The Bayesian Superorganism II: optimal foraging and the information theory of gambling

Roland J. Baddeley[1], Nigel R. Franks[2], Edmund R. Hunt[2]

[1]School of Experimental Psychology, University of Bristol, 12a Priory Road, Bristol, BS8 1TU, UK

[2]School of Biological Sciences, University of Bristol, Life Sciences Building, 24 Tyndall Avenue, Bristol, BS8 1TQ, UK

ORCIDS: RJB, 0000-0002-7431-6580; NRF, 0000-0001-8139-9604; ERH, 0000-0002-9647-124X

Authors for correspondence: email: edmund.hunt@bristol.ac.uk, roland.baddeley@bristol.ac.uk

Abstract

Part of the life cycle of an ant colony is simple: a colony collects resources; these resources are converted into more ants, and these ants in turn collect more resources. Because more ants collect more resources, this is a multiplicative process, and the expected logarithm of the amount of resources (equivalently, geometric mean (Dempster, 1955; MacArthur, 1965)) determines how successful the colony will be in the long run. Over 60 years ago, Kelly (1956) showed, using information theoretic techniques, that the rate of growth of resources for such a situation is optimised by probability matching. Thus, in the case of ants the fraction of the colony foraging at a given location should be proportional to the probability that resources will be found there, a result widely applied in the mathematics of gambling. This theoretical optimum generates predictions for which collective ant movement strategies might have evolved. In this second of three papers, we show how colony level optimal foraging behaviour can be achieved by mapping movement to Markov chain Monte Carlo methods, specifically Hamiltonian Markov chain Monte Carlo (HMC). This can be done by the ants following a (noisy) local measurement of the (logarithm of) the resource probability gradient (possibly supplemented with momentum, i.e. a propensity to move in the same direction). This maps the problem of foraging (via the information theory of gambling, stochastic dynamics and techniques employed within Bayesian statistics to efficiently sample from probability distributions) to simple models of ant foraging behaviour. This identification has broad applicability, facilitates the application of information theory approaches to understanding animal cognition, and

unifies insights from existing biomechanical, cognitive, random and optimality movement paradigms. At the cost of requiring ants to obtain (noisy) resource gradient information, we show that this model is both efficient, and matches a number of characteristics of real ant exploration.

Key words: Movement, collective behaviour, Bayesian methods, Markov chain Monte Carlo, foraging

Introduction

Life has undergone a number of major organisational transitions, from simple self-replicating molecules into complex societies of organisms (Maynard Smith and Szathmary, 1995). Social insects such as ants, with a reproductive division of labour between the egg-laying queen and non-reproductive workers whose genetic survival rests on her success, exemplify the highest degree of social behaviour in the animal kingdom: 'true' sociality or eusociality. The workers' cooperative genius is observed in diverse ways (Camazine et al., 2001) from nest engineering (Dangerfield et al., 1998) and nest finding (von Frisch, 1967), to coordinated foraging swarms (Franks, 1989) and dynamically adjusting living bridges (Reid et al., 2015). This has inspired a number of technological applications from logistics to numerical optimisation (Dorigo and Gambardella, 1997; Karaboga and Basturk, 2007). All of these behaviours may be understood as solving particular problems of information acquisition, storage and collective processing in an unpredictable and potentially dangerous world (Detrain et al., 1999). Movement (the change of the spatial location of whole organisms in time) is intrinsic to the process. Here we consider how optimal information processing is mapped to movement, at the emergent biological levels of the organism and the colony, the 'superorganism'. We develop a Bayesian framework to describe and explain the movement behaviour of ants in probabilistic, informational terms, in relation to the problem they are having to solve: the optimal acquisition of resources in an uncertain environment, to maximise the colony's geometric mean fitness (Orr, 2009). The movement models are compared to real movement trajectories from *Temnothorax* ants. This paper follows the first part of this three-part series, which was on collective probability estimation (Hunt et al., 2018b).

Operationalizing conceptual animal movement frameworks

Scientists have studied animal movement for many years and a recent framework has been proposed by Nathan et al (Nathan et al., 2008), with the aim of obtaining a more environmentally integrated understanding of it. Their framework identifies four components in a full description: the organism's internal state; motion capacity; navigational ability; and influential external environmental factors. This framework also characterises existing research as belonging to different paradigms, namely 'random' (classes of mathematical model related to the random walk or Brownian motion); 'optimality' (relative efficiency of strategies for maximising some fitness currency); 'biomechanical' (the 'machinery' of motion); and the 'cognitive' paradigm (how individuals' brains sense and respond to navigational information). However, scientists have yet to create a theoretical framework which convincingly unifies these components. Frameworks such as Nathan et al. are also focused on the individual and so for group-living organisms, especially for eusocial ones, they are incomplete. The concepts of search and uncertainty also need to be better integrated within foraging theory so that the efficiency of different movement strategies can be evaluated (Giuggioli and Bartumeus, 2010).

Here, we contend that animal foraging (movement) models should be developed with reference to the particular information processing challenges faced by the animal in its ecological niche, with information in this context referring to the distribution of fitness-relevant resources: in particular the location and quality of foraging patches, which are unknown *a priori* to the organism(s). Furthermore, an important 'module' in any comprehensive paradigm for animal movement is the role of the group and its goals in determining individual movement trajectories; there has been much research on collective behaviour in recent years, with information flow between individuals identified as an important focus of research (Sumpter, 2006). Eusocial insects like ants exhibit a highly advanced form of sociality, even being described as a 'superorganism', that is, many separate organisms working together as one (Hölldobler and Wilson, 2009). Their tremendous information processing capabilities are seen clearly in their ability to explore and exploit

collectively their environment's resources. Ants thrive in numerous ecological niches, and alone account for 15–20% of the terrestrial animal biomass on average, and up to 25% in tropical regions (Schultz, 2000).

The collective behaviour of tight-knit groups of animals like ants has been described as collective cognition (Couzin, 2009). Because a Bayesian framework seems natural for a single animal's decision-making (McNamara et al., 2006), an obvious challenge would seem to be applying its methods to describe the functioning of a superorganism's behaviour. First, we identify a simple model that describes the foraging problem the ants have evolved to solve.

Placing bets: choosing where to forage

Foraging for any organism is an ongoing process, with metabolic needs demanding to be fed. Evolution by natural selection should produce organisms that can be expected to have an efficient foraging strategy in their typical ecological context. In the case of an ant colony, although it consists of many separate individuals, each worker does not consume the food it collects and is not independent, but there is rather a colony-level foraging strategy enacted without central control that seeks to maximise colony fitness (Giraldeau and Caraco, 2000). Following the founding stage, during the ergonomic stage of a colony's life cycle (Oster and Wilson, 1978) the queen is devoted exclusively to egg-laying, while workers take over all other work, including collecting food. Thus the colony becomes a 'growth machine' (Oster and Wilson, 1978), whereby workers collect food to increase the reproductive rate of the queen, who transforms collected food into increased biomass or more numerous gene copies. Ultimately, the success or failure of this stage determines the outcome of the reproductive stage, where accumulated 'wealth' (biomass) correlates with more offspring colonies. This natural phenomenon has parallels with betting, where the winnings on a game may be reinvested to make a bigger bet on the next game. In the context of information theory, John Kelly made a connection between the rate of transmission of information over a communications channel, which might be said to noisily transmit the outcome of a game to a

gambler while bets can still be made, and the theoretical maximum exponential growth rate of the gambler's capital making use of that information (Kelly, 1956). To maximise the gambler's wealth over multiple (infinite) repeated games, it is optimal to bet only a fraction of the available capital each turn across each outcome, because although betting the whole capital on the particular outcome with the maximum expected return is tempting, any losses would quickly compound over multiple games and erode the gambler's wealth to zero. Instead, maximising logarithmic wealth is optimal, since this is additive in multiplicative games and prevents overbetting. Solving for this maximisation results in a probability matching or 'Kelly' strategy, where bets are made in proportion to the probability of the outcome. For instance, in a game with two outcomes, one of 20% probability and one of 80% probability, a gambler ought to bet 20% of his wealth on the former and 80% on the latter. This does not depend on the payoffs being fair with respect to the probabilities of the outcome, or $\sigma_i = \frac{1}{p_i}$, which in the aforementioned case would be 5 and 1.25. Instead it simply requires fair odds with respect to some distribution, or $\sum \frac{1}{\sigma_i} = 1$ where σ_i is the payoff for a bet of 1, so they could for instance be 2 to 1 or uniform odds in the case of a game with two outcomes (see supplementary Materials and Methods). For the purposes of our foraging model, we can simply impose the constraint of fair odds, and any distribution of real-life resource payoffs can be mapped to this when renormalized.

