Metabolic Basis of Complex Adaptive Systems

A generative theory

Vince Kane

George Mason University
* `vkane2@gmu.edu`

Abstract Homeostasis away from thermal equilibrium is a defining characteristic of life; metabolism is the physical process by which living systems reject entropy. I argue that a thermodynamic and information theoretic characterization of metabolic activity is sufficient to generate the principal features of complex adaptive systems (CAS) observed in the wild, to include ecosystems and social systems. Starting from first principles and applying a simple constraint of thermodynamic stability (entropy does not accumulate within the system), I take some initial steps in deriving generative mechanics for CAS and describe how these mechanics account for the phenomenological properties of CAS. Growth, diversification, and adaptation result from a minimum set of information-retaining functions operating in a dynamic environment, and recursion over that set. The complexity observed in social systems (to include non-human) is an emergent extension of basic metabolic processes at lower levels of the CAS hierarchy.

Key words: cas · complexity · adaptation · thermodynamics · information theory · metabolic · economics

1 Introduction

There is an increasing recognition of the fundamental and central role of thermodynamics in the behavior of complex systems. Even use of the term "complexity" itself is problematic, except in some relation to thermodynamics, where perhaps its most parsimonious definition may be found – complex systems are acknowledged to be characterized by "non-equilibrium" processes, with the equilibrium being referred to as thermodynamic equilibrium. Complex systems in some sense surprise us, because they are not at thermal equilibrium with their environment. There is a growing body of work applying thermodynamic analysis to specific domains; for example, a number of works in the collection [6] applies the Maximum Entropy Production Principle to a number of physical, chemical, and earth science complex systems. However, a unified theory of complex systems remains to be formulated.

When we treat with complex *adaptive* systems (CAS), the picture gets even messier, complicated by the diversity of form and process associated with CAS. This rampant diversity is itself, with accompanying non-linearity, recognized as a defining feature of CAS, but presents difficulty in formulating a unifying theory of CAS [10]. Notably, Smith [20],[21],[22] has applied thermodynamic analysis to biological systems, such as the Carnot cycle and population genetics. Krakauer [12] as well has investigated the information theoretic properties of evolution (information being the thermodynamic inverse of entropy), characterizing natural selection as the Maxwell's Demon of evolution. Schneider and Kay [17],[16] argue for a thermodynamic approach to the study of ecosystems, develop a theoretical framework for doing so, and provide data and analysis demonstrating that such an

---

* I would like to thank Rob Axtell for many thought-provoking discussions and recommended readings, and Qing Tian and Kevin Curtin for constructive criticism provided to a previous draft of this paper. I should also credit the works of Ayres, Arthur, Beinhocker, and Schrodinger for inspiring my formulation of the metabolic model.

approach does indeed yield additional insight into the behaviors of ecosystems. However, these analyses are domain-specific, and descriptive in nature (even if, in most cases, very mathematical).

In the realm of economics, there is a *body academic* of complexity-oriented economists, generally considered on the "fringe", who advocate for a more robust treatment of thermodynamics in economic theory and analysis: Georgescu-Roegen [15], Foley (and co-author Smith) [23], Arthur [1], Beinhocker [5], and Ayres ([13],[2],[3],[4], to name a handful). It is ironic that economics as a science has its genesis in analogy to the first law of thermodynamics (conservation of energy – thus "general equilibrium theory") [5], but that 2nd law thermodynamics (entropy in a system tends to increase) has made little headway in mainstream economics because the analogy to thermodynamics (specifically, general equilibrium theory) is considered played out and taken as far as it can go [23].

On the unification front, Holland [10] is perhaps the most widely-recognized as providing the most complete and general theory of CAS to date. Holland identifies seven basic characteristics that are common to all CAS – four properties: aggregation, nonlinearity, flows, and diversity – and three mechanisms: tags, internal models, and building blocks. These seven characteristics recur again and again in CAS; in Holland's opinion, they are the minimal set for describing CAS (or at least a set to which other theories reduce to). Holland's pioneering work provides a framework for identifying, analyzing, and understanding CAS. Gell-Mann [8],[9] as well provides a broad framework, though it is more heavily focused on the informational content of CAS 'schemata' (corresponding to Holland's internal models).

What all of these works have in common is that they are primarily phenomenological characterizations, and usually domain-specific in their analysis (Holland and Gell-Mann being exceptions to the latter). To be sure, observation and characterization of the regularities in natural phenomena are the first step in doing real science, as they are a requisite precursor to formulating valid explanatory models of the phenomena that have been observed. To date, however, the field of CAS lacks the follow-on explanatory theory – a universal model that generates (and thus explains) the observed phenomena from fundamental principles.

