

The Value of Information for Populations in Varying Environments

Olivier Rivoire · Stanislas Leibler

Received: 25 October 2010 / Accepted: 22 February 2011 / Published online: 10 March 2011
© Springer Science+Business Media, LLC 2011

Abstract The notion of information pervades informal descriptions of biological systems, but formal treatments face the problem of defining a quantitative measure of information rooted in a concept of fitness, which is itself an elusive notion. Here, we present a model of population dynamics where this problem is amenable to a mathematical analysis. In the limit where any information about future environmental variations is common to the members of the population, our model is equivalent to known models of financial investment. In this case, the population can be interpreted as a portfolio of financial assets and previous analyses have shown that a key quantity of Shannon’s communication theory, the mutual information, sets a fundamental limit on the value of information. We show that this bound can be violated when accounting for features that are irrelevant in finance but inherent to biological systems, such as the stochasticity present at the individual level. This leads us to generalize the measures of uncertainty and information usually encountered in information theory.

Keywords Information · Regulation · Fitness · Population dynamics · Branching processes · Control · Investment

1 Introduction

Information is a central concept in biology [1–4], which many studies have sought to formalize [5–11]. In this quest, Shannon’s theory of communication [12] has always played an

O. Rivoire (✉)
CNRS, LIPhy UMR 5588, Grenoble, France
e-mail: olivier.rivoire@ujf-grenoble.fr

S. Leibler
Laboratory of Living Matter, The Rockefeller University, New York, NY, USA

S. Leibler
School of Natural Sciences, The Simons Center for Systems Biology, The Institute for Advanced Study,
Princeton, NJ, USA

influential role. Originally, this theory was concerned with two basic problems: the problem of efficiently encoding signals, and the problem of reliably transmitting them through noisy channels. Shannon proposed a formal framework within which these questions could be addressed mathematically. By modeling information sources and communication channels in probabilistic terms, and by focusing on the asymptotic properties of long sequences of symbols, he established fundamental limits for the achievable rates of data compression and transmission [12, 13]. By virtue of the abstract nature of the model, these limits hold irrespectively of the particular material implementation. Remarkably, the same quantity, the mutual information $I(X; Y)$, a function of two random variables X and Y , emerges as a common measure of “information” in the solution of the two problems [14]. As for the related concept of entropy $H(X) = I(X; X)$, the definition of the mutual information can be axiomatized [12, 15], which has lent support to the view that this quantity represents an universal and irrefutable measure of information. The emergence of the mutual information as a central quantity in problems of point-to-point communication however rests on specific assumptions, which have to be reexamined in any other instance where a concept of “information” is to be formalized [16].

A class of problems where such a reexamination has led to identifying a different measure of information is constituted by the engineering problems of control. These problems share two essential features with biological systems: information is processed for a “function”, which confers value to the information, and feedback, whereby elements from the past are used to affect the present, is essential. Historically, this parallel between regulation in living organisms and control in engineered systems has underlaid the seminal works on control with feedback [17]. It also motivated the influential development of cybernetics, which Wiener defined as “the science of control and communication, in the animal and the machine” [18]. A law formulated in the early days of cybernetics is thus the “law of requisite variety” [19, 20], which states that the value of information for control cannot exceed the limit set by the mutual information between a disturbance and its measurement (see also [21, 22]). The issue of quantifying information in systems of control has been revisited thoroughly since this law was proposed [23]. These analyses have concurred to establish the so-called directed information [24, 25] as a measure of information more relevant than the mutual information when issues of feedback are involved. The directed information measures the causal dependence between two stochastic processes, in contrast with the mutual information which ignores any constraint of causality and only measures statistical correlations. Consistently with the law of requisite variety, the mutual information however appears as an upper bound for the value of information for control when the later is measured by a directed information.

The parallel between living organisms and engineered systems provides interesting insights but fails to account for two other essential features of living organisms: their organization into populations, and the need to evaluate performance in terms of “fitness”, i.e., in terms of an appropriate measure of reproductive value. Viewing the problem of control from the standpoint of populations of reproducing individuals indeed introduces new options for coping with unpredictable variations of the environment. Most importantly, a “bet-hedging strategy” [26] can be implemented through the diversification of the population. An analogy with financial problems of risk management has been noticed many times [27–30], including from the perspective of information processing [31–34]. Both problems involve a growing population facing an unpredictable future: in the financial problem, the population is composed by the capital of an investor, which is distributed between different assets. These assets are analogous to the phenotypes of biological organisms, and may respond differently to different environmental perturbations. The problem of quantifying the value of information in this context was first analyzed by Kelly [35], who found that the mutual information

appeared as a natural measure. His results were later expanded [36–39] showing that, in general, the relevant measure for the value of information must incorporate characteristic features of the individuals, such as their multiplication rates. A result analogous to the law of requisite variety however still holds: the value of the information that an investor may collect remains bounded by the mutual information between this information and the actual state of the environment (the stock market in the financial context) [38].

The analogy between biological populations and problems of financial investment has also its limitations. The main conceptual difference is that the financial problem is supervised by a goal-oriented investor, who centralizes the information and the decisions, while information processing is distributed between potentially independent individuals in biological populations. A first implication is that the biological problem may not correspond to an optimization at the population level, as it does by definition in finance. In any case, the justification of a criterion of optimality must involve a non-arbitrary objective function that emerges from the dynamics of the population instead of being a priori defined. The distributed nature of the biological problem also introduces a level of individual stochasticity that is absent in finance: even if every individual has the same sensor and has access to the same information, stochastic noise in individual sensor can lead to the perception of non-identical signals. This aspect of the problem of information processing, which has not been previously examined from an information theoretic standpoint, also leads to a measure of the value of information that differs from the mutual information. In this case, the law of requisite variety may be violated: the value of the relevant measure of information can exceed the value indicated by the mutual information. A population may thus effectively acquire, in a distributed form, a more accurate information than any of its members.

We shall discuss each of these points in the context of a mathematical model of growing populations in a varying environment. This model is defined in Sect. 2 and its main elements are represented in Fig. 1. It deals with two types of biological information: the information inherited by an individual from its parents, and the information directly acquired from the environment. To exploit the analogies with the engineering problem of control and the financial problem of investment (see Table 1), we introduce in Sect. 3 a suitable “fitness function”. Our presentation is then organized around three simplifying assumptions: assumption (A1) that individuals have no memory, assumption (A2) that individuals all perceive the same information from the environment, and assumption (A3) that only individuals perfectly adapted to their environment can survive. While under the conjunction of these three restrictions, the value of information is expressed by a mutual information (Sect. 4) [35], relaxing any of these assumptions exposes a different limitation of the mutual information. Relaxing (A1) introduces the possibility of feedback, in which case constraints of causality not accounted for by the mutual information need to be incorporated (Sect. 5) [40]. Relaxing (A2) introduces the possibility for individuals to perceive different signals from their common environment, which also requires generalizing the mutual information (Sect. 6). Finally, relaxing (A3) introduces the possibility of different environmental states having non-exclusive “meaning”, where the source of meaning, encapsulated in the values of the multiplication rates of the individuals, needs to be taken explicitly into account in the measure of information (Sect. 7) [14]. Different expressions for quantifying the value of information are thus obtained, which are summarized in Table 2.

Besides the question of quantifying the value of information, our model also addresses a second question, the question of characterizing the evolutionary stable strategies that optimize fitness. We shall show that, under the assumptions (A2) and (A3), these strategies amount to a Bayesian computation, as conjectured for instance in [41]. When these assumptions are not satisfied, however, we find that population-level features can make the implementation of a Bayesian computation irrelevant.

2 Model

Our approach to investigating the nature and value of information in biological systems is based on an abstract mathematical model. Expressions for the value of information will result from analyzing this model. Specifically, our model seeks to incorporate the following features, which appear to be commonly shared by all living organisms:

- (i) Living organisms change (as a result of development, phenotypic plasticity, learning, ...);
- (ii) Living organisms can generate other living organisms;
- (iii) The faculties (i) and (ii) are affected by the state of the organism and the state of its environment;
- (iv) The environment of living organisms varies.

The issue of regulation arises when constraints are present which prevent the organisms from perfectly anticipating environmental changes. Here, we focus on constraints due to limited information (see Sect. 8 for extensions):

- (v) Changes within a living organism take place in absence of complete knowledge of the forthcoming environmental states that will affect survival and reproduction.

To account for (iv), the environment is described by a discrete-time and discrete-state Markov chain, with transition matrix $p(x_t|x_{t-1})$. This Markov chain is assumed to be stationary and ergodic. We shall expand on the notion of ergodicity in Sect. 3, but, in essence, it requires that any environmental state can be reached from any other state in finite time and with finite probability [42]. Ergodic Markov chains tend asymptotically to a unique stationary distribution $p_s(x_t)$, irrespectively of their initial state, where $p_s(x_t)$ satisfies $p_s(x_t) = \sum_{x_{t-1}} p(x_t|x_{t-1})p_s(x_{t-1})$. We assume here that the environmental process is stationary. A particular case of interest is when the successive environmental states are uncorrelated and described by independently and identically distributed (i.i.d.) random variables, each having a probability $p(x_t)$, corresponding to $p(x_t|x_{t-1}) = p(x_t) = p_s(x_t)$.

Each individual organism is characterized by an internal state σ_t , to which we will refer as its current “type”; in general, it corresponds to a distinct phenotype, but may also be associated with a distinct genotype. To account for (ii), the number $f(\sigma_t; x_t)$ of offsprings generated by an individual organism at time t depends both on its type σ_t , and on the current state x_t of the environment; in particular, the individual may die if $f(\sigma_t; x_t) = 0$ or survive without reproducing if $f(\sigma_t; x_t) = 1$. As a simplifying assumption, we assume here that all offsprings inherit the type σ_t of their parent. More generally, a non-integer value of $f(\sigma_t; x_t)$ will represent the expected number of offsprings of an individual of type σ_t in environment x_t ; $f(\sigma_t; x_t)$ will therefore be called a multiplication rate. To account for (iii) and (v), the current type σ_t can depend both on the ancestral type σ_{t-1} of the individual, and on a signal y_t derived from the environment x_t . Following the example of communication theory [12], this dependence is described probabilistically, with a transition matrix $\pi(\sigma_t|\sigma_{t-1}, y_t)$ giving the probability to end up in state σ_t given (σ_{t-1}, y_t) . In the language of information theory, such a transition matrix is a “communication channel”, here with input (σ_{t-1}, y_t) and output σ_t ; mathematically, it must satisfy two basic properties:

$$\pi(\sigma_t|\sigma_{t-1}, y_t) \geq 0, \quad \text{for all } \sigma_t, \sigma_{t-1}, y_t, \quad \text{and} \quad \sum_{\sigma_t} \pi(\sigma_t|\sigma_{t-1}, y_t) = 1 \quad \text{for all } \sigma_{t-1}, y_t. \quad (1)$$

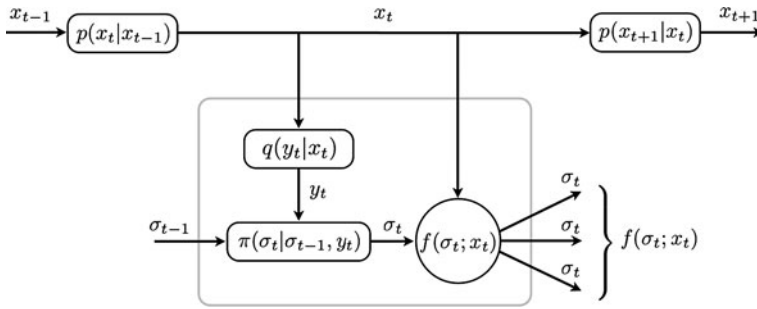


Fig. 1 Schematic representation of the relation between the environment (*upper part*) and an individual (*lower part*). The current environmental state x_t affects the number $f(\sigma_t; x_t)$ of offsprings than an individual of type σ_t generates (or, more generally, the expected number of offsprings given σ_t and x_t). The type σ_t is determined probabilistically by the transition matrix $\pi(\sigma_t|\sigma_{t-1}, y_t)$ which depends on the ancestral type σ_{t-1} and on the perception y_t that the individual has of the current environment. This perception y_t derives from x_t through a transition matrix $q(y_t|x_t)$

The relation between the signal y_t and its source x_t is also specified by a transition matrix $q(y_t|x_t)$ representing the “sensor” by which individuals acquire information from their current environment (see Fig. 1). Importantly, while the environment x_t is common to all individuals, the perceived signal y_t generally varies from individual to individual in the population.

Equation for the Conditional Mean Population Size \mathcal{N}_t The model is defined at the level of individual organisms, but selection may also act at the level of the population; for instance, a diversification between different types may confer an advantage in the long term when the environmental changes are unpredictable. An important implication is that the problem of regulation in a varying environment should not be treated by isolating an individual from the population. Here, the population is characterized by the numbers $Z_t(\sigma)$ of individuals of each type σ , which define a population vector Z_t whose norm $|Z_t| \equiv \sum_{\sigma} Z_t(\sigma)$ is the total population size. This vector Z_t is a random variable from two standpoints: it depends on the environmental sequence $\bar{x} \equiv (x_1, \dots, x_t, \dots)$, and for a given \bar{x} , it is subject to the stochasticity at the individual level, generated through the transition matrices $q(y_t|x_t)$ and $\pi(\sigma_t|\sigma_{t-1}, y_t)$ (and possibly also through the fluctuations in the number of offsprings if $f(\sigma_t; x_t)$ represents a multiplication rate). Our analysis will focus on the conditional mean

$$\mathcal{N}_t(\sigma) \equiv \langle Z_t(\sigma) \rangle \quad (\text{average taken for a given } \bar{x}), \tag{2}$$

which follows a simple recursion:

$$\mathcal{N}_t(\sigma_t) = f(\sigma_t; x_t) \sum_{\sigma_{t-1}, y_t} \pi(\sigma_t|\sigma_{t-1}, y_t) q(y_t|x_t) \mathcal{N}_{t-1}(\sigma_{t-1}). \tag{3}$$

This recursion can also be written with a vectorial notation:

$$\mathcal{N}_t = \mathcal{A}^{(x_t)} \mathcal{N}_{t-1}, \quad \text{with } \mathcal{A}_{\sigma'\sigma}^{(x_t)} \equiv f(\sigma'; x_t) \sum_{y_t} \pi(\sigma'|\sigma, y_t) q(y_t|x_t). \tag{4}$$

Here, the current environment x_t is a “quenched” variable, which is fixed independently of the dynamics of the population. From a mathematical standpoint, (4) indicates that studying \mathcal{N}_t amounts to studying the product of random matrices $\mathcal{A}^{(x_t)} \mathcal{A}^{(x_{t-1})} \dots \mathcal{A}^{(x_1)}$, which is

function of the environmental sequence \bar{x} . Considering only $\mathcal{N}_t = \langle Z_t \rangle$ overlooks the discrete nature of the population, and thus fails to account for possible events of extinction; a population of discrete individuals is indeed not infinitely divisible, and the stochasticity of the process of reproduction may lead to $|Z_t| = 0$ at some time t , after which any possibility of recovery is excluded. Remarkably however, the results presented in Sect. 3 indicate that the basic asymptotic behavior of $|Z_t|$ can be derived from the properties of $|\mathcal{N}_t|$, which will justify that our analysis concentrates on (3).

Two Basic Questions The transition matrix $\pi(\sigma_t | \sigma_{t-1}, y_t)$ specifies the “strategy” for responding to the signals that individuals inherit and acquire. A basic problem is to provide a framework for estimating the relative performance of different strategies. In some particular cases, the “best” strategy is clear: if perfect information is available, a sensible action is indeed for every individual to adopt at time t the type σ that maximizes $f(\sigma; x_t)$, thus leading to an homogeneous population. Perfect information correspond to noiseless channels, represented by the identity transition matrix δ such that $\delta(y_t | x_t) = 1$ if $y_t = x_t$, and 0 otherwise. In general, however, the communication channel q will reveal incomplete information about x_t , and a non-deterministic response, leading to a diversified population, may be more advantageous. Two basic questions thus arise:

- (Q1) What strategy, i.e., choice of the transition matrix $\pi(\sigma_t | \sigma_{t-1}, y_t)$, is the most advantageous?
 (Q2) What is the value of the information acquired through $q(y_t | x_t)$?

Answering these two questions require defining a measure of “fitness”, so as to give a precise meaning to the notions of “advantage” and “value”. In decision theory, this usually involves the introduction of an *ad-hoc* loss-function.¹ We argue however in the next section that a measure of adaptation emerges in the long-term limit, which defines a non-arbitrary fitness function.

Three Simplifying Assumptions As the model is not analytically solvable in its most general form, it is of interest to analyze it under several simplifying assumptions. Three simplifying assumptions will play a crucial role in our analysis:

- (A1) No information is inherited between successive generations, i.e., $\pi(\sigma_t | \sigma_{t-1}, y_t) = \pi(\sigma_t | y_t)$;
 (A2) Any information y_t acquired from the environment is common to all individuals in the population;
 (A3) The multiplication rates have a diagonal form: $f(\sigma; x) = f(x) > 0$ if $\sigma = x$, and $f(\sigma; x) = 0$ otherwise.

Inherited information becomes useful in presence of correlations between successive environmental states and assumption (A1) is therefore restrictive only when the environment is not i.i.d. In the language of control theory, presented in Table 1 and developed in Sect. 8, assumption (A1) of no memory corresponds to an open-loop mode of control where feedback is absent. As developed in Sect. 6, assumption (A2) of common acquired information amounts to restricting to models which can be interpreted in financial terms, with no fluctuation in the signals perceived by the individuals. Assumption (A3) of diagonal multiplication rates describes the situation where in any environmental state, there is only one type able

¹From this standpoint, our model is related to the so-called partially observable Markov decision processes studied in the operations-research literature [43].

to survive; in particular, this assumption assumes that the number of environmental states is the same as the number of types for the individuals.

The model defined by the conjunction of the three assumptions plays a special role, because, as explained in Sect. 4, the two questions (Q1) and (Q2) have simple answers in terms of the standard measures of uncertainty and information from communication theory, the entropy $H(X_t)$ and the mutual information $I(X_t; Y_t)$ where X_t and Y_t refers to the random variables associated with the environmental state x_t and the signal y_t ($H(X_t)$ and $I(X_t; Y_t)$ are defined below). As we shall show, relaxing any of these assumptions introduces generalizations of these two quantities. The models satisfying all three assumptions were also the first models of population growth to be analyzed from the standpoint of information theory [35]. These models were originally interpreted as models of gambling in horse races, with $f(x_t)$ viewed as the pay-off when horse $\sigma = x_t$ wins and $p(x_t|x_{t-1})$ as the probability for it to happen, given that horse x_{t-1} won the previous race. Generalizations to models of investment in the stock market, involving relaxation of the assumptions (A1) and (A3) have subsequently been considered from the same standpoint [14]. The relevance of this approach to understanding the adaptive value of strategies of diversification and the value of information in biological populations has also been previously noticed [32, 33], although always under the restrictive assumption (A2) that information is acquired with no individual stochasticity.

3 Fitness and Optimization

The problem of defining a relevant measure of fitness is related to a problem that has been thoroughly studied in economics: the problem of defining a suitable “utility function”, i.e. a measure of performance that rational agents should seek to optimize. In the present biological context, the arguments recalled in Appendix B lead to considering the growth rate of the population, defined by

$$\Lambda_{p;f}^{(q)}(\pi) = \lim_{t \rightarrow \infty} \frac{1}{t} \ln |\mathcal{N}_t|, \tag{5}$$

as an univocal measure of fitness in the long-term.

