

Universal Darwinism and the Origins of Order

John O. Campbell
Victoria, British Columbia
jockocampbell@gmail.com

Michael E. Price
Department of Life Sciences
Brunel University London
michael.price@brunel.ac.uk

Forthcoming in *Evolution, Development and Complexity: Multiscale Evolutionary Models of Complex Adaptive Systems*, edited by G. Georgiev, C. L. F. Martinez, M. E. Price, & J. Smart. Springer Publishing.

Working draft, not a final version
Online version 1.0

Abstract

In this chapter we describe a ‘universal Darwinism’ framework which proposes the following. The observable universe results from two types of processes: (1) disorder’s tendency to increase in isolated systems (the second law of thermodynamics), and (2) Darwinian selection, which produces orderly entities that can withstand the second law. Darwinian processes generate complex order not just in the biological domain but in all five domains of nature. These domains exist in a nested hierarchy as follows (in order of decreasing fundamentalness): cosmological, quantum, biological, neural, and cultural. Each qualifies as a distinct domain because each is characterized by a distinct ‘knowledge repository’, that is, a cumulative store of information about existence requirements in that domain (e.g. in the biological domain, a genome). Knowledge repositories are probabilistic models which make guesses about how to exist, guesses which are then tested for accuracy by the ‘embodied adapted system’ (e.g. phenotype) encoded by the knowledge repository. The repository then undergoes a Bayesian update based on test results, and thus becomes less ignorant and less entropic. These natural inferential systems evolve according to ‘variance – inheritance – selection’ Darwinian dynamics, with wiser knowledge repositories leaving more copies behind. Each new domain’s knowledge repository computationally transforms the substrate of the earlier domain (e.g., cultural repositories orchestrate the neural substrate), to generate innovative ways of overcoming the second law in the new domain. We conclude that Darwinian theory, as an explanation for the origins of complex order in the universe, may be far more fundamental than is conventionally supposed.

1. Introduction: The two processes that shape the universe

In this chapter we will propose that features of the observable universe result from two general processes. One is the entropic process of ever-increasing disorder in isolated systems, as described by the second law of thermodynamics. The other is the selection process of Darwinian evolution, which produces orderly entities that are able to exist for extended periods of time, despite the chronic dissipative threat of the entropic process. We will suggest that the selection process accounts for the origins of order not just biologically but in all scientifically observable domains, from the cosmological and quantum to the biological and cultural. The selection process favours the evolution of order in general because, across domains: (1) it favours entities which are relatively capable at solving problems of existence, and at passing their problem-solving advantages on to future generations (e.g. via survival and reproduction); (2) solving problems of existence, against the powerful current of the second law, requires problem-solving mechanisms that are highly-ordered. The ‘universal Darwinism’ framework we are proposing is currently unconventional, and we acknowledge that it is highly speculative. But we also see it as an unusually compelling framework, particularly because of its high parsimony, explanatory power, and consilience with other bodies of scientific knowledge.

Our main goal in this chapter is to demonstrate how Darwinian selection could provide a unifying framework for natural domains that, superficially, may seem highly disparate. Universal Darwinism could enable such unification because the selection process seems to involve the same key components, in no matter what domain it operates. We have identified five components as particularly important, and will describe them in detail later in the chapter: (1) a knowledge repository, (2) an embodied adapted system, (3) a process of local entropy reduction, (4) a computational inferential system, and (5) a system of variance, inheritance, and selection. To make it clear how these five key components could enable the evolution of order in a variety of natural domains, we will structure this chapter as follows. In Part 2 below, we will provide an overview of universal Darwinism, outlining the broad contours of this framework’s applicability across domains. In Part 3 we will focus more closely on each of the natural domains to which this framework applies. We identify five such domains, existing in a nested hierarchy as follows (in order of decreasing fundamentalness): cosmological, quantum, biological, neural, and cultural. In Part 4 we will summarize this universal Darwinism framework.

2. An overview of universal Darwinism

A reasonable way to approach universal Darwinism initially is to consider the following fact: throughout the entire scope of scientific subjects, Darwinian theories have been used to describe the fundamental processes by which subject-specific phenomena come into existence and subsequently evolve. We will describe many of these theories in more detail below, and for now will just note that well-developed Darwinian theories have flourished in scientific fields as diverse as cosmology (Smolin 1992, 1997), quantum physics (Zurek 2009), biology (including evolutionary psychology) (Darwin 1859; Tooby & Cosmides 1992), neuroscience (Fernando, Szathmary & Husbands 2012; Friston 2013), anthropology and archaeology (O’Brien & Lyman 2003; Richerson & Boyd 2005; Gibson &

Lawson 2014), economics (R. Frank 2012; Nelson & Winter 1982), political science (Pietraszewski et al. 2015; Thayer 2004), organizational behaviour and the business sciences (Arvey & Colarelli 2015; Saad 2011), and epistemology (Campbell 1974; Popper 1972). The ubiquitous power of the Darwinian framework can lead to a deeper understanding of the universe, if we attempt to answer the fundamental questions this power implies. These questions include ‘Why does nature use this single process across so many different domains?’, and ‘What common features are shared by Darwinian processes, in all these domains?’.

We will propose tentative answers to these questions, answers which suggest a fundamental unification in terms of common evolutionary processes operating across many natural domains. At their core, these processes all involve the familiar ‘variance – inheritance – selection’ dynamics that are required of any Darwinian system (Dennett 1995): a population of entities must come in various forms, some forms must be more advantageous than others for overcoming problems of existence, and these advantages must be heritable and selectively retained. However, we will flesh out this basic ‘selective retention of better-adapted forms’ Darwinian model, and present it in terms of an evolved inferential system which learns, implements, and retains *knowledge* about strategies for existence.

2.1 The accumulated knowledge repository as a central concept in universal Darwinism

What does it take to achieve persistent existence, in a universe dominated by the relentlessly destructive power of the second law of thermodynamics? Universal Darwinism proposes that existence requires strategies that ‘know’ how to exploit loop-holes in the second law. These strategies produce entities capable of converting lower-entropy energy into higher-entropy waste, in order to maintain their own stable, complex, low-entropy structure (e.g. in biological organisms, homeostasis). In this view, Darwinian selection is the strongest known antidote to the dissipative tendencies of the second law; no other natural process can so capably produce local states of complex existence.

Central to our argument is the contention that existence is a rare and fragile state. Many complex systems, from those described by the fundamental laws and parameters of physics to those existing in the biological and cultural domains, display a remarkable degree of ‘fine-tuning’ for existence. This fine-tuning may be explained as the result of well-developed and highly specialized knowledge repositories. David Deutsch (2011), one of the founders of quantum computation, describes the centrality of knowledge in a general and optimistic light: “everything that is not forbidden by laws of nature is achievable, given the right knowledge”. In this chapter we will describe several natural processes of knowledge accumulation, which serve to make existence achievable.

This concept of knowledge is a central aspect of universal Darwinism’s ability to achieve a deeper unification of scientific domains (Campbell 2015, 2016). In order to understand this centrality, it helps to broaden our interpretation of Darwinian ‘fitness’. In evolutionary biology, fitness is conventionally understood in terms of genetic replication, achieved via survival, reproduction, and kin selection (Hamilton 1964). We do not take issue with this interpretation at the biological domain specifically, but we do propose that across domains, fitness can be more broadly understood in terms of the accumulation of

knowledgeable strategies for achieving existence. As we explain below, by generalizing from replicating genes to accumulated knowledge, we are able to describe universal Darwinism in terms of information theory and the mathematics of Bayesian inference. This descriptive framework is particularly compelling, because it seems to apply equally well across a wide variety of natural domains.

We suggest that nature has utilized the Darwinian process – the selective retention of better-adapted forms – to accumulate repositories of evolutionary knowledge within several domains of nature. This knowledge is autopoietic (i.e., it enables systems which can maintain and reproduce themselves), and capable of bringing complex forms into existence. We will propose and discuss five types of autopoietic knowledge repositories, one in each domain of nature: (1) in the cosmological domain, this knowledge repository is represented by the laws and parameters of physics; (2) in the quantum domain, by quantum wave functions; (3) in the biological domain, by genomes (potentially acting in concert with epigenetic effects); (4) in the neural domain, by learned neural models; and (5) in the cultural domain, by cultural models. It is no accident that each natural domain possesses its own repository of accumulated knowledge, because the presence of such a knowledge repository is in fact our criterion for identifying each domain. Our strategy for ‘carving nature at the joints’, in other words, is to grant the status of ‘natural domain’ only to those regions of nature that are characterized by their own accumulated knowledge repository.

Although there is a resemblance between our hierarchy of domains and Auguste Comte's (1848) well-known hierarchy of scientific fields, Comte inserts chemistry between the levels of ‘physics’ (labelled as ‘quantum’ in our hierarchy) and ‘biology’. Why do we not regard chemistry as a separate domain? For reasons discussed below, we see chemistry as part of the cosmological and quantum domains: chemistry is orchestrated according to quantum principles, and arises from the laws and parameters of physics. Chemistry is the outcome of cosmological and quantum accumulated knowledge repositories, in other words, and as it possesses no such repository of its own, we do not consider it a discrete domain.

