

Quantum Probabilistic Description of Dealing with Risk and Ambiguity in Foraging Decisions

Peter Wittek¹, Ik Soo Lim², and Xavier Rubio-Campillo³

¹ University of Borås, Borås, Sweden

² Bangor University, Bangor, United Kingdom

³ Barcelona Supercomputing Centre, Barcelona, Spain

Abstract. A forager in a patchy environment faces two types of uncertainty: ambiguity regarding the quality of the current patch and risk associated with the background opportunities. We argue that the order in which the forager deals with these uncertainties has an impact on the decision whether to stay at the current patch. The order effect is formalised with a context-dependent quantum probabilistic framework. Using Heisenberg’s uncertainty principle, we demonstrate the two types of uncertainty cannot be simultaneously minimised, hence putting a formal limit on rationality in decision making. We show the applicability of the contextual decision function with agent-based modelling. The simulations reveal order-dependence. Given that foraging is a universal pattern that goes beyond animal behaviour, the findings help understand similar phenomena in other fields.

1 Introduction

Studying foraging strategies has been a successful approach in understanding animal decision making [1]. Foraging patterns are not restricted to animal behaviour, similar patterns occur in other scenarios, such as searching in semantic memory [2], evaluating options by humans [3], and consumer behaviour [4].

Optimal foraging theory (OFT) studies foraging behaviour claiming that organisms aim to maximise their net energy intake per unit time [5, 6]. Food sources are available in patches, which vary in quality. Furthermore, switching between patches comes with a cost. The forager faces uncertainty while making decisions about staying at a patch or moving on to the next one. Following Knight, 1921 and Ellsberg, 1961, we distinguish between two fundamental types of uncertainty: ambiguity and risk [7, 8]. We associate ambiguity with the estimation of the quality of a patch. Risk, on the other hand, means the potential of the quality of other patches, the loss or gain by not foraging elsewhere as opposed to foraging in the current patch.

Decisions in this model are bound to be sequential: the forager must make decisions patch by patch; this assumption is not uncommon [9–11]. We argue that the order in which the forager deals with risk and ambiguity has an impact on the decision, which in turn influences net energy intake. We introduce a

contextual probabilistic framework familiar from quantum mechanics to model the decision making process.

A growing number of projects in Social Sciences use computer simulation as their main research tool. A simulation using an agent-based model (ABM) defines the behaviour of any entity of a system that involves decision-making processes known as agents. The generation of emergent properties that arise from the definition of individual agents include both quantitative and qualitative concepts, combining behaviour aspects and data. Thus, the explanation provided by an ABM is closer to how knowledge is acquired in Social Sciences. We rely on an ABM simulation to find evidence of order dependence and context sensitivity in decision making.

2 Foraging decisions, uncertainty, and context dependence

Risk and ambiguity are factors in various extensions of OFT, and they have been experimentally verified (Section 2.1). Context dependence, preferring one factor over the other is also a common behaviour pattern in various animal species (Section 2.2). These observations provide the foundations for our model.

2.1 Risk and ambiguity

Stochastic variants of OFT are successful in describing strategies that deal with ambiguity, with numerous experimental validations [12, 1]. Actual foraging strategies include simple heuristics such as the fixed-time strategy, in which the forager devotes the same amount of time to each patch irrespective of the patch quality. More intricate models of patch utilization include the Bayesian decision process. In this model, animals have an a priori assessments of food distributions, and their foraging decisions are influenced by experience [13, 1].

The Bayesian foraging strategy relies on the following assumptions [14]:

1. Perfect knowledge of patch-type distribution (a priori).
2. No instant identification of the quality of a particular patch, resulting in a sample.

The second assumption corresponds to *ambiguity* as a form of uncertainty.

Foraging decision is formulated by an a posteriori distribution made using the sample. An estimator keeps track of the mean value of the current patch, which is either under- or oversampled compared to the actual patch quality. This decision making process leads to density-dependent resource harvest. In the Bayesian model, the forager is allowed to make sequential decisions that vary according its current state, which is affected by the outcome of previous decisions.