For an arbitrary environmental processes, the growth rate may depend on the particular sequence of environmental states x_1, x_2, \dots, x_t that arises. Stationary ergodic processes, however, benefit from a self-averaging property: particular realizations of such processes tend with time to share common statistical features—features that reproduce those obtained by averaging over many particular sequences; this property is also known as the asymptotic equipartition property in information theory, where it plays an equally fundamental role and underlies the choice of considering infinitely long sequences of symbols [14]. For independent environments, ergodicity amounts to the law of large numbers (see (59) in Appendix B): almost all long sequences comprise a same fraction $p(x)$ of each state x . More generally, assuming that the environmental process is stationary and ergodic, and that $\mathbb{E}[\max(0, \ln \mathcal{A}_{\sigma\sigma'}^{(x_t)})] < \infty$ for all σ, σ' , where $\mathcal{A}_{\sigma\sigma'}^{(x_t)}$ is defined as in (4), it can be shown [44, 45] that the limit

$$\Lambda_{p;f}^{(q)}(\pi) \equiv \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \ln |\mathcal{N}_t| \tag{6}$$

exists,² and that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln |\mathcal{N}_t| = \Lambda_{p,f}^{(q)}(\pi) \quad \text{with probability 1.} \tag{7}$$

No simple analytical formula is available for $\Lambda_{p,f}^{(q)}(\pi)$, also known as a Lyapunov exponent, in the most general case, but an important exception is in absence of inherited information, when $\pi(\sigma_t | \sigma_{t-1}, y_t) = \pi(\sigma_t | y_t)$ [assumption (A1) of absence of memory], in which case

$$\Lambda_{p,f}^{(q)}(\pi) = \sum_{x_t} p_s(x_t) \ln \left(\sum_{\sigma_t, y_t} f(\sigma_t; x_t) \pi(\sigma_t | y_t) q(y_t | x_t) \right). \tag{8}$$

Typical vs Mean Population Sizes Importantly, not only $\Lambda_{p,f}^{(q)}(\pi)$ describes the growth rate of the conditional mean $\mathcal{N}_t = \langle Z_t \rangle$, but also, under fairly general conditions, the growth rate of the size $|Z_t|$ of a typical population. An essential condition, however, is that the population does not become extinct. The probability of survival of a population with an arbitrary initial composition can always be expressed in terms of the probabilities $Q(\sigma | \bar{x})$ of extinction of a population starting from a single individual of type σ : if starting with $\mathcal{N}_0(\sigma)$ individuals in each state σ , the probability of survival is indeed $1 - \prod_{\sigma} Q(\sigma | \bar{x})^{\mathcal{N}_0(\sigma)}$, because each individual generates its own independent subpopulation. Here, we assume that either all the types have a non-zero probability to survive, i.e., $\mathbb{P}(Q(\sigma | \bar{x}) < 1, \forall \sigma) = 1$, or none of them survive, i.e., $\mathbb{P}(Q(\sigma | \bar{x}) = 1, \forall \sigma) = 1$; if this is not the case, we can always ignore the types that inevitably become extinct. Under this condition of regularity and a further technical condition of stability presented in Appendix A, the following classification theorem holds [46]:

$$\begin{aligned} \text{(i)} \quad & \Lambda_{p,f}^{(q)}(\pi) < 0 \implies \mathbb{P}[Q(\sigma | \bar{x}) = 1, \forall \sigma] = 1; \\ \text{(ii)} \quad & \Lambda_{p,f}^{(q)}(\pi) > 0 \implies \lim_{t \rightarrow \infty} \frac{1}{t} \ln |Z_t| = \Lambda_{p,f}^{(q)}(\pi) \tag{9} \\ & \text{almost surely conditionally to non-extinction.} \end{aligned}$$

The second case where there is a non-zero probability of non-extinction is known as the supercritical case and is obviously the one of interest here. Remarkably, this theorem indicates that the growth of branching processes is controlled by the properties of the product of random matrices $\mathcal{A}^{(x_t)} \dots \mathcal{A}^{(x_1)}$ which governs the evolution of the conditional mean \mathcal{N}_t .³

Reformulation of the Two Basic Questions Based on the mathematical results presented in this section, the questions (Q1) and (Q2) of the previous section can be stated formally. Taking the long-term growth rate $\Lambda_{p,f}^{(q)}(\pi)$ as a measure of fitness, (Q1) becomes the problem of finding a matrix $\hat{\pi}$ that maximizes it for given parameters p, f, q (while the optimal

²We use two different symbols for representing two different types of averages: $\langle Z_t(\sigma) \rangle$ for the average conditionally to the environmental sequence \bar{x} , and $\mathbb{E}[Z_t]$ for the average over environmental sequences as well.

³Even the condition of stability, detailed in Appendix A, bears on properties of this product: it basically requires that its columns all grow at a same rate so as to prevent too large fluctuations in the population size. Also note that both this condition of stability and the other condition of regularity relative to the probability of extinction are trivially satisfied for a single-type population, to which our model can be reduced in absence of inherited information, when $\pi(\sigma_t | \sigma_{t-1}, y_t) = \pi(\sigma_t | y_t)$ [assumption (A1)], by noticing that a recursion can be written directly for $|N_t| = \sum_{\sigma} N_t(\sigma)$, as for instance in (54).

growth rate $\hat{\Lambda}_{p;f}^{(q)}$ is unique, “the” optimal strategy $\hat{\pi}$ may not be). Based on the same principle, (Q2) becomes the problem of estimating $\hat{\Lambda}_{p;f}^{(q)} - \hat{\Lambda}_{p;f}^{(\eta)}$, where $\hat{\Lambda}_{p;f}^{(q)} = \Lambda_{p;f}^{(q)}(\hat{\pi})$ denotes the optimal growth rate in presence of the channel q , and $\hat{\Lambda}_{p;f}^{(\eta)}$ the optimal growth rate in its absence, with η standing for an informationless channel. (Q1) and (Q2) thus amount to estimating the following quantities:

- (Q1) $\hat{\pi} \equiv \arg \max_{\pi} \Lambda_{p;f}^{(q)}(\pi)$ (optimal strategy);
- (Q2) $\hat{\Lambda}_{p;f}^{(q)} - \hat{\Lambda}_{p;f}^{(\eta)}$ (value of the information conveyed by q).

In the next section, we show that under the assumptions (A1), (A2) and (A3), the cost of uncertainty, defined as $\hat{\Lambda}_{p;f}^{(\delta)} - \hat{\Lambda}_{p;f}^{(q)}$ where δ stands for a noiseless channel, and the value of acquired information, defined as $\hat{\Lambda}_{p;f}^{(q)} - \hat{\Lambda}_{p;f}^{(\eta)}$, correspond respectively to the conditional entropy $H(X_t|Y_t)$ and the mutual information $I(X_t; Y_t)$. In Sects. 5 and 6, we show that upon relaxing the assumptions (A1) or (A2), the cost of uncertainty and value of acquired information are still independent of the multiplication matrix $f(\sigma; x)$, and define two quantities $H_p^{(q)}$ and $I_p^{(q)}$ that generalize the notions of conditional entropy and mutual information. Finally, we show in Sect. 7, that, in absence of any assumption, the statistical quantities $I_p^{(q)}$ and $H_p^{(q)}$ are respectively bounds for the cost of uncertainty $\hat{\Lambda}_{p;f}^{(\delta)} - \hat{\Lambda}_{p;f}^{(q)}$ and the value of acquired information $\hat{\Lambda}_{p;f}^{(q)} - \hat{\Lambda}_{p;f}^{(\eta)}$, which generally also depend on f . These results are summarized in Table 2.

4 Model Under the Assumptions (A1), (A2) and (A3): Kelly’s Horse Races

As originally shown by Kelly for a model of horse races [35], under the joint assumptions (A1), (A2) and (A3) stated in Sect. 2, a simple connection is found between the long-term growth rate and information theoretic quantities. We start by assuming that, in addition to the restrictions imposed by (A1), (A2) and (A3), the environment is i.i.d. with probability $p(x_t)$ for the environmental states.

Value of Information and Cost of Uncertainty in Absence of Acquired Information In absence of acquired information (informationless channel $q = \eta$), the long-term growth rate is given by (8):

$$\Lambda_{p;f}^{(\eta)}(\pi) = \sum_x p(x) \ln (f(x)\pi(x)). \tag{10}$$

Taking into account the constraint $\sum_x \pi(x) = 1$ with a Lagrange multiplier, the strategy $\hat{\pi}$ maximizing $\Lambda_{p;f}^{(\eta)}(\pi)$ is found to be

$$\hat{\pi}(x) = p(x), \quad \forall x. \tag{11}$$

This strategy is called proportional betting and has the remarkable property of not depending on the values of the multiplication rates $f(x)$. It yields an optimal growth rate that can be broken down in two terms:

$$\hat{\Lambda}_{p;f}^{(\eta)} = \sum_x p(x) \ln f(x) + \sum_x p(x) \ln p(x) = \hat{\Lambda}_{p;f}^{(\delta)} - H_p^{(\eta)}. \tag{12}$$

The first term, $\hat{\Lambda}_{p;f}^{(\delta)} = \mathbb{E}[\ln f] = \sum_x p(x) \ln f(x)$, corresponds to the best conceivable growth rate in a typical sequence of environments: it is achieved if all individuals behave optimally in presence of perfect information, corresponding to a noiseless sensor $q = \delta$. The second term,

$$H_p^{(n)} \equiv \hat{\Lambda}_{p;f}^{(\delta)} - \hat{\Lambda}_{p;f}^{(n)} = - \sum_x p(x) \ln p(x), \tag{13}$$

is independent of f and corresponds to Shannon’s entropy for the random variable X_t , usually denoted $H(X_t)$ or $H(p)$ (we follow the common usage of representing by X_t the random variable and by x_t one of its values). The entropy quantifies the cost of uncertainty when the frequencies $p(x)$ are known but not the particular sequence of outcomes that occurs (see Appendix C). Since $\hat{\pi} = p$, a population well-adapted to a varying environment must, in this model, have evolved an “internal model of the environment” that encodes its statistical properties [33]; in this sense, the matrix π can be viewed as information about the environment that is common knowledge in the population (see also Fig. 3).

Cost of Non-optimal Strategies If the frequencies $p(x)$ are not estimated correctly, suggesting a suboptimal strategy $\pi \neq \hat{\pi}$, an additional cost is incurred,

$$\Lambda_{p;f}^{(n)}(\pi) = \hat{\Lambda}_{p;f}^{(n)}(\hat{\pi}) - D(\hat{\pi} \parallel \pi). \tag{14}$$

This cost involves another quantity playing a fundamental role in communication theory [14], the so-called relative entropy, or Kullback-Leibler divergence, which is defined by

$$D(\hat{\pi} \parallel \pi) \equiv \sum_x \hat{\pi}(x) \ln \frac{\hat{\pi}(x)}{\pi(x)}. \tag{15}$$

It measures the deviation of the distribution $\hat{\pi}$ from the distribution π and obeys the inequality $D(\pi \parallel \hat{\pi}) \geq 0$, with equality if and only if $\pi(x) = \hat{\pi}(x)$ for all x .

Value of Information and Cost of Uncertainty in Presence of Acquired Information We now assume that an information y_t is available about the environmental state x_t , through a communication channel characterized by the transition matrix $q(y_t|x_t)$. The strategy π can now depend on the signal y_t . For instance, there may be n possible signals, in which case, no information would correspond to $q(y_t|x_t) = 1/n$, and perfect information to $q(y_t|x_t) = 1$ if $y_t = x_t$ and 0 otherwise. In general, some noise may cause $q(y_t|x_t)$ to be non-zero even if $y_t \neq x_t$. The expression for the growth rate is now

$$\Lambda_{p;f}^{(q)}(\pi) = \sum_{x,y} q(y|x) p(x) \ln (f(x)\pi(x|y)), \tag{16}$$

where $q(y|x)p(x)$ represents the joint probability $\mathbb{P}_{X_t, Y_t}(x, y)$ that the environmental state is x and the perceived signal is y . By conditioning with respect to the received signal y , the problem can be reduced to the case with no information:

$$\begin{aligned} \Lambda_{p;f}^{(q)}(\pi) &= \sum_{x,y} \mathbb{P}_{X_t, Y_t}(x, y) \ln(f(x)\pi(x|y)) \\ &= \sum_y \mathbb{P}_{Y_t}(y) \left[\sum_x \mathbb{P}_{X_t|Y_t}(x|y) \ln(f(x)\pi(x|y)) \right]. \end{aligned} \tag{17}$$

For any given y , the optimization problem is therefore solved as before, with $\mathbb{P}_{X_t}(x) = p(x)$ replaced by $\mathbb{P}_{X_t|Y_t}(x|y)$. The optimal strategy is thus ‘‘conditional proportional betting’’:

$$\hat{\pi}(x|y) = \mathbb{P}_{X_t|Y_t}(x|y) = \frac{\mathbb{P}_{Y_t|X_t}(y|x) \mathbb{P}_{X_t}(x)}{\mathbb{P}_{Y_t}(y)} = \frac{q(y|x) p(x)}{\sum_{x'} q(y|x') p(x')}. \tag{18}$$

It exactly amounts to a Bayesian computation [41]. The optimal value of the growth rate can again be broken down in two terms

$$\hat{\Lambda}_{p:f}^{(q)} = \hat{\Lambda}_{p:f}^{(\delta)} - H_p^{(q)}. \tag{19}$$

$H_p^{(q)}$ is a generalization of the entropy $H_p^{(n)}$ known as the conditional entropy, usually denoted $H(X_t|Y_t)$ in communication theory [14]. It measures the residual unpredictability of X_t given Y_t and is given by

$$H(X_t|Y_t) = \sum_y \mathbb{P}_{Y_t}(y) H(X_t|Y_t = y) = - \sum_{x,y} \mathbb{P}_{X_t,Y_t}(x,y) \ln \mathbb{P}_{X_t|Y_t}(x|y). \tag{20}$$

With perfect information, $Y_t = X_t$, and the entropic cost is eliminated, $H(X_t|X_t) = 0$, leaving only $\hat{\Lambda}_{p:f}^{(\delta)} = \mathbb{E}[\ln f]$. The gain in predictability due to the signal is obtained by comparing the situations with and without information,

$$I_p^{(q)} \equiv \hat{\Lambda}_p^{(q)} - \hat{\Lambda}_p^{(n)} = H_p^{(n)} - H_p^{(q)} = H(X_t) - H(X_t|Y_t) \equiv I(X_t; Y_t). \tag{21}$$

The quantity $I_p^{(q)} = I(X_t; Y_t)$ is another important measure of information in communication theory, the mutual information [14] (see Appendix C for some of its properties).

Conclusion We assumed so far that the environment was i.i.d. but under the assumption (A1) that no information can be inherited, the results of this section can be extended to Markov environments, and more generally to ergodic and stationary environmental processes, by replacing $p(x_t)$ by the stationary distribution $p_s(x_t)$ of the environmental process. To sum up, the growth rate for the model under the assumptions (A1), (A2), (A3), can be decomposed as

$$\Lambda_{p:f}^{(q)}(\pi) = \hat{\Lambda}_{p:f}^{(\delta)} - H_p^{(q)} - D(\hat{\pi}||\pi), \tag{22}$$

or, equivalently, as

$$\Lambda_{p:f}^{(q)}(\pi) = \hat{\Lambda}_{p:f}^{(\delta)} - H_p^{(n)} + I_p^{(q)} - D(\hat{\pi}||\pi). \tag{23}$$

$\hat{\Lambda}_{p:f}^{(\delta)} = \mathbb{E}[\ln f]$ represents the optimal growth rate with perfect information, and $D(\hat{\pi}||\pi)$ the cost for following a strategy π differing from the optimal strategy $\hat{\pi}$, with $D(\hat{\pi}||\pi) = 0$ if and only if $\pi = \hat{\pi}$. The first expression makes apparent the cost of uncertainty $H_p^{(q)}$, which corresponds here to a conditional entropy:

$$H_p^{(q)} = H(X_t|Y_t), \tag{24}$$

with $\mathbb{P}_{X_t,Y_t}(x_t, y_t) = q(y_t|x_t)p_s(x_t)$. The second expression introduces $H_p^{(n)} = H(X_t)$, the entropy of the environmental variable X_t , with $\mathbb{P}_{X_t}(x_t) = p_s(x_t)$. It makes explicit the value of acquired information $I_p^{(q)}$, which corresponds to the mutual information between the environment X_t and the acquired information Y_t :

$$I_p^{(q)} = I(X_t; Y_t). \tag{25}$$

In the next three sections, we examine how these relations are modified when relaxing any of the assumptions (A1), (A2) and (A3) on which they rely.

5 Relaxing (A1): Causal Constraints and Inherited Information

We first consider the consequences of relaxing the assumption (A1) that no information is inherited: the strategy can now depend not only on the acquired information y_t but also on the past state of the individual σ_t . Under the assumptions (A2) and (A3), the past state σ_{t-1} of an individual alive at time t is necessarily x_{t-1} , and the argument used to derive (17) can be invoked to infer that the Bayesian strategy, given by $\hat{\pi}(\sigma | x_{t-1}, y_t) = \mathbb{P}_{X_t|X_{t-1}, Y_t}(\sigma | x_{t-1}, y_t)$, is optimal [14], with an associated cost of uncertainty independent of f and given by

$$H_p^{(q)} \equiv \hat{\Lambda}_{p:f}^{(\delta)} - \hat{\Lambda}_{p:f}^{(q)} = H(X_t | X_{t-1}, Y_t). \tag{26}$$

Here, following the definition of (20), $H(X_t | X_{t-1}, Y_t)$ is given by

$$H(X_t | X_{t-1}, Y_t) = - \sum_{x_t, x_{t-1}, y_t} q(y_t | x_t) p(x_t | x_{t-1}) p_s(x_{t-1}) \ln \left(\frac{q(y_t | x_t) p(x_t | x_{t-1})}{\sum_z q(y_t | z) p(z | x_{t-1})} \right). \tag{27}$$

Value of Information and Cost of Uncertainty in Absence of Acquired Information In absence of acquired information ($q = \eta$), the uncertainty cost reduces to $H_p^{(\eta)} = H(X_t | X_{t-1})$. This cost is smaller than the uncertainty cost incurred in absence of inherited information, which was shown in the previous section to be $H(X_t)$. The difference is the mutual information $I(X_t; X_{t-1}) = H(X_t) - H(X_t | X_{t-1})$, which thus quantifies the value of inherited information in this context.

The uncertainty cost $H_p^{(\eta)} = H(X_t | X_{t-1})$ can also be interpreted as the entropy rate $\mathcal{H}(X)$ of the environmental process: denoting $X^T \equiv (X_1, \dots, X_T)$, the entropy rate of the environmental process X is generally defined by $\mathcal{H}(X) = \lim_{T \rightarrow \infty} H(X^T) / T$ [14]. The limit always exists for a stationary ergodic process and it corresponds to $\mathcal{H}(X) = H(X_t | X_{t-1})$ for a Markov chain, and $\mathcal{H}(X) = H(X_t)$ for an i.i.d. process.

Value of Information in Presence of Acquired Information In presence of both acquired and inherited information, it follows from (26) that the value of acquired information is given by

$$I_p^{(q)} \equiv H_p^{(\eta)} - H_p^{(q)} = H(X_t | X_{t-1}) - H(X_t | X_{t-1}, Y_t) \equiv I(X_t; Y_t | X_{t-1}), \tag{28}$$

where the last equality defines the conditional mutual information $I(X_t; Y_t | X_{t-1})$. Using a conditioned version of the general relation $I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$, this conditional mutual information can also be written $I(X_t; Y_t | X_{t-1}) = H(Y_t | X_{t-1}) - H(Y_t | X_t)$. It is instructive to compare this quantity with the rate of mutual information between the processes X and Y , which is defined by $\mathcal{I}(X; Y) = \lim_{T \rightarrow \infty} I(X^T; Y^T) / T$, where we use again the notations $Y^T = (Y_1, \dots, Y_T)$ and $X^T = (X_1, \dots, X_T)$. Given that $I(X^T; Y^T) = H(Y^T) - H(Y^T | X^T)$, the rate of mutual information corresponds here to $\mathcal{I}(X; Y) = \mathcal{H}(Y) - H(Y_t | X_t)$, where $\mathcal{H}(Y)$ represents the entropy rate for the process Y (as an hidden Markov chain derived from a stationary ergodic chain, Y has indeed a well-defined entropy rate). From $H(Y_t | X_t) \leq \mathcal{H}(Y)$, it follows that $I_p^{(q)} = I(X_t; Y_t | X_{t-1}) \leq \mathcal{I}(X; Y)$, where the inequality is generically strict if the environmental process is not i.i.d. The value

of acquired information, $I_p^{(g)}$, is thus *not* given by the rate of mutual information $\mathcal{I}(X; Y)$, except in special cases such as when no correlations are present between successive environmental states (i.i.d. environment).