2.2 Knowledge repositories generate embodied adapted systems

A vital observation about knowledge repositories is that they instantiate, via developmental processes, embodied adapted systems (e.g. in the biological domain, a genome instantiates a phenotype). There is a duality, in other words, between the knowledge repository and its embodied adapted system, and we will attempt to identify this duality in each of nature’s domains. Nature’s knowledge repositories tend to be relatively timeless and are copied, through the process of heredity, in a largely conserved and cumulative manner. The adapted systems they instantiate, on the other hand, tend to be short-lived mortal entities; they function to test the soundness of the knowledge which brought them into being, and to update this knowledge when appropriate. If a new phenotypic design encoded by a genetic mutation succeeds at survival and reproduction, for example, then the genome will have passed this test: the mutation will have proven itself to constitute accurate knowledge about how to exist, and more copies of it will now exist in the population.

Because an embodied adapted system functions to test knowledge about how to solve problems of existence, it must be composed of at least one mechanism that attempts to solve some such problem – that is, at least one *adaptation*. An adapted system may be composed of

many more than just one adaptation, however, and some of these systems are best characterised as integrated bundles of adaptations. The most familiar example would be a complex biological organism, composed of a vast number of functionally-specialized adaptations, each ‘designed’ by selection to solve a specific problem (or set of problems) related to survival and reproduction. An embodied adapted system may be composed of a diverse array of problem-solving devices, because achieving existence can itself entail a diverse array of problems.

Given that embodied adaptive systems evolve to solve problems of existence, the fundamental ‘design principle’ of any such system will be to be organized into one or more problem-solving adaptations. However, it will not always be appropriate to regard all aspects of an adapted system as adaptations. These aspects may be (1) adaptations, (2) by-products of adaptations, or (3) random noise (Williams 1966; Tooby & Cosmides 1992; Buss et al. 1998). Adaptations can be considered to be any traits that have been selected to solve some problem(s) of existence (note that by this definition, the category of ‘adaptations’ subsumes that of so-called ‘exaptations’ [Gould & Vrba 1982; see also Price, this volume]). A by-product is a trait that is incidentally coupled with some adaptation, but that itself serves no adaptive function. For example, while nipples in female mammals can be considered adaptations for feeding offspring, male nipples seem best explained as by-products of female nipples. Noise represents stochastic, selectively neutral residual effects of the evolutionary process; because noise has no bearing on the adapted system’s ability to exist, it is invisible to selection. An example would be a nipple’s precise size and shape (e.g., diameter in nanometres).

In general, adaptations can be distinguished from by-products and noise by their improbable complexity, that is, their exceptionally low entropy (Price 2017). Adaptations are the aspects of an adapted system that would be least likely to arise by chance (i.e., as the result of stochastic effects of the evolutionary process), in the absence of a non-random process of selection. A trait’s improbability is often best assessed in terms of its functional complexity, that is, the evidence it displays of ‘special design’ (Williams 1966): the more well-engineered a trait seems for solving a specific adaptive problem – like a key that has been precisely designed to fit a specific lock – the less likely it is to have arisen by chance, and the more likely it is to be an adaptation. That is not to say that selection will tend to favour a more complex form of an adaptation over a less complex form, if both forms fulfil the same function equally well. On the contrary, the simpler form would tend to be favoured, as the more efficient solution. However, more complex forms often do represent superior solutions, and are therefore often selected (Heylighen 1999; Orr 2000; Price, this volume). Evolutionary processes can entail stochastic elements (e.g., genetic drift), but selection is the only aspect of evolution that is expected to produce improbable complexity, and that can produce adaptations. Improbable (functional) complexity is therefore the hallmark of selection, and the characteristic that best distinguishes adaptations from by-products and noise (Williams 1966; Tooby & Cosmides 1992; Price 2017).

2.3 Darwinian knowledge repositories represent probabilistic models, which reduce their own entropy via learning

How do evolved knowledge repositories learn adaptively? These repositories can be

thought of as probabilistic models that attempt ‘guesses’ about how to successfully exist in the world. They experiment, that is, with different strategies for existence. These models possess not just the property of knowledge but also its inverse, ignorance (that is, uncertainty). And ignorance, as the ‘father of information theory’ Claude Shannon (1948) pointed out, is mathematically equivalent to entropy. This idea was developed further by E. T. Jaynes, who demonstrated the equivalence of informational and thermodynamic entropy.

Jaynes (1957, 1965) noted that the models used in statistical mechanics to derive thermodynamic entropy are fairly ignorant: no matter how much statistical evidence is measured regarding a thermodynamic system, the best scientific models remain vastly ignorant of the system’s exact microstate. That’s because the evidence supporting these models are measures of macrostates, such as density or temperature, based on average properties of molecules in the system. The density of a volume of gas, for example, is the average number of molecules in the volume, and the temperature is the molecules’ average kinetic energy. So a model that knows such macrovariables has only a small amount of statistical knowledge about the system’s exact microstate. Thermodynamic entropy is a measure of that ignorance, and it is equivalent to information entropy (except that when information theory is applied to thermodynamic entropy, a conversion factor called the Boltzmann constant is used). In this sense entropy can be thought of as the number of bits of information required to move the model from its current state of uncertainty to a state of complete certainty, where it would exactly describe the complete microstate. Entropy is a measure, in other words, of a model’s current state of ignorance.

The relationship between knowledge and entropy may be clarified further if we define knowledge as 2^{-S} , where S is entropy in bits. This definition implies that knowledge increases when entropy/ignorance decreases, and it represents the probability that a model will make a correct guess about the world. For instance, if the entropy of a probabilistic model is 4 bits, then by this definition its knowledge is 2^{-4} , which is equal to $1/16$. The model’s entropy, in other words, is equivalent to that of a uniform distribution having 16 members. This is true regardless of the actual characteristics of the particular distribution, including its average and variance. (For example, if we asked you to guess what number we were imagining from 1 to 16, and you had no reason to favor one choice over another, you would have a $1/16^{\text{th}}$ chance of guessing correctly). If the model then encounters some new information about the nature of this distribution, and this update reduces its entropy to 3 bits, then it can be said to have increased its knowledge to $1/8$. The model’s entropy is now equivalent to a uniform distribution having 8 members. (Continuing with the above example, if we informed you that the number between 1 and 16 was an even number, you’d then have a $1/8^{\text{th}}$ chance of guessing correctly). The model’s knowledge ($1/8$) can be considered the probability of success of a ‘random’ guess, where there are 8 possibilities and the model has no evidence favoring any one possibility over another.

The inverse relationship between knowledge and entropy is a vital aspect of universal Darwinism. This relationship can be understood in terms of a fundamental physical principle, related to the second law of thermodynamics, known as the ‘principle of maximum entropy’: systems will move to the highest state of entropy allowed by the constraints acting upon them (Jaynes 2003). Thus, we may understand a low-entropy dynamical system in terms of the constraints it imposes on the proliferation of disorder. Throughout nature’s domains, a commonly employed tactic is to utilize knowledge repositories to impose such constraints,

and to thus enable the existence of low-entropy systems. A knowledge repository – acting in concordance with an embodied adaptive system – operates to rule out some possible states that would otherwise be included in its own evolutionary trajectory, and thus reduces its own entropy. For example, in the biological domain, some genetic knowledge repositories code for proteins in the form of enzymes. Enzymes serve to allow only specific chemical pathways, ones which would be extremely unlikely to exist in the absence of the enzymes. In this manner the knowledge repository imposes constraints upon the chemical pathways making up the organism, thus allowing itself (and the organism it encodes) to exist as a low-entropy system.

2.4 Knowledge repositories learn via Bayesian updates

A model's ignorance is reduced when evidence is used to update its assigned probabilities in a Bayesian manner. The Bayesian update is the mathematical mechanism which specifies how the prior probabilities composing a model are updated to posterior probabilities; it occurs through the application of a likelihood function, which quantifies how well each possibility is supported by the available evidence. In this view, knowledge is relieved of any anthropomorphic taint: knowledge is not a property just of human minds containing probabilistic models, but of probabilistic models in general. Humans gain knowledge in the same manner as other evolved entities: by employing evidence to update their models.

Universal Darwinism focuses on a specific kind of probabilistic model: knowledge repositories that must guess which strategies will work best for solving problems of existence. These models assign a probability to each possible strategy, based on evidence that the strategy itself merely exists; this probability is, in other words, the frequency of that strategy's existence. Thus, all possibilities of the model that have non-zero probabilities achieve some degree of existence, and the sum of their probabilities equals 1, forming a probability distribution. The relative frequency of each possible strategy may change between generations or iterations, and these changes constitute a Bayesian update of the model (Campbell 2016). In this manner we may understand Darwinian selection as a process which favours the ability to achieve existence. Those traits or entities which can best achieve existence, in other words, are the possibilities of the model that are most frequently selected. To illustrate, imagine that different members of a species have slightly different possible genomes (i.e., biological knowledge repositories), due to random mutations that effect each organism's chance of reproducing. These genomes can be thought of as probabilistic models that test out different reproductive strategies; they make 'guesses' about which strategies will most likely succeed. The genomes that make better guesses will come to predominate in the population, and the species' genome will thus be updated based on new information about the effectiveness of different strategies.

Darwinian selection may be best understood, in a statistical sense, as the change in the frequency of traits in a population between generations (S. Frank 2012). As traits are implemented from instructions contained in the model, we can understand the model as containing one possibility for each existing trait, with the probability assigned to each possibility being equal to its frequency in existence. As some traits become extinct and exit existence, others are introduced, often through random perturbations of the model (such as genetic mutation). These perturbations introduce new variations into the Darwinian selection

process, which are then themselves tested for their ability to achieve existence, and assigned a probability equal to their frequency of existence within the next generation. Heredity (or copying) is inherent to this approach, in the form of the transition from the prior to the posterior distribution of possibilities. Posterior distributions tend to be similar to prior ones, with most strategies described by the model being inherited by the next generation. The extent and type of any changes to the posterior distribution are the result of encounters with surprising new evidence (e.g. novel selection pressures).

