., confidence judgements) are affected by magnitude-sensitivity and how?
- Has magnitude-sensitivity evolved as an artefact of imperfect implementation (as in the input-dependent noise account) or as an

adaptive feature (as in the temporal discount account) that directly optimises speed-value trade-offs?

- How do aneural organisms implement magnitude-sensitivity?

Glossary

Collapsing boundaries: decision thresholds that are not fixed but that approach each other (collapse) over time. This mechanism ensures that, late in the trial, less evidence is needed in order to make a decision. Collapsing boundaries *per se* do not necessarily predict magnitude-sensitivity.

Equal alternatives: alternatives that have the same physical/reward and perceived difference (null difference). Keeping physical/reward difference constant and increasing magnitude introduces well-known psychophysical transformations (in the case of perceptual stimuli) and/or unknown transformations (in the case of stimuli such as food images). Using equal alternatives removes the issue of how to scale differences (absolute, ratios, etc.) and makes it possible to test magnitude-sensitivity directly, in the absence of confounds.

Magnitude: the sum of (goal-relevant) intensities or rewards of all

stimuli under consideration during a trial.

Magnitude-sensitivity: the phenomenon that performance in decision-making is affected by the magnitude of the stimuli under consideration; in particular, as magnitude increases, reaction times decrease and the response is made more inaccurately (when a correct response exists). The term 'magnitude-sensitivity' also refers to the mechanism that allows resolution of deadlocks adaptively by maximising value over accuracy of a decision.

Multiplicative cost of time: the value of an alternative is multiplied by a discount factor that decreases with time. An example (among many functional forms possible) of multiplicative discounting: a reward one time step in the future is discounted by rate $\gamma < 1$, two time steps in the future by γ^2 and so on. At each time step there is a cost that is proportional to the overall value of the alternatives.

Perceptual decision-making: in perceptual decision-making, decisions are made according to an objective, perceptual criterion. Examples include: deciding which of two stimuli is brighter, or which of two stimuli is bigger. It is postulated that decision-makers integrate noisy physical evidence over time.

Speed-accuracy trade-off: in decision-making, response times covary with accuracy so that fast responses are made more inaccurately, while correct responses are made slower. When decision-makers opti-

mise speed-accuracy trade-offs, a correct response is expected in the shortest time possible.

Speed-value trade-off: in decision-making, response times covary with value so that fast (more inaccurate) responses are made with high-value alternatives, while low-value alternatives responses are made more slowly. This is a trade-off that characterises value-based decisions (e.g., food choices) in naturalistic environments.

Value-based decision-making: decision-makers compare the values of rewarding alternatives, such as food items or lotteries. Decisions are made on the basis of a subjective criterion, such as which food item is preferred. Similar to perceptual decision-making, it is postulated that decision-makers integrate endogenous noisy evidence regarding the true value of the stimuli over time. In value-based decisions, the subject's reward is the value of the option chosen, rather than a function of whether they chose the best or not.

Acknowledgements

We are grateful to Andrei Teodorescu (University of Haifa) for helpful comments on a previous (significantly different and unpublished) manuscript that inspired part of this work.

A.P. and F.G. gratefully acknowledge funding from the European Research Council (ERC-ADG-835002—GEMS).

J.A.R.M. and A.R. gratefully acknowledge funding from the European Research Council (ERC-CoG-647704—DiODe).

A.R. also acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Chargé de Recherches.

Declaration of Interests

No interests to declare.

References

[1] Dussutour, A. et al. (2019). Phenotypic variability predicts decision accuracy in unicellular organisms. *Proc. R. Soc. B* 286, 20182825.

[2] Teodorescu, A. R. and Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychol. Rev.* 120, 1–38.

[3] Pirrone, A. et al. (2018). Evidence for the speed–value trade-off: Human and monkey decision making is magnitude sensitive. *Decision* 5, 129–142.

[4] Reina, A. et al. (2017). Model of the best-of-n nest-site selection process in honeybees. *Phys. Rev. E* 95, 052411.

[5] Reina, A. et al. (2018). Effects of spatiality on value-sensitive decisions made by robot swarms. In *Distributed Autonomous Robotic Systems (DARS)* (pp. 461–473). Springer volume 6 of SPAR.

[6] Talamali, M. S. et al. (2021). When less is more: robot swarms adapt better to changes with constrained communication. *Sci. Robot.* 6.