For the remainder of this paper, I propose such an explanatory theory. I begin with a qualitative description of the theory, and follow it with a more formalized model. A brief sketch of the theory is as follows:

1. The most universal characteristic of life on earth is far-from-equilibrium existence (Schrodinger [18]); that is to say, organisms avoid the condition of thermal equilibrium.
2. Metabolism is the process by which organisms do so (again, Schrodinger [18]).
3. Metabolism may be modeled as a [composite] function that harvests energy and ransforms it into information. Information (negative entropy) compensates for entropy generated internally and introduced by the environment.
4. Because of entropy, replication is a necessary function for the persistence of information.
5. Variation introduced during replication enables adaptation.
6. Entropy is a universal natural selector.
7. Diversity is the accumulation of mutual information in ensembles of CAS.

2 A qualitative description

2.1 The 1st and 2nd Laws of Thermodynamics

The laws of thermodynamics have been termed “the consititution of the universe” – we have yet to observe any physical process that violates them. They govern the behavior of all physical systems and processes. This of course includes complex systems, insofar as they are instantiated in the real world, versus simulated in silico.

The first law of thermodynamics states that the energy flow into a system is exactly balanced by the work performed by that system and the heat that it outputs (conservation of energy – it can be neither created nor destroyed). An intuitive interpretation of the second law is that closed systems – systems that do not exchange heat or matter with the external environment – tend towards thermodynamic equilibrium; more explicitly, that the entropy of a closed system always increases: $\Delta S > 0$, with reference to the initial and final entropy state of a process.

Most will be familiar with the first law; however, familiarity with and understanding of the second law has been limited by misuse, misapplication, or general misunderstanding of the term entropy itself, and specifically how it relates to physical processes. It is colloquially understood as “the amount of disorder”, but can be a difficult concept to grasp, and the easiest approach to an intuitive understanding may – paradoxically – be through statistical mechanics, rather than the classical thermodynamic approach relating heat and temperature, as formulated by Clausius. Without recapitulating the statistical mechanical derivation here, that approach concludes that the entropy of a system is proportional (by Boltzmann’s constant) to the natural logarithm of the number of quantum states the system may be found in.

It is a measure of disorder of the system, if by “disorder” we mean the difficulty of identifying or specifying the microscopic state of the system (the quantum states of all its particles), because of the large number of possible states. Thus, its entropy is a measure of uncertainty – the amount of information that must be obtained to know its complete microstate (rather than just its macroscopic properties such as temperature and volume). If we observe two systems, each with the same schedule of particles (that is, quantity vs. atomic element), but in the first we observe that some of the particles are found in crystalline form, we have some sense that there is more order in this system; conversely, the second has more entropy. We may characterize this difference in entropy as the additional (missing) information required to specify the states of all particles in the second system. Equivalently, we conclude that the first system exhibits stored information telling us something about its micostate: negative entropy, relative to the second. Thus, there is a direct relationship between the thermodynamic entropy and Shannon’s measure of information; see Tribus et al. [25] for an exposition on this topic.

In physical terms, a system at thermal equilibrium has reached a uniform temperature – the average kinetic energy of its particles is the same throughout; entropy is at a maximum (for the given temperature), and stored information is at a minimum. Generally speaking, such a system is inert, and uninteresting – and within the context of living organisms, certainly dead (though in fact organism death may occur long before thermal equilibrium – to wit, fossilized remains, which convey information to paleontologists).

2.2 Thermodynamics and Life

In his monograph "What is Life?", Schrodinger [18] proposes a thermodynamic-based definition of life: life "metabolizes" energy in its environment to maintain homeostasis away from thermodynamic equilibrium. Schrodinger argues (in Chapter 6) that the primary function of the metabolic processes of life is to avoid decay towards thermal equilibrium by incorporating negative entropy from an organism's environment, and rejecting its own entropy to the environment. We may say that an organism's metabolism is an engine: it consumes energy, turning some of that energy into work, and rejecting the rest as heat. The work that it does is replacing the entropy that builds up within the organism with negative entropy (which as we have seen, is information), and moving that entropy into the environment.

Life is, of course, a paramount example of a complex system (and is in particular a complex adaptive system). I believe that characterizing the metabolic processes of life in the language of thermodynamics is not only necessary – since life, as a physical process, is unequivocally governed by the laws of thermodynamics – but gives us a generic tool to use in reasoning about the nature of complex systems.

If we accept Schrodinger's characterization that avoidance of decay into thermal equilibrium is the prime function of life's metabolic processes, we may now point to entropy as the most generic "natural selection" force in the universe: the structures and processes that persist are precisely (by definition now) those that maintain negative entropy within to counterbalance – or even exceed – the constant onslaught of entropy. This further leads us to two insights: 1) we may now propose a formal definition of complex systems based on thermodynamics, and 2) we may expand our understanding of the term "metabolism" beyond the wetware-based biochemical processes that happen within the cell wall or in the intercellular network within an individual organism.

