Ignorance can be evolutionarily beneficial

Jared M. Field*
Wolfson Centre for Mathematical Biology,
Mathematical Institute, University of Oxford,
Oxford OX2 6GG, United Kingdom
†

Michael B. Bonsall
Mathematical Ecology Research Group, Department of Zoology,
University of Oxford, Oxford, UK

(Dated: November 10, 2018)

Information is increasingly being viewed as a resource used by organisms to increase their fitness. Indeed, it has been formally shown that there is a sensible way to assign a reproductive value to information and it is non-negative. However, all of this work assumed that information collection is cost-free. Here, we account for such a cost and provide conditions for when the reproductive value of information will be negative. In these instances, counter-intuitively, it is in the interest of the organism to remain ignorant. We link our results to empirical studies where Bayesian behaviour appears to break down in complex environments and provide an alternative explanation of lowered arousal thresholds in the evolution of sleep.

Keywords: Statistical Decision Theory — Bayes' Theorem — Information — Ignorance — Sleep

In all areas of biology observations or cues can elicit a change in behaviour or phenotype. For example, the duration of daylight hours affects the flowering time of plants [2], chemotactic gradients provide a way for bacteria to locate favourable environments [1] and the sighting of a predator may cause an animal to flee. In each case, the observation aids in the choice of an action that will benefit the organism. In this way, if there is no cost in collecting information then organisms should always collect it. Indeed, borrowing from economic theory, but replacing utilities with reproductive values, it has been shown that the reproductive value of information is always non-negative [14, 17]. This remarkable result suggests that organisms can never decrease their fitness by being more informed. This has far reaching implications in areas of biology as diverse as public goods games, foraging, collective behaviour and sleep. However, in reality, information comes at a cost [10]. In this paper, we investigate the consequences of formally including such costs in the current theoretical framework. While information may be inherently valuable in decision making, we challenge the view that it always should be or indeed is collected.

In the context of organismal biology, information use has been approached from two very distinct angles. The first employs the information theory pioneered by Shannon and Weaver [20]. The alternative, which we take, makes use of statistical decision theory [5, 13, 15]. The former focuses on uncertainty reduction whereas the latter considers how information updating, via Bayes rule, explicitly affects fitnesses. Only recently was the connection between these two approaches shown. Strikingly, mutual information (an information-theoretic measure which quantifies the uncertainty of an outcome after an observation) provides an upper bound on the value of information (a decision-theoretic measure expressed explicitly in terms of fitnesses) [6]. However, this is only the case when the fitness measure used is long-term lineage growth rate. Here, we use individual reproductive values. In this way, the value of information is defined by taking the difference in expected optimal reproductive values before and after collecting information. The literature on information use in biology is vast so we do not try to cover it here. However, for useful reviews see [5, 18, 19, 22]

We start by adopting the same framework as in [14]. Following this, we derive conditions, in the absence of a cost, for when the reproductive value of information is precisely equal to zero. In this case, an organism that collects information will be no better off than one that does not. Next, we account for a cost of collecting information. We do this by discounting the previous reproductive values associated with certain actions and beliefs prior to information collection. This way, the cost is due solely to the collection of information itself. In the proceeding section, we show that the same conditions that we derived previously now lead to the reproductive value of information being negative. In this instance, an organism that remains ignorant will have a higher fitness than one that does not. This has implications for many empirical studies whereby, in increasingly complex environments, organisms were found to stop behaving in a Bayesian manner [9]. It may be that the cost of collection

---

* jared.field@maths.ox.ac.uk
† Also at Mathematical Ecology Research Group, Department of Zoology, University of Oxford, Oxford, UK

in these complex environments outweighs any benefits. Additionally, we consider the ramifications for evolutionary problems such as the evolution of sleep. In particular, we suggest that the lowered arousal thresholds associated with sleep, often explained as the by-product of other vital functions, may instead be accounted for by our results [11]. We finish by summarising our findings and suggesting future directions of work relating ignorance at the level of the individual to behaviour at the level of the group [4].