$I_p^{(g)}$ does not indeed correspond to the rate of mutual information, but to the rate of directed information [24, 25, 40], generally defined by

$$I(Y^T \rightarrow X^T) \equiv \sum_{t=1}^T I(X_t; Y^t | X^{t-1}). \tag{29}$$

For a Markov environmental process, conditioning with respect to X^{t-1} is equivalent to conditioning with respect to X_{t-1} and $I(X_t; Y^t | X_{t-1}) = I(X_t; Y_t | X_{t-1})$, so that the generic term of the sum equates the conditional mutual information obtained in (28). If $\mathcal{I}(Y \rightarrow X) \equiv \lim_{T \rightarrow \infty} I(Y^T \rightarrow X^T)/T$ denotes the rate of directed information, we have therefore $I_p^{(g)} = \mathcal{I}(Y \rightarrow X)$. To understand the origin of the difference between $\mathcal{I}(Y \rightarrow X)$ and the rate of mutual information $\mathcal{I}(X; Y)$, we may similarly expand the mutual information $I(X^T; Y^T)$ using the chain rule [14]:

$$I(X^T; Y^T) = \sum_{t=1}^T I(X_t; Y^t | X^{t-1}). \tag{30}$$

In this expression, Y^t in (29) is replaced by $Y^T = (Y^t, Y_{t+1}, \dots, Y_T)$. Consequently, $I(Y^T \rightarrow X^T) \leq I(X^T; Y^T)$ and the difference may be interpreted as the information that would be gained about the current environmental state x_t from knowing the future signals y_{t+1}, \dots, y_T ; these signals are indeed informative about x_t , since they are correlated to x_t through x_{t+1}, \dots, x_T , although they are not accessible at time t for a strategy $\pi(\sigma | x_{t-1}, y_t)$ which relies only on the current signal y_t (keeping memory of the past signals y^{t-1} does not make a difference in the present context where x_{t-1} is available). The mutual information $I(X^T; Y^T)$ thus accounts for all statistical correlations between X^T and Y^T , while the directed information $I(Y^T \rightarrow X^T)$ accounts only for the correlations that are consistent with the constraints of causality imposed on $\pi(\sigma | x_{t-1}, y_t)$. Consistently with this interpretation, the difference $I(X^T; Y^T) - I(Y^T \rightarrow X^T)$ can be shown to be $I(X^{T-1} \rightarrow Y^T)$.

Cost of Uncertainty in Presence of Acquired Information Similarly, the uncertainty cost in presence of side-information, given in (26), does not correspond to the rate $\mathcal{H}(X|Y)$ of the conditional entropy $H(X^T|Y^T)$, but instead to the rate $\mathcal{H}(X\|Y)$ of the causally conditional entropy $H(X^T\|Y^T)$ [40, 47], which is generally defined by

$$H(X^T\|Y^T) \equiv \sum_{t=1}^T H(X_t | X^{t-1}, Y^t). \tag{31}$$

For comparison, $H(X^T|Y^T)$ can be similarly expressed with Y^t replaced by Y^T in each term of the sum, thus showing that $H(X^T\|Y^T) \geq H(X^T|Y^T)$. In the context of our model, $H(X_t | X^{t-1}, Y^t) = H(X_t | X_{t-1}, Y_t)$, and hence (26) indicates that $H_p^{(g)} = \mathcal{H}(X\|Y)$. As the conditional entropy is related to the mutual information by $I(X^T; Y^T) = H(X^T) - H(X^T|Y^T)$, the causally conditional entropy is related to the directed information by $I(Y^T \rightarrow X^T) = H(X^T) - H(X^T\|Y^T)$, or, in terms of rates, $\mathcal{I}(Y \rightarrow X) = \mathcal{H}(X) - \mathcal{H}(X\|Y)$.

Conclusion The conclusions that the uncertainty cost is given by the rate of a causally conditional entropy, which is greater than the rate of a conditional entropy,

$$H_p^{(q)} = \mathcal{H}(X\|Y) \geq \mathcal{H}(X|Y), \quad (32)$$

and the value of acquired information by the rate of directed information, which is smaller than the rate of mutual information,

$$I_p^{(q)} = \mathcal{I}(Y \rightarrow X) \leq \mathcal{I}(X; Y), \quad (33)$$

can be extended beyond Markov processes to more general ergodic stochastic processes, provided one allows for an arbitrary long memory, i.e., strategies of the form $\pi(\sigma|x^{t-1}, y^t)$ [40]. More generally, the notion of directed information appears as the relevant generalization of the notion of mutual information when causal relations, and not merely statistical relations, must be taken into account [24, 25]; for instance, while the capacity of memoryless channels is expressed in terms of a mutual information, the capacity of channels with feedback involves a directed information [48].

Coming back to our model, in absence of the simplifying assumption (A3) that the multiplication matrix $f(\sigma; x)$ is diagonal, problems involving both acquired and inherited information are generally difficult to solve; in particular, no closed-form expression for the growth rate generalizing (8) is available. Horse race models are an exception, due to the fact that the history of past types of any individual mirrors the history of past environmental states, only individuals with $\sigma_t = x_t$ survive at time t . This reduces the problem to an effectively feedforward problem, where $\pi(\sigma_t|\sigma_{t-1}, y_t)$ does not actually depend on the “control variable” σ_{t-1} , but only on the “primary variables” x_{t-1} and y_t . An other solvable case, for essentially the same reason, is the limit where any given environmental state lasts long enough for a single type to dominate the population [33]: we show in Appendix E how the problems of delay and timing that generally arise in correlated environments with inherited information can be treated in this case.

6 Relaxing (A2): Individual Stochasticity and Distributed Information

Retaining the assumptions (A1) and (A3) but now relaxing (A2) by allowing each individual in the population to perceive a different signal from the environment leads to a different generalization of the definitions of entropy and mutual information. To analyze the differences between common and individual information, we refine here the model of Fig. 1 by distinguishing two levels of information processing: a level common to the population and a level specific to individuals. This is achieved by decomposing the communication channels $q(y_t|x_t)$ in two consecutive communication channels, as represented in Fig. 2: a first communication channel attached to the environment, $q_{\text{env}}(x'_t|x_t)$, whose output is a cue x'_t common to all individuals in the population, and a second communication channel attached to each individual, $q_{\text{in}}(y_t|x'_t)$, whose output is the signal y_t (see also Fig. 3).⁴ For instance, if considering a population of bacteria, x_t may represent the chemicals constituting the medium at time t , x'_t the subset of those chemicals for which the bacteria have a sensor, and y_t the chemicals that a particular bacterium actually detects at time t , which may vary from bacterium to bacterium due to imperfect sensors. The difference between x_t , the environmental

⁴The decomposition into x'_t and y_t , although not unique, cannot be arbitrary.

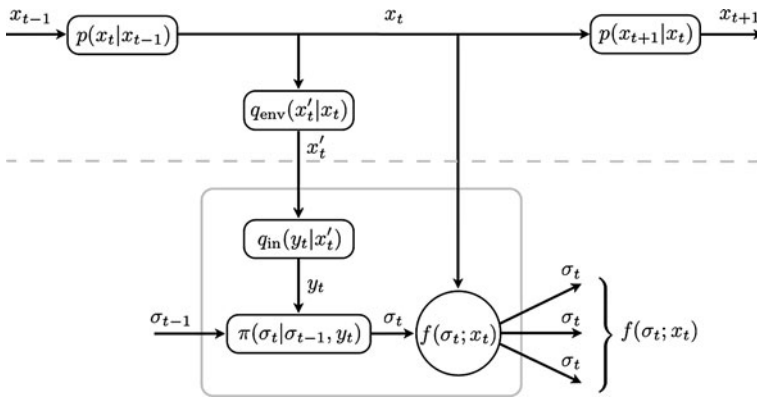


Fig. 2 Same model as in Fig. 1, but with the sensory channel $q(y_t|x_t)$ decomposed in two distinct channels: an “environmental” channel $q_{env}(x'_t|x_t)$, which specifies a perceptible signal x'_t common to the whole population, and an “individual” channel $q_{in}(y_t|x'_t)$, which specifies, independently for each individual, a perceived signal y_t . This second channel is noiseless in the financial interpretation of the model for which no stochasticity is present at the individual level; in this case $y_t = x'_t$ (noiseless individual channel $q_{in} = \delta$). In general, the channel $q(x_t|x_t)$ of Fig. 1 is related to q_{env} and q_{in} by $q = q_{env} * q_{in}$, i.e., $q(y_t|x_t) = \sum_{x'_t} q_{in}(y_t|x'_t)q_{env}(x'_t|x_t)$

state affecting the multiplication rate $f(\sigma_t; x_t)$, and x'_t , the environmental cue, may also represent a delay between sensing and reproduction.⁵

Financial Interpretation The decomposition of the channel of acquired information into an environmental channel $q_{env}(x'_t|x_t)$ and an individual channel $q_{in}(y_t|x'_t)$ is further motivated by the financial interpretation of our model. In this interpretation, $\mathcal{N}_t(\sigma_t)$ represents the number of currency units that an investor on the stock market invests in asset σ_t on day t , $\pi(\sigma_t|\sigma_{t-1}, y_t)$ represents the fraction of money transferred from asset σ_{t-1} to asset σ_t , based possibly on some information y_t available about the current state x_t of the market, and $f(\sigma; x_t)$ represents the return of asset σ on day t , a non-negative but non-necessarily integer quantity (see Table 1). Equations (3) and (4) then describe the evolution of \mathcal{N}_t in a scenario where the money is entirely reinvested every time. The essential difference with the biological case is the absence of stochasticity at the level of individuals, which are strictly identical currency units: $\pi(\sigma_t|\sigma_{t-1}, y_t)$ results in finance from the decision of an investor who centralizes the information used for manipulating each of the currency unit constituting the “population”. This has two implications: (i) y_t has to be common to the population, i.e., $q_{in} = \delta$ and as a result $y_t = x'_t$; (ii) the only source of stochasticity is the environment, which operates at the level of the population,⁶ i.e., $Z_t = \mathcal{N}_t$. In contrast, the necessity for biological populations to process information at the level of individual organisms introduces an extra level of stochasticity and heterogeneity, which underlies the qualitative differences with problems of financial investment examined in this section.

⁵For instance, we may consider $x_t = (e_t, e_{t-1})$ and $x'_t = e_{t-1}$, with e_t described by a Markov chain with transition matrix $b(e_t|e_{t-1})$ and $p(x_t|x_{t-1}) = b(e_t|e_{t-1})$.

⁶Note however that the finite divisibility of the currency unit is not accounted for if considering only \mathcal{N}_t .

Table 1 Correspondence between the terms used in biology, finance and control theory. The engineering problem is concerned with a single system, and therefore involves no multiplication rate or population. The financial problem is defined for an agent supervising any information processing, and the notion of information inherited by the individuals has therefore no obvious counterpart

Biology	Finance	Control theory	Notation
individual	currency unit	system	–
population	capital	–	–
phenotype	asset	state	σ_t
environment	market	disturbance	x_t
multiplication rate	return	1	$f(\sigma_t, x_t)$
acquired info.	side-information	feedforward information	y_t
inherited info.	–	feedback information	σ_{t-1}
strategy	portfolio	policy	$\pi(\sigma_t \sigma_{t-1}, y_t)$
fitness	utility	loss-function	$\Lambda(\pi)$

	<i>acquired</i>	<i>inherited</i>
<i>population</i>	perceptible x'_t	transition matrix π
<i>individual</i>	perceived y_t	type σ_t

Fig. 3 Four different notions of “information” contained in the model. Information has two sources, the environment and the ancestor of the individual, corresponding to acquired or inherited information, and is defined at two levels, the individual or the population. The transition matrix π may be viewed as information about the environment encoded in the organisms (see e.g. (11)). If mutations could lead to the unreliable transmission of π , an extra level of description would be introduced, with subpopulations characterized by different values of π

Cost of Uncertainty When q is decomposed into q_{env} and q_{in} , the expression for the growth rate, (8), becomes

$$\Lambda_{p:f}^{(q_{env}, q_{in})}(\pi) = \sum_{x, x'} q_{env}(x'|x) p_s(x) \ln \left(\sum_y f(x) \pi(x|y) q_{in}(y|x') \right). \tag{34}$$

Following the derivation presented in Sect. 4, its optimal value can be decomposed in two terms,

$$\hat{\Lambda}_{p:f}^{(q_{env}, q_{in})} = \hat{\Lambda}_{p:f}^{(\delta, \delta)} - H_p^{(q_{env}, q_{in})}. \tag{35}$$

The second term is again independent of f ,

$$H_p^{(q_{env}, q_{in})} \equiv \min_{\pi} \sum_{x, x'} q_{env}(x'|x) p_s(x) \ln \left(\sum_y \pi(x|y) q_{in}(y|x') \right)^{-1}, \tag{36}$$

and generalizes the notions of entropy $H_p^{(\eta, \delta)} = H(X_t)$ and conditional entropy $H_p^{(q_{env}, \delta)} = H(X_t | Y_t)$ obtained for horse races in Sect. 4.⁷ From the concavity of the logarithm (Jensen’s

⁷ $H_p^{(q_{env}, q_{in})}$ can also be written




Fig. 4 Models of communication channels for q_{env} and q_{in} when the environment alternates between two states 1 and 2. *Left*: the binary erasure channel q_e (BEC), which erases the input with probability ϵ and transmits it faithfully with probability $1 - \epsilon$. *Right*: the binary symmetric channel q_s (BSC), which exchanges the input with probability ϵ and transmits it unchanged with probability $1 - \epsilon$. The channels δ and η represent extreme cases: the noiseless channel δ corresponds to the BEC or the BSC with $\epsilon = 0$, and the informationless channel η to the BEC with $\epsilon = 1$, or the BSC with $\epsilon = 1/2$

inequality),

$$H(X_t|X'_t) \leq H_p^{(q_{env}, q_{in})} \leq H(X_t|Y_t), \tag{37}$$

where, following the usual notations, X'_t refers to the random variable for the component x'_t of the signal defined at the population level, and Y_t to the random variable for the signal y_t effectively perceived by an individual.

As an illustration of the properties of this generalized entropy, showing in particular that, generically, $H_p^{(q_{env}, q_{in})} < H(X_t|Y_t)$, we compare in Figs. 5 and 6 the benefits of the same channel q located either at the population level, $q_{env} = q$, $q_{in} = \delta$, or at the individual level, $q_{env} = \delta$, $q_{in} = q$, taking for q two classical examples of communication channels defined in Fig. 4 (the details of the calculations are presented in Appendix D).

Value of Information The fact apparent in Figs. 5–6 that the same communication channel q induces less uncertainty when located at an individual level than at a population level, i.e., $\hat{\Lambda}_{p:f}^{(q, \delta)} \leq \hat{\Lambda}_{p:f}^{(\delta, q)}$, holds generally, again as a consequence of Jensen’s inequality:

$$\hat{\Lambda}_{p:f}^{(q_{in} * q_{env}, \delta)} \leq \hat{\Lambda}_{p:f}^{(q_{env}, q_{in})} \leq \hat{\Lambda}_{p:f}^{(\delta, q_{in} * q_{env})}, \tag{38}$$

where $q_{in} * q_{env}$ denotes the convolution of q_{in} and q_{env} , i.e., $q_{in} * q_{env}(y|x) = \sum_{x'} q_{in}(y|x')q_{env}(x'|x)$. An important implication is that the mutual information between the source X_t and the perceived signal Y_t does *not* represent an upper bound for the value of acquired information. From the relation $H_p^{(q_{env}, q_{in})} + I_p^{(q_{env}, q_{in})} = \hat{\Lambda}_{p:f}^{(\delta, \delta)} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$, we have indeed a relation dual to (37) for the value of information $I_p^{(q_{env}, q_{in})} \equiv \hat{\Lambda}_{p:f}^{(\eta, \eta)} - \hat{\Lambda}_{p:f}^{(q_{env}, q_{in})}$:

$$I(X_t; Y_t) \leq I_p^{(q_{env}, q_{in})} \leq I(X_t; X'_t). \tag{39}$$

Informally, we may say that the value of the information acquired collectively by the population exceeds the value of the information acquired by any of its members. This result contrasts with the law of requisite variety derived in other contexts which states that the mu-

$$H_p^{(q_{env}, q_{in})} = H(X_t|X'_t) + \min_{\pi} \mathbb{E}_{X'_t} \left[D \left(\mathbb{P}_{X_t|X'_t}(x|X'_t) \parallel \pi * q_{in}(x|X'_t) \right) \right],$$

where $\pi * q_{in}(x|x') = \sum_y \pi(x|y)q_{in}(y|x')$ and where $\mathbb{P}_{X_t|X'_t}(x|x')$ is the optimal strategy for the same problem where $q_{in} = \delta$, i.e., $\mathbb{P}_{X_t|X'_t}(x|x') = q_{env}(x'|x)p(x) / (\sum_z q_{env}(x'|z)p(z))$.

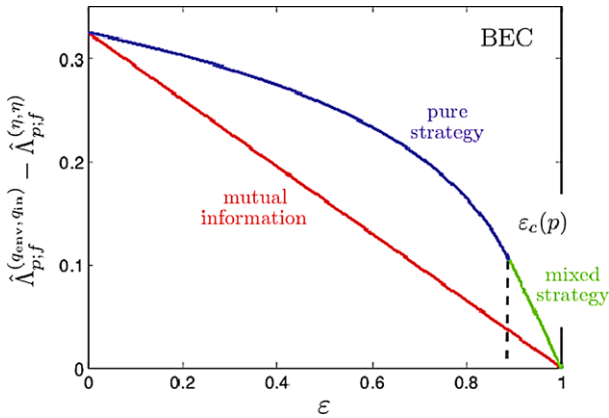


Fig. 5 (Color online) Value of information for a two-state model with information transmitted through a binary erasure channel situated either at the population level (*red curve*) or at the individual level (*blue and green curve*). The probability of the environmental state 1 is here fixed to $p(1) = 0.1$. The binary erasure channel q_e , defined in Fig. 4, has a probability of erasure ϵ which is varied along the x -axis. When the channel is at the population level, the value of information $\hat{\Lambda}_{p:f}^{(q_e, \delta)} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$ corresponds to the mutual information between the input and output signals which, for the binary erasure channel, is a linear function of ϵ (*red line*). When the same channel is at the individual level, the value of information $\hat{\Lambda}_{p:f}^{(\delta, q_e)} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$ is generally higher (*blue and green curve*), and a transition occurs at $\epsilon_c(p) = (1 - 2p(1))/(1 - p(1)) = 0.88$: for $\epsilon < \epsilon_c(p)$, the optimal strategy is a pure strategy with $\hat{\pi}(1|*) = 0$ (*blue part*), while for $\epsilon > \epsilon_c(p)$, it becomes a mixed strategy with both $\hat{\pi}(1|*) > 0$ and $\hat{\pi}(2|*) > 0$ (*green part*). The calculations are detailed in Appendix D.2

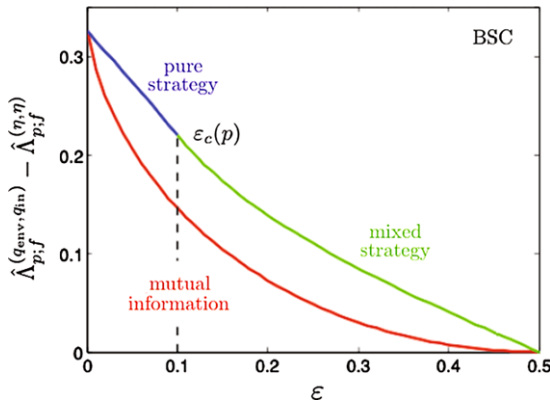


Fig. 6 (Color online) Value of information for a two-state model with information transmitted through a binary symmetric channel situated either at the population level (*red curve*) or at the individual level (*blue and green curve*). The probability of the environmental state 1 is here fixed to $p(1) = 0.1$. The binary symmetric channel q_s , defined in Fig. 4, has a probability of error ϵ which is varied along the x -axis. When the channel is at the population level, the value of information $\hat{\Lambda}_{p:f}^{(q_s, \delta)} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$ corresponds to the mutual information between the input and output signals (*red curve*). When the same channel is at the individual level, the value of information $\hat{\Lambda}_{p:f}^{(\delta, q_s)} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$ is generally higher (*blue and green curve*), and a transition occurs at $\epsilon_c(p) = p(1) = 0.1$: for $\epsilon < \epsilon_c(p)$, the optimal strategy is to adopt a pure strategy when receiving either of the two possible signals 1 and 2, i.e., $\hat{\pi}(1|1) = \hat{\pi}(2|2) = 1$ (*blue part*), while for $\epsilon > \epsilon_c(p)$ the optimal strategy is a mixed strategy with $\hat{\pi}(2|2) = 1$ but $0 < \hat{\pi}(1|1) < 1$ (*green part*). The calculations are detailed in Appendix D.3

tual information $I(X_t; Y_t)$ between an environmental fluctuation X_t and the signal Y_t sets an upper limit on the value of information for control [20, 21]. In comparison with the mutual information $I(X_t; Y_t)$, $I_p^{(q_{env}, q_{in})}$ is not symmetrical in X_t and Y_t , although it similarly satisfies $I_p^{(q_{env}, q_{in})} = 0$ if and only if X_t and Y_t are independent.

