

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A distributed-consensus mechanism of decision making explains economically irrational behaviours

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1695730> since 2020-02-21T16:14:25Z

Published version:

DOI:10.1016/j.anbehav.2019.04.016

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

A distributed-consensus mechanism of decision making explains

2 economically irrational behaviours

Abstract

4 In humans, as in many other animals, preferences between options can be reversed when
other irrelevant options are added to the choice set. Heuristics theories view such puzzling
6 departure from economic rationality as evidence that decisions rely on simple rules that make
a context-dependent use of the available information. The computational mechanisms
8 underlying these rules, however, remain largely unresolved. Using a “sequential-sampling”
model of decision making, I show that irrational decisions may arise when an information-
10 processing mechanism that works optimally in one-choice tasks is co-opted in multiple-
choice contexts. The model supports the assumption that different heuristics may sometimes
12 be the elusive expression of a single general mechanism and that natural selection, rather than
promoting the evolution of different mechanisms and rules, may favour the parsimonious use
14 of bounded computational resources.

16 KEYWORDS

Decision-making, Decoy effect, Evo-mecho, Heuristics, Irrational behaviour, Rule-of-thumb,
18 Sequential sampling models.

20 Introduction

Humans, as many other animals, make decisions by integrating different sources of
information about the expected benefits and costs of the available options. Economic
theories of consumer behaviour assume that decisions maximize the chooser's perceived
utility (Simon 1955). Evolutionary theory predicts a positive association between the
perceived utility of an option and its expected consequences on the Darwinian fitness of the
chooser (Parker, Maynard Smith 1990). While the option may vary along multiple
dimensions, the perceived utility is a one-dimensional representation of its economic value
(Kacelnik 2006). Decisions that optimize economic utilities obey the two axioms of
economic rationality: transitivity (if A is preferred to B and B to C, then A is preferred to C)
and regularity (the probability of choosing A over B is the same, independent of the presence
or the absence of a third option, C) (Kacelnik 2006).

Humans, and other animals as well, however, often show cognitive biases, which cause
systematic departures from economic rationality (Tversky and Kahneman 1974). A well-
known example is the “decoy effect”, whereby the preference between two options depends
on the presence of a third option (Huber, Payne and Puto 1982, Lea and Ryan 2015).
According to some authors, these violations, though economically irrational and sub-optimal,
could be rational and optimal from an evolutionary point of view (Kacelnik 2006). For
example, models of optimal foraging show that intransitive choices and decoy effects can be
favoured by natural selection, if the fitness benefits they provide depend on either the spatial
(Houston et al. 2007, Trimmer 2013) or the temporal patterns (McNamara et al. 2014) of the

available options. According to this view, cognitive biases are evidence that decisions do not
42 rely on complex and accurate representations of perceived utilities, but on “fast and frugal”
heuristics (or rules of thumb), which may not perform optimally in all situations, but
44 satisfactory in the ecological contexts where decisions are actually made (Gigerenzer and
Goldstein 1996).

46 Both optimality and heuristic models focus on the decision rules. The optimality models first
define the decision problem and the currency that the decision maker is expected to
48 maximize, then, they search for the rule that maximizes that currency, often, under the
assumption that decision makers are omniscient, fully-rational agents (Stephens and Krebs
50 1986). The heuristic models, in contrast, relax the rational assumption and search for rules
that can make “satisficing” decisions, with a parsimonious use of time, information, and
52 computation. Heuristics are not optimal, but approximate optimality. They are not strictly
rational, but ecologically rational, because they exploit the environmental structure to yield
54 simple, efficient rules, which could be plausibly implemented by rationally-bounded agents.
Despite these relevant differences, both optimality and heuristic models explicitly ignore the
56 computational mechanisms of decision rules and this neglect of mechanisms has been argued
to have impoverished the functional analyses of behaviours (McNamara and Houston 2009,
58 Kacelnik 2012). For example, Kacelnik et al. (2011) show that when European starlings are
given a choice between two foraging options, they adopt the same computational mechanism
60 independent of whether the alternatives are presented sequentially or simultaneously. These
results not only conflict with the “toolbox” hypothesis of heuristics as narrow-domain
62 decision devices (Gigerenzer, Todd 1999), but they also suggest that mechanisms are
important constraints in the evolution of optimal decision rules and, thus, that they should

play also an important role in the functional studies of behavioural flexibility (McNamara and Houston 2009).