[7] Reina, A. et al. (2018). Psychophysical laws and the superorganism. *Sci. Rep.* 8, 4387.

[8] Pais, D. et al. (2013). A mechanism for value-sensitive decision-making. *PLOS ONE* 8.

[9] Pirrone, A. et al. (2018). Single-trial dynamics explain magnitude sensitive decision making. *BMC Neurosci.* 19, 1–10.

[10] Bose, T. et al. (2020). Comparison of magnitude-sensitive sequential sampling models in a simulation-based study. *J. Math. Psychol.* 94, 102298.

[11] Teodorescu, A. R. et al. (2016). Absolutely relative or relatively absolute: violations of value invariance in human decision making. *Psychon. Bull. Rev.* 23, 22–38.

[12] Ratcliff, R. et al. (2018). Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. *Cogn. Psychol.* 103, 1–22.

[13] Steverson, K. et al. (2019). Sensitivity of reaction time to the magnitude of rewards reveals the cost-structure of time. *Sci. Rep.* 9, 1–14.

[14] Hunt, L. T. et al. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* 15, 470–476.

[15] Kvam, P. D. and Pleskac, T. J. (2016). Strength and weight: The determinants of choice and confidence. *Cognition* 152, 170–180.

[16] Smith, S. M. and Krajbich, I. (2019). Gaze amplifies value in decision making. *Psychol. Sci.* 30, 116–128.

[17] Niwa, M. and Ditterich, J. (2008). Perceptual decisions between multiple directions of visual motion. *J. Neurosci.* 28, 4435–4445.

[18] Shevlin, B. and Krajbich, I. (2020). Attention as a source of variability in decision-making: Accounting for overall-value effects with diffusion models. *PsyArXiv* Published online July 29, 2020. https://psyarxiv.com/rewtq/.

[19] Bose, T. et al. (2019). Inhibition and excitation shape activity selection: effect of oscillations in a decision-making circuit. *Neural Comput.* 31, 870–896.

[20] Zajkowski, W. et al. (2019). Reward certainty and preference bias selectively shape voluntary decisions. *bioRxiv* Published online November 07, 2019. https://doi.org/10.1101/832311.

[21] Frömer, R. et al. (2019). Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nat. Commun.* 10, 1–11.

[22] Otto, A. R. and Daw, N. D. (2019). The opportunity cost of time modulates cognitive effort. *Neuropsychologia* 123, 92–105.

[23] Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922.

[24] Bogacz, R. et al. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765.

[25] Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychon. Bull. Rev.* 22, 38–53.

[26] Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* 8, 150.

[27] Stafford, T. et al. (2020). Quantifying the benefits of using decision models with response time and accuracy data. *Behav. Res. Methods* 52, 2142–2155.

[28] Gwinn, R. and Krajbich, I. (2020). Attitudes and attention. *J. Exp. Soc. Psychol.* 86, 103892.

[29] Yau, Y. et al. (2020). Neural correlates of evidence and urgency during human perceptual decision-making in dynamically changing conditions. *Cereb. Cortex* 30, 5471–5483.

[30] Evans, N. J. et al. (2020). The role of passing time in decision-making. *J. Exp. Psychol. Learn. Mem. Cogn.* 46, 316–326.

[31] Voskuilen, C. et al. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *J. Math. Psychol.* 73, 59–79.

[32] Oud, B. et al. (2016). Irrational time allocation in decision-making. *Proc. R. Soc. B* 283, 20151439.

[33] Pirrone, A. et al. (2014). When natural selection should optimize speed-accuracy trade-offs. *Front. Neurosci.* 8, 73.

[34] Marshall, J. A. R. (2019). Comment on 'Optimal Policy for Multi-Alternative Decisions'. *bioRxiv* Published online December 24, 2019. https://doi.org/10.1101/2019.12.18.880872.

[35] Fawcett, T. W. et al. (2014). The evolution of decision rules in complex environments. *Trends Cogn. Sci.* 18, 153–161.

[36] Krajbich, I. et al. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298.

[37] Pins, D. and Bonnet, C. (1996). On the relation between stimulus intensity and processing time: Piéron's law and choice reaction time. *Percept. Psychophys.* 58, 390–400.