RESULTS

Framework

Following [14], we start by assuming there are n possible true states of nature, collected in the vector $\Theta = (\theta_1, .., \theta_n)$, about which a given organism is unsure. This organism believes, with probability p_i, that θ_i is the true state. In this way the vector $P = (p_1, ..., p_n)$ summarises the organism's (imperfect) knowledge of nature.

Further, we assume that this organism must take a certain action, given its beliefs, which will affect its fitness. In particular, a certain action u will have reproductive value $V_i(u)$ given the true state of nature is θ_i. With this set-up, we find the expected reproductive value of action u to be given by

$$V(u, P) = \sum_{i=1}^{n} p_i V_i(u). \tag{1}$$

Then for any vector P, we define the optimal action u^* so that

$$V(u, P) \leq V(u^*, P), \tag{2}$$

for any action u.

We now suppose that this organism can gather information thereby updating its knowledge. We assume it does this by sampling a random variable X which depends on Θ. We denote the probability the observed value of X is x given θ_i by $f(x|\theta_i)$. With this, we can interpret p_i to be a prior probability and use the observation to form its posterior using Bayes rule, which we denote by $q_i(x)$. Doing so gives

$$q_i(x) = \frac{p_i f(x|\theta_i)}{\sum_{j=1}^{n} p_j f(x|\theta_j)}. \tag{3}$$

With this extra information, the organism's knowledge is now summarised by $Q(X) = (q_1, ..., q_n)$.

With this set-up we define the expected reproductive value (taken over observations X) of information I such that [14]:

$$I = \mathbf{E}[V(\tilde{u}(X), Q(X))] - V(u^*, P). \tag{4}$$

The first term of this quantity is the reproductive value after collecting information, optimised over actions, and averaged over observations. The second term is the optimal reproductive value if information is not collected. In this way, the difference quantifies the benefit of collecting information.

Expected value of information is non-negative

While inequality (2) is expressed in terms of P it is also true for the random vector $Q(X)$. In particular,

$$V(u, Q(X)) \leq V(\tilde{u}(Q(X)), Q(X)). \tag{5}$$

Taking the expected value with respect to X we then have that

$$\mathbf{E}[V(u, Q(X))] \leq \mathbf{E}[V(\tilde{u}(Q(X)), Q(X))]. \tag{6}$$

The left hand side of inequality (6) can be rewritten as

$$\begin{aligned}\mathbf{E}[V(u, Q(X))] &= \mathbf{E}[\sum_{i=1}^{n} q_i(X) V_i(u)] && (7)\\ &= \sum_{i=1}^{n} \mathbf{E}[q_i(X) V_i(u)] && (8)\\ &= \sum_{i=1}^{n} \mathbf{E}[q_i(X)] V_i(u) && (9)\end{aligned}$$

However, as the expectation is taken over observations x, we have

$$\mathbf{E}[q_i(X)] = \sum_{x} q_i(x) f(x) \tag{10}$$

where $f(x)$ is the probability that the observed value of X is x. Using the law of total probability, $f(x)$ can in fact be written as

$$f(x) = \sum_{j=1}^{n} p_j f(x|\theta_j), \tag{11}$$

so that, coupled with (3), (10) can be reexpressed as

$$\begin{aligned}\mathbf{E}[q_i(X)] &= p_i \sum_{x} f(x|\theta_j), && (12)\\ &= p_i, && (13)\end{aligned}$$

as $f(x|\theta_j)$ is a probability distribution.

Now, using (13) in (9) we find that

$$\begin{aligned}\mathbf{E}[V(u, Q(X))] &= \sum_{i=1}^{n} p_i V_i(u) && (14)\\ &= V(u, P), && (15)\end{aligned}$$

by (1). This statement is true for any action u so that it is in particular true of the action $u^*(P)$ that optimises $V(u, P)$.

Finally, using (15) in inequality (6) and rearranging we find that

$$0 \leq \mathbf{E}[V(\tilde{u}(Q(X)), Q(X))] - V(u^*, P). \tag{16}$$

Otherwise put, the expected reproductive value of information I is non-negative.