Here I present a race model of decision making, which describes decisions as the bounded accumulation of noisy sensory information (Bogacz 2006, Ratcliff et al. 2016, Vickers 1970).

Following Kacelnik et al. (2011), I also assume the choice mechanism to work in both the sequential and the comparative contexts. Most race models assume that all types of

supporting evidence for an option are integrated into a single accumulator, which thus provides a one-dimension, dynamic representation of the option value (Ratcliff et al. 2016).

My model relaxes this assumption and, consistently with some recent neurophysiological studies of action choice (Cisek and Kalaska 2010, Cisek 2012, Engel et al. 2013), it assumes

the parallel, multi-dimension processing of the option attributes. The decision mechanism is equivalent to a race in a multi-dimension plane, which terminates as soon as the accumulating

evidence enters either a rejection or an acceptance consensus area. This general mechanism can explain decisions in both one-, two-, and multiple-choice contexts. However, while in

one- and two-choice contexts, the mechanism predicts economically rational decisions, in multi-choice contexts, it introduces systematic errors and predicts irrational choices, such as

those described by the “decoy effect”.

The model

The model considers the behaviour of a hypothetical insect, which is hovering above the grass, collecting nectar from flowers. The insect spots a bunch of brightly coloured flowers

just a few meters on the left and it decides to turn left. While approaching the bunch, its attention is captured by a yellow daisy on the left and a purple poppy on the right. The insect

86 decides to turn right and to land on the poppy. The behaviour of this foraging insect is the
result of two types of choice. The first choice depends only on the quality of the bunch of
88 flowers being inspected, which is considered to be worthy of further inspection. I call this
“sequential choice” (Kacelnik et al. 2011), because it is the type of decisions that an animal
90 makes when it encounters options sequentially. In contrast, the second type of decision
depends on the relative quality of the purple poppy, which is considered more worthy than
92 the yellow daisy. I call it “simultaneous choice”. The model assumes that the sequential and
the simultaneous choice are not the expression of different heuristics, but of the same
94 cognitive machinery designed by natural selection to make both types of decision effectively.
This cognitive machinery processes sensory information and provides an internal
96 representation of the benefits and costs of either exploiting or rejecting the prospective
resource.

98 Benefits and costs may be viewed as the two axes of a Cartesian information plane (Fig. 1):
the y-axis describes the perceived benefits; the x-axis describes the perceived costs. Benefits
100 and costs are a function (either a linear, monotonic or unimodal) of the characteristics of the
prospective resource. For example, the benefits depend on the quantity and quality of the
102 nectar provided by the flower, whereas the costs depend on the time required to collect the
nectar, which depends on the type and the location of the flower and on the presence of other
104 competitors. Benefits and costs, however, should not be viewed as the internal representation
of the economic value of the resource, but as the supporting evidence for either exploiting or
106 rejecting it (see also Appendix 1).

During the decision process, the insect accumulates over time the noisy sensory information
108 about the costs and benefits of an action, which is executed only when the accumulating

evidence reaches a given threshold. From a mathematical point of view, the dynamic of
 110 decision making can be described as a random-walk process in the two-dimensional
 informational plane, with transition probabilities that depend only on the type of stimulus
 112 processed (Forstmann, Ratcliff and Wagenmakers 2016, Ratcliff et al. 2016). Depending on
 the choice context, the insect can assess a single option at a time (sequential choice) or
 114 process simultaneously and in a parallel fashion several options (simultaneous choice).

First, I consider the dynamic of the decision making in a sequential-choice context, for
 116 example, when the insect spots the brightly coloured flowers on its left and decides to visit
 them. During the evaluation of action A (to visit the flowers), the insect is assumed to collect
 118 a sequence of noisy information about the amount of energy it is expected to obtain from the
 flowers $(Q^A(1), Q^A(2), \dots, Q^A(n))$ and the amount of time it is expected to spend for
 120 handling them $(H^A(1), H^A(2), \dots, H^A(n))$. Since information is noisy, Q and H are assumed
 to be two normally-distributed stochastic variables with parametric means q^A and h^A and
 122 covariance matrix $G = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix}$. By transforming and integrating over time each piece of
 sensory information, the insect obtains a sequence of accumulating evidence for the benefits
 124 and the costs of the evaluated action. After n samples, the evaluated action will be
 represented in the informational plane by a point with coordinates $[B^A(n), C^A(n)]$, where
 126 $B^A(n) = B^A(n-1) + b(Q^A(n))$ represents the amount of evidence for choosing the option
 and $C^A(n) = C^A(n-1) + c(H^A(n))$ represents the amount of evidence for rejecting it. In
 128 Appendix 1, I provide a formal description of the functions $b(Q)$ and $c(H)$, that is, of how
 sensory information is converted into the perceived benefits and the perceived costs of an
 130 action.