So the mere existence of a non-equilibrium system requires the following conditions:

1. The system is "open" – it is allowed (and must be able) to exchange energy and matter with its environment.
2. $\frac{dS_{internal}}{dt} \leq 0$. That is, entropy generation within the boundary is either steady-state, or decreasing. This is stated in differential form, versus difference form ($\varDelta S$), in explicit recognition of a continuous process (rather than a process with well-defined initial and final states).
3. $T_{external} < T_{internal}$. This is merely Carnot's stipulation that there must be a cooler temperature reservoir to dump heat (reject entropy) into, if work is to be accomplished.

The first condition is necessary because the 2nd law holds for closed systems, without caveat; only an open system allows condition 2, which is merely a restatement that complex systems are characterized by a persistence of negative entropy. The first condition allows a temperature gradient across which work may be performed, and the third condition establishes the direction of that gradient. Condition 2 is in direct contravention to the 2nd law – the apparent paradox of which was a subject of great debate subsequent to Clausius' formulation, and which Schrodinger devoted Chapters 6 and 7 in an attempt to resolve. The resolution of the paradox is quite simple, though, and is contained in conditions 1 and 3, allowing localized deviation from the 2nd law. Furthermore, it is precisely those instances where we perceive the most apparent contradiction to the 2nd law – the far-from-equilibrium cases – that we label as having the most complexity.

A direct consequence of these conditions is that there must be net positive entropy flux across the boundary (integrated over its surface), from the system into its environment. Note, however, that this does not necessitate that $\frac{dS_{external}}{dt} > \frac{dS_{internal}}{dt}$: complex systems may be nested within one another, and the external system may be rejecting entropy into yet another cooling reservoir (external system), proceeding so on until eventually the heat is – one hopes – radiated into space.

I would like to propose a corollary: for complex adaptive systems (CASs, such as life), $\frac{dS_{internal}}{dt} < 0$ (entropy generation within the system's boundary is strictly negative); that is, over time, such systems grow: in order, negative entropy, information – whichever term of equivalency you choose to characterize the increasing "complexity" of the system. An interesting feature of CASs is that a fractal organization follows from their adaptive nature: at the highest level, the earth's biome is a single CAS; the biome itself is composed of interacting ecosystems (coupled both physically, at their borders, and indirectly through the atmosphere), with ecosystems composed of interacting species populations, societies of individuals, the individuals themselves, etc. Much of the complexity of each of these sub-CASs exists as mutual information; in other words, the information they contain is information about the environment they reside in (including interaction with other CASs). Their adaptive nature means that all of the CASs, at all levels, are in constant readjustment to one another, constantly competing for the energy required to metabolize into negative entropy. True stasis – equilibrium, if not even thermal equilibrium – is a rarity in such an ecosystem, for that quickly leads to thermal equilibrium.

2.3 An expanded scope for "metabolism"

Since Schrodinger did not provide an explicit thermodynamc definition of the term metabolism, I will here summarize his arguments and characterizations into an explicit one: the metabolic processes of a system are those which capture energy and pre-existing negative entropy in the system's environment, using the energy available for work (the exergy, as established by the temperature gradient and the system's efficiency) to assimilate or generate negative entropy, and to reject the system's [positive] entropy into the environment. Traditionally, the system under consideration is understood to be, at most, an individual organism. But need it be restricted to such? The system referred to as an individual organism is generically a CAS, at a specific level in the entire, complex hierarchy that makes up the CAS of the earth's biome. Can we refer generically to the metabolism of a CAS?

Many organisms exist within the context of a society, and examples of social organisms abound across the domains of life: bacteria, protozoans (such as some forms of amoeba), fungi, insects, fish, birds, all sorts of mammals, and of course humans, and even examples of cross-species and cross-kingdom societies, such as lichen. When we refer to a society of individuals, much more than a mere grouping is implied: there is also a subcontext of a set of interactional behaviors between the individuals of that society, typically designated as cooperative, and unique to that society, which have evolved to help ensure "survival of the species".

And what is survival of the species but a persistence of negative entropy, a persistence of the information bound within its structures – the unbound energy from the sun, turned into the stored energy of chemical bonds? Metabolism, then, in a more general sense applied to the level of a society includes the processes used by a society – the entire set of social behaviors of its individuals – to capture the energy and information in its environment and use it to maintain, and even grow, its negative entropy.

This persistence of negative entropy is not merely the enhanced reproductive success of all its members; in many social insects for example, there are only one or a few reproductively active members. The individual members of these societies cannot long survive, or persist their genetic information, when separated from their society – and specifically, from the metabolic processes that that society, as a CAS in and of itself, utilizes to persist its negative entropy. Persistence of the individual's negative entropy is in many cases subordinated to and dominated by the society's.

In this context then, metabolism is no longer just an analogy of a society's processes to an organism's internal metabolic processes: there is direct equivalence in a real, thermodynamic sense. The society performs energetic transformations of entropy that are inaccessible to its individual members, in both form and scale; this is the thermodynamic perspective on the phenomena that complexity theorists often label as "emergence".

