The fitness value of information

Michael Lachmann

Max Planck Inst. for Evol. Anthropology

Deutscher Platz 6

D-04103 Leipzig, Germany

email: `lachmann@eva.mpg.de`

Matina C. Donaldson-Matasci

Max Planck Inst. for Evol. Anthropology

Deutscher Platz 6

D-04103 Leipzig, Germany

email: `donaldso@eva.mpg.de`

Carl T. Bergstrom

Department of Biology

University of Washington

Seattle, WA 98195-1800, USA

email: `cbergst@u.washington.edu`

February 19, 2009

1 Introduction

Information is a central organizing concept in our understanding of biological systems at every scale, from the digital coding in DNA to the long-range calls of cetaceans [1]. In some fields of biology, such as neurobiology, information is naturally and usefully measured using the information-theoretic quantities of entropy and mutual information [2, 3, 4]. However, information theoretic measures have seen substantially less use in evolutionary biology, behavioral ecology, and related areas. Why is this? One problem is that measures of entropy do not directly address information quality; they do not distinguish between relevant and irrelevant information. When we think about fitness consequences we care very much about relevant and irrelevant information. For example, from an information-theoretic standpoint one has the same amount of information if one knows the timing of sunrise on Mars as one has if one knows the timing of sunrise on Earth. Yet individuals of few if any Earthbound species find the timing of sunrise on Mars relevant to their survival. Information measures based on entropy have therefore been deemed irrelevant to the evolutionary ecology of information. Instead evolutionary biologists and behavioral ecologists tend to focus on decision-theoretic measures such as the expected value of perfect information or the expected value of sample information[5, 6, 7, 8, 9], with value often measured in terms of fitness consequences [10, 11, 12].

The disconnect between information-theoretic and decision-theoretic measures is perplexing. Entropy and mutual information appear to measure information quantity while reflecting nothing about fitness consequences; the expected value of information measures fitness consequences but has nothing to do with the actual length or information quantity of a message. But early work in population genetics [13, 14, 15] and recent analyses of evolution in fluctuating environments [16, 17] hint at a possible relation between information and fitness. What is this relation? Information theorists since Kelly [18] have observed that in special circumstances, information value and information-theoretic measures may be related. Here we argue that these special circumstances are exactly those about which biologists should be most concerned: they include the context of evolution by natural selection. We present a simple model of evolution in an uncertain environment, and calculate the increase in Darwinian fitness that is made possible by responding

to a cue conveying information about the environmental state. We show that in certain cases this *fitness value of information* is exactly equal to the mutual information between the cue and the environment. More generally, we find that this mutual information, which seemingly fails to take anything about organismal fitness into account, nonetheless imposes an upper bound on the fitness value of information.

2   Two measures of information

Environmental cues can help organisms living in an uncertain environment predict the future state of the environment, and thereby can allow them to choose an appropriate phenotype for the conditions that they will face. We will consider a population of annual organisms living in a variable environment. The environmental state E and the environmental cue C are correlated random variables independently drawn every year; both are common to all individuals, so that all individuals encounter the same conditions in a given year.

We might measure the information conveyed by the environmental cue C in two different ways. The typical approach in statistical physics, communication engineering, neurobiology, and related fields is to use an information-theoretic measure such as the mutual information. The mutual information describes the extent to which a cue reduces the uncertainty about the environment, measured in terms of entropy. Following [4], we define the *entropy* of the random variable E representing the environmental state as

$$H(E) = -\sum_e p(e) \log p(e),$$

where $p(e)$ is the probability of observing the state e. Similarly, we define the *conditional entropy* of the environment E, once the random variable C representing the cue has been observed, as

$$H(E|C) = -\sum_c p(c) \sum_e p(e|c) \log p(e|c),$$

where $p(c)$ is the probability of observing cue c, and $p(e|c)$ is the conditional probability that the environment is in state e, given that cue c has been observed. Throughout the paper, we will take logarithms to base 2, as is the convention in information theory.

Definition: The *mutual information* between a cue C and a random environmental state E measures how much the cue reduces the uncertainty about the state E: thus $I(E;C) = H(E) - H(E|C)$.

Another approach, common in decision theory, economics, behavioral ecology, and related fields, is to look at the expected value of information: how information improves the expected payoff to a decision-maker [8]. We write the maximal fitness obtainable without a signal as $F(E) = \max_x \sum_e p(e) f(x, e)$, where x is a strategy, and $f(x, e)$ is the fitness of that strategy when the environmental state e occurs. Similarly, we write the optimal fitness attainable with a signal $F(E|C) = \sum_c p(c) \max_{x_c} \sum_e p(e|c) f(x_c, e)$ where x_c is a strategy used in response to the cue c.

Definition: The *decision-theoretic value of information* that a cue C provides about the state of the world E is defined as the difference between the maximum expected payoff or fitness that a decision-maker can obtain by conditioning on C, and the maximum expected payoff that could be obtained without conditioning on C. This is written as $\Delta F(E;C) = F(E|C) - F(E)$.

2.1 An illustrative example

To illustrate the difference between the two measures of information described above, we start with a simplified model. The environment has two possible states, such as wet and dry years, and each organism has two possible phenotypes into which it can develop. Fitness depends on the match of phenotype to environment, as follows:

	Phenotype φ_1	Phenotype φ_2
Environment e_1	5	1
Environment e_2	1	3

Given that the probability of environmental state e_1 is p, and of state e_2 is $(1 - p)$, what should these individuals do in the absence of information about the condition of the environment? The organism maximizes its single-generation expected fitness by developing into phenotype φ_1 if

$p > 1/3$, and into phenotype φ_2 otherwise. This is optimal because $5 \cdot p + 1 \cdot (1-p) > 1 \cdot p + 3 \cdot (1-p)$ when $p > 1/3$. The payoff earned with this strategy would be $F(E) = \max[5p + (1-p), p + 3(1-p)]$.

$$F(E) = \begin{cases} 5p + (1-p) & \text{for } p \geq \frac{1}{3} \\ p + 3(1-p) & \text{for } p < \frac{1}{3} \end{cases} \tag{1}$$

Now we suppose that there is a perfectly informative environmental cue C which accurately reveals the state E of the environment. How do we measure the information provided by this cue?

Within an information-theoretic framework, we measure the amount of information in the cue by calculating mutual information between the cue and the environment; this measures how much the cue tells us about the environment. Since the cue is perfectly informative, the conditional uncertainty about the environment once the cue is observed is zero: $H(E|C) = 0$. The mutual information is therefore

$$I(E; C) = H(E) - H(E|C) = H(E) = -p \log p - (1-p) \log(1-p). \tag{2}$$

Figure 1A plots the mutual information between cue and environment as a function of the probability p that environment e_1 occurs.