The other aspect of uncertainty, *risk*, is also present in OFT. If foraging decisions are influenced by past history, the variations of any foraging parameter affect the expected rate of food gain, and hence the optimal foraging strategy [1]. Variance in a parameter is associated with risk. Foraging decisions are risk-sensitive, as empirical and theoretical proofs show [15, 16]. Risk sensitivity should have a sequential component, but it is often overlooked [9]. Whether a simultaneous or a sequential decision making model follows reality closer depends on the degree to which a forager commits itself when making a choice. For instance, in a sequence of choices, immediate rewards are more valuable than delayed ones: the time saved is used to pursue further rewards [10].

2.2 Contextuality

Context-dependent decision rules consider both aspects of uncertainty. We understand the human decision making is context-dependent [17, 18], but the phenomenon is less understood in animals. As Freidin and Kacelnik, 2011 put it, “context-dependent utility results from the fact that perceived utility depends on background opportunities” [11]. Spatially or temporally distinct patches are contexts that differ in quality. The sequence of optimal decisions depends on the attributes of the present opportunity and its background options.

Ample experimental evidence shows context-dependent decisions. Honey bees have an intransitive behaviour pattern, they perform a comparative evaluation of flowers depending on the context [19]. Gray jays have a similar behaviour, hinting at a complex decision process involving context-dependent assessment of each fitness-related value of the options [20]. Rufous humming birds change their risk preferences depending on whether binary or trinary choices are available to them [21]. The foraging choices of European starlings are better explained by context-dependent utility [22], the birds being more risk averse at lower temperatures [23]. Another study confirmed these findings, pointing out that sequential food encounters are more likely in an animal’s natural environment [11]. Simultaneous decision making is important in many species, for instance, humans are able to alternate between the two models of choice [24].

3 Quantum probabilistic description of foraging decisions

We turn to quantum probability theory to derive a formal decision model of a context-dependent optimal foraging strategy that considers both risk and ambiguity in a sequential setting. We use the quantum probabilistic description purely as a mathematical device, we do not conjecture that a forager’s context-dependent decision making process is the result of a macroscopic escalation of quantum

effects starting at a sub-atomic level. This investigation is part of a recent trend which claims quantum-like behaviour of systems is in fact not uncommon [25, 26].

Our efforts are not the first to bridge contextual probability and animal behaviour. Competing lizard communities and population dynamics also show quantum-like behaviour [27, 28].

3.1 Quantum probability and decisions

The fundamental difference between classical and quantum probability is that the event algebra in the latter is non-commutative. That is, given two events, A and B , $p(A \cap B) \neq p(B \cap A)$. To understand how non-commutativity leads to context-dependence, let us consider two hypotheses that describe the decision space of a forager:

- h_1 : Stay at the current patch.
- h_2 : Leave the patch.

These are complementary and mutually exclusive (this example resembles the one described in Trueblood and Busemeyer, 2011 about a medical decision on whether a patient has urinary tract infection [29]).

Consider the following events:

- A : Current patch quality with two possible outcomes: a_1 – the patch quality is good; a_2 – the patch quality is bad.
- B : Quality of other patches. A collective observation across all other patches with two possible outcomes.

A corresponds to ambiguity, whereas B corresponds to risk, the opportunity cost. We will argue that A and B are incompatible observations on a system, leading to non-commutativity and context-dependence. To achieve that, first we demonstrate how quantum probability theory derives probabilistic outcomes. In a quantum framework, the forager's state of belief is described by a state vector in superposition. Under observation A , this superposition is written as

$$|\psi\rangle = \sum_{i,j} \alpha_{ij} |A_{ij}\rangle \quad (1)$$

In this case, for instance, A_{h_1,a_1} means the event that the forager will stay at the patch and the current patch quality is good, and A_{h_2,a_2} corresponds to the event in which the forager will leave the patch and the patch has low quality. These elementary events form an orthonormal basis. The α_{ij} values are the

corresponding probability amplitudes. Since the norm of the state vector must be one, we have

