OPEN ACCESS Freely available online

PLOS ONE

A Mechanism for Value-Sensitive Decision-Making

Darren Pais[1], Patrick M. Hogan[2], Thomas Schlegel[3], Nigel R. Franks[3], Naomi E. Leonard[1], James A. R. Marshall[2]*

1 Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, New Jersey, United States of America, **2** Department of Computer Science and Kroto Research Institute, University of Sheffield, Sheffield, South Yorkshire, United Kingdom, **3** School of Biological Sciences, University of Bristol, Bristol, Bristol, United Kingdom

Abstract

We present a dynamical systems analysis of a decision-making mechanism inspired by collective choice in house-hunting honeybee swarms, revealing the crucial role of cross-inhibitory 'stop-signalling' in improving the decision-making capabilities. We show that strength of cross-inhibition is a decision-parameter influencing how decisions depend both on the difference in value and on the mean value of the alternatives; this is in contrast to many previous mechanistic models of decision-making, which are typically sensitive to decision accuracy rather than the value of the option chosen. The strength of cross-inhibition determines when deadlock over similarly valued alternatives is maintained or broken, as a function of the mean value; thus, changes in cross-inhibition strength allow adaptive time-dependent decision-making strategies. Cross-inhibition also tunes the minimum difference between alternatives required for reliable discrimination, in a manner similar to Weber's law of just-noticeable difference. Finally, cross-inhibition tunes the speed-accuracy trade-off realised when differences in the values of the alternatives are sufficiently large to matter. We propose that the model, and the significant role of the values of the alternatives, may describe other decision-making systems, including intracellular regulatory circuits, and simple neural circuits, and may provide guidance in the design of decision-making algorithms for artificial systems, particularly those functioning without centralised control.

Citation: Pais D, Hogan PM, Schlegel T, Franks NR, Leonard NE, et al. (2013) A Mechanism for Value-Sensitive Decision-Making. PLoS ONE 8(9): e73216. doi:10.1371/journal.pone.0073216

Editor: Matjaz Perc, University of Maribor, Slovenia

Received July 11, 2013; **Accepted** July 16, 2013; **Published** September 2, 2013

Funding: JARM and NRF acknowledge Biotechnology and Biological Sciences Research Council grant BB/G02166X/2 (www.bbsrc.ac.uk). NEL acknowledges Office of Naval Research grant N00014-09-1-1074 (http://www.onr.navy.mil). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors note that co-author JARM is a PLOS ONE Editorial Board member. This does not alter the authors' adherence to all the PLOS ONE policies on sharing data and materials.

* E-mail: James.Marshall@sheffield.ac.uk

Introduction

Animals constantly make decisions, yet decision-making mechanisms and their evolution are still poorly understood in many cases. Recent years have seen a convergence of several research fields aiming to improve our understanding of general decision-making principles. Behavioural ecologists have argued for the need to combine the traditional study of animal behaviour through the lens of optimality arguments [1], with an increased understanding of the mechanisms underlying behaviour and their evolution [2]. At the same time psychologists and neuroscientists, who focus on understanding the mechanistic bases of behaviour, are increasingly focussing attention on how these mechanisms can implement optimal behaviour (*e.g.* [3-5]). Behavioural ecologists in the burgeoning subfield of collective animal behaviour are also interested in mechanisms, in terms of interaction rules and patterns, that generate sophisticated group decisions [6].

Some researchers have noted the parallels between these apparently disparate fields, by observing that the interaction patterns of neurons in brain circuits and animals in groups appear to be very similar [7-10], and also that tools and concepts from psychology and neuroscience may usefully be imported into the study of collective animal behaviour [11,12]. These ideas have been made concrete in modelling studies where, for example, optimality analyses from neuroscience [9] or decision-making tests from psychology [8] have been applied to models of collective decision-making by social insect colonies of ants and honeybees, and in experimental studies where the parallels have successfully guided the search for decision-making mechanisms in honeybees [13,14].

In this paper we present a comprehensive analysis of our previous empirically-motivated model of decision-making by house-hunting honeybees swarms [13], and argue that its decision-making properties may in turn guide the study of decision-making systems at other levels of biological complexity, up to individual brains, and down to intracellular decision-making circuits, as well as inform the design of artificial, decentralized decision-making systems. Our previous analysis showed that the particular pattern of `stop-signalling' observed in swarms allows them to adaptively avoid deadlock by choosing randomly when presented with two potential nest sites of equal quality, and to converge on choosing the best of two potential nest sites when there is a sufficiently large difference in their quality [13,14].

Here, we show further aspects of value-sensitive decision-making that arise from cross-inhibitory stop-signalling. We analyse a model whose decision-dynamics are characterised by fast attraction to a one-dimensional decision manifold, followed by slower time-evolution along this manifold. We leverage a time-scale separation to reveal how the strength of cross-inhibition critically determines the decision-system response to both the

difference in value and the mean value of the two alternatives. These analytic results considerably extend our previous initial analysis of this model's decision dynamics [13].

We show that stronger cross-inhibition yields a greater minimum difference in value required for discrimination between the alternatives. When the difference in value is below this minimum, the alternatives are treated as equal or nearly equal, and the cross-inhibition determines whether or not the alternatives are of sufficiently high value to warrant breaking decision deadlock. A stronger cross-inhibition increases the minimum mean value of the alternatives above which a decision deadlock is broken and the system randomly chooses one of the alternatives. When the (nearly) equal alternatives have mean value below the minimum mean value threshold, deadlock is maintained, allowing for the arrival of information on other, possibly more valuable, alternatives.

We show that cross-inhibition strength determines the minimum detectable difference in the value of alternatives, as a function of their mean value, in a manner similar to Weber's law as arising from psychological studies. We further show that for decisions over alternatives that do differ sufficiently in quality, that the stochastic decision dynamics exhibit a speed-accuracy trade-off in decision-making that depends critically on the difference in value and mean value of the alternatives, with dependence controlled by the strength of the cross-inhibition. The speed-accuracy trade-off is qualitatively similar to the statistically-optimal trade-off of the drift-diffusion model of decision-making, although we present evidence that decision-making does not achieve optimality under the parameterisations we consider here.

Model

The decision-making model we study is an extension of our previous empirically-motivated deterministic model [13] to include stochastic fluctuations in the relevant recruitment and interaction rates. Although we shall initially describe the model in terms of house-hunting honeybees, the formulation is general and could describe any decision-maker in which accumulators compete to reach a decision threshold, are activated, decay, and inhibit each other according to the values of the alternatives they represent. For the simplest case of a decision over two alternatives, the time-evolution of the general model is described entirely by a two-dimensional system of coupled stochastic ordinary differential equations as

$$\begin{cases} \mathrm{d}y_A &:= (y_U\gamma_A - y_A(\alpha_A - y_U\rho_A + y_B\sigma_B))\mathrm{d}t \\ &\quad + k\sqrt{y_U^2 + y_A^2 + y_U^2 y_A^2}\,\mathrm{d}W_A \\ \mathrm{d}y_B &:= (y_U\gamma_B - y_B(\alpha_B - y_U\rho_B + y_A\sigma_A))\mathrm{d}t \\ &\quad + k\sqrt{y_U^2 + y_B^2 + y_U^2 y_B^2}\,\mathrm{d}W_B \end{cases} \quad (1)$$

where y_A and y_B are the proportion of scout bees recruiting to potential nest sites A and B respectively, and $y_U = 1 - y_A - y_B$ is the proportion of uncommitted scouts in the colony. Since y_A and y_B represent accumulated commitment to the alternatives, in general we refer to them as *accumulators* as is typical in theoretical neuroscience, for example [15]. Greek letters are used to denote parameters of the colony's decision-making system, that could be tuned by evolution. Latin letters are used to denote parameters of the decision problem faced by the colony that are outside of its control. Here, γ_i is the rate at which scouts independently discover and begin recruiting to potential nest site i, α_i is the rate at which scouts spontaneously abandon their commitment to site i, ρ_i is the rate at which scouts committed to site i recruit uncommitted scouts via the`waggle dance' [16], and $\sigma_i$ is the rate at which scouts committed to site $i$ convert scouts recruiting for the competitor site to a state of non-commitment, using the`stop-signal' to disrupt waggle-dancing bees [13,17]. Our previous experimental work has shown that this signal is delivered in a targeted manner, in that bees committed to a particular site deliver stop-signals primarily to bees dancing for competitor sites [13].

A collective decision is reached when one of the scout populations reaches a (variable) quorum threshold ω. We assume that all of the rates depend on the value v_i of the relevant potential nest site. As in previous work we set $\gamma_i = \rho_i = v_i$ and $\alpha_i = 1/v_i$ [13]. Moving beyond the model of [13], we further assume that these crucial decision rates γ_i, α_i and ρ_i are subject to some stochastic variability, due to the inherently noisy evaluations of nest site quality v_i undertaken by individual scout bees; since our earlier work [13] showed that stop-signal strength should be independent of nest-site value, and since we are interested primarily in how sensory noise is processed by the decision-making system, no noise is added to rate $\sigma = \sigma_A = \sigma_B$. We assume independent white-noise (Wiener) processes added to the value-dependent rates, with identical variances k^2. As described in Text S1 (section S.2) independent Wiener processes can be combined into a single noise term with larger variance. This is captured in the $\mathrm{d}W_A$ and $\mathrm{d}W_B$ terms in Eq. 1 in which $\mathrm{d}W_i$ is a normally-distributed increment of the Wiener process i, with mean 0 and variance 1. Thus the parameter k controls the noisiness, or difficulty, of the decision problem, where higher k means noisier evaluations. This approach to capturing sensory noise in an infinite-population model is standard in theoretical neuroscience (*e.g.* [4]) and has previously been used to model collective behaviour of social insects (*e.g.* [9]). Note that noise captured in the Wiener processes of Eq. 1 is thus sensory noise, rather than intrinsic noise arising from finite populations of scout honeybees; correct derivation of intrinsic noise requires approaches based on the Master equation (*e.g.* [18]) and is beyond the scope of the present paper. For our dynamical systems analyses, we will set $k = 0$ in Eq. 1, recovering the noise-free dynamics of [13], while for our stochastic decision dynamics analyses, we will set $k > 0$.