One may expand this discussion to include CASs composed of multiple societies of different types, cooperating (or competing) with each other, and coevolving, and so on, all the way up to the earth's biome itself as a CAS. All the processes of organisms that we consider "living", all of their varied levels of sophistication, all of the complexities of their metabolic processes, their cognitive abilities (where they exist), and their societies may be understood as a continual battle between continuity of information, and entropic decay to thermal equilibrium, all of it fueled (ultimately) by the solar flux. At some point in the distant past some very localized clumping of negative entropy happened to be able to persist, and eventually (or simultaneously) to grow, and the accumulation of negative entropy has been the story of life on earth ever since.

3 A formal description of metabolism

3.1 Thermodynamic stability

We may characterize the metabolic feature of life using the language of both thermodynamics and of information theory, and in fact both characterizations are equivalent. Thermodynamically we say that an entity is an open system that maintains $\frac{dS_{internal}}{dt} \leq 0$, when its metabolic process is functioning. The "open system" qualification allows local deviation from the 2nd law.

From an information theory perspective, we characterize the entity as a communication channel in time; it takes as input its current state σ^t and outputs its next state $\sigma^{t+\delta t}$ (σ^{t+}, hereafter). We denote the information carried by the current state as $H(\sigma^t)$, and by the next state $H(\sigma^{t+})$. The term mutual information I(X,Y) describes the average information each variable provides about the other; it is the information that is shared by the two variables. For a noisy channel with noise signal ν, $Y = X + \nu$ and $I(X,Y) = H(Y)–H(\nu)$ [24]. Thus, $I(\sigma^t, \sigma^{t+}) = H(\sigma^t)–H(\nu)$. The noise term ν reflects a change in the state configuration σ. For real-world systems, the term $-H(\nu)$ represents the thermodynamic entropy generated by the entity itself and incident upon it from the environment – a loss of information between successive states.

The minimum requirement for avoiding decay to thermal equilibrium is that $dS/dt = 0$, which is a stasis implying that $I(\sigma^t, \sigma^{t+}) = H(\sigma^t) = H(\sigma^{t+})$ (the second relation due to symmetry of mutual information) – in other words, the system does not lose information, and the information shared by successive states is maximized. Let us call this constraint, $dS/dt = 0$, "thermodynamic stability" - a steady-state in which neither thermodynamic entropy nor information accumulates. Since we are more interested in long-term behavior than instantaneous results (or discrete time results over a fixed $\varDelta t$), we may loosen this restriction somewhat by requiring only that the *average* entropy accumulation is steady state: $\overline{dS/dt} = 0$. Schneider and Kay [17] note that the notion of stability evaporates in real-world ecosystems, due to their inherent adaptivity and consequential diversity; however, it is useful as an idealized theoretical constraint to derive some minimal properties for complex adaptive systems.

Formally, we define the communication channel between states as a transformation, $\sigma^{t+} = f_\mu(\sigma^t)$. Note that while I have used functional notation here, f_μ is not necessarily a function in the strict mathematical sense; abstractly, it is a process that moves a system from state to state, and could as well be described by a state machine. For thermodynamic stability, we do not require that the states themselves are equal, but that the information measure is conserved (though if the states are equal, information conservation necessarily follows). Let us make some observations about f_μ.

Observation 1 An entropy-reduction function must exist.

At a minimum, f_μ must eliminate the introduction of entropic noise. First, it may prevent or reduce the introduction of an external noise signal ν_{ext}. This would come in the form of a physical boundary that isolates the system from the external environment, for example the cell wall. Unfortunately, we must acknowledge the dictates of the 2nd law: if the boundary completely isolates the system, internal entropy will still take its toll, $\overline{dS/dt} > 0$, and the entity decays. However, it may recover the lost information by rebuilding reject any entropy that is introduced; that is, it may "pump" the noise signal to its environment (thus, if a boundary exists, it must be at least semi-permeable to allow outward flux). Third, it may actively correct the informational error term $-H(\nu)$ introduced by the total noise signal. In other words, it rebuilds the information that is lost due to entropy. Fourth, it may "negate" the noise signal. We shall return later to these latter two methods.

Observation 2 An energy-collection function must exist.

No matter which method it uses for thermodynamic entropy elimination, f_μ does work, and as a physically realized machine, it must therefore consume energy. Also, as a physical machine (engine), with efficiency less than 100%, it generates entropy internally. Therefore f_μ performs at least two fundamental operations: 1) it must capture energy from its environment to activate entropy elimination, and 2) it must eliminate thermodynamic entropy. In other words f_μ is a composite function. Note that the first operation must harvest energy at a high enough rate to power at least both operations (i.e., including the harvesting and transport of energy), and the second must reject the entropy generated by both operations at a high enough rate to prevent its accumulation.