$$\sum_{i,j} |\alpha_{ij}|^2 = 1. \quad (2)$$

To measure the probability of one combination, we use a projection operator. The projector onto the event $h_1 \wedge a_1$ is given by

$$P_{11} = P(h_1, a_1) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (3)$$

If we apply the projector on the state vector, the square norm of the projected vector will be the quantum probability of $h_1 \wedge a_1$: $\|P_{11}|\psi\rangle\|^2 = |\alpha_{11}|^2$. Similarly, if we are interested in the probability of whether the forager would stay at the current spot, we need to project onto the first two basis vectors, and the result will be $|\alpha_{11}|^2 + |\alpha_{12}|^2$.

Switching to observation B , the state of belief is a superposition of four basis vectors: $|\psi\rangle = \sum_{i,j} \beta_{ij} |B_{ij}\rangle$. These basis vectors are not identical to previous ones: the same hypotheses, the decision whether to stay or leave, are now expressed in a different basis. This is equivalent to saying that the same problem is studied from a different perspective: under observation A , the forager bases its decision on local information, whereas under B , it looks at a global perspective. If the two perspectives can be expressed in the same basis, they are called compatible, otherwise they are incompatible; in the present case we take A and B incompatible. Since the state vector can be expressed in both basis, we have

$$|\psi\rangle = \sum_{i,j} \alpha_{ij} |A_{ij}\rangle = \sum_{i,j} \beta_{ij} |B_{ij}\rangle. \quad (4)$$

To switch from one perspective to the other, to change the basis, we need to apply a unitary rotation. To change from perspective A to B , we need to apply the unitary transformation U_{AB} , whereas to change from B to A , we need to apply U_{BA} .

Context sensitivity arises from the quantum equivalent of Bayes' rule, also known as Lüders' rule [30]. Suppose that event A is true, the patch is of good quality. This changes the original state $|\psi\rangle$ to $|\psi_A\rangle = P_{a1}|\psi\rangle/\|P_{a1}|\psi\rangle\|$. The denominator re-normalises the state vector to meet the condition described in Eq. (2). The probability of event B given that A is true will be $\|P_{b1}|\psi_A\rangle\|^2 = \|P_{b1}P_{a1}|\psi\rangle\|^2/\|P_{a1}|\psi\rangle\|^2$. Generally speaking, two projectors are not commutative, that is, $P_{a1}P_{b1} \neq P_{b1}P_{a1}$. Therefore the probability of event A given that

B is true will be different. Measuring the perceived quality of other patches, the forager may deem those more desirable, and the quality of the current patch loses its importance.

Since the result of an earlier decision changes the context of a new decision, a simple Bayesian inference model has difficulty accounting for order effects [29]. A key concept of quantum probability is *incompatibility*. Not all elementary events can be measured simultaneously, incompatible events can only be measured in a sequence. The first assumption in the Bayesian decision process, which assumes the forager has an a priori knowledge of the probability distribution of the entire event space (Section 2.1), is too strong. For instance, the forager has an assessment of the quality of the current patch, which updates its a priori estimate of the distribution. To sample the frequency of such patches, it must move on to the next patch, abandoning the current one. The quality of the current patch and the quality of the other patches are incompatible observations. The latter is related to the risk the forager faces. Incompatibility is a source of order effects on judgements, and this is the critical point at which quantum probabilities differ from classic probabilities [30]. We show that this approach formally introduces a limit on rational decisions by applying the uncertainty principle to foraging decisions involving risk and ambiguity.

3.2 Inherent uncertainty in sequential decisions

If two observables do not commute, a state cannot be a simultaneous eigenvector of the two observables in general [31, p.233]. This leads to a form of uncertainty relation, similar to the one the inequality found by Heisenberg in his analysis of sequential measurements of position and momentum. This original relation states that there is a fundamental limit to the precision with which the position and momentum of a particle can be known.

