

Bee reverse-learning behavior and intra-colony differences: simulations reveal benefits of diversity.

A. G. Dyer^{a,b,c}, A. Dorin^d, V. Reinhardt^c, J. E. Garcia^{a,e}, M.G.P. Rosa^b

^a, Media and Communication, RMIT University, Australia 3001

^b, Department of Physiology, Monash University, Australia 3800

^c Johannes Gutenberg-Universität, Mainz, Germany 55099

^d, Faculty of Information Technology, Monash University, Australia 3800

^e, School of Applied Sciences, RMIT University, Australia 3001

Author for correspondence:

Assoc. Prof. Alan Dorin

Monash University,

Building 63, Faculty of IT,

Wellington Rd.

Clayton,

Australia 3800

Ph + 61 3 9905 3576

alan.dorin @ monash.edu

Running heading: Individual differences in color reverse-learning

Abstract

Foraging bees use color cues to help identify rewarding from unrewarding flowers. As environmental conditions change, bees may require behavioral flexibility to reverse their learnt preferences. Learning to discriminate perceptually similar colors takes bees a long time, and thus potentially poses a difficult task to reverse-learn. We trained free-flying honeybees to learn a fine color discrimination task that could only be resolved (with about 70% accuracy) following extended differential conditioning. The bees were then tested for their ability to reverse-learn this visual problem. Subsequent analyses potentially identified individual behavioral differences that could be broadly classified as: ‘Deliberative-decisive’ bees that could, after several flower visits, decisively make a large change to learnt preferences; ‘Fickle-circumspect’ bees that changed their preferences by a small amount every time they received a reward, or failed to receive one, on a particular color; and ‘Stay’ bees that did not change from their initially learnt preference. To understand the ecological implications of the observed behavioral diversity, agent-based computer simulations were conducted by systematically varying parameters describing flower reward switch oscillation frequency, flower handling time, and fraction of defective ‘target’ stimuli that contained no reward. These simulations revealed that when the frequency of reward reversals is high, Fickle-circumspect bees are more efficient at nectar collection, but as reward reversal frequency decreases, the performance of Deliberative-decisive bees becomes most efficient. As the reversal frequency continues to fall, Fickle-circumspect and Deliberative-decisive strategies approach one another in efficiency. In no tested condition did Stay bees outperform the other groups. These findings indicate there is a fitness benefit for honeybee colonies containing individuals exhibiting different strategies for managing changing resource conditions.

Keywords: individual behavior, learning, flower, pollinator, visual ecology, honeybee, individual-based simulation

Highlights:

- Explores individual behavioral differences in plant-pollinator interactions
- Reveals behavioral differences for real bees’ complex colour decision-making

- Bee / flower model reveals conditions where specific decision strategies are superior
- Model reveals conditions that could have led to evolution of intra-colony variety

1 Introduction

To understand decision-making in bees for difficult visual tasks, it is useful to combine approaches of behavioral testing with computer modeling. This multidisciplinary approach allows for the interpretation of the ecologically relevant factors that may influence how and why individuals make certain decisions, and how this potentially benefits the colony (Burns, 2005; Burns and Dyer, 2008). In particular, situations in which behavior varies between individuals, or local environmental conditions influence individual decision-making, agent-based models (ABMs; also called individual-based models) offer a powerful approach for understanding the intricate interactions and emergent outcomes of complex systems in the context of behavioral ecology (DeAngelis and Mooij, 2005; Dorin et al., 2008; Grimm, 1999; Grimm and Railsback, 2005; Grimm et al., 2005; Huston et al., 1988; Judson, 1994). ABMs have been used to understand the ecology of bee behavior since the 1980s (Hogeweg and Hesper, 1983). For example, ABMs have been used to understand bee foraging strategies with respect to recruitment, homing and memory of food source location, with the assumption of a homogeneous population (de Vries and Biesmeijer, 1998).

Agent-based models have also demonstrated that the benefits of recruitment by honeybees are dependent on the density and distribution of flowers within an environment. For instance, individual honeybees use a symbolic dance language to communicate the likely vector location of profitable food resources to nest mates (Leadbeater and Chittka, 2007; Seeley, 1985). There are individual differences in how nest mates within a hive respond to a signaled dance language. These depend on factors like a bee's individual experience and the strength of the signal that can indicate resource quality (Biesmeijer and Seeley, 2005; Leadbeater and Chittka, 2007). Understanding the potential hive benefits of the symbolic dance language communication has been possible with ABMs, revealing that effects like flower resource distribution and density have significant influences on whether it is

beneficial for a hive to have bees that follow communication signals, or individuals that rely on individual foraging capacity (Dornhaus et al., 2006).

Foraging for nutrition in the form of nectar in natural environments presents a variety of potential dilemmas for free-flying bees. For example, whereas there may be a number of flowers possessing similar identifying cues that offer nectar as a nutritional reward (Dyer and Chittka, 2004a), there could also be mimics like orchids that offer no reward (Dafni, 1984). It is also possible that the amount of reward offered by a certain species of flower varies over time (Chittka et al., 1997; Townsend-Mehler and Dyer, 2012; Waddington et al., 1981; Waddington and Heinrich, 1981), and that plants that usually have rewarding flowers will present empty flowers, simply because they have been recently visited by other foragers (Chittka and Schürkens, 2001; Giurfa et al., 1994; Heinrich, 1979; Townsend-Mehler and Dyer, 2012). An additional complexity for foraging bees is introduced by the flowers of the legume *Desmodium setigerum*, as these flowers even have a capacity to change color within 2 hours to potentially manipulate the behavior of pollinators (Willmer et al., 2009).

Many social bees, like honeybees and bumblebees, tend to exhibit flower constancy and typically remain constant to one type of rewarding flower, as long as it continues to present rewards (Chittka et al., 1999; Raine and Chittka, 2007b), although in complex natural environments flower constancy may break down (Raine and Chittka, 2007b). This type of pollinator behavior can be evaluated in reverse-learning experiments (Mota and Giurfa, 2010; Pavlov, 1927), and has been investigated in honeybees using discrimination tasks based on olfactory cues (Komischke et al., 2002; Mota and Giurfa, 2010), tactile cues (Scheiner et al., 1999; Scheiner et al., 2001) and saliently different color discrimination (Menzel, 1969; von Helversen, 1974) tasks; and for bumblebees on both color (Raine and Chittka, 2012) and sensorimotor learning tasks (Chittka, 1998). When considering different colors like 'orange' and 'blue', free-flying honeybees can quickly learn within five trials to choose a rewarding color with accuracy greater than 80%, and then quickly switch these learnt preferences after a further 1-2 trials if the reward contingency is reversed. However, with this short training, honeybees can only reverse decisions up to three times before discrimination falls to chance levels (Menzel, 1969). This finding was

confirmed in a separate study that trained honeybees to discriminate between saliently different ‘blue’ and ‘yellow’ color stimuli that were learnt in three trials to an accuracy greater than 80%, and the bees could then quickly switch preferences after 1-2 further trials when the reward paradigm was switched (von Helversen, 1974). In this case of a short learning opportunity, honeybees also chose between the color stimuli at random levels if the reverse training continued for more than three reversals; however, if the training was extended to 10 rewards on a particular color stimulus then reverse-learning was very robust for at least nine reversals (von Helversen, 1974). This indicates that length of training is important to the capacity of honeybees to robustly reverse-learn a salient color task.

Recent work examining how harnessed honeybees reverse-learn olfactory stimuli has revealed that different individual bees may possess different strategies for reverse learning (Mota and Giurfa, 2010). If harnessed honeybees are presented with two different odorants to discriminate between using a standard proboscis extension reaction (PER) experimental setup (Bitterman et al., 1983), some honeybees can reverse-learn the discrimination up to three times (Mota and Giurfa, 2010), which is consistent with the work on salient color discrimination with free-flying honeybees (Menzel, 1969; von Helversen, 1974). However, an important difference in the recent work on olfactory reverse learning was the observation of individual differences between how honeybees were able to perform the reverse switching task (Mota and Giurfa, 2010). It has been proposed that there are three categories of honeybees: ‘efficient’ reversers that could quickly change preferences when experimental conditions changed; bees that did learn the initial discrimination task but then appeared unable to reverse learn the task; and a third category that failed to learn the initial olfactory discrimination task (Mota and Giurfa, 2010). The existence of the last category of bees implies that there was a reasonable degree of perceptual difficulty involved in the learning of this olfactory discrimination task. This finding of individual differences for perceptually difficult olfactory learning in harnessed honeybees agrees with other recent work reporting differences in performance levels when individual free-flying honeybees solve perceptually difficult color discrimination tasks (Burns and Dyer, 2008; Muller and Chittka, 2008).