2.5 Computation

Computation, another key component of universal Darwinism, essentially involves the application of logic (Spencer-Brown 1979), and logic statements can be built up from just a couple of atomic operations (Bishop 1967). It has been demonstrated that any two-state system – such as ‘true/false’, ‘on/off’, or ‘negative/positive’ – can be used to form all possible logical operations (Spencer-Brown 1979). All of mathematics can be constructed from logical statements, as can universal Turing machines (computers). Universal computers, when running properly constructed algorithms (which are themselves built from atomic logical operations), may output any pattern of applied mathematics that describes a natural process.

A relatively recent computer technique, called evolutionary computation (De Jong 2006), uses the Darwinian selection process to solve problems of existence. Solutions may be selected according to a programmer-provided fitness function, but the programmer need not specify the detailed steps by which a Darwinian algorithm can find these solutions. The programmer instead just provides an initial algorithm, and the fitness function which evaluates the optimality of solutions. The program then randomly varies the algorithm, tests the output of these variations against the specifications of the fitness function, and selects the variations that best succeed. It then makes small variations in some of the selected algorithms, and repeats the process until a solution is found that sufficiently satisfies the fitness function.

We could view the evolutionary computation process described above as an artificial execution of Darwinian selection. Equally, however, we could consider natural executions of Darwinian selection as forms of computation. This latter perspective may help explain the pervasiveness of Darwinian selection in nature. A computational substrate lies at the foundations of physical reality: all quantum dynamics may be precisely coded in an algorithm written in qubits (Deutsch 1985; Lloyd 2013), so physical reality at the quantum level is essentially computational. We will attempt to demonstrate below that each natural domain may be considered in terms of a knowledge repository that functions as a complex algorithm, and that each domain-specific algorithm orchestrates the dynamics of its substrate (i.e., of material from lower-lying domains). This orchestration produces new, mathematically describable complex forms that are capable of existence. As this computational process outputs mathematical patterns, it may go some way towards explaining what Wigner (1960) described as the ‘unreasonable effectiveness of mathematics in the natural sciences’.

Given that computation may arise from any simple two-state substrate, it is perhaps unremarkable that the substrates of many natural domains are so amenable to computation. Indeed it has been demonstrated that computation may be performed using the substrates of

qubits (Deutsch 1985), chemistry (McGregor al. 2012), DNA (Ignatova, Marinez-Perez & Zimmermann 2008), and neurons (Zylberberg et al. 2011). A more interesting question might be: ‘Why does nature exploit this potential by employing computation across its many domains?’ We suggest that the answer provided by Lloyd (2013), in reference to the quantum domain, may be generalized to other domains as well: “The reason is that many complex, ordered structures can be produced from short computer programs, albeit after lengthy calculations”. In other words, it may be much more efficient for evolutionary processes – which are in the business of discovering complex, ordered structures that are capable of existence – to search for algorithms that can produce such structures, instead of searching for the structures directly.

2.6 A system of variance, inheritance, and selection

Now that we have outlined the key components of this relatively novel universal Darwinism framework, we wish to clarify how this framework can be understood in more traditional Darwinian terms. It is commonly observed that for Darwinian selection to occur in general, a population of entities must be subject to a system of variance, inheritance, and selection (Dennett 1995). That is, entities must vary in form, some forms must be better than others at solving problems of existence, and these solutions must be heritable and selectively retained.

We propose that in the context of universal Darwinism, these ‘variance – inheritance – selection’ dynamics unfold as follows. Each natural domain is orchestrated by a heritable knowledge repository: a probabilistic model that contains knowledge about possible autopoietic strategies, or algorithms, for achieving existence in that domain. No copying process is exact, and variations introduced in the hereditary process create slight variations in these algorithms. The execution of these algorithms, which occurs through developmental processes, produces embodied adaptive systems which then test the algorithms for their ability to promote existence. The evidence produced by these tests is used to select the algorithms that succeed, and this selection process in turn updates the probabilistic model (i.e., the knowledge repository) in a Bayesian manner. From this perspective, we see how the general ‘variance – inheritance – selection’ model of the evolutionary process may be incorporated into the specific universal Darwinism framework outlined above.

3. The five domains of universal Darwinism

In the overview of universal Darwinism presented above, we describe how the key components of this framework operate to generate complex order across natural domains. We will now zoom in on each domain, and see how these components function in these more specific contexts. But we will first provide a brief justification for the method by which we have drawn boundaries between these domains.

It is obvious that a good deal of complexity has evolved in the universe since the Big Bang. The growth of this complexity has been cumulative, with new, additional forms of complexity repeatedly emerging from substrates provided by older forms of complexity. Within this hierarchy of progressive complexity, it is a challenge to identify clear conceptual boundaries between the various forms or domains of complexity that have evolved. We believe a strong case can be made, however, for proposing the following five natural

domains: cosmological, quantum, biological, neural, and cultural. Our criterion for regarding each of these domains as distinct is that each seems to possess its own knowledge repository, which has evolved in order to orchestrate newer complex structures from older complex substrates. In the cosmological domain, the knowledge repository is represented by the laws and parameters of physics; in the quantum domain, by the quantum wave function; in the biological domain, by the genome (potentially in concert with epigenetic effects); in the neural domain, by learned mental models; and in the cultural domain, by cultural models. The natural domains are organised in a nested pattern; for example, culture is composed of a neural substrate, which is composed of a biological substrate, and so on. The boundaries between these domains arise naturally, whenever a new form of knowledge repository emerges that can orchestrate existing substrates into new forms of complex existence. The reason we have decided not to consider chemistry as a separate domain (in contrast to Comte's scientific hierarchy, as noted above) is because there does not seem to be a knowledge repository that can be seen as chemistry-specific, and distinct from those of quantum physics and biology.

In the manner described above, nature can be viewed as a nested hierarchy of domains. We will now consider how the universal Darwinism framework can be applied to each domain in turn, ordered from most to least fundamental.

3.1 Cosmological selection

Our use of 'cosmological' is perhaps slightly unusual, in that we mean the physical processes already in place at the beginnings of the universe. These processes, in the form of the laws and parameters of physics, gave rise to what is usually thought of as the cosmos. In this sense the cosmological domain is fundamental in providing the common substrate, and the rules of nature, from which all subsequent domains emerge.

One of the most basic cosmological paradoxes is the 'Goldilocks principle' (sometimes called the 'fine tuning problem'): the proposition that the conditions that allow complexity to exist in the universe can occur only when approximately 32 universal dimensionless constants of physical theory lie within a very narrow range. If any fundamental constants were even only slightly different, the universe would not be conducive to the establishment and development of matter, astronomical structures, chemical diversity, or life. A question therefore arises: 'Why should the universe's physical laws and parameters have the precise values necessary to develop complexity?'

Cosmological natural selection, with and without intelligence

One scientific answer to this puzzle that has received fairly widespread attention is Smolin's theory of cosmological natural selection, or CNS (Smolin 1992, 1997, 2007). Smolin's CNS assumes the perspective, which is by now fairly conventional among cosmologists, that our universe is just one of many in an extremely large population of universes – a multiverse – in which new universes are being constantly generated. Smolin adds to this perspective the idea that in such a multiverse, universe designs that were better at reproducing themselves would become better-represented. In other words, Smolin's CNS is modelled on Darwin's theory of natural selection, and it explains the universe's finely-tuned aspects as evolved adaptations for reproduction. These adaptations cause universes to develop

black holes, from which emerge offspring universes that are separate and inaccessible from their parent universe (Gardner & Conlon 2013). In short, this theory explains the fine-tuned laws and parameters of our universe as a knowledge repository. This repository produces adaptations in the form of complex structures, such as chemistry and stars, that are necessary to produce black holes. This theory embodies the Darwinian paradigm of variance, inheritance and selection as it posits that child universes, born in the process of black hole formation, inherit slightly varied physical laws and parameters from their parent. Over many generations, this process selects for universe designs which prove most prolific at black hole production. As the laws of physics have become extremely precise, allowing little latitude for deviations, they impose tight constraints on the entropy of the physical processes that lead to universe reproduction. We thus may understand cosmological selection as the primary anti-entropic example of Darwinian selection in the nested hierarchy of domains.

The major theoretical variant of Smolin's CNS is cosmological natural selection with intelligence, or CNSI (Crane 1994/2010; Harrison 1995; Gardner 2000; Smart 2009; Vidal 2014; Price 2017), which proposes that intelligent life is an adaptation for universe reproduction. CNSI expects that in the future, intelligent life will acquire the technical expertise that would allow it to create new universes, which would replicate the laws and parameters of its natal universe (to which it would already be adapted). That is not to suggest that black holes could not also enable universe reproduction; both they and intelligence could fulfil this function. However, intelligence may represent the more recently-evolved of these adaptations, and could perhaps prove superior to black holes at the task of high-fidelity universe reproduction (Price 2017, this volume).

The perspective of CNSI fits nicely within the context of universal Darwinism: intelligence is the ability to discover knowledge, and as noted above, anything allowed by the laws of nature – including the creation of new universes – is achievable, given sufficient knowledge (Deutsch 1997). As knowledge repositories are themselves physical structures, both they and the intelligence which produces them may be considered fundamental aspects of evolution. In line with this perspective, it has been suggested that intelligent life is the most complexly improbable aspect of the universe's embodied adapted system, or 'phenotype'. Therefore, according to the adaptationist logic spelled out in section 2.2 above, intelligent life represents the aspect of this phenotype that is most likely to be an adaptation, as opposed to a by-product or noise (Price 2017, this volume). From this perspective, human biocultural evolution would represent a developmental subroutine of cosmological evolution, targeted towards production of a species that would be intelligent enough to reproduce its own universe.