Within a decision theoretic framework, we measure the value of the information in the cue by calculating how much the ability to detect the cue improves the expected fitness of the organism. When an organism receives the cue, it can always develop the appropriate phenotype for the environment, and thus obtain a single-generation expected fitness of $F(E|C) = 5p + 3(1-p) = 3 + 2p$. The decision-theoretic value of information in this case is therefore

$$\begin{aligned} \Delta F(E; C) &= F(E|C) - F(E) \\ &= \begin{cases} 2 - 2p & \text{for } p \geq \frac{1}{3} \\ 4p & \text{for } p < \frac{1}{3} \end{cases} \end{aligned} \tag{3}$$

By using the cue, the organism increases its single-generation expected fitness by the decision-theoretic measure of the value of information. This quantity, illustrated in Figure 1B, differs considerably from the information-theoretic measure of mutual information shown in Figure 1A. Not only do the graphs take on different forms, but their units of measurement differ. Mutual information is measured in bits, whereas the value of information is measured in fitness units.
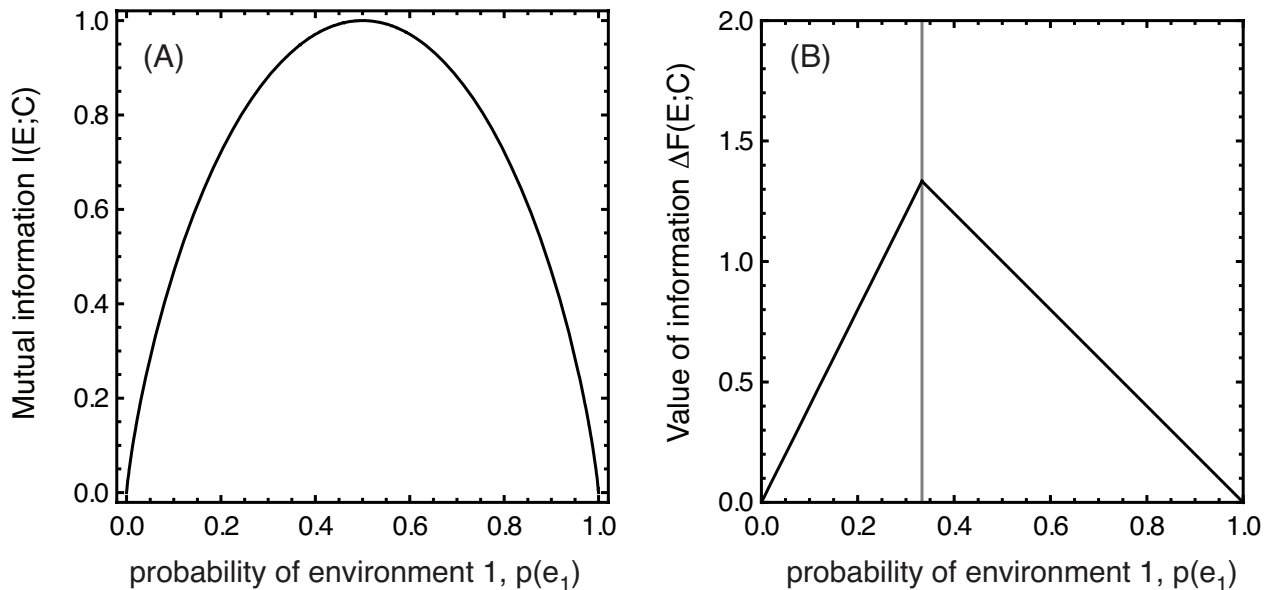
Figure 1: Two different measures of the information in a cue are commonly used. (A) The mutual information between the cue and the environment measures the reduction in environmental uncertainty once the cue has been observed. (B) The decision-theoretic value of information measures the change in expected fitness that is made possible by using the cue.

3 The fitness value of information

In the previous section, we measured the value of information by its effect on expected fitness over a single generation. But as many authors have shown [19, 20, 21, 22, 23, 24], organisms will not always be selected to use a strategy that maximizes their fitness in a single generation. Instead, a better proxy for the likely outcome of evolution is to think of organisms as maximizing the long-term growth rate of their lineage. This distinction is critical when the environment changes from one generation to the next, and affects all individuals within one generation in the same way. Under these circumstances, maximizing the long-term growth rate over a very large number of generations is equivalent to maximizing the expected value of the logarithm of the fitness in a single generation.

From an evolutionary perspective, it therefore makes sense to define the value of a cue not in terms of single-generation fitness consequences, but rather in terms of the increase in long-term growth rate it makes possible. Let $g(x) = \sum_e p(e) \log f(x, e)$ be the expected long-term growth rate of a strategy x. We write the maximum long-term growth rate obtainable without a

cue as $G(E) = \max_x \sum_e p(e) \log f(x, e)$. Similarly, we write the maximum long-term growth rate attainable with a cue as $G(E|C) = \sum_c p(c) \max_{x_c} \sum_e p(e|c) \log f(x_c, e)$.

> **Definition:** The *fitness value of information* $\Delta G(E; C)$ associated with a cue or signal is the greatest fitness decrement or cost that would be favored by natural selection in exchange for the ability to detect and respond to this cue: $\Delta G(E; C) = G(E|C) - G(E)$.

3.1 Proportional betting

To explore the connection between environmental uncertainty and long-term growth rate, we will first look at an even simpler example, where the organism survives only if it matches its phenotype to the environment perfectly.

	Phenotype φ_1	Phenotype φ_2
Environment e_1	7	0
Environment e_2	0	7/2

In the short run, individuals maximize expected fitness by employing the highest-payoff phenotype only. But we can immediately see that the long-run fitness of lineage is not maximized in the same way: playing only one strategy will inevitably lead to a year with zero fitness and consequent extinction for the lineage. Thus organisms will be selected to *hedge their bets*, randomizing which of the two phenotypes they adopt [25, 26].

If environment e_1 occurs with probability p and environment e_2 occurs with probability $(1-p)$, with what probability should an individual adopt each phenotype? As we consider a larger and larger span of generations, a larger and larger fraction of the probability is taken up by "typical sequences" of environments, in which environment e_1 occurs around Np times, and environment e_2 occurs around $N(1-p)$ times [4]. A strategy that maximizes the growth rate over these typical sequences will, with very high probability, be the one that is observed as the result of natural selection [27]. To find this genotype, let us assume a genotype that develops with probability x into phenotype φ_1 and with probability $(1-x)$ into phenotype φ_2; the population growth over a typical sequence of N events will be $(7x)^{Np}(7/2(1-x))^{N(1-p)}$. Maximizing the population growth

is equivalent to maximizing the per-generation exponent of growth, or the log of the expression divided by N:

$$\begin{aligned} g(x) &= p\log(7x) + (1-p)\log(7/2\,(1-x)) \\ &= p\log(x) + (1-p)\log(1-x) \,+\, p\log(7) + (1-p)\log(7/2). \end{aligned} \quad (4)$$

