

Shannon Information and Biological Fitness

Carl T. Bergstrom
Department of Biology
University of Washington
Seattle, WA 98195-1800, USA
e-mail: cbergst@u.washington.edu

Michael Lachmann
Max Planck Inst. for Evol. Anthropology
Deutscher Platz 6
D-04103 Leipzig, Germany
email: lachmann@eva.mpg.de

Abstract — When studying information, biologists and behavioral scientists often eschew Shannon entropy. Instead, they commonly use a decision-theoretic measure of the value of information, on the grounds that Shannon’s measure draws no distinction between useful and useless information. Here we show that these two measures are intimately related in the context of biological evolution. We present a simple model of evolution in an uncertain environment, and calculate the increase in Darwinian fitness that is made possible by information about the environmental state. This fitness increase — the *fitness value of information* — is a composite of both the Shannon entropy and the decision-theoretic measure of information value. Furthermore, the Shannon entropy of the environment, which seemingly fails to take anything about Darwinian fitness into account, nonetheless imposes an upper bound on the fitness value of information.

I. INTRODUCTION

Information appears in almost every area of biology — from the mating signals transmitted between the sexes, to environmental cues used by plants to adapt to their environment, to digital storage of information in the DNA. Nonetheless, information theoretic measures such as Shannon entropy or mutual information are seldom used in many of the areas of biology that aim to understand how organisms have evolved to deal with information, including behavioral biology and evolutionary ecology.

The problem is that Shannon entropy and mutual information do not directly address information quality; they do not distinguish between relevant and irrelevant information. Thus decision theorists, economists, and behavioral biologists typically measure the value of information by its effect on expected payoff or expected fitness [1, 2, 3, 4, 5].

Definition: The *value of information* associated with a cue or signal X is defined as the difference between the maximum expected payoff or fitness that a decision-maker can obtain by conditioning on X and the maximum expected payoff that could be obtained without conditioning on X .

The dissonance between Shannon entropy and the value of information has long puzzled biologists in general and the authors of this paper in particular. Shannon entropy (and mutual information) appear to measure information quantity while reflecting nothing about fitness consequences; the value of information measures fitness consequences but has nothing to do with the actual length or information quantity of a message. What, if any, are the relations between them? Information theorists since Kelly [7] have observed that in

special circumstances, information value and Shannon’s measures may be related. Here we argue that these “special circumstances” are exactly those about which biologists should be most concerned: the context of evolution by natural selection. We address the question “how much is information worth to living organisms?” and show that the answer combines both Shannon entropy and the decision-theoretic value of information.

II. A BASIC MODEL

How should a biologist measure the cost of uncertainty¹ or the value of information? In biology, Darwinian fitness is paramount: as a first approximation, we expect the evolved traits of organisms to increase the fitness of those who exhibit them. Thus when we look at biological adaptations for acquiring or processing information, the relevant value of this information is measured in the currency of Darwinian fitness. Put simply, we want to know how the information effects fitness. The natural measure the “worth” of information to a biological organism is the following: the *fitness value of information* is the greatest fitness decrement or cost that would be favored by natural selection in exchange for the ability to attain the information.

We and others have previously demonstrated that like stock brokers or habitual gamblers, biological organisms faced with uncertain conditions are selected to behave as if they are concerned with long-term growth rates. Thus the fitness value of information to biological organisms is best measured in terms of the consequences of this information on the long-term growth rates of organismal lineages. These long-term growth rates can be measured as the expected value of the logarithm of the growth rate in a single generation [9] (as opposed to the expected value of the growth rate itself).

To illustrate these results and to develop an intuition about the value of information in biological systems, consider the following simple model of organisms living in a variable environment². The environment has two possible states, state 1 and

¹Numerous studies in population ecology and genetics have shown that fitness and population growth in uncertain environments depend on the exact nature of the uncertainty; they depend both on the distribution of individual reproductive successes, and on the correlations in individual successes (reviewed in ref. [8]). One can capture this complexity by distinguishing between two types of uncertainty or risk [9]. *Idiosyncratic risk* is independent of that faced by other individuals, whereas *aggregate risk* is perfectly correlated among individuals. For example, predation imposes largely idiosyncratic risk on a herd of herbivores, whereas drought imposes largely aggregate risk. In this paper, we focus exclusively on aggregate risk. We will address mixed aggregate and idiosyncratic risk in a subsequent report.