When is the expected value of information equal to zero?

So far we have shown (and also in [14]) that the expected value of information is non-negative. However, if it is equal to zero then an organism that collects information will be no better off than an organism that does.

To investigate this, we start by bounding inequality (6) from above. Similarly to equations (7) and (8) but for the right hand side of (6) we have

$$\mathbf{E}[V(\tilde{u}(Q(X)),Q(X))] = \mathbf{E}[\sum_{i=1}^{n} q_i(X)V_i(\tilde{u}(X)] \tag{17}$$

$$= \sum_{i=1}^{n} \mathbf{E}[q_i(X)V_i(\tilde{u}(X))]. \tag{18}$$

However this time as the action $u(X)$ taken depends on the observation we must take the expectation directly. In particular, we have

$$\mathbf{E}[V(\tilde{u}(Q(X)),Q(X))] = \sum_{i=1}^{n}\sum_{x} q_i(x)V_i(\tilde{u}(x))f(x). \tag{19}$$

We now suppose that there exists some $\hat{x}$ such that

$$\sum_{i=1}^{n}\sum_{x} q_i(x)V_i(\tilde{u}(x))f(x) \leq \sum_{i=1}^{n}\sum_{x} q_i(x)V_i(\tilde{u}(\hat{x}))f(x) \tag{20}$$

Using (3), this then becomes

$$\sum_{i=1}^{n}\sum_{x} q_i(x)V_i(\tilde{u}(x))f(x) \leq \sum_{i=1}^{n} p_i V_i(\tilde{u}(\hat{x}))\sum_{x} f(x|\theta_i) \tag{21}$$

$$= V(\tilde{u}(\hat{x}),P), \tag{22}$$

as $f(x|\theta_i)$ is a probability distribution. Hence, if such an $\hat{x}$ does exist, we have showed that

$$V(u^*,P) \leq \mathbf{E}[V(\tilde{u}(Q(X)),Q(X))] \leq V(\tilde{u}(\hat{x}),P). \tag{23}$$

However, from the very definition of u^* it must be that

$$V(\tilde{u}(\hat{x}),P) \leq V(u^*,P), \tag{24}$$

so that in fact

$$\mathbf{E}[V(\tilde{u}(Q(X)),Q(X))] = V(u^*,P), \tag{25}$$

and so

$$I = 0. \tag{26}$$

For such an $\hat{x}$ to exist we need, from (20), that

$$\sum_{x} q_i(x)V_i(\tilde{u}(x))f(x) \leq \sum_{x} q_i(x)V_i(\tilde{u}(\hat{x}))f(x), \tag{27}$$

for each i. Using (3) this is equivalent to requiring

$$\sum_{x} f(x|\theta_i)V_i(\tilde{u}(x)) \leq V_i(\tilde{u}(\hat{x}))\sum_{x} f(x|\theta_i) \tag{28}$$

$$= V_i(\tilde{u}(\hat{x})), \tag{29}$$

which is in turn equivalent to

$$f(\hat{x}|\theta_i)V_i(\tilde{u}(\hat{x})) + \sum_{x\neq\hat{x}} f(x|\theta_i)V_i(\tilde{u}(x)) \leq V_i(\tilde{u}(\hat{x})). \tag{30}$$

This statement will clearly be true in three cases. Either

$$V_i(\tilde{u}(x)) = 0, \tag{31}$$

for each $x \neq \hat{x}$ and for each i or

$$f(x|\theta_i) = 0, \tag{32}$$

for each $x \neq \hat{x}$ and for each i or, finally, a mixture of these two previous extreme cases. In particular, we can have instances where there is a particular $x \neq \hat{x}$ so that

$$V_i(\tilde{u}(x)) \neq 0, f(x|\theta_i) = 0, \tag{33}$$

for each i or *vice versa* so that

$$V_i(\tilde{u}(x)) = 0, f(x|\theta_i) \neq 0, \tag{34}$$

for each i.