Optimal Strategy and Bayesian Inference Another remarkable feature displayed in Figs. 5 and 6 is the possible existence of a critical level of noise $\varepsilon_c(p)$ below which a stochastic response is not required for achieving optimal growth. This contrasts with the horse race model of Sect. 4, where a non-deterministic response is required not only to achieve an optimal growth, but even more fundamentally to avoid extinction. Here, the diversification of the population caused by the deterministic response of individuals perceiving stochastic signals is optimal at low error rates. Although estimation and decision can be separated in principle [49], and although a Bayesian computation, as in (18), would provide an optimal estimation, the simplest implementation of the optimal strategy involves here no computation at all: when $\varepsilon < \varepsilon_c(p)$, the individual can process the signal as if it was perfectly reliable.⁸ This situation is analogous to the situation with optimal source-channel communication: although in principle a solution can always be obtained by treating separately the problems of source compression and channel coding [12], a computationally much simpler solution may be available, which in some cases does not involve any coding at all [51]. Living systems are unlikely to solve stochastic control problems by relying on the estimation-decision separation principle, as they are unlikely to solve communication problems by relying on the source-channel separation principle [8].

7 Relaxing (A3): General Multiplication Rates and Functional Information

Retaining the assumption (A1) of no inherited information but relaxing assumption (A3) that the multiplication rates $f(\sigma; x)$ have a diagonal form leads to a different departure from the usual concepts of communication theory. Now the cost of uncertainty $\hat{\Lambda}_{p:f}^{(\delta)} - \hat{\Lambda}_{p:f}^{(q)}$ and the value of acquired information $\hat{\Lambda}_{p:f}^{(q_{env}, q_{in})} - \hat{\Lambda}_{p:f}^{(\eta, \eta)}$ are no longer necessarily independent of the multiplication rates $f(\sigma; x)$, and cannot therefore be written as statistical quantities $H_p^{(q_{env}, q_{in})}$ or $I_p^{(q_{env}, q_{in})}$ depending only on the transition matrices $p(x_t|x_{t-1})$, $q_{env}(x'_t|x_t)$ and $q_{in}(y_t|x'_t)$.

Uncertainty Cost A very special feature of models satisfying assumption (A3), i.e., models where the multiplication rates have a diagonal form, with $f(\sigma; x) = f(x)$ if $\sigma = x$ and 0 otherwise, is that the environmental states x and the individual types σ are in one-to-one correspondence. The environment is however generally defined independently of any reference to the internal states of the individuals of the population. We should therefore not expect a quantity like the entropy rate of the environmental process, $\mathcal{H}(X)$, to correctly capture the cost of uncertainty, which depends essentially on the definition of the internal states of the individuals. The environmental states may indeed specify details that are irrelevant to the growth of the population, say the positions of distant stars, which inflate arbitrarily the entropy rate $\mathcal{H}(X)$ without influencing the uncertainty cost $\hat{\Lambda}_{p:f}^{(\delta, \delta)} - \hat{\Lambda}_{p:f}^{(q_{env}, q_{in})}$. As an example, consider an horse race where each distinct environmental state corresponds to a

⁸The optimal strategy does not require a Bayesian computation, but it nevertheless follows a Bayesian logics, in the sense of the word given in stochastic adaptive control theory [50].

distinct ordered list of arrival of all the horses participating to the race; this description includes useless information if only the first horse has a non-zero pay-off.⁹ More generally, the uncertainty cost ignores any stochastic element of the environment that is irrelevant for the growth of the population, but nevertheless contributes to the entropy rate $\mathcal{H}(X)$.

In addition, when several types σ have non-zero multiplication rates $f(\sigma; x) > 0$ in a given environmental state x , the non-optimal but yet surviving types contribute to the growth although they are not associated with the exact prediction of the optimal type σ for the given environment x . Again, this implies that an entropic measure based only on the environmental process tends to overestimate the uncertainty cost. We show in Appendix F that the following bound holds:

$$\hat{\Lambda}_{p;f}^{(\delta,\delta)} - \hat{\Lambda}_{p;f}^{(q_{env},q_{in})} \leq H_p^{(q_{env},q_{in})}. \tag{40}$$

Here, the generalized entropy $H_p^{(q_{env},q_{in})}$ is the uncertainty cost for a horse race model with same channels $p(x_t|x_{t-1})$, $q_{env}(x'_t|x_t)$ and $q_{in}(y_t|x'_t)$. As defined in the previous sections, this generalized entropy is independent of the value of the multiplication rates f . The quantity $\hat{\Lambda}_{p;f}^{(\delta,\delta)} - \hat{\Lambda}_{p;f}^{(q_{env},q_{in})}$ can be seen as a measure of uncertainty that refines $H_p^{(q_{env},q_{in})}$ by accounting for the effective reduction of uncertainty due to the redundancy between environments and types encoded in f . This is a further refinement over the concept of entropy $H_p^{(q_{env},q_{in})}$, which itself can be seen as refining Hartley’s measure $\ln n$ by accounting for the effective reduction of uncertainty due to the unequal probabilities of the different environmental states. In the two cases, the refinement takes the form of an inequality, with equality if f is diagonal in the first case, and if all the environmental states are equiprobable in the second case.

Value of Information A corresponding inequality holds for the value of information, with

$$\hat{\Lambda}_{p;f}^{(q_{env},q_{in})} - \hat{\Lambda}_{p;f}^{(\eta,\delta)} \leq I_p^{(q_{env},q_{in})}. \tag{41}$$

In particular, under the assumptions (A1) and (A2) such that $I_p^{(q_{env},\delta)}$ is given by the mutual information $I(X_t, Y_t)$, the value of information is bounded by $I(X_t, Y_t)$. The deviation of $\hat{\Lambda}_{p;f}^{(q_{env},\delta)} - \hat{\Lambda}_{p;f}^{(\eta,\delta)}$ from $I_p^{(q_{env},\delta)}$, when the multiplication rates are non-diagonal, can be interpreted as arising from the fact that the environmental states have no longer an exclusive “meaning”, in the sense that the same environment can be beneficial to different types, and different environments to the same type. A noticeable feature of models with non-diagonal multiplication matrices is also that the optimal strategy may actually exclude some types σ , i.e., we may have $\hat{\pi}(\sigma|y) = 0$ for some σ . A trivial example is when two types σ and σ' are present, for which $f(\sigma; x) > f(\sigma'; x)$ in any environmental state x , in which case the optimal strategy will never populate σ' . A less trivial, yet analytically solvable class of models which display the same feature is defined by extending (A3) to the case where the off-diagonal terms of the matrix $f(\sigma; x)$ are non-zero but constant, i.e., $f(\sigma; x) = f(x)$ for $\sigma = x$ and $f(\sigma; x) = \tilde{f}(x) < f(x)$ for $\sigma \neq x$ (see Appendix D); in particular for $n = 2$ states, the model is solvable for arbitrary matrices f , as illustrated in Figs. 7 and 8. An-

⁹In such a case, we may still capture the uncertainty cost by a statistical quantity by partitioning the environmental states into exclusive sets $\mathcal{X}(\sigma)$ grouping the lists where horse σ is first, such that $f(\sigma; x) = f(\sigma)$ if $x \in \mathcal{X}(\sigma)$ and 0 otherwise: assuming i.i.d. races, the uncertainty cost then correspond to the entropy of the coarse-grained description $H(\mathcal{X}_t)$ rather than the entropy $H(X_t)$, with obviously $H(X_t) \geq H(\mathcal{X}_t)$ (see also Appendix F).




Fig. 7 (Color online) Nature of the optimal strategy for a model with 2 types and 2 environmental states whose probabilities $p(1)$ and $p(2) = 1 - p(1)$ are varied along the x -axis, and with multiplication rates whose ratio $\gamma = f(1; 1)/f(2; 1) - 1 = f(2; 2)/f(1; 2) - 1$ is varied along the y -axis (this ratio may be thought as quantifying the dissimilarity between the two types). No information is assumed to be available. The blue curves delineate the regions of the parameter space where the optimal strategy involves switching, with $0 < \hat{\pi}(1) < 1$, from the regions where one of the two types is excluded from the optimal strategy, corresponding to homogeneous populations, with $\hat{\pi}(1) = 1$ or $\hat{\pi}(2) = 1$. The location of the transitions is given by $\gamma_c^{(1)} = (1 - p(1))/p(1) - 1$ and $\gamma_c^{(2)} = p(1)/(1 - p(1)) - 1$. The calculations are detailed in Appendix D.1




Fig. 8 (Color online) In green: Optimal growth rate $\hat{\Lambda}_p^{(\eta, \eta)}$ for a model in absence of information with 2 types and 2 environmental states, as a function of the probability $p(1)$ of the first environmental state. The values of multiplication rates are $f(1; 1) = 3$, $f(2; 2) = 2$, $f(2; 1) = 2$, and $f(1; 2) = 1$. The dashed lines represent transitions between mixed and pure strategies: for $p(1) < p_c^{(1)} = 0.2$ the type $\sigma = 1$ is excluded from the optimal strategy, while for $p(1) > p_c^{(2)} = 0.6$ this is the case for $\sigma = 2$. In red: Optimal growth rate $\hat{\Lambda}_{p;f}^{(\delta, \delta)}$ for the same model in presence of complete information. In blue: Optimal growth rate $\hat{\Lambda}_{p;f}^{(\delta, \delta)} - H(X_t)$ for a corresponding horse race model with $f(1; 1) = 3$, $f(2; 2) = 2$ and $f(2; 1) = f(1; 2) = 0$ and no information. This example shows that the uncertainty cost $\hat{\Lambda}_p^{(\delta, \delta)} - \hat{\Lambda}_p^{(\eta, \eta)}$ (difference between the red and green curves) is generally strictly smaller than the entropy of the environment $H(X_t)$ (difference between the red and blue curves). The calculations are detailed in Appendix D.1

Table 2 Expressions for the value of information and cost of uncertainty under different assumptions. The top row corresponds to the most general model and each subsequent row involves an additional assumption, indicated in the second column. The last row thus defines the most restrictive model, which is the horse model from which we started in Sect. 4. We then presented the implications of relaxing successively the various assumptions that it involves, thus moving up in this table. The different measures of information are related by $\hat{\Lambda}_{p;f}^{(q_{env},q_{in})} - \hat{\Lambda}_{p;f}^{(\eta,\eta)} \leq I_p^{(q_{env},q_{in})} \leq \mathcal{I}(X' \rightarrow X) \leq I(X_t; X'_t)$ and $\mathcal{I}(Y \rightarrow X) \leq I_p^{(q_{env},q_{in})}$, where $\mathcal{I}(X' \rightarrow X) = I_p^{(q_{env},\delta)}$ and $\mathcal{I}(Y \rightarrow X) = I_p^{(q_{env}*q_{in},\delta)}$: see (33), (38), (41). Similarly, the different measures of uncertainty are related by $\hat{\Lambda}_{p;f}^{(\delta,\delta)} - \hat{\Lambda}_{p;f}^{(q_{env},q_{in})} \leq H_p^{(q_{env},q_{in})} \leq H(X_t|Y_t) \leq \mathcal{H}(X\|Y)$ and $\mathcal{H}(X\|X') \leq H_p^{(q_{env},q_{in})}$, where $\mathcal{H}(X\|X') = H_p^{(q_{env},\delta)}$ and $\mathcal{H}(X\|Y) = H_p^{(q_{env}*q_{in},\delta)}$: see (32), (37), (40)

Assumptions	Extra assumption	Value of information	Cost of uncertainty	Sect.
	–	$\hat{\Lambda}_{p;f}^{(q_{env},q_{in})} - \hat{\Lambda}_{p;f}^{(\eta,\eta)}$	$\hat{\Lambda}_{p;f}^{(\delta,\delta)} - \hat{\Lambda}_{p;f}^{(q_{env},q_{in})}$	VII
(A3)	no survival for $\sigma_t \neq x_t$			
	$f(\sigma_t; x_t) = f(x_t)\delta(x_t \sigma_t)$	$I_p^{(q_{env},q_{in})}$	$H_p^{(q_{env},q_{in})}$	VI
(A2) (A3)	no individuality	directed information	causally conditional entropy	
	$q_{in}(y_t x'_t) = \delta(y_t x'_t)$	$\mathcal{I}(Y \rightarrow X)$	$\mathcal{H}(X\ Y)$	V
(A1) (A2) (A3)	no feedback	mutual information	conditional entropy	
	$\pi(\sigma_t \sigma_{t-1}, y_t) = \pi(\sigma_t y_t)$	$I(X_t; Y_t)$	$H(X_t Y_t)$	IV

other important solvable class of models is when a separation of time scales allows for an adiabatic approximation, as presented in Appendix E.¹⁰

General Conclusion To sum up the results of the last three sections, the relaxations of the assumptions (A1), (A2) and (A3) lead to generalizations of the notions of entropy and mutual information in three different directions: (i) to account for the constraints of causality (Sect. 5); (ii) to account for the level at which information is processed (Sect. 6); (iii) to account for the meaning of information encoded in the matrix $f(\sigma; x)$ (this section). In the cases (i) and (iii), which had been previously studied from the standpoint of financial investment, the mutual information appears as an upper limit for the value of acquired and inherited information, consistently with Ashby’s law of requisite variety [20]; yet this limit cannot generally be reached, and a tighter and achievable upper bound is provided by $I_p^{(q_{env},q_{in})}$. In the case (ii), which is specific to the biological interpretation of the model, the fundamental limit $I_p^{(q_{env},q_{in})}$ can be greater than the mutual information¹¹ $I(X_t; Y_t)$. In general, all three assumptions (A1), (A2) and (A3) may be jointly violated, and the uncertainty cost and value of information need to be measured accordingly. These conclusions are summarized in Table 2.

¹⁰Yet another interesting subclass of models is when $f(\sigma; x) = g(x)\phi(x|\sigma)$ where ϕ is a transition matrix, i.e., $\sum_x \phi(x|\sigma) = 1$; the growth rate can then be written $\Lambda_{p;f}^{(\eta,\delta)} = \mathbb{E}_X[\ln g(X)] - H(X) - D(p\|\phi * \pi)$ and the optimal strategy is given by the minimization of $D(p\|\phi * \pi)$. The problem to be solved to find the optimal strategy $\hat{\pi}$ is then equivalent to a problem of blind source identification, i.e., the problem of inferring the source π of the inputs of a communication channel ϕ given the distribution p of its outputs.

¹¹Ashby’s law of requisite variety can be rescued by noticing that it should be applied, not to an individual but to the population. Indeed, the population as a whole has only access to X'_t , and (39) verifies that the benefits conferred by this signal is limited by $I(X_t; X'_t)$, in agreement with Ashby’s law.




Fig. 9 An alternative representation to Fig. 2, where the strategy $\pi(\sigma_t|\sigma_{t-1}, y_t)$ is viewed as a controller. The controller receives both feedforward information y_t through the sensor $q_{in}(y_t|x'_t)$, and feedback information σ_{t-1} , subject to delay (see also Table 1)

The problem of measuring the degree of adaptation of a population with given communication channels q_{env} and q_{in} can be treated as well. As shown in Appendix G, the identity involving the relative entropy that emerged from the analysis of horse race models in Sect. 4, (14), is more generally replaced by an inequality.

8 Generalizations

Regulation is a general requirement for the sustainability and optimization of systems facing uncertainties. It forms the core issue of control in engineering, where acquired information is referred to as feedforward information and inherited information as feedback information (see Table 1 and Fig. 9). Quantifying the value of limited information is a long-standing open conceptual problem in control theory [23, 49]. For growing populations, the law of large numbers and its extension, ergodicity, make the problem well-posed by introducing in the long-term, or infinite horizon limit in the language of control theory, an unambiguous loss-function, the growth rate of the population (see Sect. 3 and Appendix B). Uncertainties are however generally not only due to limited information, and regulation must typically be made in presence of other constraints. These constraints can generally be classified in three categories: (i) constraints on estimation, i.e., on the acquisition of information about the current internal and external states, σ_{t-1} and x_t , with the constraints on acquired information considered so far being an example; (ii) constraints on decision, i.e., on the computation of σ_t from σ_{t-1} and y_t ; (iii) constraints on actuation, i.e., on the implementation of the switch from σ_{t-1} to σ_t . Biological constraints of the later type for instance arise when the types correspond to different developmental stages, in which case constraints of irreversibility are common.¹² While constraints on the organisms limit the ability of a population to control its growth rate, it is interesting to notice that constraints on the environment, such as constraints

¹²Such constraints may be taken into account in our model by specifying a graph whose nodes are the types σ_t and whose links are the possible transitions. An age structure can for instance be enforced by constraining the transitions matrices to have the form of Leslie matrices [52]. More generally, the constraints may restrict the graph of connectable types, as considered for instance in [53]. Our model can also account for the presence of an unreliable “actuator” by constraining $\pi(\sigma_t|\sigma_{t-1}, x_t)$ to be of the form $\pi(\sigma_t|\sigma_{t-1}, y_t) = \sum_{\sigma'_t} a(\sigma_t|\sigma'_t)\pi(\sigma'_t|\sigma_{t-1}, y_t)$; this assumes that the inherited type σ_t is an output of the actuator $a(\sigma_t|\sigma'_t)$, otherwise, if $\pi(\sigma_t|\sigma_{t-1}, y_t) = \sum_{\sigma'_t} a(\sigma'_t|\sigma_t)\pi(\sigma_t|\sigma_{t-1}, y_t)$ with σ'_t controlling the multiplication rate $f(\sigma'_t; x_t)$ but not being inherited, the problem becomes equivalent to a model with effective multiplication rate $\tilde{f}(\sigma_t, x_t) = \sum_{\sigma'_t} f(\sigma'_t; x_t)a(\sigma'_t|\sigma_t)$.