²While this model or variants thereof appear commonly in the biology literature, the inspiration for this paper came largely from J. L. Kelly’s application of information theory to gambling problems

state 2, that occur with probability p_1 , and $p_2 = 1 - p_1$ respectively. All individuals encounter exactly the same environment in a given period. At the beginning of its development, each organism makes an important developmental decision to adopt one of two alternative physiologies or *phenotypes*: one suited to environment 1, or one suited to 2. The organism survives to reproduce only if its physiology properly matches the demands of the current environment. The organism's fitness is given by the following matrix:

	Phenotype 1	Phenotype 2
Environment 1	w_1	0
Environment 2	0	w_2

What should these individuals do in the absence of information about the condition of the environment? In the short run, individuals maximize expected fitness by playing the highest-payoff strategy only. This yields an expected single-generation fitness of $\max[p_1 w_1, p_2 w_2]$.

But in the long run, playing only one strategy will inevitably lead to a year with zero fitness and subsequent extinction. Thus natural selection will favor not the short run maximization above, but rather a maximization of long-term fitness. These organisms will be selected to *hedge their bets* [9, 10], developing into phenotype 1 with some probability and phenotype 2 otherwise. As the number of generations N grows large, natural selection is overwhelmingly likely [9] to favor the strategy that maximizes the growth rate for typical sequences [6], in which environment 1 occurs $N p_1$ times, and environment 2 occurs $N p_2$ times. For a genotype that develops with probability x into phenotype 1, the growth rate for such a sequence of events will be $(w_1 x)^{N p_1} (w_2 (1 - x))^{N p_2}$ and the log of this growth rate will be maximized when $N(p_1 \log(w_1 x) + p_2 \log(w_2 (1 - x)))$ is maximized. This occurs when $p_1/x = p_2/(1 - x)$ or when $x = p_1$. Thus for almost all sequences of environments, the strategy that develops with probability p_1 into phenotype 1 will maximize expected log growth rate and thus take over the population. For this strategy, the expected log growth rate will be $p_1 \log(w_1 p_1) + p_2 \log(w_2 p_2)$.

We have set up a simple biological model where uncertainty critically affects Darwinian fitness. What is the fitness value of information here? Suppose that individuals are able to detect a cue that they can use to forecast (100% accurately) the state of the environment. In this case the organism will use phenotype 1 in environment 1, and phenotype 2 in environment 2. What is the fitness value of [being able to obtain] this cue?

First, we can look at how the cue improves the short-run expected fitness. With the cue, individuals can always develop the appropriate phenotype for the environment, and obtain short-run expected fitness $p_1 w_1 + p_2 w_2$. Thus in the short run, the expected value of information is $p_1 w_1 + p_2 w_2 - \max[p_1 w_1, p_2 w_2] = \min[p_1 w_1, p_2 w_2]$. This is exactly the decision-theoretic value of information.

[7]. In this section, we follow Cover and Thomas's (1991) presentation; these authors offer a parable about a sempiternal gambler who perpetually reinvests his entire winnings at the horse track. Their gambling story can be recast quite naturally as a model of organisms evolving by natural selection to match their physiologies to uncertain environmental conditions. In subsequent sections, we extend these results to further explore our model of biological evolution.

But natural selection will not maximize this quantity; instead as discussed above it maximizes the expected log growth rate. Without the cue, expected log growth rate is $p_1 \log(p_1 w_1) + p_2 \log(p_2 w_2)$. With the cue it is $p_1 \log w_1 + p_2 \log w_2$. The biological fitness value of information — namely, the difference between log growth with and without the cue — is exactly the Shannon entropy of the environment $H = p_1 \log p_1 + p_2 \log p_2$. The payoffs w_i have dropped out. For this very simple example, the fitness value of information has nothing to do with the fitnesses obtained in different environments, but instead depends exclusively on the Shannon entropy³ of the system [7, 6].