When considering color stimuli, recent work on honeybees (Avarguès-Weber et al., 2010; Giurfa, 2004; Reser et al., 2012) and bumblebees (Dyer and Chittka, 2004c) has revealed that the difficulty of a task can be controlled by varying the perceptual similarity of color stimuli. Specifically, the probability with which color differences can be judged by bees follows a sigmoidal-type function (Dyer, 2012b; Dyer and Neumeyer, 2005; Dyer et al., 2008b). Color differences can be conveniently specified in a color space like a Hexagon color model, which allows for the Euclidean distance between stimuli to be quantified (Chittka, 1992). Recent research has shown that when either honeybees (Avarguès-Weber et al., 2010; Avarguès-Weber et al., 2011; Dyer, 2012b; Giurfa, 2004) or bumblebees (Dyer and Chittka, 2004c; Dyer et al., 2011) learn color information in isolation (termed absolute conditioning), they only demonstrate a coarse level of color discrimination [between colors separated by about 1.5 hexagon units (Dyer and Chittka, 2004c; Dyer and Murphy, 2009)]. In comparison, when bees learn a target color in the presence of perceptually similar distractor stimuli (termed differential conditioning), they can master relatively fine color discriminations [<0.10 hexagon units (Dyer and Chittka, 2004c; Dyer and Murphy, 2009)]. However, learning color discrimination with differential conditioning takes considerably longer for bees, than with absolute conditioning. For example, learning a color distance of about 0.04-0.08 hexagon units with 75% accuracy typically takes honeybees or bumblebees about 50-60 visits (Burns and Dyer, 2008; Dyer and Chittka, 2004a; Dyer and Chittka, 2004b, c; Dyer and Murphy, 2009). There is evidence that this type of perceptually difficult discrimination places increased load on the information processing since individual bees will slow down to maintain accuracy when facing fine color discrimination problems (Chittka et al., 2003; Dyer and Chittka, 2004b), and will not perform at a high level of accuracy unless incorrect choices are punished with a bitter tasting substance (Avarguès-Weber et al., 2010; Chittka et al., 2003; Dyer, 2012a; Rodriguez-Gironés et al. 2013). Since perceptually similar color stimuli potentially place increased cognitive load on bee color judgments, it is important to understand the extent to which bees can reverse-learn such fine discrimination tasks, and to determine if there are differences in decision-making behavior between individuals processing similar colors.

In this current study we combine behavioral testing of free-flying bees that had to solve perceptually difficult color reversal tasks, with the use of ABMs, to interpret the potential colony level benefits of the diversity of observed behaviors.

2 METHODS: Behavioral experiments

Behavioral experiments were conducted with free-flying honeybees (*Apis mellifera* Linnaeus) at the biological gardens of Johannes Gutenberg University (Mainz, Germany) between July and September 2009. A hive of honeybees was maintained 10 m from a gravity feeder that provided 5% (vol.) sucrose solution. Individual bees were collected from the feeder site on a small Plexiglas spoon, and were transferred to a test site situated 15 m from the feeder, and 20 m from the hive. At the test site individual bees were marked with a color code on their thorax. Each bee was tested individually, and testing typically took 5-6 h/bee.

The training apparatus consisted of a vertical rotating screen of 50 cm diameter. A photograph of the rotating screen is presented in a previous study (Dyer et al., 2005). This apparatus presented stimuli on four 6×8 cm hangers (e.g. *ibid* figure 1), each with a small landing platform. The spatial position of stimuli was continuously changed during training, and hangers were exchanged for fresh ones to exclude olfactory cues (Dyer et al., 2005; Dyer et al., 2008a). Two of the hangers presented target stimuli, and two hangers presented distractor stimuli, in random positions. This apparatus represents a complex natural scenario where bees have to forage from several flowers of potentially similar color; however, the hangers enable well-defined data collection by counting choices (touches to the landing stage of stimuli). In an experiment, bees were rewarded with 25% sucrose for making correct choices on a designated target stimulus (a rewarded stimulus during a particular component of an experiment), whilst a distractor stimulus only presented plain water. When a bee landed on a target stimulus it was presented with an additional sucrose drop on a Plexiglas spoon so that it could be moved 1 m away to behind a small screen, so that stimuli could be exchanged (Dyer et al., 2005; Dyer et al., 2008a).

Stimuli were 6×6 cm colored cards that were of turquoise and blue appearance (Tonpapier no.s 32 & 37 respectively, Baehr, Germany) to a human observer, respectively. Stimuli spectral properties were measured with a spectrophotometer

(Spectro 320; Instrument Systems, München, Germany), in the range between 300 and 650 nm. The color visual system of the honeybee is different to human vision, and is based on three spectrally different types of photoreceptors maximally sensitive in the ultraviolet (UV), blue (B) and green (G) regions of the electromagnetic spectrum (Dyer et al., 2011). To quantify the color difference between the turquoise and blue stimuli we used a color hexagon model designed for hymenopteran trichromatic vision (Chittka, 1992), considering the previously measured electrophysiological recordings for honeybee photoreceptors (Peitsch et al., 1992), data for standard daylight conditions converted to photon flux units (Wyszecki and Stiles, 1982), and assuming the visual system was adapted to the grey plastic background of the training apparatus (Chittka, 1992). The color distance was 0.06 hexagon units. This color discrimination task thus involves colors that are sufficiently close so as to require differential conditioning to be learnt by bees (Avarguès-Weber et al., 2010; Dyer and Neumeyer, 2005).

A total of 32 bees were tested. The experiment was counterbalanced in a randomized fashion so that the ‘turquoise’ and ‘blue’ stimuli were each initial targets for half of the bees. Each bee was first provided with a form of absolute conditioning to the target stimulus for 30 decisions, where the distractor hangers only presented the grey background material. This ensured the bees were familiar with the experimental apparatus, and also enabled within-subject testing of whether absolute conditioning enabled any color learning of the perceptually difficult color discrimination task.

Following this initial absolute conditioning, each bee received differential conditioning to the similar turquoise and blue stimuli for 60 decisions. At the conclusion of the 60th decision, each test bee was satiated on the initial target stimulus and allowed to return to the hive. When a bee returned to the apparatus, all four hangers next presented it with the initial distractor stimulus to which this particular bee was exposed (e.g. blue or turquoise), now offering a reward, to allow for priming to the reverse contingency. The bee was allowed to collect sucrose from four landings and was satiated on the 4th hanger, so that it would return to the hive. When the test bee returned again it received differential conditioning with the initial reward situation switched, i.e. what was a distractor is now a target and *vice versa*. This reverse training lasted for 20 decisions. The bee was primed, satiated and the

rewards reversed again for another 20 decisions as described above. Next, in the final phase of the experiment, the rewards were reversed once more for 60 decisions. Thus each bee had to solve a complex set of reverse-learning tasks that involved perceptually difficult discriminations where individual discrimination decisions were potentially ambiguous. For statistical analysis data were checked for normal distribution and arcsine transformed where necessary. Graphical representations of the reversal described above are provided with the results.

3 RESULTS: Behavioral experiments

To evaluate if the bees had learnt the target stimulus following the initial absolute conditioning for 30 decisions, we considered the frequency of correct choices for the target in the first 10 decisions of the differential conditioning phase. For the ‘turquoise’ target group the frequency of correct choices for the target was 52.4% (± 16.9 s.d.), which was not significantly different from chance levels (1-sample t-test, $N=16$ bees, $t=0.560$, $df15$, $p=0.584$). For the ‘blue’ target group the frequency of correct choices for the target was 47.9% (± 15.5 s.d.), which was not significant (1-sample t-test, $N=16$ bees, $t=0.533$, $df15$, $p=0.602$). Finally, these pooled results were not significantly different from chance expectation (independent sample t-test, $t=0.774$, $df30$, $p=0.455$). Thus the two color stimuli were initially perceived or treated as very similar by the bees.

With differential conditioning, the bees learnt the visual task gradually, to the point where the mean frequency of correct choices in the interval from choice 50 – 60 was $73.5\% \pm 14.2$ s.d. (Fig. 1). With the first reversal there was an indication that some bees could switch their decisions quickly (Fig. 1, 2), but the choices for the initial target ($45.0\% \pm 23.7$ s.d.) were not what would be expected if the bees had completely switched their initially learnt color preference ($100-73.5 = 26.5\%$). Thus, considering previous evidence that there could potentially be individual differences between honeybees in how they reverse learn a perceptually difficult task (Mota and Giurfa, 2010), we next tested for evidence of bees having different abilities to reverse-learn.

To analyze the decision making of the bees we considered the variation in choices at a particular reversal (R_{1-3} , where R_1 denotes the first reversal etc.) for a given color, when it changed from being rewarding (a target) to being unrewarding (a distractor). For each bee a value representing the change in its individual frequency of correct choices was calculated for the three respective reversals (R_1 for 60-70, R_2 for 80-90 and R_3 for 100-110 choices). To statistically examine the magnitude of the reversal values (R_{1-3}) we considered the standard deviation (s.d.) of all bees during the last 10 decisions of the initial learning phase (decisions 1-60) of the visual task (s.d. was 14.1 and 14.2 for the ‘turquoise’ and ‘blue’ groups respectively) as the measure of behavioral variability. We then considered three categories of decision-making. Category 1 was no significant variation in choices (R_{1-3}) following a particular reversal, defined as within 1.645 s.d. (90% probability of a trained bee having data within this range by chance); these bees were termed ‘Stay’ bees. Category 2 was a change greater than 1.645 s.d. and less than 1.960 s.d.; these bees were termed ‘Fickle-circumspect’ bees. Category 3 was a change greater than 1.960 s.d. in choices during a reversal; these bees were termed ‘Deliberative-decisive’ bees. We used these values because they are well-accepted and robust statistical criteria such that 90% of data should lie within 1.645 sd. from the mean, and 95% of data within 1.960 s.d. of the mean (Sokal and Rohlf, 1981).

Subsequent analysis of the experiment revealed that 12% of bees (4/32) were consistent Deliberative-decisive bees (Fig. 2). There were four bees that did not reverse-learn. When the first experimental reward reversal occurred these four still chose the initial target with a mean accuracy of $72.6\% \pm 9.0$ s.d.. If these bees were Deliberative-decisive their choices for the initial target should have dropped to less than 30%. Therefore 12% of bees (4/32) ended up being classified as consistent ‘Stay’ bees. Throughout the three (multiple) reversals, these remained loyal to the color they initially learned (Fig. 3).

There were 24 bees that did not always remain faithful to just one strategy. Of these, seven bees initially exhibited a Stay bee strategy but moved to a Deliberative-decisive bee strategy, and nine bees were initially Deliberative-decisive and moved to a Stay bee strategy. There were eight bees that showed no clear pattern of loyalty to a strategy, some of which sometimes fell into the Fickle-circumspect strategy. Figure

4a shows a plot of the rank order of all the 32 bees tested considering the mean percentage change in preference for the designated target color that occurred for the three reversals.