In the case of ants choosing where to forage, the probability matching strategy can be directly mapped onto their collective behaviour. With two available foraging patches having a 20% and 80% probability of food being present at any one time, the superorganism should match this probability by deploying 20% and 80% of foragers to the two sites. Regardless of the particular payoff σ_i available at each site, provided $\sum \frac{1}{\sigma_i} = 1$ this strategy is optimal over the long term, with the evolutionary time scale of millions of years favouring its selection. Fig. 1 shows a comparison of the Kelly strategy where probabilities of receiving a resource payoff are matched, regardless of the

payoff, with a strategy that allocates foragers proportional to the one-step expected return $p_i\sigma_i$, which does take the payoff into account.

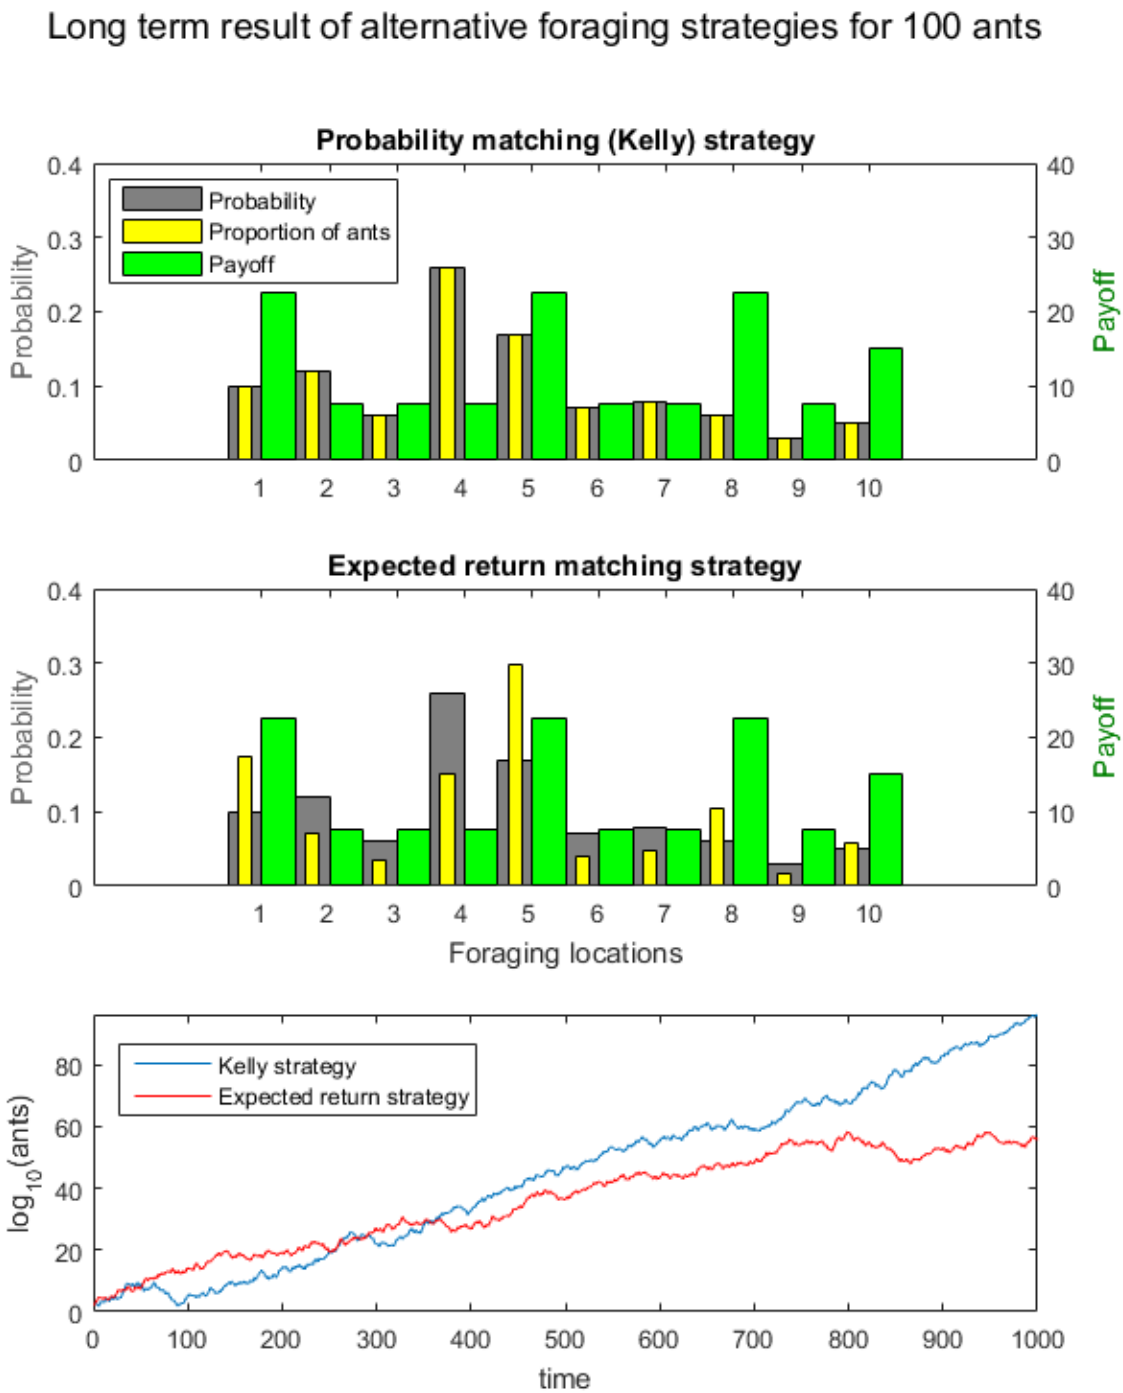

Fig. 1. A comparison of the Kelly strategy with an expected return matching strategy, over the long term (identical one-step payoffs for a 'win' in both cases). In the top pane (probability matching) the proportion of ants 'bet' (yellow bars) matches the probability of success (grey). The Kelly strategy increasingly outperforms any other strategy as time goes by.

Previous analysis of the behaviour of Bayesian foragers versus those modelled using the marginal value theorem indicated that, rather than abandoning a patch when instantaneous food intake rate equals foraging costs, a forager should consider the potential future value of a patch before moving on, even when the current return is poor (Olsson et al., 2006). The priority of resource reliability over immediate payoff in our model, when long-term biomass maximisation is the goal, is itself an interesting finding about superorganismal behaviour; but here we go further and specify models of movement to operationalize this strategy.

Certain methodologies designed to sample from probability distributions – Markov chain Monte Carlo (MCMC) methods – may be used as models of movement that also achieve a probability matching (Kelly) strategy. Exploring the environment and sampling from complex probability distributions can be understood as equivalent problems. MCMC methods aim to build a Markov chain of samples that draw from each region of probability space in correct proportion to its density. A well-mixed Markov chain is analogous to a probability matching strategy. Once the Markov chain has converged on its equilibrium distribution (the target probability distribution, or resource quality distribution in our ant model) it spends time in each location proportional to the quality or value (probability) of each point.

Ant trajectory data

We use our data (Hunt et al., 2016a) from previous work examining the movement of lone *Temnothorax* ants in an empty arena outside their colony's nest (Hunt et al., 2016b). *Temnothorax albipennis* ants have been used as a model social system for study in the laboratory, because information flow between the environment and colony members, and among colony members, can be rigorously studied. The ants typically have one queen and up to 200-400 workers (Franks et al., 2006); they inhabit fragile rock crevices and find and move into a new nest when its nest is damaged. With workers being only about 2mm long, the trajectories of individuals can be tracked easily. Behavioural state-based models have been developed that account for the flow of individuals

between states with differential equations (Sumpter and Pratt, 2003; Pratt et al., 2005), but these lack an account of the ants' movement processes.

In our data (Hunt et al., 2016a) there are two experimental regimes, one in which the foraging arena was entirely novel to exploring ants, and one in which previous traces of the ants' activities remained. We use the data from the former treatment, where each ant encounters a cleaned arena absent of any pheromones or cues from previous exploring ants. We restrict our analysis to the first minute of exploration, well before any of the ants have an opportunity to reach the boundary of the arena. Log-binned root mean square displacement is calculated and a linear regression made against log time. A gradient equal to a half indicates a standard diffusion process (Brownian motion) whereas greater than a half indicates superdiffusive movements. This approach to characterising ant search behaviour has been taken in e.g. Franks et al. (2010).

We run simulations of our models in MATLAB 2015b (supplementary material).

Each new model is introduced to explain an important additional aspect of the ants' empirical behaviour. A supplementary Materials and Methods section is at the end of the paper.

Results

We develop a sequence of models in turn, along with results from simulations.

Movement in two dimensions: Markov chain Monte Carlo methods (MCMC)

We can consider the problem of sampling from probability distributions of two continuous dimensions to obtain a model of movement. The development of MCMC methods from the 1950s onward, to become more efficient, might be considered to parallel the evolutionary history of animal foraging strategies. Some more details on the methods are found in the Materials and Methods section.

Basic model: Metropolis-Hastings

The first MCMC method to be developed was the Metropolis-Hastings (M-H) algorithm (Metropolis et al., 1953; Hastings, 1970), which is straightforward to implement and still commonly used today.