Further and most important, the model assumes that the benefits and the costs of an action interact additively in determining its overall value (see Appendix 1). Suppose that, under certain conditions, there exist two optimal decision thresholds, T^* and L^* (with both $T > 0$ and $L > 0$), so that the insect would maximize its long term rate of energy intake by exploiting the resource if the amount of supporting evidence is larger than T^* or by rejecting the resource if it is lower than $-L^*$. One decision mechanism could be simply to compute the difference between benefits and costs and to use this new variable as the decision variable: the insect would exploit the resource if $B^A(n) - C^A(n) \geq T^*$, it would reject the resource if $B^A(n) - C^A(n) \leq -L^*$, or it would continue the assessment if none of the two previous conditions is satisfied. From a mathematical point of view, the decision mechanism is a 1-dimensional random walk (1-DRW), because for each time step of equal length, the decision variable can only increase (if the perceived benefits overcome the costs) or decrease (if the perceived costs overcome the benefits).

There exists, however, an alternative and functionally equivalent mechanism that saves the computational costs of computing the difference between benefits and costs. Mathematically speaking, the alternative mechanism is a 2-Dimension Random Walk (2-DRW). Each dimension has a lower (θ_C^l, θ_B^l) and an upper (θ_C^u, θ_B^u) threshold, which is dynamically adjusted during the evaluation of an option according to the following rules:

$$\begin{aligned}
 \theta_B^u(n) &= T + C^A(n) \\
 \theta_C^l(n) &= -T + B^A(n) \\
 \theta_B^l(n) &= -L + C^A(n) \\
 \theta_C^u(n) &= L + B^A(n)
 \end{aligned}
 \tag{Eq. 1}$$

150 The insect would execute the action when BOTH $B^A(n) \geq \theta_B^u(n)$ and $C^A(n) \leq \theta_C^l(n)$, that is, as soon as the random walk reaches the threshold line $b = T + c$. The insect would reject the resource when BOTH $B^A(n) \leq \theta_B^l(n)$ and $C^A(n) \geq \theta_C^u(n)$ (Fig. 1a, see also the Jupyter notebook in the supplementary materials). The two parameters (T, L) of Eq. 1 describe the insect's decision strategy and directly affect its fitness, by influencing both the response times and the error probabilities.

156 Now, let us consider the decision process when the insect has to choose between two ($k = 2$) or more ($k > 2$) alternative actions. For example, suppose the insect has to decide whether to land on the yellow daisy or on the purple poppy (Fig. 3a and the Jupyter Notebook in the supplementary materials). In this case, the 1-DRW assumes, first, that the insect computes, for each option, the difference between benefits and costs; then, that it integrates over time these differences into two decision variables; and, finally, that it chooses the resource, whose decision variable first reaches the acceptance threshold T . Unlike the 1-DRW, the 2-DRW assumes that evidence for either the D-action (approach the daisy) or the P-action (approach the poppy) is accumulated in two decision vectors, $[B^D, C^D]$ and $[B^P, C^P]$, and that the decision depends on which of them first enters the “consensus” area, where $B \geq \theta_B^u$ and $C \leq \theta_C^l$. While in one-choice context, the thresholds defining the consensus area are a function of the accumulating benefits and costs of the single assessed resource (Eqn. 1), in two- ($k = 2$) and in multiple-choice contexts ($k > 2$), the thresholds are assumed to be a function of the average benefits and costs of all the assessed resources. Specifically, the insect is assumed to dynamically adjust the four thresholds according to the following set of equations:

$$\begin{aligned}
\theta_B^u &= T + \bar{C}(n) \\
\theta_C^l &= -T + \bar{B}(n) \\
\theta_B^l &= -L + \bar{C}(n) \\
\theta_C^u &= L + \bar{B}(n)
\end{aligned}
\tag{Eq. 2}$$

172 Where $\bar{B}(n) = \frac{1}{k} \sum_i^k B_i(n)$ and $\bar{C}(n) = \frac{1}{k} \sum_i^k C_i(n)$ are the average benefits and costs after having processed n samples of the k resources.

