



Cite this article: Seoane LF, Solé RV. 2018
Information theory, predictability and the
emergence of complex life. *R. Soc. open sci.*
5: 172221.
<http://dx.doi.org/10.1098/rsos.172221>

Received: 20 December 2017

Accepted: 24 January 2018

Subject Category:

Biology (whole organism)

Subject Areas:

evolution/theoretical biology/computational
biology

Keywords:

complexity, emergence, computation,
evolution, predictability

Author for correspondence:

Luís F. Seoane

e-mail: seoane@mit.edu

Information theory, predictability and the emergence of complex life

Luís F. Seoane^{1,2,3} and Ricard V. Solé^{2,3,4}

¹Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

²ICREA-Complex Systems Lab, Universitat Pompeu Fabra (GRIB), Dr Aiguader 80, 08003 Barcelona, Spain

³Institut de Biologia Evolutiva, CSIC-UPF, Pg Maritim de la Barceloneta 37, 08003 Barcelona, Spain

⁴Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

LFS, 0000-0003-0045-8145; RVS, 0000-0001-6974-1008

Despite the obvious advantage of simple life forms capable of fast replication, different levels of cognitive complexity have been achieved by living systems in terms of their potential to cope with environmental uncertainty. Against the inevitable cost associated with detecting environmental cues and responding to them in adaptive ways, we conjecture that the potential for predicting the environment can overcome the expenses associated with maintaining costly, complex structures. We present a minimal formal model grounded in information theory and selection, in which successive generations of agents are mapped into transmitters and receivers of a coded message. Our agents are guessing machines and their capacity to deal with environments of different complexity defines the conditions to sustain more complex agents.

1. Introduction

Simple life forms dominate our biosphere [1] and define a lower bound of embodied, self-replicating systems. But life displays an enormously broad range of complexity levels, affecting many different traits of living entities, from their body size to their cognitive abilities [2]. This creates somewhat a paradox: if larger, more complex organisms are more costly to grow and maintain, why is not all life single-celled? Several arguments help provide a rationale for the emergence and persistence of complex life forms. As an instance, Gould [1] proposes that complexity is not a trait explicitly favoured by evolution. A review of fossil records convinces Gould that, across genera, phyla and the whole biosphere, we observe the expected random fluctuations around the more successful adaptation to life. In this big picture,

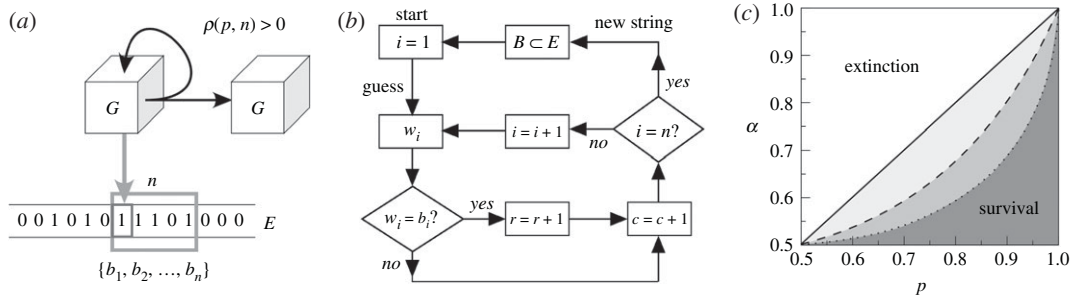


Figure 1. Predictive agents and environmental complexity. (a) An agent G interacts with an external environment E that is modelled as a string of random bits. These bits take value 0 with probability p and value 1 otherwise. The agent tries to guess a sequence of n bits at some cost, with a reward bestowed for each correctly guessed bit. The persistence and replication of the agent can only be granted if the balance between reward and cost is positive ($\rho_E^G > 0$). (b) For a machine attempting to guess n bits, an algorithmic description of its behaviour is shown as a flow graph. Each loop in the computation involves scanning a random subset of the environment $B = (b_1, \dots, b_n) \subset E$ by comparing each $b_i \in B$ to a proposed guess w_i . (c) A mean field approach to a certain kind of 1-guesser (modelled in the text through equations (1.1)–(1.3)) in environments of infinite size renders a boundary between survival ($\rho_E^G > 0$) and death ($\rho_E^G < 0$) as a function of the cost–reward ratio (α) and of relevant parameters for the 1-guesser model (p in this case).

bacteria are the leading life form and the complexity of every other living system is the product of a random drift. Complex life would never be explicitly favoured, but a complexity wall exists right below bacteria: simpler forms fail to subsist. Hence, a random fluctuation is more likely to produce more complex forms, falsely suggesting that evolution promotes complexity.

Major innovations in evolution involve the appearance of new types of agents displaying cooperation while limiting conflict [3,4]. A specially important innovation involved the rise of cognitive agents, namely those capable of sensing their environments and reacting to their changes in a highly adaptable way [5]. These agents were capable of dealing with more complex, non-genetic forms of information. The advantages of such cognitive complexity become clear when considering their potential to better predict the environment, thus reducing the average hazards of unexpected fluctuations. As pointed out by Francois Jacob, an organism is ‘a sort of machine for predicting the future—an automatic forecasting apparatus’ ([6, p. 9]; see also [7,8]). The main message is that *foreseeing* the future is a crucial ability to cope with uncertainty. If the advantages of prediction overcome the problem of maintaining and replicating the costly structures needed for inference, more complex information-processing mechanisms might be favoured under the appropriate circumstances.

This kind of problem has been addressed within ecological and evolutionary perspectives. One particularly interesting problem concerns the potential of some types of organisms to develop cognitive potential for prediction. Are all living systems capable to develop such feature? What is the limit of predictive power for a given group, and how is it affected by the lifestyle? Plants, for example, have no nervous system but exhibit some interesting capacities for decision making, self/non-self discrimination, or error correction [9]. Studies involving the evolution of prediction in simulated plants reveal that increased predictability of available resources was achieved by a proper assessment of environmental variability [10]. Some of the molecular mechanisms that pervade plant responses seem to deal with switch-like changes triggered by genetic networks [11]. In this case, growth, seed production or germination correlate with the degree of environmental fluctuations. A key conclusion of relevance to our paper is that some reproduction strategies are not selected in given environments due to a lack of predictability.

Here we aim at providing a minimal model that captures these trade-offs. In doing so, we characterize thoroughly an evolutionary driver that can push towards evermore complex life forms. We adopt an information theory perspective in which agents are inference devices interacting with a Boolean environment. For convenience, this environment is represented by a tape with ones and zeros, akin to non-empty inputs of a Turing machine (figure 1a). The agent G locates itself in a given position and tries to predict each bit of a given sequence of length n —hence it is dubbed an n -guesser. Each attempt to predict a bit involves some cost c , while a reward r is received for each successful prediction. 1-guessers are simple and assume that all bits are uncorrelated, while ($n > 1$)-guessers find correlations and can get a larger benefit if some structure happens to be present in the environment. A whole n -bit prediction cycle can be described as a program (figure 1b). A survival function ρ depends on the number of attempts to guess bits and the number of correct predictions. Successful guessers have a positive balance between

reward and prediction cost. They get replicated and pass on their inference abilities. Otherwise, the agent fails to replicate and eventually dies.

As a simple illustration of our approach, consider a 1-guesser living in an infinitely large environment E where uncorrelated bits take value 1 with probability p , and 0 with probability $1 - p$. The average performance of a guesser G when trying to infer bits from E is given by \bar{p}_E^G ; this is, the likelihood of emitting a correct guess:

$$\bar{p}_E^G = p^G(1)p + p^G(0)(1 - p), \quad (1.1)$$

where $p^G(k)$ is the frequency with which the guesser emits the bit value $k \in \{0, 1\}$. The best strategy possible is to emit always the most abundant bit in the environment. Then

$$\bar{p}_E^G = \max\{p, 1 - p\}. \quad (1.2)$$

Without loss of generality, let's assume 1 is the most common bit. Then the average reward minus cost obtained by such a guesser reads:

$$\rho_E^G = pr - c = (p - \alpha)r. \quad (1.3)$$

This curve trivially dictates the average survival or extinction of the optimal 1-guessers in infinite, unstructured environments as a function of the cost–reward ratio $\alpha \equiv c/r$ (grey area subtended by the solid diagonal line in figure 1c). Note that this parameter α encodes the severity of the environment—i.e. how much does a reward pay off given the investment needed to obtain it. Further, note that any more complex guessers (like the ones described in successive sections) would always fare worst in this case: they would potentially pay a larger cost to infer some structure where there is none. This results in narrower survival areas qualitatively represented by shades of grey subtended by the discontinuous lines in figure 1c. Again, these reduced niches for more complex bit-guessers would come along because there is not structure to be inferred; but that can change if correlations across environmental bits appear, as will be seen below.

The idea of autonomy and the fact that predicting the future implies performing some sort of computation suggests that a parsimonious theory of life's complexity needs to incorporate reproducing individuals (and eventually populations) and information (they must be capable of predicting future environmental states). These two components define a conflict and an evolutionary trade-off. Being too simple means that the external world is perceived as a source of noise. Unexpected fluctuations can be harmful, and useful structure cannot be harnessed in your benefit. Becoming more complex (hence able to infer larger structures, if they exist) implies a risk of not being able to gather enough energy to support and replicate the mechanisms for inference. As will be shown below, it is possible to derive the critical conditions to survive as a function of the agent's complexity and to connect these conditions to information theory. As advanced above, this allows us to characterize mathematically a scenario in which a guesser's complexity is explicitly selected for. Actual living beings will embody the necessary inference mechanisms in their morphology, or in their genetic or neural networks. Instead of developing specific models for each of these alternative implementations, we resort to mathematical abstractions based on bit-strings, whose conclusions will be general and apply broadly to any chosen strategy.