Results

General Decision Dynamics — Separation of Timescales

Here we present analytic results on the general decision dynamics of the model. A well-established technique for studying models of binary decision-making similar to that described in Eq. 1 is to reduce the system of equations to a one-dimensional description of the decision dynamics (*e.g.* [4,9]). Denote the mean value of alternatives $\bar{v} := (v_A + v_B)/2$ and the difference in value of alternatives $\Delta v := v_A - v_B$. For large $\bar{v}$ and small $\Delta v/\bar{v}$, it can be shown that there is a separation of timescales; a singular perturbation analysis of the zero-noise ($k = 0$) dynamics (Text S1, Figure S1) reveals fast convergence, dominated by the dynamics of the uncommitted population y_U, to a stable one-dimensional decision manifold, followed by slow time-evolution, dominated by the relative dynamics of the accumulators y_A and y_B, along this manifold as illustrated in Figure 1. We note that the slow manifold, defined implicitly by (Text S1)

$$y_A y_B = \frac{2\bar{v}}{\sigma}\frac{y_U(1+y_A)(1+y_B)}{3-y_U}, \quad (2)$$

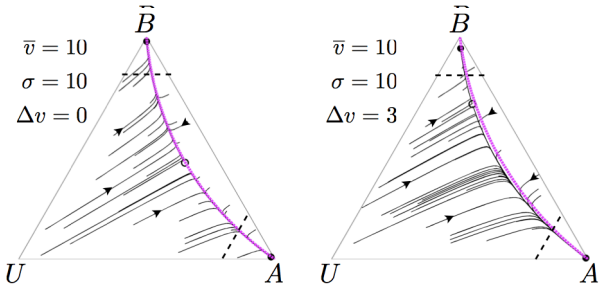

Figure 1. Decision-making dynamics on the unit simplex with vertex U corresponding to a fully uncommitted decision-maker ($y_U = 1$), vertex A to a decision-maker fully committed to alternative A ($y_A = 1$, and vertex B to a decision-maker fully committed to alternative B ($y_B = 1$). When the accumulator for alternative A or B (y_A or y_B) surpasses a decision threshold, illustrated with a dashed line, the corresponding alternative is selected by the decision-maker. Flow lines indicate sample noise-free trajectories over time, demonstrating fast convergence to a slow, invariant manifold. A singular perturbation analysis (Text S1) proves this separation of timescales, and gives the expression Eq. 2 for the slow manifold (magenta line), which is independent of Δv (thus, the slow manifold is the same in the right and left plots). The dynamics on the slow manifold depend on parameters of the decision problem $\bar{v}$ and Δv and of the cross-inhibition rate σ; stable attractors (filled circles) can co-exist with unstable saddle-nodes (hollow circles) on the slow manifold. Thus, decision-making can be reduced to a single decision-variable; this is the form of several classic models of decision-making, including those implementing provably optimal statistical tests.
doi:10.1371/journal.pone.0073216.g001

depends on $\bar{v}$ and σ but not on Δv, whereas the dynamics along the slow manifold depend explicitly on $\bar{v}$, σ, and Δv (Text S1). The analytically-calculated slow manifold is superimposed on the simulated decision-making dynamics in Figure 1 and in Figure S2, where it can be seen that the slow manifold approximates the slow dynamics well over a range of parameter values, deteriorating only when Δv is on the order of $\bar{v}$.

Thus, analysing the stochastic decision dynamics along the stable one-dimensional manifold will give a good understanding of the decision-making properties of the system as a whole. This is particularly relevant because the reduced dynamics resemble classical models of binary decision-making. For example, the general one-dimensional stochastic differential equation

$$\mathrm{d}x := (a+bx)\mathrm{d}t + c\mathrm{d}W, \tag{3}$$

where $\mathrm{d}W$ is the Wiener increment as in Eq. 1, includes Orstein-Uhlenbeck processes (OU — $a=0$, $b\neq 0$) and the drift-diffusion model (DDM — $a\neq 0$, $b=0$) as special cases. In these models as applied to decision-making, a may in certain cases correspond to the signal in the stimulus presented to the decision-maker, and c the noise in that stimulus. The decision-variable x models the tendency to choose one of two alternatives where a decision is made in favor of one alternative when x crosses a positive threshold, and the other alternative when x crosses a negative threshold. In the statistically-optimal DDM parameterisation, x represents the log likelihood ratio of the alternatives so that $x=0$ corresponds to equal evidence for each alternative.

Bogacz *et al.* previously recovered O-U processes and the DDM from two-dimensional connectionist models of choice in the visual cortex, while we recovered the DDM from two-dimensional models of nest-site selection by social insect colonies [9]. The DDM [19] is of particular interest to researchers studying decision-making because it corresponds to the statistically-optimal test for compromising between speed and accuracy of decision-making, and gives the best fits to reaction-time and error-rate distributions of subjects undertaking psychophysical decision tasks [4]. The analyses of [4] and [9] were facilitated by studying equations that converged to a linear stable manifold, whereas the stable manifold for Eq. 1 is clearly non-linear (Figure 1; Text S1). Nevertheless approximations to this manifold, as well as stochastic simulations, will enable us to analyse decision-making along it.

Minimum Value of Acceptable Equal Alternatives

Our previous analysis showed that the decision-making model of Eq. 1 with $k=0$, when alternatives are of equal value ($v := v_A = v_B$), exhibits a pitchfork bifurcation as a function of increasing cross-inhibition rate σ and value v [13]. In the pre-bifurcation case, a single attractor exists at which each accumulator is of equal size, whereas in the post-bifurcation case this attractor becomes an unstable saddle point, and attractors corresponding to each alternative emerge. That is, there is a critical level of cross-inhibition σ^* below which the decision-maker remains deadlocked between the two equal alternatives, but above which it converges to choosing one alternative at random. This threshold, plotted in Figure 2, was previously [13] calculated as

$$\sigma^* = \frac{4v^3}{(v^2-1)^2}. \tag{4}$$

Figure 2 demonstrates a further very useful decision-making property, that of value sensitivity. To illustrate the general principle, consider the particular case of a honeybee swarm that has discovered two equally poor potential nest sites. If both of these alternatives are of such low value to the swarm, through having insufficient volume to allow for effective colony growth and reproduction in the future, for example, then the swarm would be better off waiting to see if its scouts can discover other, higher value, alternatives in the vicinity. Figure 2 shows that, if the value of the alternatives v is sufficiently low given the swarm's rate of cross-inhibition σ then this is precisely what happens; the recruiter populations for the two alternatives A and B become deadlocked at equal commitment, while leaving a proportion of the swarm in the uncommitted state U and thus available to discover alternatives through independent exploration of the environment (Figure 2; bottom-left inset). Figure 3 presents stochastic simulations of a scenario illustrating this behaviour (see Text S1), in which two equal but poor quality alternatives are discovered, and stable deadlock persists between them until a third superior alternative is discovered and subsequently chosen. This late selection of an alternative differs qualitatively from earlier models [20], in which no mechanism for adaptive deadlock was considered; in [20] a lower recruitment rate for a poor alternative gives enough time for a late-discovered good alternative to overtake the poor and reach the decision threshold first. Although we have not presented them, our model with a single-discovered alternative, in which no cross-inhibition would occur, would exhibit similar dynamics. There is experimental evidence, however, that for honeybee swarms even with only two alternatives available for discovery, times-to-discovery relative to time-to-decision are sufficient to ensure that both alternatives are discovered and a competition between them occurs [13]. The results of our model also agree qualitatively with experimental evidence that honeybee swarms are able to choose a good-quality nest site over four other medium-quality nest sites [21], which

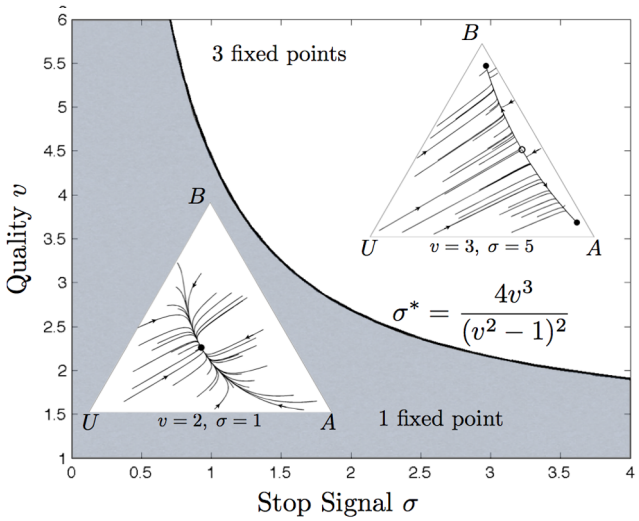

Figure 2. Value-dependent decision-making over equal alternatives. A critical cross-inhibition level σ^* can be calculated, below which stable decision-deadlock results due to a single stable attractor on the $y_A = y_B$ line. Increasing the strength of cross-inhibition above the critical threshold σ^*, this attractor becomes unstable and two stable attractors, one for each alternative, emerge from it and rapidly move apart [13]; in this situation one alternative will thus be chosen at random by the system. As the equation and plot for σ^* make clear, the level of cross-inhibition required to break deadlock decreases with increasing value v of the two alternatives. Thus, holding cross-inhibition level constant, decisions over equal but low value alternatives can result in deadlock, while decisions over equal but high value alternatives can result in a random choice. This can lead to sophisticated decision dynamics (Figs. 3 and S3).
doi:10.1371/journal.pone.0073216.g002

presumably requires an adaptive deadlock to be maintained between discovered medium-quality sites, until discovery of the good-quality site enables its selection.

If, however, for the same rate of cross-inhibition σ the value of the equal alternatives is sufficiently high, then the dynamics bifurcate so that the decision-maker converges on choosing one of the two alternatives at random (Figure 2; top-right inset). This illustrates a very sophisticated decision-making strategy; if information about only two alternatives is available but neither is very valuable then waiting to see if a better alternative is discovered could be sensible, whereas if the two alternatives are both of sufficient quality then quickly choosing one at random rather than wasting further time waiting for alternatives would be appropriate. Evolution could tune the level of cross-inhibition σ in a decision-maker to set the acceptance threshold for the value $\bar{v}$ of equal alternatives to an appropriate level, given the needs of an organism and the quality of alternatives typically available in an environment, as Figure 2 illustrates.