Condition (31) asserts that all other possible actions will have a reproductive value of zero. Though unlikely, this may be realised if presented with an extreme situation (such as certain death in the face of predation) when only one action will lead to survival.

Alternatively, via (11) condition (32) implies

$$f(\hat{x}) = 1. \tag{35}$$

The interpretation here is that only the observation $\hat{x}$ will be observed. In this way, this condition says something about the environment in which an organism finds itself. If the organism finds itself in a period of relative constancy (such as certain safety due to a temporary lack of predators), this condition will be fulfilled.

The blended conditions (33) and (34) represent the more realistic situations where there is still flexibility in actions with nonzero reproductive value but the observation to which it corresponds will not be observed or where an observation will have nonzero probability of being observed but the associated reproductive value is zero, respectively. Otherwise put, an organism may be adapted and able to respond to a certain cue however currently is in a situation where that cue will not be observed. Alternatively, the organism may observe a certain cue but is unable to adequately respond.

Accounting for a cost of information collection

We now assume that the collection of information comes at a cost c. The reproductive value associated with any given action u taken after gathering information will then be discounted by this cost. Hence, if an organism collects information the reproductive value of action u given state θ_i will be

$$\bar{V}_i(u) = V_i(u) - c. \tag{36}$$

As before, we have that

$$\mathbf{E}[\bar{V}(u, Q(X))] \leq \mathbf{E}[\bar{V}(\tilde{u}(Q(X)), Q(X))], \tag{37}$$

for any u. Replacing $V_i(u)$ with $\bar{V}_i(u)$ in (9), the left hand side of this inequality can be written as

$$\mathbf{E}[\bar{V}(u, Q(X))] = V(u, P) - c, \tag{38}$$

which is, in particular, true for u^* so that we have

$$V(u^*, P) - c \leq \mathbf{E}[\bar{V}(\tilde{u}(Q(X)), Q(X))] \tag{39}$$

Similarly, replacing occurrences of $V_i(\cdot)$ with $\bar{V}_i(\cdot)$ in the entire preceding section it follows that

$$\mathbf{E}[\bar{V}(\tilde{u}(Q(X)), Q(X))] \leq V(u^*, P) - c, \tag{40}$$

so that, if there is a cost to collecting information we find that

$$\mathbf{E}[\bar{V}(\tilde{u}(Q(X)), Q(X))] = V(u^*, P) - c. \tag{41}$$

Hence

$$I = -c. \tag{42}$$

This result, as before, relies on one of the four conditions (31)-(34) being true. In this case, an organism that collects information will be worse off than an organism that does not.

Alternative measures for valuing information

With some important recent exceptions, it is not generally believed that organisms are behaving in a strictly Bayesian manner [3, 9, 12]. Instead, it is suspected that they follow Rules of Thumb that approximate optimal Bayesian strategies [13, 15]. These rules are based on either the experience of an organism's ancestors (and so is genetically encoded), the experience of the organism itself or a combination of both [7, 9, 15] . When these rules are based on the experience of the organism they will be informed by typical observations. In some circumstances, the expected value (such as with the analysis performed above) will be a good indication of typical sampled values.

However, in other cases the expected value is in fact a very poor measure of typical values. To be more concrete, consider $G(X)$ defined such that

$$G(X) = V(\tilde{u}(Q(X)), Q(X)) - V(u^*, P) \tag{43}$$

with probability density function $g(x)$. Note that this random variable has the property that $\mathbf{E}[G(X)] = I$.

If $g(x)$ is a unimodal symmetric distribution then I will be a decent indication of typical values. However, if $g(x)$ is, say, multimodal or skewed then I will be a poor measure. Moreover, if $g(x)$ is positively skewed then typical values of G may be negative. In this case, *even without a cost associated with the act of information collection*, it will not be beneficial for an organism operating under a Rule of Thumb to collect information.

Whether or not $g(x)$ is skewed will depend on $f(x)$, the distribution describing X. However, the random variable X describes the environment the organism is in. For this reason, it is highly plausible that X and so $g(x)$ will change as a function of time. For example, it may be that predators are more prevalent at a certain time over the course of one day. In this case, the probability of observing a predator will also change over that period. So too, then, will $g(x)$.