2. Evolution and information theory

Key aspects of information theory relate deeply to formulations in statistical physics [12–14] and there have been several calls to further integrate information theory in biological research [15–23]. This theory shall play important roles in population or ecosystems dynamics, in regulatory genomics, and in chemical signal processing among others [7,24–41], but a unifying approach is far from complete. Given its generality and power, information theory has also been used to address problems that connect Darwinian evolution and far from equilibrium thermodynamics [42–46]. In its original formulation, Shannon's information theory [47,48] considers symbols being conveyed from a transmitter to a receiver through a channel. Shannon only deals with the efficiency of the channel (related to its noise or reliability) and the entropy of the source. This theory ignores the content of the emitted symbols, despite the limitations of such an assumption [18,49].

A satisfactory connection between natural selection and information theory can be obtained by mapping our survival function ρ into Shannon's transmitter–receiver scheme. To do so, we consider replicators at an arbitrary generation T attempting to 'send' a message to (i.e. getting replicated into) a later generation $T + 1$. Hence, the older generation acts as a transmitter, the newer one becomes a receiver, and the environment and its contingencies constitute the channel through which the embodied message must be conveyed (figure 2a). From a more biological perspective, we can think of a genotype as

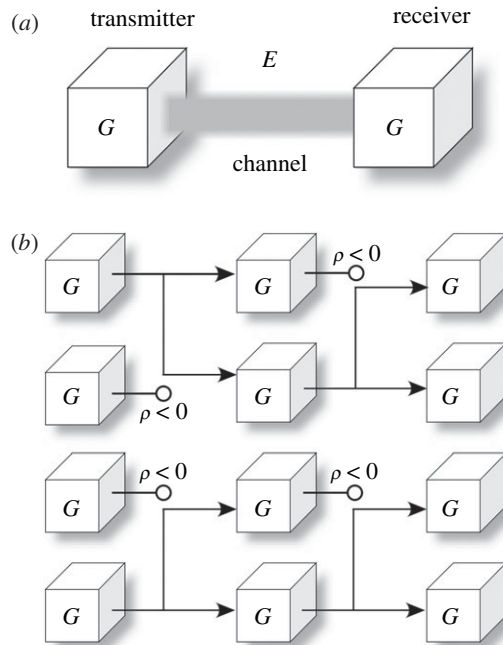


Figure 2. Information and evolution through natural selection. (a) The propagation of a successful replicator can be understood in terms of a Shannon-like transmission process from one generation to the next in which older generations play the role of a transmitter, younger generations that of a receiver and the environment constitutes a noisy channel. (b) A simple diagram of the underlying evolution of a population of bit-guessers. The survival and replication of a given agent G is indicated by branching, whereas failure to survive is indicated with an empty circle as an endpoint.

a generative model (the instructions in an algorithm) that produces a message that must be transmitted. That message would be embodied by a phenotype and it includes every physical process and structure dictated by the generative model. As discussed by von Neuman & Burks [50], any replicating machine must pass on a physically embodied copy of its instructions—hence the phenotype must also include a physical realization of the algorithm encoded by the genotype.¹ Finally, any evolutionary pressure (including the interaction with other replicating signals) can be included as contrivances of the channel.

Following a similar idea of messages being passed from one generation to the next one, Maynard-Smith [18] proposes that the replicated genetic message carries *meaningful information* that must be protected *against* the channel contingencies. Let us instead depart from a replicating message devoid of meaning. We realize that the channel itself would convey more reliably those messages embodied by a phenotype that better deals with the environmental (i.e. channel) conditions. Dysfunctional messages are removed due to natural selection. Efficient signals get more space in successive generations (figure 2b). Through this process, *meaningful* bits of environmental information are *pumped* into the replicating signals, such that the information in future messages will *anticipate* those channel contingencies. In our picture, meaningful information is not protected against the channel conditions (including noise), but emerges naturally from them.

2.1. Messages, channels and bit-guessers

Let us first introduce our implementation of environments (channels), messages and the replicating agents. The latter will be dubbed *bit-guessers* because efficient transmission will be equivalent to accurately predicting channel conditions—i.e. to correctly guessing as many bits about the environment as possible. The notation that follows may seem arid, so it is good to retain a central picture (figure 3): guessers G possess a generative model I^G that must produce messages that fare well in an environment

¹Note that many of the phenotypic structures built in order to get replicated are later dismissed (think e.g. about the germ versus somatic cell lines). We present a clear division between genotype and phenotype for sake of illustration. We are aware of the murky frontier between these concepts.

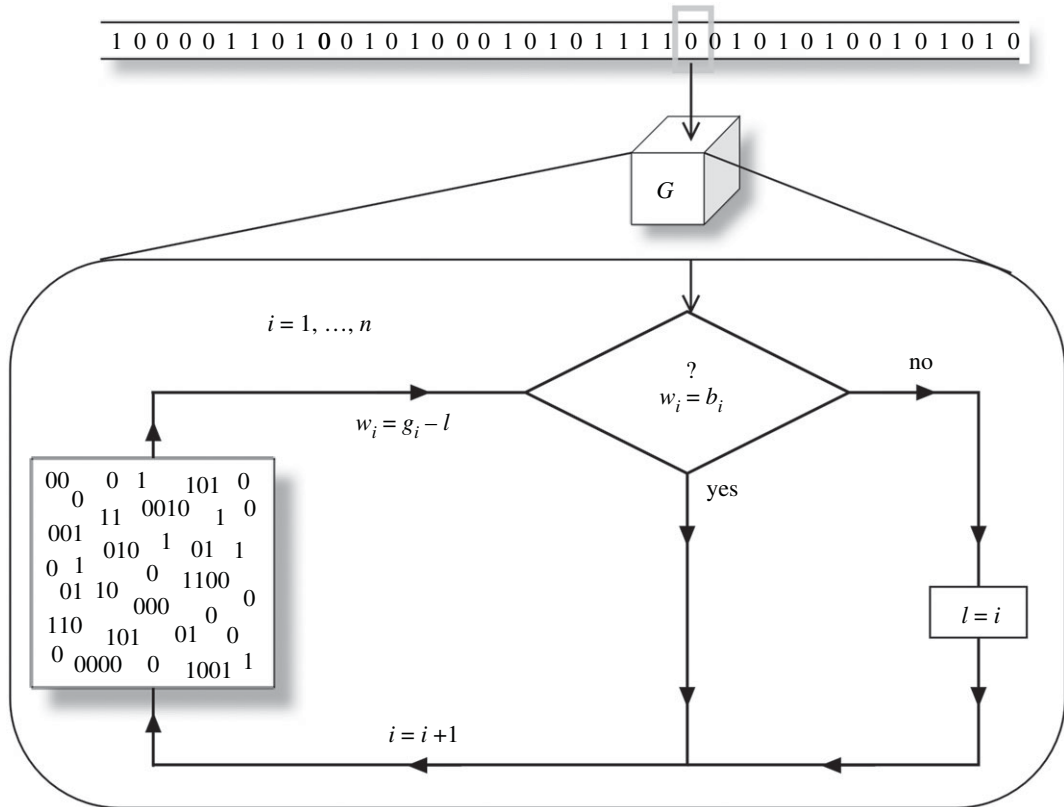


Figure 3. From a generative model to inference about the world. A diagrammatic representation of the algorithmic logic of the bit guessing machine. Our n -guesser contains a generative model (represented by a pool of words) from which it draws guesses about the environment. If a bit is successfully inferred, the chosen conjecture is pursued further by comparing a new bit. Otherwise, the inference is reset.

E. Both these messages and the environments are modelled as strings of bits. What follows is a rigorous mathematical characterization of how the different bit sequences are produced.

Let us consider m -environments, strings made up of m sorted random bits. We might consider one single such m -environment—i.e. one realization E of m sorted random bits ($e_i \in E, i = 1, \dots, m; e_i \in \{0, 1\}$). Alternatively, we might work with the ensemble E^m of all m -environments—i.e. all possible environments of the same size ($e_{i,l} \in E_l, i = 1, \dots, m; \text{where } E_l \in E^m, l = 1, \dots, 2^m$)—or we might work with a sample \hat{E}^m of this ensemble ($E_l \in \hat{E}^m, l = 1, \dots, |\hat{E}^m|; \text{where } \hat{E}^m \subset E^m$). We might evaluate the performance of our bit-guessers in single m -environments, in a whole ensemble, or in a sample of it.

These m -environments model the channels of our information theory approach. Attempting to transmit a message through this channel will be implemented by trying to guess n -sized words from within the corresponding m -environment. More precisely, given an n -bit message W (with $n < m$) which an agent tries to transmit, we extract an n -sized word ($B \subset E$) from the corresponding m -environment. Therefore, we choose a bit at a random position in E and the successive $n - 1$ bits. These make up the $b_i \in B$, which are compared with the $w_i \in W$. Each w_i is successfully transmitted through the channel if $w_i = b_i$. Hence attempting to transmit messages effectively becomes an inference task: if a guesser can anticipate the bits that follow, it has a greater chance of sending messages through. Messages transmitted equal bits copied into a later generation, hence increasing the fitness of the agent.