Since observations A and B do not commute, there must be a similar limit in foraging decisions corresponding to Heisenberg's uncertainty principle. The forager needs to leave the current patch to assess the quality of other patches: there is an inherent uncertainty in the decision irrespective of the quantity of information gained about either A or B . Since the observables in foraging do not have a strict physical meaning, the physical constants of the canonical commutation relation are not present. Yet, as long as the operators do not commute, an uncertainty principle with a similar lower bound will hold. With regards to risk and ambiguity, we state that

$$\sigma_A \sigma_B \geq c, \tag{5}$$

where $c > 0$ is a constant. The constant itself will depend on actual foraging scenarios, we do not believe there is a universal way of determining its value.

4 Simulation by agent-based modelling

We use an ABM simulation to model contextuality of a decision function, which is not an entirely novel idea. Kitto, Boschetti, and Bruza (2012) proposed an ABM to simulate changing attitudes in social decision making [32]. The context in that case evolves with time, global attitudes and the individuals' own local attitudes change over time, making a case for a contextual decision making function. We do not require temporal evolution of contexts, we assume an order dependence emphasising the two different aspects of dealing with uncertainty.

We use Pandora [33, 34], an open-source ABM framework designed to accomplish a realistic simulation environment for social scientists. The source code of the simulation is available online⁴.

4.1 Model

In the simulation model, we generate a small map with heterogeneous resources available to the forager. An agent is placed at a random position with certain *requirements*: the quantity of resources needed at every time step. The *starvation rate* is the agent's accumulated percentage of past time steps when requirements were not achieved. The inverse of this value is the *net food intake*. The agent has a limited time *horizon*, the number of time steps into the future the agent is exploring possible scenarios. The agent has a *knowledge* of the patch distribution, it keeps track of the values and knowledge quality of patches already visited.

An agent is parametrized by the following characteristics:

- Ambiguity aversion: a value between 0 to 1 specifying the preference of the agent to avoid ambiguity. Ambiguity aversion is related to observation A .
- Risk aversion: a value between 0 to 1 specifying the preference of the agent to avoid risk. Risk aversion is related to observation B .

The sum of risk and ambiguity preferences is equal to one. An agent at every time step chooses between two actions: either explore and forage in the current patch, resulting in a decrease in ambiguity; or move and explore an adjacent patch, resulting in a decrease in risk.

Possible scenarios are explored using a Markov decision process that allows to explore the decisions of an agent within a fully observable stochastic state model. In our simulation, the model was solved using the UCT algorithm [35], capable of defining optimal policies for problems with a large set of possible states and finite horizon. The parameters of the algorithm were the horizon of

⁴ Code is available at <https://github.com/xrubio/pandora>. This example is in the `pandora/examples/quantumForaging` folder.

the agent and the width (i.e. the number of states explored during the process), the latter was defined at 300 for all the experiments.

The risk and ambiguity of the resources are varied, and also the agent's preference to minimizing risk (R) or minimizing ambiguity (A). The decision function is a normalised linear combination of the two options for minimizing uncertainty, combined with the cost associated with starving. Thus, the final cost C of an action a given current state s is as follows:

$$C_{s,a} = \frac{m}{M} + (1 - a_A)A + (1 - a_R)R \quad (6)$$

where m is the resources consumed in this time step, M corresponds to the resource requirements, a_A is the decrease of ambiguity of the action, and a_R is the decrease of risk of the action. The value of decrease is 1 if the knowledge of the patch is not complete, and 0 all the other cases.

Once an action is chosen, the knowledge of the patch where it occurs is increased by 1 if the maximum is not already reached. At the same time, the quality of the patch in the knowledge map of the agent is updated to the real value if a random value between 0 and 10 is lesser than current knowledge. Finally, the starvation rate of the agent is updated comparing the requirements with foraged resources; this will be always 0 when choosing a Movement action.

4.2 Experiments

We analysed the trade-off between ambiguity and risk with different time horizon thresholds for a resource map of 20×20 cells, and a time span of 1,000 steps. Every scenario contains a single agent, and the exploration of the parameter space is defined terms of time horizon values (1, 3, 5, and 7 steps), and risk aversion ranging from 0 to 1 in 0.1 increments. Each scenario was run 10 times for a total of 440 experiments. This configuration was used in two different experiments:

1. Gradual resource map. The value of each patch is its x relative coordinate multiplied by 10.
2. Random resource map. The value of each patch is defined using a random uniform distribution.