The preceding analysis assumes an evolutionarily hard-wired level of cross-inhibition, but further sophistication is possible if one considers what might happen to our hypothetical decision-maker, considering two equal but low value alternatives, if it waits too long. Any decision-maker has finite time and resources available to make decisions; in the case of a honeybee swarm members have finite energy reserves, since they load up with honey before swarming and do not resume foraging until the swarm has found a suitable nest site [10]. If after a long period of time the swarm still only has information about the two low-value alternatives then it is reasonable to assume that no better alternatives are available as they would likely have been discovered and, in any case, the resources of the swarm are being rapidly depleted. In this scenario it would be better for the swarm to choose one of the low value nest sites than none at all. This can be achieved by progressively increasing the cross-inhibition rate σ; as Figure 2 indicates, by

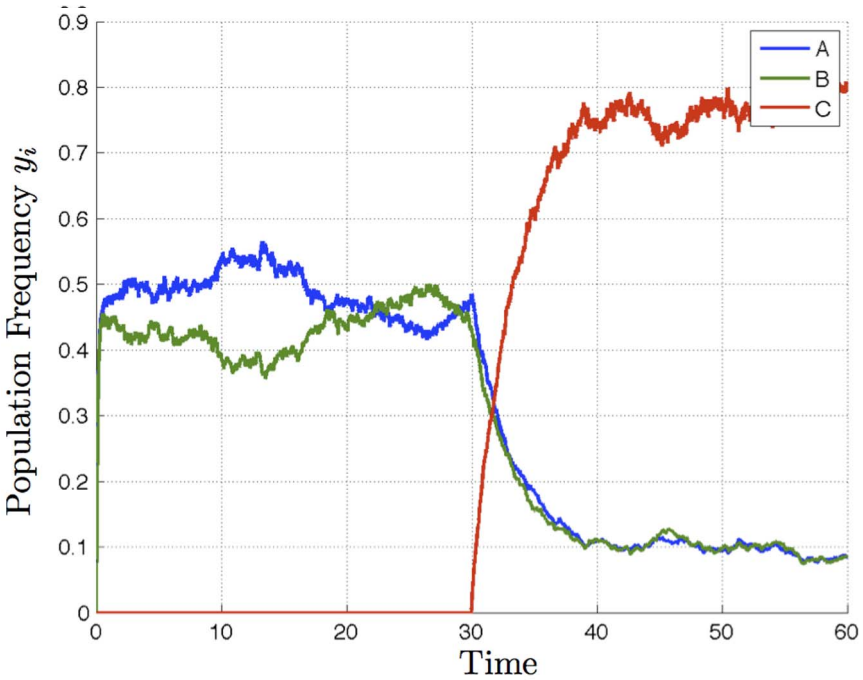

Figure 3. Stochastic simulation shows that for two sufficiently poor but equal alternatives, deadlock between the two persists until a third, superior alternative is discovered (at time $t = 30$), at which point it is selected by the decision-maker. The three-alternative model simulated here is a simple extension of the two-alternative model of Eq. 1, as described in section S.2 of Text S1. Noise parameter $k = 0.05$.
doi:10.1371/journal.pone.0073216.g003

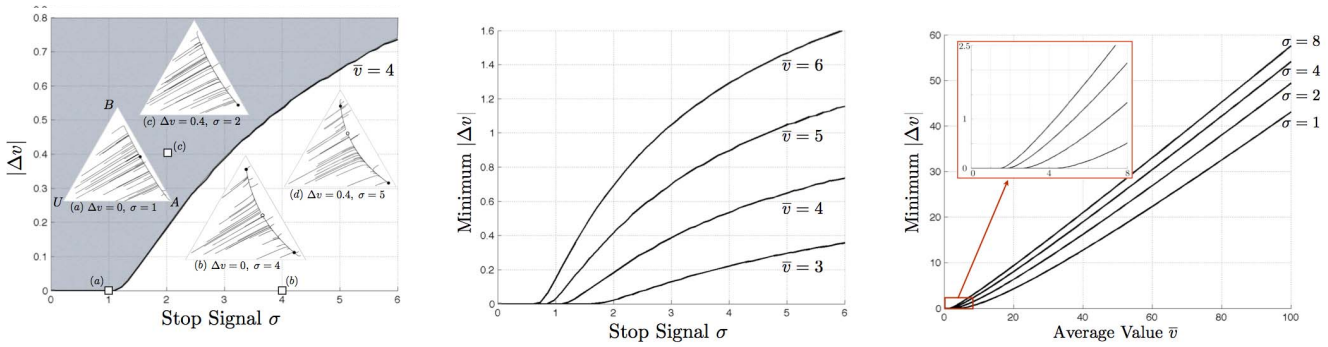

Figure 4. Dependence of decision-making over alternatives on absolute difference in value of alternatives $|\Delta v|$, cross-inhibition strength σ, and mean value of alternatives $\overline{v}$. (Left) Bifurcation set as a function of σ and $|\Delta v|$, for fixed $\overline{v}=4$. This generalises the result of Figure 2, for which $|\Delta v|=0$. The grey region corresponds to parameters where the decision dynamics have a single stable attractor (pre-bifurcation), whereas the white region corresponds to those having two stable attractors and one saddle node (post-bifurcation). Sample phase-portraits illustrate how the positions of these fixed points change according to σ and $|\Delta v|$. Plots (a) and (b) illustrate the results of Figure 2, in which $|\Delta v|=0$. Increasing $|\Delta v|$ moves the stable attractor towards the superior alternative in the pre-bifurcation case (see plot (c)), although it may still correspond to a population state in which threshold is reached for neither alternative; whereas increasing $|\Delta v|$ in the post-bifurcation case moves the saddle point towards the inferior alternative, thereby increasing the basin of attraction for the superior alternative (see plot (d)). Thus for a decision with given $|\Delta v|$ that is too low to precipitate a threshold decision, increasing σ precipitates a decision, in which the more valuable alternative is more likely to be selected. (Middle) The relationship between σ and the minimum $|\Delta v|$ required for a unique attractor for the best alternative depends on $\overline{v}$. (Right) The relationship between $\overline{v}$ and the minimum $|\Delta v|$ required for a single alternative to unambiguously be considered the best converges on a linear relationship, with slope determined by σ. This is similar to Weber's law of just noticeable difference, observed in psychological studies, with σ determining the Weber coefficient.
doi:10.1371/journal.pone.0073216.g004

doing so a point is reached at which the value of the alternatives v, which previously resulted in stable deadlock between them, is suddenly sufficient to precipitate a random choice between the two. A stochastic simulation illustrates this process in Figure S3.

Minimum Relevant Differences Between Equal Alternatives

The decision dynamics of the model are sensitive not only to the value of the available alternatives but also to the absolute *difference* $|\Delta v|$ in the values of the alternatives, as illustrated in Figure 4. First, the results of Figure 2 generalize to non-zero $|\Delta v|$; an increase in the rate of cross-inhibition σ leads to a bifurcation resulting in two stable attractors, one for each alternative.

As Figure 4 (Left) shows, for small $|\Delta v|$ the stable deadlock point (pre-bifurcation) is moved towards the better of the two alternatives (plot (c) in 4(Left)), but may still be placed such that neither alternative reaches threshold and thereby is selected. However, for small $|\Delta v|$, as in the case of equal value alternatives, increasing cross-inhibition ensures at least that a decision is reached; two stable attractors, one for each alternative, are introduced at the bifurcation with a saddle node between them.

For larger values of $|\Delta v|$, the saddle node (post-bifurcation) moves towards the inferior alternative, thereby increasing the chances that the better alternative is selected (point (d) in 4(Left)). For $|\Delta v|$ sufficiently large relative to the mean value $\overline{v}$ of the alternatives, the (pre-bifurcation) single stable attractor corresponding to the best alternative will be such that the decision-maker can reach the decision threshold required to select that alternative. Figure 4 (Middle) illustrates the minimum $|\Delta v|$ required to retain a (pre-bifurcation) single attractor for the best alternative as a function of σ for a given $\overline{v}$.

In Figure 4 (Right) the minimum $|\Delta v|$ required to retain a single attractor for the best alternative is plotted as a function of $\overline{v}$. The situation in which a single attractor exists is precisely the situation in which the decision-maker could be thought of as unambiguously identifying one superior alternative from the two available, since when two attractors exists, one for each alternative, some decision trajectories lead the system towards selecting the worst of the two alternatives. Notably, the minimum $|\Delta v|$ converges on a linear relationship with $\overline{v}$, with slope determined by σ (Figure 4 (Right)). This is analogous to Weber's law of just-noticeable difference, formulated in psychology, which states that the minimum difference in stimulus intensity required to discriminate between two sources varies linearly with their mean intensity as

$$\frac{\Delta v}{\overline{v}} = K, \tag{5}$$

where K is an empirically-determined constant. From Figure 4 (Right) it is evident that K in Eq. 5 is a function of cross-inhibition rate σ. Thus cross-inhibition controls the Weber coefficient with lower rates σ corresponding to lower Weber coefficients K, leading to a shallower increase of decision difficulty with mean value of alternatives in the decision.

Full Dynamics Classification

Figure 5 illustrates the full set of dynamical regimes that the stop-signal model of Eq. 1 can exhibit, as its parameters are changed. Figure 5 (Left) shows the pitchfork bifurcation with increasing cross-inhibition σ in the $|\Delta v|=0$ case. The dynamics in the $|\Delta v|\neq 0$ case exhibit a saddle-node bifurcation as a function of cross-inhibition rate σ (Figure 5 (Middle)). The dynamics also exhibit a hysteretic effect as a function of difference in value of the two alternatives Δv (Figure 5 (Right)). For a given value of $\overline{v}$, the bifurcations of the dynamics of Eq. 1, in two parameters Δv and σ, are qualitatively identical to the cusp catastrophe [22]. The plots in Figure 5 represent three slices through this cusp catastrophe bifurcation set. Each of these regimes is illustrated with stochastic simulations in the movies S1, S2, S3; the hysteresis loop implied by Figure 5 (Right) is illustrated in movie S4.