We are trying to sample from the target probability distribution (resource quality distribution) $P(x)$ which can be evaluated (observed) for any x, at least to within a multiplicative constant. This means we can evaluate a function $P^*(x)$ such that $P(x) = P^*(x)/Z$. There are two challenges that make it difficult to generate representative samples from $P(x)$. The first challenge is that we do not know the normalising constant $Z = \int d^N x P^*(x)$, and the second is that there is no straightforward way to draw samples from P without enumerating most or all of the possible states. Correct samples will tend to come from locations in x-space where $P(x)$ is large, but unless we evaluate $P(x)$ at all locations we cannot know these in advance (Mackay, 2003).

The M-H method makes use of a proposal density Q (which depends on the current state x) to create a new proposal state to potentially sample from. Q can be simply a uniform distribution: in a discretized environment these can be $x^{(t)} + [-1,0,1]$ with equal probability. After a given proposed movement is generated, the animal compares the resource quality at this new location with the resource quality at the previous location. If the new location is superior, it stays in its new location. In contrast, if the resource quality is worse, it randomly 'accepts' this new location, or 'rejects' this location based on a very simple formula based on the ratio of resource quality (if it is far worse, the animal very rarely fails to return, whereas if it is not much worse, it often accepts this mildly inferior location – see also Materials and Methods). What is important about this extremely simple algorithm is that, as long as the environment is ergodic (all locations can be potentially reached), given time, the exploring animal will visit each location eventually. Visits will be made with a probability proportional to its resource quality: it will execute an optimal Kelly exploration strategy. The problem here, however, is the time taken. Whilst the M-H method is widely used for sampling from high-dimensional problems, it has a major disadvantage in that it explores the

probability distribution by a random walk, and this can take many steps to move through the space, according to $\sqrt{T}\epsilon$ where T is the number of steps and ϵ is the step length. *Temnothorax* ants were found to be engaged in a superdiffusive search in an empty arena (Materials and Methods), and similarly MCMC methods also have been developed to explore probability space more efficiently.

Introducing momentum: Hamiltonian Monte Carlo (HMC)

Random walk behaviour is not ideal when trying to sample from probability distributions, since it is more time-consuming than necessary. One popular method for avoiding the random walk-like exploration of state space is Hybrid Monte Carlo (Duane et al., 1987), also known as Hamiltonian Monte Carlo (HMC). This simulates physical dynamics to preferentially explore regions of the state space that have higher probability.

Unlike the M-H model of movement, HMC makes use of local gradient information such that the walker (ant) tends to move in a direction of increasing probability. How *T. albipennis* may measure this is explored in the Discussion. For a system whose probability can be written in the form

$$P(x) = \frac{1}{Z}\exp\left[-E(x)\right]$$

the gradient of $E(x)$ can be evaluated and used to explore the probability space more efficiently. This is defined as:

$$E(x) = -\ln P(x)$$

Using this definition the local gradient $\nabla E(x)$ can be calculated numerically.

The Hamiltonian is defined as $H(\boldsymbol{x}, \boldsymbol{p}) = E(\boldsymbol{x}) + K(\boldsymbol{p})$, where $K(\boldsymbol{p})$ is a 'kinetic energy' which can be defined as:

$$K(\boldsymbol{p}) = \boldsymbol{p}^{\mathrm{T}}\boldsymbol{p}/2$$

In HMC, this momentum variable $\boldsymbol{p}$ augments the state space $\boldsymbol{x}$ and there is an alternation between two types of proposal. The first proposal randomises the momentum variable, with $\boldsymbol{x}$ unchanged,

and the second proposal changes both $\boldsymbol{x}$ and $\boldsymbol{p}$ using simulated Hamiltonian dynamics. The two proposals are used to create samples from the joint density

$$P_H(\boldsymbol{x},\boldsymbol{p}) = \frac{1}{Z_H}\exp[-H(\boldsymbol{x},\boldsymbol{p})] = \frac{1}{Z_H}\exp[-E(\boldsymbol{x})]\exp\,[-K(\boldsymbol{p})]$$

As shown, this is separable, so the marginal distribution of $\boldsymbol{x}$ is the desired distribution $\exp[-E(\boldsymbol{x})]\,/Z$, and the momentum variables can be discarded and a sequence of samples $\{\boldsymbol{x}^{(t)}\}$ is obtained that asymptotically comes from $P(\boldsymbol{x})$ (Mackay, 2003).

We set the variable number of leapfrog steps (see supplementary Materials and Methods and Brooks et al., 2011)) to $L = 10$; after following the Hamiltonian dynamics for this number of steps a new momentum is randomly drawn and a new period of movement begins. This behaviour of moving intermittently in between updating the walker (ant) behaviour captures the behaviour observed in real ants (Hunt et al., 2016b). We set the leapfrog step length $\varepsilon = 0.3$ (supplementary Materials and Methods).

For $N = 18$ simulated HMC 'ants' sampling from the same sparse probability distribution as before, for 600 iterations, the r.m.s. displacement was again found and its log was regressed on log time. The gradient was found to be 0.567, confidence interval (0.528 0.606), which is significantly greater than 0.5, so in this respect it is more similar to the superdiffusive search found in real ants (Franks et al., 2010).

We can also examine the correlation of velocities between successive movement periods. Since momentum $\boldsymbol{p} = m\boldsymbol{v}$ is a vector in two-dimensional space, we can set $m = 1$ and find a magnitude for the momentum to determine the 'speed' of each movement (over the course of $L = 10$ leapfrog steps). In previous research on ant movements (Hunt et al., 2016b) the correlation between successive average event speeds in the cleaning treatment was found to be 0.407 ± 0.039 (95% CI). As expected for the HMC model, because the momentum is discarded and replaced with a new random momentum after each movement, the correlation of successive event speeds is equal

to zero in this model. We can make the HMC model more 'ant-like' – and potentially more efficient – by only partially refreshing this momentum variable after the end of a movement period.

Increasing correlations between steps: partial momentum refreshment (PMR)

HMC with one leapfrog step is referred to as Langevin Monte Carlo after the Langevin equation in physics (e.g. Kennedy, 1990) and was first proposed by Rossky et al. (1978). However, these methods do not require $L = 1$, so we use $L = 10$ to enhance comparability with the previous HMC model.

The momentum at the end of each movement can be updated according to the equation $p' = \alpha p + (1 - \alpha^2)^{1/2} n$, where p is the existing momentum, p' the new momentum, α is a constant in the interval $[-1,1]$ and n is a standard normal random vector. With α less than one p' is similar to p but with repeated iterations it becomes almost independent of the initial value. This technique of partial momentum refreshment (PMR) was introduced by Horowitz (1991). Such models are well-described in Brooks et al. (2011). Setting α equal to 0.65 (for example) and simulating with $N = 18$ results in speed correlations equal to 0.387 ± 0.012 (95% CI) which overlaps with the confidence interval for the real ant data.

The PMR method can be compared to an ant moving with a certain momentum (direction and speed) and then intermittently updating this momentum in response to its changing position in the physical and social environment. The momentum changes as per the HMC method along a single trajectory, according to its subjective perception of foraging quality and potentially influenced by the pheromonal environment. If at the end of the trajectory it does not find itself in a more attractive region than before, it returns to its previous position: and with the correct model parameters (stepsize and number of leapfrog steps) this should be a relatively infrequent occurrence (see methodological discussion in Brooks et al. (2011)). Its starting momentum in a particular direction is maintained to some degree but with some randomness mixed in – and so its previous tendency to move toward regions of high probability (quality) is not discarded as in HMC but used to make more

informed choices about which direction to move in next. This is because foraging patches are likely to show some spatial correlation in their quality, with high quality regions more likely to neighbour other high quality regions (Klaassen et al., 2006; Van Gils, 2010). Previous empirical research (Hunt et al., 2016b) found evidence that ant movements are predetermined to some degree in respect of their duration. This implies that periods of movement are followed by a more considered sensory update and decision about where to move to next. A series of smaller movements (like 10 leapfrog steps) followed by a larger momentum update, as in the PMR model, would seem to correspond well with this intermittent movement behaviour.

Measuring the performance of MCMC foraging models

The performance of the three MCMC models developed here can be measured in the following way. As discussed, the foraging ants should pursue a probability matching strategy, whereby they allocate their numbers across the environment in proportion to the probability that it will return (any) payoff. This will maximise the long-term rate of growth of the colony, or from a gene-level perspective, maximise the number of copies of the ants' genes in the environment. Matching the probability distribution of resources in the environment can be understood as minimising the distance between it and the distribution of resource gatherers. In the domain of information theory, the difference between two probability distributions is measured using the cross-entropy

$$H(p,q) = H(p) + D_{KL}(p||q)$$

Where $H(p) = -\sum_i p(i)\log\ p(i)$ is the entropy of p and $D_{KL}(p||q)$ is the Kullback-Leibler (K-L) divergence of q from p (also known as the relative entropy of p with respect to q). This is defined as

$$D_{KL}(p||q) = \sum_i p(i)\log\frac{p(i)}{q(i)}$$

If we take p to be a fixed reference distribution (the probability of collecting resources in the environment), cross entropy and K-L divergence are identical up to an additive constant, $H(p)$, and is minimised when $q = p$, where the K-L divergence is equal to zero. Cross-entropy minimisation is a

common approach in optimization problems in engineering, and in the present case can be used to represent the task the ant foragers are trying to perform: match their distribution q with the distribution p of resources in the environment. The magnitude and rate of reduction of the cross-entropy is therefore used to compare the effectiveness of the MCMC models (M-H, HMC, PMR) presented here. However, as noted later, for dynamic environments (where the distribution of resource probabilities p is not fixed), K-L divergence is the suitable cost function to minimise.