The only part of this equation that depends on x is $p\log(x)+(1-p)\log(1-x)$, so any dependence on the fitnesses when the organism properly matches the environments (i.e. on the values "7" and "7/2") has dropped out. The maximum occurs when $x = p$, a strategy called *proportional betting*. With this strategy, organisms develop into the two phenotypes in proportion to the probabilities of the two environments — and these optimal proportions do not depend on the fitness benefits [4]. The optimal growth rate is thus

$$\begin{aligned} G(E) &= p\log(p) + (1-p)\log(1-p) \,+\, p\log 7 + (1-p)\log(7/2) \\ &= p\log 7 + (1-p)\log(7/2) - H(E). \end{aligned}$$

$$(5)$$

3.2   Uncertainty and optimal growth

To understand the connection between environmental uncertainty and optimal growth more fully, it is instructive to generalize the simple model above to include several different environments and phenotypes. Let us assume that there are n environments, and that for each environment there is one optimal phenotype. The payoff for phenotype φ_e in environment e is d_e, and the payoff for any other phenotype in environment e is 0. Let the probability of environment e occurring be $p(e)$, and the probability of developing into phenotype φ_e be $x(e)$. Then the growth rate of the lineage over a typical sequence of environments, in which environment e occurs approximately $Np(e)$ times, is $\prod_e (d_e x(e))^{Np(e)}$. Again, instead of maximizing the above expression, we can maximize its log, divided by N:

$$\begin{aligned} g(x) &= \sum_e p(e)\log(d_e x(e)) \\ &= \sum_e p(e)\log d_e + \sum_e p(e)\log x(e) \end{aligned} \quad (6)$$

The left part does not depend on the strategy x, so we just need to maximize the right part, which as before is independent of the fitness values d_e in the different environments. But instead of simply giving the solution, let us rewrite the above expression as

$$
\begin{aligned}
g(x) &= \sum_e p(e) \log d_e + \left(\sum_e p(e) \log p(e) - \sum_e p(e) \log p(e) \right) + \sum_e p(e) \log x(e) \\
&= \sum_e p(e) \log d_e + \sum_e p(e) \log p(e) - \sum_e p(e) \log \frac{p(e)}{x(e)} \\
&= \sum_e p(e) \log d_e - H(E) - D_{KL}(p||x). \tag{7}
\end{aligned}
$$

The term $D_{KL}(p||x) \equiv \sum_e p(e) \log(p(e)/x(e))$ is the Kullback-Leibler divergence (K-L divergence) between the distribution of environments and the distribution of phenotypes. The K-L divergence, also known as the *relative entropy*, quantifies how greatly a given distribution (here $x(e)$) varies from a reference distribution (here $p(e)$). In coding theory, for example, the K-L divergence between two distributions P and X tells us how much longer a message will be on average if use a code optimized for distribution X when the true distribution is P. The K-L divergence is minimized when $x(e) = p(e)$ for all e, in which case $D_{KL}(p||x) = 0$.

The representation of the growth rate given in Equation 7 allows us to generalize and interpret the results of the previous section. We note that the term $D_{KL}(p||x)$ is the only part of the expression for $g(x)$ that depends on the organism's choice of strategy; this is zero when $x(e) = p(e)$. Therefore the optimal assignment of phenotypes will be a proportional betting strategy, and and will achieve a growth rate of

$$
G(E) = \sum_e p(e) \log d_e - H(E), \tag{8}
$$

similar to Equation 5 above. When the distribution of phenotypes is $x(e)$ instead of the optimal $p(e)$, then the growth rate per generation will be reduced by $D_{KL}(p||x)$.

Furthermore, if the organism knew exactly what the environment would be at every generation, it would choose the optimal phenotype every time. The growth rate would then be $\prod_e (d_e)^{Np(e)}$, corresponding to an average log growth rate of $\sum_e p(e) \log d_e$. Comparing this quantity to Equation 8, we see that the growth rate associated with the proportional betting strategy is equal to the growth rate that could be achieved with full information, minus the entropy of the unknown environmental state. At least for this special case, environmental uncertainty reduces the log growth

rate by an amount equal to the entropy of the environment!

3.3   The illustrative example, revisited

In the previous two sections, we assumed that organisms can only survive when they choose exactly the right phenotype for the environment. In general, however, choosing the wrong phenotype need not be fatal. How does this affect the optimal growth rate, and the fitness value of information?

Let us return to the first example presented in Section 2.1; in this example the wrong phenotype had a fitness of 1 instead of 0. Now the expected log growth rate in the absence of a cue is $g(x) = p \log(4x + 1) + (1 - p) \log(3 - 2x)$ and here, organisms will not follow a strict proportional-betting strategy. Instead, the optimal strategy x^* will be to always develop into a single phenotype when p is near 0 or 1, and to hedge bets when p takes on an intermediate value:

$$x^*(p) = \begin{cases} 0 & \text{for } p \leq \frac{1}{7} & \text{always phenotype } \varphi_2 \\ \frac{7p-1}{4} & \text{for } \frac{1}{7} < p < \frac{5}{7} & \text{bet-hedging between 1 and 2} \\ 1 & \text{for } \frac{5}{7} \leq p & \text{always phenotype } \varphi_1 \end{cases} \tag{9}$$

This yields a growth rate without the cue of

$$G(E) = \begin{cases} (1 - p) \log 3 & \text{for } p \leq \frac{1}{7} \\ p \log 7 + (1 - p) \log \frac{7}{2} - H(E) & \text{for } \frac{1}{7} < p < \frac{5}{7} \\ p \log 5 & \text{for } p \geq \frac{5}{7} \end{cases} \tag{10}$$

If organisms could sense a cue that perfectly revealed the state of the environment, all individuals could develop as the phenotype which matches the environment, and we would see a log growth rate of $G(E|C) = p \log(5) + (1 - p) \log(3)$. The difference in expected log fitness — and thus the fitness value of information — is given by:

$$\Delta G(E; C) = \begin{cases} p \log 5 & \text{for } p \leq \frac{1}{7} \\ H(E) + p \log 5/7 + (1 - p) \log 6/7 & \text{for } \frac{1}{7} < p < \frac{5}{7} \\ (1 - p) \log 3 & \text{for } p \geq \frac{5}{7} \end{cases} \tag{11}$$

10

In the central region, where $\frac{1}{7} < p < \frac{5}{7}$, the fitness value of information is equal to the entropy of the environment plus a linear function of the probability of each environment: $p \log 5/7 + (1 - p) \log 6/7$. Outside the range, when the optimal strategy invests in only one of the phenotypes, the value of the cue depends linearly on the probability of each environment: $p \log 5$ or $(1 - p) \log 3$. Since in this region the strategy used is identical to the strategy optimizing fitness in a single generation, the value of the cue is just like the difference in long term growth rate when optimizing over only a single generation. Calculus reveals that the function is continuous and once continuously differentiable everywhere. The fitness value of information for the cue appears to be related to the mutual information between the cue and environment. In the next sections, we present a general model that will allow us to quantify and interpret that relationship.