4 METHODS: Simulation experiments

The analyses described above suggest that individual bees within a colony apply different strategies. These differences appear to follow a continuum, but often the same individuals retain one of several strategies over multiple reversals (Figure 4). We thus next used agent-based computer simulations to test the biological advantage of the colonies having bees with different strategies for dealing with switching between perceptually similar flower colors. The hypothesis for this ABM testing was that if a proposed strategy was biologically advantageous in nature, then the simulation experiments should find a case where this strategy outperforms the competing strategies. The null hypothesis is that a particular strategy does not perform better in any simulated condition than its competitors; this includes a scenario where individuals are just poor performers (Mota and Giurfa 2010).

The general purpose of our simulation code is to act as a framework in which to explore the interactions between insect pollinators and the flowers they visit. In tackling the specific hypothesis above, the spatial component of the model we describe below could be replaced with the sequential presentation of randomly selected flowers to each forager since our experiments only need model flowers distributed uniform-randomly. However in general, environmental spatial arrangement has the potential to impact on insect behavior; a thesis discussed previously in relation to our simulation framework (Bukovac *et al* 2013).

The agent-based model used for our simulation represents a foraging patch as a grid with dispersed flower-agents, and a colony of bee agents that search the patch to collect flower nutritional rewards. The simulation and its agents were updated synchronously and in discrete time steps. We conducted experiments by varying parameters for the distribution of foraging strategies of bees in a colony and for the availability and reliability of flower nutritional rewards. An explanation of the

simulation follows (see Fig. 5a for an overview), and the parameters are listed in Table 1.

4.1 Flower distribution and reward availability

We modeled two flower types, T1 and T2 of equal abundance, uniform randomly distributed, each in its own grid cell, across the foraging patch to simulate a temperate environment in which resources are typically randomly spaced (Dornhaus and Chittka, 2004; Visscher and Seeley, 1982). Flower agents may offer a reward unit to visiting bee-agents, or not, depending on the experiment. If a reward was present it was replenished after a bee-agent collected it, however, an individual bee-agent cannot return to a flower-agent within a single bout (foraging session) so that a particular reward is only available once for each bee-agent in a simulated bout. This assumption allowed us to isolate bee-agents from one another to avoid forager population density effects that were not central to the particular research questions under consideration in the main experiments. In control experiments we introduced ‘defective’ flower-agents that did not offer a reward even when their type was expected to do so. Thus by varying the relative percentage of defective flowers in an experiment, we additionally isolated population density effects.

4.2 The bee-agent and its colony

A flowchart of the algorithm each bee-agent follows is provided in Fig. 5b. Each bee-agent is mobile on the foraging grid, and has variable flower-agent landing probabilities. These are preferences (0-100%) for the respective flower-agent types (sum = 100%). Bee agents also have visual acuity and memory parameters that determine whether the bee realizes it has seen a flower before, and whether or not it correctly identifies a flower when detected. (See below and Table 1.)

To eliminate locality artefacts bee-agents were randomly placed on the grid to begin and end foraging bouts from any location (i.e. no hive location was assigned). Since we are interested in relative foraging strategy performance, flight speeds were irrelevant and fixed at one (orthogonal) grid cell per simulation time step. Bee-agents cannot co-occupy a grid cell or a flower it contains. Potential collisions are resolved in favour of a random bee-agent.

If a bee-agent occupies a grid cell containing a flower-agent it determines (using a strategy discussed below) whether to land and attempt reward collection. Each bee-agent memorises the flower-agents it just visited and will not revisit a flower-agent in a particular bout. After a foraging bout, bee-agents ‘forget’ the locations of previous visits. The justification for this is that animals have different memory phases and are likely to just use short-term memory for recalling recent choices made in complex foraging environments (Menzel, 2001). Apart from for the four most recently visited flowers, a bee-agent must enter a flower-agent’s grid cell to determine if it has already been visited. To ensure that flower-agents are not repeatedly re-approached bee-agents were modelled with accurate memory of the location of the last four flower-agents approached (Pyke and Cartar, 1992; Soltz, 1986).

At each simulation time step, each bee-agent examines its neighbourhood for flower-agents. The bee-agent randomly chooses an unapproached flower-agent as its destination, or, if no unapproached flower-agents are detected, the bee-agent chooses a random neighbouring cell. We recognise that bee foraging path selection is probably not entirely random (Pyke and Cartar, 1992), however, we are interested only in relative flower selection strategy performance, not absolute success of particular navigational approaches. In previous work we tested the random walk used in this simulation against a navigational strategy based on empirical data about a bee’s preferred direction of flight when departing a flower (Bukovac *et al* 2013). This showed that changes to the movement pattern affected foraging only by reducing overall success by a fraction of total nectar collected. Hence, as long as all bee-agents choose a path in the same way – randomly in our case – we can understand the relative success of different foraging strategies.

4.3 Flower detection

Bees can use multiple cues (vision and olfaction) to find flowers (Leonard *et al.*, 2011; Streinzer *et al.*, 2009). Our model only considered visual cues. Bee spatial acuity is relatively poor compared with a vertebrate lens eye (Land, 1997). In real life bumblebees can only detect a plant’s cluster of 3-5 flowers (each flower of 2.5 cm diameter) at a distance of < 0.7 m (Dyer *et al.*, 2008b; Wertlen *et al.*, 2008). Detection appears to be a step function (Dyer *et al.*, 2008b; Wertlen *et al.*, 2008) so we modeled bee-agent acuity as distances > 0.7 m — not detected, distances < 0.7 m — 95%

chance of detecting flowers. Thus a bee-agent can detect the presence of a flower-agent in a neighboring grid cell (up to $0.35\text{ m} + 0.35\text{ m} = 0.7\text{ m}$ away) with 95% accuracy, but cannot detect more distant flowers (Dyer et al., 2008b; Wertlen et al., 2008).

If a bee-agent is sharing a grid cell with a detected flower-agent the bee-agent has a probability of accurately discriminating a rewarding from an unrewarding flower-agent. This probability is modeled as the learnt percentage preference for a particular flower-agent type. This value represents the probability that the bee brain correctly determines if a given flower is a target or distractor. With these values it is possible that a bee-agent detects a target flower-agent, but chooses not to land if its preference for that type causes it to mistakenly perceive it as a distractor. Conversely a bee may choose to land on a distractor due to a perceptual error causing it to mistakenly perceive it as a target. This models the probabilistic way that bees discriminate between similar flower colors depending upon experience (Dyer 2012b).

All bee-agents undergo two phases of flower-agent preference adjustment behavior based on experimental results for our behavioral experiments. These phases allow a bee-agent to adjust its internal preferences for flower-agents based on previous successful and failed reward collection attempts. During phase 1, bee-agents conduct standard learning behavior that was modeled to be directly comparable to the learning behavior described above for real bees (Fig. 1: phase 1). During phase 2, bee-agents adopt either a Deliberative-decisive, Fickle-circumspect or Stay behavior based on the strategies we outlined above for the real bees (Fig. 1-3). The simulation of these phases is discussed below.

4.4 Simulation phase 1 – bee-agent preliminary information acquisition

All bee-agents commence phase 1 with equal 50% preferences for each flower type T1 and T2. Thus, initially a bee-agent has a 50% probability of choosing to land on the first flower-agent it encounters. This corresponds directly to how real bees initially generalize similar color flowers (Dyer, 2012b; Dyer et al., 2011) and is consistent with our behavioral data reported above.

During phase 1, each time a bee-agent visits a rewarding flower-agent its preference for that type (always T1 in phase 1) increases by 1% and its preference for the other

type (always T2 in phase 1) drops by 1%. Each time a bee-agent visits an unrewarding flower-agent preference for its type decreases by 1% and its preference for the other type of flower-agent increases by 1%. This learning model fits how real bees have been shown to learn in experiments (Dyer, 2012b; Dyer et al., 2011). Bee-agent's flower-agent preferences are clamped between 80 and 20% in keeping with realistic learning behavior of bees for perceptually difficult tasks (Dyer, 2012b; Dyer et al., 2011) and our own behavioral results (Fig. 1).

At the conclusion of phase 1, entry into phase 2 occurs, modeling the potential reward reversal behavior observed in our behavioral experiments, and thus triggering a change in bee-agent foraging behavior.

4.5 Transition from simulation phase 1 to phase 2

A simulation always begins in a scenario where T1 flower-agents are rewarding and T2 are not. This situation holds throughout phase 1. Phase 1 continues until an average of 70% preference for T1 (and 30% preference for T2) is reached for the entire bee-agent colony. Then phase 2 begins.

4.6 Simulation phase 2 – Deliberative-decisive, Fickle-circumspect or Stay behavior

Throughout phase 2, each bee-agent adopts one of the following foraging strategies depending on the experiment.

Deliberative-decisive. A Deliberative-decisive bee-agent decisively reverses its preferences if, over a period of deliberation time its preferred flower type is consistently found to have become unrewarding and another type is detected as being consistently rewarding. When a switch occurs in the rewards offered by a flower-agent type (e.g. T1 flower-agents were rewarding and suddenly become unrewarding), a Deliberative-decisive bee-agent that encounters four unrewarding T1 and four rewarding T2 flower-agents, can invert its preference accuracy. For instance, if a Deliberative-decisive bee-agent had learnt through experience to have 70% preference for T1 and 30% for T2 flower-agents, after encountering four unrewarding T1 and four rewarding T2 this can become a 70% preference for T2 and 30% for T1. This bee-agent has discriminatory ability and sufficient neural flexibility to adjust to a new reward situation. This bee-agent's foraging closely fits the profile of the free-flying bees shown in

Figure 2.