Example simulations for the three models sample from a target distribution p with three simulated resources patches. This example distribution is generated by combining a gamma-distributed background noise (shape parameter=0.2, scale parameter=1) on a 100×100 grid given a Gaussian blur ($\sigma = 3$, filter size 100×100), in equal 50% proportion with three patches of resources, which are single points of increasing relative magnitude of 1, 2, and 3 that have been given a Gaussian blur ($\sigma = 10$, filter size 100×100). The distribution p is thus also on a 100×100 grid. The simulations are run for 50,000 time steps, a reasonable period of time to explore this space of 10,000 points. Figure 1 shows the M-H model, which converges rather slowly on the target environment p. Figure 2 shows the performance of the HMC and PMR models, which show an improvement in the convergence rate because they avoid the random walk behaviour of the M-H model.

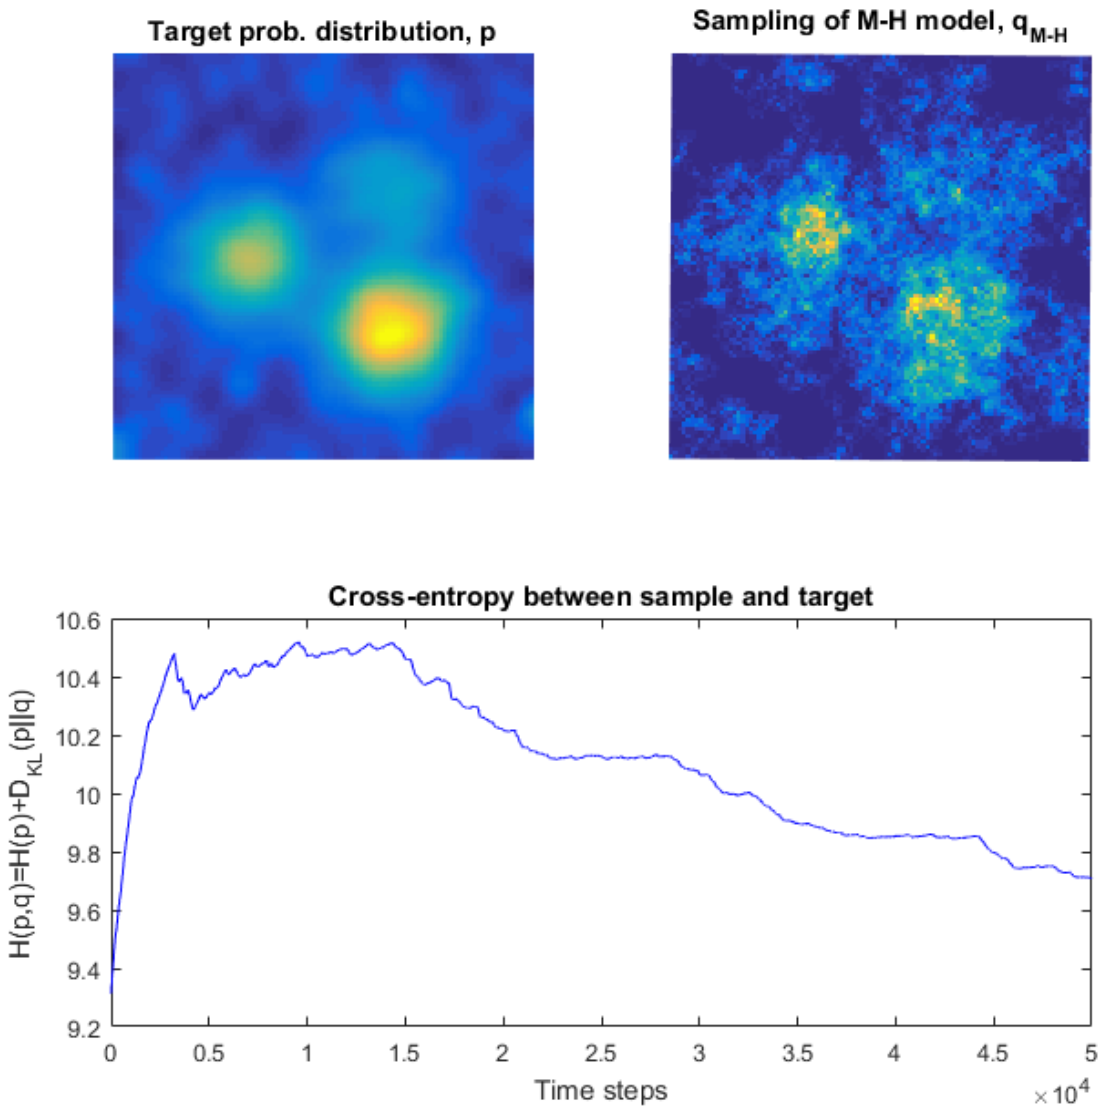

Figure 2. Performance of the M-H model as it generates a sample distribution q that approximates the target distribution p, the location of resources in the environment.

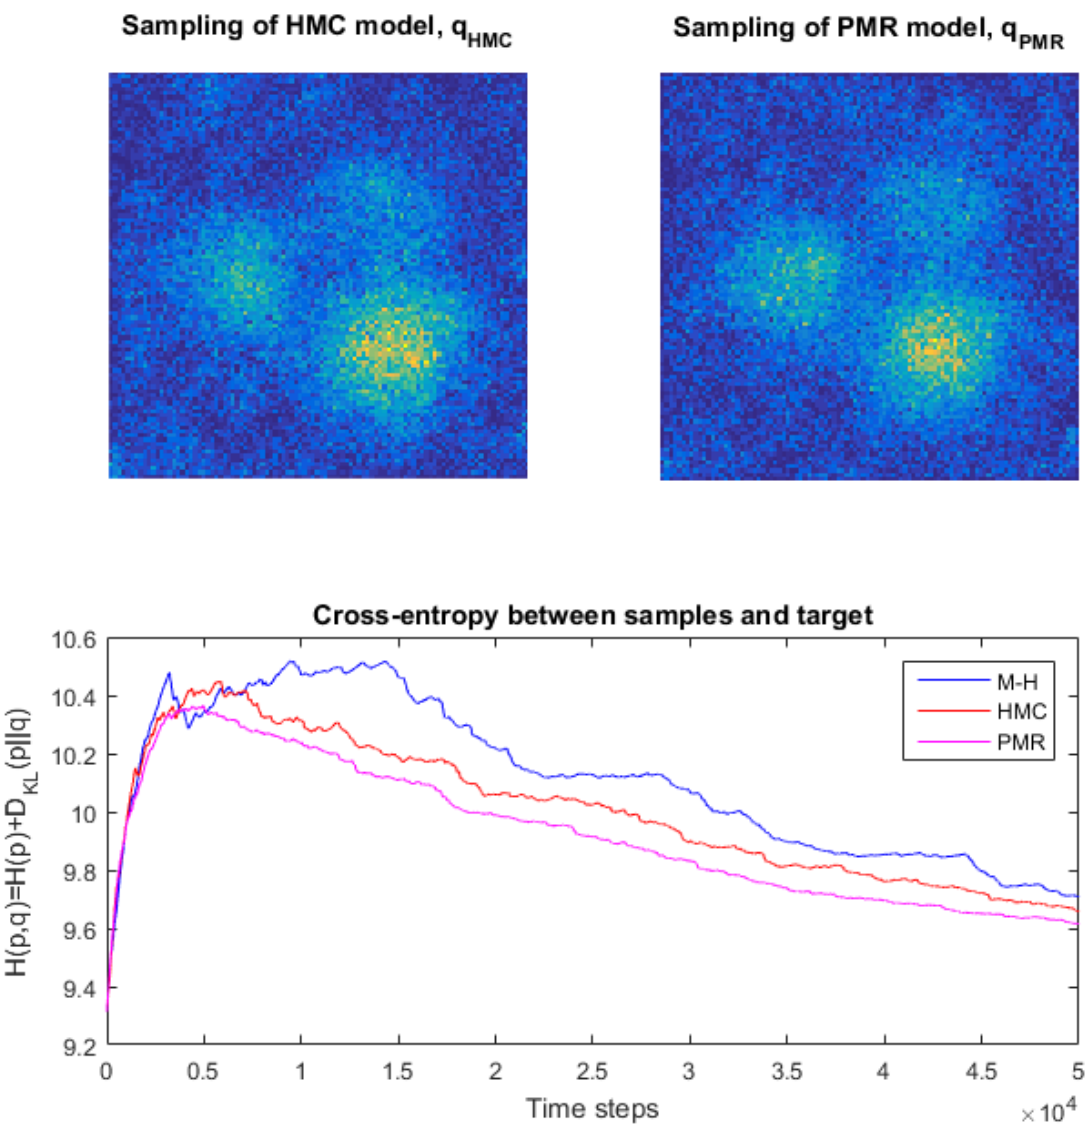

Figure 3. Performance of the HMC and PMR models, compared to that of M-H. In general, PMR and HMC outperform M-H because random walk type exploration of probability space is avoided, by following local gradient information and making larger steps. Their performance depends on the nature of the target distribution and choosing suitable values for step length ε and number of steps L.