The saddle-node bifurcation of Figure 5 (Middle) clearly shows two features of the cross-inhibition rate σ. First, even for small differences in the value of alternatives relative to their mean value,

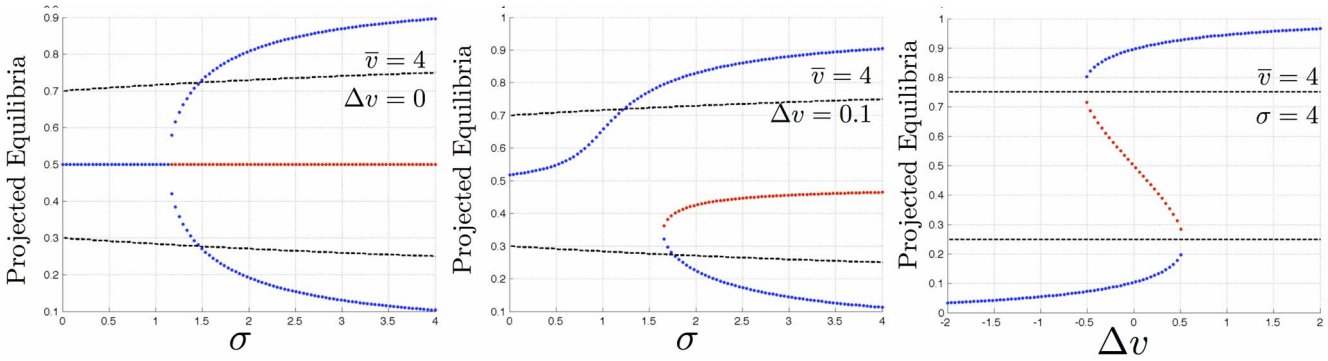

Figure 5. Full bifurcation behaviour of the stop-signal model of Eq. 1. According to parameterisation of the decision problem and decision-maker, the dynamics include **(i)** pitchfork bifurcation as a function of cross-inhibition rate σ in the equal alternatives case, **(ii)** saddle-node bifurcation as a function of cross-inhibition σ in the unequal alternatives case, and **(iii)** hysteresis as a function of difference in value of alternatives Δv. Fixed points are projected onto the [0,1] line as described in Text S1 and Figure S4. Blue dots indicate stable attractors, and red indicate unstable saddle points. Decision thresholds at $y_A = y_B = 0.7$ are indicated by dashed lines.
doi:10.1371/journal.pone.0073216.g005

increasing cross-inhibition σ improves decision-making by moving the (pre-bifurcation) single stable attractor further and further towards the state in which there is a more highly-activated accumulator for the superior alternative. If the decision threshold, defined by dashed lines, is set to an appropriate value, increasing the cross-inhibition would therefore amplify the differences in the qualities of the the alternatives sufficiently to precipitate a decision for the better alternative, on average.

Figure 5 (Middle) also shows that too high a rate of cross-inhibition σ can be detrimental. If the cross-inhibition rate is increased then a stable attractor for the *inferior* alternative suddenly appears in a saddle-node bifurcation, with an unstable saddle point between it and the original stable attractor. This can be helpful to ensure a decision if a threshold is not reached pre-bifurcation; however, in the case that a threshold is reached pre-bifurcation for the superior alternative, the bifurcation might not be helpful because post-bifurcation the superior alternative is no longer a unique solution. Further increase in the cross-inhibition rate σ moves the inferior attractor further toward or beyond the decision threshold for the inferior alternative, and moves the saddle point closer towards equal-magnitude accumulators for each alternative (0.5 on the y-axis of Figure 5 (Middle)). Thus increasing cross-inhibition too much changes the dynamics such that there may be an increasing risk of the decision-maker converging on choosing the inferior of the two alternatives. However, as we show below higher levels of cross-inhibition can have benefits for speed-accuracy trade-offs.

In Figure 5 (Right), there is a hysteretic effect as difference in the quality of alternatives Δv is smoothly increased and then decreased over time; this is illustrated in an animation of stochastic simulations in Text S1. While Δv is increasing, from an initially low level, over the interval of Δv in which three fixed points co-exist (approximately -0.5 to $+0.5$ in the figure) the system will be in the vicinity of the lower of the two stable attractors. At a sufficiently high value of Δv (approximately 0.5), the system will jump to the other, upper stable attractor. If Δv is then reduced over the same interval, the system will remain in the vicinity of the upper, stable attractor until Δv is less than approximately -0.5. While for a bee swarm, values of alternatives are unlikely to change smoothly over time in this way, this may be the case for other decision-makers, where exploitation of an alternative degrades its value, as in the example of intracellular decisions on activation of metabolic pathways considered in the Discussion. For neural decision-circuits, as also mentioned in the Discussion, laboratory experiments may be able to vary stimuli over time in this way. In both these cases the hysteretic effect of Figure 5 (Right) could act as a diagnostic that the decision-circuit used is similar in form to that described in Eq. 1.

Other authors have previously presented similar bifurcation results in different contexts for different models. For example [23] examines error rate and reaction times in connectionist models with non-linear interactions between accumulators, where these interactions serve to act as priming biases for decisions. Cell-fate decisions are analysed in [24] with respect to speed of intracellular signalling change, using the tools of bifurcation analysis. Foraging by social insect colonies, which differs from decision-making in that optimal foragers should exploit resources proportionally to their quality [25], has also been studied in this way [26], as has accuracy of collective decisions in such models [27]. While these previous studies do not, as we do, consider decisions in which a single decision-maker must choose only one option whose value they are rewarded by, they do highlight the importance of nonlinear interactions between accumulators in enabling the kinds of bifurcation behaviour presented here. In particular, nonlinear interaction between accumulators is not necessary for such behaviour; indirect nonlinear interaction, through accumulator populations competing for a finite pool of uncommitted individuals [26], for example, is sufficient.

Speed-Accuracy Trade-offs

As noted above, several classical models of decision-making, including the DDM and the (un)stable O-U process, are described using equations of stochastic motion on a line. The separation of timescales result presented above demonstrates that the decision dynamics converge rapidly to a line, along which they slowly diffuse. Of particular interest in decision-making models are speed-accuracy trade-offs [28-30], and the optimal compromise between these two quantities [4,9]. We therefore undertook preliminary numerical investigations (described in the Text S1) into the stochastic behaviour of the decision system under different parameterisations, once the system has converged to the stable decision-manifold, and until it crosses a decision-threshold.

Figure 6 presents a classic speed-accuracy trade-off, for a parameterisation that results in only a single attractor for the best alternative available. In Figure S5 in Text S1 we present numerical analyses of other cases, which highlight further interesting decision dynamics; in particular, we show for certain parameterisations that having an attractor for the incorrect

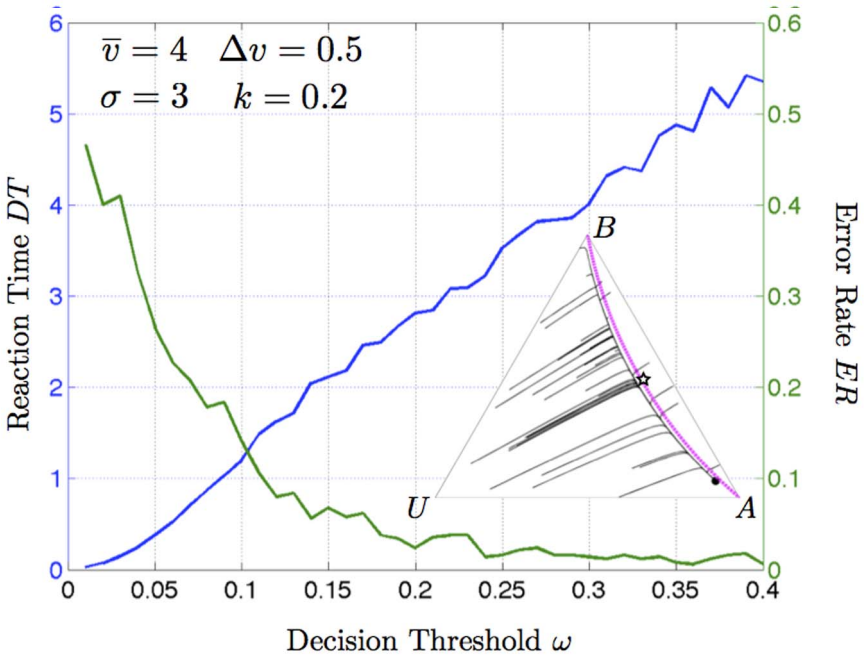

Figure 6. Speed-accuracy trade-off when differences between alternatives are sufficiently large that a single attractor for the best alternative exists; the observed speed-accuracy trade-off is qualitatively similar to that realised by the statistically-optimal drift-diffusion model of decision-making (see Text S1). When two attractors for alternatives of different values exist, however, the presence of the unstable saddle-node can improve error rate without compromising reaction time (see Text S1 and Figure S5).
doi:10.1371/journal.pone.0073216.g006

alternative can actually *improve* reaction time, without compromising decision accuracy (compare top left and top right plots of Figure S5). The fact that this improvement is possible indicates that decision-making along the stable manifold with a single attractor is *not* a statistically-optimal drift-diffusion process under the parameterisations studied here.

Discussion

Although motivated by and presented in terms of decision-making by house-hunting honeybee swarms, our model exhibits a number of beneficial decision-making qualities that we might expect other organisms to exhibit. At the heart of our analysis is the observation that, in a choice, an animal is typically rewarded by the value of the chosen alternative, rather than whether or not it chose the best. In particular the model decision-maker displays a sensitivity to the absolute as well as the relative value of the alternatives under consideration; this enables it to wait for information on better alternatives to arise when considering equally poor alternatives, but to spontaneously choose one equal alternative at random when both are good enough relative to a crucial decision-making parameter, the rate of stop-signalling, or cross-inhibition, σ. The decision-maker exhibits other properties observed in psychological studies, such as speed-accuracy trade-offs, and Weber's law of just-noticeable difference. The increasing rate of cross-inhibition may also improve the energetic costs of decision-making, although possibly at the expense of decision accuracy (as discussed in Text S1 and Figure S6). Our investigation has focussed on analytic treatment of the noise-free equations and stochastic simulations of speed-accuracy trade-offs and decision dynamics for binary decision-problems. Much work remains to be done in extending these analyses, for example to decisions over more than two alternatives.