Fig. 10 Generalization of the models presented in Figs. 1 and 2, which correspond to particular forms of the transition matrix $\tau(\sigma_t|\sigma_{t-1}, x_t)$

on the possible states that may follow the current environmental state, render the future more predictable and have therefore the opposite effect of enhancing this ability.¹³

To encompass more general forms of constraints, our model can be extended to the model represented in Fig. 10, which considers a population with internal states σ_t and environmental states x_t described by

$$\mathcal{N}_t(\sigma_t) = f(\sigma_t; x_t) \sum_{\sigma_{t-1}} \tau(\sigma_t|\sigma_{t-1}, x_t) \mathcal{N}_{t-1}(\sigma_{t-1}), \tag{42}$$

where $\tau(\sigma_t|\sigma_{t-1}, x_t)$ is a transition matrix, and where the environment follows as before a Markov chain $p(x_t|x_{t-1})$. Imposing constraints on control formally amounts to restricting $\tau(\sigma_t|\sigma_{t-1}, x_t)$ to a subset \mathcal{C} of the set of conceivable transition matrices. Different “information patterns” [49], specifying “who knows what and when”, can thus be enforced. For instance, excluding feedback information corresponds to restricting $\tau(\sigma_t|\sigma_{t-1}, x_t)$ to the form $\tau(\sigma_t|x_t)$, and excluding feedforward information to restricting it to the form $\tau(\sigma_t|\sigma_{t-1})$. The model with constraints on acquired information presented in Sect. 6 can be formulated in this more general framework, by considering $\tilde{x}_t = (x_t, x'_t)$ for the environmental states, $\tilde{\sigma}_t = (\sigma_t, y_t)$ for the internal states, and $\tilde{p}(\tilde{x}_t|\tilde{x}_{t-1}) = q_{\text{env}}(x'_t|x_t)p(x_t|x_{t-1})$ for the transition matrix between environmental states; the transition matrix $\tau(\tilde{\sigma}_t|\tilde{\sigma}_{t-1}, \tilde{x}_t)$ must then be constrained to the form $\tau((\sigma_t, y_t)|(\sigma_{t-1}, y_{t-1}), (x_t, x'_t)) = \pi(\sigma_t|\sigma_{t-1}, y_t)q_{\text{in}}(y_t|x'_t)$, which defines a subset $\mathcal{C}^{(q_{\text{in}})}$ of admissible transition matrices. Several extensions of our model can similarly be formulated. For instance, the sensor q_{in} may be taken to depend on the type σ_t , thus allowing for different phenotypes to have different abilities to sense the environment: $\tau(\sigma_t|\sigma_{t-1}, x_t)$ must then be constrained to the form $\tau(\sigma_t|\sigma_{t-1}, x_t) = \sum_{y_t} \pi(\sigma_t|\sigma_{t-1}, y_t)q_{\text{in}}(y_t|x_t, \sigma_{t-1})$. Another possible extension is to consider that the types are transmitted with some errors by constraining $\tau(\sigma_t|\sigma_{t-1}; x_t)$ to the form $\tau(\sigma_t|\sigma_{t-1}; x_t) = \sum_{z_t} \pi(\sigma_t|z_t, x_t)\mu(z_t|\sigma_{t-1})$, where $\mu(z_t|\sigma_{t-1})$ represents a given “mutational” transition matrix.

The two basic questions (Q1) and (Q2) formulated in Sect. 3 can be addressed in this more general framework by taking again the growth rate $\Lambda_{p;f}(\tau)$ as a fitness function. A first point of comparison is provided by $\hat{\Lambda}_{p;f}$, the optimal growth rate in absence of any constraint, obtained after optimization over τ . The optimal transition matrix, $\hat{\tau}(\sigma_t|\sigma_{t-1}; x_t)$, is

¹³Another intriguing duality, between control and knowledge, was noted by Shannon: “we may have knowledge of the past and cannot control it; we may control the future but have no knowledge of it” [13].

easily characterized: at time t , it converts all the population to one of the types σ maximizing $f(\sigma; x_t)$, irrespectively of the type σ_{t-1} inherited from the previous generation. This optimal strategy is, however, generally excluded by the presence of constraints, characterized by the subset \mathcal{C} to which $\tau(\sigma_t|\sigma_{t-1}, x_t)$ must belong. Given \mathcal{C} , question (Q1) becomes the problem of finding a transition matrix $\hat{\tau}$ which maximizes $\Lambda_{p,f}(\tau)$ subject to the constraint $\tau \in \mathcal{C}$. This defines an optimal growth rate under constraints, $\hat{\Lambda}_{p,f}^{(\mathcal{C})}$. The arguments of the previous sections can then be repeated *mutatis mutandis*. For instance, under the assumptions (A1) and (A3), the solution $\hat{\tau}$ is independent of f and is associated with a generalization of Shannon’s entropy, $H_p^{(\mathcal{C})} \equiv \hat{\Lambda}_{p,f} - \hat{\Lambda}_{p,f}^{(\mathcal{C})}$, given by

$$H_p^{(\mathcal{C})} = \min_{\tau \in \mathcal{C}} \sum_x p(x) \ln \frac{1}{\tau(x|x)}. \tag{43}$$

More generally, this quantity provides an upper bound for the cost of the constraints, $\hat{\Lambda}_{p,f} - \hat{\Lambda}_{p,f}^{(\mathcal{C})}$. On the other hand, question (Q2) pertains to the value of relaxing a constraint \mathcal{C} to a lesser constraint $\mathcal{C}' \supset \mathcal{C}$, and amounts to estimating the quantity $I_{p:f}^{(\mathcal{C};\mathcal{C}')} = \hat{\Lambda}_{p:f}^{(\mathcal{C}')} - \hat{\Lambda}_{p:f}^{(\mathcal{C})}$, which generalizes the notion of mutual information $I_p^{(q_{env}, q_{in})}$ obtained when \mathcal{C} corresponds to the presence of the channels (q_{env}, q_{in}) , and \mathcal{C}' to the absence of any channel.

The major problem not addressed in the present framework is the specification of the constraints and, more broadly, the characterization of the costs for implementing any particular strategy. For instance, when analyzing the value of acquired information, not only should we take into account the benefit provided by the communication channel q_{in} , but also the cost for producing and operating it. This cost $c(q_{in})$ is to be measured in terms of growth rate, and its value will determine whether the sensor q_{in} has an adaptive value [53]. More generally, a trade-off between cost and accuracy will arise if $c(q_{in})$ is taken to be an increasing function of the accuracy of q_{in} . From this point of view, imposing constraints in the form of a subset \mathcal{C} of achievable transition matrices corresponds to assuming that some strategies have infinite costs while some other are cost-less. Costs are also generally present not only in the estimation step, but also in the decision and actuations steps; for instance, there may be a cost for switching between types, as there are transaction costs in finance [54].

Several other extensions can also be considered to explore other features of regulation in biological populations. For instance, from the standpoint of understanding the origin of diversification in a population, a key aspect of biological environments is their spatial heterogeneities. This feature may be incorporated at a mean-field level (not taking into account any geometrical properties of space) by making not only the acquired information y_t specific to individuals, but also the environmental factor z_t affecting their multiplication rates. We may thus assume that a “micro-environment” (y_t, z_t) derives independently for each individual from the “macro-environment” (x_t, x'_t) , through a transition matrix $v(y_t, z_t|x_t, x'_t)$ attached to each individual. The dynamics of the population is then described by

$$\mathcal{N}_t(\sigma_t) = \sum_{z_t} f(\sigma_t; z_t) \sum_{y_t, \sigma_{t-1}} \pi(\sigma_t|\sigma_{t-1}, y_t) v(y_t, z_t|x_t, x'_t) \mathcal{N}_{t-1}(\sigma_{t-1}), \tag{44}$$

where x_t and x'_t are again quenched environmental variables defined through $p(x_t|x_{t-1})$ and $q_{env}(x'_t|x_t)$. We recover our previous model when $v(y_t, z_t|x_t, x'_t) = q_{in}(y_t|x'_t) \delta(z_t|x_t)$, i.e., $z_t = x_t$. More generally, if y_t and z_t are conditionally independent, i.e., $v(y_t, z_t|x_t, x'_t) = q_{in}(y_t|x_t) u(z_t|x_t)$ for some transition matrix $u(z_t|x_t)$, then the model can be reduced to a model without spatial heterogeneity but with an effective multiplication rate $\tilde{f}(\sigma_t, x_t) =$

$\sum_{z_t} f(\sigma_t, z_t)u(z_t|x_t)$. Note that this effective multiplication rate will generally be non-integer, even when $f(\sigma_t, z_t)$ represents an actual number of offsprings. $\bar{f}(\sigma_t; x_t)$ can also be non-diagonal even though $f(\sigma_t; x_t)$ is diagonal, so that in this case uncertainty is not measured by Shannon entropy even under the restrictive assumptions (A1), (A2), (A3). Note also that while the relevant temporal average of the multiplication rates is the geometric mean, the relevant spatial average in presence of spatial heterogeneities is an arithmetic mean. Another type of spatial heterogeneity is when several patches of population are present and each patch experiences independently an environmental sequence \bar{x} described by the same Markov chain $p(x_t|x_{t-1})$. In the limit of infinitely many patches, the growth of the overall population is then not described by the quenched Lyapunov exponent $\Lambda_{\text{quenched}} = \lim_{t \rightarrow \infty} (1/t) \mathbb{E} \ln |\mathcal{N}_t|$ introduced in Sect. 3, but by the annealed Lyapunov exponent $\Lambda_{\text{annealed}} = \lim_{t \rightarrow \infty} (1/t) \ln \mathbb{E} |\mathcal{N}_t|$, which averages over all environmental sequences instead of focusing on typical ones. These two growth rates, sometimes called the “stochastic growth rate” and the “megamatrix growth rate” in the ecological literature [55], satisfy the general relation $\Lambda_{\text{quenched}} \leq \Lambda_{\text{annealed}}$, which reflects the fact that with many independent patches, the overall population benefits from the few patches that experience atypical but particularly favorable environmental sequences.

Finally, we may mention briefly several other generalizations. A relatively straightforward one, which preserves a close connection to communication theory, is to consider continuous environmental and organismal states [56]; an interesting phenomenon of discretization whereby the optimal distribution of phenotype is actually discrete has then been described [57]. The extension to continuous time is also relatively straightforward (see e.g. [33]). Models where both time and space are continuous are also commonly considered in finance, and can be treated with the tools of stochastic calculus [58]. Another kind of generalization is to introduce a longer time scale at which the transmission of the matrix π is itself subject to mutations, thus allowing to address the issue of the evolution of π towards $\hat{\pi}$. Finally, a more challenging extension is to account for interactions between individuals. For instance, it would be interesting to consider situations where the environment of one population is determined by another population, and to include the possibilities of communication between individuals, or of sexual reproduction.

9 Conclusion

Applications of the concepts of information theory to biology have often been criticized on two main grounds [59]: their failure to account for the directionality of information (the statistical problem of causality), and their failure to account for the value of information (the semantic problem of meaning). Following treatments of analogous problems in engineering and finance, we presented and analyzed a model in which these two features could be integrated. The analysis revealed another limitation of the usual concepts of information theory: their failure to account for the different levels at which information may be processed in a population, which led us to new generalizations of the entropy and mutual information.

Acknowledgements We thank O. Feinerman, G. Iyengar and E. Kussell for comments and discussions. O.R. was supported by a Simons Foundation fellowship from the Rockefeller University.

Appendix A: Definition and Properties of the Model

In mathematical terms, our model belongs to the class of Athreya-Karlin models of multi-type branching process in random environments [42, 60]. Without seeking the highest level

of rigor and generality, it can be described as follows. Let the environmental process be a discrete time, stationary and ergodic, stochastic process $\bar{x} = (x_1, \dots, x_t, \dots)$ with a finite set of states. Let \mathcal{S} be a finite set of admissible internal states σ (types) of the individuals in a population. Let $\xi_{\sigma|\sigma';i}^{(t)}$ be the random variable giving the number of offsprings of type σ that a particular individual i of type σ' generates at time t . The reproductive process, identical and independent for each individual i , is described by the joint distribution $\mathbb{P}_t(\xi_{\sigma|\sigma'}^{(t)})$, which is conditional on the environmental state x_t , and therefore dependent on t . To complete the definition, we may consider starting at time $t = 1$ with a single individual in the given type $\sigma_0 \in \mathcal{S}$, but the asymptotic results, conditional on non-extinction, will not depend on this initial composition. The number $Z_t(\sigma)$ of individuals of type σ at time t is a random variable in terms of which the branching process can recursively be defined as

$$Z_t(\sigma) = \sum_{\sigma' \in \mathcal{S}} \sum_{i=1}^{Z_{t-1}(\sigma')} \xi_{\sigma|\sigma';i}^{(t)}, \tag{45}$$

where different values of i correspond to different realizations of the same random variable.

Our most basic model, with no information processing, assumes that the reproductive process has the particular form

$$\mathbb{P}_t \left(\xi_{\sigma|\sigma'}^{(t)} = \xi \right) = R(\xi|\sigma; x_t) \pi(\sigma|\sigma'; x_t), \tag{46}$$

where $R(\xi|\sigma; x_t)$ is generally a transition matrix, with $\xi \in \mathbb{N}$. Equation (3) is obtained by taking the expectation in (45) with respect to the random variables $\bar{\xi} = \{\xi_{\sigma|\sigma';i}^{(t)}\}_{t,\sigma,\sigma',i}$ for a given environmental sequence $\bar{x} = (x_1, \dots, x_t, \dots)$:

$$\mathcal{N}_t(\sigma) = \mathbb{E}_{\bar{\xi}} [Z_t(\sigma)|\bar{x}]. \tag{47}$$

$\mathcal{N}_t(\sigma)$ depends on $R(\xi|\sigma; x_t)$ only through the multiplication rates defined by

$$f(\sigma; x) = \mathbb{E}_{\bar{\xi}} [\xi R(\xi|\sigma; x)|x] = \sum_{\xi=0}^{\infty} \xi R(\xi|\sigma; x). \tag{48}$$

Other properties, not considered here, such as the probability of extinction, may depend on the fluctuations in the number of offsprings.

In the case of a constant environment, only two events can happen with positive probability [42, 60]: either the population goes extinct, i.e., $|Z_t| \equiv \sum_{\sigma} Z_t(\sigma) = 0$ for some t , or it explodes, i.e., $Z_t \rightarrow \infty$ with $t \rightarrow \infty$. There are therefore two essential questions: (1) What is the probability of extinction? (2) What is the growth rate in the case of explosion? The answer to these questions is contained in the matrix \mathcal{A} given by $\mathcal{A}_{\sigma\sigma'} = \mathbb{E}_{\xi_{\sigma|\sigma'}} [\xi_{\sigma|\sigma'}] = f(\sigma)\pi(\sigma|\sigma')$. Assuming that \mathcal{A} is irreducible and aperiodic, it follows from the Perron-Frobenius theorem that \mathcal{A} has an unique largest real eigenvalue $\lambda = \exp(\Lambda)$ with a corresponding eigenvector v having strictly positive components, which can be normalized so that $\sum_{\sigma} v_{\sigma} = 1$. Let $Q(\sigma)$ be the probability of extinction when initiating the population with a single individual of type σ ; then the answer to (1) is:

$$Q(\sigma) < 1, \forall \sigma \quad \text{if and only if} \quad \Lambda > 0. \tag{49}$$

When $\Lambda > 0$, the branching process is said to be supercritical, and we assume that our model is in this regime to prevent almost sure extinction. For such processes, the answer to (2) is

given by Kesten-Stigum theorem [61]:

$$e^{-t\Lambda} Z_t \rightarrow Wv \quad \text{with probability 1,} \tag{50}$$

where W is a scalar random variable with the following property:

$$\mathbb{P}(W > 0) > 0, \quad \text{if and only if} \quad \mathbb{E} \left[\sum_{\sigma, \sigma'} \xi_{\sigma|\sigma'} \max(0, \ln \xi_{\sigma|\sigma'}) \right] < \infty. \tag{51}$$

We assume that our model satisfies this $X \log X$ condition, i.e., $\mathbb{E}_{\xi_{\sigma|\sigma'}} [\xi_{\sigma|\sigma'} \max(0, \ln \xi_{\sigma|\sigma'})] = \sum_{\xi} (\xi \ln \xi) R(\xi|\sigma) \pi(\sigma|\sigma') < \infty$. Thus, under the assumption of non-extinction, the distribution of the population is asymptotically described by the distribution of the first moments $\mathcal{N}_t(\sigma) = \mathbb{E}_{\xi} [Z_t(\sigma)|x]$, i.e., by the dominant eigenvector v of \mathcal{A} [61]:

$$\lim_{t \rightarrow \infty} \frac{Z_t(\sigma)}{\sum_{\sigma'} Z_t(\sigma')} = v_{\sigma} \quad \text{almost surely conditionally on non-extinction of the population.} \tag{52}$$

In the case of varying environments, there is no longer necessarily convergence of the composition of the population as in (50), but the Lyapunov exponent for the product of random matrices $\mathcal{A}^{(x_t)} \dots \mathcal{A}^{(x_1)}$ still corresponds to the typical growth rate of growing populations: $\Lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \ln(\sum_{\sigma} Z_t(\sigma))$ almost surely, as indicated in (9). The stability condition required for this result to hold is that, with probability one [46],

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\ln \min_{\sigma} \sum_{\sigma'} (\mathcal{A}^{(x_t)} \mathcal{A}^{(x_{t-1})} \dots \mathcal{A}^{(x_1)})_{\sigma\sigma'} \right] = \limsup_{t \rightarrow \infty} \frac{1}{t} \ln \|\mathcal{A}^{(x_t)} \mathcal{A}^{(x_{t-1})} \dots \mathcal{A}^{(x_1)}\|, \tag{53}$$

where $\|M\|$ represents a matrix norm, for instance $\|M\| = \sum_{\sigma, \sigma'} |M_{\sigma, \sigma'}|$.

Appendix B: Growth Rate as a Measure of Fitness

The question of defining an optimal strategy for a game or a financial investment whose outcome is uncertain has a long history, dating back from the earliest days of probability theory. We review here some of the solutions that have been proposed in this context before turning to their relevance for biological populations. We start by assuming that the environmental process is i.i.d. and that no information is acquired. In a model with no correlations between successive environmental states, no gain can be expected from knowing the previous state, and we can assume without restriction that the optimal transition matrix $\pi(\sigma_t|\sigma_{t-1})$ is of the form $\pi(\sigma_t|\sigma_{t-1}) = \pi(\sigma_t)$. Under these assumptions, we need only consider the total size of the population, $|\mathcal{N}_t| = \sum_{\sigma} \mathcal{N}_t(\sigma)$ rather than the population vector \mathcal{N}_t . Indeed,

$$|\mathcal{N}_t| = A^{(x_t)} |\mathcal{N}_{t-1}|, \quad \text{with} \quad A^{(x_t)} \equiv A^{(x_t)}(\pi) \equiv \sum_{\sigma'} \mathcal{A}_{\sigma'\sigma}^{(x_t)} = \sum_{\sigma'} f(\sigma'; x_t) \pi(\sigma'). \tag{54}$$

With these notations, a population of initial size $|\mathcal{N}_0|$ acquires after T time steps a size $|\mathcal{N}_T|$ determined by a product of T scalar random variables:

$$|\mathcal{N}_T| = A^{(x_T)} A^{(x_{T-1})} \dots A^{(x_1)} |\mathcal{N}_0|. \tag{55}$$

Arithmetic Mean The difficulty of defining an optimal strategy π basically stems from the fact that $|\mathcal{N}_T|$ is a random variable whose value depends on the particular sequence of environments (x_1, \dots, x_T) : to each such sequence corresponds an optimal strategy π , but in general no strategy is optimal for every environmental sequence. A naïve solution would be to maximize the expected return. When the successive environmental states are independent, this corresponds to

$$\mathbb{E}[A^{(x_T)} A^{(x_{T-1})} \dots A^{(x_1)}] = (\mathbb{E}[A^{(x_r)}])^T, \tag{56}$$

with $\mathbb{E}[A^{(x_r)}] = \mathbb{E}[A^{(x_r)}(\pi)] = \sum_x p(x) \sum_{\sigma} f(\sigma; x_r) \pi(\sigma).$

\mathbb{E} denotes here the expectation with respect to the fluctuations of the environment. This leads to selecting the portfolio maximizing the so-called arithmetic mean of the return,

$$\max_{\pi} \mathbb{E}[A^{(x_r)}(\pi)] \quad (\text{max. arithmetic mean}). \tag{57}$$

This strategy may however be very risky, as illustrated by the following example: consider a horse race involving only two horses a and b having equal probability of winning, with returns given by $f(a; a) = 3, f(b; a) = 0$ when horse a wins, and $f(a; b) = 0, f(b; b) = 1/2$ when horse b wins. The expected return, $f(a; a)\pi(a) + f(b; b)\pi(b)$, where $\pi(a) + \pi(b) = 1$, is clearly optimized by betting everything on horse a , i.e., $\pi(a) = 1$ and $\pi(b) = 0$. But following this strategy in a sequence of races where the gains are systematically reinvested almost surely leads to bankruptcy. Indeed, if horse b ever wins, everything is lost, and this happens with probability $1 - (1/2)^T$, which tends to 1 as T increases. The maximum expected return is indeed optimal only when averaging over all possible sequences of outcomes, in which case the gain resulting from the only environmental sequence where a never fail to win more than compensate for the loss experienced with all the other sequences of outcomes. When dealing with a single sequence of outcomes, such an average over different environmental sequences is however not relevant.