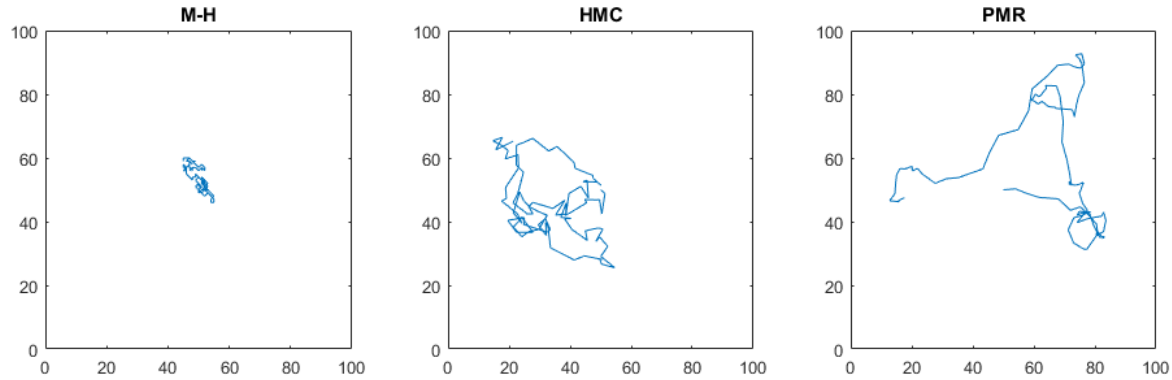

Figure 4. Comparison of example trajectories for the 3 MCMC models: they become increasingly superdiffusive, and the HMC and PMR models preferentially follow the local quality gradient.

Figure 4 shows example trajectories produced through 100 timesteps of the 3 models, sampling from the same probability distribution as shown in Figure 2. The random walk behaviour of the M-H model is evident, while the greater tendency to make longer steps in one direction is evident in the PMR model in comparison to the HMC model.

Optimal foraging and Lévy flights

We have presented a new class of foraging models based on MCMC methods, which operationalise movement for a Kelly strategy (probability matching) in two-dimensional space. There is an extensive theoretical and empirical literature examining the distribution of step lengths for foraging animals that considers the hypothesis that a Lévy distribution is optimal (Bartumeus, 2007; Benhamou, 2007; Humphries et al., 2010; Viswanathan et al., 2008). Lévy flights are a particular form of superdiffusive random walk where the distribution of move step-lengths fits an inverse power law such that the probability of a move of length l is distributed like $(l) \approx l^{-\mu}$, where $1 < \mu \leq 3$.

We use the method of (Humphries et al., 2013) to identify individual movement steps in two dimensional data, treating monotonic movements in a certain direction in one dimension (i.e., x or y) as a step. We estimate the exponent using maximum likelihood estimation (White et al., 2008). The distribution of ranked step length sizes in both real and simulated data is shown in Figure 5.

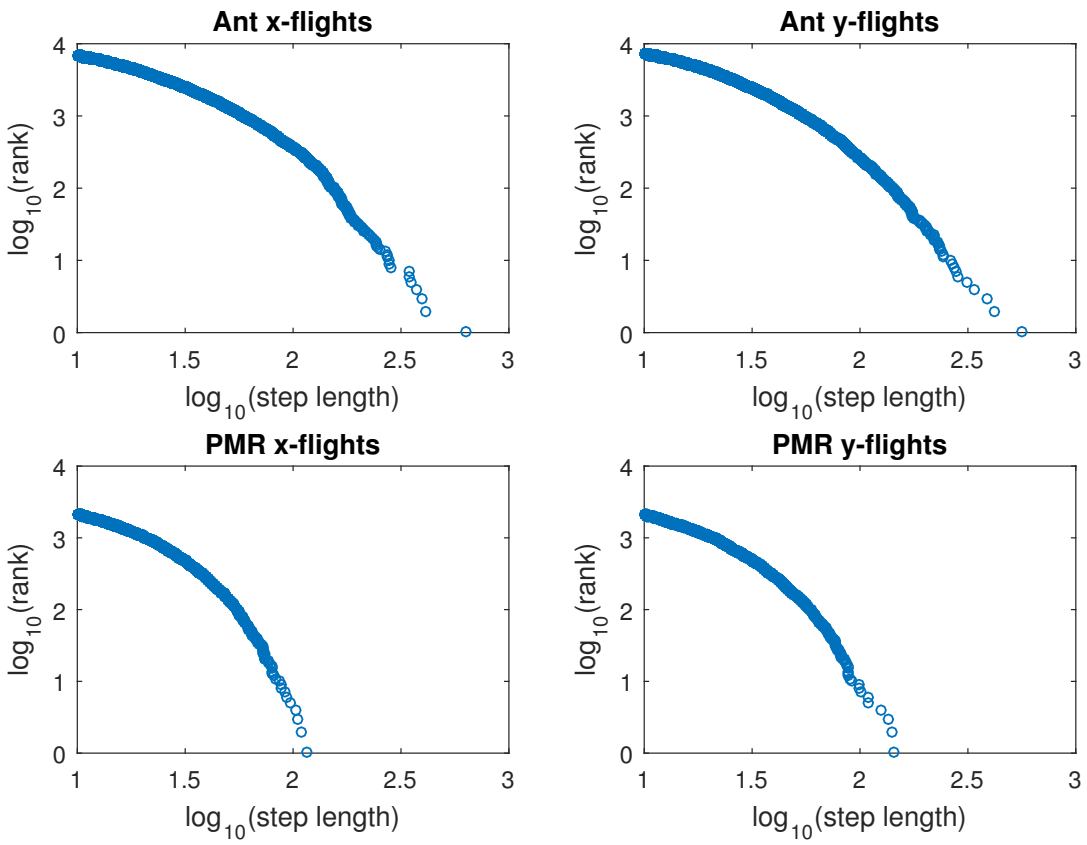

Figure 5. The (apparently) power-law distributed step lengths for both real ants and simulated PMR walkers.

There are similar exponents estimated for both the real data in an empty arena and PMR trajectories sampling from a sparse probability distribution (randomly generated from a gamma distribution with shape parameter 0.2 and scale parameter 1 (Table 1). The exponent μ in both cases is in the right region for a Lévy flight $1 < \mu \leq 3$. This would seem to be evidence for a Lévy strategy in the ants, but we suggest an alternative in the next section of this paper.

Dimension of steps	Maximum likelihood estimate of exponent, truncated Pareto distribution	
	Empirical data	PMR data
x	2.41	2.22
y	2.55	2.33

Table 1. Power-law exponents in both empirical and PMR simulated trajectories potentially indicate a Lévy walk

368 **Discussion**

369 ***MCMC models and existing movement paradigms***

370 Our framework is an important step in integrating key features of the movement paradigm such as
371 articulated by Nathan et al. (2008). It incorporates elements of randomness, producing correlated
372 random walks in certain environments; it quantifies optimality in respect of foraging strategies via
373 cross-entropy (Kullback-Leibler divergence); it includes an important aspect of common animal
374 behaviour, namely intermittent movement (Kramer and McLaughlin, 2001), and specifically for the
375 ants' neural and/or physiological behaviour, apparent motor planning (Hunt et al., 2016a); and it
376 makes explicit the information used by the animal step-by-step. Finally, and crucially, it explicates
377 cognition at the emergent group level, because individual movement is at the service of a group-
378 level Kelly strategy. One component of Nathan et al.'s framework is the internal state of the
379 organism. This is not included in the models here, though as discussed in the first part of this series
380 (Hunt et al. 2018b) state-dependent behaviours such as tandem running (Franks and Richardson,
381 2006) could be included by analogy with particle filtering (e.g. Gordon et al., 1993), for instance. Our
382 use of the Markov assumption (movement being memoryless, depending only on the current
383 position) is justifiable with respect both to the worker ant's individual cognitive capacity, and its
384 single-minded focus on serving the colony through discovering and exploiting resources, rather than
385 seeking mates for itself, for instance. Its motion capacity is linked to the specification of a partial
386 momentum refreshment model; while we specify the navigation capacity in its ability to measure
387 the quality gradient, which is also an externally determined factor.