174 Results

In one-choice decision tasks, the 1-DRW and 2-DRW are mathematically equivalent, because
176 a 2-dimensional random-walk described by two vectors X and Y , with threshold line $y = T + x$ can be reduced to a 1-dimensional random walk, described by the vector $Z = (X - Y)/\sqrt{2}$, with
178 threshold $T' = T/\sqrt{2}$. From a computational point of view, this means that 1-DRW and DCM are equivalent both in terms of error probabilities and of average number of observations
180 needed to make a decision (Fig. 1b).

In contrast, in two-choice tasks, the 2-DRW is less efficient than the 1-DRW, because, for a
182 given level of accuracy, it takes longer to make a decision. Fig. 2 shows results of two series of simulations, which use the 1-DRW and the 2-DRW to choose between a high- and a low-
184 quality resource. Both the 1-DRW and the 2-DRW simulations predict preferences for the high-quality option to increase asymptotically with the increase of T and, consequently, with
186 the increase of the decision time. However, the 1-DRW performs better than the 2-DRW, because at any decision time, it predicts stronger preferences and, thus, more accurate
188 decisions. Furthermore, in two-choice contexts, the 2-DRW introduces a weak bias in the pattern of choice, as shown by the divergent pattern of the iso-probability lines of Figure 3b.

190 When a focal option (i.e. the poppy, represented by the star marker of Fig. 3b, with Benefits
 $b(q^P) = 5$ and Costs $c(h^P) = -1$) is compared against an alternative of equal overall value
 192 (i.e. the daisy, with $b(q^D) - c(h^D) = 6$), the 2-DRW algorithm predicts always random
 choice (i.e. the slope of the 0.5 iso-probability line of Fig. 3b is 1). However, when the
 194 alternatives show different overall values ($b(q^D) - c(h^D) \neq 6$), preferences depend on the
 $b(q^D)$ and $c(h^D)$ and not just on their sum. For example, in the simulations of Fig. 3b, when
 196 the alternative option has $b(q^D) = 0$ and $c(h^D) = -4$, the probability that it is chosen is
 0.13, but it is twice as large (preference = 0.26) when it has $b(q^D) = 4$ and $c(h^D) = 0$.
 198 The performance of 2-DRW decreases even more when the number of items to choose among
 increases. Fig. 4a, b show what happens when a third option, a lily (L), is added to the choice
 200 set. In this example, the poppy and the daisy differ in both benefits and costs, but have the
 same economic value (i.e. the same difference between benefits and costs): $b(q^P) =$
 202 $-c(h^D) = 5$ and $b(q^D) = -c(h^P) = 1$. For this reason, in a dyadic choice, the insect does
 not show any preference either for the poppy or for the daisy. To simulate the decoy effect, I
 204 further assume the lily to be asymmetrically dominated by the daisy, that is, $c(h^L) =$
 $c(h^D) - \delta$, with $0 < \delta \ll c(h^D)$, and $b(q^L) = b(q^D)$.
 206 When the lily is introduced in the set of the available options, the positions of both the
 decision thresholds and the consensus area change. In the two-choice test, the accumulating
 208 evidence for either the daisy or the poppy causes the benefit threshold to decrease and the
 cost threshold to increase at a similar rate. As a consequence, during the process, the
 210 minimum distance between the consensus area and either the poppy or the daisy was similar,
 and the two options had the same probability of winning the race. With the lily, however, the
 212 two thresholds no longer change at similar rates and the consensus area no longer maintains

the same minimum distance from the options (see also the Jupyter notebook in the

supplementary materials). In fact, since the lily has very low costs and moderate energetic

benefits, during the assessment, the benefit threshold (θ_B^u) decreases faster than the cost

threshold (θ_C^u), because $T + \bar{C}(n) < T - \bar{B}(n)$. For this reason, with the lily, the option with

the lowest costs (the daisy) is closer to and more likely to enter the consensus area than the

option with the highest benefits and costs (the poppy) (Fig. 4a).