Taken together, these codependent operations make the composite nature of f_μ a non-trivial statement: the minimum possible system to maintain thermodynamic stability is a two node operator network, whose nodes are mutually dependent on each other to keep the network coherent. It is a simple economy of thermodynamic commerce: one node collects energy in exchange for which the other rejects entropy (see Figure 1). But indeed, we might identify that a simple form of emergence has occurred: a system behavior (dS/dt=0) has appeared from the mutually dependent interaction of two entities, neither of which can perform the system function – even scaled down – on its own.

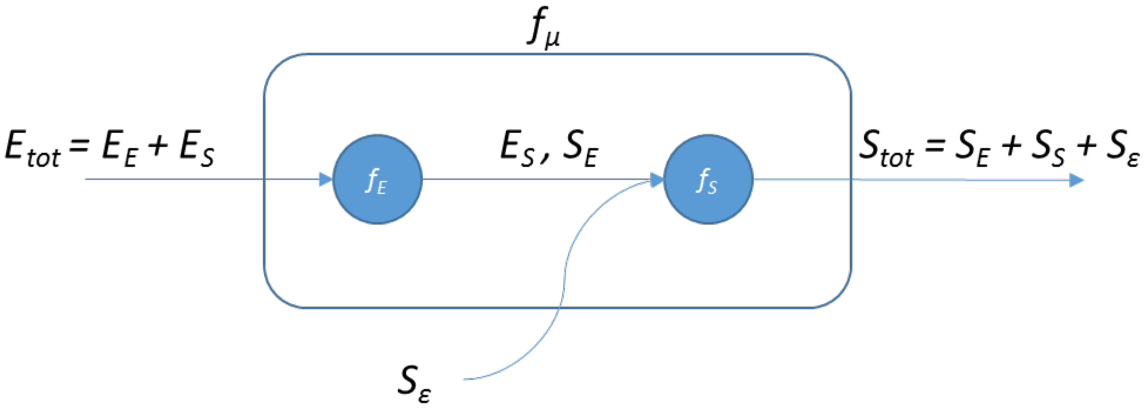

Figure 1: A 2-node network implementing the process f_μ maintaining thermodynamic stability. $S_{tot} = -H(\nu)$. f_E is the energy-collecting node, consuming E_E energy and generating entropy S_E. f_S is the entropy-eliminating node, consuming E_S energy and generating entropy S_S. S_ε represents the entropy flux from the environment, ε.

3.2 Repair Mechanisms

As noted before, the term $-H(\nu)$ represents the loss of system information that occurs during a metabolic transformation f_μ between times t and $t + \Delta t$. A thermodynamically stable system maintains $\overline{dS/dt} = 0$ by adding $H(\nu)$ information back in to the system, recovering the information that has been lost to entropy – in other words, to repair itself. In order to do this, all states $\sigma \in \Sigma_\mu$ (Σ_μ being the set of all states accessible to each other via f_μ) must contain within their physical instantiation a complete description of f_μ. This description $\langle f_\mu \rangle$ (the angle brackets denote "description of") must be complete because the noise signal is not guaranteed to be constant (even if $H(\nu)$ is) – things fall apart. If there were a description, say $\langle f_\mu{}^- \rangle$, that was not sufficient to reconstruct all of f_μ, eventually the part of f_μ corresponding to the missing description would degenerate, and the system fails.

In turn, f_μ, and specifically its subcomponent f_S (the information-rebuilding function), must be able to generate all components of f_μ from the given description. And so we have the classic definition of a recursive Turing machine [19] and the self-reproducing feature so characteristic of life. In the terminology of computational theory, a machine f_S is able to read an input tape ($\langle f_\mu \rangle \subset \sigma$) and construct another machine f_μ that can write an output state σ_{t+} that contains within it $\langle f_\mu \rangle$.

Note also that the 2-node-network of Figure 1 is brittle; in fact, it tolerates no failure ($-H(\nu)$) whatsoever: the erasure/degeneracy of either of its functions collapses the system altogether. Since these are physical processes that take time to complete, it must have multiple copies of f_E and f_S so that some minimum number remain operable after introduction of entropy. This is CAS' implementation of Shannon's noisy channel theorem.

In any case, thermodynamic stability is achieved when the recursion rate is equal to the entropy generation rate. In effect, this is the bare minimum of routine maintenance – the system rebuilds state configuration information at just the right rate to counterbalance entropy. However, some implicit – and overly restrictive – assumptions were made to hold $dS/dt = 0$. For example, the system must be able to diagnose the nature of information loss, requiring yet another function; additionally, it must be selective in its repair function, regenerating only those subfunctions that have degenerated. While these functions are quite commonly found in the more complex varieties of CAS in nature, they are quite specialized, with a non-trivial amount of rationality to them. It turns out that they are not necessary for our generative theory, and consequently cannot be expected to be universal. A much simpler strategy is unrestrained recursion; i.e., growth.