3.4   Effective proportional betting

In order to explore the connection between fitness value and information, we will develop a general model based on the one presented in Section 3.2, but which relaxes the assumption that developing the wrong phenotype is always fatal. Let us again assume that an organism has to make a developmental decision between n possible phenotypes, each of which is a best match to one of n environments. Each environmental state e occurs with probability $p(e)$, and the fitness of phenotype φ in environment e is $f(\varphi, e)$. The best match for environment e is phenotype φ_e; that is, $\max_\varphi[f(\varphi, e)] = f(\varphi_e, e)$. The strategy x defines the probability $x(\varphi)$ that an individual will develop into phenotype φ. We want to find the strategy that maximizes the expected long-term growth rate, $g(x) = \sum_e p(e) \log \sum_\varphi f(\varphi, e) x(\varphi)$.

In a previous paper [28], we introduced a method for calculating optimal bet-hedging strategies that will prove useful in the present analysis. What follows is a very brief outline of the method; full details are provided in the previous paper.

We first define a set of hypothetical "extremist" phenotypes which fit the model of Section 3.2, so that each hypothetical phenotype φ'_e is ideally adapted to one environment e, where it has fitness d_e, but fails to survive in any other environment. We next aim to describe each actual phenotype φ as a bet-hedging strategy combining the hypothetical phenotypes φ'_e. That is, we

would like to find fitnesses d_e for the extremist phenotypes and mixing strategies $s(e|\varphi)$ across extremist phenotypes such that

$$f(\varphi, e) = s(e|\varphi)d_e \tag{12}$$

for all environments e and phenotypes φ. For each phenotype φ, the strategy $s(e|\varphi)$ describes the bet-hedging proportions of hypothetical phenotypes φ'_e that would produce the same fitness, measured separately for each environment. This problem is equivalent to defining the fitness matrix $\mathbf{F}$, with entries $F_{\varphi e} = f(\varphi, e)$, as a product of two unknown matrices: $\mathbf{S}$, with entries $S_{\varphi e} = s(e|\varphi)$, and $\mathbf{D}$, with diagonal entries $D_{ee} = d_e$ and 0 elsewhere. Solving for these two matrices is straightforward and can almost always be done uniquely [28].

The advantage of this approach is that it will allow easy comparison to the simplified model presented in Section 3.2, which highlights the connection between growth rate and uncertainty. We can express a strategy x as a row vector $\overrightarrow{x}$, with each element $x_\phi = x(\phi)$ representing the probability of developing phenotype ϕ. To describe the strategy's fitness in any particular environment e, we need simply look at the eth element of the vector $\overrightarrow{x}\mathbf{F}$:

$$f(x, e) = [\overrightarrow{x}\mathbf{SD}]_e = [\overrightarrow{x}\mathbf{S}]_e d_e = y(e)d_e. \tag{13}$$

This means that the strategy x, which produces each phenotype φ with probability $x(\varphi)$, is exactly equivalent to a strategy y that produces each hypothetical phenotype φ'_e with probability $y(e) = [\overrightarrow{x}S]_e$. We can write down the long-term growth rate for a lineage that uses strategy x by calculating the growth rate for the equivalent strategy y:

$$
\begin{aligned}
g(x) &= \sum_e p(e) \log f(x, e) \\
&= \sum_e p(e) \log y(e)d_e \\
&= \sum_e p(e) \log d_e - H(E) - D_{KL}(p||y). \tag{14}
\end{aligned}
$$

This equation is very similar to Equation 7, except that instead of measuring the Kullback-Leibler divergence of the strategy x from the environmental distribution, we measure the divergence of the effective strategy y from the environmental distribution. The maximum growth rate that can be achieved occurs when $y(e) = [\overrightarrow{x}S]_e = p(e)$, in which case $D_{KL}(p||y) = 0$ (see Figure 2A.) If

12

there is no strategy x that can achieve this, then the strategy that minimizes the Kullback-Leibler divergence is optimal (see Figure 2B.) If we think of the effective strategy y as representing the effective bets the strategy is placing on each environment, then we see that the optimal strategy effectively does proportional betting—or as close as it can get.

The optimal growth rate for a strategy that effectively does proportional betting is therefore

$$G(E) = \sum_e p(e) \log d_e - H(E) \tag{15}$$

just as it was for a diagonal fitness matrix (Equation 8.) Thus, even when choosing the wrong phenotype is not fatal, the optimal growth rate is limited by the entropy of the environment. However, the first term is no longer the growth rate that could be achieved if individuals could predict the environment perfectly. Instead, it is the growth rate that could be achieved if individuals could predict each environment e perfectly, and instead of using the actual phenotype φ_e with fitness $f(\varphi_e, e)$, they could use the higher-fitness hypothetical phenotype φ'_e with fitness d_e. The value of perfect information is therefore the reduction in entropy it facilitates, $H(E)$, plus a negative term that reflects the fitness cost due to the fact that individuals are in practice restricted to the actual phenotypes rather than the hypothetical ones, $\sum_e p(e) \log f(\varphi, e)/d_e$ (for example, see Equation 11.)

3.5   Information and fitness value

Until now, we have considered only cues that allow individuals to predict the state of the environment perfectly. We would now like to calculate the value of a partially informative cue.

All individuals within a generation observe the same environmental cue c, which occurs with probability $p(c)$. Once that cue has been observed, the probability of each environmental state is given by the conditional probability distribution, $p(e|c)$. A conditional strategy x specifies the probability of developing into each phenotype φ, after observing the cue c: $x(\varphi|c)$. This can be represented as a matrix $\mathbf{X}$, with entries $X_{c\varphi} = x(\varphi|c)$. To describe the strategy's fitness in a particular environment e, after a cue c has been observed, we can look at the c-th row and the e-th column of the matrix $\mathbf{XF}$:

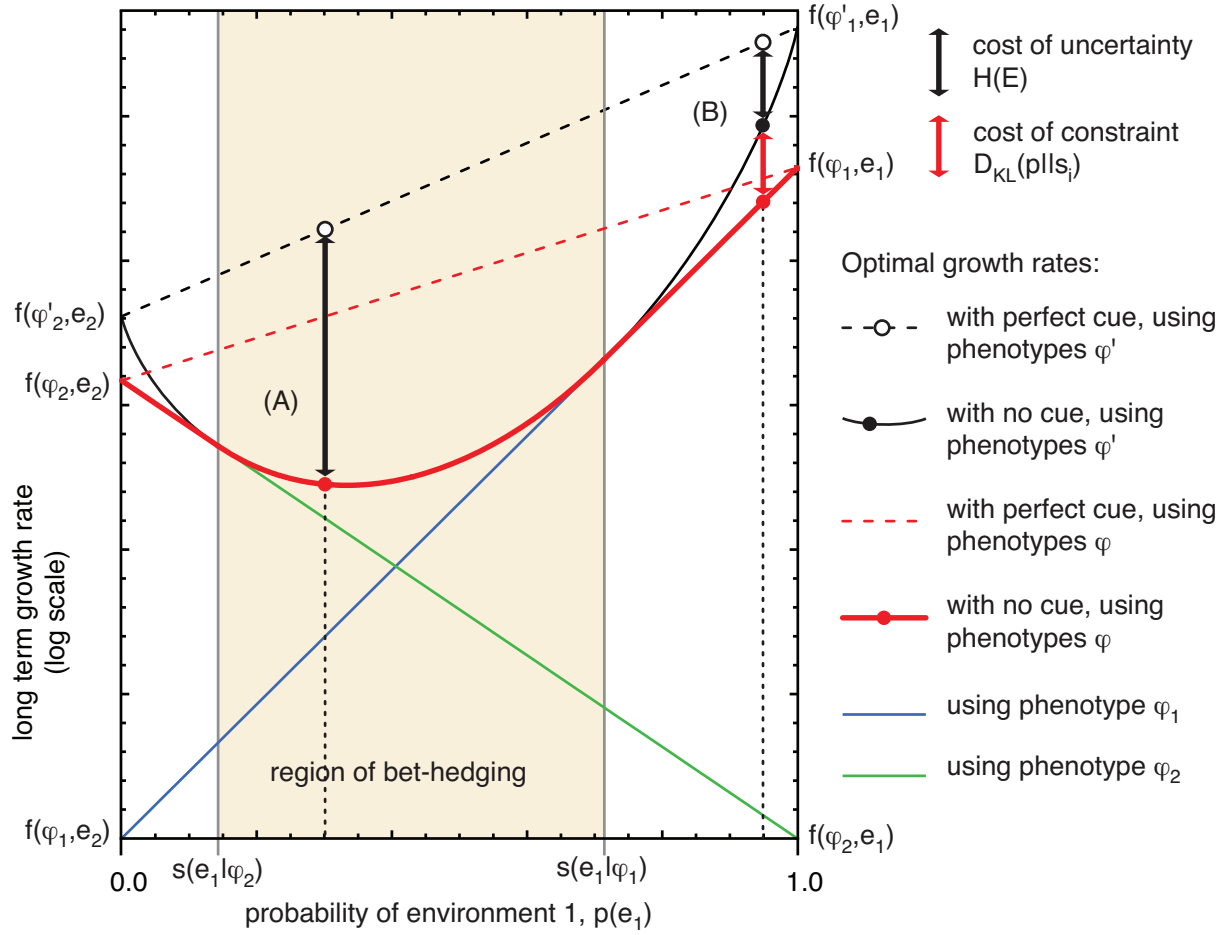$$f(x, e) = [\mathbf{XSD}]_{ce} = [\mathbf{XS}]_{ce} d_e = y(e|c) d_e. \tag{16}$$

Figure 2: Maximal growth rate when choosing the wrong phenotype is not fatal. The red solid line indicates the maximal growth rate using the two phenotypes φ_1 and φ_2. The black solid line is the growth rate that could be achieved with an unconstrained strategy, using the hypothetical phenotypes φ'_1 and φ'_2. Since the hypothetical phenotypes are fatal in the wrong environment, the optimal unconstrained strategy always uses proportional betting. In contrast, the optimal constrained strategy bet-hedges only in an intermediate range of environmental frequencies, labeled the "region of bet-hedging". Outside that region, it uses the single phenotype that does best on average. (A) Within the region of bet-hedging, the constrained strategy does just as well as the unconstrained strategy. Compared to an unconstrained strategy that can perfectly predict the environment, both strategies incur a cost of uncertainty equal to the entropy of the environment $H(E)$ (see Equations 8 and 15.) (B) Outside the region of bet-hedging, the constrained strategy does worse than the constrained strategy, since it cannot bet-hedge anymore. It always uses phenotype φ_1 on the right of the region, where environment e_1 is more common, and φ_2 on the left. The growth rate achieved is therefore exactly as with the decision-theoretic strategy, optimizing fitness in just one generation. In this case, the constrained strategy pays not only the cost of uncertainty, $H(E)$, but also a cost of constraint that arises from the inability to bet strongly enough on the most common environment. This constraint further reduces the growth rate by the Kullback-Leibler divergence $D_{KL}(p||s_i)$, which gets larger as we get farther from the boundary of the region of bet-hedging (see Equation 14.)

14

This shows that a conditional strategy y which produces the hypothetical phenotype φ'_e with probability $y(e|c) = [\mathbf{XS}]_{ce}$, conditional on observing the cue c, is exactly equivalent to the conditional strategy x. The growth rate of the strategy x can therefore be written as:

$$
\begin{aligned}
g(x) &= \sum_c p_c \sum_e p(e|c) \log \sum_\varphi x(\varphi|c) f(\varphi, e) \\
&= \sum_c p_c \sum_e p(e|c) \log y(e|c) d_e \\
&= \sum_e p(e) \log d_e - H(E|C) - D_{KL}(p(e|c)||y(e|c)), \quad (17)
\end{aligned}
$$

which is like a conditional version of Equation 14. Instead of the uncertainty of the environment $H(E)$, we have the conditional uncertainty after observing a cue, $H(E|C)$. Instead of the relative entropy $D_{KL}(p||y)$, we measure the conditional relative entropy $D_{KL}(p(e|c)||y(e|c))$, which reflects the difference between the bets the strategy effectively places on environments and the environmental probabilities, conditional on which cue is observed. As usual, the best strategy is effective proportional betting, conditional on the cue, but this may not always be possible.

What is the fitness value of the cue C? First of all, consider the situation where a bet-hedging strategy can effectively do proportional betting, both without the cue and with each possible cue. Then the Kullback-Leibler divergence terms in Equations 14 and 17 are always zero. We can therefore write:

$$
\begin{aligned}
\Delta G(E;C) &= G(E|C) - G(E) \\
&= \left(\sum_e p(e|c) \log d_e - H(E|C) \right) - \left(\sum_e p(e) \log d_e - H(E) \right) \\
&= H(E) - H(E|C) = I(E;C). \quad (18)
\end{aligned}
$$

The value of receiving a cue—when effective proportional betting is possible—is exactly the mutual information between the cue and the environment (see Figure 3.)