Discussion

This model follows the long tradition of sequential-sampling models in theoretical

psychology (Busemeyer and Townsend 1993, Ratcliff et al. 2016, Vickers 1970), in that it

views decision making as a process of noisy accumulation of evidence over time. According

to these models, decision makers use sensory information of the available options to build a

dynamic internal representation of their perceived value (i.e. economic “utility”). They make

options to compete against each other and choose the option whose perceived value first

reaches a strategically-set decision threshold. The competition can be in the form of either a

“race” (if options are valued independently of the others) or a “tug-of-war” (if options are

valued relatively to the others) (Kacelnik et al. 2011). In both types of models, the perceived

values are assumed to be one-dimensional and the decision rules to be context-independent,

in that they are strategically adjusted before the process, but kept fixed throughout. The one-

dimension assumption is a sufficient (but not necessary) condition for economic rational

choices and for the two axioms of rationality, transitivity and regularity. Since humans and

other animals (both vertebrates and invertebrates) often choose irrationally, some authors

have proposed modified versions of the “tug-of-war” model, with ad-hoc (non-linear)

computational mechanisms that could reconcile the one-dimension assumption with the
236 reversal effects in multi-alternative, multi-attribute choice (review in Tsetsos et al. 2010).

Although my model shares some similarities with the “race” models, in that it assumes the
238 perceived values to be context independent, it shows two important differences. First, the
decision rules are context-dependent, because the thresholds are dynamically adjusted during
240 the decision process. Second, the internal representation is multi-dimensional, because
decision makers are not assumed to integrate the option attributes into a one-dimension
242 representation of its value, but to accumulate different kinds of evidence along different
levels of representations. For these reasons, the decision process might be no longer viewed
244 as a one-direction race against a fixed finish line, but as a multi-direction race towards a
consensus area, whose boundaries vary dynamically during the process. This view is
246 consistent with some recent findings in the neurophysiology of decision making: for example,
Cisek (2012) reviews neurophysiological evidence that sensorimotor neurons are not simply
248 the executors of a command, but the main actors of the deliberative process of decision
making. According to this view, alternative actions are simultaneously represented in the
250 sensorimotor regions of the brain and compete against each other for execution (Gallivan et
al. 2015). The competition occurs at multiple levels and involves the parallel processing of
252 different types of information, such as the subjective benefits and costs of the options, the
biomechanical costs of actions, and the geometric relationships between the options (Cisek
254 and Pastor-Bernier 2014). The competition terminates when the different levels of
representations (i.e. different axes of the decision plane) reach a consensus.

256 A distributed-consensus mechanism of decision-making requires thresholds to be
dynamically adjusted during the decision process. In one-choice decision tasks, where the

options are assessed in terms of their intrinsic perceived benefits and costs, the optimal calibration can be achieved by simple inhibitory mechanisms (see eq. 1), which guarantee that decision is made as soon as the difference between benefits and costs is greater than a given (optimal) threshold. This independent dynamic adjustment of thresholds makes the 2-DRW functionally equivalent to the 1-DRW model, because it requires the same amount of evidence to make a decision. From a computational point of view, however, the two models are not equivalent, because the 2-DRW model is fully described by two variables (B , benefit, and C , costs), whereas the 1-DRW requires a third variable, which is the difference of the other two. One may be tempted to speculate that the 2-DRW algorithm might have been favoured by natural selection over the 1-DRW alternative, because of its greater computational (and neuro-anatomical) parsimony: in fact, the 2-DRW can work as efficiently as the 1-DRW, but with fewer computational units (i.e. neurons).

In this paper, I have explored the hypothesis that the 2-DRW model, which performs optimally in one-choice decision tasks, might have been co-opted in two-choice contexts, where it still works satisfactorily, and in multiple-choice contexts, where it sometimes causes irrational behaviour. According to this hypothesis, thus, there exists a single, general mechanism of decision making. Sequential and simultaneous choice are not the expression of different heuristics, but of the same flexible mechanism, used under different contexts. This mechanism is not merely computational, but psychological, because it describes the process of decision making. If we are to understand the functional role of flexible behaviour, in fact, we need to study the evolution of the psychological mechanisms that control behavioural flexibility (McNamara and Houston 2009, Castellano 2015, Castellano and Cermelli 2011, 2015). I argue that heuristic theories are limited in this respect, because they neglect the

mechanisms underlying the decision process. Although heuristics have inspired much
theoretical and empirical work on the rules that govern behavioural flexibility, it is time to
move on. Heuristics describe the decision rules, but not their underlying computational
mechanisms. They are in the mind of the beholder, not in the mind of the decision maker. In
the recent years, our understanding of the neurophysiology of decision making has advanced
to the point where evolutionary theories can be extended beyond the normative models of
optimal behaviours (Barron et al. 2015, Mobbs et al. 2018). We are now at the point where
we can address questions about the optimal mechanisms of information processing and, in the
process, achieve a much deeper understanding of the evolution of behavioural flexibility.