Having suggested that our model might describe adaptive decision-making in general, what are the prospects for finding similar decision-making networks in other species? The form of the model equations is that of chemical reaction kinetics, in which interactions between chemical species are described by `mass action' terms. Therefore, there is the potential for intra-cellular regulatory networks to implement these decision-dynamics quite easily, for example in deciding for which of a number of available substances to activate the associated metabolic pathway. Evidence that single-cells can, for example, implement Bayesian-estimation through intra-cellular signalling [31], or exhibit Weber's law in gene regulatory pathways [32,33] indicates that such decision-making at the cell level is entirely plausible. Mutual inhibition also features in models of transcription in cell-fate decisions [24].

Another obvious class of decision models that invite comparison are those developed to describe neural networks for decision-making in simple perceptual decision tasks, such as those that take place in the primate visual cortex. A variety of accumulator models have been studied for their ability to fit experimental data, as well as implement optimal decision strategies (*e.g.* [4]). Optimal parameterisation of many such models requires evidence-accumulating pathways to interact [4], which the cross-inhibition mechanism in our model also implements. While optimality analyses in these models do take account of variable rewards for correct choices (*e.g.* [34,35]), they do not typically account for the fact that in real animals incorrect choices over, for example, food items still result in a reward, albeit one that is not the best available. Recently however, there is increasing interest in combining ideas from psychophysics, such as the Drift-Diffusion Model (DDM) [19], with the study of value integration processes (*e.g.* [36,37]).

Many accumulator models, like the classic DDM, also struggle with the correct choice when presented with zero net evidence, such as when choosing between two stimuli of equal average magnitude, and thus cross decision thresholds only through random drift. When choice of either alternative would result in

an equal reward, such behaviour is clearly sub-optimal. Proposals to deal with this include implementing `urgency signals' or collapsing decision thresholds over time [38,39], and the use of time-dependent sensory gain, and asymmetric inhibition between evidence pathways [38]. Our model differs from these proposed mechanisms, in that it spontaneously exhibits behaviour like that of an unstable Ornstein-Uhlenbeck process in order to break deadlock, according to the value of the alternatives under consideration and the strength of cross-inhibition.

Our non-linear model differs from the linear formulation of accumulator models underlying many analyses (*e.g.* [4]). The non-linear interaction terms of our model can, however capture neural activation dynamics; the logistic activation curve for neural populations in an accumulator model, used in [23], are qualitatively similar to`activation patterns' in the stop-signalling model, and [23] derives behaviour qualitatively similar, although not identical, to the stable-deadlock and deadlock-breaking results presented above. It is not unreasonable to expect convergent evolution to arrive at the same simple solution to the problem of value-dependent decision-making, in systems as diverse as single cells, honeybee swarms, and vertebrate nervous systems.

Supporting Information

Figure S1 Level curves in x,z coordinates.
(TIFF)

Figure S2 Comparison between the analytically computed slow manifold $h(x)$ plotted in magenta and simulations of the stop-signaling dynamics (S1). The match between the analytical slow manifold and the simulations is excellent, except for the case $\Delta v = \mathcal{O}(\bar{v})$, $\sigma = \mathcal{O}(\bar{v})$. For this set of plots, $\bar{v} = 10$, $\mathcal{O}(\bar{v}) \equiv 10$ and $\mathcal{O}(1) \equiv 1$.
(TIFF)

Figure S3 Simulations of the stochastic dynamics (S20) with time-varying stop-signal. A deadlocked population is able to converge to a decision for one of two equal alternatives by slowly ramping up the stop-signal; the critical value of stop-signal for the pitchfork bifurcation is marked on the bottom plot. Noise parameter $k = 0.05$.
(TIFF)

Figure S4 Illustration of equilibrium (y_{Aeq}, y_{Beq}) projected orthogonally onto $y_U = 0$. These projected equilibria are plotted in Figure 5 of the main text.
(TIFF)

Figure S5 Increasing stop-signalling rate σ has energetic benefits, as the total number of individuals involved in decision-making at any point in time is reduced. However, given the wisdom-of-the-crowds effect, this may have an adverse effect on collective accuracy, as fewer individual value estimates are pooled.
(TIFF)

Figure S6 Error Rate (ER, green) and Reaction Time (RT, blue) for the stochastic decision-making dynamics (S20) stopsde as a function of decision threshold ω. Parameters are indicated on each plot: (a) standard parameterisation from figure 6 in main text, (b) bistability with difference in value of alternatives, resulting from stronger cross-inhibition, (c) monostability for larger difference in value of alternatives and stronger cross-inhibition, (d) symmetric bistability when alternatives are equal in value.
(TIFF)

Matlab Code S1 Matlab code for stochastic simulation models.
(ZIP)

Movie S1 Locations of fixed-points, and simulated stochastic trajectories, as a function of varying stop-signal level (equal alternatives).
(MP4)

Movie S2 Locations of fixed-points, and simulated stochastic trajectories, as a function of varying stop-signal level (unequal alternatives).
(MP4)

Movie S3 Locations of fixed-points, and simulated stochastic trajectories, as a function of varying difference in quality of alternatives.
(MP4)

Movie S4 Hysteretic effect as result of smoothly varying difference in quality of intervals repeatedly over a fixed interval.
(MP4)

Text S1 Further information on analytic and simulation results.
(PDF)

Acknowledgments

We thank Jonathan Cohen, Jochen Ditterich, Phil Holmes, Jeremy Niven, Peter Swain and David Sumpter for discussion and comments.

Author Contributions

Conceived and designed the experiments: DP PMH NEL JARM. Performed the experiments: DP. Analyzed the data: DP PMH NEL JARM. Contributed reagents/materials/analysis tools: DP PMH JARM. Wrote the paper: DP PMH TS NRF NEL JARM.

References

1. Parker GA, Smith JM (1990) Optimality theory in evolutionary biology. Nature 348: 27–33.
2. McNamara JM, Houston AI (2009) Integrating function and mechanism. Trends in Ecology & Evolution 24: 670–675.
3. Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nature Neuroscience 2: 176–185.
4. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two - alternative forced-choice tasks. Psychological Review 113: 700–765.
5. Bogacz R (2007) Optimal decision-making theories: linking neurobiology with behaviour. Trends in Cognitive Sciences 11: 118–125.
6. Sumpter DJT (2010) Collective Animal Behavior. Princeton University Press.
7. Visscher PK (2007) Group decision making in nest-site selection among social insects. Annual Reviews Entomology 52: 255–275.
8. Passino KM, Seeley TD, Visscher PK (2008) Swarm cognition in honey bees. Behavioral Ecology and Sociobiology 62: 401–414.
9. Marshall JAR, Bogacz R, Dornhaus A, Planqué R, Kovacs T, et al. (2009) On optimal decision-making in brains and social insect colonies. Journal of the Royal Society: Interface 6: 1065–1074.
10. Seeley TD (2010) Honeybee Democracy. Princeton University Press.
11. Couzin ID (2009) Collective cognition in animal groups. Trends in Cognitive Sciences 13: 36–43.
12. Marshall JAR, Franks NR (2009) Colony-level cognition. Current Biology 19: R395–R396.
13. Seeley TD, Visscher PK, Schlegel T, Hogan PM, Franks NR, et al. (2012) Stop signals provide cross inhibition in collective decision-making by honeybee swarms. Science 335: 108–111.
14. Niven JE (2012) How honeybees break a decision-making deadlock. Science 335: 43–44.
15. Usher M, McClelland JL (2001) The time course of perceptual choice: The leaky, competing accumulator model. Psychological Review 108: 550–592.
16. Von Frisch K (1967) The Dance Language and Orientation of Bees. Harvard University Press.

17. Nieh JC (2010) A negative feedback signal that is triggered by peril curbs honey bee recruitment. Current Biology 20: 310–315.
18. van Kampen NG (2001) Stochastic Processes in Physics and Chemistry. North-Holland, third edition.
19. Ratcliff R (1978) A theory of memory retrieval. Psychological Review 85: 59.
20. Britton NF, Franks NR, Pratt SC, Seeley TD (2002) Deciding on a new home: how do honeybees agree? Proceedings of the Royal Society of London Series B: Biological Sciences 269: 1383–1388.
21. Seeley TD, Buhrman SC (2001) Nest-site selection in honey bees: how well do swarms implement the ?-of-? rule? Behavioral Ecology and Sociobiology 49: 416–427.
22. Arnol'd VI (1992) Catastrophe Theory. Springer-Verlag, 3rd edition.
23. Brown E, Holmes P (2001) Modeling a simple choice task: stochastic dynamics of mutually inhibitory neural groups. Stochastics and Dynamics 1: 159–191.
24. Nené NR, Garca-Ojalvo J, Zaikin A (2012) Speed-dependent cellular decision making in nonequilibrium genetic circuits. PloS one 7: e32779.
25. Goulson D (1999) Foraging strategies of insects for gathering nectar and pollen, and implications for plant ecology and evolution. Perspectives in Plant Ecology, Evolution and Systematics 2: 185–209.
26. Nicolis SC, Deneubourg JL (1999) Emerging patterns and food recruitment in ants: an analytical study. Journal of Theoretical Biology 198: 575–592.
27. Sumpter DJT, Zabzina N, Nicolis SC (2012) Six predictions about the decision making of animal and human groups. Managerial and Decision Economics 33: 295–309.
28. Marshall JAR, Dornhaus A, Franks NR, Kovacs T (2006) Noise, cost and speed-accuracy trade-offs: decision-making in a decentralized system. Journal of the Royal Society: Interface 3: 243–254.
29. Pratt SC, Sumpter JT (2006) A tunable algorithm for collective decision-making. Proceedings of the National Academy of Sciences 103: 15906–15910.
30. Passino KM, Seeley TD (2006) Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. Behavioral Ecology and Sociobiology 59: 427–442.
31. Perkins TJ, Swain PS (2009) Strategies for cellular decision-making. Molecular Systems Biology 5.
32. Goentoro L, Shoval O, Kirschner MW, Alon U (2009) The incoherent feedforward loop can provide fold-change detection in gene regulation. Molecular Cell 36: 894–899.
33. Ferrell JE (2009) Signaling motifs and weber's law. Molecular Cell 36: 724–727.
34. Feng S, Holmes P, Rorie A, Newsome WT (2009) Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? PLoS Computational Biology 5: e1000284.
35. Gao J, Tortell R, McClelland JL (2011) Dynamic integration of reward and stimulus information in perceptual decision-making. PloS one 6: e16749.
36. Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. Nature Neuroscience 13: 1292–1298.
37. Tsetsos K, Chater N, Usher M (2012) Salience driven value integration explains decision biases and preference reversal. Proceedings of the National Academy of Sciences 109: 9659–9664.
38. Ditterich J (2006) Stochastic models of decisions about motion direction: behavior and physiology. Neural Networks 19: 981–1012.
39. Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. The Journal of Neuroscience 31: 6339–6352.