Now let us consider the more general situation, where effective proportional betting may not be possible. Let y^* be the best possible effective betting strategy when no cue is available, and let y^*_C be the best possible effective betting strategy when the cue C is available. The fitness value of information is then

$$
\begin{aligned}
\Delta G(E;C) &= G(E|C) - G(E) \\
&= I(E;C) - \left(D_{KL}(p(e|c)||y^*_C(e|c)) - D_{KL}(p||y^*) \right). \quad (19)
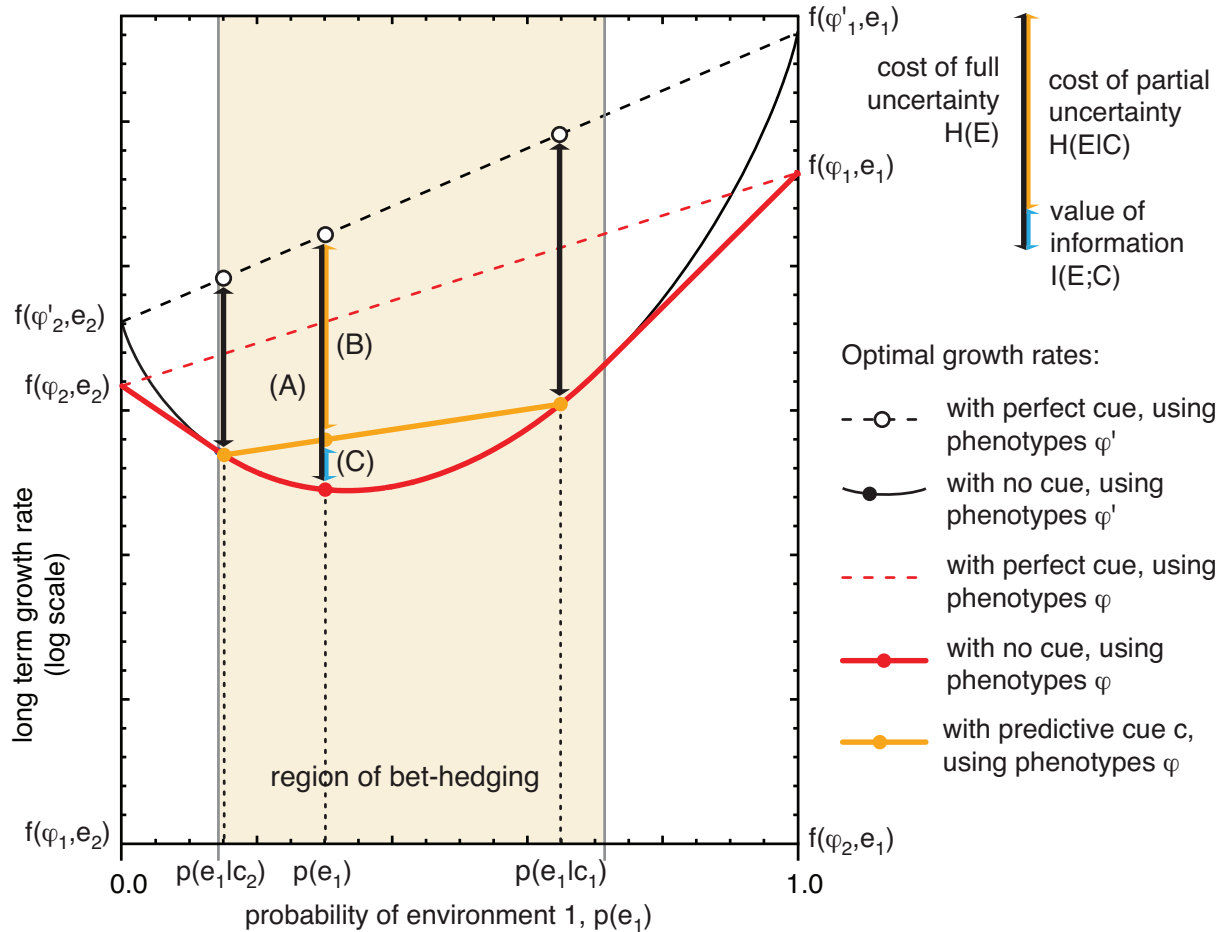\end{aligned}
$$

15

Figure 3: When the optimal strategy is to bet-hedge, both with and without a cue, the fitness value of information is equal to the mutual information between the cue and the environment. We calculate the value of a partially informative cue by looking at the reduction in growth rate, as in Figure 2, relative to a perfectly informed, unconstrained strategy. (A) With no cue at all, the cost of uncertainty is equal to the entropy of the environment $H(E)$. (B) Once a particular cue c_i has been observed then the reduction in growth rate is just the cost of uncertainty, $H(E|c_i)$. Averaging across the different cues, the reduction in growth rate for a strategy using a partially informative cue is simply the conditional entropy $H(E|C)$. (C) The fitness value of information is, in this case, the amount by which the cue reduces uncertainty about the environment — that is, exactly the mutual information between the cue and the environment (see Equation 18.)

We would like to show that the mutual information $I(E;C)$ is an upper bound for the fitness value of information $\Delta G(E;C)$. That means we need to show that the right-hand term is never negative: the cost of constraining the unconditional strategy cannot be greater than the cost of constraining the conditional strategy. We'll do this in two steps. First of all, we define the unconditional strategy $y_C^*(e) = \sum_c p(c) y_C^*(e|c)$ as the strategy an observer would see, watching someone play $y_C^*(e|c)$ in response to the cues c, but without observing the cues. The first step is to show that this marginal strategy can be no farther from the marginal distribution of environments than the conditional strategy is from the conditional distribution of environments. We can write the Kullback-Leibler divergence between the two joint distributions over cues and environments in two different ways:

$$
\begin{aligned}
D_{KL}(p(e,c)||y_C^*(e,c)) &= D_{KL}(p(c|e)||y_C^*(c|e) + D_{KL}(p(e)||y_C^*(e)) \\
&= D_{KL}(p(e|c)||y_C^*(e|c) + D_{KL}(p(c)||y_C^*(c)).
\end{aligned}
\tag{20}
$$

However, the marginal distribution over cues is the same for the two distributions, because y_C^* is defined in terms of the way it responds to cues generated according to the distribution p. This means the last term is zero, so

$$
\begin{aligned}
D_{KL}(p(e|c)||y_C^*(e|c) &= D_{KL}(p(c|e)||y_C^*(c|e) + D_{KL}(p(e)||y_C^*(e)) \\
D_{KL}(p(e|c)||y_C^*(e|c) - D_{KL}(p(e)||y_C^*(e)) &= D_{KL}(p(c|e)||y_C^*(c|e) \geq 0
\end{aligned}
\tag{21}
$$

as desired. Finally, we note that since y^* is defined as the optimal unconditional strategy for the environmental distribution p, $D_{KL}(p(e)||y^*(e)) \leq D_{KL}(p(e)||y_C^*(e))$. This shows that the fitness value of a cue cannot exceed the mutual information between that cue and the environment:

$$
\begin{aligned}
\Delta G(E;C) &= I(E;C) - (D_{KL}(p(e|c)||y_C^*(e|c)) - D_{KL}(p||y^*)) \\
&\leq I(E;C) - (D_{KL}(p(e|c)||y_C^*(e|c)) - D_{KL}(p||y_C^*)) \\
&\leq I(E;C).
\end{aligned}
\tag{22}
$$