III. EXTENDING THE MODEL

Thus far we have been looking at a very special case in which the fitness of the organism is zero when the wrong phenotype is adopted. A more realistic model would allow the possibility of non-zero fitness even when the organism develops to the wrong phenotype. Let us now assume that an organism has to make a developmental decision between n possible phenotypes, each of which is optimally adapted to one of n environments. The environments occur with probabilities p_i and the fitness to phenotype j in environment i is w_{ij} .

How should an organism respond? To maximize short-run expected payoff, an organism should simply develop the phenotype with the highest expected fitness. Expected fitness then will be $E[w] = \max_j \left[\sum_i p_i w_{ij} \right]$.

What about long-term payoff? Let us look at a strategy that produces phenotype i with probability x_i . The organism will be selected to maximize the expected log growth rate for $\sum_i x_i = 1$ so we want to find the strategy that maximizes the log growth rate $\sum_i p_i \log \sum_j w_{ij} x_j$ subject to the constraint that $\sum_i x_i = 1$. The Lagrangian for this problem is

$$L(x_1, x_2, \dots, x_n, \lambda) = \sum_i p_i \log \sum_j w_{ij} x_j - \lambda \left(\sum_i x_i - 1 \right) \quad (1)$$

Since the constraint function is a linear function, it immediately satisfies the constraint qualification that the partials of the constraint function at the constrained maximizer are not all zero. We maximize the Lagrangian by taking partial derivatives and setting to zero. The partials with respect to x_k yield a set of n equations:

$$\frac{\partial}{\partial x_k} \left[\sum_i p_i \log \sum_j w_{ij} x_j - \lambda \sum_i x_i \right] = 0 \quad (2)$$

Assuming that W (the matrix whose (i, j) entry is w_{ij}) is invertible, we can write $y_i = \sum_j w_{ij} x_j$, and $V = W^{-1}$, so that $x_j = \sum_i v_{ji} y_i$. Then we can solve

$$\frac{\partial}{\partial y_k} \left[\sum_i p_i \log \sum_j y_j - \lambda \sum_{ij} v_{ji} y_i \right] = 0 \quad (3)$$

³More generally, we should replace the Shannon entropy of the environment with the mutual information between the environment and whatever cues are available. When the cues are perfectly accurate, as in all the examples treated herein, these quantities are identical. When they are not, mutual information is the appropriate measure. Kelly [7] has shown that imperfect side information allows a gambler to increase his earnings rate by exactly the mutual information between the side information and the actual outcome. Similarly, in our model the fitness value of an imperfect cue will be exactly the mutual information between cue and environment.

for all k , which gives for all k :

$$\frac{p_k}{y_k} - \lambda \sum_j v_{jk} = 0 \quad (4)$$

Now we can solve for the constraint $\sum_i x_i = 1$, which gives $\lambda = 1$ and thus we have $y_k = p_k / \sum_j v_{jk}$. Substituting this into the equation for the log growth rate, to get the maximal log growth rate, gives

$$\sum_i p_i \log\left(\frac{p_i}{\sum_j v_{ji}}\right) = \sum_i p_i \log(p_i) - \sum_i p_i \log\left(\sum_j v_{ji}\right) \quad (5)$$

What are the values of information in each of these cases?

The log growth rate with a cue that reveals the exact environment is $\sum_i p_i \log(w_{ii})$, so that the value of the cue is

$$-\sum_i p_i \log(p_i) + \sum_i p_i \log\left(w_{ii} \sum_j v_{ji}\right) \quad (6)$$

This is exactly the Shannon entropy of the cue plus a linear transform of the probabilities p_i . We note that this equation holds only for interior solutions where the organism develops into all n phenotypes with positive probability, i.e., when all x_i in $(0, 1)$. Outside of this interior, the fitness value of information will be otherwise, as we will see in the following section.