3.3 Growth mechanisms

We have already remarked on the necessity of redundancy if a system is to be fault-tolerant in the presence of entropy. That redundancy is most easily achieved by simply eliminating the requirement for selective diagnosis and repair (or selective reproduction of parts) – having already a mechanism (f_S) for repairing entropic damage to the basic metabolic process f_μ, simply churning out entire copies of f_μ requires less computational sophistication (Kolmogorov's algorithmic complexity [11] – the shortest length algorithmic description able to compute a given function) and should therefore be favored.

Since we are attempting to outline a fundamental theory of generative mechanisms for complex behaviors, we need not (should not) assume the *a priori* existence of more "complex" behaviors (specifically, more functions) where simpler ones (fewer functions) will suffice. Applying this version of Occam's Razor, then, the most general version of f_S is one whose sole function is to generate more copies of f_μ, or its components f_E and f_S. (Presently, I purposefully neglect discussion of whether f_μ is generated as a singular structural entity or as a network of structurally distinct but cooperating f_E and f_S entities,as this would delve too deeply into specific domains.)

We should note, however, that f_μ is a physical machine, and the information building component also requires more matter to completely reconstruct it, whereas a repair mechanism could operate over the material already present in the degraded f_μ. We could therefore add a third function to our network: a matter harvesting node. I neglect explicit consideration of this function, because – while a necessary component of growth in natural systems – it is not critical in constructing the general outline of the theory. So we assume for the present that the requisite material distribution is sufficient for f_S to perform its task.

Thus, the simplest version of the information building function repeatedly generates f_μ as energy is made available by the energy-harvesting component. And here we see in this basic formulation that entropy is acting as a natural selection force: in order to make copies of itself, f_μ must operate at a fast enough rate to compensate for $H(\nu)$. Presuming that the average information-generation rate exceeds the average entropy generation rate, the resulting growth is exponential in time, until the limits of available resources are reached. In the meantime, energy consumption will be exponential as well. Admittedly, whole-sale self-reproduction is a brute force strategy to counter entropy, but it is one that works, so long as the energy resource holds out.

Finally, we note that the presence of growth leads us to a result that the average entropy generation $\overline{dS/dt} < 0$. If the energy consumption of f_μ expands to the limit of the environmental availability, then its $\overline{dS/dt}$ will fluctuate around the energy rate of change (zero at steady-state).

3.4 Variation and adaptation

Up to this point, we have not introduced variation, and consequently adaptation. Also, the f_μ metabolic model has been analyzed in the context of an energetically (and materially) unconstrained environment. From the beginning however, we have characterized f_μ as a noisy communication channel, with input state σ_t and output state σ_{t+}. Formally, our model thus far:

$$f_\mu : \Sigma_\mu, -H(\nu) \mapsto \Sigma_\mu$$

and

$$f_S : \Sigma_\mu, -H(\nu) \mapsto f_\mu$$

I.e., the states $\sigma \in \Sigma_\mu$ could be depicted as a fully-connected directed graph / state machine with all edges f_μ, and as a generator f_S that could output f_μ from the common description $\langle f_\mu \rangle$ encoded within each state in Σ_μ. However, the entropy term was assumed to provide variation within the joint information $H(\Sigma_\mu)$ of the state-space; in other words, one stayed within the fully-connected portion of the graph. Two alternatives are possible: 1) a degraded transformation ${f_\mu}^-$ exits the state-space to a dead-end node (system collapse), or 2) the information building function is modified to some ${f_S}^*$ that maps to a new state transformation function ${f_\mu}^* : {\Sigma_\mu}^* \mapsto {\Sigma_\mu}^*$.

The new function ${f_\mu}^*$ either provides more entropy-compensation than f_μ, or it provides less: entropy is the universal fitness selector. If the latter, it could be the case that the new system's average entropy accumulation rate is now positive, and the system decays in time. If the former, it will grow at a faster rate than the previous system characterized by f_μ. In that case, it must also follow that energy consumption keeps apace with the growth rate.

One possibility for variation is a new energy-harvesting function ${f_E}^*$ that is able to exploit a hitherto untapped energy resource; a prime example would be the stored chemical energy in the structural bonds of some species (biological or chemical) extant in the environment. Of course, it may well be the case that this species is another CAS, with metabolic function $f_{\mu\prime}$, with energy bound as information in its state $\sigma\prime$. Thus, one entity's entropy is another's information. Regardless, we may characterize this exploitation as mutual information that the entity's component function ${f_E}^*$ has with respect to its environment: $I_{mutual}(\langle {f_E}^* \rangle, \varepsilon)$. Likewise, we may characterize an entity's survival strategies *against* exploitation by environmental threats as $I_{mutual}(\langle {f_S}^* \rangle, \varepsilon)$.