Expected Utility An argument often given in the economic literature, which dates back from D. Bernoulli’s analysis of the famous St Petersburg paradox [62], is that the criterion based on the arithmetic mean fails to recognize that small losses may represent more “utility” for the gambler than large gains. According to this view, utility of losses and gains depends on the gambler, and may for instance vary with the initial wealth $|\mathcal{N}_0|$. At any given time, each investor should be considered as having his own utility function u that he seeks to optimize,

$$\max_{\pi} \mathbb{E}[u(A^{(x_r)}(\pi))] \quad (\text{max. expected utility } u). \tag{58}$$

The choice of $u(x)$ is critical, since it quantifies the notion of risk. Based on the postulate that an increase in wealth should result in an increase in utility inversely proportionate to the quantity of goods already possessed, Bernoulli proposed $u(x) = \ln x$ as a sensible form of the utility function. In finance, where the problem arises when selecting a diversified portfolio of assets, the risk is often measured by the expected return variance. The return of a given asset σ at time t corresponds in our model to the multiplication rate $f(\sigma; x_t)$, and the expected return of a portfolio can be written vectorially as $\pi^T R$ where π is the vector of portfolio weights $\pi(\sigma)$, and R is the vector of expected returns $R(\sigma) = \mathbb{E}[f(\sigma; x)]$. Following a proposition made by Markowitz [63], the risk is usually measured by $\pi^T \Sigma \pi$

where $\Sigma(\sigma, \sigma') = \mathbb{E}[f(\sigma; x)f(\sigma'; x)] - R(\sigma)R(\sigma')$ represents the covariance matrix of the returns. The portfolio vectors π maximizing $\pi^\top R - \lambda \pi^\top \Sigma \pi$ then defines a family of efficient portfolios parametrized by λ , a parameter fixing the degree of risk that the investor is ready to undertake (λ can also be interpreted as a Lagrange multiplier for the maximization of $\pi^\top R$ at fixed level of risk $\pi^\top \Sigma \pi$, or alternatively, for the minimization of the risk for a fixed expected return). Except for the fact that the covariance matrix Σ is used rather than the correlation matrix $C(\sigma, \sigma') = \mathbb{E}[f(\sigma; x)f(\sigma'; x)]$, Markowitz criterion is essentially similar to maximizing a quadratic utility function $u(x) = x - \lambda x^2$. This function may be viewed as the second-order approximation of a more general utility function, where the approximation is justified by the difficulty of estimating higher-order moments of the returns. Despite their widespread use, criteria based on utility theory and its variants however present a fundamental problem: they are based on *ad-hoc* definitions of risk.

Geometric Mean An independent line of inquiry, initiated by Kelly [35, 64], has promoted the optimization of the geometric mean as an objective criterion. It is based on the observation that if $|\mathcal{N}_T|$ is indeed a random variable whose value depends on the particular sequence of outcomes, for large T most sequences lead to a common, typical, value of the compound return. This can be seen as resulting from the strong law of large numbers applied to $\ln(|\mathcal{N}_T|/|\mathcal{N}_0|) = \sum_t \ln A^{(x_t)}$, which, as a sum of i.i.d. random variables, satisfies

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln |\mathcal{N}_t| = \mathbb{E}[\ln A^{(x_t)}(\pi)] \quad \text{with probability 1.} \tag{59}$$

This result motivates the maximum geometric mean return strategy,

$$\max_{\pi} \mathbb{E}[\ln A^{(x_t)}(\pi)] \quad (\text{max. geometric mean}). \tag{60}$$

This criterion is formally equivalent to optimizing a logarithmic utility function, $u(x) = \ln x$, as originally proposed by Bernoulli in the framework of utility theory [62]. From this standpoint, it may appear as an arbitrary criterion [65], but the argument given here does not rely on the notion of utility function: it relies instead on a fundamental mathematical result, the strong law of large numbers.

Strategies $\hat{\pi}$ corresponding to maximizing the geometric mean as in (60) have, besides (59), a number of other attractive properties [14, 64] (a hat over a quantity, such as $\hat{\pi}$, will always indicate an optimized value of the quantity). From a biological point of view, a particularly important property of $\hat{\pi}$ is asymptotical optimality in an even stronger sense than indicated by (59): $\hat{\pi}$ outperforms any other strategy π (which may vary in time) for almost every sequence of outcomes [14], i.e.,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \ln \frac{|\mathcal{N}_t(\pi)|}{|\mathcal{N}_t(\hat{\pi})|} \leq 0 \quad \text{with probability 1.} \tag{61}$$

From a biological standpoint, the strategy $\hat{\pi}$ is an evolutionary stable strategy [66]: a population characterized by $\hat{\pi}$ cannot be outnumbered by a population with a different π . In other words, if one were to start with a variety of species characterized by different π , one would almost surely end up with a population dominated by the species with largest geometric mean $\mathbb{E}[\ln A^{(x_t)}(\pi)]$. This justifies using the growth rate, given by the geometric mean $\mathbb{E}[\ln A^{(x_t)}(\pi)]$, as an unambiguous measure of adaptation, or “fitness”. Fitness is often informally defined as the expected number of descendants of an individual in a given environment [67], which, in our model, would correspond to $\mathbb{E}[f(\sigma; x)]$ if considering the

descendants after one generation. In general, however, the definition of a fitness function must be supplemented with the references to an “horizon” T and to a particular sequence of future environmental states [68]. In our model, the growth rate emerges as a unique measure of fitness when considering the long-term limit $T \rightarrow \infty$, but, if considering a finite “horizon”, there may be a different strategy π that outperforms $\hat{\pi}$; for instance, at the scale of a single time step, a better strategy may be to optimize the expected multiplication rate, which essentially amounts to an optimization of the arithmetic mean. Note also that our measure of fitness for long-term adaptation is not attached to a particular individual but rather to a trait propagated in a population, the trait defined by the strategy π . An implication of the fact that the fitness function is defined at the population level is that we should not seek to interpret the behavior of the members of the population in terms of the maximization of an individual utility function [69].

Appendix C: Rudiments of Information Theory

Origins of the Entropy The entropy $H_p^{(n)}$ introduced in Sect. 4 appears in source coding theory as the optimal rate of lossless compression for the memoryless source p [12]. To understand why the same quantity occurs in coding theory and models of population dynamics, consider a sequence of T environmental states: if there are n possible states, the number of such sequences is $n^T = e^{T \ln n}$, and $\ln n$, the rate at which the number of possible sequences increases with T , provides a first plausible measure of uncertainty. This measure, originally proposed by Hartley [70], does not account for the fact that some states may be less probable than others, thus effectively reducing the uncertainty. If $p(1), \dots, p(n)$ represent the probabilities of the n different states, the law of large numbers indeed indicates that, almost surely, long environmental sequences are in state x a fraction $p(x)$ of the time. The entropy $H_p^{(n)} = -\sum_x p(x) \ln p(x)$ corresponds to the rate of increase of the number of these typical sequences. The number of typical sequences is indeed $\Upsilon_p^{(T)} \equiv T! / [(p(1)T)! \dots (p(n)T)!]$ which, for $T \rightarrow \infty$, satisfies $(\ln \Upsilon_p^{(T)})/T \rightarrow -\sum_x p(x) \ln p(x) = H_p^{(n)}$. Since the typical sequences are all equiprobable, the entropy also characterizes the probability $e^{-T H_p^{(n)}}$ of observing a particular typical sequence; this property of asymptotic equipartition, which generalizes beyond i.i.d. processes, is central to information theory [14]. The entropy satisfies $0 \leq H_p^{(n)} \leq \ln n$, with $H_p^{(n)} = \ln n$ if and only if no reduction of uncertainty can be gained from the fact that some states are less probable than others, which is the case only when all states are equiprobable, i.e., $p(x) = 1/n$ for all $x = 1, \dots, n$. In the other extreme case where only one state can occur, say $p(1) = 1$, the entropy takes its minimal value $H_p^{(n)} = 0$, corresponding to an absence of uncertainty.¹⁴

Origins of the Mutual Information The mutual information $I(X; Y)$ appears in channel coding theory, the theory of reliable transmission through noisy channels [12], where the capacity of the noisy channel q is given by $C^{(q)} = \max_p I_p^{(q)}$, and in rate-distortion theory, the theory of lossy data compression [13], where the optimal compression rate to describe a source p within a mean distortion D is given by $R_p(D) = \min_q \{I_p^{(q)} : \mathbb{E}[d(x, y)] \leq D\}$, where $\mathbb{E}[d(x, y)] = \sum_{x,y} q(y|x) p(x) d(x, y)$ is the mean distortion for a given distance function $d(x, y)$ between the symbol x from the original data and the symbol y from the

¹⁴ $\ln p(x)$ is not defined when $p(x) = 0$, but $p(x) \ln p(x) \rightarrow 0$ when $p(x) \rightarrow 0$, and it is therefore understood in the definition of $H_p^{(n)}$ that $p(x) \ln p(x) = 0$ whenever $p(x) = 0$.

compressed data. When $X = Y$, the mutual information $I(X; X)$ is nothing but the entropy $H(X)$. The mutual information between two random variables X and Y can also be expressed as $I(X; Y) = H(X) + H(Y) - H(X; Y)$, or as the relative entropy between the joint distribution of (X, Y) and the product of their marginal distributions, i.e., $I(X; Y) = D(\mathbb{P}_{X,Y} \parallel \mathbb{P}_X \mathbb{P}_Y)$. It shows that $I(X; Y)$ is a symmetric function of its variables, and is always non-negative, with $I(X; Y) = 0$ if and only if X and Y are independent, i.e., $\mathbb{P}_{X,Y}(x, y) = \mathbb{P}_X(x)\mathbb{P}_Y(y)$. The mutual information is thus a measure of statistical dependence between random variables.

Appendix D: Analytically Solvable Models

Here, we present the analysis of the model of Fig. 2 in a few simple cases where a solution can be obtained analytically. Beyond horse race models, which are defined by the assumptions (A1), (A2), (A3) introduced in Sect. 2, a general class of solvable model is when (A3) is relaxed to allow for non-zero multiplication rate of the form $f(\sigma; x) = f(x)$ if $\sigma = x$, and $f(\sigma; x) = \tilde{f}(x) < f(x)$ otherwise (horse race models correspond to the case where $\tilde{f}(x) = 0$). Under the assumptions (A1) and (A2) that the environment is i.i.d. and that $q_{in} = \delta$, the mathematical simplicity of these models stems from the fact that $\pi(x|x')$ contributes only to one term indexed by x in the following sum:

$$\Lambda_{f;p}^{(q_{env}, q_{in})}(\pi) = \sum_{x,x'} q_{env}(x'|x)p(x) \ln \left(\tilde{f}(x) + (f(x) - \tilde{f}(x))\pi(x|x') \right). \tag{62}$$

The case where q_{in} is a binary erasure channel as defined in Fig. 4 has also the same property. A subclass of this class of model is when the organisms and the environment have only two states, as in Figs. 5, 6, and 7. We present the details of the analysis of this two-state model below, always assuming that the environment is i.i.d.

D.1 Two-state Generic Model with no Information

In absence of information,

$$\Lambda_{f;p}^{(\eta, \eta)}(\pi) = p(1) \ln(f(1; 1)\pi(1) + f(2; 1)\pi(2)) + p(2) \ln(f(1; 2)\pi(1) + f(2; 2)\pi(2)). \tag{63}$$

It is convenient to introduce the variables

$$\gamma_1 = \frac{f(1; 1) - f(2; 1)}{f(2; 1)}, \quad \gamma_2 = \frac{f(2; 2) - f(1; 2)}{f(1; 2)} \tag{64}$$

which, without loss in generality, can be assumed to be positive. Using $\pi(1) + \pi(2) = 1$, the expression for the growth rate then becomes

$$\Lambda_{f;p}^{(\eta, \eta)}(\pi) = \hat{\Lambda}_p^{(\delta, \delta)} + p(1) \ln \frac{1 + \gamma_1 \pi(1)}{1 + \gamma_1} + p(2) \ln \frac{1 + \gamma_2 \pi(2)}{1 + \gamma_2}. \tag{65}$$

Since $\pi(1) + \pi(2) = 1$, the optimization involves only one independent variable, say $\pi(1)$, subject to the constraints $0 \leq \pi(1) \leq 1$. As a function of $p(1)$, we thus obtain the following

solution:

$$\hat{\Lambda}_{f;p}^{(\eta,\eta)} = \begin{cases} \hat{\Lambda}_{f;p}^{(\delta,\delta)} - p(1) \ln(1 + \gamma_1) & \text{if } p_1 \leq p_c^{(1)}, \\ \hat{\Lambda}_{f;p}^{(\delta,\delta)} - H_p^{(\eta,\eta)} + p(1) \ln\left(1 + \frac{\gamma_1}{(1+\gamma_1)\gamma_2}\right) \\ \quad + p(2) \ln\left(1 + \frac{\gamma_2}{(1+\gamma_2)\gamma_1}\right) & \text{if } p_c^{(1)} \leq p(1) \leq p_c^{(2)}, \\ \hat{\Lambda}_{f;p}^{(\delta,\delta)} - p(2) \ln(1 + \gamma_2) & \text{if } p(1) \geq p_c^{(2)}, \end{cases} \tag{66}$$

$$\text{with } p_c^{(1)} = \frac{\gamma_2}{\gamma_1 + \gamma_2 + \gamma_1\gamma_2}, \quad p_c^{(2)} = 1 - \frac{\gamma_1}{\gamma_1 + \gamma_2 + \gamma_1\gamma_2}. \tag{67}$$

The first and third cases correspond respectively to $\hat{\pi}(1) = 0$ and $\hat{\pi}(1) = 1$, when not switching is optimal, and the intermediate case to $\hat{\pi}(1) = [p(1)(\gamma_1 + \gamma_2 + \gamma_1\gamma_2) - \gamma_2]/(\gamma_1\gamma_2)$. The location of the transitions between the different cases is represented in Fig. 7 for $\gamma_1 = \gamma_2$, and the optimal growth rate $\hat{\Lambda}_p^{(\eta,\eta)}$ as a function of $p(1)$ in Fig. 8 for a particular choice of the parameters.

D.2 Two-state Diagonal Model with a Binary Erasure Channel

We assume here that assumption (A3) holds, i.e., f is diagonal, but consider that assumption (A2) does not hold, and q_{in} is the binary erasure channel q_e defined in Fig. 4. The optimization problem to be solved is $H_p^{(\delta,q_e)} = -\max_{\pi} \Upsilon_p^{(q_e)}(\pi)$ with

$$\Upsilon_p^{(q_e)}(\pi) = p(1) \ln(\pi(1|1)(1 - \varepsilon) + \pi(1|*)\varepsilon) + p(2) \ln(\pi(2|2)(1 - \varepsilon) + \pi(2|*)\varepsilon). \tag{68}$$

Clearly $\hat{\pi}(1|1) = \hat{\pi}(2|2) = 1$, so that we have a single independent variable over which to optimize, say $\pi \equiv \pi(1|*) = 1 - \pi(2|*)$. If we introduce $\gamma = \varepsilon/(1 - \varepsilon)$ then

$$\Upsilon_p^{(q_e)}(\pi) = \ln(1 - \varepsilon) + p(1) \ln(1 + \gamma\pi) + p(2) \ln(1 + \gamma(1 - \pi)). \tag{69}$$

This is formally equivalent to the optimization performed above in absence of information (this formal equivalence extends beyond two-state models). As a function of the level of noise ε , the solution for the binary erasure channel is

$$H_p^{(\delta,q_e)} = \begin{cases} -\min(p(1), p(2)) \ln(1 - \varepsilon) & \text{if } \varepsilon \leq \varepsilon_c(p) \\ H_p^{(\eta,\eta)} - \ln(2 - \varepsilon) & \text{if } \varepsilon \geq \varepsilon_c(p) \end{cases}$$

$$\text{with } \varepsilon_c(p) = \max\left(\frac{1 - 2p(1)}{1 - p(1)}, \frac{1 - 2p(2)}{1 - p(2)}\right). \tag{70}$$

If, for instance, we assume that $p(1) \leq p(2)$, then $\varepsilon_c(p) = (1 - 2p(1))/(1 - p(1))$ and $\hat{\pi} = 0$, $H_p^{(\delta,q_e)} = -p(1) \ln(1 - \varepsilon)$ for $\varepsilon < \varepsilon_c(p)$. An illustration is given in Fig. 5 where we compare for $p(1) = 0.1$ the individual information $I_p^{(\delta,q_e)} = H_p^{(\eta,\eta)} - H_p^{(\delta,q_e)}$ with the mutual information $I_p^{(q_e,\delta)} = H_p^{(\eta,\eta)} - H_p^{(q_e,\delta)} = (1 - \varepsilon)H_p^{(\eta,\eta)}$ for the same binary erasure channel.

D.3 Two-state Model Diagonal with a Binary Symmetric Channel

We assume here that assumption (A3) holds, i.e., f is diagonal, but consider that assumption (A2) does not hold, and q_{in} is the binary symmetric channel q_s defined in Fig. 4. The binary

symmetric channel corresponds to $q_s(1|1) = q_s(2|2) = 1 - \varepsilon$ and $q_s(2|1) = q_s(1|2) = \varepsilon$, where, without loss in generality, we can assume that $0 \leq \varepsilon \leq 1/2$. The optimization problem to be solved is $H_p^{(\delta, q_s)} = -\max_{\pi} \Upsilon_p^{(q_s)}(\pi)$ with

$$\Upsilon_p^{(q_s)}(\pi) \equiv p(1) \ln(\pi(1|1)(1 - \varepsilon) + \pi(1|2)\varepsilon) + p(2) \ln(\pi(2|2)(1 - \varepsilon) + \pi(2|1)\varepsilon). \tag{71}$$

We have here two independent parameters over which to optimize, $\pi_1 = \pi(1|1)$ and $\pi_2 = \pi(2|2)$, since $\pi(2|1) = 1 - \pi(1|1) = 1 - \pi_1$ and $\pi(1|2) = 1 - \pi(2|2) = 1 - \pi_2$. If we introduce $\gamma = \varepsilon/(1 - \varepsilon)$, the function to optimize becomes

$$\Upsilon_p^{(q_s)}(\pi_1, \pi_2) \equiv \ln(1 - \varepsilon) + p(1) \ln(\pi_1 + (1 - \pi_2)\gamma) + p(2) \ln(\pi_2 + (1 - \pi_1)\gamma). \tag{72}$$

The calculation shows that the only case where we can have both $0 < \hat{\pi}_1 < 1$ and $0 < \hat{\pi}_2 < 1$ is the blind case where $\gamma = 1$ and $\varepsilon = 1/2$, for which we have the proportional betting solution $\hat{\pi}_1 = p(1)$ and $\hat{\pi}_2 = p(2)$. In any other case, $\hat{\pi}_1 = 1$ or $\hat{\pi}_2 = 1$, which reduces the problem to an optimization over a single variable. The solution is

$$H_p^{(\delta, q_s)} = \begin{cases} -\ln(1 - \varepsilon) & \text{if } \varepsilon \leq \varepsilon_c(p) \\ H_p^{(0)} - \varepsilon_c(p) \ln((1 - \varepsilon)/\varepsilon) & \text{if } \varepsilon \geq \varepsilon_c(p) \end{cases}, \quad \text{with } \varepsilon_c(p) = \min(p(1), p(2)). \tag{73}$$

If for instance $p(1) \leq p(2)$, when $\varepsilon_c(p) < p(1)$ we have both $\hat{\pi}_1 = 1$ and $\hat{\pi}_2 = 1$ while for $\varepsilon_c(p) > p(1)$, we have $\hat{\pi}_2 = 1$ but $\hat{\pi}_1 < 1$. An illustration is given in Fig. 6 where we compare for $p(1) = 0.1$ the individual information $I_p^{(\delta, q_s)}$ with the mutual information $I_p^{(q_s, \delta)}$ for the same binary symmetric channel.