In this paper, we allow bit-guessers a minimal ability to react to the environment. Hence, instead of attempting to transmit a fixed word W , they are endowed with a generative model Γ^G . This mechanism (explained below) builds the message W as a function of the broadcast history:

$$w_i = w_i(w_1, \dots, w_{i-1}; b_1, \dots, b_{i-1}).$$

Hence, the fitness of a generative model is rather based on the ensemble of messages that it can produce.² To evaluate this, our guessers attempt to transmit n -bit words many (N_g) times through a same channel.

²There is a compromise worth investigating between the fidelity of the message that an agent tries to convey and its ability to react to environmental conditions in real time. Exploring this trade-off is left for future work. By now, the reaction capabilities of our bit-guessers will be kept to a minimum.

For each one of these broadcasts, a new n -sized word $B^j \subset E$ (with $b_i^j \in B^j$ for $j = 1, \dots, N_g$ and $i = 1, \dots, n$) is extracted from the same m -environment; and the corresponding W^j are generated, each based on the broadcast history as dictated by the generative model (see below).

We can calculate different frequencies with which the guessers or the environments present bits with value $k, k' \in \{0, 1\}$:

$$p^G(k; i) = \frac{1}{N_g} \sum_{j=1}^{N_g} \delta(w_{i'}^j, k), \quad (2.1)$$

$$p_E(k'; i) = \frac{1}{N_g} \sum_{j=1}^{N_g} \delta(b_{i'}^j, k'), \quad (2.2)$$

$$p_{G,E}(k, k'; i) = \frac{1}{N_g} \sum_{j=1}^{N_g} \delta(w_{i'}^j, k) \delta(b_{i'}^j, k'), \quad (2.3)$$

$$p_E^G(i) = \frac{1}{N_g} \sum_{j=1}^{N_g} \delta(w_{i'}^j, b_{i'}^j) \Rightarrow \quad (2.4)$$

and
$$\Rightarrow \bar{p}_E^G = \frac{1}{n} \sum_{i=1}^n p_E^G(i); \quad (2.5)$$

with $\delta(x, y)$ being Dirac's delta. Note that $p^G(k; i)$ has a subtle dependency on the environment (because G may react to it) and that \bar{p}_E^G indicates the average probability that guesser G successfully transmits a bit through channel E .

Thanks to these equations we can connect with the cost and reward functions introduced before. For every bit that attempts to be transmitted, a cost c is paid. A reward $r = c/\alpha$ is cashed in only if that bit is successfully received. α is a parameter that controls the pay-off. The survival function reads

$$\rho_E^G(\alpha) = (\bar{p}_E^G - \alpha)r \quad (2.6)$$

and \bar{p}_E^G can be read from equation (2.5). As a rule of thumb, if $\bar{p}_E^G > \alpha$ the given guesser fares well enough in the proposed environment.

It is useful to quantify the entropy per bit of the messages produced by G :

$$H(G) = -\frac{1}{n} \sum_{i=1}^n \sum_k p^G(k; i) \log(p^G(k; i)) \quad (2.7)$$

and the mutual information between the messages and the environment:

$$I(G : E) = \frac{1}{n} \sum_{i=1}^n \sum_{k, k'} p_{G,E}(k, k'; i) \times \log \left(\frac{p_{G,E}(k, k'; i)}{p^G(k; i) p_E(k'; i)} \right). \quad (2.8)$$

To evaluate the performance of a guesser over an ensemble \hat{E}^m of environments (instead of over single environments), we attempt N_g broadcasts over each of N_e different environments ($E_l \in \hat{E}^m$, $l = 1, \dots, N_e \equiv \|\hat{E}^m\|$) of a given size. For simplicity, instead of labeling $b_{i'}^j$, we stack together all $N_g \times N_e$ n -sized words W^j and B^j . This way $b_i^j \in B^j$ and $w_i^j \in W^j$ for $i = 1, \dots, n$ and $j = 1, \dots, N_g N_e$. We have $p^G(k; i)$, $p_{\hat{E}^m}(k'; i)$, $p_{G, \hat{E}^m}(k, k'; i)$, $p_{\hat{E}^m}^G(i)$ and $\bar{p}_{\hat{E}^m}^G$ defined just as in equations (2.2)–(2.5), only with j running through $j = 1, \dots, N_g N_e$. Also as before, we average the pay-off across environments to determine whether a guesser's messages get successfully transmitted or not given α and the length m of the environments in the ensemble

$$\rho_{\hat{E}^m}^G(\alpha) = (\bar{p}_{\hat{E}^m}^G - \alpha)r. \quad (2.9)$$

Note that

$$I(G : \hat{E}^m) = \frac{1}{n} \sum_{i=1}^n \sum_{k, k'} p_{G, \hat{E}^m}(k, k'; i) \times \log \left(\frac{p_{G, \hat{E}^m}(k, k'; i)}{p^G(k; i) p_{\hat{E}^m}(k'; i)} \right) \quad (2.10)$$

is different from

$$\langle I(G : E) \rangle_{\hat{E}^m} = \frac{1}{N_e} \sum_{l=1}^{N_e} I(G : E_l). \tag{2.11}$$

We use $\langle \cdot \rangle_{\hat{E}^m}$ to indicate averages across environments of an ensemble \hat{E}^m .

Finally, we discuss the generative models at the core of our bit-guessers. These are mechanisms that produce n -sized strings of bits, partly as a reaction to contingencies of the environment. Such message-generating processes Γ^G could be implemented in different ways, including artificial neural networks (ANNs) [51], spiking neurons [52], Bayesian networks [53,54], Turing machines [55], Markovian chains [56], ϵ -machines [57], random Boolean networks (RBNs) [58], among others. These devices elaborate their guesses through a series of algorithms (e.g. back-propagation, message passing or Hebbian learning) provided they have access to a sample of their environment.

In the real world, trial and error and evolution through natural selection would be the algorithm wiring the Γ^G (or, in a more biological language, a genotype) into our agents. The dynamics of such evolutionary process are very interesting. However, in this paper, we aim at understanding the limits imposed by a channel's complexity and the cost of inference, not the dynamics of how those limits may be reached. Therefore, we assume that our agents perform an almost perfect inference given the environment where they live. This best inference will be hard-wired in the guesser's generative model Γ^G as explained right ahead.

A guesser's generative model usually depends on the environment where it is deployed, so we note $\Gamma^G \equiv \Gamma_E^G$. This Γ_E^G will consist of a pool of bits $g_i \in \Gamma_E^G$ (figure 3) and a series of rules dictating how to emit those bits: either in a predetermined order or as a response to the channel's changing conditions. Whenever we pick up an environment $E = \{e_i, i = 1, \dots, m\}$, the best first guess possible will be the bit (0 or 1) that shows up with more frequency. Hence

$$\Gamma_E^G(1) \equiv g_1 = \max_k \{p_E(k; 1)\}. \tag{2.12}$$

If both 0 and 1 appear equally often we choose 1 without loss of generality. If the agent succeeds in its first guess, its safest next bet is to emit the bit (0 or 1) that more frequently follows g_1 in the environment. We proceed similarly if the first two bits have been correctly guessed, if the first three bits have been correctly guessed, etc. We define $p_{B|\Gamma}(k; i)$ as the probability of finding $k = \{0, 1\}$ at the i th position of the B^j word extracted from the environment, provided that the guess so far is correct:

$$p_{B|\Gamma}(k^j; i) = \frac{1}{Z(i)} \sum_{j=1}^m \delta(b_{i'}^j, k^j) \prod_{i'=1}^{i-1} \delta(b_{i'}^j, g_{i'}). \tag{2.13}$$

The index j , in this case, labels all n -sized words within the environment $(b_{i'}^j \in B^j) \subset E$ and $Z(i)$ is a normalization constant that depends on how many words in the environment match Γ_E^G up to the $(i - 1)$ th bit:

$$Z(i) = \sum_{j=1}^m \prod_{i'=1}^{i-1} \delta(b_{i'}^j, g_{i'}). \tag{2.14}$$

It follows

$$\Gamma_E^G(i = 2, \dots, n) \equiv g_i = \max_k \{p_{B|\Gamma}(k^j; i)\}. \tag{2.15}$$

Note that the pool of bits in Γ_E^G consists of an n -sized word, which is what they try to emit through (i.e. it constitutes the guess about) the channel. If a guesser would not be able to react to environmental conditions, the word W that is actually generated at every emission would be the same in every case and $w_i^j = g_i$ always; but we allow our guessers a minimal reaction if one of the bits fails to get through (i.e. if one of the guesses is not correct). This minimal reaction capacity by our guessers results in:

$$w_i^j = \Gamma_E^G(i - l) = g_{i-l}, \tag{2.16}$$

where l is the largest i at which $w_i^j \neq b_i^j$. This means that a guesser restarts the broadcast of Γ_E^G whenever it makes a mistake.³

³Note that more elaborated guessers would not only reset their guess. They might browse through a tree with conditional instructions at every point. Besides an extended memory to store the growing number of branches, they would also require nested *if-else* instructions. On the other hand, ANNs or Bayesian networks might implement such tree-browsing without excessive *if-else* costs.

Altogether, our guesser consists of a generative model I^G that contains a pool of bits and a simple conditional instruction. This is reflected in the flow chart in figure 3.

We have made a series of choices regarding how to implement environmental conditions. These choices affect how some randomness enters the model (reflected in the fact that, given an environment E , a guesser might come across different words $B^l \subset E$) and also how we implement our guessers (including their minimal adaptability to wrong guesses). We came up with a scheme that codes guessers, environments (or channels), and messages as bit-strings. This allows us a direct measurement of information-theoretical features which are suitable for our discussion, but the conclusions at which we arrive should be general. Survival will depend on an agent's ability to embody *meaningful* information about its environment. This will ultimately be controlled by the underlying cost–efficiency trade-off.