The cosmological knowledge repository

It remains an open question which of the two models of cosmological evolution presented above – Smolin's CNS, or CNSI – is more plausible. Regardless, the information that allows cosmological evolution to happen must have a physical representation. We therefore might ask: 'Where and in what form is the cosmological knowledge repository physically located?' There is no consensus yet about the specifics of the answer, but there is about its general outline. The holographic principle (Stephens, 't Hooft & Whiting 1993) has been widely adopted, particularly the AdS/CFT correspondence (Maldacena 1998), which stipulates a duality between an informational model and physical processes. The model may

be considered to exist on the boundary of physical space, and is an exact dual of the physics which occur within that space. An emerging consensus is that this informational model is composed of quantum entanglements (Orus 2014; Verlinde 2010). In particular, it has been demonstrated that the gravitational field equations emerge from this entanglement (Swingle & Van Raamsdonk 2014). This view fits well within universal Darwinism, as it identifies the knowledge repository which undergoes the Darwinian process of variance, heredity and selection. Some researchers have described this knowledge repository in biological terms, even going so far as to call it the ‘DNA’ of physical reality (Orus 2014). In this view, the actual physics of the universe arose through a developmental process specified by the knowledge repository, and the physical structure of the universe may thus be considered an embodied adapted system in the domain of cosmological selection.

The laws of nature support universal computation, that is, computational processes capable of generating the entire range of patterns or states that are mathematically describable. This basic computational substrate of the laws of nature, which itself may be the product of a Darwinian evolutionary process, may be understood as the foundation from which all subsequent evolutionary processes and domains have emerged. In this sense, we may consider those emergent processes as subroutines of cosmological selection, developed from a primary algorithm or substrate which takes the form of nature’s laws.

This brief sketch allows us to consider the cosmological domain in terms of an inferential system. It contains a knowledge repository, accumulated over evolutionary time, that in each generation of universes is updated with evidence concerning what does and what does not facilitate existence. In each generation, a variety of autopoietic algorithms encoded by the knowledge repository are executed and tested. Those that produce outputs which form a mechanism for reproduction – such as atoms, chemistry, stars, and black holes – are selected, and they retain updated heritable models of existence, variants of which will in turn be tested in future generations.

3.2 Quantum selection

The cosmological domain, orchestrated by the laws and parameters of physics, provides a fecund substrate for the evolution of complexity. To understand why it does so – assuming we accept that physics is governed by quantum laws, and that gravity is emergent from quantum phenomena – we can ask: ‘What special properties do quantum laws possess that facilitate the emergence of complexity?’ We suggest that two widely accepted quantum properties may be involved.

First, quantum theory may be interpreted as a description of inferential systems (Rovelli 1996; Campbell 2010; Hoehn 2017a). A proposed set of information theory-type axioms for quantum theory describe constraints on the acquisition of information in the quantum domain. It has been demonstrated that any observer – from a quantum observer to a human one – who processes information as stipulated by these axioms, and who uses the mathematics of Bayesian inference, will infer the descriptions of quantum theory (Hoehn 2017b). In other words, the best description of quantum phenomena that can be inferred by an observer, given the possible evidence that they could gather, is given by quantum theory.

The information which a quantum system may convey to other quantum systems in its environment, including human-constructed measuring devices, is highly constrained. In

particular, information conveying superposed quantum states – which includes the vast majority of all quantum states – cannot exist outside its own system, and so cannot be conveyed to an observer (Zurek & Zwolak 2013). Instead, only classical (or Holevo) information can survive this transition in measurable quantities. This transition occurs via a selection process dubbed quantum Darwinism (Zurek 2009).

Quantum Darwinism: Classical reality as an embodied adapted system

According to quantum Darwinism, observers will experience a quantum system in terms of the classical information – such as position, momentum, and charge – which can survive the transition to, and thus achieve existence in, classical reality. The extended network formed by this information constitutes classical reality itself (Zurek 2014). We may consider classical reality to be an embodied adapted system, constructed from the quantum knowledge repository – that is, from the quantum system’s wave function, which describes the quantum state of the system. Although the wave function is usually portrayed as a mathematical abstraction, quantum theory tells us that it must also have a physical form (Landauer 1996): quantum systems must include a physical implementation of an information processing model, equivalent to the evolution of a state vector in Hilbert space. The wave function’s physical form, however, may exist at near the Planck scale and therefore be as yet undetectable. A near consensus has only recently arisen that the wave function must have a physical incarnation (that is, an ‘ontic’ rather than just a possible ‘epistemic’ form) (Pusey, Barrett & Randolph 2012; Hardy 2013), and ‘t Hooft (2016) has speculated on its possible physical forms, near the Planck scale. We are left in a situation similar to that of biology prior to the discovery of DNA in 1953. Although there was consensus agreement at that time about the computational effectiveness of Mendel’s genes, the genes’ physical form was still unknown (and many doubted there was a physical form).

In order to understand how quantum Darwinism occurs as a system of variance, inheritance, and selection, it may help to review the dynamics of a quantum interaction. This interaction – in which information is transferred between one quantum system and another in its environment – involves the processes of entanglement and decoherence. Before two quantum systems become entangled, each is described by its own wave function. Entanglement represents a merger of the two systems: during entanglement the two systems share a single wave function, and the individual systems in effect cease to exist. Then, in the process of decoherence, the entangled system again splits into two. Each system is again described by its own wave function, in terms of a pointer or classical state, but these functions now contain a high degree of mutual information. Each system then begins to evolve separately from the other, in accordance with the Schrodinger equation. This evolution re-emerges each system into the quantum realm: the re-introduction of superpositions result in a smearing of each system’s predictive abilities, over a probability distribution of possible outcomes.

As two quantum systems becomes entangled and exchange information, and then subsequently undergo decoherence, each quantum system’s wave function is updated based on information it receives from the environment. Following the ‘quantum jump’ that results from the interaction, the system’s wave function accurately predicts (with probability 1) the outcome of the same immediately repeated interaction. In other words, the ‘quantum jump’ may be interpreted as a Bayesian update, which brings the quantum model to certainty concerning the outcome of a particular interaction. We may therefore view this information,

which the system attempts to copy to its environment, as an experimental test of a strategy for existence in classical reality. Those tests that achieve existence serve to update the quantum knowledge repository (the wave function).

Quantum decoherence essentially describes a process by which information is copied or transferred from a quantum system to its environment. Much of the varied information contained in the state vector is copied, but most has extremely short periods of survival. The fittest quantum states, or pointer states, are selected (Zurek 1998). The selected quantum states are those capable of the greatest reproductive success (Zurek 2014). Clearly quantum Darwinism is formulated in a manner consistent with three-step Darwinian dynamics: variance, inheritance, and selection. Quantum interactions which select classical (i.e., pointer) states, therefore, may be considered a Darwinian process. These interactions minimize the quantum system's entropy change, thus conforming to the paradigm in which the selection process counters the entropic process. As Zurek and colleagues describe it (Dalvit, Dziarmaga & Zurek 2005), "pointer states are obtained by minimizing system entropy and demanding that the answer be robust when varying within a reasonable range of time".

Classical or pointer states of quantum systems are selected for their ability to survive in the environment, and are robust in the face of decoherence. Quantum systems may merge, through entanglement, but again those more complex systems that survive intact are adapted to their environments and therefore able to exist. This evolutionary process explores the space of possible merged designs, and the resulting adapted systems describe an evolutionary tree: sub-atomic particles, atoms, molecules, and complex chemistry. As Zurek writes (parenthetical text added) (Zurek 2007):

In the real world observers find out pointer states of systems... by letting natural selection take its course: Pointer states are the "robust species", adapted to their environments. They survive intact its monitoring. More importantly, multiple records about S (the system) are deposited in E (the environment). They favor pointer states, which are the "fittest" in the Darwinian sense - they can survive and multiply.

Although this Darwinian description of quantum dynamics is novel, it is firmly grounded in the standard quantum axioms (Dirac 1958; Gillespie 1970). The first two of these axioms describe the wave function as continuously evolving in Hilbert space, in accordance with Schrodinger's equation. We may interpret these axioms as describing variable forms of the wave function. Two axioms describe the 'measuring process' in which information is passed to the environment in a process analogous to heredity. These axioms apply measurement operators and Born's rule to the wave function, to select a specific outcome eigenvector and a probability distribution of output eigenvalues. As Zurek demonstrates, this selection process is entirely consistent with a Darwinian interpretation in terms of heredity-variation-selection. A final axiom states that measurements or interaction result in the updating of the probabilistic wave function to greater predictive accuracy. This may be considered a Bayesian update, so the quantum system may be considered an inferential system.

In this view, these quantum axioms – accepted as standard for sixty years – readily lend themselves to a Darwinian interpretation. More recently, quantum dynamics have been shown to be consistent with the inferential system interpretation of universal Darwinism (Zurek 2014; Hoehn 2017b; Campbell 2010) that we present here.