Competing interests

I have no competing interests.

References

Barron, A. B., Gurney, K. N., Meah, L. F. S., Vasilaki, E., & Marshall, J. A. R. (2015).

Decision making and action selection in insects: inspiration from vertebrate-based
theories. *Frontiers in Behavioral Neuroscience*, 9, 216.

Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour.
Trends in Cognitive Sciences, 11, 118-125.

Bussemeyer, J. R. , & Townsend, J. T. (1993). Decision Field Theory - A dynamic cognitive
approach to decision-making in an uncertain environment. *Psychological Review*,
100, 432-459.

Castellano, S. (2015). Bayes' rule and bias roles in the evolution of decision making.
302 *Behavioral Ecology*, 26, 282-292.

Castellano, S., & Cermelli, P. (2011). Sampling and assessment accuracy in mate choice: A
304 random-walk model of information processing in mating decision. *Journal of*
 Theoretical Biology, 274, 161-169.

306 Castellano, S., & Cermelli, P. (2015). Prey's exploitation of predator's fear: when the
 caterpillar plays the Gruffalo. *Proceedings of the Royal Society B: Biological*
308 *Sciences*, 282, 20151786.

Cisek, P. (2012). Making decisions through a distributed consensus. *Current Opinion in*
310 *Neurobiology*, 22, 927-936.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of
312 action choices. *Annual Review of Neuroscience*, 33, 269-298.

Engel, A., Maye, A., Kurthen, M., & Konig, P. (2013). Where's the action? The pragmatic
314 turn in cognitive science. *Trends in Cognitive Sciences*, 17, 202-209.

Forstmann, B. U., Ratcliff, R., & Wagenmakers, E. J. (2016). Sequential Sampling Models in
316 Cognitive Neuroscience: Advantages, Applications, and Extensions. *Annual Review of*
 Psychology, 67, 641-666.

318 Gallivan, J. P., Barton, K. S., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2015).
 Action plan co-optimization reveals the parallel encoding of competing reach
320 movements. *Nature Communications*, 6, 7428.

Gigerenzer, G., & Goldstein, G. D. (1996). Reasoning the fast and frugal way: models of
322 bounded rationality. *Psychological Review*, 103, 650–669.

- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: the adaptive toolbox. In G.
 324 Gigerenzer G, P. M. Todd, & ABC Research Group (Eds.), *Simple Heuristics that
 make us smart* (pp. 3-34). Oxford: Oxford University Press.
- 326 Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives:
 violations of regularity and similarity hypothesis. *Journal of Consumer Research*, 9,
 328 90–98.
- Houston, A. I, McNamara, J. M., & Steer, M. (2007). Violations of transitivity under fitness
 330 maximization. *Biology Letters*, 3, 365–367.
- Kacelnik, A. (2006). Meanings of rationality. In M. Nudds, & S. Hurley (Eds). *Rational*
 332 *animals?* (p. 87-106). Oxford: Oxford University Press.
- Kacelnik, A. (2012). Putting mechanisms into behavioral ecology. In P. Hammerstein, & S. J.
 334 Stevens (Eds). *Evolution and the mechanisms of decision making* (p. 21-38).
 Cambridge, MA: MIT Press.
- 336 Kacelnik, A., Vasconcelos, M., Monteiro, T., & Aw, J. (2011). Darwin’s “tug-of-war” vs.
 starlings’ “horse-racing”: how adaptations for sequential encounters drives
 338 simultaneous choice. *Behavioral Ecology and Sociobiology*, 65, 547-558.
- Lea, A. M., & Ryan, M. J. (2015). Irrationality in mate choice revealed by Tungara frogs.
 340 *Science*, 349, 964-966.
- McNamara, J. M., & Houston, A. I. (2009). Integrating function and mechanism. *Trends in*
 342 *Ecology and Evolution*, 24, 670-675.
- McNamara, J. M., Trimmer, P. C., & Houston, A. I. (2014). Natural selection can favour
 344 ‘irrational’ behaviour. *Biology Letters*, 10, 20130935.