Because of the minimal implementation discussed, all bit-guessers of the same size are equal. Environmental ensembles of a given size are also considered equivalent. Hence, the notation is not affected if we identify guessers and environments by their sizes. Accordingly, in the following we substitute the labels G and E by the more informative ones n and m respectively. Hence $\rho_{E_m}^G(\alpha)$ becomes $\rho_m^n(\alpha)$, \bar{p}_E^G becomes \bar{p}_m^n , etc.

3. Results

The question that motivates this paper relates to the trade-off between fast replication versus the cost of complex inference mechanisms. To tackle this we report a series of numerical experiments. Some of them deal with guessers in environment ensembles of fixed size, others allow guessers to switch between environment sizes to find a place where to thrive.

Our core finding is that the complexity of the guessers that can populate a given environment is determined by the complexity of the latter. (In information-theoretical terms, the complexity of the most efficiently replicated message follows from the predictability of the channel.) Back to the fast replication versus complexity question, we find environments for which simple guessers die off, but in which more complex life flourishes—thus offering a quantifiable model for real-life excursions in biological complexity.

Besides verifying mathematically that the conditions for complex life exist, our model allows us to explore and quantify when and how guessers may be pushed to m -environments of one size or another. We expect to use this model to investigate this question in future papers. As neat examples, at the end of this paper we report (i) the evolutionary dynamics established when guessers are forced to compete with each other and (ii) how the fast replication versus complexity trade-off is altered when resources can be exhausted. These are two possible evolutionary drivers of complex life, as our numerical experiments show.

3.1. Numerical limits of guesser complexity

Figure 4 shows \bar{p}_m^n , the average probability that n -guessers correctly guess 1 bit in m -environments. The 1-guesser (that lives off maximally decorrelated bits given the environment) establishes a lower bound. More complex machines will guess more bits on average, except for infinite environment size $m \rightarrow \infty$, at which point all guessers have equivalent predictive power.

As m grows, environments get less and less predictable. Importantly, the predictability of shorter words decays faster than that of larger ones, thus enabling guessers with larger n to survive where others would perish. There are 2^n possible n -words, of which m are realized in each m -environment. When $m \gg 2^n$, the environment implements an efficient, ergodic sampling of all n -words—thus making them maximally unpredictable. When $n \lesssim m < 2^n$ the sampling of n -sized words is far from ergodic and a non-trivial structure is induced in the environment because the symmetry between n -sized words is broken—they cannot be equally represented due to finite size sampling effects.

This allows that complex guessers (those with the ability to contemplate larger words, keep them in memory and make choices regarding information encoded in larger strings) can guess more bits, on average, than simpler agents. In terms of messages crossing the channel, while shorter words are meaningless and basically get transmitted (i.e. are correctly guessed) by chance alone, larger words might contain meaningful, non-trivial information that get successfully transmitted because they cope with the environment in an adequate way.

Note that this symmetry breaking to favour predictability of larger words is just a mechanism that allows us to introduce correlations in a controlled and measurable way. In the real world, this mechanism

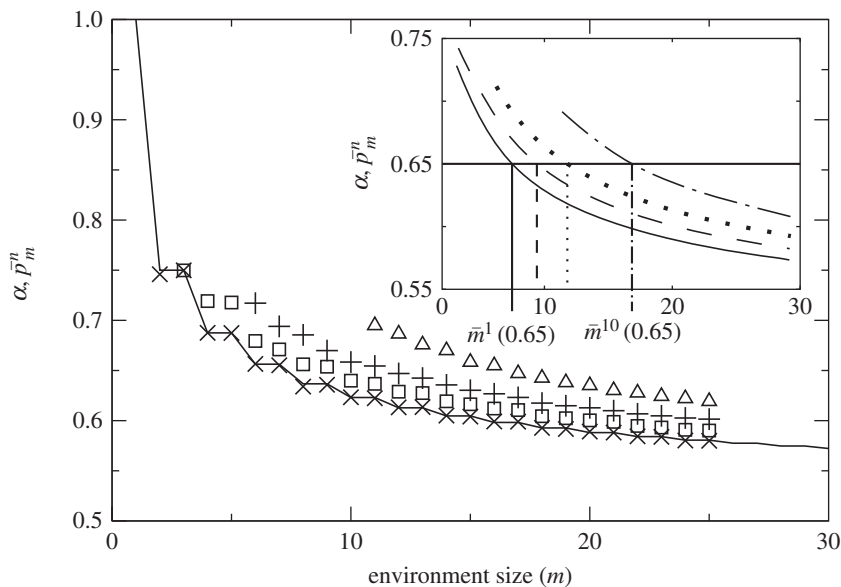


Figure 4. Probability of correctly guessing a bit in environment ensembles of constant size. \bar{p}_m^n , average probability that n -guessers correctly guess 1 bit in m -environments for different n values. Here \bar{p}_m^1 can be computed analytically (solid line in the main plot) and marks an average, lower predictability boundary for all guessers. In the inset, the data have been smoothed and compared with a given value of α (represented by a horizontal line). At the intersection between this line and \bar{p}_m^n we find $\bar{m}^n(\alpha)$, the environment size at which n -sized agents guess just enough bits to survive given α . Note that n -guessers are evaluated only in environments of size $m \geq n$.

might correspond to asymmetries between dynamical systems in temporal or spatial scales. Although our implementation is rather ad hoc (suitable to our computational and conceptual needs), we propose that similar mechanisms might play important roles in shaping life and endowing the universe with meaningful information. Indeed, it might be extremely rare to find a kind of environment in which words of all sizes become non-informative simultaneously.

The mutual information between a guesser’s response and the environment (i.e. between broadcast messages and channel conditions) further characterizes the advantages of more complex replicators. Figure 5a shows $I(G : E_m)$ and $\langle I(G : E) \rangle_{E_m}$. As we noted above, these quantities are not the same. Let us focus on 1-guessers for a moment to clarify what these quantities encode.

Given an m -environment, 1-guessers have got just one bit that they try to emit repeatedly. They do not react to the environment—there is not room for any reaction within one bit, so their guess is persistently the same. The mutual information between the emitted bit and the arbitrary words $B \subset E$ that 1-guessers come across is precisely zero, as shown in the inset of figure 5a. Hence, $\langle I(G : E) \rangle_{E_m}$ captures the mutual information due to the slight reaction capabilities of guessers to the environmental conditions.

While the bits emitted by 1-guessers do not correlate with $B \subset E$, they do correlate with each given E as they represent the most frequent bit in the environment. Accordingly, the mutual information between a 1-guesser and the aggregated environments (reflected by $I(G : E_m)$) is different from zero (main panel of figure 5a). To this quantity contribute both the reaction capability of guessers and the fact that they have hard-wired a near-optimal guess in Γ_E^G , as explained in §2.1.

We take the size of a guesser n as a crude characterization of its complexity. This is justified because larger guessers can store more complex patterns. $\langle H(G) \rangle_{E_m}$ indicates that more complex guessers look more entropic than less complex ones (figure 5b). Larger guessers come closer to the entropy level of the environment (black thick line in figure 5b), which itself tends rapidly to $\log(2)$ per bit. Better performing guessers appear more disordered to an external observer even if they are better predictors when considered within their context. Note that $\langle H(G) \rangle_{E_m}$ is built based on the bits actually emitted by the guessers. In biological terms, this would mean that this quantity correlates with the complexity of the phenotype. For guessers of fixed size n , we observe a slight decay of $\langle H(G) \rangle_{E_m}$ as we proceed to larger environments.

The key question is whether the pay-off may be favourable for more complex guessers provided that they need a more costly machinery in order to get successfully reproduced. As discussed above, if we would use e.g. ANN or Bayesian inference graphs to model our guessers, a cost could be introduced for