Computing the quantum domain

A second property of quantum phenomena that enables emergent complexity is the duality between quantum physical forms and computation (Markopoulou 2012). This duality is stated as the Church-Turing-Deutsch principle (Deutsch 1985; Lloyd 2013): Any physical (quantum) phenomena may be simulated to any degree of accuracy, with an algorithm written in qubits and run on a quantum computer. As Lloyd (2007) describes it, the universe computes “its own dynamical evolution; as the computation proceeds, reality unfolds”.

The computational nature of the quantum domain raises the question: ‘Why did cosmological evolution settle on a computational substrate in its quest to evolve complexity?’ The answer suggested by Lloyd is that a quantum algorithm specifying a complex form may be much less complex – and thus more easily discovered, through a Darwinian search – than the form itself (Lloyd 2013). Thus, a computational substrate, running Darwinian algorithms, may facilitate the discovery of complex forms that are capable of existence.

Quantum computations are performed within a system of entangled qubits, but the answer to the computation is outputted through an interaction with another quantum system (e.g., a measurement) (Gavrilov 2012). In this manner we may consider classical reality as the output of quantum computation, or equivalently, as the calculation of forms capable of existence within classical reality.

3.3 Biological selection

Biological evolution by natural selection is widely understood and non-controversial, and throughout Part 2 above we provided numerous examples of how it operates in the context of universal Darwinism. We will therefore sketch biological selection only briefly in the current section, with a focus on aspects of it that we have not already emphasized.

Natural selection is generally understood as acting on a genetic knowledge repository (that may act in concert with epigenetic effects), which orchestrates the material of quantum chemistry into the many structures of life. We use this distinction to define the transition between the quantum and biological domains, following Davies and Walker (2012):

We therefore identify the transition from non-life to life with a fundamental shift in the causal structure of the system, specifically, a transition to a state in which algorithmic information gains direct, context-dependent, causal efficacy over matter.

Crucially, the mathematical model of natural selection is extremely general, substrate independent, and applicable to any population of traits which vary in frequency between generations according to a fitness function. Because any process described by this generalized formula meets the criteria of a generalized Darwinian system, the formula serves as a mathematical framework for universal Darwinism. Steven Frank (2012) demonstrates this simple, general model:

A simple model starts with n different types of individuals. The frequency of each type is q_i . Each type has w_i offspring, where w expresses fitness. In the simplest case, each type is a clone producing w_i copies of itself in each round of reproduction. The frequency of each type after selection is

$$q'_i = q_i \frac{w_i}{w}$$

Where $w = \sum_1^n q_i w_i$ is the average fitness of the trait in the population. The summation is over all of the n different types indexed by the i subscripts.

Fitness functions may be interpreted as likelihood functions in Bayesian inference. In the example above, the frequency before and after are the prior and posterior probabilities, and relative fitness is the likelihood function of a Bayesian update. Several researchers, including S. Frank (2012) and Friston (2013), have developed mathematical theories of evolutionary change equivalent to the inferential system model developed here (Campbell 2016).

In the biological case, the q_i form a probability distribution that describes the frequency of success enjoyed by a range of strategies for existence. This probabilistic model refers to autopoietic strategies that are tested as phenotypes (i.e., embodied adapted systems). Those strategies which successfully contribute to existence in the next generation – those, that is, that display fitness – update the strategy frequencies found in the next generation. As the phenotypes produced are highly complex, low-entropy systems, this process illustrates the general ability of inferential systems to overcome the forces of entropic dissipation.

A computational interpretation of biology has been developing over the past several decades. Dawkins (1982) endorsed this view decades ago, when computation was starting to become a pervasive cultural metaphor:

I shall use the word 'program' in exactly the same sense as Maynard Smith uses 'strategy'... A program (or strategy) is a recipe for action, a set of notional instructions that an animal seems to be 'obeying', just as a computer obeys its program.

A bit later, Nobel Laureate Sydney Brenner (1999) urged researchers to consider biology in computational terms:

One way of looking at the problem is to ask whether we can compute organisms from their DNA sequences. This computational approach is related to Von Neumann's suggestion that very complex behaviours may be explicable only by providing the algorithm that generates that behaviour, that is, explanation by way of simulation... A proper simulation must be couched in the machine language of the object, in genes, proteins and cells

More recently, Davies and Walker (2012) have presented biology as an essentially algorithmic process, whereby a largely genetic algorithm orchestrates its chemical substrate into the complex output of life:

As we have presented it here, the key distinction between the origin of life and other emergent transitions is the onset of distributed information control, enabling context-dependent causation, where an abstract and non-physical systemic entity (algorithmic information) effectively becomes a causal agent capable of manipulating its material substrate.

The above views fit well within the universal Darwinism framework. Specifically, they complement the notion that one natural domain evolves from another when it forms a knowledge repository (i.e., algorithm) that can orchestrate the substrate of the previous domain into new complex forms that are capable of existence.

The algorithmic nature of genetics has been illuminated through the Encode project, which found that the majority of the genetic repository codes for regulatory functions rather than for the direct construction of proteins (Brydon 2012):

It is thought that these “non-coding” RNA transcripts act something like components of a giant, complex switchboard, controlling a network of many events in the cell by regulating the processes of

replication, transcription and translation – that is, the copying of DNA and the making of proteins based on information carried by messenger RNAs.

We may think of this regulation in terms of computer programs containing many nested ‘IF, THEN, ELSE, END’ statements. For example, “IF the concentration of molecule A is above a given level, THEN produce more molecule B, ELSE produce more molecule C, END.” Such algorithms appear to be methods widely used by life to compute biological output from chemical input.

3.4. Neural selection

Biological natural selection – understood as the evolution of principally genetic knowledge repositories, capable of instantiating and testing phenotypic adapted systems – has produced a wide range of adaptations. Biological selection has fashioned these adaptations from a variety of organismal tissue types, and for hundreds of millions of years this variety has included neural tissue. An adaptation constructed from neural tissue is not less ‘biological’ than one constructed from, say, skeletal or muscular tissue, and the category of ‘biological adaptations’ includes many genetically-encoded, psychological, information-processing adaptations. The field of evolutionary psychology (Tooby & Cosmides 1992) is devoted to the identification and functional analysis of species-typical adaptations of this kind. We emphasize these points to make clear that in the context of the ‘neural selection’ domain, we will *not* be focusing on genetically-encoded psychological adaptations, because we consider these to be biological adaptations. The neural selection domain is concerned, rather, with neuronal systems that are organized primarily via experience and learning.

To further clarify the distinction between the psychological embodied adapted systems that exist in the biological versus the neural domain, consider spoken human language. Human nature appears to be equipped with genetically-encoded psychological language-acquisition mechanisms (Pinker & Bloom 1992). These adaptations enable all normally-developing individuals to acquire, relatively automatically, whatever language is spoken in their early childhood developmental environment. Because these mechanisms are species-typical, all human cultures exhibit complex spoken language. There is of course diversity, however, in the linguistic rules and words used across cultures; just because you know one language does not mean you will know any others. Language acquisition, therefore, depends on two important kinds of knowledge repositories. The first is the ‘biological’ knowledge repository, and it provides species-specific, genetically-encoded language acquisition mechanisms, which enable children to acquire whatever language is spoken in their population. The second is the ‘neural’ knowledge repository, which stores the child’s knowledge about the specific rules and words of their own spoken language.

It is crucial to emphasize that ‘biological’ knowledge repositories are genetic and ‘neural’ ones are non-genetic. It is also important, however, to note that from a certain perspective it is non-sensical to refer to neural knowledge repositories as non-genetic. That’s because if it weren’t for genes, these neural structures would, of course, not exist at all. It is ultimately genes, for example, which encode for the language acquisition mechanisms that enable a child to acquire a specific language, even if the genes do not specify what language this will be. Although the terminology of ‘genetic’ versus ‘non-genetic’ presents this risk of confusion, however, this confusion itself reflects a familiar pattern in universal Darwinism: a new kind of natural domain has emerged, because the material of a substrate domain has been

orchestrated into a new kind of knowledge repository. In this case, the biological substrate of genetically-encoded neurons has been orchestrated to create a new, neural-based knowledge repository, which is organized primarily via experience and learning.

In the domain of neural selection, if the neural structure represents the knowledge repository, then the embodied adapted system would be the behavioral output of this neural structure. For instance, in the case of acquired language the embodied adapted system would be acts entailing the use of language to achieve one's goals. Another example would be an organism that has to learn which species in its environment constitute edible prey (such as the southern toad in the example below). The neural information-processing system that enabled such learning would take environmental input into consideration (e.g., which prey species tasted best), and produce learned predatory behavior as output. Learning-based neural systems thus provide a new level of information processing, to coordinate an organism's repertoire of neural-based sensory information with its repertoire of neural-based behaviors.

Bayesian brains

The cerebral cortex is largely responsible for this sensory/behaviour mapping in mammals, but the evolutionary origins of such mapping may extend back to near the genesis of multi-cellular organisms (Tomer 2010). The cerebral cortex is universally composed of sophisticated computational units called 'cortical columns' and a near consensus is emerging that the computations performed are a type of Bayesian inference (Lee 2015), as for example with the predictive coding model (Cain et al. 2016). This convergence of neuroscientific understanding to a view of the cerebral cortex as a type of Bayesian computer is sometimes called the Bayesian Brain theory. Karl Friston and colleagues summarize this computational algorithm executed by the Bayesian brain (Bastos et al. 2012):

It might be thought impossible to specify the computations performed by the brain. However, there are some fairly fundamental constraints on the basic form of neuronal dynamics. The argument goes as follows—and can be regarded as a brief summary of the free energy principle.