IV. TWO ILLUSTRATIVE EXAMPLES

In this section, we consider a pair of examples in which the fitness of the organism is non-zero when it develops to the “wrong” phenotype. We start with a two-environment, two-phenotype example. Since the players have no control over the state of the environment, we can study the decision-theoretic behavior of the players without loss of generality using the following matrix where $1 > a > b$:

	Phenotype 1	Phenotype 2
Environment 1	1	b
Environment 2	a	1

If the player invests x in phenotype 1 and $1-x$ in phenotype 2, her expected log growth rate will be $p \log[x + a(1-x)] + (1-p) \log[bx + (1-x)]$. In the absence of information about which environmental state is realized, the choice of $x^*(p)$ that maximizes expected log growth given the probability p of environment 1 is:

$$x^*(p) = \begin{cases} 0 & \text{for } p \leq \frac{a(1-b)}{1-ab} \\ 1 & \text{for } p \geq \frac{1-b}{1-ab} \\ \frac{p+ab(1-p)-a}{(1-a)(1-b)} & \text{for } \frac{a(1-b)}{1-ab} < p < \frac{1-b}{1-ab} \end{cases} \quad (7)$$

If the player knows exactly the state of the environment, she will match her phenotype to the environment always, for an expected log growth rate of $\log 1 = 0$. If she does not know the state of the environment, her fitness when following strategy $x^*(p)$ will be:

$$\begin{cases} p \log[a] & \text{for } p \leq \frac{a(1-b)}{1-ab} \\ (1-p) \log[b] & \text{for } p \geq \frac{1-b}{1-ab} \\ (1-p) \log[1-p] + p \log[p] + \log[1-ab] - (1-p) \log[1-a] - p \log[1-b] & \text{Otherwise} \end{cases} \quad (8)$$

Since log growth is zero whenever the environment is known, the value of the information of a cue that reveals the environmental state is simply minus the payoff of the selected strategy when the environment is unknown.

In the central region $\frac{a(1-b)}{1-ab} < p < \frac{1-b}{1-ab}$, the fitness value of information is equal to the Shannon entropy of the environment plus a linear function of the probability of each environment. Outside the range, when the optimal strategy invests in only one of the phenotypes, the value of the cue is $-p \log[a]$ or $-(1-p) \log[b]$. This value is exactly equal to the short run value of information that one would get when maximizing the fitness in one generation.

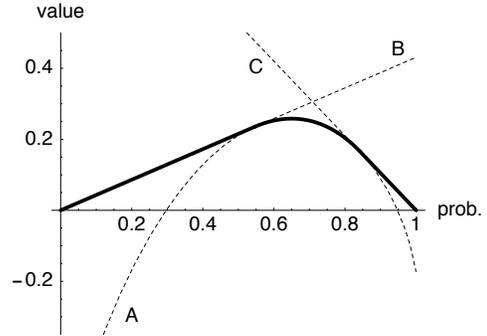


Figure 1. The value of information as a function of environmental probabilities p (heavy solid curve) is a composite of three value functions: Curve (A) is the Shannon-like curve $(1-p) \log[1-p] + p \log[p] + \log[1-ab] - (1-p) \log[1-a] - p \log[1-b]$. Curves (B) and (C) are the linear functions $p \log[a]$ and $(1-p) \log[b]$ respectively. Parameter values: $a = 0.65$, $b = 0.35$.

Simple calculus reveals that despite being composed of a linear component and a logarithmic component, the value of information curve is not only continuous but also once continuously differentiable.

To get a better intuition of how the value of information relates to the evolutionarily optimal strategy in the absence of information, we move to the case of 3 environments that occur with probabilities p_1 , p_2 , and $(1-p_1-p_2)$. While the principles generalize to larger numbers of environments and less-symmetric payoffs, three symmetric environments are far easier to represent graphically than some of the more complicated alternatives. Thus we consider the following payoff structure where $k > 1$:

	Phenotype 1	Phenotype 2	Phenotype 3
Env. 1	k	1	1
Env. 2	1	k	1
Env. 3	1	1	k

Using the approach sketched out above, we can compute the fractional investment x_1, x_2, x_3 in each strategy that maximizes long-term growth rate:

$$\begin{aligned} x_1 &= (p_1(1+w) - p_2 - p_3)/(w-1) \\ x_2 &= (p_2(1+w) - p_1 - p_3)/(w-1) \\ x_3 &= (p_3(1+w) - p_1 - p_2)/(w-1) \end{aligned} \quad (9)$$