Appendix E: A Solvable Model in Non i.i.d. Environments

We connect here the model proposed in [33] to the framework of this paper and discuss how uncertainties in timing can thus be quantified. We will thus make explicit the time scales involved in the trade-off between short-term adjustment to the current environmental conditions, and longer term anticipation of changes of these conditions.

Two time scales govern short-term adjustment: the time $\tau(\epsilon)$ that an environmental state ϵ lasts, and the time $\alpha(\epsilon; \epsilon')$ that it takes for the type with largest multiplication rate in ϵ to dominate the population; this later adjustment time depends on the composition of the population at the time of the environmental change, and the notation $\alpha(\epsilon; \epsilon')$ indicates that we consider a population initially adjusted to some other environmental state ϵ' . In the so-called adiabatic regime where the population has always time to adjust to the current environment, i.e., $\alpha(\epsilon; \epsilon') \ll \tau(\epsilon)$ for all ϵ and $\epsilon' \neq \epsilon$, the dynamics of the population has a common feature with horse race models: at the end of an environmental period, most of the population shares a common type, much as in horse race models where, at the end of a time step, only the money invested in the winning horse yields a non-zero payoff. In horse race models, the particular form of the multiplication rates, where $f(\sigma; \epsilon) = 0$ whenever $\sigma \neq \epsilon$, implies that the uncertainty cost $\hat{\Lambda}_{\omega; f} - \hat{\Lambda}_{\omega; f}^{(C)}$ can be measured by the entropic function $\hat{H}_{\omega}^{(C)}$ defined in (43), where the transition matrix for the environment is here denoted $\omega(\epsilon_t | \epsilon_{t-1})$. We shall see that the same function contributes to the uncertainty cost of systems in the

adiabatic regime; in particular, the entropy of the environmental process can account for part of the uncertainty cost, as first noticed in [33].

The characteristic time that an environmental state ϵ lasts can be defined as the mean time $\tau(\epsilon)$ spent in ϵ

$$\tau(\epsilon) = \frac{1}{1 - \omega(\epsilon|\epsilon)}. \tag{74}$$

To simply define an adjustment time $\alpha(\epsilon; \epsilon')$, we assume that each environmental state ϵ is associated with a different optimal type denoted with the same symbol $\sigma = \epsilon$, i.e., $f(\epsilon; \epsilon) > f(\sigma; \epsilon)$ for all $\sigma \neq \epsilon$. We also assume that, in the course of a single time step, an individual is more likely to stay in its current type than to adopt a new one, i.e., $\pi(\sigma'|\sigma; \epsilon) \ll 1$ for $\sigma' \neq \sigma$; these two assumptions ensure that, in a constant environment ϵ , an optimal type $\sigma = \epsilon$ can indeed dominate the population if given sufficient time. Under these assumptions, a population initially composed of \mathcal{N}_0 individuals adjusted to environment ϵ' , has, after a time t spent in environment ϵ , a number $\mathcal{N}_t(\epsilon)$ of individuals of type ϵ which is given by

$$\mathcal{N}_t(\epsilon) \simeq f(\epsilon; \epsilon)^t Q(\epsilon; \epsilon') \mathcal{N}_0. \tag{75}$$

Here, $Q(\epsilon; \epsilon')$, which satisfies $0 < Q(\epsilon; \epsilon') < 1$, can be interpreted as a (non-symmetric) overlap between the compositions of the population before and after the environmental change from ϵ' to ϵ ; as shown in Appendix H, it is given for $\epsilon \neq \epsilon'$ by

$$Q(\epsilon; \epsilon') = \Delta(\epsilon; \epsilon') \pi(\epsilon|\epsilon'; \epsilon') + \Delta(\epsilon'; \epsilon) \pi(\epsilon|\epsilon'; \epsilon) \quad \text{with} \quad \Delta(\epsilon; \epsilon') = \frac{f(\epsilon; \epsilon')}{f(\epsilon'; \epsilon') - f(\epsilon; \epsilon')}. \tag{76}$$

The adjustment time $\alpha(\epsilon; \epsilon')$ can then be defined as the time at which the sub-population of type ϵ starts to overtake the sub-population of type ϵ' , $\mathcal{N}_t(\epsilon) \sim \mathcal{N}_t(\epsilon')$; given (75) and $\mathcal{N}_t(\epsilon') \sim f(\epsilon'; \epsilon)^t \mathcal{N}_0$, this leads to

$$\alpha(\epsilon; \epsilon') = \frac{1}{\ln f(\epsilon; \epsilon) - \ln f(\epsilon'; \epsilon')} \ln \left(\frac{1}{Q(\epsilon; \epsilon')} \right). \tag{77}$$

In the ‘‘adiabatic regime’’ where environmental periods exceed the adjustment times, i.e., $\alpha(\epsilon; \epsilon') \ll \tau(\epsilon)$ for all ϵ and $\epsilon' \neq \epsilon$, we obtain from (75) a simple expression for the Lyapunov exponent (see Appendix H),

$$\Lambda_{\omega; f}^{(\text{adiabatic})}(\pi) \simeq \sum_{\epsilon} \omega(\epsilon) \ln f(\epsilon; \epsilon) - \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon') \omega(\epsilon') \ln \left(\frac{1}{Q(\epsilon; \epsilon')} \right), \tag{78}$$

with the convention that, when $\epsilon' \neq \epsilon$, $Q(\epsilon; \epsilon) = \pi(\epsilon|\epsilon; \epsilon)$. The first term on the right hand side corresponds to the optimal Lyapunov exponent, $\hat{\Lambda}_{\omega; f}$, and the second, when optimized over π , to the uncertainty cost $\hat{\Lambda}_{\omega; f} - \hat{\Lambda}_{\omega; f}^{(C)}$. This second term depends, via $Q(\epsilon; \epsilon')$, on both the transition matrix π , and the values of the multiplication rates f . These two contributions are, however, set apart when the transition matrix π can be factorized in (76), which occurs in two notable cases. One case is in absence of a sensor, when $\pi(\epsilon|\epsilon'; \epsilon') = \pi(\epsilon|\epsilon'; \epsilon) = \pi(\epsilon|\epsilon')$, and, therefore,

$$Q(\epsilon; \epsilon') = \Gamma(\epsilon; \epsilon') \pi(\epsilon|\epsilon') \quad \text{with} \quad \Gamma(\epsilon; \epsilon') = \Delta(\epsilon; \epsilon') + \Delta(\epsilon'; \epsilon). \tag{79}$$

Another case is in presence of a reliable sensor, when $\pi(\epsilon|\epsilon'; \epsilon') \ll \pi(\epsilon|\epsilon'; \epsilon)$, and, therefore,

$$Q(\epsilon; \epsilon') \simeq \Gamma(\epsilon; \epsilon')\pi(\epsilon|\epsilon'; \epsilon) \quad \text{with } \Gamma(\epsilon; \epsilon') = \Delta(\epsilon'; \epsilon). \tag{80}$$

In both of these cases, the second term of the right-hand side of (78) can be decomposed as

$$\begin{aligned} \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \left(\frac{1}{Q(\epsilon; \epsilon')} \right) &= \sum_{\epsilon, \epsilon' \neq \epsilon} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \Gamma(\epsilon; \epsilon') \\ &\quad - \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \pi(\epsilon|\epsilon'; \epsilon). \end{aligned} \tag{81}$$

The term $-\sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \pi(\epsilon|\epsilon'; \epsilon)$ is analogous to the term obtained for a horse-race model in a Markov environment. The other term, involving $\Gamma(\epsilon; \epsilon')$, may be interpreted as the cost of the delay for transferring the majority of the population from one type to the next (such a term is absent in horse race models where transfers of capital can occur instantaneously prior to the environmental change).

Since only the term $-\sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \pi(\epsilon|\epsilon'; \epsilon)$ depends on the transition matrix π , the optimal strategy is the one that minimizes it, and it has exactly the same features as in horse race models. If we consider for instance the situation with no information, we find $\hat{\pi}^{(n)}(\epsilon|\epsilon') = \omega(\epsilon|\epsilon')$, the proportional betting strategy, and the optimal value of the last term in (81) is the entropy rate

$$H_{\omega}^{(n)} = - \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon')\omega(\epsilon') \ln \omega(\epsilon|\epsilon'). \tag{82}$$

It is instructive to make here explicit the characteristic times $\tau(\epsilon)$ giving the mean duration in each environmental state ϵ . This is done by introducing $\tilde{\omega}(\epsilon|\epsilon') = \omega(\epsilon|\epsilon')/\tau(\epsilon')$, the probability that the environment changes from state ϵ' to state $\epsilon \neq \epsilon'$, given that it does change its state, and $\tilde{\omega}(\epsilon) = \omega(\epsilon)\tau(\epsilon)/\tau$, the probability to end up in state ϵ when such an environmental change occurs, with $\tau = \sum_{\epsilon} \tilde{\omega}(\epsilon)\tau(\epsilon)$ representing the mean duration of a period of constant environment. With these definitions, it can indeed be shown that

$$H_{\omega}^{(n)} = \frac{1}{\tau} H_{\tilde{\omega}}^{(n)} + \sum_{\epsilon} \frac{\tau(\epsilon)}{\tau} \tilde{\omega}(\epsilon) H_{b(1/\tau(\epsilon))}^{(n)} \tag{83}$$

where $b(1/\tau(\epsilon))$ refers to the Bernoulli distribution with parameter $1/\tau(\epsilon)$ whose entropy is

$$H_{b(1/\tau(\epsilon))}^{(n)} = - \frac{1}{\tau(\epsilon)} \ln \frac{1}{\tau(\epsilon)} - \left(1 - \frac{1}{\tau(\epsilon)} \right) \ln \left(1 - \frac{1}{\tau(\epsilon)} \right). \tag{84}$$

Equation (83) shows that the uncertainty has two components, each of which measurable by an entropy: an uncertainty about the nature of the next environment, captured by $H_{\tilde{\omega}}^{(n)}$, and an uncertainty about the timing of environmental changes, captured by $H_{b(1/\tau(\epsilon))}^{(n)}$. As in horse race models, the maladjustment cost has, in the ‘‘adiabatic’’ limit and in absence of information, the form of a relative entropy which can also be decomposed in two terms; when the corresponding expressions are expanded for large $\tau(\epsilon)$, the formulas presented in [33] are thus recovered.

Appendix F: Proof of the Entropic Bound

We prove here the bound on the uncertainty cost for the model of Fig. 2,

$$\hat{\Lambda}_{p;f}^{(\delta,\delta)} - \hat{\Lambda}_{\omega;f}^{(q_{\text{env}},q_{\text{in}})} \leq H_p^{(q_{\text{env}},q_{\text{in}})}, \tag{85}$$

for i.i.d. environments but arbitrary multiplication rates f .

By defining $\epsilon = (x, x')$, $\omega(\epsilon) = q_{\text{env}}(x'|x)p(x)$ and $q(y|\epsilon) = q_{\text{in}}(y|x')$, the Lyapunov exponent $\Lambda_{p;f}^{(q_{\text{env}},q_{\text{in}})}(\pi)$ is more concisely, but equivalently, written

$$\Lambda_{\omega;f}^{(q)}(\pi) = \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_{\sigma,y} f(\sigma; \epsilon) \pi(\sigma|y) q(y|\epsilon) \right). \tag{86}$$

For each environmental state ϵ , let $\phi(\epsilon)$ be one of the types σ with maximal multiplication rate, such that $f(\sigma; \epsilon) \leq f(\phi(\epsilon); \epsilon)$ for all σ . An assumption is here that $f(\phi(\epsilon); \epsilon) > 0$ for all ϵ , which is necessary for the population not to go extinct, assuming that $\omega(\epsilon) > 0$ for all ϵ .

From the definition of $\hat{\Lambda}_{\omega;f}^{(q)}$ as $\max_{\pi} \Lambda_{\omega;f}^{(q)}(\pi)$, for any probability matrix $\pi(\sigma|y)$, we have

$$\hat{\Lambda}_{\omega;f}^{(\delta)} - \hat{\Lambda}_{\omega;f}^{(q)} \leq \sum_{\epsilon} \omega(\epsilon) \ln f(\phi(\epsilon); \epsilon) - \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_{\sigma,y} f(\sigma; \epsilon) \pi(\sigma|y) q(y|\epsilon) \right) \tag{87}$$

$$= - \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_{\sigma,y} \frac{f(\sigma; \epsilon)}{f(\phi(\epsilon); \epsilon)} \pi(\sigma|y) q(y|\epsilon) \right) \tag{88}$$

$$= - \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_y \left(\pi(\phi(\epsilon)|y) + \sum_{\sigma \neq \phi(\epsilon)} \frac{f(\sigma; \epsilon)}{f(\phi(\epsilon); \epsilon)} \pi(\sigma|y) \right) q(y|\epsilon) \right) \tag{89}$$

$$\leq - \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_y \pi(\phi(\epsilon)|y) q(y|\epsilon) \right). \tag{90}$$

Let $\hat{\rho}(\epsilon|y)$ be a transition matrix that achieves the minimum in the definition of $H_{\omega}^{(q)}$, which is

$$H_{\omega}^{(q)} = \min_{\rho} \sum_{\epsilon} \omega(\epsilon) \ln \left(\sum_y \rho(\epsilon|y) q(y|\epsilon) \right)^{-1}. \tag{91}$$

Under the assumption that ϕ is injective, by taking $\pi(\sigma|y) = \hat{\rho}(\epsilon|y)$ if $\sigma = \phi(\epsilon)$, and $\pi(\sigma|y) = 0$ if there is no ϵ for which $\sigma = \phi(\epsilon)$, we define a probability matrix π for which the right-hand side of (90) corresponds exactly to $H_{\omega}^{(q)}$. Hence

$$\hat{\Lambda}_{\omega;f}^{(\delta)} - \hat{\Lambda}_{\omega;f}^{(q)} \leq H_{\omega}^{(q)}. \tag{92}$$

If ϕ is non-injective, an even tighter upper bound can be designed. To this end, we go back to the variables (x, y') and define a coarse-grained environmental process whose states are the equivalent classes for the relation $\phi(x) = \phi(z)$, and whose probability distribution is defined on the quotient set by $\tilde{p}(\tilde{x}) = \sum_{z \in \tilde{x}} p(z)$ for every equivalent class

$\tilde{x} = \{z : \phi(z) = \phi(x)\}$. Introducing also $\tilde{q}_{\text{env}}(x'|\tilde{x}) = \sum_{x \in \tilde{x}} q_{\text{env}}(x'|x)p(x)/\tilde{p}(\tilde{x})$, the expression in (90) becomes

$$\begin{aligned} & \sum_{x',x} q_{\text{env}}(x'|x)p(x) \left(\sum_y \pi(\phi(x)|y)q(y|x') \right) \\ &= \sum_{x',\tilde{x}} \tilde{q}_{\text{env}}(x'|\tilde{x})\tilde{p}(\tilde{x}) \ln \left(\sum_y \pi(\phi(\tilde{x})|y)q(y|x') \right), \end{aligned} \tag{93}$$

where $\tilde{\pi}(\phi(\tilde{x})|y) = \pi(\phi(x)|y)$ is well-defined by definition of the equivalence relation. We are then reduced to the injective case, and can therefore conclude

$$\hat{\Lambda}_{p:f}^{(\delta,\delta)} - \hat{\Lambda}_{p:f}^{(q_{\text{env}},q)} \leq H_{\tilde{p}}^{(q_{\text{env}},q)}. \tag{94}$$

Finally, it follows from the definition of generalized entropy $H_p^{(q_{\text{env}},q)}$ and from the concavity of the logarithm that coarse-graining always reduces the entropy, i.e., $H_{\tilde{p}}^{(q_{\text{env}},q)} \leq H_p^{(q_{\text{env}},q)}$, thus proving the entropic bound in the general case.

Appendix G: Proof of the Maladjustment Bound

We prove here a bound on the cost incurred for following a non-optimal strategy in the model of Fig. 2. This bound generalizes the bound established for models of financial investments where $q_{\text{in}} = \delta$ [14]. We consider here an i.i.d. environment but arbitrary multiplication rates f . If $\hat{\pi}_{p'}$ denotes an optimal strategy for the i.i.d. environment with probability $p'(x)$ rather than $p(x)$, we show that

$$\hat{\Lambda}_{p:f}^{(q_{\text{env}},q_{\text{in}})} - \Lambda_{p:f}^{(q_{\text{env}},q_{\text{in}})}(\hat{\pi}_{p'}) \leq D(p\|p'), \tag{95}$$

where $D(p\|p') = \sum_x p(x) \ln[p(x)/p'(x)]$ is the relative entropy between the environmental distributions p and p' .

Writing the Lyapunov exponent $\Lambda_{p:f}^{(q_{\text{env}},q_{\text{in}})}(\pi)$ as in (86), we show that

$$\hat{\Lambda}_{\omega:f}^{(q)} - \Lambda_{\omega:f}^{(q)}(\hat{\pi}_{\omega'}) \leq D(\omega\|\omega'), \tag{96}$$

from which (95) follows by taking $\epsilon = (x, x')$ and $\omega(\epsilon) = q_{\text{env}}(x'|x)p(x)$ and $\omega'(\epsilon) = q_{\text{env}}(x'|x)p'(x)$.

The proof relies on the characterization of $\hat{\pi}_{\omega'}$ in terms of the so-called Kuhn-Tucker conditions [14], which generalizes to inequality constraints the method of Lagrange multipliers. These conditions imply here the existence of a set of $\lambda_y \geq 0$ satisfying

$$\sum_{\epsilon} \omega'(\epsilon) \left(\frac{f(\sigma; \epsilon)q(y|\epsilon)}{\sum_{y',\sigma'} f(\sigma'; \epsilon)\hat{\pi}_{\omega'}(\sigma'|y')q(y'|\epsilon)} \right) \begin{cases} = \lambda_y & \text{if } \hat{\pi}_{\omega'}(\sigma|y) > 0, \\ \leq \lambda_y & \text{if } \hat{\pi}_{\omega'}(\sigma|y) = 0, \end{cases} \tag{97}$$

and $\sum_y \lambda_y = 1$.

After noticing that by taking the union of the two environmental state spaces if necessary, we can assume that the two processes described by ω and ω' have same states, we generalize

a proof presented in [14] by considering the following series of inequalities:

$$\begin{aligned}
 \hat{\Lambda}_{\omega;f}^{(q)} - \Lambda_{\omega;f}^{(q)}(\hat{\pi}_{\omega'}) &= \sum_{\epsilon} \omega(\epsilon) \ln \left(\frac{\sum_{y,\sigma} f(\sigma; \epsilon) \hat{\pi}_{\omega}(\sigma|y) q(y|\epsilon)}{\sum_{y',\sigma'} f(\sigma'; \epsilon) \hat{\pi}_{\omega'}(\sigma'|y') q(y'|\epsilon)} \right) \\
 &= \sum_{\epsilon} \omega(\epsilon) \ln \left(\frac{\omega'(\epsilon) \sum_{y,\sigma} f(\sigma; \epsilon) \hat{\pi}_{\omega}(\sigma|y) q(y|\epsilon)}{\omega(\epsilon) \sum_{y',\sigma'} f(\sigma'; \epsilon) \hat{\pi}_{\omega'}(\sigma'|y') q(y'|\epsilon)} \right) + D(\omega\|\omega') \\
 &\leq \ln \left(\sum_{\epsilon} \frac{\omega'(\epsilon) \sum_{y,\sigma} f(\sigma; \epsilon) \hat{\pi}_{\omega}(\sigma|y) q(y|\epsilon)}{\sum_{y',\sigma'} f(\sigma'; \epsilon) \hat{\pi}_{\omega'}(\sigma'|y') q(y'|\epsilon)} \right) + D(\omega\|\omega') \\
 &= \ln \left(\sum_{y,\sigma} \hat{\pi}_{\omega}(\sigma|y) \sum_{\epsilon} \omega'(\epsilon) \left(\frac{f(\sigma; \epsilon) q(y|\epsilon)}{\sum_{y',\sigma'} f(\sigma'; \epsilon) \hat{\pi}_{\omega'}(\sigma'|y') q(y'|\epsilon)} \right) \right) \\
 &\quad + D(\omega\|\omega') \\
 &\leq \ln \left(\sum_{y,\sigma} \hat{\pi}_{\omega}(\sigma|y) \lambda_y \right) + D(\omega\|\omega') \\
 &= \ln \left(\sum_y \lambda_y \right) + D(\omega\|\omega') \\
 &= D(\omega\|\omega'), \tag{98}
 \end{aligned}$$

where the first inequality follows from the concavity of the logarithm (Jensen’s inequality), and the second from (97).