388 ***The mechanism of gradient sensing***

389 We may consider further the ability of ants to use local gradient information, as in the HMC and
390 PMR models, with respect to the ants' sensory system. *T. albipennis* is well-known for relying heavily
391 on visual information in movement (McLeman et al., 2002; Pratt et al., 2001) and in common with
392 most (or perhaps all) ants on olfactory information. It may be that the intermittent movement

examined in Hunt, Baddeley et al (Hunt et al., 2016b) is associated with limitations in the quality of sensory information when moving (Kramer and McLaughlin, 2001). We suggest that *T. albipennis* workers have relatively good eyesight for their small size, having around 80 ommatidia in each compound eye (Hunt et al., 2018a) and may be assumed conservatively to have an angle of acuity of 7 degrees (Pratt et al., 2001). Therefore, movement would seem unlikely to make much difference to how well they can see. Since our model highlights the importance of gradient following, this may be more difficult to measure for the olfactory system during movement. Indeed, in Hunt et al. (2016b) we suggest that social information from pheromones or other cues is only fully attended to during periods of stopping because of motor planning, with the duration of movements being predetermined by some endogenous neural and/or physiological mechanism. This may be therefore a mechanistic reason for the stepwise movement in the PMR model, in addition to its informational efficiency which is its evolutionary origin. Even more sophisticated MCMC models that rely on the second derivative of the probability distribution, such as the Riemann Manifold Langevin method (Girolami and Calderhead, 2011), may be relevant, because this property (the rate of change of the gradient) may be only measured with adequate accuracy when the ant is at rest.

Lévy step length distributions indicate a world with little gradient information

In a 'flat' quality landscape, or sparse world, our model generates Lévy-like behaviour as seen for instance in Reynolds et al., 2013. This remains an adaptive response, but it is not a true Lévy distribution, because there is a finite variance. The empirical distribution of momentums provides insight into the length-scales on which the world remains smooth. Much interest has been generated by Lévy flight based foraging models which theoretically optimise mean resource collection for certain random worlds; and this would seem to be evidence for just such a strategy in *T. albipennis* ants. Scale-free reorientation mechanisms have indeed been suggested as a response to uncertainty in invertebrates (Bartumeus and Levin, 2008). Yet the generation of a Lévy -like distribution from our gradient-following model suggests that such observations may not really be

418 scale-free. Even within the same species, Lévy models only fit the empirical observations well under
419 certain scenarios (Benhamou, 2007).

420 ***Measuring efficiency, selecting for unpredictability***

421 The rate of resource collection can be straightforwardly calculated by finding the cross-entropy
422 (Kullback-Leibler divergence) between the spatial distribution of resources, and the realised foraging
423 distribution resulting from the foraging strategy. The distribution of resources is seen from a 'genes-
424 eye' view of the animal or superorganism, with respect to maximising the long-term biomass or
425 number of copies of genes in the environment: this focuses on a location's probability of yielding
426 resources, or reliability, as opposed to the one-off payoff. The foraging strategy is that chosen by
427 natural selection. Minimising the cross-entropy (Kullback-Leibler divergence) is achieved by obeying
428 the matching law: foraging proportional to the probability finding the best resources at each
429 available location. This strategy is especially suited to a superorganism like an ant colony, because it
430 can forage in multiple locations simultaneously by allocating worker ants in numbers proportional to
431 the location's reliability, through self-organisation (Camazine et al., 2001).

432 There has been some intimation before that MCMC could be a model for biological
433 processes (Neal, 1993), with some query about whether the requisite randomness is possible in
434 organisms. We think that not only is spontaneous (i.e. non-deterministic, 'random') behaviour
435 present, it is necessary for survival in terms of being unpredictable around predators, prey or
436 competitors (Brembs, 2010), or for finding food using a 'strategy of errors' (Deneubourg et al.,
437 1983). Indeed, the Bayesian framework developed here allows predictions to be developed
438 regarding the optimal amount of 'randomness' in behaviour at both the level of the individual and
439 colony (in the partial momentum refreshment model, this is adjusted with the α parameter) that can
440 be tested in future empirical research. Further predictions arise from the momentum reversal step in
441 MCMC PMR (Brooks et al., 2011), which may be compared to observations of U-turning in ants
442 (Beckers et al., 1992; Hart and Jackson, 2006). Recent literature (Hoffman and Gelman, 2014) has

developed methods to adjust the path length dynamically, while removing the need to have a parameter L for the number of leapfrog steps. Observing how ants (and other organisms) adjust their step lengths according to different resource distributions will be instructive of their underlying movement model.

Selection for collective foraging phenotypes

We began by noting that life has from the beginning been selected for its information processing capabilities, and the major evolutionary transitions represent successive leaps forward in its efficiency. The Bayesian framework developed here permits the ready simulation and quantification of alternative strategies for organisms existing at the various points on the scale of biological organisation, and generates predictions for the evolution of new individual and collective movement algorithms in environments of whatever specified resource distribution. Mathematically, just as Fisher's fundamental law of genetics noted that positive variances imply either increase or a plateau in fitness (Orr, 2009), the ultimate story of life's explosion of complexity may be summarised in our framework in the simple statement that for a movement strategy to be favoured under natural selection:

$$D_{KL}(p||q_{new}) < D_{KL}(p||q_{old}) \quad (1)$$

i.e. the Kullback-Leibler divergence (measuring the similarity of two distributions) between a potential (genetically accessible) collective movement strategy that results in the equilibrium distribution of foragers q_{new}, and the organism's resource environment p, has to be lower than under the current strategy found in the population that results in distribution q_{old}. This reduction may indeed be achieved by more sophisticated, coordinated, collective behaviour, notwithstanding higher individual energetic cost. This information processing focused expression suggests that emergent, collective behavioural phenotypes may be causally prior over the preservation of an individual organism's genotype. Under competition, larger improvements in efficiency associated with major organisational innovations may be favoured; and this pressure for enhanced information

processing, seen in the progressively advancing informational division of labour between biological components, is most manifest in group living organisms such as cooperatively foraging insects. Their strategies may be instructive for biomimetic improvements to existing MCMC methods.

Conclusion

We described the foraging problem as a repeated multiplicative game, where an ant colony has to place 'bets' on which foraging patches to visit, with an ultimate payoff of more colonies or copies of their genes being created. Ants are very successful in terms of their terrestrial biomass (Schultz, 2000), and so it would seem likely that they are following a highly evolved strategy. We suggest the theoretical optimum is a 'Kelly' or probability matching strategy, which maximises the long term 'wealth' or biomass of the colony rather than the resource collection of single ants. By mapping the foraging problem to a set of methods designed to effectively sample from probability distributions, we present models of ant movement that achieve this matching behaviour. These MCMC-based models thus provide spatially explicit predictions for movement that describe and explain how colonies optimally explore and exploit their environment for food resources. We also show how Lévy -like step length distributions can be generated by following a local gradient that is uninformative, suggesting that contrary to being an evolved strategy, Lévy flight behaviour may be a spontaneous phenomenon. While we do not include interactions between ants in the model, past theoretical analysis of information use in collective foraging suggests that totally independent foraging is actually optimal for a broad range of model parameters when the environment is dynamic. This is because information about food patches may not be worth waiting for if they are short lived (Dechaume-Moncharmont et al., 2005).

Understanding the logic of information flow at the level of the gene and the cell has been identified as a priority (Nurse, 2008). However, given that no level of organisation is causally privileged in biology (Noble, 2011), explicating this at the organismal and super-organismal level

should also advance our understanding. Our Bayesian framework operationalizes earlier proposed frameworks (such as that of Nathan et al., 2008) for movement in a coherent and logical way, accounting for the uncertainty in both the individual ant and the colony's cognition in relation to the foraging problem. It also allows quantification of the system's emergent information processing capabilities and hypothesis generation for different organisms moving in different environments. Our MCMC models can be used as a foundation upon which further organismal and ecological complexity can be explained in future research.

Author contributions. RJB proposed the Bayesian framework and its specific theoretical concepts; NRF advised on social insect biology; ERH drafted the paper, ran the simulations, and suggested eqn. (1); all authors contributed to the present draft.

Funding statement. E.R.H. thanks the UK Engineering and Physical Sciences Research Council (grant no. EP/I013717/1 to the Bristol Centre for Complexity Sciences).

Ethics statement. The ant *Temnothorax albipennis* is not subject to any licencing regime for use in experiments. The ants were humanely cared for throughout the experiment.

Competing interests. We have no competing interests.

Supplementary Materials and Methods

Further details on Kelly betting

For full treatment see chapter 6 of Cover and Thomas (2006). A probability matching strategy is also optimal with 'superfair' odds when $\sum \frac{1}{\sigma_i} < 1$. In both of these cases, the full capital or proportion of foragers should be allocated across foraging sites to maximise the long-term wealth, although the Kelly strategy may also be followed by betting a fixed proportion α of full bets. This so-called α-Kelly strategy has a much slower rate of growth in wealth (MacLean et al., 2010), however, short-term

volatility in wealth (or biomass) is reduced by a factor α^2 (Rising et al., 2012). Animals are often risk averse to variable rewards (Kacelnik and Bateson, 1996) and this may be seen at the level of the colony; in practice only a proportion of ants forage, with some staying purposefully inactive for significant periods of time (Charbonneau and Dornhaus, 2015). In a colony of, say, 200 workers, 50 might forage, which would correspond to $\alpha = 0.25$; but for the purpose of this paper we do not consider this further.