[38] Stafford, T. and Gurney, K. N. (2004). The role of response mecha-

nisms in determining reaction time performance: Piéron's law revisited. *Psychon. Bull. Rev.* 11, 975–987.

[39] Palmer, J. et al. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404.

[40] Polanía, R. et al. (2014). Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron* 82, 709–720.

[41] Van Maanen, L. et al. (2012). Piéron's law and optimal behavior in perceptual decision-making. *Front. Neurosci.* 5, 143.

[42] Simen, P. et al. (2016). Scale (in) variance in a unified diffusion model of decision making and timing. *Psychol. Rev.* 123, 151–181.

[43] Turner, W. et al. (2021). Perceptual change-of-mind decisions are sensitive to absolute evidence magnitude. *Cogn. Psychol.* 124, 101358.

[44] Ratcliff, R. and McKoon, G. (2018). Modeling numerosity representation with an integrated diffusion model. *Psychol. Rev.* 125, 183–217.

[45] Smith, S. M. and Krajbich, I. (2021). Mental representations distinguish value-based decisions from perceptual decisions. *Psychon. Bull. Rev.* 28, 1413–1422.

[46] Krajbich, I. et al. (2015). A common mechanism underlying food choice and social decisions. *PLOS Comput. Biol.* 11, e1004371.

[47] Gilovich, T. et al. (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press.

[48] Gigerenzer, G. and Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart* (pp. 3–34). Oxford University Press.

[49] Alonso, E. and Mondragón, E. (2012). Uses, abuses and misuses of computational models in classical conditioning. In *Proceedings 11th International Conference on Cognitive Modeling* (pp. 96–100).

[50] Usher, M. and McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592.

[51] Turner, B. M. et al. (2017). The dynamics of multimodal integration: The averaging diffusion model. *Psychon. Bull. Rev.* 24, 1819–1843.

[52] Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.

[53] Usher, M. and McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.* 111, 757–769.

[54] Marshall, J. A. R. et al. (2009). On optimal decision-making in brains and social insect colonies. *J. R. Soc. Interface* 6, 1065–1074.

[55] Ratcliff, R. and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367.

[56] Khodadadi, A. and Townsend, J. T. (2015). On mimicry among sequential sampling models. *J. Math. Psychol.* 68, 37–48.

[57] Marshall, J. A. R. et al. (2021). Magnitude-sensitive reaction times reveal non-linear time costs in multi-alternative decision-making. *BioRxiv* Published online September 15, 2021. https://doi.org/10.1101/2021.05.05.442775.

[58] Kirkpatrick, R. P. et al. (2021). Equal evidence perceptual tasks suggest a key role for interactive competition in decision-making. *Psychol. Rev.* Advance online publication. https://doi.org/10.1037/rev0000284.

[59] Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics* 13, 37–58.

[60] Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cogn. Psychol.* 57, 153–178.

[61] Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108.

[62] Ratcliff, R. et al. (2016). Diffusion decision model: Current issues and history. *Trends Cogn. Sci.* 20, 260–281.

[63] Barron, A. B. et al. (2015). Decision-making and action selection in insects: inspiration from vertebrate-based theories. *Front. Behav. Neurosci.* 9, 216.

[64] Smith, P. L. and Van Zandt, T. (2000). Time-dependent poisson counter models of response latency in simple judgment. *Br. J. Math. Stat. Psychol.* 53, 293–315.

[65] Ratcliff, R. and Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356.

[66] Mormann, M. M. et al. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgm. Decis. Mak.* 5, 437–449.

[67] Drugowitsch, J. et al. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628.

[68] Sepulveda, P. et al. (2020). Visual attention modulates the integration of goal-relevant evidence and not value. *eLife* 9, e60705.

[69] Mormann, M. and Russo, J. E. (2021). Does attention increase the value of choice alternatives? *Trends Cogn. Sci.* 25, 305–315.

[70] Pirrone, A. and Gobet, F. (2021). Is attentional discounting in value-based decision making magnitude sensitive? *J. Cogn. Psychol.* 33, 327–336.

[71] Callaway, F. et al. (2021). Fixation patterns in simple choice reflect optimal information sampling. *PLOS Comput. Biol.* 17, e1008863.

[72] Krajbich, I. and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci. U.S.A.* 108, 13852–13857.

[73] Rich, E. L. and Wallis, J. D. (2016). Decoding subjective decisions from orbitofrontal cortex. *Nat. Neurosci.* 19, 973–980.