Figure 3 illustrates a case where the fitness value of the cue is exactly equal to the information it conveys; Figure 4 illustrates two cases in which the fitness value of the cue is strictly less than the information it contains.
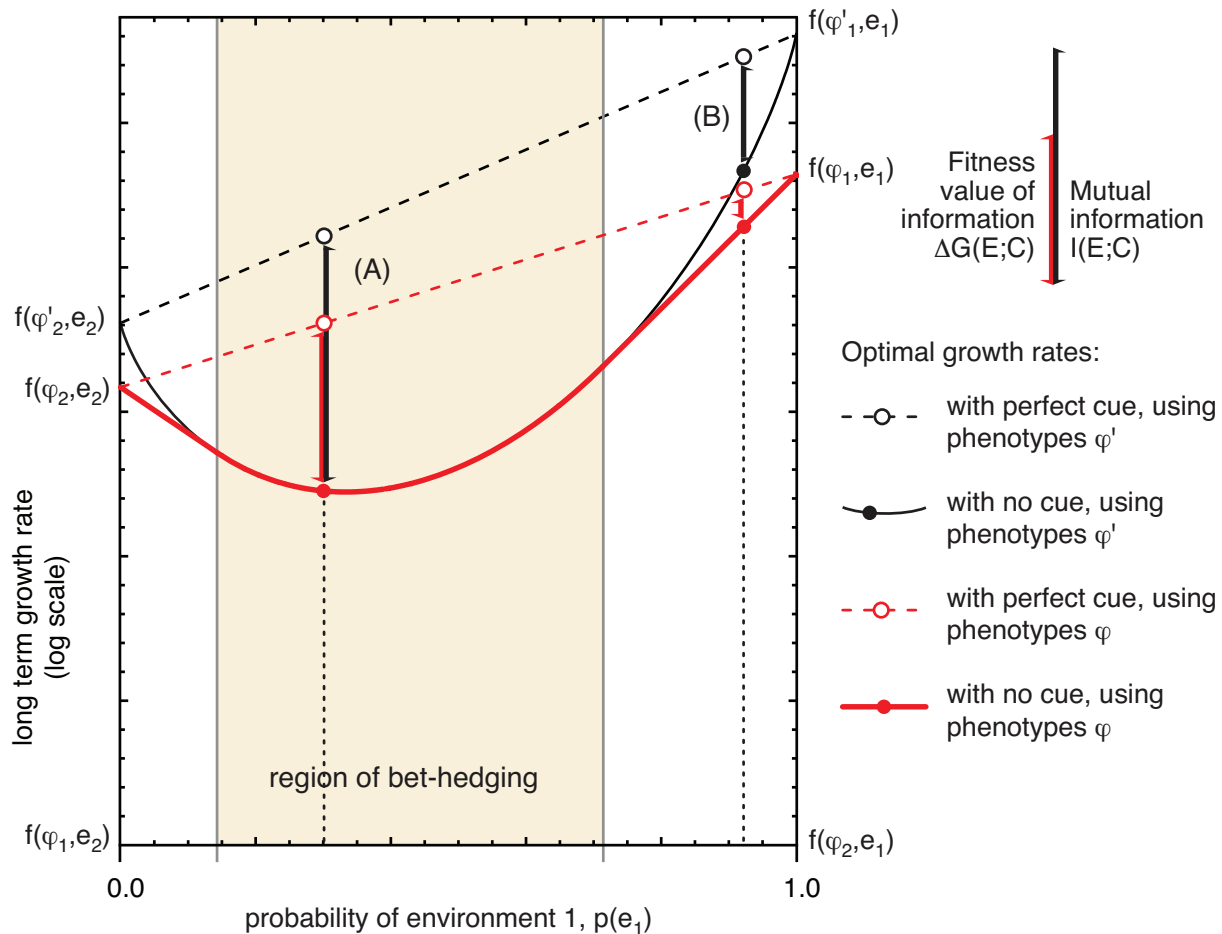
Figure 4: The mutual information between a cue and the environment is an upper bound on the fitness value of that information. For an unconstrained strategy, using the extremist phenotypes φ', the value of a cue is exactly equal to the information it conveys. We illustrate two cases of a constrained strategy where the value of information is strictly less than the amount of information. (A) For example, say that the optimal strategy without a cue would be to bet-hedge; the optimal response to a perfectly informative cue would be to choose the single best phenotype φ_1 or φ_2. For an unconstrained strategy, the value of this perfect cue would be equal to the mutual information. The fitness value for the constrained strategy is lower, because although it can achieve just the same growth rate as the unconstrained strategy without information, once information is available the unconstrained strategy can do better. (B) If there is no bet-hedging even without a cue, then the constrained strategy does worse both with and without the cue. The value of the cue using a constrained strategy is thus not directly comparable to the value of the cue when using an unconstrained strategy. However, we prove in the text that the difference in growth rates for the constrained strategy cannot exceed the difference in growth rates for the unconstrained strategy (see Equation 22.)

4 Discussion

Organisms can use bet-hedging to deal with fluctuating environmental conditions — or they can use information from cues that correlate with the environment. In this paper, we have shown that the fitness value of using information in this way is intimately related to an information-theoretic measure, the mutual information between cue and environment. Under appropriate circumstances, the fitness benefit of being able to detect and respond to that cue is exactly equal to this mutual information. More generally, the mutual information between cue and environment provides an upper bound on the fitness value of responding to the cue. These results are surprising, in that the mutual information measure seemingly takes into account nothing about the fitness structure of the environment.

Why do we observe this connection between the fitness value of information and the mutual information? To answer that question, it helps take a closer look at information-theoretic definition of information: information is the reduction of uncertainty, where uncertainty measures the number of states a system might be in. Thus mutual information between the world and a cue is the fold reduction in uncertainty about the world after the cue is received. For example, if a system could be in any of six equiprobable states, and a cue serves to narrow the realm of possibility to just three of these, the cue provides a twofold reduction in uncertainty. For reasons of convenience, information is measured as the logarithm of the fold reduction in uncertainty. Logarithmic units ensure that the measure is additive, so that for example we can add the information received by two successive cues to calculate the total information gained [29, 30, 3]. Thus while information concepts are often thought to be linked with the famous sum $-\sum p \log(p)$, the fundamental concept is not a particular mathematical formula. Rather, it is the notion that information measures the fold reduction in uncertainty about the possible states of the world.