DISCUSSION

The notion of putting a value on information has existed in economic theory for quite some time [8]. Despite this, analogous work in the context of animal behaviour has taken a little longer to catch up. Recently, however, it was formally shown that the reproductive value of information is always non-negative when there is no cost in its collection [14]. From this, it has been concluded that organisms should always collect information. This conclusion has since not been investigated much further in the literature. While costs associated with sampling have been considered in some evolutionary games (see [16], where the costs are crucial to maintain a mix of strategies), the focus has not been on the whether or not information should be collected in the first place.

Here, using the same framework, we derived particular conditions for when the value of information will be equal to zero. In such a case, an organism that collects information will in fact be no better off than one that does not. Following this, we explicitly accounted for a cost of collecting information. In reality, this may come in the form of energy, time or resources that could otherwise be spent on other vital biological functions [10]. In this instance, we found that under the same conditions as before, the value of information will now be negative. Thus, there will be times when the collection of information will have a negative impact on an organism's fitness.

These conditions can be broadly organised into three groups. First, if an organism finds itself (potentially temporarily) in a situation whereby only one action will lead to a non-zero reproductive value, then the value of information will be negative. Second, if it is such that only one observation will be observed (again, potentially temporarily), then an organism will not benefit from collecting information. Both of these cases, in the absence of a cost of collecting information, are emphasised informally in [17]. The third, a mixture of the previous two conditions, is not. In this case, an organism can still be flexible in its actions, and there can still be variance in observations, and yet the value of information will be less than zero. To be more concrete, this third condition will be fulfilled if an organism may observe a certain cue but is unable to adequately respond or if it is adapted and able to respond to a certain cue however currently is in a situation where it will not be observed. In either case, an organism will do better if they remain relatively ignorant.

Though the results presented here are quite broad they, and in particular the last condition, have strong implications for the evolution of sleep and sleep-like states. Broadly speaking, sleep can be defined physiologically (characterised by certain brain activities) or behaviourally (characterised by inactivity and arousal thresholds). Some organisms, like dolphins, sleep according to one definition but not the other [21]. It is however the behavioural definition, and in particular the lowered arousal thresholds, that presents the evolutionary puzzle. This vulnerable disconnect from an organism's environment is most often explained by assuming *a priori* that sleep serves some vital function for which this behavioural shutdown is necessary [11]. Here, however, we have shown that such an assumption is not necessary. If there is a cost to collecting information, then there will be times when it should not be collected. We are not, of course, suggesting that vital functions of sleep do not exist. However, our result opens up the possibility that the vital functions evolved after behavioural shutdowns and not the other way around.

In an empirical review of animal Bayesian updating, it was found that in simple environments all but one species performed consistently with Bayesian predictions [9]. In complex environments, this was not found to be true. The explanation put forward was that in these complex environments it is more difficult, or takes longer, for the organism in question to successfully learn prior distributions. An alternative and complementary explanation, suggested by our results, is that the cost of collecting information may be too high in these environments. In other words, they are making a Bayesian decision to remain ignorant. It would be exceedingly interesting to test this hypothesis experimentally.

More generally, we have shown that periods of ignorance can lead to fitness benefits at the level of the individual. For future theoretical work, it would be fruitful to understand how this may translate to behaviour at the level of the group. Recently, it has been shown how uninformed individuals can help democratic consensus be reached in the face of internal conflicts [4]. This particular study however makes no reference to individual fitnesses but is, in essence, mechanistic. It would be particularly interesting to see if, taking an evolutionary approach, similar conclusions can be found.

COMPETING INTERESTS

We have no competing interests.

AUTHORS' CONTRIBUTIONS

JMF carried out the research. JMF and MBB wrote the manuscript.

FUNDING

JMF is funded by the Charles Perkins Scholarship with additional financial support from UTS, Sydney.

ACKNOWLEDGMENTS

We thank Thomas W. Scott and John M. McNamara for valuable comments and discussion.