Fickle-circumspect. A bee-agent may learn to make accurate decisions about flower-agent types based on experience as described for experimental phase 1. The Fickle-circumspect bee type continues to learn in phase 2 with a 1% change per flower-agent visit as it did initially in phase 1. Whenever the reward situation changes, this bee type incrementally adjusts its preferences. It is fickle in the sense that it is always willing to adjust preferences based on new evidence. It is circumspect in the sense that the changes it makes are slight. This bee can adapt, but only gradually. Fickle-circumspect bee preferences are clamped between 80 and 20% in keeping with the findings of our behavioral experiments (Fig. 1).

Stay. A Stay bee's preference 'stays' loyal to the flower-agent type for which it originally acquired a strong preference, even if that flower-agent type changes from rewarding to unrewarding. For example, once a Stay bee-agent reaches ~70% preference for a flower-agent type, this preference becomes hard-wired, even if the availability of rewards offered by the flower-agents changes. This bee-agent's foraging closely fits the profile of the free-flying bees shown in Figure 3.

4.7 Simulation experiments

We conducted experiments to determine the impact of varying different conditions on the foraging success of colonies, each containing 60 bee-agents, utilizing the following four foraging strategy compositions: all Fickle-circumspect bees, all Stay bees, all Deliberative-decisive bees, 1/3 mix of each bee type.

Alternating reward experiments

Real flowers typically offer rewards for only a period of time, and a flower that is rewarding in the morning may not be later in the day (and *vice versa*) (Chittka *et al.*, 1997; Townsend-Mehler and Dyer, 2012; Waddington *et al.*, 1981; Waddington and Heinrich, 1981; Willmer *et al.*, 2009). Thus we considered scenarios where the variation in the availability of rewards offered by flower-agents of type T1 and T2 alternates cyclically. The length of phase 2 was fixed at 14,400 time steps – a value chosen because it has many even factors and is large enough to potentially allow full

exploration of differences in bee-agent strategy – for all alternating reward experiments. This allowed us to use only even factors of the length of phase 2 as reward swapping periods to ensure that the number of time steps that T1 and T2 were rewarding was identical in every run, and to ensure that only complete oscillation periods were simulated (i.e. no reward periods were truncated by the end of the simulation). Simulated periods ranged from 120 time steps per cycle (exactly 120 changes of the flower reward situation during phase 2) to 7200 time steps per cycle (exactly 2 changes of the flower reward situation during phase 2).

Throughout phase 1 of these experiments, T1 is rewarding and T2 is not. From the start of phase 2 for n time steps, T2 is rewarding and T1 is not. For the next n time steps T1 is rewarding and T2 is not. This switching process continues throughout phase 2 where n is an even factor of 14,400 between 120 and 7200. (See Figure 6 for a generic run.)

Defective reward control experiments

When bees forage with coworkers or bees from other colonies operating in the same flower patch, there is some likelihood that a bee lands on a flower that should be offering a reward at the time of the visit, but the reward has recently been acquired by another bee (Chittka et al., 1997; Townsend-Mehler and Dyer, 2012; Waddington et al., 1981; Waddington and Heinrich, 1981). By varying the fraction of a rewarding flower-agent type that is defective (in the sense that they do not offer rewards when they should), we tested the relative impact of these conditions on bee-agents' foraging strategies. Phase 1 of these defective reward experiments was run as described above. During phase 2 of these experiments the proportion of the T1 flower-agents that were defective and contained no reward was systematically varied. T2 flower-agents remained unrewarding throughout the experiment.

Handling time control experiments

Flower handling time may impact on the relative success of different bee foraging strategies (Burns and Dyer, 2008; Chittka et al., 1997). For example, the requirement of switching between different flower species can influence the handling time of bumblebees (Laverty, 1994; Laverty and Plowright, 1988; Raine and Chittka, 2007b). Even if handling times are identical for rewarded and unrewarded landings, bees

making poor landing decisions could be expected to waste more time on flower handling for no reward (Burns and Dyer, 2008; Muller and Chittka, 2008). We thus tested for the impact of flower handling times by instigating a single reward swap at the transition from phase 1 to phase 2, using a series of flower handling times from 0 to 20 time steps. In this experiment bee-agents sat idle for a number of simulation time steps corresponding to flower handling time when they visited a flower-agent, regardless of whether or not a reward was collected.

Dependent variable

The dependent variables of interest are the relative amounts of nutrition collected during each simulation run by the four bee colony compositions under the experimental conditions. Nectar collection rate is a variable likely to be of high importance to colony survival in natural conditions where resources may be constrained at certain times of the year (Burns and Dyer, 2008; Mattila and Seeley, 2007; Raine and Chittka, 2007a).

4.8 Simulation model verification and validation

Simulation verification involved checking the correctness of our model by ensuring that the simulation behaved in accordance with the behaviors detailed above.

Validation checks included that:

- Learning bee-agents changed preferences in $\pm 1\%$ increments appropriately.
- Bee-agent approach (short-term) and visit (long-term) memories operate correctly.
- Stay bees did not change preferences during phase 2.
- Flower distribution mean was 50% T1, 50% T2.

The simulation was tested with and without bee agent short-term memory. The short-term memory had no effect on the relative success of the different foraging strategies.

The verification process allowed us to determine some derived values for the behavior of the system that were important for interpreting our results. These are given below.

- Check that Deliberative-decisive bees make a flower-agent preference switch after receiving four rewards from sources that were not expected to offer a reward, and four failed attempts to extract a reward from flowers where a reward was expected.

There is always some delay after a flower-agent reward availability swap until a Deliberative-decisive bee-agent makes a preference switch due to its need to visit a minimum of eight flower-agents. Mean time until bee-agent preference switch was 765 time-steps (273 s.d.) after the flower-agents changed their rewards (median switch time was 756 time steps). Values computed from a sample of 60 Deliberative-decisive bee-agents.

- Check transition condition from phase 1 to phase 2 occurs when the mean bee-agent flower-agent preference for T1 reaches 70%.

We found that the mean time to reach a preference of T1 = 70% (T2 = 30%) was 1131 time steps (45 s.d.). At this time, the bee-agent colonies had foraged on average a total of 667 units of reward (18 s.d.). Values computed from 40 randomized simulation runs with 60 bee-agents per colony.

The gradual increase in a bee-agent's preference for T1 flower-agents during the initial learning period (phase 1) is illustrated in Figure 6. This learning performance is consistent with fine color learning in foraging honeybees (Avarguès-Weber et al., 2010; Avarguès-Weber et al., 2011; Dyer, 2012b; Giurfa, 2004), bumblebees (Dyer and Chittka, 2004c; Dyer et al., 2011) and our own behavioral data (Fig. 1). Reward foraged is plotted on the lower (green) lines (Fig. 6). T1 preferences are plotted on the upper (blue) lines (Fig. 6).

5 RESULTS: Simulation experiments

5.1 Alternating reward experiments

Sample results for a 'validation test' colony containing only three bees, one of each foraging strategy, are illustrated (Fig. 6). This figure highlights the alternating phases of flower reward and the responses of each bee strategy to these changes.

The impact on mean reward foraged of varying the reward-swapping period during

phase 2 of the simulations is illustrated in Figure 7 (note the x-axis *log* scale). Swapping periods from 120 time steps (120 swaps during the 14,400 time step long phase 2) to 7200 time steps (2 reward swaps during phase 2) were simulated. At no single value for swapping period in this range could we find any indication that the 100% Stay strategy hive (illustrated in green, Fig. 7, depicting approximately constant foraging performance across the board as expected) or the mixed beehive (illustrated in dotted red, Fig. 7) was most effective. However, the all Fickle-circumspect strategy hives were consistently the best performers for short reward swapping periods until the all Deliberative-decisive strategy hives clearly took over as the most effective foragers when reward swap period reached ~1800 time steps (Fig. 7). As swapping period continued to increase, the difference between Deliberative-decisive and Fickle-circumspect strategies increased also. The success of the two strategies approached one another again as the time delay for the Fickle-circumspect bees to adjust their preferences became less relevant with increasing flower reward oscillation period, and correspondingly, with a decrease in the number of reward swaps.

On the basis of these simulation results, there was no evidence to support the hypothesis that Stay bees should exist as a major group within a hive under the conditions we tested. Therefore, we tentatively conclude it is likely that the bees we initially classified as Stay bees were more likely to be Fickle-circumspect bees that, under our behavioral test conditions, had received insufficient time to learn. Figure 4b shows the same data as Figure 4a, now re-plotted in keeping with this conclusion. It distinguishes only between the Fickle-circumspect and Deliberative-decisive strategies that were observed in the behavioral experiments and whose biological advantage was supported by our ABMs.

5.2 Control experiments

Defective reward simulation runs (N=20) determine the impact of defective flower-agents on foraging strategies by providing a source of noise that could potentially confuse bee-agents as to which flower type is rewarding. We separately considered the 0, 30, 50, 70 and 90% defective flower distributions by testing how the four different hive compositions collected nutrition for each distribution.

There was no significant difference in the dependent variable of mean reward

collected for any of the defective reward experiments [One-way ANOVA, d.f. (3,76); 0% condition $F = 1.992$, $p=0.122$; 30% $F = 0.886$, $p=0.452$; 50% $F = 1.181$, $p=0.323$; 70% $F = 0.087$, $p=0.967$; 90% $F = 1.277$, $p=0.268$], showing that the relative foraging of the bee-agent hives with different strategies was independent of defective flower density effects.

During handling time simulation runs ($N = 20$), a single reward swap occurred at the conclusion of simulation phase 1. From this time point, all bee-agents enacted their phase 2 change-of-preference strategies. In these runs, every landing on a flower caused a bee-agent to sit idle during the designated flower handling time.