With this view, it is easier to see why information bears a close relation to biological fitness: For simplicity, consider an extreme example in which individuals survive only if their phenotype matches the environment exactly, and suppose that there are ten possible environments that occur with equal probability. In the absence of any cue about the environment, the best the organism can do is randomly choose one of the ten possible phenotypes with equal probability. Only one

tenth of the individuals will then survive, since only a tenth will match the environment with their phenotype. If a cue conveys 1 bit of information and thus reduces the uncertainty about the environment twofold, the environment can be only in one of five possible states. The organism will now choose randomly one of five possible phenotypes, and now a fifth of the population will survive — a twofold increase in fitness, or a gain of 1 bit in the log of the growth rate.

What happens when the environments are not equiprobable? In this case we can understand the connection between information and fitness by looking to long sequences of environments and the theory of typical sequences. The theory tells us that almost surely one of the "typical sequences" — those sequences in which the environments occur in their expected frequencies — will occur [4]. Moreover, all typical sequences occur with equal probability. Thus a lineage is selected to divide its members equally among all typical sequences. Since any one mistake in phenotype is lethal, only a fraction of these lineages, namely those that have just the right sequence, will survive. The number of typical sequences in this case is exactly $2^{NH(E)}$ where N is the number of generations in the sequence and $H(E)$ is the entropy of the environment. Correspondingly, the fraction of surviving lineages will be $2^{-NH(E)}$. If a cue C is received that reduces the uncertainty of the environments by $I(E;C)$, then the fraction of surviving lineages can be increased by exactly $2^{NI(E;C)}$. This is analogous to the situation in communication: if we need to encode a string of symbols that are not equiprobable, we turn to a long sequence of such symbols. Our code then needs only to be efficient for representing typical sequences of symbols, and those typical sequences occur with equal probability. The number of such sequences is 2^{NH}, where N is the length and H is the entropy of the symbols. If the message recipient also obtains side information related to the message itself, then the mutual information I between the message and the side information measures the reduction in the number of possible messages that need to be encoded by the transmitter. This number of messages is reduced exactly 2^{NI}-fold by the presence of the side information.

Finally, what happens when having the wrong phenotype is not lethal, but simply decreases fitness? In this case, we can no longer simply count the number of lineages that have the correct sequence of phenotypes to determine the fraction that survive. However, we can transform the system into one where this is possible, by constructing an alternate set of hypothetical phenotypes that survive in just one environment, and expressing everything in terms of those phenotypes.

We imagine that, instead of an individual developing a single phenotype, it develops a certain combination of the alternate phenotypes; instead of following lineages of individuals, we follow lineages of these alternate phenotypes. The fraction that survive without information is, at best, $2^{-NH(E)}$, while the fraction that survive with information is, at best, $2^{-NH(E|C)}$. The mutual information $I(E;C)$ places an upper limit on the fold increase of lineages that survive when a cue is available.

We can now see why the concept of information is the same across different disciplines. In communication theory, the transmission of information is the reduction of uncertainty about what signals will come through a channel, from an initial set of all possible signals down to the *post hoc* set of signals actually received. In thermodynamics, a decrease in entropy refers to the fold reduction in the number of states that a system can be in. In evolutionary biology, fitness refers to the fold increase in the number of surviving lineages.

Acknowledgments

References

[1] John Maynard Smith. The idea of information in biology. *Quarterly Review of Biology*, 74(4):395–400, 1999.

[2] N. Wiener. *Cybernetics.* John Wiley and Sons, New York, 1948.

[3] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423; 623–656, 1948.

[4] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* John Wiley and Sons, New York, 1991.

[5] L. J. Savage. *The Foundations of Statistics.* Wiley, New York, 1954.

[6] I. J. Good. On the principle of total evidence. *British Journal for the Philosophy of Science*, 17:319–321, 1966.

[7] R. L. Winkler. *Introduction to Bayesian Inference and Decision.* Holt, Rinehard, and Winston, Inc., 1972.

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

[9] 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.

[10] D. W. Stephens and J. R. Krebs. *Foraging Theory*. Princeton University Press, Princeton, 1986.

[11] D. W. Stephens. Variance and the value of information. *American Naturalist*, 134(1):128–140, 1989.

[12] M. Lachmann and C. T. Bergstrom. The disadvantage of combinatorial communication. *Proceedings of the Royal Society of London Series B*, 271:2337–2343, 2004.

[13] J. B. S. Haldane. The cost of natural selection. *Journal of Genetics*, 55:511–524, 1957.

[14] M. Kimura. Natural selection as the process of accumulation of genetic information in adaptive evolution. *Genetical Research*, 2:127–140, 1961.

[15] J. Felsenstein. Macroevolution in a model ecosystem. *American Naturalist*, 112:177–195, 1978.

[16] C. T. Bergstrom and M. Lachmann. Shannon information and biological fitness. In *IEEE Information Theory Workshop 2004*, pages 50–54. IEEE, Oct. 2004.

[17] E. Kussell and S. Leibler. Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309:2075–2078, 2005.

[18] J. L. Kelly. A new interpretation of information rate. *Bell Systems Technical Journal*, 35:917–926, 1956.

[19] E. R. Dempster. Maintenance of genetic heterogeneity. *Cold Spring Harbor Symposium on Quantitative Biology*, 20:25–32, 1955.

[20] J. B. S. Haldane and S. D. Jayakar. Polymorphism due to selection of varying direction. *Journal of Genetics*, 58:237–242, 1963.

[21] D. Cohen. Optimizing reproduction in a randomly varying environment. *Journal of Theoretical Biology*, 12:119–129, 1966.

[22] R. C. Lewontin and D. Cohen. On population growth in a randomly varying environment. *Proceedings of the National Academy of the Sciences*, 62:1056–1060, 1969.

[23] J. H. Gillespie. Polymorphism in random environments. *Theoretical Population Biology*, 4:193–195, 1973.

[24] J. Yoshimura and V. A. A. Jansen. Evolution and population dynamics in stochastic environments. *Res. Popul. Ecol.*, 38(2):165–182, 1996.

[25] W. S. Cooper and R. H. Kaplan. Adaptive "coin-flipping"; a decision-theoretic examination of natural selection for random individual variation. *Journal of Theoretical Biology*, 94(94):135–151, 1982.

[26] Seger and Brockmann. What is bet-hedging? *Oxford Surverys in Evolutionary Biology*, 4:182–211, 1987.

[27] 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.

[28] M. C. Donaldson-Matasci, M. Lachmann, and C. T. Bergstrom. Phenotypic diversity as an adaptation to environmental uncertainty. *Evolutionary Ecology Research*, 10:493–515, 2008.

[29] H. Nyquist. Certain factors affecting telegraph speed. *Bell Systems Technical Journal*, pages 324–346, 1924.

[30] R. V. Hartley. Transmission of information. *Bell Systems Technical Journal*, 7:535–563, 1928.