---

[1] Julius Adler. Chemotaxis in bacteria. *Science*, 153(3737):708–716, 1966.

[2] Richard Amasino. Seasonal and developmental timing of flowering. *The Plant Journal*, 61(6):1001–1013, 2010.

[3] Jay M Biernaskie, Steven C Walker, and Robert J Gegear. Bumblebees learn to forage like bayesians. *The American Naturalist*, 174(3):413–423, 2009.

[4] Iain D Couzin, Christos C Ioannou, Güven Demirel, Thilo Gross, Colin J Torney, Andrew Hartnett, Larissa Conradt, Simon A Levin, and Naomi E Leonard. Uninformed individuals promote democratic consensus in animal groups. *science*, 334(6062):1578–1580, 2011.

[5] Sasha RX Dall, Luc-Alain Giraldeau, Ola Olsson, John M McNamara, and David W Stephens. Information and its use by animals in evolutionary ecology. *Trends in ecology & evolution*, 20(4):187–193, 2005.

[6] Matina C Donaldson-Matasci, Carl T Bergstrom, and Michael Lachmann. The fitness value of information. *Oikos*, 119(2):219–230, 2010.

[7] Luc-Alain Giraldeau, JR Krebs, and NB Davies. The ecology of information use. *Behavioural ecology: an evolutionary approach*, pages 42–68, 1997.

[8] John P Gould. Risk, stochastic preference, and the value of information. *Journal of Economic Theory*, 8(1):64–84, 1974.

[9] Thomas J Valone. Are animals capable of bayesian updating? an empirical review. *Oikos*, 112(2):252–259, 2006.

[10] Simon B Laughlin, Rob R de Ruyter van Steveninck, and John C Anderson. The metabolic cost of neural information. *Nature neuroscience*, 1(1):36–41, 1998.

[11] Steven L Lima and Niels C Rattenborg. A behavioural shutdown can make sleeping safer: a strategic perspective on the function of sleep. *Animal Behaviour*, 74(2):189–197, 2007.

[12] Philippe Louapre, Joan Van Baaren, Jean-Sébastien Pierre, and JJM Van Alphen. Information gleaned and former patch quality determine foraging behavior of parasitic wasps. *Behavioral Ecology*, page arr090, 2011.

[13] John McNamara and Alasdair Houston. The application of statistical decision theory to animal behaviour. *Journal of Theoretical Biology*, 85(4):673–690, 1980.

[14] John M McNamara and Sasha RX Dall. Information is a fitness enhancing resource. *Oikos*, 119(2):231–236, 2010.

[15] John M McNamara, Richard F Green, and Ola Olsson. Bayes theorem and its applications in animal behaviour. *Oikos*, 112(2):243–251, 2006.

[16] John M McNamara, Philip A Stephens, Sasha RX Dall, and Alasdair I Houston. Evolution of trust and trustworthiness: social awareness favours personality differences. *Proceedings of the Royal Society of London B: Biological Sciences*, 276(1657):605–613, 2009.

[17] Rebecca K Pike, John M McNamara, and Alasdair I Houston. A general expression for the reproductive value of information. *Behavioral Ecology*, page arw044, 2016.

[18] Kenneth A Schmidt, Sasha RX Dall, and Jan A Van Gils. The ecology of information: an overview on the ecological significance of making informed decisions. *Oikos*, 119(2):304–316, 2010.

[19] Janne-Tuomas Seppänen, Jukka T Forsman, Mikko Mönkkönen, and Robert L Thomson. Social information use is a process across time, space, and ecology, reaching heterospecifics. *Ecology*, 88(7):1622–1633, 2007.

[20] C. E. Shannon and W. Weaver. The mathematical theory of communication. 1949.

[21] Jerome M Siegel. Sleep viewed as a state of adaptive inactivity. *Nature Reviews Neuroscience*, 10(10):747–753, 2009.

[22] Thomas J Valone. From eavesdropping on performance to copying the behavior of others: a review of public information use. *Behavioral ecology and sociobiology*, 62(1):1–14, 2007.