We did not identify any condition where the mixed beehive composition was most effective for the flower handling times tested. Flower handling time did not impact on the ranking of the different hive compositions for any test condition.

6 DISCUSSION

The behavioral experiments suggest a complex picture for individual honeybees having to reverse-learn a difficult color discrimination task. Whilst some honeybees were loyal to one of several strategies, other bees appeared to have flexibility to change between different strategies. A classic problem in understanding decision-making in insects is, why have colonies evolved different strategies for solving perceptually difficult problems (Burns, 2005; Burns and Dyer, 2008)? Agent-based simulations allowed us to test hypotheses elicited by this question, using empirical results of behavioral experiments as a basis.

In the current study, the empirical analyses initially suggested that there were three types of bee strategy for collecting nutrition, in a situation involving similarly colored flowers that could alternate in the availability of rewards offered. The potential strategies identified were Deliberative-decisive, Fickle-circumspect and Stay strategies. We found that Stay strategy honeybees never outperformed the other two strategies under the test conditions, often to the detriment of the hive's overall efficiency (Fig. 7). This leads us to question what ecological factors *might* maintain Stay bees, or if Stay bees really exist in nature. Stay bees may be useful by persevering with their hard-won knowledge, being effective foragers in situations where there are only occasional unrewarding target flowers. However, the control

experiment on defective flowers allows us to discount this hypothesis. Even when many defective target flowers were introduced to the ABM, Stay bees never outperformed the Deliberative-decisive and Fickle-circumspect bees.

Another possibility is that the bees that we classified as Stay bees following our initial analysis of the empirical data may have actually been Fickle-circumspect bees that learn very slowly (Fig. 3). The individual-based simulation allowed us to conclude that, under the conditions tested, bee colonies constrained to the requirement of collecting the maximum amount of nutrition per unit time should contain bees with only two types of decision-making strategy for complex color tasks (Figures 4 & 7). Following the learning of a perceptually difficult task, our results suggest that under repeated and frequent changes in a reward situation, it is best to make frequent and small changes to preferences in order to keep abreast of the changing environment, without missing an opportunity to fine-tune the perceptual system. Even though this fine-tuning may not bring the perceptual system fully into line with the new reward situation (Fig. 6), the minor adjustments take little time to be swayed in the reverse direction, and so the bee may simply revert to earlier preferences when needed using the same mechanism. However, under situations involving less frequent changes, this type of fine-tuning is slow to capitalize. Instead, once sufficient evidence has been gathered that the reward situation has changed, a large and decisive change of preferences is a more effective way to bring a bee's perceptually based decisions into line with new foraging conditions. The difference between these two strategies decreases as the time taken to switch preferences becomes less significant (i.e. when there are less reward swaps and they have long stable periods between them). Interestingly, there doesn't appear to be just two distinct categories in nature. Figure 4 shows that, considering the mean change in target flower preferences by individual bees following a reward reversal, there is a continuum of performance differences from the bees that do just exhibit solely Fickle-circumspect, or Deliberative-decisive strategies. Indeed many bees lie in between these two proposed categories. Previous work on both bumblebees (Chittka et al. 2003) and honeybees (Burns and Dyer 2008) has reported that while there is evidence of fast-inaccurate and slow-accurate individual bees for solving perceptually difficult tasks, there appears to also be a continuum between these 'different' strategies. This suggests that differences in individual animal differences for solving

perceptually difficult tasks are variable through the bee population within a hive. Future work may consider if such observations are constant within individuals (Chittka et al. 2003), or may be influenced by factors like age or level of experience. However, in bumblebees some capacity to modulate decisions within individual has been observed (Chittka et al. 2003), and the current study also observed that honeybees varied their strategies (Fig. 4).

Future work could consider whether there are any real-world conditions for which the presence of Stay bees is of benefit to a beehive. We note that the better the Deliberative-decisive bees learn during phase 1 to prefer T1 flowers, the longer it will take them to reverse their preferences after a reward swap in favor of T2, or any subsequent reversals. This occurs because their preferences for the currently unfavorable flower will be so low that they seldom visit to gain evidence for any newly instigated favorability. Perhaps the presence of Stay bees in a hive, together with Deliberative-decisive bees, may be useful if the bees can communicate with one another in some circumstances (Dornhaus et al., 2006). In particular, a Stay bee loyal to one flower type might act as a ‘watch-bee’ consistently monitoring flowers that have been previously found rewarding, ready to inform the hive once they become viable targets.

In complex real world environments there is often no way for animals in general to know which conditions will be encountered, and thus it could be beneficial for a colony to possess individuals with multiple strategies, in agreement with bet hedging hypotheses (Burns and Dyer, 2008; Muller and Chittka, 2008). Figure 7 supports this hypothesis in the case of bees from both our behavioral and simulation experiments.

Classically it has been appreciated that one form of decision-making in humans can be described as ‘Rationalistic’ in which a human subject “becomes aware of a problem, posits a goal, carefully weighs alternative means, and chooses among them according to his estimates of the respective merit with regard to the state of affairs he prefers” (Etzioni, 1967; Tarter and Hoy, 1998). Although the above description of the mental processes underlying this strategy is not applicable to insects, the outcomes, in terms of measured behavior, parallel those of bees employing the Deliberative-decisive strategy. Alternatively, decision-making in human subjects may use an *Incrementalist* strategy. This strategy “seeks to adapt decision-making strategies to the limited cognitive capacities of decision-makers and to reduce the scope and cost

of information collection and computation” (Etzioni, 1967; Tarter and Hoy, 1998), by making a continual stream of micro-adjustments. This parallels the behavior we observed in Fickle-circumspect bees. In the field of Artificial Intelligence, which is concerned with decision-making agents, these different strategies are understood in terms of the concepts of *simple reflex agents* that select their behavior based purely on their current situation; and *model-based reflex* agents that maintain an internal model of the part of the world that has been visible to them in the past in order to take this into account in their decision making (Russell and Norvig, 2010, pp. 48-52).

Our finding that bees demonstrate a variety of decision-making strategies is therefore consistent with theories of decision making in both humans and artificial intelligence. This suggests that results from the current honeybee and simulation experiments have widespread implications across a number of fields. Indeed the beehive may, by maintaining a diversity of individual level behavior, be acting as a super-organism (Hölldobler and Wilson, 2009; Seeley, 1989) with respect to its implementation of decision-making for complex problems that have no obvious, complete *a priori* solution. In addition, the behavioral experiments found some evidence that individual bees have a capacity to modulate their own decision-making approach, and modulation of decisions has previously been reported for bumblebees making speed/accuracy judgments for stimuli of similar color (Chittka et al., 2003; Dyer and Chittka, 2004b). However, there may be differences between visual processing and decision-making in honeybees and bumblebees (Dyer et al., 2008b; Morawetz and Spaethe, 2012; Raine and Chittka, 2012). This would be a fertile topic for further investigation.

We conclude that for a perceptually difficult task, honeybees demonstrate a variety of methods for adjusting their flower preferences when environmental conditions change. Our agent-based simulations of this behavior reveal that this diversity of strategies allows a hive to efficiently collect resources in complex ecological conditions.

Acknowledgements

We thank Kevin Korb for discussions. A.G.D. is grateful to the Alexander von Humboldt Foundation and the ARC DP0878968/0987989/130100015 for funding support. A.D. is grateful to ARC DP1094064/130100015 for funding support.

REFERENCES

- Avarguès-Weber A, de Brito Sanchez MG, Giurfa M, Dyer AG. Aversive reinforcement improves visual discrimination learning in free flying honeybees. *PLOS One* 2010; 5:e15370.
- Avarguès-Weber A, Deisig N, Giurfa M. Visual cognition in social insects. *Annu Rev Entomol* 2011; 56:423–43.
- Biesmeijer J, Seeley TD. The use of waggle dance information by honeybees throughout their foraging careers. *Behav Ecol Sociobiol* 2005; 59:133-42.
- Bitterman ME, Menzel R, Fietz A, Schäfer S. Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *J Comp Psychol* 1983; 97:107-19.
- Bukovac, Z, Dorin, A, Dyer AG, A-Bees See: A Simulation to Assess Social Bee Visual Attention During Complex Search Tasks, in Proceedings of the 12th European Conference on Artificial Life (ECAL 2013), Liò et al (eds.) Taormina, Italy, 2-6 Sept., MIT Press, pp. 276-283.
- Burns JG. Impulsive bees forage better: the advantage of quick, sometimes inaccurate foraging decisions. *Anim Behav* 2005; 70:e1-e5.
- Burns JG, Dyer AG. Diversity of speed accuracy strategies benefits social insects. *Curr Biol* 2008; 18:R953–R954.
- Chittka L. The color hexagon: a chromaticity diagram based on photoreceptor excitations as a generalized representation of colour opponency. *J Comp Physiol A* 1992; 170:533-43.
- Chittka L. Sensorimotor learning in bumblebees: long term retention and reversal training. *J Exp Biol* 1998; 201:515-24.
- Chittka L, Dyer AG, Bock F, Dornhaus A. Bees trade off foraging speed for accuracy. *Nature* 2003; 424:388.
- Chittka L, Gumbert A, Kunze J. Foraging dynamics of bumble bees: correlates of movements within and between plant species. *Behav Ecol* 1997; 8:239-49.
- Chittka L, Schürkens S. Successful invasion of a floral market. *Nature* 2001; 411:653.
- Chittka L, Thomson JD, Waser NM. Flower constancy, insect psychology, and plant evolution. *Naturwiss* 1999; 86:361-77.
- Dafni A. Mimicry and deception in pollination. *Ann Rev Ecol Syst* 1984; 15:259-78.
- de Vries H, Biesmeijer JC. Modelling collective foraging by means of individual behaviour rules in honey-bees. *Behavioural Ecology and Sociobiology* 1998; 44:109-24.
- DeAngelis DL, Mooij WM. Individual-based modelling of ecological and evolutionary processes. *Annu Rev Ecol Evol Syst* 2005; 36:147-68.
- Dorin A, Korb KB, Grimm V. Artificial-Life Ecosystems: What are they and what could they become? In: Eleventh International Conference on Artificial Life (Bullock S, Noble J, Watson RA, Bedau MA, eds): MIT Press 2008; 173-80.
- Dornhaus A, Chittka L. Why do honeybees dance. *Behavioural Ecology and Sociobiology* 2004; 55:395-401.