- Biological systems are homeostatic (or allostatic), which means that they minimize the dispersion (entropy) of their interoceptive and exteroceptive states.
- Entropy is the average of surprise time, which means that biological systems minimize the surprise associated with their sensory states at each point in time.
- In statistics, surprise is the negative logarithm of Bayesian model evidence, which means biological systems—like the brain—must continually maximize the Bayesian evidence for their (generative) model of sensory inputs.
- Maximizing Bayesian model evidence corresponds to Bayesian filtering of sensory inputs. This is also known as predictive coding.

This formulation incorporates senses and perceptions as inputs to a Bayesian inferential system, and behaviors as outputs (Friston et al. 2017):

This allows one to frame behavior as fulfilling optimistic predictions, where the optimism is prescribed by prior preferences or goals. In other words, action realizes predictions that are biased toward preferred outcomes. More specifically, the generative model entails beliefs about future states and policies, where policies that lead to preferred outcomes are more likely. This enables action to realize the next (proximal) outcome predicted by the policy that leads to (distal) goals. This behavior emerges when action and inference maximize the evidence or marginal likelihood of the model generating predictions.

The generative neural models produced through the processing of learned information entail beliefs about future states, and these models are updated through evidence-based (i.e. sensory) experience to become more accurate in reaching preferred goals. The sum of these generative models in an organism may be considered as its total knowledge repository in the neural domain. The selective process, described in terms of preferences or goals, is fully consistent with Darwinian mathematics in terms of a fitness function. A range of variable generative models is selected, based on fitness goals; these somewhat random behavioural variations are experimented with, especially among immature organisms. The behaviours that are selected and retained are those which maximize ‘fitness’ (goal achievement), that is, those which produce evidence in support of predicted success. Successful behaviours are learned and able to be reproduced in appropriate future circumstances, constituting a type of heredity. Thus the Bayesian Brain theory, cast in terms of the free energy principle, may be interpreted as a Darwinian process.

Over the past decade, Friston has demonstrated how the Bayesian Brain model emerges from a general principle of evolutionary dynamics called the free energy minimization principle. This principle states that dynamic systems act to reduce both the prediction error of their internal models and the complexity of those models. When applied to neuroscience, this principle states that internal models of the world are updated, in a Bayesian manner, to reduce discrepancies between the evidence the model receives in the form of sensory inputs and the model’s predictions about those inputs. The brain achieves this accuracy with a low level of complexity, as is evident when we consider that the cerebral cortex accomplishes sophisticated perceptual and behavioural tasks more accurately than the current generation of super computers, despite using much less energy and computational power than these computers (Fischetti 2011).

Friston and colleagues have recently extended the free energy paradigm to encompass the dynamic systems of biology and culture (Friston 2013; Ramstead, Badcock & Friston 2017). The free energy minimization principle has been shown to be equivalent to the inferential systems model of evolutionary dynamics (Campbell 2016); systems tend to evolve more accurate internal models, detailing possible forms of existence, by executing variations of existing strategies and recording the evidence of their relative success. For example, the southern toad will learn through experience to avoid eating bumblebees (Tinbergen 1965), and this learned behaviour involves an evidence-based update of the toad’s internal models regarding edible insects.

Such refinements of internal learned mental models lead to more fine-tuned and less ignorant behavior, or equivalently, to lower-entropy behavior. We may thus consider the domain of neural selection to represent another case of nature using evolutionary systems to counter the entropic process.

3.5 Cultural selection

Homo sapiens is unique, among the species that evolution has brought into existence on our planet, in producing an evolving culture. While a handful of other species display some rudimentary cultural traits, only our species participates in a culture that accumulates new knowledge in a continuous fashion. Even our nearest and now extinct relatives in the

genus *Homo* appear to have maintained a relatively static kit of artefacts over the duration of their species. *Homo sapiens*, on the other hand, in less than a hundred thousand years succeeded in migrating out of Africa, colonizing practically the entire planet, setting foot on the moon, and placing robotic instruments throughout the solar system.

All these accomplishments have been due to human cultures, which have adapted through the continuous accumulation of new shared knowledge. This cultural knowledge has orchestrated the behaviour of group members into a new level of adapted behaviour, which has in turn tended to promote cultural success (i.e., the propagation of cultural adaptations). The domain of cultural selection, in other words, continues a pattern that should by now seem familiar, as it is repeated throughout the universal Darwinism framework: the substrate of the preceding domain (here, neural-based learned behavior) is orchestrated towards the production of new kind of knowledge repository, thus defining the boundary of a new selective domain (here, the cultural domain). From this perspective, we are well positioned to unravel many mysteries surrounding the emergence of the cultural knowledge repository and to understand, at least in broad outline, how it has emerged.

Cultural knowledge is rooted in the neural-based social learning we share with our primate relatives and which, in turn, has been enabled through the biological evolution of structures such as neurons, brains, and specialized neural circuitry. In particular, *Homo sapiens* inherited biological and neural sensory abilities evolved by our animal ancestors since the Cambrian explosion. These abilities have been honed and fine-tuned over a vast evolutionary period, to provide the brain with accurate and reliable information about what is going on in the outside world. In addition to our more primitive sensory brain regions, we possess a uniquely enormous cerebral cortex, constituting roughly eighty per cent of our brain mass. This elaborate evolutionary addition to our brain has enabled our species to engage in novel and relatively powerful types of imaginative thinking; we can consider a plethora of possible explanations for outcomes in the world, and many possible ways in which these outcomes could be related. The possibilities we can imagine include those we cannot verify by our senses: we can think about angels and demons, and about scientific puzzles which we are currently unable to conclusively solve. Such imaginative thoughts can be compelling but are, alas, prone to inaccuracy; we can easily imagine things which are untrue.

However, these two functional areas of our brain – the sensory and the imaginative – may interact synergistically, whereby speculations generated in our cerebral cortex are judged and selected using sensory information. Such synergy would constitute essentially Bayesian, evidence-based hypothesis selection. It appears likely that the brain's pervasive neural architecture of cortical columns is a specific adaptation designed to perform this type of Bayesian computation, on inputs originating from the senses and the cerebral cortex (Cain et al. 2016). We suggest that this synergy is what powers cultural evolution.

The imaginative process produces variations on cultural beliefs and practices which may be inherited through social learning or imitation, with selection being made among these variations based on the sensory evidence. Once again, that is, we see an equivalence between the Bayesian and Darwinian formulations of evolutionary dynamics. For example, archaeologists are increasingly using techniques developed within evolutionary biology, such as evolutionary trees or cladistics, to describe the phylogeny of cultural artefacts such as arrowheads. Some have described this program of placing archaeological artefacts within a

biologically inspired phylogeny as the ‘fundamental issue of evolutionary archaeology’ (O’Brien & Lyman 2003). This research suggests that cultures experiment with many variations on artefact design, and that evidence about design effectiveness is retained within the repository of cultural knowledge in a Darwinian manner, similar to natural selection.

Nowhere is the Darwinian/Bayesian nature of cultural practice clearer than in the case of science. In science, alternative hypotheses are generated by our imaginative minds, and experimental tests are devised to provide empirical, sensory evidence that becomes the basis for selecting among these hypotheses. Through science, cultural evolution accumulates evidence-based knowledge and thus reduces ignorance or entropy. We thus regard culture as the most recently-emerging domain of nature in which accumulated knowledge serves to defy the spirit of the second law.

Written language enabled the cultural knowledge repository to come into its own

Early ancestral members of our species lived in small hunter gatherer societies, sharing a common body of cultural knowledge but lacking a great deal of specialization among its members. Social learning was facilitated between members and between generations using a complexly descriptive spoken language. At this early stage of societal evolution, the cultural knowledge repository was constrained to the sum of the cultural knowledge possessed by living individual members. As agriculture provided the means for larger groups with greater specializations, however, individualized knowledge became a constraint on an evolving, widely shared body of cultural knowledge. Human memory is fallible and the recording of important historical facts, such as resource accounting, required a more objective, robust and widely accessible knowledge repository. This was achieved by transferring parts of the cultural knowledge repository to the external medium of written information, where it could be accessed by multiple members and survive the mortality of individual knowledge-bearers. The earliest-known example of such writing are cuneiform clay tablets, used in Mesopotamia for accounting purposes at least 5,000 years ago.

With the invention of the printing press and numerous other knowledge repositories such as those involved with computers, our external cultural knowledge repository is growing at an exponential rate. This fairly well-documented transition of the cultural knowledge repository from internal models within the minds of individuals to shared externally recorded models illustrates the dynamic evolutionary transitions which knowledge repositories can undergo.

How Bayesian is cultural evolution?

We may view cultural evolution as a largely Bayesian process, where hypotheses concerning possible forms of cultural existence are tested by evidence. There is a vast range of cultural practices, however, between those that operate in an ‘approximate’ Bayesian manner and those that follow Bayesian mathematics more exactly.

For example, the evolution of agriculture seems likely to have been a largely Bayesian process, with many different practices designed to increase food production being attempted and those that have worked being preferentially selected. In this manner, humans have been influencing the evolutionary fitness of many species for at least tens of thousands of years.

Cultural practices have apparently influenced the phylogenetic history of cereal crop plants, for instance, for over thirty thousand years (Allaby 2017).