Supplementary Information for 'A Mechanism for Value-Sensitive Decision-Making'

D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard and J. A. R. Marshall

S.1 Timescale Separation

This section comprises the details of the timescale separation calculation for the stop-signalling dynamics given by

$$\begin{aligned}\frac{dy_A}{dt} &= \frac{-2y_A}{2\overline{v}+\Delta v} + \left(\overline{v}+\frac{\Delta v}{2}\right) y_U(1+y_A) - \sigma y_A y_B \\ \frac{dy_B}{dt} &= \frac{-2y_B}{2\overline{v}-\Delta v} + \left(\overline{v}-\frac{\Delta v}{2}\right) y_U(1+y_B) - \sigma y_A y_B.\end{aligned} \tag{S1}$$

We perform a nonlinear coordinate transformation of the dynamics (S1) and apply singular perturbation theory to separate timescales and derive an analytic expression for the slow manifold (heteroclinic connections). Assuming large $\overline{v}$, define $\epsilon := \frac{1}{\overline{v}}$ as the small parameter for the timescale separation calculations. We follow the notation used in Chapter 11 of [3] and apply Tikhonov's Theorem (Theorem 11.1 in [3]) to prove the timescale separation.

Standard Singular Perturbation Model

Consider the coordinate transformation $(y_A, y_B) \mapsto (x, z)$ given by

$$\begin{aligned} z &= (1+y_A)(1+y_B) \\ x &= \frac{1+y_A}{1+y_B}. \end{aligned} \tag{S2}$$

Define the 2-simplex $\Delta_2 = \big\{(x,y) \in \mathbb{R}^2 | x \geq 0,\ y \geq 0 \text{ and } x+y=1\big\}$. The above transformation is well-defined on the domain $(y_A, y_B) \in \Delta_2$ since the Jacobian of the linearization has non-zero determinant $\frac{-2(1+y_A)}{1+y_B}$ on Δ_2. The inverse transformation is given by

$$\begin{aligned} y_A &= \sqrt{zx} - 1 \\ y_B &= \sqrt{\frac{z}{x}} - 1. \end{aligned} \tag{S3}$$

Note that $(y_A, y_B) \in \Delta_2$ implies that $z \in \left[1, \frac{9}{4}\right]$ and $x \in \left[\frac{1}{2}, 2\right]$. Level curves in x, z coordinates on the simplex are illustrated in Fig. S2.

Define the functions $\alpha(x,z)$ and $\beta(x,z)$ as

$$\alpha(x,z) = \sqrt{zx} + \sqrt{\frac{z}{x}} = 2 + y_A + y_B \tag{S4}$$

$$\beta(x,z) = \sqrt{zx} - \sqrt{\frac{z}{x}} = y_A - y_B. \tag{S5}$$

Then,

$$y_A y_B = z + 1 - \alpha,\ y_A + y_B = \alpha - 2, \text{ and } y_U = 3 - \alpha. \tag{S6}$$

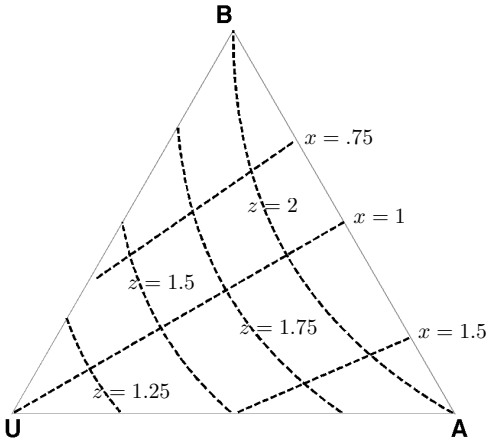

Figure S1: Level curves in x,z coordinates.

Computing the time derivative of z and substituting from (S1), (S4), (S5) and (S6) we get

$$\begin{aligned}\dot{z} &= \dot{y}_B(1+y_A) + \dot{y}_A(1+y_B) \\ &= \frac{2}{4\overline{v}^2 - \Delta v^2}\left[\Delta v\ \beta(x,z) - 2\overline{v}(2z - \alpha(x,z))\right] \\ &+ 2\overline{v}z(3-\alpha(x,z)) - \sigma\alpha(x,z)(z+1-\alpha(x,z)).\end{aligned} \tag{S7}$$

Computing the time derivative of x and substituting from (S1), (S4), (S5) and (S6) we get

$$\begin{aligned}\dot{x} &= \frac{\dot{y}_A(1+y_B) - \dot{y}_B(1+y_A)}{(1+y_B)^2} \\ &= \frac{x\Delta v(4z - 2\alpha(x,z)) - 4x\overline{v}\beta(x,z)}{4\overline{v}^2 z - z(\Delta v)^2} + x\Delta v(3-\alpha(x,z)) \\ &+ \frac{\sigma x\beta(x,z)}{z}(z+1-\alpha(x,z)).\end{aligned} \tag{S8}$$

Let $d = \sigma/\overline{v}$. The transformed dynamics (S7), (S8) can be written as a singular perturbation problem in standard form with small parameter $\epsilon = 1/\overline{v}$,

$$\begin{aligned}\epsilon\frac{dz}{dt} &= 2z(3-\alpha) - d\alpha(z+1-\alpha) \\ &+ \frac{2\epsilon^3\beta\Delta v}{4-\epsilon^2(\Delta v)^2} - \frac{4(2z-\alpha)\epsilon^2}{4-\epsilon^2(\Delta v)^2} \qquad =: g(x,z,\epsilon)\end{aligned} \tag{S9}$$

$$\begin{aligned}\frac{dx}{dt} &= x(3-\alpha)\Delta v + \frac{\sigma\beta x}{z}(z+1-\alpha) \\ &- \frac{4\beta x\epsilon}{4z - z\epsilon^2(\Delta v)^2} + \frac{x(4z-2\alpha)\epsilon^2\Delta v}{4z - z\epsilon^2(\Delta v)^2} \qquad =: f(x,z,\epsilon).\end{aligned} \tag{S10}$$

To ensure that the $\epsilon \to 0$ limit is well-defined we assume

$$\lim_{\overline{v}\to\infty}\frac{\Delta v}{\overline{v}} = \lim_{\epsilon\to 0}\epsilon\Delta v = 0. \tag{S11}$$

Slow Manifold Calculation

The slow manifold is given by the root of $g(x, z, 0) = 0$.

$$\begin{aligned} g(x,z,0) = 0 &\implies 2z(3-\alpha(x,z)) - d\alpha(x,z)(z+1-\alpha(x,z)) = 0 \\ &\implies 2zy_U - d\alpha(x,z)(z+1-\alpha) = 0 \\ &\implies 2(1+y_A)(1+y_B)y_U = d(3-y_U)(y_A y_B) \\ &\implies \frac{d}{2} y_A y_B = \frac{y_U(1+y_A)(1+y_B)}{3-y_U}. \end{aligned} \tag{S12}$$

Equation (S12) is an implicit expression for the slow manifold. Define the function $\hat{x} = \sqrt{x} + \frac{1}{\sqrt{x}}$. Then $\alpha(x, z) = \sqrt{z}\hat{x}$. In order to obtain an explicit expression we rewrite (S12) in the (x, z) coordinates as follows,

$$\begin{aligned} g(x,z,0) = 0 &\implies 2z(3-\alpha(x,z)) - d\alpha(x,z)(z+1-\alpha(x,z)) = 0 \\ &\implies 2z(3-\sqrt{z}\hat{x}) - d\sqrt{z}\hat{x}(z+1-\sqrt{z}\hat{x}) = 0 \\ &\implies 6\sqrt{z} - 2z\hat{x} - d\hat{x}z - d\hat{x} + d\hat{x}^2\sqrt{z} = 0 \\ &\implies (2+d)\hat{x}z - (d\hat{x}^2+6)\sqrt{z} + d\hat{x} = 0. \end{aligned} \tag{S13}$$

Equation (S13) is quadratic in $\sqrt{z}$. The solutions to the quadratic are given by

$$\sqrt{z} = \frac{d\hat{x}^2 + 6 \pm \sqrt{D}}{2(2+d)\hat{x}}, \tag{S14}$$

where the discriminant $D = d^2\hat{x}^4 + 36 + 4d\hat{x}^2 - 4d^2\hat{x}^2$. Hence we have two distinct solutions for the slow manifold given by

$$z = \left(\frac{d\hat{x}^2 + 6 + \sqrt{D}}{2(2+d)\hat{x}}\right)^2, \left(\frac{d\hat{x}^2 + 6 - \sqrt{D}}{2(2+d)\hat{x}}\right)^2. \tag{S15}$$

The second solution in (S15) lies outside the feasible domain $z \in [1, \frac{9}{4}]$ (and correspondingly $(y_A, y_B) \in \Delta_2$) and is hence rejected.

To summarize, the slow manifold is given by

$$\boxed{\begin{array}{l} z = \left(\dfrac{d\hat{x}^2 + 6 + \sqrt{D}}{2(2+d)\hat{x}}\right)^2 =: h(x) \\ \text{where } \hat{x} = \sqrt{x} + \dfrac{1}{\sqrt{x}} \text{ and } D = (d\hat{x}^2+6)^2 - 4d\hat{x}^2(2+d). \end{array}} \tag{S16}$$

Attractivity of the Slow Manifold

The boundary layer dynamics are given by

$$\frac{dz}{d\tau} = g(x, z + h(x), 0), \tag{S17}$$

where x is treated as fixed parameter. Stability of the boundary layer dynamics requires the exponential stability of its origin, uniformly in the fixed parameter x [3]. To test for exponential stability of the origin, we compute the Jacobian of the dynamics (S17) evaluated at the origin

$$\frac{\partial}{\partial z} g(x, z + h(x), 0)\Big|_{z=0} = -\frac{1}{12}\sqrt{1296 + \frac{(1+x)^2\sigma\left(24x + (x-1)^2\sigma\right)}{x^2}} \tag{S18}$$

and note that $\frac{\partial}{\partial z} g(x, z + h(x), 0)\big|_{z=0} < 0$ for all $x \in [\frac{1}{2}, 2]$.