- Dornhaus A, Kluegl F, Oechslein C, Puppe F, Chittka L. Benefits of recruitment in honey bees: effects of ecology and colony size in an individual-based model. *Behav Ecol* 2006; 17:336-44.
- Dyer AG. The mysterious cognitive abilities of bees: why models of visual processing need to consider experience and individual differences in animal performance. *J Exp Biol* 2012a; 215:387-95.
- Dyer AG. Psychophysics of Honey Bee Color Processing in Complex Environments. In: *Honeybee Neurobiology and Behavior*. Netherlands: Springer 2012b; 303-14.
- Dyer AG, Chittka L. Biological significance of discriminating between similar colours in spectrally variable illumination: bumblebees as a study case. *J Comp Physiol A* 2004a; 190:105-14.
- Dyer AG, Chittka L. Bumblebees (*Bombus terrestris*) sacrifice foraging speed to solve difficult colour discrimination tasks. *J Comp Physiol A* 2004b; 190:759-763.
- Dyer AG, Chittka L. Fine colour discrimination requires differential conditioning in bumblebees. *Naturwiss* 2004c; 91:224-7.
- Dyer AG, Murphy AH. Honeybees choose "incorrect" colors that are similar to target flowers in preference to novel colors. *Isr J Plant Sci* 2009; 57:203-10.
- Dyer AG, Neumeyer C. Simultaneous and successive colour discrimination in the honeybee (*Apis mellifera*). *J Comp Physiol A* 2005; 191:547-57.
- Dyer AG, Neumeyer C, Chittka L. Honeybee (*Apis mellifera*) vision can discriminate between and recognise images of human faces. *J Exp Biol* 2005; 208:4709-14.
- Dyer AG, Paulk AC, Reser DH. Colour processing in complex environments: insights from the visual system of bees. *Proc Royal Soc B* 2011; 278:952-9.
- Dyer AG, Rosa MGP, Reser DH. Honeybees can recognise images of complex natural scenes for use as potential landmarks. *J Exp Biol* 2008a; 211:1180-6.
- Dyer AG, Spaethe J, Prack S. Comparative psychophysics of bumblebee and honeybee colour discrimination and object detection. *J Comp Physiol A* 2008b; 194:617-27.
- Etzioni A. Mixed-Scanning: A "Third" Approach to Decision Making. *Publ Admin Rev* 1967; 27:385-92.
- Giurfa M. Conditioning procedure and color discrimination in the honeybee *Apis mellifera*. *Naturwiss* 2004; 91:228-31.
- Giurfa M, Núñez J, Backhaus W. Odour and colour information in the foraging choice behaviour of the honeybee. *J Comp Physiol A* 1994; 175:773-9.
- Grimm V. Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecological Modelling* 1999; 115:129-48.
- Grimm V, Railsback SF. *Individual-based modeling and ecology*. Princeton, N.J.: Princeton University Press; 2005.
- Grimm V, Revilla E, Berger U, Jeltsch F, Mooij WM, Railsback SF, Thulke HH, Weiner J, Wiegand T, DeAngelis D. Pattern-Oriented Modelling of Agent-

- Based Complex Systems: Lessons from Ecology. Science (New York, NY) 2005; 310:987-91.
- Heinrich B. Bumblebee economics. Cambridge: Harvard University Press; 1979.
- Hogeweg P, Hesper B. The ontogeny of the interaction structure in bumble bee colonies: A MIRROR Model. . Behav Ecol Sociobiol 1983; 12:271-83.
- Hölldobler B, Wilson EO. The Superorganism: The Beauty, Elegance, and Strangeness of Insect Societies. New York: WW Norton & Company; 2009.
- Huston M, DeAngelis D, Post W. New Computer Models Unify Ecological Theory. BioSci 1988; 38:682-91.
- Judson OP. The rise of the individual-based model in ecology. Trends in Ecology and Evolution 1994; 9:9-14.
- Komischke B, Giurfa M, Lachnit H, Malun D. Successive olfactory reversal learning in honeybees. Learning and memory 2002; 9:122-9.
- Land MF. Visual acuity in insects. Ann Rev Entomol 1997; 42:147-77.
- Laverty TM. Bumble bee learning and flower morphology. Anim Behav 1994; 47:531-45.
- Laverty TM, Plowright RC. Flower handling by bumblebees: a comparison of specialists and generalists. Anim Behav 1988; 36:733-40.
- Leadbeater E, Chittka L. Social learning in insects - From miniature brains to consensus building. Curr Biol 2007; 17:R703-R713.
- Leonard AS, Dornhaus A, Papaj DR. Flowers help bees cope with uncertainty: signal detection and the function of floral complexity. Journal of Experimental Biology 2011; 214:113-21.
- Mattila HR, Seeley TD. Genetic diversity in honey bee colonies enhances productivity and fitness. Science (New York, NY) 2007; 317:362-4.
- Menzel R. Das Gedächtnis der Honigbiene für Spektralfarben. II. Umlernen und Mehrfachlernen. Z Vergl Physiol 1969; 63:290-309.
- Menzel R. Behavioral and neural mechanisms of learning and memory as determinants of flower constancy. In: Cognitive Ecology of Pollination (Chittka L, Thomson JD, eds). Cambridge: University Press; 2001. p21-40.
- Morawetz L, Spaethe J. Visual attention in a complex search task differs between honeybees and bumblebees. J Exp Biol 2012; 215:2515-23.
- Mota T, Giurfa M. Multiple reverse olfactory learning in honeybees. Frontiers Behav Neurosci 2010; 4:1-9.
- Muller H, Chittka L. Animal Personalities: The Advantage of Diversity. Curr Biol 2008; 18:R961-R963.
- Pavlov IP. Lectures on conditioned reflexes. New York: International Publishers. 1927.
- Peitsch D, Fietz A, Hertel H, de Souza J, Ventura DF, Menzel R. The spectral input systems of hymenopteran insects and their receptor-based colour vision. J Comp Physiol A 1992; 170:23-40.

- Pyke GH, Cartar RV. The flight directionality of bumblebees: do they remember where they came from? *Oikos* 1992; 65:321-7.
- Raine NE, Chittka L. The adaptive significance of sensory bias in a foraging context: floral colour preferences in the bumblebee *Bombus terrestris*. *PLoS One* 2007a; 2:e556.
- Raine NE, Chittka L. Flower constancy and memory dynamics in bumblebees (Hymenoptera: Apidae: *Bombus*). *Entomol Gen* 2007b; 29:179-99.
- Raine NE, Chittka L. No trade-off between learning speed and associative flexibility in bumblebees: a reversal learning test with multiple colonies. *PLoS One* 2012; 7:e45096.
- Reser DH, Wijesekara Witharanage R, Rosa MGP, Dyer AG. Honeybees (*Apis mellifera*) Learn Color Discriminations via Differential Conditioning Independent of Long Wavelength (Green) Photoreceptor Modulation. *PLoS One* 2012; 7:e48577.
- Rodriguez-Gironés MA, Trillo A, Corcobado G. Long term effects of aversive reinforcement on colour discrimination learning in free-flying bumblebees. *PLoS One* 2013; *PLoS One* 8: e71551.
- Russell SJ, Norvig P. *Artificial Intelligence, a Modern Approach*, 3rd ed. New Jersey: Pearson Education/Prentice Hall. 2010.
- Scheiner R, Erber J, Page RE. Tactile learning and the individual evaluation of the reward in honey bees (*Apis mellifera* L.). *J Comp Physiol A* 1999; 185:1-10.
- Scheiner R, Weiss A, Malun D, Erber J. Learning in honey bees with brain lesions: how partial mushroom-body ablations affect sucrose responsiveness and tactile antennal learning. *Animal Cognition* 2001; 3:227-35.
- Seeley TD. *Honeybee ecology. A study of adaptation in social life*. Princeton: Princeton University Press. 1985.
- Seeley TD. The honey bee colony as a superorganism. *Am Sci* 1989; 77:546-53.
- Sokal RR, Rohlf FJ. *Biometry*, 2nd ed., WH Freeman & Co., NY. 1981.
- Soltz RL. Foraging path selection in bumblebees: hindsight or foresight? *Behav* 1986; 99:1-21.
- Streinzer M, Paulus HF, Spaethe J. Floral colour signal increases short-range detectability of a sexually deceptive orchid to its bee pollinator. *J Exp Biol* 2009; 212:1365-70.
- Tarter cJ, Hoy WK. Toward a contingency theory of decision making. *Journal of Educational Administration* 1998; 36:212 - 28.
- Townsend-Mehler JM, Dyer FC. An integrated look at decision making in bees as they abandon a depleted food source. *Behav Ecol Sociobiol* 2012; 66:275-86.
- Visscher PK, Seeley TD. Foraging strategy of honeybee colonies in a temperate deciduous forest. *Ecology* 1982; 63:1790-1801.
- von Helversen O. Dressurversuche mit Bienen. *Biol Unserer Zeit* 1974; 4:119-124.