- Mobbs, D., Trimmer, P. C., Blumstein, D. T., & Dayan, P. (2018). Foraging for foundations
 346 in decision neuroscience: insights from ethology. *Nature Reviews Neuroscience*, 19,
 419-427.
- 348 Parker, G. A., & Maynard Smith, J. (1990). Optimality theory in evolutionary biology.
Nature, 348, 27-33.
- 350 Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision models:
 Current issues and history. *Trends in Cognitive Sciences*, 20, 260-281.
- 352 Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of*
Economics, 69, 99-118.
- 354 Stephens, D. W., & Krebs, J. K. (1986). *Foraging theory*. Princeton, NJ: Princeton University
 Press.
- 356 Trimmer, P. C. (2013). Optimal behaviour can violate the principle of regularity. *Proceedings*
of the Royal Society B: Biological Sciences, 280, 1763.
- 358 Tsetsos, K., Usher, M., & Chater, N. (2010). Preference Reversal in Multiattribute Choice.
Psychological Review, 117, 1275-1293.
- 360 Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases.
Science, 185, 1124-1131.
- 362 Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination.
Ergonomics, 13, 37-58.

364 Appendix 1

366 The insect decides to exploit a resource if the perceived benefits are greater than the
 368 perceived costs. Here I explain how sensory information about a flower is converted into the
 368 perceived benefits and costs of an action (i.e. exploiting or avoiding the resource).

Within the inspecting patch, the predation risk is assumed to be independent of the types of
 370 flower. It follows, thus, that the costs of a resource are only due to the loss of opportunities.
 Suppose that the insect can assess the energetic value of a resource (q_i) and its handling time
 372 (h_i) and that it knows which is the average quality (\bar{q}), the average handling time (\bar{h}), and the
 average searching time (\bar{s}) of the resources within the patch. Under these assumptions, the
 374 insect is assumed to exploit a resource i if its rate of energy intake (q_i/h_i) is greater than that
 expected if it decided to reject it and to search for the next resource.

$$376 \quad \frac{q_i}{h_i} > \frac{\bar{q} - \gamma \bar{s}}{\bar{s} + \bar{h}} \quad \text{Eq. A1}$$

The right hand side of the inequality represents the opportunity costs: in the numerator, there
 378 is the expected quality of the $(i + 1)$ resource (that is, \bar{q} , the average quality of the patch),
 discounted by the average cost of searching ($\gamma \bar{s}$) (where γ , a constant, is the cost of searching
 380 per unit time). In the denominator, there is the expected time for exploiting the $(i + 1)$
 resource, which is the sum of the average searching (\bar{s}) and handling (\bar{h}) times.

382 From Eq. A1, it derives that the expected value (in units of energy) of the i resource is the
 difference between what the insect is expected to gain and to lose if it decided to exploit the
 384 resource:

$$e_i = q_i - \frac{(\bar{q} - \gamma \bar{s})}{\bar{s} + \bar{h}} h_i \quad \text{Eq. A2}$$

Eqn. A2 can be rescaled in order to obtain an a -dimensional representation of the perceived economic value of the i resource:

$$u_i = \frac{e_i}{(\bar{q} - \gamma \bar{s})} = \frac{q_i - (\bar{q} - \gamma \bar{s})}{(\bar{q} - \gamma \bar{s})} - \frac{h_i - (\bar{s} + \bar{h})}{\bar{s} + \bar{h}}. \quad \text{Eq. A3}$$

The 1-DRW and the 2-DRW assume that the noisy sensory information about the energetic quality and the handling time of a resource is converted into the perceived benefits and costs by two linear filters, $b(Q)$ and $c(H)$:

$$b(Q) = \frac{Q - (\bar{q} - \gamma \bar{s})}{(\bar{q} - \gamma \bar{s})}, \quad \text{Eq. A4a}$$

$$c(H) = \frac{H - (\bar{s} + \bar{h})}{\bar{s} + \bar{h}}, \quad \text{Eq. A4b}$$

where Q and H are the two normally-distributed stochastic variables with expectations q and h , and standard deviation σ (see main text). Notice that from Eqn. A1 it derives that b and c are linear functions of the incoming sensory information, It might be emphasized, however, that linearity between sensation and perception is by no means a necessary condition of the model (see for example Castellano 2015) The b and the c functions convert different types of sensory information into the same currency (patterns of neural stimulation) and it is on this common ground that different features of one, or more than one, option can be compared.