Appendix H: Perturbative Approximation

With \mathcal{N}_t representing the population vector whose components $\mathcal{N}_t(\sigma)$ are the mean number of individuals of type σ , and assuming here that no information is acquired, (4) can be written

$$\mathcal{N}_{t+1} = \left(\mathcal{A}_0^{(\epsilon_t)} + \mathcal{A}_1^{(\epsilon_t)} \right) \mathcal{N}_t \tag{99}$$

where, using a bracket notation, the elements of the matrices $\mathcal{A}_0^{(\epsilon)}$ and $\mathcal{A}_1^{(\epsilon)}$, are

$$\langle \sigma' | \mathcal{A}_0^{(\epsilon)} | \sigma \rangle = \begin{cases} f(\sigma; \epsilon) & \text{if } \sigma' = \sigma, \\ 0 & \text{if } \sigma' \neq \sigma, \end{cases} \tag{100}$$

$$\langle \sigma' | \mathcal{A}_1^{(\epsilon)} | \sigma \rangle = \begin{cases} -f(\sigma; \epsilon)(1 - \pi(\sigma|\sigma; \epsilon)) & \text{if } \sigma' = \sigma, \\ f(\sigma'; \epsilon)\pi(\sigma'|\sigma; \epsilon) & \text{if } \sigma' \neq \sigma. \end{cases} \tag{101}$$

The rationale for this decomposition is that \mathcal{A}_1 is a perturbation for \mathcal{A}_0 when the assumption is made that individuals are less likely to switch to a new type than to retain their current type. To simplify the discussion, we also assume that each environmental state ϵ is associated with a unique optimal type $\sigma = \epsilon$, satisfying $f(\epsilon; \epsilon) > f(\sigma; \epsilon)$ for all $\sigma \neq \epsilon$.

Under these assumptions, we can derive approximate expressions for the eigenvalues and eigenvectors of $\mathcal{A}_1 + \mathcal{A}_0$ by a perturbative expansion. If $\lambda_{\sigma}^{(\epsilon)}$, $|\psi_{\sigma}^{(\epsilon)}\rangle$ and $\langle \psi_{\sigma}^{(\epsilon)}|$ denote

respectively the eigenvalue, and the right and left eigenvectors of the matrix $\mathcal{A}_0^{(\epsilon)} + \mathcal{A}_1^{(\epsilon)}$, we have, to first order in the perturbative expansion:

$$\lambda_\sigma^{(\epsilon)} = f(\sigma; \epsilon)\pi(\sigma|\sigma; \epsilon), \tag{102}$$

$$|\psi_\sigma^{(\epsilon)}\rangle = |\sigma\rangle + \sum_{\sigma' \neq \sigma} \frac{f(\sigma'; \epsilon)}{f(\sigma; \epsilon) - f(\sigma'; \epsilon)} \pi(\sigma'|\sigma; \epsilon)|\sigma'\rangle, \tag{103}$$

$$\langle\psi_\sigma^{(\epsilon)}| = \langle\sigma| + \sum_{\sigma' \neq \sigma} \frac{f(\sigma'; \epsilon)}{f(\sigma; \epsilon) - f(\sigma'; \epsilon)} \pi(\sigma|\sigma'; \epsilon)\langle\sigma'|. \tag{104}$$

The domain of validity of this approximation can be estimating by comparing the first and second order contributions to the eigenvalues, and we thus get the condition

$$\left| \sum_{\sigma' \neq \sigma} \frac{f(\sigma; \epsilon)f(\sigma'; \epsilon)}{f(\sigma; \epsilon) - f(\sigma'; \epsilon)} \pi(\sigma'|\sigma; \epsilon)\pi(\sigma|\sigma'; \epsilon) \right| \ll f(\sigma; \epsilon)(1 - \pi(\sigma|\sigma; \epsilon)), \quad \forall \sigma, \epsilon. \tag{105}$$

Given that $1 - \pi(\sigma|\sigma; \epsilon) = \sum_{\sigma' \neq \sigma} \pi(\sigma'|\sigma; \epsilon)$, a sufficient condition for (105) to hold is

$$\pi(\sigma|\sigma'; \epsilon) \ll \left| \frac{f(\sigma; \epsilon) - f(\sigma'; \epsilon)}{f(\sigma'; \epsilon)} \right|, \quad \forall \sigma, \sigma', \epsilon (\sigma \neq \sigma'). \tag{106}$$

This shows that the underlying assumption behind the perturbative expansion is that changes in composition of the population should primarily be due to differences in multiplication rates, rather than be due to switches to new types.

For the dynamics to be in the ‘‘adiabatic regime’’, it is furthermore necessary that the environment stays long enough in any given state ϵ . When this is the case, the population vector \mathcal{N}_t is, at the end of a period spent in state ϵ , quasi aligned along the dominant eigenvector $|\psi_\epsilon^{(\epsilon)}\rangle$ that corresponds to the most favorable type for environment ϵ , $|\psi_\epsilon^{(\epsilon)}\rangle \simeq |\epsilon\rangle$. If ϵ' was the environmental state preceding the current state ϵ , the system is described by $|\psi_{\epsilon'}^{(\epsilon')}\rangle$ at $t = 0$ and, at $t = \tau(\epsilon)$, when the environmental state becomes ϵ , we require that

$$\left| \sum_{\sigma \neq \epsilon} \langle\psi_\sigma^{(\epsilon)}|\psi_{\epsilon'}^{(\epsilon')}\rangle (\lambda_\sigma^{(\epsilon)})^{\tau(\epsilon)} \right| \ll \langle\psi_\epsilon^{(\epsilon)}|\psi_{\epsilon'}^{(\epsilon')}\rangle (\lambda_\epsilon^{(\epsilon)})^{\tau(\epsilon)}. \tag{107}$$

This condition can be made explicit by using the perturbative formulas

$$\langle\psi_\sigma^{(\epsilon)}|\psi_{\epsilon'}^{(\epsilon')}\rangle = \begin{cases} 1 & \text{if } \sigma = \epsilon', \\ \frac{f(\sigma; \epsilon')}{f(\epsilon'; \epsilon') - f(\sigma; \epsilon')} \pi(\sigma|\epsilon'; \epsilon') + \frac{f(\epsilon'; \epsilon)}{f(\sigma; \epsilon) - f(\epsilon'; \epsilon)} \pi(\sigma|\epsilon'; \epsilon) & \text{if } \sigma \neq \epsilon'. \end{cases} \tag{108}$$

Given ϵ , the longest delay time $\alpha(\epsilon; \epsilon')$ is therefore when the preceding environment ϵ' corresponds to the second largest eigenvalue of $\mathcal{A}_1^{(\epsilon)}$, that is, when $\epsilon' = \sigma$ such that $f(\sigma; \epsilon) = \max_{\sigma' \neq \epsilon} f(\sigma'; \epsilon)$. Denoting $Q(\epsilon; \epsilon') = \langle\psi_\epsilon^{(\epsilon)}|\psi_{\epsilon'}^{(\epsilon')}\rangle$ we thus obtain the condition $f(\epsilon'; \epsilon)^{\tau(\epsilon)} \ll f(\epsilon; \epsilon)Q(\epsilon; \epsilon')$ or, equivalently,

$$\tau(\epsilon) \ll \frac{1}{\ln f(\epsilon; \epsilon) - \ln f(\epsilon'; \epsilon)} \ln \frac{1}{Q(\epsilon; \epsilon')}, \tag{109}$$

where the right-hand side can be taken as a definition for the adjustment time $\alpha(\epsilon; \epsilon')$.

Let now denote $\tilde{\omega}(\epsilon|\epsilon')$ the probability for the environment to change from state ϵ' to state $\epsilon \neq \epsilon'$, given that it does change its state. This is given by $\tilde{\omega}(\epsilon|\epsilon') = \omega(\epsilon|\epsilon')/(1 - \omega(\epsilon'|\epsilon'))$, where $\tau(\epsilon') = 1/(1 - \omega(\epsilon'|\epsilon'))$ also corresponds to the mean time spent in state ϵ' . The unconditional probability to end up in state ϵ when an environmental change occurs is $\tilde{\omega}(\epsilon) = \omega(\epsilon)\tau/\tau(\epsilon)$ where $\tau = \sum_{\epsilon} \tilde{\omega}(\epsilon)\tau(\epsilon)$ represents the mean duration of a period of constant environment. In terms of these quantities, the growth rate is

$$\Lambda_{\omega}^{(\text{adiabatic})}(\pi) = \frac{1}{\tau} \sum_{\epsilon, \epsilon' \neq \epsilon} \tilde{\omega}(\epsilon|\epsilon') \tilde{\omega}(\epsilon') \ln \left((\lambda_{\epsilon}^{(\epsilon)})^{\tau(\epsilon)} \langle \psi_{\epsilon}^{(\epsilon)} | \psi_{\epsilon'}^{(\epsilon')} \rangle \right), \quad (110)$$

which is also equivalent to

$$\Lambda_{\omega}^{(\text{adiabatic})}(\pi) = \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon') w(\epsilon') \ln f(\epsilon; \epsilon) - \sum_{\epsilon, \epsilon'} \omega(\epsilon|\epsilon') w(\epsilon') \ln \frac{1}{Q(\epsilon; \epsilon')}, \quad (111)$$

with the convention that $Q(\epsilon; \epsilon) = \pi(\epsilon|\epsilon; \epsilon)$, while $Q(\epsilon; \epsilon')$ for $\epsilon \neq \epsilon'$ is given by (76).

References

1. Maynard-Smith, J.: The concept of information in biology. *Philos. Sci.* **67**(2), 177–194 (2000)
2. Jablonka, E.: Information: Its interpretation, its inheritance, and its sharing. *Philos. Sci.* **69**(4), 578–605 (2002)
3. Szostak, J.W.: Functional information: molecular messages. *Nature* **423**, 689 (2003)
4. Nurse, P.: Life, logic and information. *Nature* **454**, 424–426 (2008)
5. Quastler, H. (ed.): *Essays on the Use of Information Theory in Biology*. University of Illinois, Urbana (1953)
6. Rashevsky, N.: Life, information theory, and topology. *Bull. Math. Biol.* **17**, 229–235 (1955)
7. Atlan, H.: *L'organisation biologique et la théorie de l'information*. Hermann, Paris (1972)
8. Berger, T.: Living information theory. *IEEE Inf. Theory Soc. NewsL.* **53**, 1–19 (2003)
9. Adami, C.: Information theory in molecular biology. *Phys. Life Rev.* **1**, 3–22 (2004)
10. Taylor, S.F., Tishby, N., Bialek, W.: Information and fitness. *arXiv:0712.4382* (2007)
11. Polani, D.: Information: currency of life? *HFSP J.* **5**, 307–316 (2009)
12. Shannon, C.E.: A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948)
13. Shannon, C.E.: Coding theorems for a discrete source with a fidelity criterion. *IRE Natl. Conv. Rec.* **7**, 142–163 (1959)
14. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley-Interscience, New York (1991)
15. Csiszár, I.: Axiomatic characterizations of information measures. *Entropy* **10**, 261–273 (2008)
16. Shannon, C.: The bandwagon. *Trans. Inf. Theory* **2**, 3 (1956)
17. Rosenbluth, A., Wiener, N., Bigelow, J.: Behavior, purpose and teleology. *Philos. Sci.* **10**(1), 18–24 (1943)
18. Wiener, N.: *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge (1948)
19. Ashby, W.R.: *An Introduction to Cybernetics*. Chapman & Hall Ltd, London (1956)
20. Ashby, W.R.: Requisite variety and its implications for the control of complex systems. *Cybernetica* **1**, 83–99 (1958)
21. Touchette, H., Lloyd, S.: Information-theoretic limits of control. *Phys. Rev. Lett.* **84**(6), 1156–1159 (2000)
22. Touchette, H., Lloyd, S.: Information-theoretic approach to the study of control systems. *Physica A* **331**, 140–172 (2004)
23. Mitter, S.K.: Control with limited information. *Eur. J. Control* **7**(2–3), 122–131 (2001)
24. Marko, H.: The bidirectional communication theory—a generalization of information theory. *IEEE Trans. Inf. Theory* **21**, 1345–1351 (1973)
25. Massey, J.L.: Causality, feedback and directed information. In: *Proc. Intl. Symp. Info. Theory Applic. (ISITA-90)*, pp. 303–305 (1990)
26. Seger, J., Brockmann, H.J.: What is bet-hedging? *Oxf. Surv. Evol. Biol.* **4**, 182–211 (1987)

27. Lewontin, R.C., Cohen, D.: On population growth in a randomly varying environment. *Proc. Natl. Acad. Sci. USA* **62**, 1056–1060 (1969)
28. Real, L.A.: Fitness uncertainty and the role of diversification in evolution and behaviour. *Am. Nat.* **115**, 623–638 (1980)
29. Stearns, S.C.: Daniel Bernoulli (1738): evolution and economics under risk. *J. Biosci.* **25**, 221–228 (2000)
30. Wagner, A.: Risk management in biological evolution. *J. Theor. Biol.* **225**, 45–57 (2003)
31. Stephens, D.W.: Variance and the value of information. *Am. Nat.* **134**, 128–140 (1989)
32. Bergstrom, C.T., Lachmann, M.: Shannon information and biological fitness. In: *Information Theory Workshop IEEE '04*, San Antonio, Texas, pp. 50–54 (2004)
33. Kussell, E., Leibler, S.: Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078 (2005)
34. Donaldson-Matasci, M.C., Bergstrom, C.T., Lachmann, M.: The fitness value of information. *Oikos* **119**, 219–230 (2010)
35. Kelly, J.: New interpretation of information rate. *Bell Syst. Tech. J.* **35**, 917–926 (1956)
36. Breiman, L.: Optimal gambling systems for favorable games. In: *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, pp. 65–78. University of California Press, Berkeley (1961)
37. Algoet, P.H., Cover, T.M.: Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *Ann. Probab.* **16**, 876–898 (1988)
38. Barron, A.R., Cover, T.M.: A bound on the financial value of information. *IEEE Trans. Inf. Theory* **34**, 1097–1100 (1988)
39. Cover, T.M.: Shannon and investment. *IEEE Inf. Theory Soc. Newsl. (Special Golden Jubilee Issue)* 10–11 (1998)
40. Permuter, H.H., Kim, Y.-H., Weissman, T.: On directed information and gambling. In: *Proc. International Symposium on Information Theory (ISIT)*, Toronto, Canada (2008)
41. Perkins, T.J., Swain, P.S.: Strategies for cellular decision-making. *Mol. Syst. Biol.* **5**, 326 (2009)
42. Karlin, S., Taylor, H.M.: *A First Course in Stochastic Processes*. Academic Press, New York (1975)
43. Pack Kaelbling, L., Littman, M.L., Cassandra, A.R.: Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99–134 (1998)
44. Furstenberg, H., Kesten, H.: Products of random matrices. *Ann. Math. Stat.* **31**, 457–469 (1960)
45. Kingman, J.F.C.: Subadditive ergodic theory. *Ann. Probab.* **1**, 883–899 (1973)
46. Tanny, D.: On multitype branching processes in a random environment. *Adv. Appl. Probab.* **13**, 464–497 (1981)
47. Kramer, G.: Directed information for channels with feedback. PhD thesis, Swiss Federal Institute of Technology (ETH), Zurich (1998)
48. Kim, Y.-H.: A coding theorem for a class of stationary channels with feedback. *IEEE Trans. Inf. Theory* **25**, 1488–1499 (2008)
49. Witsenhausen, H.S.: Separation of estimation and control for discrete time systems. *Proc. IEEE* **59**, 1557–1566 (1971)
50. Berger, J.O.: *Statistical Decision Theory and Bayesian Analysis*. Springer, Berlin (1985)
51. Gastpar, M., Rimoldi, B., Vetterli, M.: To code, or not to code: lossy source-channel communication revisited. *IEEE Trans. Inf. Theory* **49**, 1147–1158 (2003)
52. Leslie, P.: On the use of matrices in certain population mathematics. *Biometrika* **33**, 183–212 (1945)
53. Kussell, E., Leibler, S., Grosberg, A.: Polymer-population mapping and localization in the space of phenotypes. *Phys. Rev. Lett.* **97**, 068101 (2006)
54. Iyengar, G.N., Cover, T.M.: Growth optimal investment in horse race markets with costs. *IEEE Trans. Inf. Theory* **46**, 2675–2683 (2000)
55. Tuljapurkar, S., Horvitz, C.C., Pascarella, J.B.: The many growth rates and elasticities of populations in random environments. *Am. Nat.* **162**, 489–502 (2003)
56. Haccou, P., Iwasa, Y.: Optimal mixed strategies in stochastic environments. *Theor. Popul. Biol.* **47**, 212–243 (1995)
57. Sasaki, A., Ellner, S.: The evolutionarily stable phenotype distribution in a random environment. *Evolution* **49**, 337–350 (1995)
58. Merton, R.C.: Lifetime portfolio selection under uncertainty: the continuous-time case. *Rev. Econ. Stat.* **51**, 247–257 (1969)
59. Godfrey-Smith, P.: Information in biology. In: Hull, D., Ruse, M. (eds.) *The Cambridge Companion to the Philosophy of Biology*, pp. 103–119. Cambridge University Press, Cambridge (2007)
60. Athreya, K.B., Ney, P.E.: *Branching Processes*. Springer, Berlin (1972)
61. Kurtz, T.G., Lyons, R., Pemantle, R., Peres, Y.: A conceptual proof of the Kesten-Stigum theorem for multi-type branching processes. In: Athreya, K.B., Jagers, P. (eds.) *Classical and Modern Branching Processes. IMA Volumes in Mathematics and its Applications*, vol. 84, pp. 181–185. Springer, New York (1997)

62. Bernoulli, D.: Specimen theoriae novae de mensura sortis. Pap. Imp. Acad. Sci. St. Petersburg **5**, 175–192 (1738)
63. Markowitz, H.: Portfolio selection. *J. Finance* **7**, 77–91 (1952)
64. Morsegaard Christensen, M.: On the history of the growth optimal portfolio. Unpublished (2005)
65. Samuelson, P.A.: The fallacy of maximizing the geometric mean in long sequences of investing or gambling. *Proc. Natl. Acad. Sci. USA* **68**, 2493–2496 (1971)
66. Maynard Smith, J., Price, G.: The logic of animal conflict. *Nature* **248**, 15–18 (1973)
67. Mills, S., Beatty, J.: The propensity interpretation of fitness. *Philos. Sci.* **46**, 263–286 (1979)
68. Beatty, J., Finsen, S.: Rethinking the propensity interpretation: a peek inside Pandora's box. In: Ruse, M. (ed.) *What the Philosophy of Biology Is, Essays Dedicated to David Hull*. Kluwer Academic, Dordrecht (1989)
69. Robson, A.J.: A biological basis for expected and non-expected utility. *J. Econ. Theory* **68**, 397–424 (1996)
70. Hartley, R.V.L.: Transmission of information. *Bell Syst. Tech. J.*, 535 (1928)