Over agriculture's long history, humans have learned how to exert increasingly greater control over the variables influencing crop productivity. Agriculture has thus evolved to become a more exact Bayesian method in the minds of its practitioners, as they have learned to juggle the many variables under their control with the intent of increasing crop yields. Recent refinements in agricultural techniques have allowed human culture to achieve breakthroughs in biochemical designs, which have overcome limitations imposed by the biology of crop species. For example, unaided biology does not appear to possess an enzyme capable of gene substitution in an intact DNA strand, but this feat has recently been accomplished by researchers applying the agricultural principles of artificial selection to the chemical evolution of enzymes in a process named *DNA-Templated Organic Synthesis* (DTS) (Li & Lui 2004). As these researchers note, "such an enormous sequence space can in principle be explored efficiently by iterated cycles of DTS-based translation, selection for desired binding or catalytic properties, template amplification by PCR, and template diversification by mutagenesis or recombination, representing a true evolutionary process." It is clear, in other words, that such methods of modern agriculture – and of scientific practice more broadly – employ the Darwinian notion of a "true evolutionary process".

However, processes of cultural evolution may often appear to be only approximately Bayesian. As noted above, the human mind is highly fallible as a scientific apparatus, and what constitutes evidence for Bayesian updates of cultural models will often be unclear to members of that culture. Ultimately, however, evidence about how well a practice contributes to cultural success will influence the fate of that culture, even if not easily calculated by its members themselves. In other words, whereas nature is precisely Bayesian in how it weighs the evidence (e.g., about the influence of an agricultural practice on a culture's migratory success), this process may be perceived as more opaque and approximate to members of the culture itself. It is nevertheless human agency which determines the evolutionary trajectory: our ability to imagine and test the innovative possibilities, and to select and retain those which contribute to cultural success, will determine the fate of our own cultures.

Only with the emergence of science have humans come to fully understand the power of the Bayesian synergy that is performed by the sensory and imaginative components of our brains. Indeed, Bayesian probability has been hailed as 'the logic of science' (Jaynes 2003). All cultural knowledge may be typified as being produced by inferential systems, but science has vastly accelerated the process of cultural evolution because it aspires to a more precise, less approximate form of Bayesian inference. Scientific hypotheses are constructed to be testable by evidence, which is gathered through experiments that have been carefully designed and conducted for exactly that purpose. This ensures that the Bayesian likelihood function, which describes how well each competing hypothesis is supported by the evidence, is relatively straightforwardly assessed and easily computed.

In cultural practices that are less explicitly scientific – for example, early agriculture or arrowhead design – competing hypotheses about innovation effectiveness, and the ways in which evidence should be brought to bear, may be less clear. If an innovation succeeds, for example, was it due to the practitioner's status, or to their skill, or to the type of incantation they used? Science brings to the forefront of our consciousness the synergistic potential

between imagination and evidence, unleashing a whole new level of human knowledge accumulation that better conforms to nature's timeless method of precise Bayesian inference.

The algorithmic and autopoietic nature of cultural knowledge has been noted by many, but perhaps none as poetically as George Spencer-Brown (1979):

It may be helpful at this stage to realize that the primary form of mathematical communication is not description, but injunction. In this respect it is comparable with practical art forms like cookery, in which the taste of a cake, although literally indescribable, may be conveyed to a reader in the form of a set of injunctions called a recipe. Music is a similar art form, the composer does not even attempt to describe the set of sounds he has in mind, much less the set of feelings occasioned through them, but writes down a set of commands which, if they are obeyed by the reader, can result in a reproduction, to the reader, of the composer's original experience... Even natural science appears to be more dependent upon injunction than we are usually prepared to admit. The professional initiation of the man of science consists not so much in reading the proper textbooks, as in obeying injunctions such as 'look down this microscope'.

The individual in a cultural context is seldom left to puzzle things out on their own, but rather is 'taught' or 'instructed' in the algorithms of cultural practices. Spencer-Brown provides some examples above, and we may also recognize their pertinence to the examples we have previously explored such as arrowhead design ('always start flaking the blank like this') and agriculture ('always plant this crop on the first full moon after the spring equinox'). Such algorithms, when faithfully followed, produce outputs of cultural adaptations that tend to propagate themselves successfully.

That much of cultural practice may be considered computational is made evident by the extent to which computers can 'automate' them. This automation encompasses not only business, communications, and manufacturing, but extends even to games such as chess and go. In particular, the ability of computation to mimic the evolutionary Darwinian/Bayesian paradigm, through an exponentially growing number of applications involving evolutionary computation, has revolutionized the practice of science (Ignatova, Marinez-Perez & Zimmermann 2008).

4. Summary and Conclusion

In this chapter we have endeavoured to show how, across all the domains of nature, the process of Darwinian selection acts to creatively generate complex order, against the antagonistic tide of the second law. In all these domains, the selective process is characterized by the same key components. Each domain is defined by its own knowledge repository, a probabilistic model which puts forward guesses – instantiated in the form of an embodied adapted system – about what it takes to exist in that domain. The guesses are tested for accuracy by the embodied adapted system, and test results are fed back to the knowledge repository. The knowledge repository is then updated for greater accuracy in a Bayesian manner, thus reducing its ignorance and entropy. These entropy-reducing inferential systems operate in the context of 'variance – inheritance – selection' Darwinian dynamics: some knowledge repositories are smarter than others, so they produce more copies of themselves and thus competitively outperform their more ignorant rivals.

We propose that these Darwinian inferential systems are responsible for all of the complex order that exists in the five natural domains of the cosmological, quantum, biological, neural, and cultural. These domains exist in nested hierarchy, and those higher in the hierarchy can be considered subroutines of those that are lower. Each domain's knowledge repository computationally transforms the substrate of the preceding domain, to generate novel forms of rebellion against the second law.

Overall, this universal Darwinism framework can be seen as jigsaw puzzle that, while still incomplete, has already taken recognizable shape. It has some important missing pieces, but is complete enough to provide good ideas of how the missing pieces are shaped, and of what the completed puzzle is going to look like. For example, with regard to the knowledge repositories in the cosmological and quantum domains, we are in a position similar to that of biologists in the decades preceding the discovery of DNA: we are yet to directly observe the physical forms of these knowledge repositories, but we suspect that these forms must exist, and have plausible ideas about what they might look like. It will be exciting to watch as the puzzle implied by universal Darwinism becomes increasingly complete in the years ahead. We expect that as progress towards completion is made, there will be increasing support for the notion that as an explanation for the origins of complex order, Darwin's theory is far more fundamental than has generally been appreciated.

References

- Allaby, R. G. et al. (2017). Geographic mosaics and changing rates of cereal domestication. The Royal Society.
- Arvey R. & Colarelli S. M. (Eds). (2015). *Biological foundations of organizational behavior*. University of Chicago Press.
- Bastos A. M. et al. (2012). Canonical microcircuits for predictive coding. *Neuron*, Vol. 76(4).
- Bishop, E. (1967). *Foundations of Constructive Analysis*. New York: Academic Press. ISBN 4-87187-714-0.
- Brenner, S. (1999). Theoretical biology in the third millennium. *Philosophical Transactions of the Royal Society*.
- Brydon, E. (2012). In massive genome analysis ENCODE data suggests 'gene' redefinition. Cold Springs Harbor Laboratory. Downloaded 2nd Dec. 2013 from <http://www.cshl.edu/Article-Gingeras/massive-genome-analysis-by-encode-redefines-the-gene-and-sheds-new-light-on-complex-disease>.
- Buss, D. M., Haselton, M. G., Shackelford, T. K., Bleske, A. L., & Wakefield, J. C. (1998). Adaptations, exaptations, and spandrels. *American Psychologist*, 53(5), 533.

Cain N. et al. (2016). The Computational Properties of a Simplified Cortical Column Model. *PLOS computational biology*, Vol. 12(9).

Campbell, D. T. (1974). Evolutionary epistemology. In P.A. Schilpp (Ed.), *The Philosophy of Karl R. Popper*. The Library of Living Philosophers. LaSalle, IL: Open Court Publishing Company, Volume 14-1, 413–463.

Campbell, J. O. (2010). Quantum Darwinism as a Darwinian process. arXiv preprint.

Campbell, J. O. (2015). *Darwin does physics*. s.l.: CreateSpace.

Campbell, J. O. (2016). Universal Darwinism as a process of Bayesian inference. *Front. Syst. Neurosci.* doi: 10.3389/fnsys.2016.00049.

Comte, A. (1848). *A General View of Positivism*. Translation by Bridges, J. H. (1865), Trubner and Co.; reissued by Cambridge University Press (2009; ISBN 978-1-108-00064-2).

Crane, L. (1994/2010). Possible implications of the quantum theory of gravity: an introduction to the meduso-anthropic principle. arXiv:hep-th/9402104v1. Reprinted in Crane, L. (2010). Possible implications of the quantum theory of gravity: an introduction to the meduso-anthropic principle. *Foundations of Science*, vol. 15, pp. 369–373.

Dalvit, D. A. R., Dziarmaga, J. & Zurek, W. H. (2005). Predictability sieve, pointer states, and the classicality of quantum trajectories. *Phys. Rev. A*, Vol. 72, 062101.

Darwin, C. R. (1859). *On the Origin of Species*. London: John Murray.

Davies, P. C. W. & Walker, S. I. (2012). The Algorithmic Origins of Life. *Journal of the Royal Society*, Vol. 10.

Dawkins, R. (1982). *The Extended Phenotype: The Long Reach of the Gene*. Oxford University Press.

De Jong, K. A. (2006). *Evolutionary computation: a unified approach*. Cambridge MA: MIT Press.

Dennett, D. (1995). *Darwin's Dangerous Idea*. New York: Schuster.

Deutsch, D. (1985). Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London*, Vol. A 400 pg. 97-117.