[74] Kacelnik, A. et al. (2011). Darwin's "tug-of-war" vs. starlings'"horse-racing": how adaptations for sequential encounters drive simultaneous choice. *Behav. Ecol. Sociobiol.* 65, 547–558.

[75] Brunton, B. W. et al. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340, 95–98.

[76] Lu, Z.-L. and Dosher, B. A. (2008). Characterizing observers using external noise and observer models: assessing internal representations with external noise. *Psychol. Rev.* 115, 44–82.

[77] Louie, K. et al. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci. U.S.A.* 110, 6139–6144.

[78] Bose, T. et al. (2020). Frequency-sensitivity and magnitude-sensitivity in decision-making: Predictions of a theoretical model-based study. *Comput. Brain Behav.* 3, 66–85.

[79] Bogacz, R. et al. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos. Trans. R. Soc. B* 362, 1655–1670.

[80] Strait, C. E. et al. (2014). Reward value comparison via mutual inhibition in ventromedial prefrontal cortex. *Neuron* 82, 1357–1366.

[81] Strait, C. E. et al. (2015). Signatures of value comparison in ventral striatum neurons. *PLOS Biol.* 13, e1002173.

[82] Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* 34, 333–359.

[83] Pastor-Bernier, A. et al. (2012). Dorsal premotor cortex is involved in switching motor plans. *Front. Neuroeng.* 5, 5.

[84] Louie, K. et al. (2011). Reward value-based gain control: divisive normalization in parietal cortex. *J. Neurosci.* 31, 10627–10639.

[85] Koyama, M. and Pujala, A. (2018). Mutual inhibition of lateral inhibition: a network motif for an elementary computation in the brain. *Curr. Opin. Neurobiol.* 49, 69–74.

[86] Roxin, A. and Ledberg, A. (2008). Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. *PLOS Comput. Biol.* 4, e1000046.

[87] Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968.

[88] Ditterich, J. et al. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* 6, 891–898.

[89] Busemeyer, J. R. and Townsend, J. T. (1992). Fundamental derivations from decision field theory. *Math. Soc. Sci.* 23, 255–282.

[90] Peterson, J. C. et al. (2021). Using large-scale experiments and machine learning to discover theories of human decision-making. *Science* 372, 1209–1214.

[91] Vasconcelos, M. et al. (2015). Irrational choice and the value of information. *Sci. Rep.* 5, 1–12.

[92] Seeley, T. D. et al. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* 335, 108–111.

[93] Cisek, P. et al. (2009). Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* 29, 11560–11571.

[94] Carland, M. A. et al. (2016). Evidence against perfect integration of sensory information during perceptual decision making. *J. Neurophysiol.* 115, 915–930.

[95] Trueblood, J. S. et al. (2021). Urgency, leakage, and the relative nature

of information processing in decision-making. *Psychological Review* 128, 160–186.

[96] Evans, N. J. et al. (2017). The computations that support simple decision-making: A comparison between the diffusion and urgency-gating models. *Sci. Rep.* 7, 1–13.

[97] Tajima, S. et al. (2016). Optimal policy for value-based decision-making. *Nat. Commun.* 7, 1–12.

[98] Pirrone, A. et al. (2021). Input-dependent noise can explain magnitude-sensitivity in optimal value-based decision making. *Judgm. Decis. Mak.* 16, 1221–1233.

[99] Sugrue, L. P. et al. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* 6, 363–375.

[100] Tajima, S. et al. (2019). Optimal policy for multi-alternative decisions. *Nat. Neurosci.* 22, 1503–1511.

[101] Bhui, R. (2019). Testing optimal timing in value-linked decision making. *Comput. Brain Behav.* 2, 85–94.

[102] Fudenberg, D. et al. (2018). Speed, accuracy, and the optimal timing of choices. *Am. Econ. Rev.* 108, 3651–84.

[103] Sellitto, M. et al. (2010). Myopic discounting of future rewards after medial orbitofrontal damage in humans. *J. Neurosci.* 30, 16429–16436.

[104] McNamara, J. M. and Houston, A. I. (2009). Integrating function and mechanism. *Trends Ecol. Evol.* 24, 670–675.

[105] Stephens, D. W. and Anderson, D. (2001). The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behav. Ecol.* 12, 330–339.