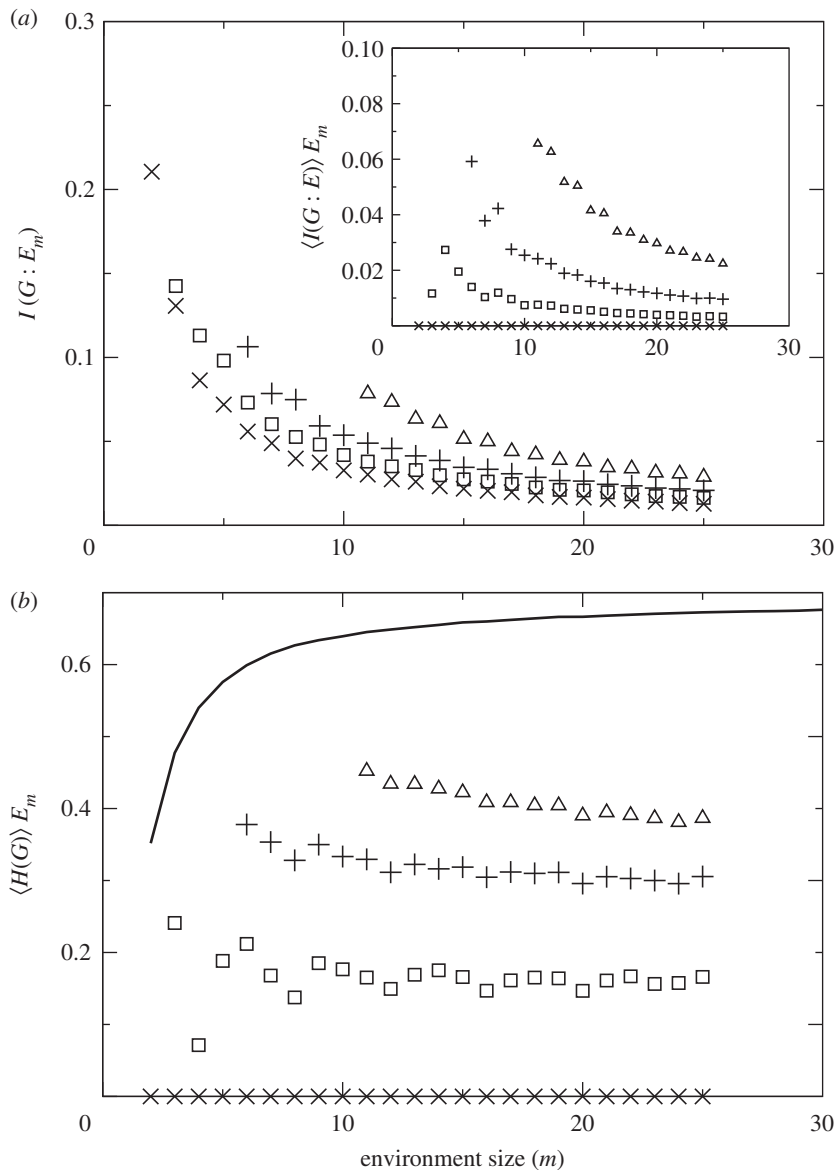


Figure 5. Mutual information and entropy. Guessers with $n = 1$ (crosses), $n = 2$ (squares), $n = 5$ (pluses) and $n = 10$ (triangles) are presented. (a) $I(G : E_m)$ and $\langle I(G : E) \rangle E_m$ (inset) quantify the different sources of information that allow more complex guessers to thrive in environments in which simpler life is not possible. (b) The entropy of a guesser’s message given its environment seems roughly constant in these experiments despite the growing environment size. This suggests an intrinsic measure of complexity for guessers. Larger guessers look more random even if they might carry more meaningful information about their environment. The thick black line represents the average entropy of the environments (which approaches $\log(2)$) against which the entropy of the guessers can be compared.

the number of units, nodes or hidden variables. These questions might be worth studying somewhere else. Here we are interested in the mathematical existence of such favourable trade-off for more complex life. To keep the discussion simple, bit-guessers incur only in a cost proportional to the number of bits that they try to transmit. Note that we do not lose generality, because such limit cost shall always exist. Equation (2.9) captures all the forces involved: the cost of transmitting longer messages versus the reward of a successful transmission.

Guessers of a given size survive in an environment ensemble if, on average, they can guess enough bits of the environment or, using the information theory picture, if they can convey enough bits through the channel (in any case, they survive if $\bar{p}_m^n > \alpha$, which implies $\rho_m^n > 0$). Setting a fixed value of α we find out graphically $\bar{m}^n(\alpha)$, the largest environment at which n -guessers survive (figure 4, inset). Because m -environments look more predictable to more complex guessers we have that $\bar{m}^n(\alpha) > \bar{m}^{n'}(\alpha)$ if $n > n'$. This guarantees that for $\alpha > 0.5$ there always exist m -environments from which simple life is banned

while more complex life can thrive—i.e. situations in which environmental complexity is an explicit driver towards more complex life forms.

This is the result that we sought. The current model allows us to illustrate mathematically that limit conditions exist under which more complex and costly inference abilities can overcome the pressure for fast and cheaper replication. Also, the model allows for explicit, information-theoretically based quantification of such a limit.

3.2. Evolutionary drivers

Despite its laborious mathematical formulation, we think that our bit-guesser model is very simple and versatile. We think that it can easily capture fundamental information-theoretical aspects of biological systems. In future papers, we intend to use it to further explore relationships between guessers and environments, within ecological communities, or in more simple symbiotic or parasitic situations. To illustrate how this could work out, we present now some minimal examples.

Let us first explore some dynamics in which guessers are encouraged to explore more complex environments, but this same complexity can become a burden. As before, let us evaluate an n -guesser $N_g \cdot N_e$ times in a sample of the m -environment ensemble. Let us also look at $\hat{\rho}_m^n(\alpha, N_g, N_e)$, the accumulated reward after these $N_g \cdot N_e$ evaluations—note that $\hat{\rho}_m^n$ is an empirical random variable now. If $\hat{\rho}_m^n(\alpha, N_g, N_e) > 0$, the n -guesser fares well enough in this m -environment and it is encouraged to explore a more complex one. As a consequence, the guesser is promoted to an $(m + 1)$ -environment, where it is evaluated again. If $\hat{\rho}_m^n(\alpha, N_g, N_e) < 0$, this m -environment is excessively challenging for this n -guesser, and it is demoted to an $(m - 1)$ -environment. Note that the n -guesser itself remains with a fixed size throughout. It is the complexity of the environment that changes depending on the reward accumulated.

As we repeatedly evaluate the n -guesser, some dynamics are established which let the guesser explore more or less complex environments. The steady state of these dynamics is characterized by a distribution $P^n(m, \alpha)$. This tells us the frequency with which n -guessers are found in environments of a given size (figure 6a). Each n -guesser has its own distribution that captures the environmental complexity that the guesser deals more comfortably with. The overlaps and gaps between $P^n(m, \alpha)$ for different n suggest that: (i) some guessers would engage in harsh competition if they needed to share environments of a given kind and (ii) there is room for different guessers to get segregated into environments of increasing complexity.

The average

$$\hat{m}^n(\alpha) = \sum_m m P^n(m, \alpha) \quad (3.1)$$

should converge to $\hat{m}^n(\alpha) \simeq \bar{m}^n(\alpha)$ under the appropriate limit. This is, if we evaluate the guessers numerically enough times, the empirical value $\hat{m}^n(\alpha)$ should converge to the mean field value $\bar{m}^n(\alpha)$ shown in the inset of figure 4. Figure 6b shows dynamically derived averages $\hat{m}^n(\alpha)$ and some deviations around them as a function of α .

It is easily justified that guessers drop to simpler environments if they cannot cope with a large complexity. It is less clear why they should seek more complicated environments if they thrive in a given one. This might happen if some external force drives them; for example, if simpler guessers (which might be more efficient in simpler environments) have already crowded the place. Let us remind, from figure 4, how given an environment size more complex guessers can always accumulate a larger reward. This might suggest that complex guessers always pay off, but the additional complexity might become a burden in energetic terms—consider, e.g. the exaggerated metabolic cost of mammal brains. It is non-trivial how competition dynamics between guessers of different size can play out. Let us gain some insights by looking at a simple model.

n -guessers with $n = 0, 1, 2, 3$ and 4 were randomly distributed occupying 100 environments, all of them with fixed size m . These guessers were assigned an initial $\hat{\rho}_i(t = 0) = n\rho_0$. Here, $i = 1, \dots, 100$ labels each one of the 100 available guessers. Larger guessers start out with larger $\hat{\rho}_i(t = 0)$ representing that they come into being with a larger metabolic load satisfied. A 0-guesser represents an unoccupied environment. New empty environments might appear only if actual ($n \neq 0$) guessers die, as we explain below. We tracked the population using $P^m(n, t)$, the proportion of 0-, 1-, 2-, 3- and 4-guessers through time.⁴

⁴These experiments were the more computationally demanding, that is why we took $n = 1, 2, 3, 4$ instead of the values $n = 1, 2, 5, 10$ used throughout the paper. The insights gained from the simulations do not depend on the actual values of n .