- Waddington KD, Allen T, Heinrich B. Floral preferences of bumblebees (*Bombus edwardsii*) in relation to intermittent versus continuous rewards. *Anim Behav* 1981; 29:779-84.
- Waddington KD, Heinrich B. Patterns of movement and floral choice by foraging bees. In: *Foraging Behavior* (Kamil AC, Sargent TD, eds). New York: Garland STPM Press; 1981. p215-30.
- Wertlen AM, Niggebrugge C, Vorobyev M, Hempel de Ibarra N. Detection of patches of coloured discs by bees. *J Exp Biol* 2008; 211:2101-4.
- Willmer, P., et al., Bidirectional Flower Color and Shape Changes Allow a Second Opportunity for Pollination. *Curr Biol*, 2009. 19: p. 919-23.
- Wyszecki G, Stiles WS. *Color science: Concepts and methods, quantitative data and formulae*. New York: Wiley. 1982.

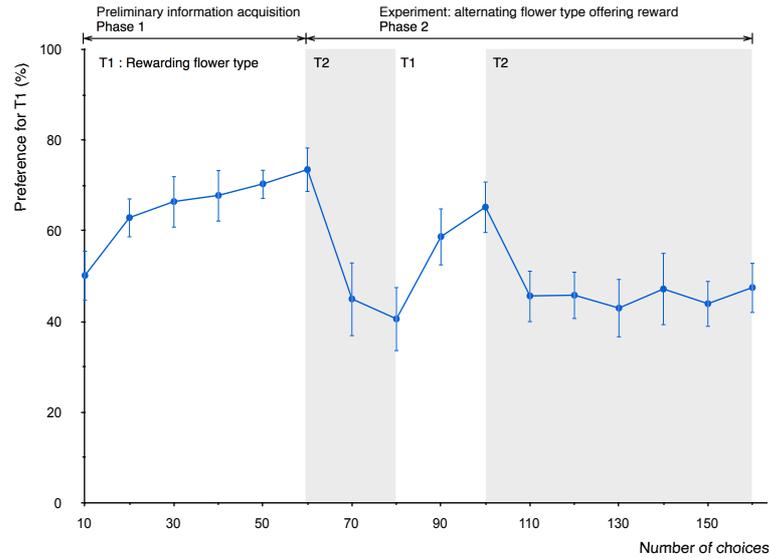


Fig. 1. Learning and switching behavior of all test bees. Mean ($\pm 95\%$ confidence interval) of 32 bees during differential conditioning and the three reversals occurring after 60, 80 and 100 choices.

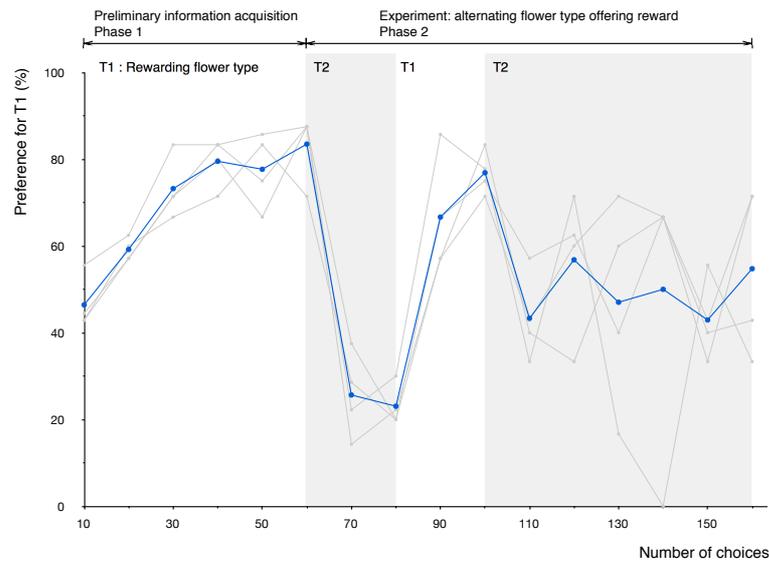


Fig. 2. Learning and switching behavior of Deliberative-decisive bees. Mean (dark blue line) of the four bees (grey lines) that always exhibited a Deliberative-decisive strategy during differential conditioning and the three reversals following 60, 80 and 100 choices.

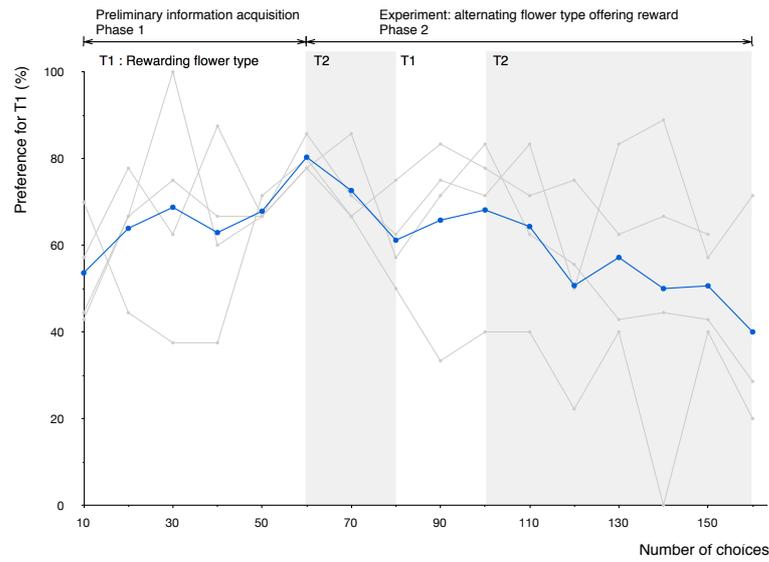


Fig. 3. Learning and switching behavior of Stay bees. Mean (dark blue line) of the four bees (grey lines) that always exhibited a Stay strategy during differential conditioning and the three reversals following 60, 80 and 100 choices.

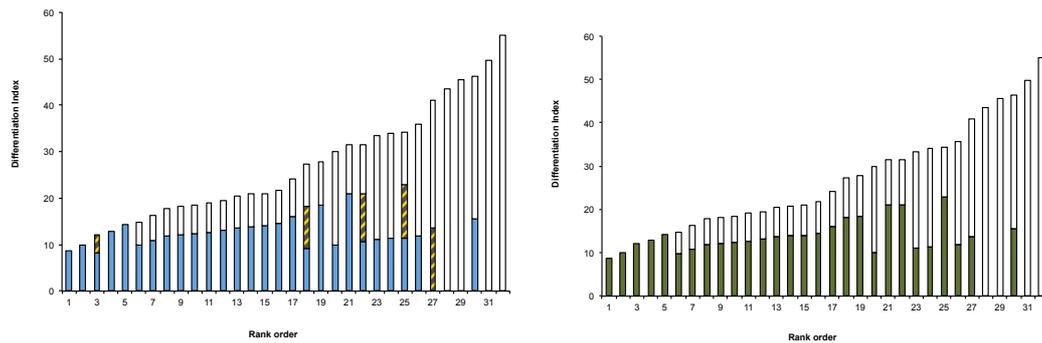


Fig. 4a & b. Rank order of 32 individual honeybees on a reverse-learn fine color discrimination task that was repeated three times. **(a)** The Y-axis shows Differentiation Index (mean absolute percentage change in a bee's preference for a given rewarding color when a switch (1-3) occurred). The X-axis shows rank order of percentage change for the individual bees. If a bee's choices following a reversal was within 90%CI for the initial target color it was classified as a Stay solution (blue bars), if a bee switched (>95%CI) it was classified as a Deliberative-decisive solution (white bars), and a bee with a change in behaviour between 90-95% CI was classified as a Fickle-circumspect bee (yellow hatched bars). Bars thus show the relative frequency of bees using the potentially different strategies. Many bees did use a mixture of strategies, but some bees were 100% loyal to a particular strategy and their respective Differential Index scores were statistically different (Compare solid white versus solid blue bars: Mann Whitney U, $p=0.021$). For the Deliberative-decisive bees (solid white bars), this observed behaviour was not due to chance distribution as there was a $1/8000$ chance of one individual successively switching three times, and using a binomial test the $\Pr(X_4) = 1 - \Pr(X < 4) = 8.75e^{-12}$ (i.e. $p\text{-value} < 0.0001$ for the observed number of bees exhibiting this multiple switching behaviour by chance), whilst in contrast some bees never switched preference (solid blue bars) and were always within initial 90% confidence limits for the initial target colour. **(b)** Following agent-based simulations to dissect why different proposed foraging strategies may exist in nature, there was little support for a Stay strategy being optimal in the conditions tested, and so possibly only two types of strategies may exist in nature. Thus honeybee behavior for reverse-learning perceptually difficult color discrimination tasks could be categorized as a Deliberative-decisive solution (white bars), or a more conservative Fickle-circumspect strategy (green bars).