When $\sum \frac{1}{\sigma_i} > 1$ subfair odds are said to be available, and it is better to hold back some proportion of the ant foragers (or a gambler's wealth) and try to bet on outcomes offering the most favourable expected return $p_i \sigma_i$. In this case the optimal betting amount is solved using a 'waterfilling' method and is dealt with in Erkip (1996).

MCMC methods – resources

Good introductions to Markov chain Monte Carlo (MCMC) methods can be found in Neal (1993), Mackay (2003), and Brooks et al. (2011).

Metropolis-Hastings – further details and comparison with ant data

The M-H method makes use of a proposal density Q (which depends on the current state x) to create a new proposal state to potentially sample from. Q can be simply a uniform distribution: in a discretized environment these can be $x^{(t)} + [-1,0,1]$ with equal probability. To decide whether to accept the new state as the next step in the sample, we compute the quantity

$$a = \frac{P^*(x')}{P^*(x^{(t)})} \frac{Q(x^{(t)}; x')}{Q(x'; x^{(t)})}$$

If $a \geq 1$ then the new state is accepted; otherwise it is accepted with probability a. If the step is accepted, we set $x^{(t+1)} = x'$; if it is rejected then $x^{(t+1)} = x^{(t)}$ (the current state is taken forward to the next step). The M-H method is widely used for sampling from high-dimensional problems, but it has a major disadvantage that it explores the probability distribution by a random walk, which can

take many steps to move through the space, according to $\sqrt{T}\epsilon$ where T is the number of steps and ϵ is the step length.

The Metropolis-Hastings algorithm was used to sample from a sparse probability distribution, randomly generated from a gamma distribution with shape parameter 0.2 and scale parameter 1, with a Gaussian blur applied to make it continuous. Using $N = 18$ simulated walkers exploring for 600 iterations, the exponent was found to be 0.487, 95% confidence interval (0.444 0.529), i.e. not significantly different to 0.5 and so engaging in a standard diffusive search, as expected. We can use the M-H walkers as a model of exploring ants, and so to assess its validity we can examine trajectory data of real ants exploring a uniform empty space (Hunt et al., 2016a) and compare it with simulations from this model, to see whether ants are engaged in a standard diffusive (Brownian) exploration of space. The root mean square (r.m.s.) displacement for $N = 18$ ants was found for the first 60 seconds of their exploration of an unfamiliar arena (600 trajectory points with a sampling period of 0.1s, a time short enough that they would not reach the edge of the arena), and its log was regressed on log time. The gradient was found to be 0.59 with a 95% confidence interval of (0.55,0.63), which is greater than 0.5 and so the ants are engaged in superdiffusive search.

Leapfrog method

In two of the MCMC methods presented, HMC and PMR, there are two key parameters ε and L that must be chosen suitably to generate a sample efficiently. See Brooks et al. (2011) for a full treatment of these issues.

Although the dynamics of a system in Hamiltonian mechanics is described using differential equations for continuous time and space, for implementation on a computer the equations must be approximated by discretizing time using a small step length ε. We may assume a Hamiltonian of the

form $H(q,p) = U(q) + K(p)$ and a kinetic energy $K(p) = p^T M^{-1} p/2$, and a diagonal M with elements $m_1, \dots, m_d$ so that $K(p) = \sum_{i=1}^{d} \frac{p_i^2}{2m_i}$. Then the leapfrog method proceeds as follows:

$$
\begin{aligned}
p_i\,(t+\varepsilon/2) &= p_i(t) - (\varepsilon/2)\,\frac{\partial U}{\partial q_i}(q(t)),\\
q_i(t+\varepsilon) &= q_i(t) + \varepsilon\,\frac{p_i(t+\varepsilon/2)}{m_i},\\
p_i(t+\varepsilon) &= p_i\,(t+\varepsilon/2) - (\varepsilon/2)\,\frac{\partial U}{\partial q_i}(q(t+\varepsilon)).
\end{aligned}
$$

We start with a half step for the momentum variables, a full step follows for the position variables using those new momentum variables, and finally we do another half step for the momentum variables using the new values for the position variables. This process is iterated a certain number of times L, and preserves volume exactly. It is also reversible by negating p and applying the same number of steps L, then negating p again. It gains its name from the use of interleaved time points such that p_i and q_i 'leapfrog' over each other.

The leapfrog step length was set to be $\varepsilon = 0.3$ in the simulations presented in the main text. Selecting a suitable step length is vital, because if it is too large there will be a very low acceptance rate for the states proposed by the simulated trajectories. If it is too small it will waste computation time, as the trajectories will make unnecessarily short progress through the space, or it can lead to slow exploration by a random walk through the space if the trajectory length εL is too small, which defeats the object of the method. In the first instance reasonably effective ε and L can be found through trial and error.

References

Bartumeus F, 2007. Levy processes in animal movement: An evolutionary hypothesis. Fractals 15:151-162. doi: 10.1142/s0218348x07003460.

Bartumeus F, Levin SA, 2008. Fractal reorientation clocks: Linking animal behavior to statistical patterns of search. Proceedings of the National Academy of Sciences of the United States of America 105:19072-19077. doi: 10.1073/pnas.0801926105.

Beckers R, Deneubourg JL, Goss S, 1992. Trails and U-turns in the selection of a path by the ant *Lasius niger*. Journal of Theoretical Biology 159:397-415. doi: 10.1016/s0022-5193(05)80686-1.

Benhamou S, 2007. How many animals really do the Levy walk? Ecology 88:1962-1969. doi: 10.1890/06-1769.1.

Brembs B, 2010. Towards a scientific concept of free will as a biological trait: spontaneous actions and decision-making in invertebrates. Proceedings of the Royal Society B: Biological Sciences 278:930-939. doi: 10.1098/rspb.2010.2325.

Brooks S, Gelman A, Jones G, Meng XLE, 2011. Handbook of Markov Chain Monte Carlo. CRC press.

Camazine S, Deneubourg JL, Franks NR, Sneyd J, Bonabeau E, Theraulaz G, 2001. Self-organization in biological systems. Princeton University Press.

Charbonneau D, Dornhaus A, 2015. Workers 'specialized' on inactivity: Behavioral consistency of inactive workers and their role in task allocation. Behavioral Ecology and Sociobiology 69:1459-1472. doi: 10.1007/s00265-015-1958-1.

Couzin ID, 2009. Collective cognition in animal groups. Trends in Cognitive Sciences 13:36-43. doi: 10.1016/j.tics.2008.10.002.

Cover TM, Thomas JA, 2006. Elements of information theory. New Jersey: John Wiley & Sons.

Dangerfield JM, McCarthy TS, Ellery WN, 1998. The mound-building termite *Macrotermes michaelseni* as an ecosystem engineer. Journal of Tropical Ecology 14:507-520. doi: 10.1017/s0266467498000364.

Dechaume-Moncharmont F-X, Dornhaus A, Houston AI, McNamara JM, Collins EJ, Franks NR, 2005. The hidden cost of information in collective foraging. Proceedings of the Royal Society B: Biological Sciences 272:1689-1695. doi: 10.1098/rspb.2005.3137.

Dempster ER, 1955. Maintenance of genetic heterogeneity. Cold Spring Harb. Symp. Quant. Biol. 20:25-32.

Deneubourg JL, Pasteels JM, Verhaeghe JC, 1983. Probabilistic behaviour in ants: A strategy of errors? Journal of Theoretical Biology 105:259-271. doi: 10.1016/s0022-5193(83)80007-1.

Detrain C, Deneubourg JL, Pasteels JM, 1999. Information processing in social insects. Springer Science & Business Media.

Dorigo, M. and L.M. Gambardella, Ant colony system: a cooperative learning approach to the traveling salesman problem. Evolutionary Computation, IEEE Transactions on, 1997. **1**(1): p. 53-66. doi: 10.1109/4235.585892.

Duane S, Kennedy AD, Pendleton BJ, Roweth D, 1987. Hybrid Monte Carlo. Physics Letters B 195:216-222. doi: 10.1016/0370-2693(87)91197-x.

Erkip E, 1996. The efficiency of information in investment: Stanford.

Franks NR, 1989. Army ants: a collective intelligence. American Scientist 77:138-145.

Franks NR, Dornhaus A, Best CS, Jones EL, 2006. Decision making by small and large house-hunting ant colonies: one size fits all. Animal Behaviour 72:611-616. doi: 10.1016/j.anbehav.2005.11.019.

Franks NR, Richardson TO, Keir S, Inge SJ, Bartumeus F, Sendova-Franks AB, 2010. Ant search strategies after interrupted tandem runs. Journal of Experimental Biology 213:1697-1708. doi: 10.1242/jeb.031880.