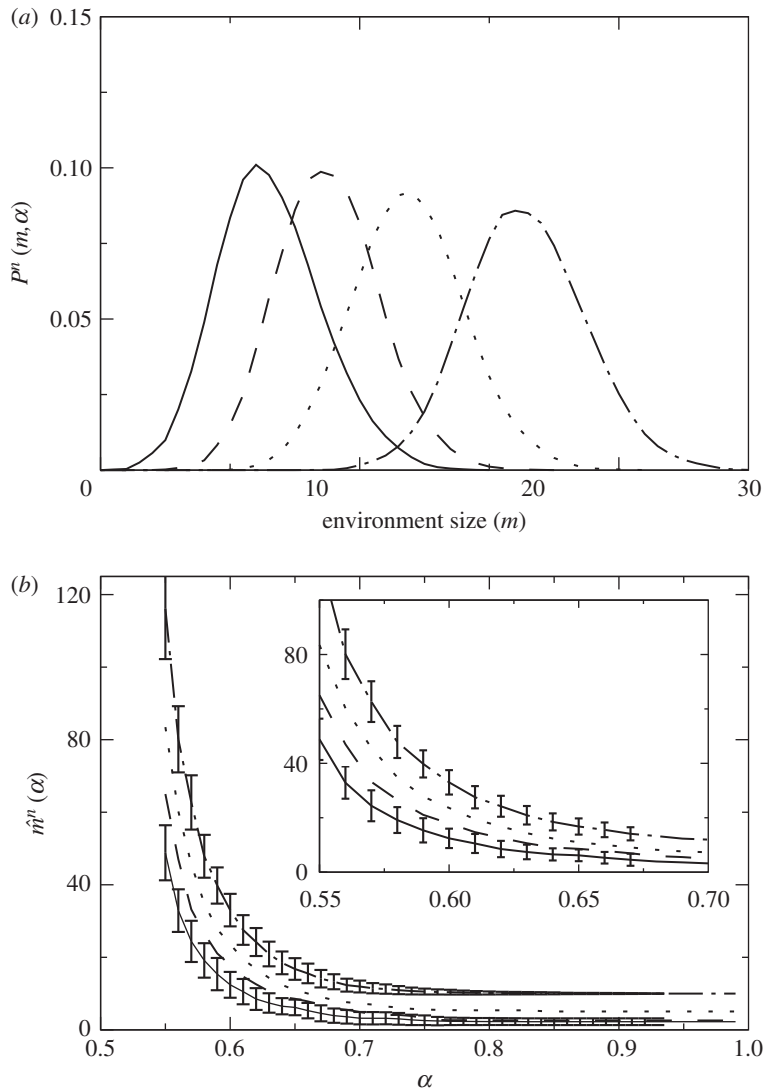


Figure 6. Dynamics around $\bar{m}^n(\alpha)$. Again, guessers with $n = 1$ (solid line), $n = 2$ (dashed line), $n = 5$ (dotted line) and $n = 10$ (dot-dashed line). (a) $P^n(m, \alpha)$ tells us how often we find n -guessers in m -environments when they are allowed to roam constrained only by their survival function ρ_m^n . The central value \hat{m}^n of $P^n(m, \alpha)$ must converge to $\bar{m}^n(\alpha)$ and oscillations around it depend (through N_g and N_e) on how often we evaluate the guessers in each environment. (b) Average \hat{m}^n for $n = 1, 2, 5, 10$ and standard deviation of $P^n(m, \alpha)$ for $n = 1, 10$. Deviations are not presented for $n = 2, 5$ for clarity. The inset represents a zoom-in into the main plot.

At each iteration, a guesser (say the i th one) was chosen randomly and evaluated with respect to its environment. Then the wasted environment was replaced by a new, random one with the same size. We ensured that every guesser attempts to guess the same amount of bits on average. This means e.g. that 1-guessers are tested twice as often as 2-guessers, etc. If after the evaluation we found that $\hat{\rho}_i(t + \Delta t) < 0$, then the guesser died and it was substituted by a new one. The n of the new guesser was chosen randomly after the current distribution $P^n(n, t)$. If $\hat{\rho}_i(t + \Delta t) > 2n\rho_0$, the guesser got replicated and shared its $\hat{\rho}_i$ with its daughter, who overrode another randomly chosen guesser. This replication at $2n\rho_0$ represents that, before creating a similar agent, parents must satisfy a metabolic load that grows with their size. There is a range ($0 < \hat{\rho}_i < 2n\rho_0$) within which guessers are alive but do not replicate.

Of course, this minimal model is just a proxy and softer constraints could be placed. These could allow e.g. for random replication depending on the accumulated $\hat{\rho}_i(t + \Delta t)$, or for larger progeny if $\hat{\rho}_i(t + \Delta t) \gg 2n\rho_0$. These are interesting variations that might be worth exploring. There are also some insights to be gained from the simple set-up considered here. We expect that more complex models will largely inherit the exploratory results that follow.

Figure 7a,b shows $P_m(n, t = 10000)$ with $\alpha = 0.6$ and 0.65 . Note that for large environments all guessers combined do not add up to 100. Indeed, they fall short of that number—i.e. mostly empty

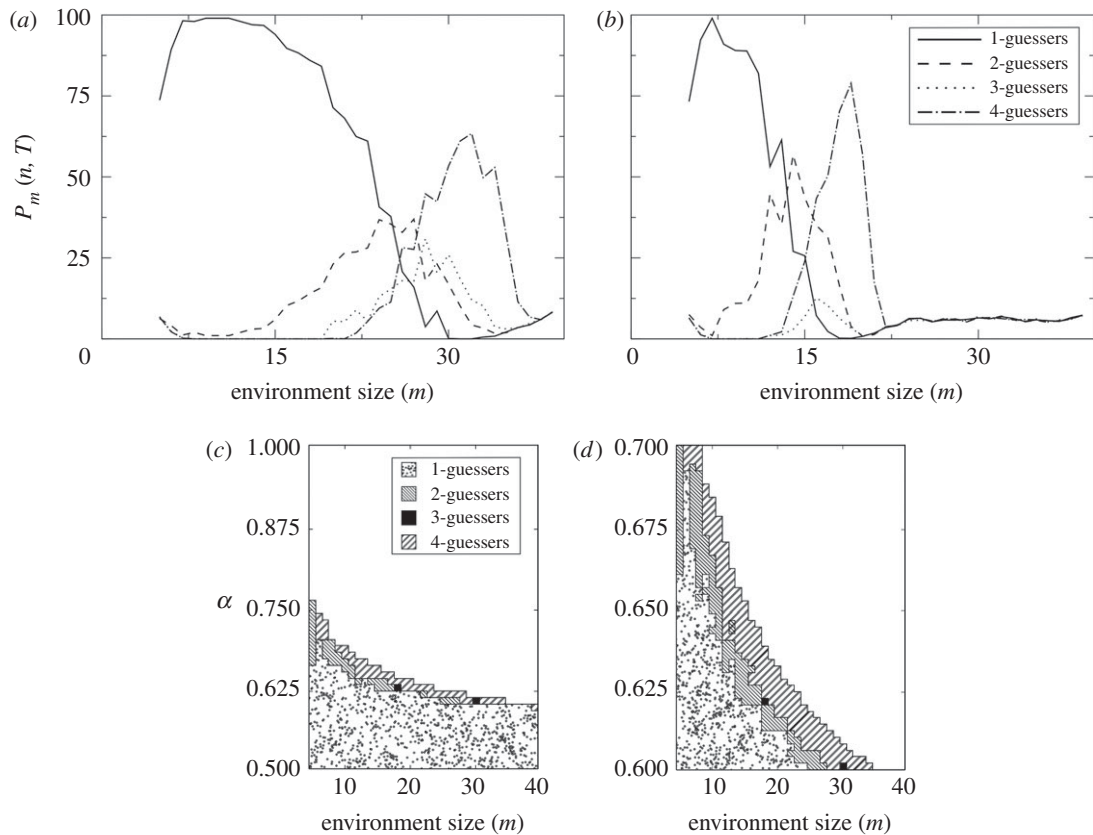


Figure 7. Evolutionary drivers: competition. Coexisting replicators will affect each other's environments in non-trivial ways which may often result in competition. We implement a dynamics in which 1-, 2-, 3- and 4-guessers exclusively occupy a finite number of environments of a given size (fixed m). The 100 available slots are randomly occupied at $t = 0$ and granted to the best replicators as the dynamics proceed. We show $P_m(n, t = 10\,000)$ for $m = 5, \dots, 39$ and $\alpha = 0.6$ (a) and $\alpha = 0.65$ (b). The most abundant guesser at $t = 10\,000$ is shown for $\alpha \in (0.5, 1)$ (c) and $\alpha \in (0.6, 0.7)$ (d). Once m is fixed, there is an upper value of α above which no guesser survives and all 100 available slots remain empty. Competition and the replication-predictability trade-off segregate guessers according to the complexity of the environment—i.e. of the transmission channel. Coexistence of different guessers seems possible (e.g. $m = 15$ in b), but it cannot be guaranteed that the dynamics have converged to a steady distribution.

slots remain. The most abundant guesser after 10 000 iterations is shown in figure 7c as a function of m and α .

These plots show how guessers naturally segregate in environments depending on their complexity, with simpler guessers crowding simpler environments as suggested above. In such simple environments, the extra reward earned by more complex guessers does not suffice to overcome their energetic cost and they lose in this direct competition. They are, hence, pushed to more complex environments where their costly inference machinery pays off.

After 10 000 iterations, we also observe cases in which different guessers coexist. This means that the mathematical limits imposed by this naive model do not imply an immediate, absolute dominance of the fittest guesser. Interesting temporal dynamics might arise and offer the possibility to model complex ecological interactions.

So far, our guessers only interacted with the environment in a passive way, by receiving the reward that the corresponding m -environment dictates. But living systems also shape their niche in return. Such interplay can become very complicated and we think that our model offers a powerful exploratory tool. Let us study a very simple case in which the actions of the guessers (i.e. their correctly guessing a bit or not) affect the reward that an environment can offer.