4.3 Discussion of results

Figure 1 shows net food intakes for the gradual resource maps, and Figure 2 shows net food intakes for the random distribution of resources.




Fig. 1. Net food intakes for different scenarios using the gradual resource map

A common pattern to all the experiments is a distinct phase transition at complete risk aversion. At risk aversion value of one, the net food intake sharply drops compared to all other values of this parameter. The phenomenon is not surprising: this behaviour means that the forager hardly ever consumes resources, and keeps visiting patch after patch to learn the overall distribution. While hardly a realistic scenario, it is worth taking note of.

For a time horizon of one time step, there is a different pattern for risk aversion below 0.5. The variance of net food intake is lowest at 0.1 in both types of resource distribution. Whereas the mean value is higher for larger values of risk aversion, a real-life forager would probably prefer reducing the variance. This behaviour indicates fairly long visits to a patch before exploring other options. To some extent, the observation holds for a time horizon of 3 steps. For risk aversion values higher than 0.5, the pattern is less obvious. While generally it pays off to reduce risk first, the exact extent of risk aversion is ambiguous.




Fig. 2. Net food intakes for different scenarios using the random resource map

If the time horizon is at least 5 steps, reducing risk first becomes far more important and leads to much higher rewards. A patch is easily depleted in five steps, so thinking ahead means a willingness to reduce risk, and the pay-off is clearly visible.

The values of risk aversion of zero and one are the extreme cases of order dependence, whereas the values in between correspond to a case of a more balanced decision making process. For a given parameter setting in the simulated environment, the highest point of the food intake is an empirical limit on the completeness of knowledge, and corresponds to the inequality in the uncertainty principle, as shown in Eq. (5).

5 Conclusions

Relying on a classical Markovian decision model, we simulated the behaviour of a forager in an environment where food resources are available in patches of varying quality. Well-defined patterns emerged that show order-dependence of decisions, and the decisions have a significant impact on net food intake. Our observations coincide with a quantum probabilistic model that considers two aspects of uncertainty, risk and ambiguity, and states that decisions relating to these two aspects are order-dependent.

6 Acknowledgement

This work was partially supported by the European Commission Seventh Framework Programme under Grant Agreement Number FP7-601138 PERICLES. Xavier Rubio-Campillo is supported by the SimulPast Project (CSD2010-00034), funded by the CONSOLIDER-INGENIO2010 program of the Ministry of Science and Innovation – Spain. We also thank the reviewers for their insights, they helped us clarify conceptual issues.

References

1. Pyke, G.: Optimal foraging theory: a critical review. *Annual Review of Ecology and Systematics* **15** (1984) 523–575
2. Hills, T., Jones, M., Todd, P.: Optimal foraging in semantic memory. *Psychological Review* (2012)
3. Rushworth, M., Kolling, N., Sallet, J., Mars, R.: Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Current Opinion in Neurobiology* (2012)
4. Rajala, A., Hantula, D.: Towards a behavioral ecology of consumption: delay-reduction effects on foraging in a simulated Internet mall. *Managerial and Decision Economics* **21**(3-4) (2000) 145–158
5. MacArthur, R., Pianka, E.: On optimal use of a patchy environment. *American Naturalist* **100**(916) (1966) 603–609
6. Charnov, E.: Optimal foraging, the marginal value theorem. *Theoretical Population Biology* **9**(2) (1976) 129–136
7. Knight, F.: Risk, uncertainty and profit. Houghton Mifflin, Boston (1921)
8. Ellsberg, D.: Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics* **75**(4) (1961) 643–669
9. Houston, A., McNamara, J.: A sequential approach to risk-taking. *Animal Behaviour* **30**(4) (1982) 1260–1261
10. Kacelnik, A., Bateson, M.: Risky theories – The effects of variance on foraging decisions. *American Zoologist* **36**(4) (1996) 402–434
11. Freidin, E., Kacelnik, A.: Rational choice, context dependence, and the value of information in European starlings (*Sturnus vulgaris*). *Science* **334**(6058) (2011) 1000–1002
12. Oaten, A.: Optimal foraging in patches: a case for stochasticity. *Theoretical Population Biology* **12**(3) (1977) 263–285