Giraldeau L-A, Caraco T, 2000. Social foraging theory: Princeton University Press.

Girolami M, Calderhead B, 2011. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. Journal of the Royal Statistical Society Series B-Statistical Methodology 73:123-214. doi: 10.1111/j.1467-9868.2010.00765.x.

Giuggioli L, Bartumeus F, 2010. Animal movement, search strategies and behavioural ecology: a cross-disciplinary way forward. Journal of Animal Ecology 79:906-909. doi: 10.1111/j.1365-2656.2010.01682.x.

Gordon NJ, Salmond DJ, Smith AF. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In: IEE Proceedings F - radar and signal processing 1993 Apr (Vol. 140, No. 2, pp. 107-113) doi: 10.1049/ip-f-2.1993.0015

Hart A, Jackson DE, 2006. U-turns on ant pheromone trails. Current Biology 16:R42-R43. doi: 10.1016/j.cub.2006.01.015.

Hastings WK, 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika 57:97. doi: 10.2307/2334940.

Hoffman MD, Gelman A, 2014. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. Journal of Machine Learning Research 15:1593-1623.

Hölldobler B, Wilson EO, 2009. The superorganism: the beauty, elegance, and strangeness of insect societies: WW Norton & Company.

Horowitz AM, 1991. A generalized guided Monte Carlo algorithm. Physics Letters B 268:247-252. doi: 10.1016/0370-2693(91)90812-5.

Humphries NE, Queiroz N, Dyer JRM, Pade NG, Musyl MK, Schaefer KM, Fuller DW, Brunnschweiler JM, Doyle TK, Houghton JDR, Hays GC, Jones CS, Noble LR, Wearmouth VJ, Southall EJ, Sims DW, 2010. Environmental context explains Levy and Brownian movement patterns of marine predators. Nature 465:1066-1069. doi: 10.1038/nature09116.

Humphries NE, Weimerskirch H, Sims DW, 2013. A new approach for objective identification of turns and steps in organism movement data relevant to random walk modelling. Methods in Ecology and Evolution 4:930-938. doi: 10.1111/2041-210x.12096.

Hunt ER, Baddeley R, Worley A, Sendova-Franks A, Franks NR, 2016a. Data from: Ants determine their next move at rest: motor planning and causality in complex systems. Royal Society Open Science 3:150534. doi: 10.5061/dryad.jk53j

Hunt ER, Baddeley RJ, Worley A, Sendova-Franks AB, Franks NR, 2016b. Ants determine their next move at rest: motor planning and causality in complex systems. Royal Society Open Science 3:150534. doi: 10.1098/rsos.150534.

Hunt ER, Dornan C, Sendova-Franks, AB, Franks NR, 2018a. Asymmetric ommatidia count and behavioural lateralization in the ant *Temnothorax albipennis*. Scientific Reports 8:5825. doi: 10.1038/s41598-018-23652-4.

Hunt ER, Franks NR, Baddeley RJ, 2018b. The Bayesian Superorganism I: collective probability estimation. bioRxiv, 468942. doi: 10.1101/468942

Kacelnik A, Bateson M, 1996. Risky theories - The effects of variance on foraging decisions. American Zoologist 36:402-434.

Karaboga, D. and B. Basturk, A powerful and efficient algorithm for numerical function optimization: artificial bee colony (ABC) algorithm. Journal of Global Optimization, 2007. **39**(3): p. 459-471.

Kelly JL, 1956. A new interpretation of information rate. Bell System Technical Journal 35:917-926.

Kennedy AD, 1990. The theory of hybrid stochastic algorithms. Probabilistic Methods in Quantum Field Theory and Quantum Gravity 224:209-223.

Klaassen RHG, Nolet BA, van Gils JA, Bauer S, 2006. Optimal movement between patches under incomplete information about the spatial distribution of food items. Theoretical Population Biology 70:452-463. doi: 10.1016/j.tpb.2006.04.002.

Kramer DL and McLaughlin RL, 2001. The behavioural ecology of intermittent locomotion. American Zoologist 41(2) doi: 10.1093/icb/41.2.137

MacArthur RH, 1965. Ecological consequences of natural selection, In: Theoretical and mathematical biology. New York: Blaisdell.

Mackay DJ, 2003. Information theory, inference and learning algorithms: Cambridge University Press.

MacLean LC, Thorp EO, Zhao Y, Ziemba WT, 2010. Medium term simulations of the full kelly and fractional kelly investment strategies. The Kelly Capital Growth Investment Criterion: Theory and Practice 15:543-562. doi: 10.1142/9789814293501_0038.

Maynard Smith J, Szathmary E, 1995. The major transitions in evolution. Oxford, England: W. H. Freeman and Co.

McLeman MA, Pratt SC, Franks NR, 2002. Navigation using visual landmarks by the ant *Leptothorax albipennis*. Insectes Sociaux 49:203-208. doi: 10.1007/s00040-002-8302-2.

McNamara, J. M., R. F. Green and O. Olsson (2006). "Bayes' theorem and its applications in animal behaviour." Oikos 112(2): 243-251. doi: 10.1111/j.0030-1299.2006.14228.x

Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E, 1953. Equation of state calculations by fast computing machines. Journal of Chemical Physics 21:1087-1092. doi: 10.1063/1.1699114.

Nathan R, Getz WM, Revilla E, Holyoak M, Kadmon R, Saltz D, Smouse PE, 2008. A movement ecology paradigm for unifying organismal movement research. Proceedings of the National Academy of Sciences of the United States of America 105:19052-19059. doi: 10.1073/pnas.0800375105.

Neal RM, 1993. Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-93-1: Dept. of Computer Science, Univ. of Toronto.

Noble D, 2011. A theory of biological relativity: no privileged level of causation. Interface Focus 2:55-64. doi: 10.1098/rsfs.2011.0067.

Nurse P, 2008. Life, logic and information. Nature 454:424-426. doi: 10.1038/454424a.

Olsson O, Brown JS, Lundberg P, 2006. The Foraging Benefits of Information and the Penalty of Ignorance. Oikos 112:260-273.

Orr HA, 2009. Fitness and its role in evolutionary genetics. Nature Reviews Genetics 10:531-539. doi: 10.1038/nrg2603.

Oster GF and Wilson EO, 1978. Caste and Ecology in the Social Insects. Princeton University Press.

Pratt SC, Brooks SE, Franks NR, 2001. The Use of Edges in Visual Navigation by the Ant *Leptothorax albipennis*. Ethology 107:1125-1136. doi: 10.1046/j.1439-0310.2001.00749.x.

Pratt SC, Sumpter DJT, Mallon EB, Franks NR, 2005. An agent-based model of collective nest choice by the ant Temnothorax albipennis. Animal Behaviour 70:1023-1036. doi: 10.1016/j.anbehav.2005.01.022.

Reid CR, Lutz MJ, Powell S, Kao AB, Couzine ID, Garnier S, 2015. Army ants dynamically adjust living bridges in response to a cost-benefit trade-off. Proceedings of the National Academy of Sciences of the United States of America 112:15113-15118. doi: 10.1073/pnas.1512241112.

Reynolds AM, Lepretre L, Bohan DA, 2013. Movement patterns of Tenebrio beetles demonstrate empirically that correlated-random-walks have similitude with a Levy walk. Scientific Reports 3. doi: 10.1038/srep03158.

Rising JK, Wyner AJ, Ieee, 2012. Partial Kelly Portfolios and Shrinkage Estimators. 2012 IEEE International Symposium on Information Theory Proceedings (ISIT).

Rossky PJ, Doll JD, Friedman HL, 1978. Brownian dynamics as smart Monte Carlo simulation. Journal of Chemical Physics 69:4628-4633. doi: 10.1063/1.436415.

Schultz TR, 2000. In search of ant ancestors. Proceedings of the National Academy of Sciences 97:14028-14029. doi: 10.1073/pnas.011513798.

Sumpter DJT, 2006. The principles of collective animal behaviour. Philosophical Transactions of the Royal Society B-Biological Sciences 361:5-22. doi: 10.1098/rstb.2005.1733.

Sumpter DJT, Pratt SC, 2003. A modelling framework for understanding social insect foraging. Behavioral Ecology and Sociobiology 53:131-144. doi: 10.1007/s00265-002-0549-0.

Van Gils JA, 2010. State-dependent Bayesian foraging on spatially autocorrelated food distributions. Oikos 119:237-244. doi: 10.1111/j.1600-0706.2009.17497.x.

Viswanathan GM, Raposo EP, da Luz MGE, 2008. Levy flights and superdiffusion in the context of biological encounters and random searches. Physics of Life Reviews 5:133-150. doi: 10.1016/j.plrev.2008.03.002.

von Frisch K, 1967. The dance language and orientation of bees. Cambridge, MA: Harvard University Press.

White EP, Enquist BJ, Green JL, 2008. On estimating the exponent of power-law frequency distributions. Ecology 89:905-912. doi: 10.1890/07-1288.1.