To do so, we rethink the bits in an environment as resources that can not only be exhausted if correctly guessed, but also replenished after enough time has elapsed. Alternatively, thinking from the message broadcasting perspective, a spot on the channel might appear crowded if it is engaged in a successful transmission. Assume that every time that a bit is correctly guessed it gets exhausted (or gets crowded)

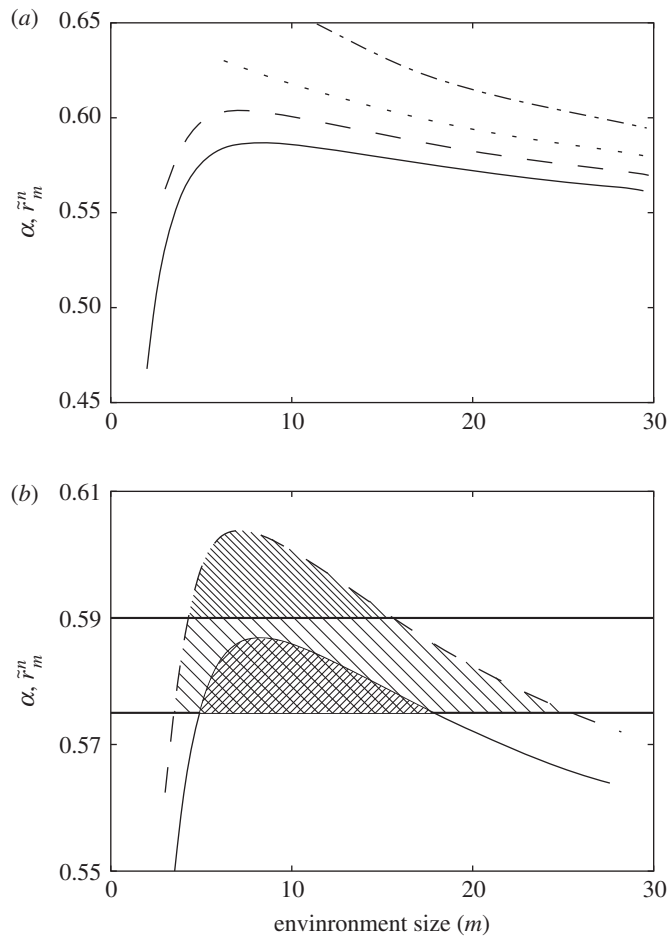


Figure 8. Evolutionary drivers: exhausted resources. Rather than monopolizing channel slots (as in figure 5), we can also conceive individual bits as valuable, finite resources that get exhausted whenever they are correctly *guessed*. Then a successful replicator can spoil its own environment and new conditions might apply to where life is possible. (a) Average reward obtained by 1-, 2-, 5- and 10-guessers in environments of different sizes when bits get exhausted with efficiency $\beta = 1$ whenever they are correctly guessed. (b) Given $\alpha = 0.575$ and $\alpha = 0.59$, 1- and 2-guessers can survive within upper and lower environment sizes. If the environment is too small, resources get consumed quickly and cannot sustain the replicators. In message transmission language, the guessers crowd their own channel. If the environment is too large, unpredictability takes over for these simple replicators and they perish.

with an efficiency β so that on average each bit cannot contribute any reward $\beta(\bar{p}_m^n/m)$ of the time. The average reward extracted by a guesser from an ensemble becomes

$$\bar{r}_m^n = \left(1 - \beta \frac{\bar{p}_m^n}{m}\right) \bar{p}_m^n r, \quad (3.2)$$

which is plotted for 1-, 2-, 5- and 10-guessers and $\beta = 1$ in figure 8.

Smaller guessers living in very small environments quickly crowd their channels (alternatively, exhaust the resources they depend on). In figure 8b (still with $\beta = 1$) given some α , 1- and 2-guessers can only survive within some lower and upper limits. Furthermore, the slope of the curves around these limits also tell us important information. If these guessers dwell in environments around the lower limit (i.e. near the smallest m -environment where they can persist), then moving to larger environments will always report larger rewards. But if they dwell close to the upper limit, moving to larger environments will always be detrimental. In other words, dynamics such as the one introduced at the beginning of this section (illustrated in figure 6a) would have, respectively, unstable and stable fixed points in the upper and lower limits of persistence.

This simple model illustrates how scarcity of resources (and, more generally, other kinds of guesser–environment interactions) might play an important role as evolutionary drivers towards more complex

life. This does not intend to be an exhaustive nor a definitive model, just an illustration of the versatility of the bit-guessers and environments introduced in this paper.

4. Discussion

In this paper, we have considered a fundamental question related to the emergence of complexity in living systems. The problem being addressed here is whether the mathematical conditions exist such that more complex organisms can overcome the cost of their complexity by developing a higher potential to predict the external environment. As suggested by several authors [6–8], the behavioural plasticity provided by the exploratory behaviour of living systems can be understood in terms of their ability to deal with environmental information [59].

Our models make an explicit approach by considering a replication–predictability trade-off under very general assumptions, namely: (i) more complex environments look more unpredictable to simpler replicators and (ii) agents that can keep a larger memory and make inferences based on more elaborated information can extract enough valuable bits from the environment so as to survive in those more challenging situations. Despite the inevitable cost inherent to the cognitive machinery, a selection process towards more complex life is shown to exist. This paves the way for explicit evolutionary pressures towards more complex life forms.

In our study we identify a transmitter (replicators at a given generation), a receiver (replicators at the next generation) and a channel (*any* environmental conditions) through which a message (ideally instructions about how to build newer replicators) is passed on. Darwinian evolution follows naturally as effective replicators transit a channel faster and more reliably thus getting more and more space in successive generations. The inference task is implicit as the environment itself codes for meaningful bits of information that, if picked up by the replicators, boost the fitness of the phenotypes embodied by the successful messages.

This view is directly inspired by a qualitative earlier picture introduced by Maynard-Smith [18]. That metaphor assigned to the DNA some external meaning that had to be preserved *against* the environmental noise. Contrary to this, we propose that, as messages attempt to travel from a generation to the next one, all channel conditions (including noise) pump relevant bits into the transmitted strings—hence there is no need to protect meaning against the channel because, indeed, meaningful information emerges out of the replicator’s interaction with such channel contingencies.

The way that we introduce correlations in our scheme (through a symmetry breaking between the information borne by short and larger words due to finite size effects) is compatible with this view. However, interestingly, it also suggests that meaningful information might arise naturally even in highly unstructured environments when different spatial and temporal scales play a relevant role. Note that our findings imply that environmental complexity is a driver of life complexity, but a question shall remain: ‘where did all that environmental complexity arise from in the first place?’ The way in which we link complexity and environmental size suggests an answer: that real living systems have an option to wander in ever larger environments (which, by sheer size, will be more complex than smaller ones). This is similar in spirit to the simulations illustrated in figure 6. Another possibility is that living systems themselves shall modify their environment.

This way of integrating information theory and Darwinian evolution is convenient to analyse the questions at hand that concern the emergence of complex life forms. But it also suggests further research lines. As discussed at the beginning of the paper, guessers and their transmissible messages might and should shape the transmission channel (e.g. by crowding it, as explored briefly in §3.2). What possible co-evolutionary dynamics between guessers and channels can be established? Are there stable ones, others leading to extinction, etc.? Do some of them, perhaps, imply open-ended evolution? Which ones? These are questions that relate tightly to the phenomenon of niche construction, and that hinge on the question posed in the previous paragraph regarding the many possible origins of environmental complexity. We propose that they can be easily modelled within the bit-guesser paradigm that we introduced in this paper. Further exploring the versatility of the model, a guesser’s transmitted message might be considered an environment in itself; thus opening the door to ecosystem modelling based on bare information theory. It also suggests the exploration of different symbiotic relationships from this perspective and how they might affect coevolution.

An important question was left aside that concerns the memory versus adaptability trade-off of bit-guessers. Here we studied guessers with a minimal adaptability to focus on the emerging hierarchy of complexity. Adaptability at faster (e.g. at behavioural) temporal scales is linked to more complex

inferences with richer dynamics. This brings in new dilemmas as to how to weight the different building blocks of complex inference—e.g. how do we compare memory and *if-else* or *while* instructions? These and other questions are left for exploration in future research.

Finally, it is interesting to contextualize our results along with two recent papers published after this work was concluded. On the one hand, Boyd *et al.* [60] discuss how a system can behave as a *thermodynamic engine* that produces work with environmental correlations as fuel. This is, we believe, a very relevant discussion for the thermodynamic limits of biophysical systems that could bring explicit meaning to our abstract costs and rewards. On the other hand, Marzen & DeDeo [34] study the relationship between an environment and the resources devoted to sensory perception. They use a utilitarian approach to discover two regimes: one in which the cost of sensory perception grows with environmental complexity, and another one in which this cost remains broadly independent of the complexity of the environment. The authors say then that *lossy compression* allows living systems to survive without keeping exhaustive track of all the information in the environment. In these two papers, replication and the pressure of Darwinian selection are not so explicitly discussed as in our research, but both works bring in interesting elements that can enrich the bit-guesser framework.

Data accessibility. This theoretical paper does not rely on empirical data. All the plots can be generated following the equations and instructions provided in the paper.

Authors' contributions. Both authors contributed equally to this work. Both L.F.S. and R.V.S. designed and implemented the theoretical framework, analysed the data, elaborated the figures and wrote the manuscript. Both authors have read and approved the final version of this manuscript.

Competing interests. We declare we have no competing interests.

Funding. This research was funded by the Botin Foundation, by Banco Santander through its Santander Universities Global Division, by the Santa Fe Institute, by the Secretaria d'Universitats i Recerca del Departament d'Economia i Coneixement de la Generalitat de Catalunya and by the European Research Council through ERC grant no. ERC SYNCOM 294294.

Acknowledgements. The authors thank the members of the Complex Systems Lab, and Jeremy Owen, Henry Lin, Jordan Horowitz and David Wolpert for very useful discussions.