FIGURE CAPTIONS

Figure 1. (a) Graphical representations of the two-dimension random-walk process of decision making (2-DRW) under a one-choice decision task. Noisy information about the expected benefits and the costs of an action is accumulated over time into two separate counters, which describe the position of the action in the informational plane. The action is executed when it enters the green zone and discarded when it enters the red zone. During the assessment, the thresholds that define the acceptance (θ_C^u, θ_B^u) and the rejection (θ_C^l, θ_B^l) areas are dynamically adjusted, as described by Eqn. 1 (see main text). The plot shows the acceptance/rejection areas in three moments of the decision process, which terminates with the execution of the action (when the random walk enter the green area). The dynamic adjustment of the two thresholds minimizes the response time, because it forces the process to stop as soon as the blue path touches the black threshold. (b) A contour plot of the preferences for a resource as a function of its costs and the benefits. Costs and benefits interact additively on the preferences. If the benefits are greater than the costs, the option would be positively selected ($P > 0.5$). If costs are greater than benefits, the option would be negatively selected ($P < 0.5$). The strength of selection is predicted to depend on only the difference between benefits and costs and to be independent of their absolute values.

Figure 2. A comparison of the speed-accuracy trade-offs in the one- (1-DRW, filled circles) and in the two-dimension random walk models (2-DRW, open circles), during a series of two-choice discrimination tests between a high-quality ($c(h_{\text{high}}) = 0.5, b(q_{\text{high}}) = 0.5$) and a low-quality resource ($c(h_{\text{low}}) = 0.0, b(q_{\text{low}}) = 0.0$), in which the acceptance ($T = \theta_B^u - \theta_C^u$) and the rejection ($L = \theta_B^l - \theta_C^l = -T$) decision thresholds were let to vary between 0.1 and 3 by

steps of 0.05. Each point is associated to a threshold value and shows the average decision
 426 latency and the preference for the high-quality option after running 10,000 random-walk
 simulations. Preferences for the high-quality option increase with the increasing latencies
 428 (that is, with the increasing T), but, at any decision latency, 1-DRW performs better than 2-
 DRW.

Figure 3. (a) A graphical representation of the bi-dimension random-walk process of
 432 decision making in a two-choice decision task. The process is the same as in Fig.1, but now it
 is run in parallel on the star and the black-circle options. The black-circle option (D, the
 434 daisy) has lower expected costs and lower perceived benefits ($[c(h^D) = -5, b(q^D) = 1]$)
 than the star option (P, the poppy) ($[c(h^P) = -1, b(q^P) = 5]$), so that their overall values
 436 (the difference between benefits and costs) are the same. Since the acceptance thresholds
 along the cost and benefit axes increase with the increase, respectively, of the average costs
 438 ($\bar{C}(n)$) and the average benefits ($\bar{B}(n)$), they oscillate midway between the optimal values of
 the two conflicting options. As a consequence, the 2-dimension random walk needs a larger
 440 amount of evidence and much more time to enter the consensus area. (b) Contour plot of the
 two-choice preference for an option, when compared against the poppy (star). The 2-DRW
 442 introduces a slight bias in the preference pattern. In fact, the slope of the iso-probability lines
 is one only when the difference between the benefits and costs of an option equals that of the
 444 poppy, as shown by the daisy (black circle).

Figure 4. (a) A bi-dimension random-walk process in a three-choice decision task. The third
 446 option (the lily, grey circle, $c(h^L) = -4.8, b(q^L) = 1$) is asymmetrically dominated by the

448 daisy, (the black-circle option, $c(h^L) = -5, b(q^L) = 1$). Its introduction modifies the
position of the consensus areas, biasing the preferences to the advantage of those items that
450 have the lowest costs. While in two-choice test, the daisy is as likely as the poppy to be
chosen by the insect, in the three-choice context, the daisy is strongly preferred over the
452 poppy. (b) Contour plot of the preferences in three-choice discrimination tests. The iso-
probability lines show the preferences for a given option (for example, the daisy, represented
454 by the black circle) with respect to a fixed alternative (the poppy, represented by the star
marker), when a third alternative (the decoy represented by the grey circle) is also present.
456 Due to the characteristics of the decoy (very low costs and moderate benefits), small changes
along the cost dimension have stronger effects on the preferences than changes of similar
458 magnitude along the benefit dimension.