The reduced dynamics on the slow manifold defined by (S16) are given by

$$\begin{aligned} \dot{x} &= f\left(x, h(x), 0\right) \\ &= \frac{\sigma x}{h(x)}\left[h(x) + 1 - \alpha\left(x, h(x)\right)\right]\beta\left(x, h(x)\right) + x(3 - \alpha\left(x, h(x)\right)\Delta v. \end{aligned} \tag{S19}$$

Comparisons to Simulations

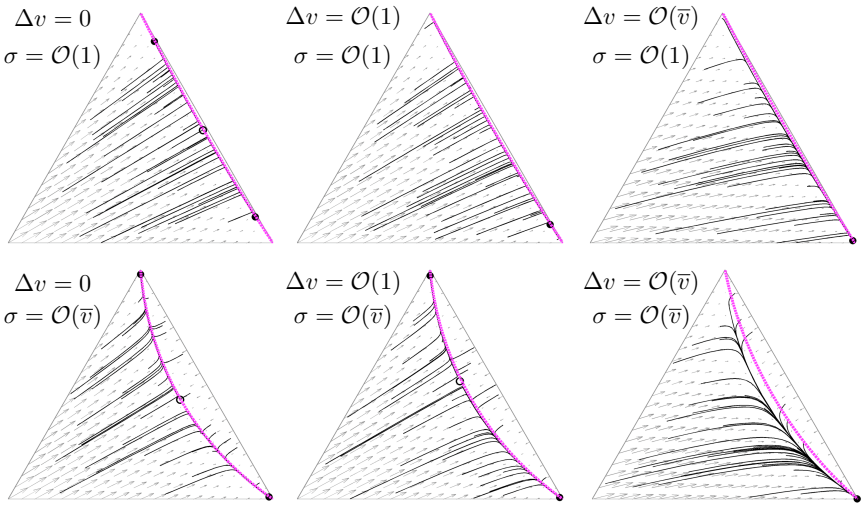

Figure S2: Comparison between the analytically computed slow manifold $h(x)$ plotted in magenta and simulations of the stop-signaling dynamics (S1). The match between the analytical slow manifold and the simulations is excellent, except for the case $\Delta v = \mathcal{O}(\overline{v})$, $\sigma = \mathcal{O}(\overline{v})$. For this set of plots, $\overline{v} = 10$, $\mathcal{O}(\overline{v}) \equiv 10$ and $\mathcal{O}(1) \equiv 1$.

In Fig. S2, we plot a comparison of the analytically derived slow manifold (S12) (or equivalently (S16)) and trajectories of the dynamics (S1) for various combinations of parameters σ and Δv, for large $\overline{v}$ (taken to be $\overline{v} = 10$ for this plot). The match between the approximation and the trajectories is excellent, except for the case when both Δv and σ are large. This case violates the limiting conditions of the singular perturbation calculation (S11).

S.2 Stochastic Simulations

The stochastic decision-making dynamics are given by

$$\begin{aligned} dy_A &= (\gamma_A y_U + \rho_A y_A y_U - \alpha_A y_A - \sigma_B y_A y_B)\,dt + k\sqrt{y_U^2 + y_A^2 + y_U^2 y_A^2}\,dW_A \\ dy_B &= (\gamma_B y_U + \rho_B y_B y_U - \alpha_B y_B - \sigma_A y_B y_A)\,dt + k\sqrt{y_U^2 + y_B^2 + y_U^2 y_B^2}\,dW_B, \end{aligned} \tag{S20}$$

for choices over two alternatives, where $y_U = 1 - y_A - y_B$, and

$$\begin{aligned} dy_A &= (\gamma_A y_U + \rho_A y_A y_U - \alpha_A y_A - \sigma_B y_A y_B - \sigma_C y_A y_C)\,dt + k\sqrt{y_U^2 + y_A^2 + y_U^2 y_A^2}\,dW_A \\ dy_B &= (\gamma_B y_U + \rho_B y_B y_U - \alpha_B y_B - \sigma_A y_A y_B - \sigma_C y_B y_C)\,dt + k\sqrt{y_U^2 + y_B^2 + y_U^2 y_B^2}\,dW_B \\ dy_C &= (\gamma_C y_U + \rho_C y_C y_U - \alpha_C y_C - \sigma_A y_A y_C - \sigma_B y_B y_C)\,dt + k\sqrt{y_U^2 + y_C^2 + y_U^2 y_C^2}\,dW_C \end{aligned} \tag{S21}$$

for choices over three alternatives, where $y_U = 1 - y_A - y_B - y_C$. The parameter k in (S20) and (S21) sets the level of stochasticity or noisiness in the decision-making process; higher values of k correspond to noisier evaluations.

To derive the noise terms of the above equations, white noise is added to each value-dependent transition rate in the original ordinary differential equations (ODEs), as motivated in the main paper, to capture sensory noise in individual valuations of alternative nest sites. As described in [1], for example, multiple independent white noise (Wiener) processes can be combined into a single Wiener process whose standard deviation is

the square root of the sum of the constituent processes' variances. Thus, for example, the rate terms for spontaneous commitment to and abandonment of a nest site, from the original noise-free ODE model, with white noise terms having standard deviation k added to these rates, can be written

$$y_U(\gamma_A dt + k dW_{\gamma_A}) - y_A(\alpha_A dt + k dW_{\alpha_A}), \tag{S22}$$

where dW_{γ_A} and dW_{α_A} are independent Wiener processes with unit variance/standard deviation, capturing normally-distributed error around, respectively, the spontaneous commitment and abandonment rates for site A, γ_A and α_A. Note that for simplicity we assume the standard deviations in the normally-distributed fluctuations around the rates are all k; relaxing this would still allow a single Wiener process to be constructed, but without a single parameter controlling the overall noisiness of the system. Since our motivation in studying the stochastic model is primarily to show that the deterministic dynamics we derived are not sensitive to noise, the former modelling approach is satisfactory.

As noted above, the sum (or difference) of two independent Wiener processes with standard deviations a and b is itself a Wiener process with standard deviation $\sqrt{a^2 + b^2}$. Applying this fact to the above expression we derive a stochastic ODE with a single Wiener process

$$(y_U\gamma_A - y_A\alpha_A)dt + k\sqrt{y_U^2 + y_A^2}dW, \tag{S23}$$

where dW is a third Wiener process having unit variance, since $\sqrt{k^2y_U^2 + k^2y_A^2} = \sqrt{k^2}\sqrt{y_U^2 + y_A^2}$. The full noise terms in the above equations of the stochastic model are thus constructed in this way by adding Wiener fluctuations to all value-dependent rates in the original deterministic ODE model and summing them into a single Wiener process; note that value-dependent rates include all the rates *except* the stop-signalling rate, which previous work has shown should not be value-dependent [4].

Two illustrative simulations of the stochastic dynamics (S20) and (S21) are shown in Fig. S3 and Fig. 3 in the main text:

- Breaking symmetric deadlock (Fig. S3): The decision-making process is between two equal alternatives with value $v_A = v_B = 3$, hence the dynamics are described by (S20). The stop-signal σ ramps up linearly from an initial value of $\sigma = 0$ to a final value of $\sigma = 3$. The increasing stop-signal enables the deadlock to be broken when $\sigma > \sigma^* \approx 1.7$ with one alternative randomly winning out. This simulation is motivated in part by the fact that elapsed time might influence the computations that underlie decision processes since prolonged deliberation can be expensive [2].

- Deadlock enables better final outcome (Fig. 3, main text): For a system tuned to intermediate levels of stop-signal ($\sigma = 1$ here), alternatives of low quality result in deadlock. This enables the system to wait for other potentially better alternatives to arrive into the mix and precipitate a decision. In this simulation, a third high-quality option ($v_C = 4$) enters at $t = 30$ and dominates. The dynamics are described by (S21)

S.3 Animations Corresponding to Fig. 5

- *FPlocs1.mp4, FPlocs2.mp4, FPlocs3.mp4*: Bifurcation plot and stochastic simulations corresponding to Fig. 4 of the main text. Each snapshot of the simplex shows 50 trajectories initialized at the origin $y_A = y_B = 0$ (correspondingly $y_U = 1$), with noise parameter $k = 0.1$, parameters $\overline{v}$, Δv and σ indicated below each simplex, and for time $t \in [0, 30]$.

  Circles are equilibria (filled stable, hollow unstable) of the deterministic dynamics (S1), and the dashed lines are quorum thresholds $y_A = 0.7$ and $y_B = 0.7$. The inset bar plots show the fraction of trajectories reaching each threshold; 'none' indicates trajectories reaching neither threshold in the given time (plotted in cyan). For $\Delta v = 0$, trajectories reaching either threshold are plotted in blue. For $\Delta v \neq 0$ trajectories reaching the correct threshold (one with higher v_i) are plotted in green, those reaching the incorrect threshold are in red.

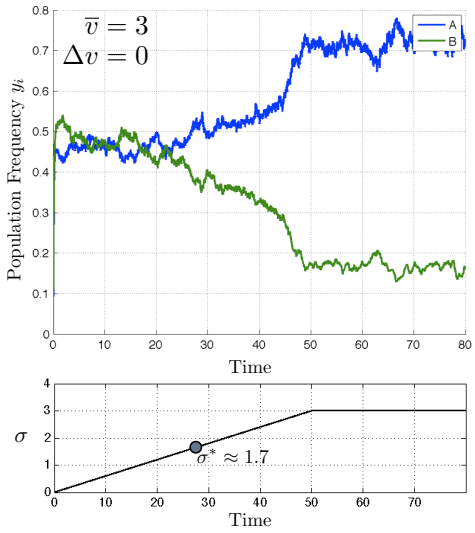

Figure S3: Simulations of the stochastic dynamics (S20) with time-varying stop-signal. A deadlocked population is able to converge to a decision for one of two equal alternatives by slowly ramping up the stop-signal; the critical value of stop-signal for the pitchfork bifurcation is marked on the bottom plot. Noise parameter $k = 0.05$.

- *hysteresis.mp4*: Cloud of points moving according to the stochastic dynamics (S20) with noise parameter $k = 0.1$ and other parameters shown on the plot. This simulation illustrates the hysteretic effect associated with continuously varying the difference between alternatives Δv for a given mean value $\overline{v}$ and stop signal σ.