13. Green, R.: Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. *Theoretical Population Biology* **18**(2) (1980) 244–256
14. Valone, T., Brown, J.: Measuring patch assessment abilities of desert granivores. *Ecology* **70**(6) (1989) 1800–1810
15. Caraco, T., Martindale, S., Whittam, T.: An empirical demonstration of risk-sensitive foraging preferences. *Animal Behaviour* **28**(3) (1980) 820–830
16. Stephens, D.: The logic of risk-sensitive foraging preferences. *Animal Behaviour* **29**(2) (1981) 628–629
17. Anand, P.: *Foundations of rational choice under risk*. Oxford University Press (1995)
18. Waksberg, A., Smith, A., Burd, M.: Can irrational behaviour maximise fitness? *Behavioral Ecology and Sociobiology* **63**(3) (2009) 461–471
19. Shafir, S.: Intransitivity of preferences in honey bees: support for 'comparative' evaluation of foraging options. *Animal Behaviour* **48**(1) (1994) 55–67
20. Waite, T.: Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Behavioral Ecology and Sociobiology* **50**(2) (2001) 116–121
21. Hurly, T., Oseen, M.: Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour* **58**(1) (1999) 59–66
22. Pompilio, L., Kacelnik, A.: Context-dependent utility overrides absolute memory as a determinant of choice. *Proceedings of the National Academy of Sciences* **107**(1) (2010) 508–512
23. Bateson, M.: Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour* **64**(2) (2002) 251–260
24. Kolling, N., Behrens, T., Mars, R., Rushworth, M.: Neural mechanisms of foraging. *Science* **336**(6077) (2012) 95–98
25. Pitowsky, I.: *Quantum Probability – Quantum Logic*. Springer (1989)
26. Khrennikov, A.: *Ubiquitous quantum structure: from psychology to finance*. Springer Verlag (2010)
27. Aerts, D., Czachor, M., Kuna, M., Sinervo, B., Sozzo, S.: Quantum probabilistic structures in competing lizard communities. arXiv preprint arXiv:1212.0109 (2012)
28. Aerts, D., Czachor, M., Kuna, M., Sozzo, S.: Systems, environments, and soliton rate equations: A non-Kolmogorovian framework for population dynamics. arXiv preprint arXiv:1303.0281 (2013)
29. Trueblood, J., Busemeyer, J.: A quantum probability account of order effects in inference. *Cognitive Science* **35** (2011) 1518–1552
30. Busemeyer, J., Pothos, E., Franco, R., Trueblood, J.: A quantum theoretical explanation for probability judgment errors. *Psychological Review* **118**(2) (2011) 193–218
31. Cohen-Tannoudji, C., Diu, B., Laloë, F.: *Quantum mechanics* (1996)
32. Kitto, K., Boschetti, F., Bruza, P.: The quantum inspired modelling of changing attitudes and self-organising societies. In: *Proceedings of QI-12, 6th International Quantum Interaction Symposium, Paris, France, Springer (June 2012)*
33. Rubio-Campillo, X.: Pandora: An HPC agent-based modelling framework. <https://github.com/xrubio/pandora/> (2013) Accessed: 2013-04-01.
34. Wittek, P., Rubio-Campillo, X.: Scalable agent-based modelling with cloud HPC resources for social simulations. In: *Proceedings of CloudCom-12, 4th IEEE International Conference on Cloud Computing Technology and Science, Taipei, Taiwan (December 2012)* 355–362
35. Bonet, B., Geffner, H.: Action selection for MDPs: Anytime AO* vs. UCT. In: *Proceedings of AAAI-12, 26th Conference on Artificial Intelligence, Toronto, Canada (July 2012)*