Since we placed no constraints on the form of ${f_\mu}^*$, we have allowed the possibilities that ${f_E}^*$ is itself a composite of one or more predecessor functions ${f_E}^i$ and likewise for ${f_S}^*$ – in other words, the building block modularity so commonly found in CAS [10]. Recognizing that within each one of these distinct ${f_E}^i$ and ${f_S}^i$ is potentially unique mutual information about the environment, we conclude that variations of ${f_\mu}^*$'s may successively build mutual information about the environment. That is to say, they have diversified.

And so with these generic mechanisms we have a recipe for multiple CAS ${f_\mu}^i$ competing with and adapting to each other *ad infinitum* in an environmental ensemble, with the potential to create all of the bewildering diversity of CAS observed on earth.

3.5 Metabolic Societies

There is another method by which we may achieve composite functionality: two entities ${f_\mu}^1$ and ${f_\mu}^2$ may cooperate. Refer to Figure 2. Here the two entities mutually benefit each other with the sharing of energy resources and information accumulation. What is of note in this relationship, versus a simple accumulation of entities, is that each of the metabolic entities is able to benefit from the other's unique information about the environment; the uniqueness of the mutual information could reside in either or both the energy-harvesting or the entropy-reducing functions of each entity.

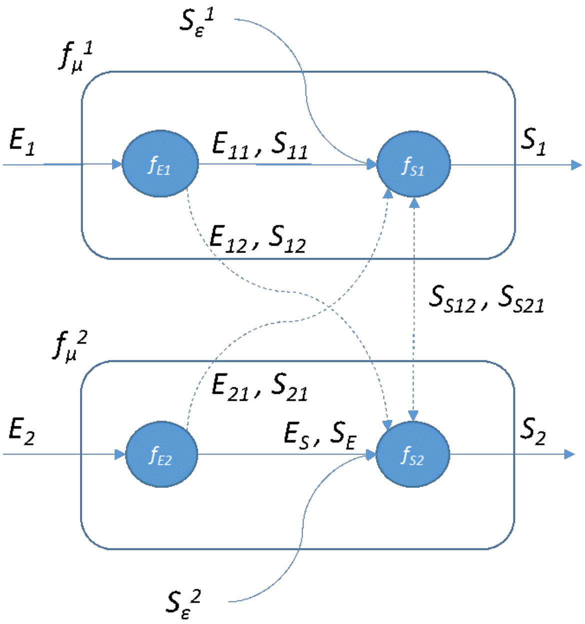

Figure 2: Cooperative Metabolism: Two distinct metabolic entities ${f_\mu}^1$ and ${f_\mu}^2$, sharing harvested energy resources and information accumulation. Each benefits the other with their respective unique mutual information with the environment. Dashed lines indicate possible inequality of the exchanged quantities.

This is a society of two: a composite metabolic process enabled by the set of relations between two distinct metabolic entities. Each subprocess (distinct metabolic entity) now has access to additional transformations between energy and information, that it does not have on its own. It can compensate for more entropy, and is thus more fit. Without loss of generality, we may scale this model along the dimension of f_μ, such that there is a set of distinct but cooperating metabolic processes $F_\mu = \{{f_\mu}^i | i \in \{1..N\}\}$, or along quantities of ${f_\mu}^i$ for a given i, to achieve societies of many individuals, all sharing the benefits of mutual information.

Figure 3 shows how a "society" of cooperating metabolic entities ${f_\mu}^i$ may be reduced to a composite metabolic entity ${f_\mu}^*$ identifiable at a higher level of hierarchy, and by the same token how a CAS at a higher level of hierarchy may be decomposed into its socially interacting constituent metabolic entities. Composites may be heaped on top of each other *ad infinitum*, to the limits of available resources. It's turtles *almost* all the way down – at least to the level of single-molecule machines.