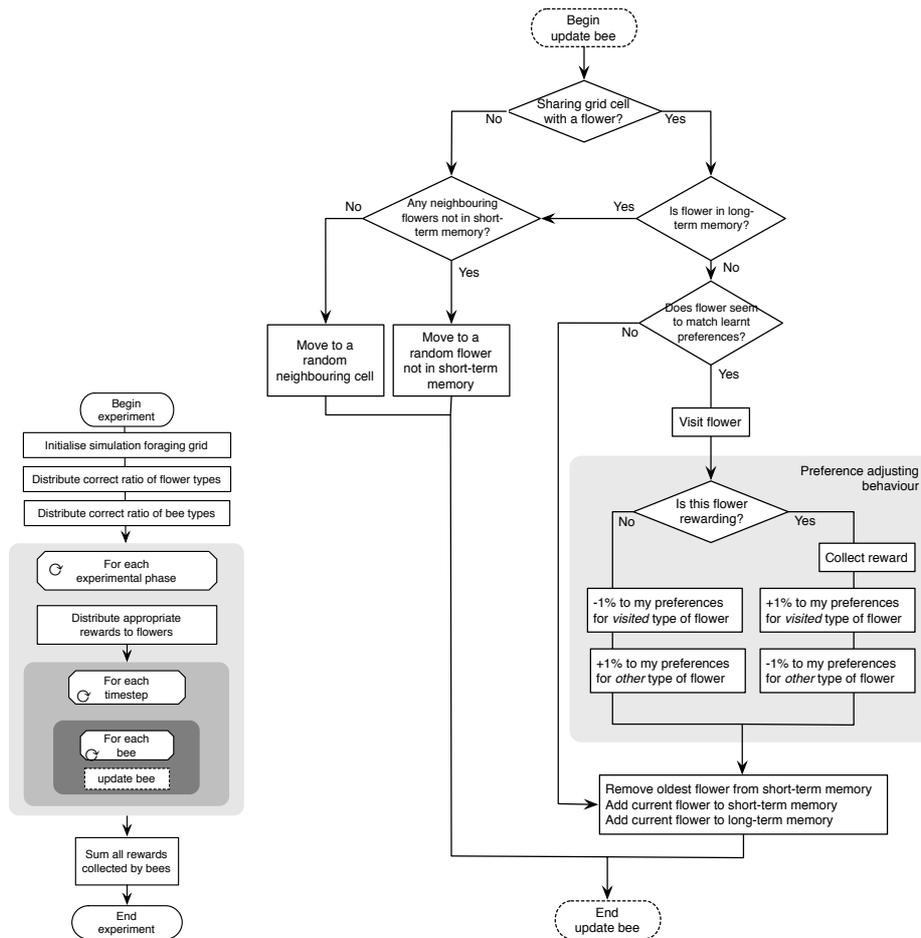


Fig. 5a & b. Flowcharts summarising the agent-based model. **(a)** The overall procedure for each experiment; **(b)** The procedure for updating an individual bee-agent every time step of a simulation. The preference adjusting behavior presented in this figure applies to all bee-agents during the first phase of every experiment, this being the preliminary information acquisition period which models the real bees' differential conditioning (section 4.4). The bee-agents' preference adjusting behavior at other times is detailed in the text (section 4.6).

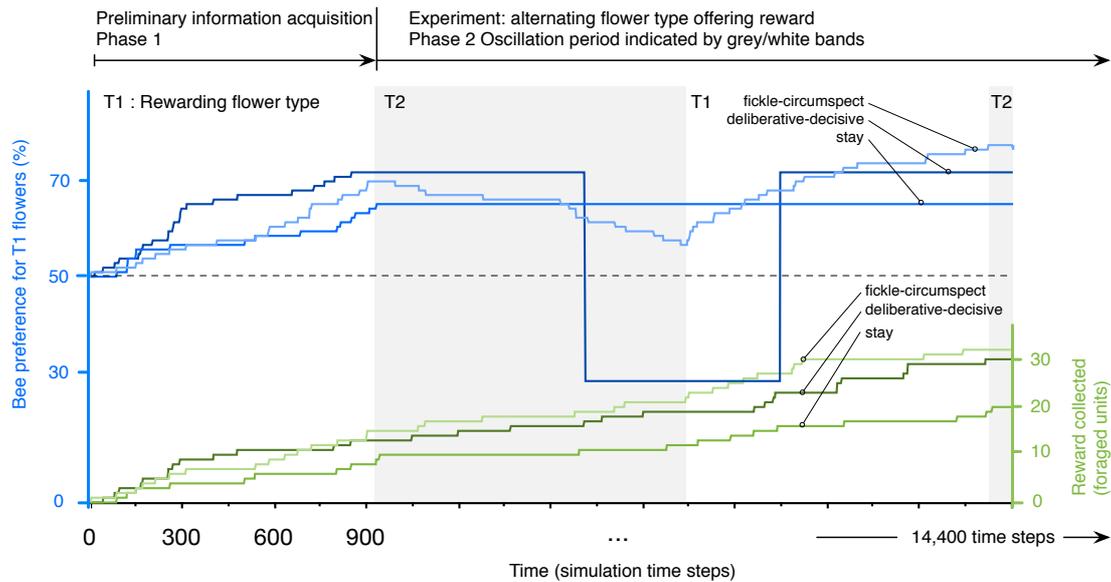


Fig. 6. Sample ‘alternating reward’ agent-based simulation experiment for single bee-agents of respective decision making strategies. The upper plots (blue) are the preferences of each bee-agent for T1 rising from an initial 50% value (vertical, left-hand axis). The lower plots show units foraged by the individual bee-agents during the depicted simulation (vertical, right-hand axis). The preliminary acquisition period during which the bee-agent colony reaches a mean preference for T1 of 70% is marked *phase 1*. From then on the reward availability is swapped periodically until the end of the simulation; this is *phase 2*. The availability of rewards from T2 and T1 during phase 2 is illustrated by alternating grey bands from the beginning of phase 2. The experiment continues like this until the end of phase 2 at time step 14,400 (not shown).

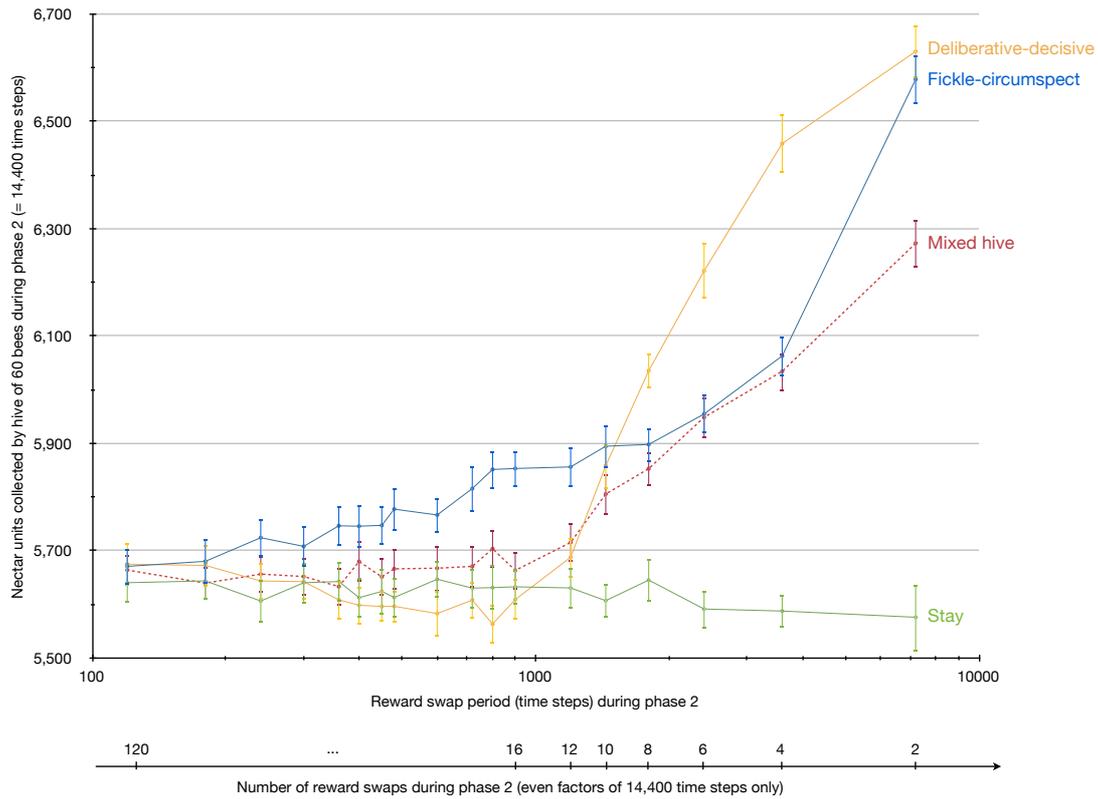


Fig. 7. Mean reward foraged by hives of different bee-agent compositions during simulation phase 2 of 14,400 time steps, as compared to the length of time (in simulation time steps plotted on a *log* scale) for which rewards are alternately offered by flower-agents of T1 and T2. Mean values based on 50 simulation runs. Error bars indicate 95% confidence intervals. X-axis data points correspond to even factors of the length of phase 2 (14,400 time steps) ensuring that T1 and T2 each offer rewards for an identical number of time steps in un-truncated periods during every simulation.

Environment	
Patch size	571 × 571 cells, torroidal boundary.
Grid cell size	0.35m x 0.35m
Flower-agents	
T1 : T2 flower relative abundance (mean)	1 : 1
Total flowers	13071
Bee-agents	
Colony sizes and compositions	60 agents/colony, composed as: all Fickle-circumspect, all Stay, all Deliberative-decisive, 1/3 mix of each.
Flower detection accuracy	95% from a neighboring cell or one shared with a flower.
Storage capacity	100 reward units.
Recently approached flower memory length.	4 most recent approaches.
Visited flower memory length.	Every flower visited on a single bout.
Learning increment (for all bees during phase 1, and for Fickle-circumspect bees during phase 2)	±1% point per flower visit.
Preference switch threshold (for Deliberative-decisive bees during phase 2)	Visit 4 surprising rewarding flowers and 4 surprising unrewarding flowers without encountering any unsurprising flowers. (A surprising visit occurs when a bee expects a reward but does not receive one, and <i>vice versa</i> .)
Simulation length	
Duration of phase 1 (variable length)	Until colony mean preference for T1 flowers ≥70%.
Duration of phase 2 (fixed length)	14,400 time steps.

Table 1. Summary table of simulation parameters.