References

- Gould SJ. 2011 *Full house*. Harvard, MA: Harvard University Press.
- Bonner JT. 1988 *The evolution of complexity by means of natural selection*. Princeton, NJ: Princeton University Press.
- Maynard-Smith J, Szathmáry E. 1997 *The major transitions in evolution*. Oxford, UK: Oxford University Press.
- Szathmáry E, Maynard-Smith J. 1997 From replicators to reproducers: the first major transitions leading to life. *J. Theor. Biol.* **187**, 555–571. (doi:10.1006/jtbi.1996.0389)
- Jablonka E, Lamb MJ. 2006 The evolution of information in the major transitions. *J. Theor. Biol.* **239**, 236–246. (doi:10.1016/j.jtbi.2005.08.038)
- Jacob F. 1998 *Of flies, mice and man*. Harvard, MA: Harvard University Press.
- Friston K. 2013 Life as we know it. *J. R. Soc. Interface* **10**, 20130475. (doi:10.1098/rsif.2013.0475)
- Wagensek J. 2000 Complexity versus uncertainty: the question of staying alive. *Biol. Phil.* **15**, 493–508. (doi:10.1023/A:1006611022472)
- Trewavas A. 2014 *Plant behaviour and intelligence*. Oxford, UK: Oxford University Press.
- Oborny B. 1994 Growth rules in clonal plants and environmental predictability—a simulation study. *J. Ecol.* **82**, 341–351. (doi:10.2307/2261302)
- Topham AT, Taylor RE, Yan D, Nambara E, Johnston IG, Bassel GW. 2017 Temperature variability is integrated by a spatially embedded decision-making center to break dormancy in *Arabidopsis* seeds. *Proc. Natl Acad. Sci. USA* **114**, 6629–6634. (doi:10.1073/pnas.1704745114)
- Jaynes ET. 1957 Information theory and statistical mechanics. *Phys. Rev.* **106**, 620. (doi:10.1103/PhysRev.106.620)
- Jaynes ET. 1957 Information theory and statistical mechanics. II. *Phys. Rev.* **108**, 171. (doi:10.1103/PhysRev.108.171)
- Parrondo JM, Horowitz JM, Sagawa T. 2015 Thermodynamics of information. *Nat. Phys.* **11**, 131–139. (doi:10.1038/nphys3230)
- Joyce GF. 2002 Molecular evolution: booting up life. *Nature* **420**, 278–279. (doi:10.1038/420278a)
- Joyce GF. 2012 Bit by bit: The Darwinian basis of life. *PLoS Biol.* **10**, e1001323. (doi:10.1371/journal.pbio.1001323)
- Krakauer DC. 2011 Darwinian demons, evolutionary complexity, and information maximization. *Chaos* **21**, 037110. (doi:10.1063/1.3643064)
- Maynard-Smith J. 2000 The concept of information in biology. *Philos. Sci.* **67**, 177–194. (doi:10.1086/392768)
- Nurse P. 2008 Life, logic and information. *Nature* **454**, 424–426. (doi:10.1038/454424a)
- Walker SI, Davies CW. 2012 The algorithmic origins of life. *J. Phys. Soc. Interface* **10**, 20120869. (doi:10.1098/rsif.2012.0869)
- Hilbert M. 2017 Complementary variety: when can cooperation in uncertain environments outperform competitive selection? *Complexity* **2017**, 5052071. (doi:10.1155/2017/5052071)
- Krakauer D, Bertschinger N, Olbrich E, Ay N, Flack JC. 2014 The information theory of individuality. (<http://arxiv.org/abs/1412.2447>)
- Schuster P. 1996 How does complexity arise in evolution? *Complexity* **2**, 22–30. (doi:10.1002/(SICI)1099-0526(199609/10)2:1<22::AID-CPLX>3.0.CO;2-H)
- Adami C. 2012 The use of information theory in evolutionary biology. *Ann. N. Y. Acad. Sci.* **1256**, 49–65. (doi:10.1111/j.1749-6632.2011.06422.x)
- Bergstrom CT, Lachmann M. 2004 Shannon information and biological fitness. In *Information Theory Workshop, 2004*, pp. 50–54. IEEE.
- Dall SR, Giraldeau LA, Olsson O, McNamara JM, Stephens DW. 2005 Information and its use by animals in evolutionary ecology. *Trends Ecol. Evol.* **20**, 187–193. (doi:10.1016/j.tree.2005.01.010)
- Dall SR, Johnstone RA. 2002 Managing uncertainty: information and insurance under the risk of starvation. *Phil. Trans. R. Soc. Lond. B* **357**, 1519–1526. (doi:10.1098/rstb.2002.1061)
- Donaldson-Matasci MC, Bergstrom CT, Lachmann M. 2010 The fitness value of information. *Oikos* **119**, 219–230. (doi:10.1111/j.1600-0706.2009.17781.x)
- Donaldson-Matasci MC, Lachmann M, Bergstrom CT. 2008 Phenotypic diversity as an adaptation to environmental uncertainty. *Evol. Ecol. Res.* **10**, 493–515.
- Evans JC, Votier SC, Dall SR. 2015 Information use in colonial living. *Biol. Rev.* **91**, 658–672. (doi:10.1111/brv.12188)
- Hidalgo J, Grilli J, Suweis S, Muñoz MA, Banavar JR, Maritan A. 2014 Information-based fitness and the emergence of criticality in living systems. *Proc. Natl*

- Acad. Sci. USA* **111**, 10 095–10 100. (doi:10.1073/pnas.1319166111)
32. Kussell E, Leibler S. 2005 Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078. (doi:10.1126/science.1114383)
 33. Marzen SE, DeDeo S. 2016 Weak universality in sensory tradeoffs. *Physical Review E* **94**, 060101. (doi:10.1103/PhysRevE.94.060101)
 34. Marzen SE, DeDeo S. 2017 The evolution of lossy compression. *J. R. Soc. Interface* **14**, 20170166. (doi:10.1098/rsif.2017.0166)
 35. McNamara JM, Houston AI. 1987 Memory and the efficient use of information. *J. Theor. Biol.* **125**, 385–395. (doi:10.1016/S0022-5193(87)80209-6)
 36. Rivoire O, Leibler S. 2011 The value of information for populations in varying environments. *J. Stat. Phys.* **142**, 1124–1166. (doi:10.1007/s10955-011-0166-2)
 37. Sartori P, Granger L, Lee CF, Horowitz JM. 2014 Thermodynamic costs of information processing in sensory adaptation. *PLoS Comput. Biol.* **10**, e1003974. (doi:10.1371/journal.pcbi.1003974)
 38. Segré D, Ben-Eli D, Lancet D. 2000 Compositional genomes: prebiotic information transfer in mutually catalytic noncovalent assemblies. *Proc. Natl Acad. Sci. USA* **97**, 4112–4117. (doi:10.1073/pnas.97.8.4112)
 39. Segré D, Shenhav B, Kafri R, Lancet D. 2001 The molecular roots of compositional inheritance. *J. Theor. Biol.* **213**, 481–491. (doi:10.1006/jtbi.2001.2440)
 40. Szathmáry E. 1989 The integration of the earliest genetic information. *Trends Ecol. Evol.* **4**, 200–204. (doi:10.1016/0169-5347(89)90073-6)
 41. Tkačik G., Bialek W. 2014 Information processing in living systems. (<http://arxiv.org/abs/1412.8752>)
 42. Drossel B. 2001 Biological evolution and statistical physics. *Adv. Phys.* **50**, 209–295. (doi:10.1080/00018730110041365)
 43. England JL. 2013 Statistical physics of self-replication. *J. Chem. Phys.* **139**, 121923. (doi:10.1063/1.4818538)
 44. Goldenfeld N, Woese C. 2010 Life is physics: evolution as a collective phenomenon far from equilibrium. *Annu. Rev. Condens. Matter Phys.* **2**, 375–399. (doi:10.1146/annurev-commatphys-062910-140509)
 45. Nicolis G, Prigogine I. 1977 *Self-organization in nonequilibrium systems*. New York, NY: Wiley.
 46. Perunov N, Marsland R, England J. 2014 Statistical physics of adaptation. (<http://arxiv.org/abs/1412.1875>)
 47. Shannon CE. 2001 A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423. (doi:10.1002/j.1538-7305.1948.tb01338.x)
 48. Shannon CE, Weaver W. 1949 *The Mathematical Theory of Communication*. Champaign, IL: University of Illinois Press.
 49. Corominas-Murtra B, Fortuny J, Solé RV. 2014 Towards a mathematical theory of meaningful communication. *Sci. Rep.* **4**, 4587. (doi:10.1038/srep04587)
 50. von Neumann J, Burks AW. 1966 Theory of self-reproducing automata. *IEEE Trans. Neural Netw.* **5**, 3–14.
 51. Hopfield JJ. 1988 Artificial neural networks. *IEEE Circuits Devices Mag.* **4**, 3–10. (doi:10.1109/101.8118)
 52. Maass W, Bishop CM. 2001 *Pulsed neural networks*. Cambridge, MA: MIT Press.
 53. Jensen FV. 1996 *An introduction to Bayesian networks*. London, UK: UCL Press.
 54. Pearl J. 1985 *Bayesian networks: a model of self-activated memory for evidential reasoning*. University of California (Los Angeles). Computer Science Department.
 55. Turing AM. 1936 On computable numbers, with an application to the Entscheidungsproblem. *J. Math.* **58**, 5.
 56. Markov A. 1971 Extension of the limit theorems of probability theory to a sum of variables connected in a chain. In *Dynamic Probabilistic Systems, Volume 1: Markov Chains* (ed. R. Howard), pp. 552–577. New York, NY: John Wiley and Sons.
 57. Crutchfield JP, Young K. 1989 Inferring statistical complexity. *Phys. Rev. Lett.* **63**, 105. (doi:10.1103/PhysRevLett.63.105)
 58. Kauffman SA. 1993 *The origins of order: self organization and selection in evolution*. Oxford, UK: Oxford University Press.
 59. Gerhart J, Kirschner M. 1997 *Cells, embryos, and evolution*. Malden, MA: Blackwell Science.
 60. Boyd AB, Mandal D, Crutchfield JP. 2017 Correlation-powered information engines and the thermodynamics of self-correction. *Phys. Rev. E* **95**, 012152. (doi:10.1103/PhysRevE.95.012152)