S.4 Projected Equilibria

The bifurcations plots in Fig. 5 of the main text are obtained by projecting equilibria of the dynamics (S1) orthogonally onto the $y_U = 0$ line, as illustrated for an equilibrium (y_{Aeq}, y_{Beq}) in Fig. S4. For any given equilibrium (y_{Aeq}, y_{Beq}), the orthogonal projection onto the $y_U = 0$ line is given by the point $(\frac{1}{2}+\frac{\Delta y_{eq}}{2}, \frac{1}{2}-\frac{\Delta y_{eq}}{2})$ where $\Delta y_{eq} = y_{Aeq} - y_{Beq}$. The location of the point along $y_U = 0$ is given by $P = \frac{1}{2} + \frac{\Delta y_{eq}}{2}$, which is the quantity that is plotted on the vertical axis of Fig. 5.

S.5 Speed-Accuracy Trade-offs

Error Rate (ER) and Reaction Time (RT) plots for the stochastic stop-signalling dynamics (S20) are shown in Fig. S6 for four sets of parameters. For each of these sets, ER and DT are plotted as a function of the decision threshold ω. We use 50 values of ω in the range $\omega \in [0, 0.4]$ shown on the x-axis to construct each plot. In each plot, for a given value of ω, 500 simulations of the stochastic dynamics (S20) are run with initial condition at the zero evidence point ($y_A = y_B =: y_0$) on the corresponding slow manifold. The slow manifold and initial condition (shown as a star at (y_0, y_0)) are plotted in the insets. Each simulation is stopped when a trajectory reaches the threshold (either $y_A > y_0 + \omega$ or $y_B > y_0 + \omega$). RT is the mean time taken by the 500 trajectories to reach either threshold. ER is the fraction of trajectories reaching the incorrect threshold (one with lower value v_i). Each of the plots in Fig. S6 requires $500 \times 50 = 25000$ runs of the stochastic dynamics (S20).

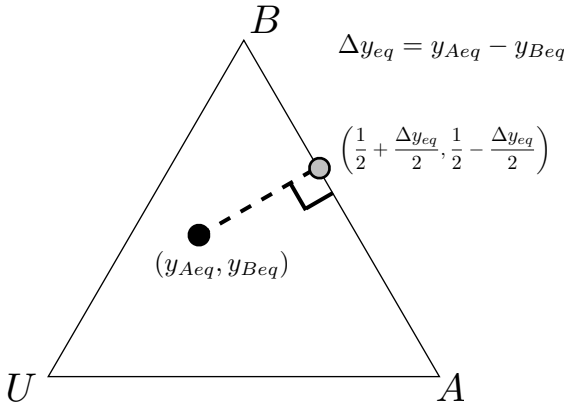

Figure S4: Illustration of equilibrium (y_{Aeq}, y_{Beq}) projected orthogonally onto $y_U = 0$. These projected equilibria are plotted in Fig. 5 of the main text.

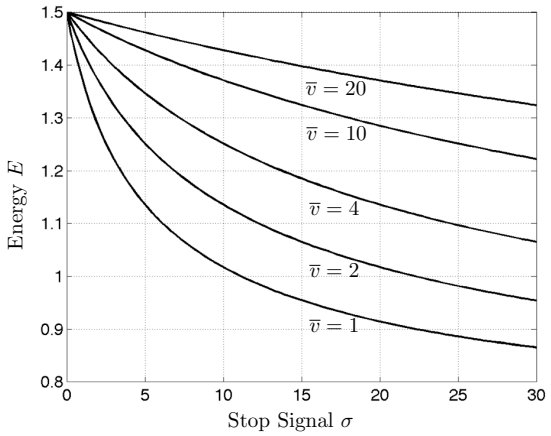

Figure S5: Increasing stop-signalling rate σ has energetic benefits, as the total number of individuals involved in decision-making at any point in time is reduced. However, given the wisdom-of-the-crowds effect, this may have an adverse effect on collective accuracy, as fewer individual value estimates are pooled.

S.6 Energetic Expenditure

Stop-signalling may also benefit the energy consumption of the swarm during decision-making. Energetic considerations are often ignored in studies of decision-making, yet are of crucial importance for decision-makers embodied in the real world. As mentioned in the main text, for example, house-hunting honeybee swarms have fixed energetic resources since members do not forage until decision-making is completed. A similar argument can be applied to decision-making in neural circuits, where activation of a neural population will consume energy over time. In Fig. S5 we show how increasing the stop-signalling rate σ improves energy consumption by the swarm by reducing the number of individuals involved in the decision-making. However, this is at the possible expense of decision accuracy, as discussed below.

Expected energetic expenditure of the swarm can be characterised as

$$\mathbb{E}(\text{energy expenditure}) \propto \int_{1/2}^{2} (y_A(x) + y_B(x)) \,\text{Prob}(x) dx \tag{S24}$$

where, $x \in [1/2, 2]$ parameterizes the the one-dimensional slow manifold (S16). $y_i(x)$ is the inverse transform

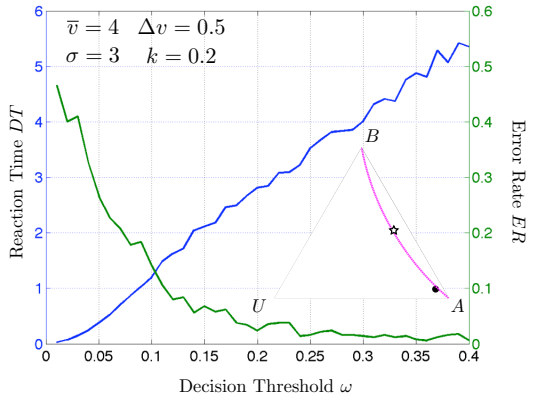

(a) Standard case from main text Fig. 6

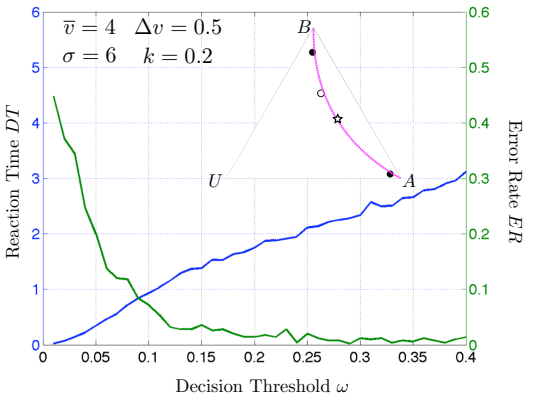

(b) Bistability, larger σ

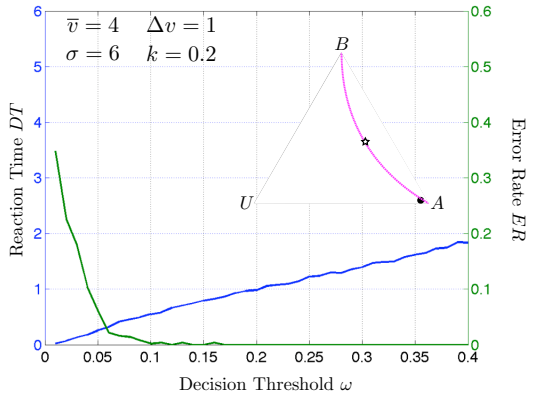

(c) Single attractor, larger Δv and larger σ

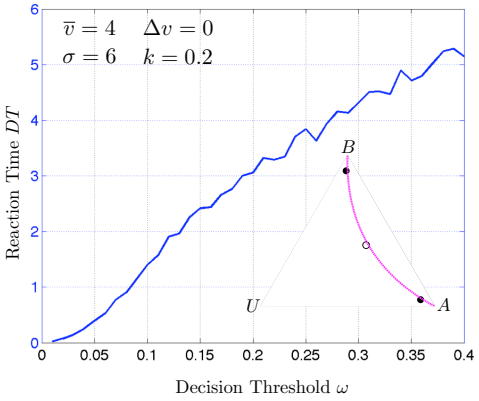

(d) Symmetric $\Delta v = 0$

Figure S6: Error Rate (ER, green) and Reaction Time (RT, blue) for the stochastic decision-making dynamics (S20) as a function of decision threshold ω. Parameters are indicated on each plot.

(S3) that recovers the number of individuals committed to site i when the decision-system is at point x on the decision manifold, and Prob(x) is the probability that the decision-system is at point x on the manifold at any given point in time. For simplicity let us assume that the swarm is equally likely to occupy all of the points x on the manifold. This is a conservative assumption, since during decision-making, and before decision thresholds are reached, the system will spend more of its time at points closer to the centre of the manifold ($x = 1$, corresponding to equal commitment of scouts to the two alternatives, *i.e.* $y_A = y_B$), and it is precisely these points whose energetic cost is reduced the most by increased stop-signalling. Thus, by assuming a uniform distribution over the stable manifold we underweight the points on the manifold that have the greatest energetic saving for increased stop-signalling. Nevertheless, plotting the energetic expenditure of (S24) as a function of σ shows that increasing the rate of stop-signalling reduces the expected energetic costs of decision-making (Fig. S5).

It should be noted, however, that although reducing energetic expenditure during decision-making is beneficial, there is a disadvantage not accounted for by our continuous model. In a decision-making system with discrete decision-components, whether they be honeybee scouts, or members of neuronal populations, reducing the number of components participating in a process will increase the noise associated with the process. By reducing the number of individuals responding to stochastic samples of the value of an alternative, the group's average evaluation will become much less reliable, a phenomenon well understood by any practitioner

of statistics attempting to work from limited sample sizes, and also familiar to those aware of the 'wisdom of crowds' effect in collective behaviour. Since our model is continuous it is beyond the scope of the present paper to investigate this trade-off in any detail.

References

[1] R. Bogacz, E. Brown, J. Moehlis, P. Holmes and J. D. Cohen (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks, *Psychological Review* 113 pp. 700–765.

[2] T. D. Hanks, M. E. Mazurek, R. Kiani, E. Hopp and M. N. Shadlen (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task, *The Journal of Neuroscience* 17 pp. 6339–6352.

[3] H. K. Khalil (1992) *Nonlinear Systems*. Macmillan, New York.

[4] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks and J. A. R. Marshall (2012) Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* 335 pp.108–111.