[106] Gershman, S. J. and Bhui, R. (2020). Rationally inattentive intertemporal choice. *Nat. Commun.* 11, 1–8.

[107] Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman.

[108] Parker, G. A. and Smith, J. M. (1990). Optimality theory in evolutionary biology. *Nature* 348, 27–33.

[109] Plonsky, O. and Erev, I. (2021). To predict human choice, consider the context. *Trends Cogn. Sci.* 10, 819–820.

[110] Nakagaki, T. et al. (2000). Maze-solving by an amoeboid organism. *Nature* 407, 470–470.

[111] Latty, T. and Beekman, M. (2011). Speed–accuracy trade-offs during foraging decisions in the acellular slime mould physarum polycephalum. *Proc. R. Soc. B* 278, 539–545.

[112] Boisseau, R. P. et al. (2016). Habituation in non-neural organisms: evidence from slime moulds. *Proc. R. Soc. B* 283, 20160446.

[113] Reid, C. R. et al. (2016). Decision-making without a brain: how an amoeboid organism solves the two-armed bandit. *J. R. Soc. Interface* 13, 20160030.

[114] Vogel, D. and Dussutour, A. (2016). Direct transfer of learned behaviour via cell fusion in non-neural organisms. *Proc. R. Soc. B* 283, 20162382.

[115] Shirakawa, T. et al. (2011). An associative learning experiment using the plasmodium of physarum polycephalum. *Nano Commun. Netw.* 2, 99–105.

[116] Reid, C. R. et al. (2012). Slime mold uses an externalized spatial "memory" to navigate in complex environments. *Proc. Natl. Acad. Sci. U.S.A.* 109, 17490–17494.

[117] Latty, T. and Beekman, M. (2011). Irrational decision-making in an amoeboid organism: transitivity and context-dependent preferences. *Proc. R. Soc. B* 278, 307–312.

[118] Vallverdú, J. et al. (2018). Slime mould: the fundamental mechanisms of biological cognition. *Biosystems* 165, 57–70.

[119] Boussard, A. et al. (2021). Adaptive behaviour and learning in slime moulds: the role of oscillations. *Proc. R. Soc. B* 376, 20190757.

[120] Couzin, I. D. (2009). Collective cognition in animal groups. *Trends Cogn. Sci.* 13, 36–43.

[121] Marshall, J. A. R. and Franks, N. R. (2009). Colony-level cognition. *Curr. Biol.* 19, R395–R396.

[122] Trianni, V. et al. (2011). Swarm cognition: an interdisciplinary approach to the study of self-organising biological collectives. *Swarm Intell.* 5, 3–18.

[123] Passino, K. M. et al. (2008). Swarm cognition in honey bees. *Behav. Ecol. Sociobiol.* 62, 401–414.

[124] Conradt, L. and List, C. (2009). Group decisions in humans and animals: a survey. *Philos. Trans. R. Soc. B* 364, 719–742.

[125] Bose, T. et al. (2017). Collective decision-making. *Curr. Opin. Behav. Sci.* 6, 30–34.

[126] Solé, R. et al. (2019). Liquid brains, solid brains. *Philos. Trans. R. Soc. B* 374, 20190040–20190040.

[127] Pratt, S. C. et al. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant leptothorax albipennis. *Behav. Ecol. Sociobiol.* 52, 117–127.

[128] Reina, A. et al. (2015). A design pattern for decentralised decision making. *PLOS ONE* 10, e0140950.

[129] Reina, A. et al. (2015). A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intell.* 9, 75–102.

[130] Talamali, M. S. et al. (2019). Improving collective decision accuracy via time-varying cross-inhibition. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 9652–9659). IEEE.

[131] Hamann, H. (2018). *Swarm robotics: A formal approach*. Springer.

[132] Dorigo, M. et al. (2020). Reflections on the future of swarm robotics. *Sci. Robot.* 5, eabe4385.

[133] Frohnwieser, A. et al. (2016). Using robots to understand animal cognition. *J. Exp. Anal. Behav.* 105, 14–22.

[134] Mitri, S. et al. (2013). Using robots to understand social behaviour. *Biol. Rev.* 88, 31–39.

[135] Garnier, S. (2011). From ants to robots and back: How robotics can contribute to the study of collective animal behavior. In *Bio-inspired self-organizing robotic systems* (pp. 105–120). Springer.