Deutsch, D. (1997). *The Fabric of Reality*. London: Penguin.

Deutsch, D. (2011). *The Beginning of Infinity*. London: Penguin.

Dirac, P.A.M. (1958). *Quantum Mechanics*. Oxford: Clarendon Press.

Fernando, C, Szathmary, E & Husbands, P. (2012). Selectionist and evolutionary approaches to brain function: a critical appraisal. *Computational Neuroscience*.

Fischetti, M. (2011). Computers versus Brains. *Scientific American*. November 1, 2011.

Frank, R. H. (2012). *The Darwin economy: Liberty, competition, and the common good*. Princeton University Press.

Frank, S. A. (2012). Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *Journal of Evolutionary Biology*, Vols. 25: 2377-2396.

Friston, K. J. (2013). Life as we know it. *Journal of the Royal Society Interface*, Vol. 10: 20130475.

Friston K. J. et al. (2017). Active Inference: A process theory. *Neural Computation*, Vol. 29(1).

Gardner, A. & Conlon, J. P. (2013). Cosmological natural selection and the purpose of the universe. *Complexity* 18: 48–56.

Gardner, J. N. (2000). The selfish biocosm. *Complexity* 5: 34-45.

Gavrilov, Z. (2012). Quantum computation. MIT. Downloaded 5th Dec. 2017 from <http://web.mit.edu/zoya/www/quantComp.pdf>.

Gibson M. A. & Lawson, D. W. (Eds). (2014). *Applied evolutionary anthropology: Darwinian approaches to contemporary world issues*. Springer.

Gillespie, D. T. A. (1970). *Quantum Mechanical Primer*. Scranton: International Textbook Company.

Gould, S. J. & Vrba, E. S. (1982). Exaptation – a missing term in the science of form. *Paleobiology* 8 (1): 4–15.

Hamilton, W. (1964). The genetical evolution of social behaviour I. *Journal of Theoretical Biology*, vol. 7, no. 1, pp. 1–16.

Hardy, L. (2013). Are quantum states real? *International Journal of Modern Physics B*.

Harrison, E. R. (1995). The natural selection of universes containing intelligent life. *Quarterly Journal of the Royal Astronomical Society* 36: 193-203.

- Heylighen F. (1999): "The Growth of Structural and Functional Complexity during Evolution", in: F. Heylighen, J. Bollen & A. Riegler (eds.) *The Evolution of Complexity* (Kluwer Academic, Dordrecht), p. 17-44.
- Hoehn, P. A. (2017a). Quantum theory from questions. *Phys. Rev.*, Vol. A 95, 012102.
- Hoehn, P. A. (2017b). Quantum theory from rules on information acquisition. *Entropy*, Vol. 19(3), 98.
- Ignatova, Z., Marinez-Perez, I. & Zimmermann, K. (2008). *DNA Computing Models*. Springer. ISBN 978-0-387-73635-8.
- Jaynes, E. T. (1957). Information Theory and Statistical Mechanics I., *Phys. Rev.*, 106, 620.
- Jaynes, E. T. (1965). Gibbs vs Boltzmann Entropies. *Am. J. Phys.*, 391.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. University of Cambridge Press.
- Landauer, R. (1996). The physical nature of information. *Physics Letters A*, Vol. 217. [https://doi.org/10.1016/0375-9601\(96\)00453-7](https://doi.org/10.1016/0375-9601(96)00453-7).
- Lee, T. S. (2015). The visual system's internal model of the world. *Proceedings of the IEEE. Institute of Electrical and Electronics Engineers*, Vols. 103(8), 1359–1378.
- Li, X. & Lui, D. (2004). DNA-Templated Organic Synthesis: Nature's Strategy for Controlling Chemical Reactivity Applied to Synthetic Molecules. *Angew. Chem. Int. Ed.*, Vol. 43, 4848.
- Lloyd, S. (2007). *Programming the Universe*. Vintage; Reprint edition.
- Lloyd, S. (2013). *The Universe as Quantum Computer*. Arxiv preprint.
- Maldacena, J. M. (1998). The large N limit of superconformal field theories and supergravity. *Adv.Theor.Math.Phys*, Vol. 2, pp. 231-252.
- Markopolou, F. (2012). *The computing spacetime*. [ed.] S. B. Cooper, et al. *How the world computes*. Berlin, Hei : Springer.
- McGregor, S. (2012). Evolution of Associative Learning in Chemical Networks. *Comput Biol* 8(11). doi:10.1371/journal.pcbi.1002739.
- Nelson, R. R. & Winter, S. G. (1982). *An evolutionary theory of economic change*. Harvard University Press. ISBN 0-674-27228-5.

- O'Brien, M. J. & Lyman R. L. (2003). Resolving Phylogeny: Evolutionary Archaeology's Fundamental Issue. [book auth.] T. L. VanPool and C. S. VanPool. *Essential Tensions in Archaeological Method and Theory*. Salt Lake City: University of Utah Press, pp. 115-135.
- Orr, H. A. (2000). Adaptation and the cost of complexity. *Evolution*, 54(1), 13-20.
- Orus, R. (2014). A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States. *Annals of Physics* 349, 117-158.
- Pietraszewski, D. Curry, O.S., Petersen, M.B., Cosmides, L. & Tooby, J. (2015). Constituents of political cognition: Race, party politics, and the alliance detection system. *Cognition*, 140, 24–39.
- Pinker, S. & Bloom, P. (1992). Natural language and natural selection. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, J. Barkow, L. Cosmides, and J. Tooby, Eds. pp. 451-493. Oxford University Press.
- Popper, K. (1972). *Objective Knowledge*. Clarendon Press.
- Price, M. E. (2017). Entropy and selection: Life as an adaptation for universe replication. *Complexity*, vol. 2017, Article ID 4745379, 4 pages, 2017. doi:10.1155/2017/4745379
- Pusey M F, Barrett J & Randolph T. (2012). On the reality of quantum states. *Nature Physics* 8, pp. 475 - 478.
- Ramstead M. J. D., Badcock P. B., & Friston K. J. (2017). Answering Schrödinger's question: A free-energy formulation. *Phys Life Rev.*
- Richerson, P. J. & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. University of Chicago Press.
- Rovelli, C. (1996). Relational Quantum Mechanics. *International Journal of Theoretical Physics*, Vols. 35, pp. 1637-78.
- Saad G. (Ed). (2011). *Evolutionary Psychology in the Business Sciences*. Berlin: Springer.
- Shannon, C. E. (1948). A mathematical theory of communications. *Bell System Technical Journal* 27(3): 379–423.
- Smart, J. M. (2009). Evo devo universe? A framework for speculations on cosmic culture. In *Cosmos and Culture: Cultural Evolution in a Cosmic Context*, S. J. Dick and M. L. Lupisella, Eds., pp. 201–295, Government Printing Office, NASA SP-2009-4802, Washington, DC.
- Smolin, L. (1992). Did the universe evolve? *Classical and Quantum Gravity*, vol. 9, no. 1, pp. 173–191.

- Smolin, L. (1997). *The Life of the Cosmos*. New York: Oxford University Press.
- Smolin, L. (2007). The status of cosmological natural selection. Arxiv preprint.
- Spencer-Brown, G. (1979). *The Laws of Form*. New York: E.P. Dutton.
- Stephens, C. R., 't Hooft, G. and Whiting, B. F. (1993). Black hole evaporation without information loss. *Classical and quantum gravity* Volume 11, Number 3.
- Swingle, B. & Van Raamsdonk, M. (2014). Universality of Gravity from Entanglement. ArXiv preprint.
- 't Hooft, G. (2016). *The cellular automaton interpretation of quantum mechanics*. Springer.
- Thayer, B. A. (2004). *Darwin and international relations: On the evolutionary origins of war and ethnic conflict*. University Press of Kentucky.
- Tinbergen, N. (1965). *Animal behaviour*. Time Inc.
- Tomer R. et al. (2010). Profiling by image registration reveals common origin of annelid mushroom bodies and vertebrate pallium. *Cell*, Vol. 142 (5): 800–809.
- Tooby J. & Cosmides L. (1992). The psychological foundations of culture. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, J. Barkow, L. Cosmides, and J. Tooby, Eds. Oxford University Press.
- Verlinde, E. (2010). On the Origin of Gravity and the Laws of Newton., arXiv:1001.0785.
- Vidal, C. (2014). *The Beginning and the End: The Meaning of Life in a Cosmological Perspective*. Springer.
- Wigner, E. P. (1960). "The unreasonable effectiveness of mathematics in the natural sciences. Richard Courant lecture in mathematical sciences delivered at New York University, May 11, 1959". *Communications on Pure and Applied Mathematics* 13: 1–14.
- Williams, G. C. (1966). *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton University Press.
- Zurek, W. H. (1998). Decoherence, Einselection and the Existential Interpretation. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, pp. 1793-1821.
- Zurek, W. H. (2007). Relative States and the Environment: Einselection, Envariance, Quantum Darwinism, and the Existential Interpretation. arXiv:0707.2832v1.
- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics*, vol. 5, pp. 181-188.

Zurek, W. H. (2014). Quantum Darwinism, classical reality and the randomness of quantum jumps. *Physics Today*, Vols. 67, 10, 44.

Zurek, W. H. & Zwolak, M. (2013). Complementarity of quantum discord and classically accessible information. *Scientific Reports* 3, Article number: 1729. doi:10.1038/srep01729.

Zylberberg, Ariel , et al. (2011). The human Turing machine: a neural framework for mental programs. *Trends Cogn Sci* 15: 293–300.