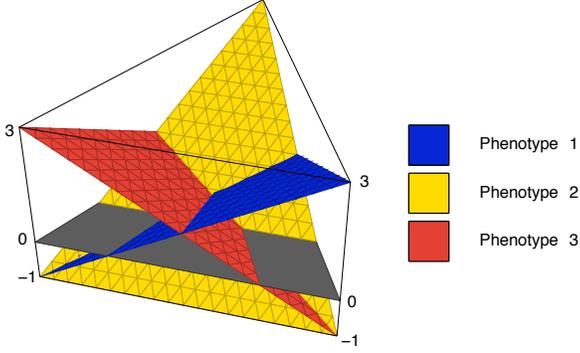


Figure 2. Fractional investment in each strategy in order to maximize long-term growth rate, displayed on the simplex $p_1 + p_2 + p_3 = 1$. Each point on the triangle corresponds to a set of environmental probabilities $\{p_1, p_2, p_3\}$ equal to where one corner represents $(1, 0, 0)$, another $(0, 1, 0)$, and the third $(0, 0, 1)$. The height of the three surfaces at any point indicates the fractional investment in each strategy at that point. Heights greater than one require corresponding bets *against* the other strategies; see text.

Here we have a curious sort of investment; the gray surface is the "invest zero" plane. When the colored surfaces drop below this, the player is effectively "betting against" the occurrence of the corresponding environment. This sort of investment may be feasible in a stock market or a horse race. But these negative bets seem to lack a biological meaning. In biological situations, we do better to look at the constrained case where the player can make only non-negative investments in each phenotype. The solution above is then only reasonable in the central region where all three bets are non-negative. This area, which we will call Region 1, is delimited by $p_i > 1/(2+w)$ for all $i = 1, 2, 3$. Outside of Region 1, we will have to compute optimal bets subject to constraints that no bet is negative. We do this below.

When one environment is sufficiently unlikely but the other two are common, an individual will invest in the phenotypes corresponding to the two common environments but the rare one. This defines three regions on the simplex given by the trio of inequalities $p_i < 1/(2+w)$, $p_j < p_k w$, and $p_k < p_j w$. In these three areas which collectively we call Region 2, investment will be given by

$$\begin{aligned} x_i &= 0 \\ x_j &= (p_j w - p_k)/(w-1) \\ x_k &= (p_k w - p_j)/(w-1) \end{aligned} \quad (10)$$

Finally, when two environments are sufficiently rare, all investment will be in the common environment. This occurs outside of the areas covered by Regions 1 and 2, in three corner areas which collectively we call Region 3.

Because of the different betting strategies in each region, the value of information in each region is computed by a different formula. We take these in turn.

In Region 1, by equation (6), a cue that indicates the state of the environment increases the expected log growth rate by

$$\log[w/(2+w)] + \sum p_i \log p_i \quad (11)$$

This is simply a constant plus the Shannon entropy of the environment.

In Region 2, let a be the phenotype never adopted by the organism. Then the cue increases the expected growth rate by

$$\log w - \sum_{i \neq a} p_i \log p_i - (1-p_a) \log[(1+w)/(1-p_a)] \quad (12)$$

In Region 3, let b be the phenotype always adopted by the organism. The cue increases expected log growth rate by

$$(1-p_b) \log w \quad (13)$$

This is simply the decision-theoretic log value of information.

Putting these all together, we get the following surface:

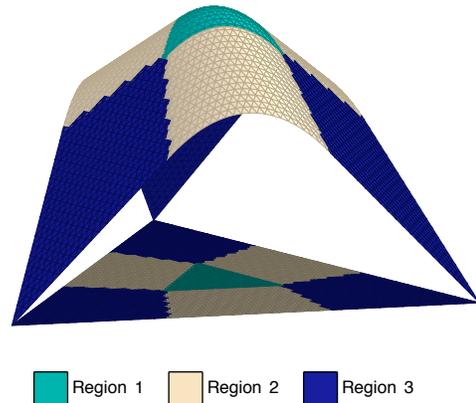


Figure 3. The value of information as a function of the probabilities that each environment occurs, for the symmetric three-environment scenario with $w = 2$, displayed on the simplex $p_1 + p_2 + p_3 = 1$.

Surprisingly, this fitness value of information surface seamlessly sews together a region described by the Shannon information (Region 1), a region described by the decision-theoretic value of information (Region 3), and an intermediate region (Region 2). Comparing the height of the surface and the gradients along the relevant edge and point boundaries, calculus reveals that this surface is again continuous and once continuously differentiable (but not twice continuously differentiable) everywhere. The fitness value of information incorporates both the Shannon measure and the decision-theoretic value — and through the continuity of the corresponding regions, we also see a fundamental connection between these two measures.

V. BOUNDING THE FITNESS VALUE OF INFORMATION

We can also show that the fitness value of information is bounded above by the Shannon entropy⁴. To do so, compare the expected log growth rate of individuals of two types. Type A individuals receive a cue x with possible values $x_1, x_2, x_3, \dots, x_n$ drawn from a distribution with probability function $P(x)$ and entropy $H(x)$. Each individual then maximizes expected log growth rate by following some investment strategy $s(x)$ that sets how to invest in the various phenotypes, given the receipt of cue x .

⁴As before, this result can be extended to replace Shannon entropy with mutual information between cue and environment

Type B individuals do not receive this cue. Instead, they follow the betting strategy $r = \sum_x P(x)s(x)$, thereby employing a probability-matching mixture of the various $s(x)$ strategies used by Type A individuals.

Represent the fitness of an individual using strategy $s(y)$ when the cue was x by $w(s(y)|x)$. The expected log growth rate for Type A individuals is then

$$\bar{w}_A = \sum_x P(x) \log[w(s(x)|x)]. \quad (14)$$

The expected log growth rate for Type B individuals is

$$\bar{w}_B = \sum_x P(x) \log \left[\sum_y P(y) w(s(y)|x) \right]. \quad (15)$$

Since fitnesses are non-negative, the $w(s(y)|x)$ terms in the summation above must be at least zero even for $y \neq x$, and therefore

$$\begin{aligned} \bar{w}_B &\geq \sum_x P(x) \log[P(x)w(s(x)|x)] \\ &= \sum_x P(x) \log[w(s(x)|x)] + \sum_x P(x) \log P(x) \\ &= \bar{w}_A - H(x). \end{aligned} \quad (16)$$

Thus the fitness value of information is at most equal to the Shannon entropy of the cue, *irrespective of the actual fitness payoffs* w .

VI. CONCLUSIONS

In this paper we have shown that two measures of information, Shannon entropy and the decision-theory value of information, are united into one single information measure when one looks at the strategies that natural selection will favor, namely those that maximize the long term growth rate of biological organisms. Furthermore, we have shown that in evolving biological systems, the fitness value of information is bounded above by the Shannon entropy. These results suggest a close relationship between biological concepts of Darwinian fitness and information-theoretic measures such as Shannon entropy or mutual information.

These results also suggest that biologists will be able to make valuable use of information theory in studying the evolution of communication. Even before knowing what a biological signal means, how it is used, or what the fitness structure of the environment may be, we have shown that one can place an upper bound on the fitness consequences of responding to that signal, simply by measuring the information content of the signalling channel.

ACKNOWLEDGMENTS

This work was supported in part by a UW RRF award to CB. The authors thank David Park for help in producing the figures.

REFERENCES

- [1] L. J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [2] I. J. Good. On the principle of total evidence. *British Journal for the Philosophy of Science*, 17:319–321, 1966.
- [3] J. P. Gould. Risk, stochastic preference, and the value of information. *Journal of Economic Theory*, 8:64–84, 1974.
- [4] D. W. Stephens. Variance and the value of information. *American Naturalist*, 134(1):128–140, 1989.
- [5] F. P. Ramsey. Weight or value of knowledge. *The British Journal for the Philosophy of Science*, 41:1–4, 1990). Published posthumously from personal notes.
- [6] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, 1991.
- [7] J. L. Kelly. A new interpretation of information rate. *Bell Systems Technical Journal*, 35:917–926, 1956.
- [8] J. Yoshimura and V. A. A. Jansen. Evolution and population dynamics in stochastic environments. *Res. Popul. Ecol.*, 38(2):165–182, 1996.
- [9] A. J. Robson, C. T. Bergstrom, and J. K. Pritchard. Risky business: Sexual and asexual reproduction in variable environments. *Journal of Theoretical Biology*, 197:541–556, 1999.
- [10] Seger and Brockmann. What is bet-hedging? *Oxford Surveys in Evolutionary Biology*, 4:182–211, 1987.