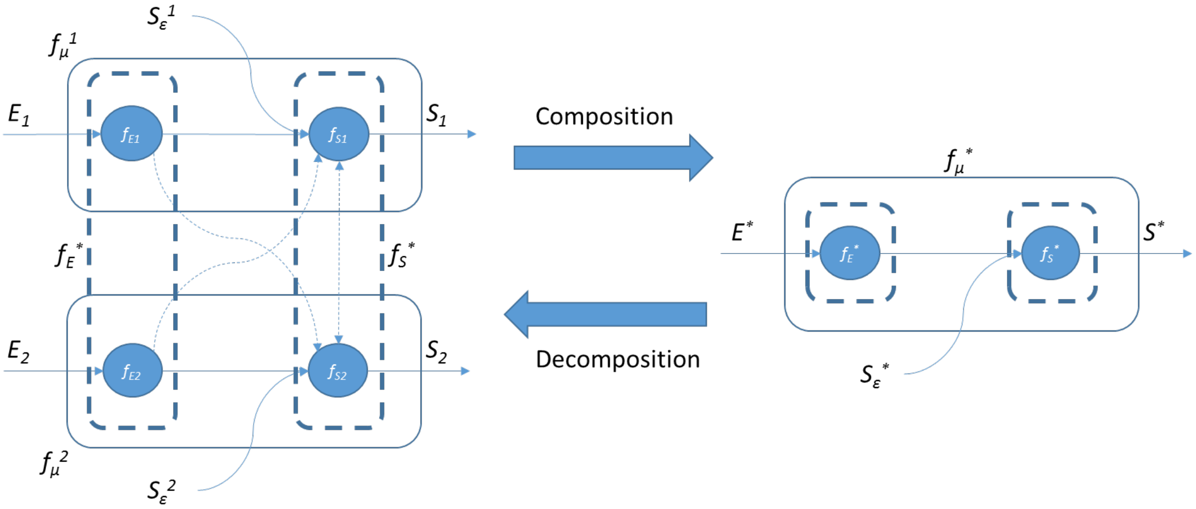

Figure 3: Metabolic Hierarchy: On the left is the constituent view of a CAS, and on the right, the composite. The cooperative metabolism of two or more distinct metabolic entities $f_\mu{}^i$ (constituent view) may be reduced to a single composite metabolic entity $f_\mu{}^*$ (composite view). Likewise, a composite metabolic entity may be decomposed into constituents $f_\mu{}^i$ that are socially interacting for mutual benefit. The dashed boxes group the functions that have collapsed or expanded between the constituent and the composite view. Differentiated $f_\mu{}^{*j}$ may cooperate to form the next higher level, etc.

With this "simple" model, we can generate some interesting scenarios found in nature – take the case of some of the social insects, ants and bees for example: a single individual, the queen, has a relatively low-activity energy-harvesting function, and her primary information-accumulation function is reproduction of all the metabolic entity types associated with the society. Her army of non-reproductive workers, however, is dedicated to gathering energy resources to support the queen's recursion habit, and building an entropy-reducing environment to protect it. The case of social insects is likely to be a scenario of metabolic divergence rather than cooperative convergence, but we placed no conditions on how the composition occurred.

We may also build models of extreme co-dependency, where the energy-providing and entropy-reducing/information accumulating functions are so tightly coupled between entities, and the mutual information contained in each is critical to both, that neither metabolic "entity" can endure without the other, and we may have trouble deciding whether to call it one entity, or multiple.

As a thought experiment, let us now build a physical boundary around an aggregation F_μ of such co-dependent entities (let's say that wall-building is the sole function of one of the $f_S{}^i$ in the F_μ ensemble), and concentrate the recursion function in another $f_S{}^j$ of the ensemble (like the social insects). At what point can we say that it is no longer a society of interacting metabolic processes, but a distinct organism – or the converse, if we traverse the hierarchy in the opposite direction? I could suggest that we leave the question for biologists to provide an academic answer to, but the question is inherently difficult because it crosses domains.

Which serves to illustrate my point: there is a continuum of societies of metabolic processes across many levels of CAS: intracellular, intercellular, the various systems comprising an organism, and of course the collective behaviors of multiple cooperating organisms, to include humans. A composite metabolic process of varieties of energy to information transformation networks is a useful general model for all of them.

3.6 Comparison to canonical description

This model incorporates Gell-Mann's descriptive model of CAS based on *schema* [8]. In the biological realm, the schema is the genotype; in the social realm, it is encapsulated in what we call "culture". In the generic metabolic model, it is the description $\langle f_\mu \rangle$ of the metabolic machinery of f_μ – a compact encoding of the structure, relations, and operations of f_μ for interacting with its environment to collect energy, accumulate information, and reject entropy; the subprocess f_S (or some portion thereof) "unpacks" the schema and instantiates it as the physical machine f_μ . Since f_μ interacts with the environment, its description $\langle f_\mu \rangle$ must perforce also contain mutual information about the environment.

I have commented throughout how the metabolic model accounts for many features commonly associated with CAS; Table 1 explicitly lists the seven basic features of Holland's framework [10] and the mechanism(s) or component(s) of the metabolic model that explain and generate them.

Table 1: Comparison to Holland's phenomenological regularities

Holland's feature	**Generative mechanism in the metabolic model**
Aggregation	Cooperative metabolism
Nonlinearity	Cooperative metabolism and co-dependency – the sharing of mutual information
Flows	Energy collecting functions and entropy-reducing functions.
Diversity	Variation during instantiation of f_μ and/or transcription of $\langle f_\mu \rangle$; selection through the filter of entropy
Tags	Mutual information and cooperative metabolism
Internal Models	Description of the metabolic machine $\langle f_\mu \rangle$; mutual information
Building Blocks	Energy collecting functions and entropy-reducing functions, and variations thereof; variations on metabolic entities themselves

4 Application Notes

In the interest of maintaining as abstract and general a model as possible, I have tried, throughout the formal description, to abstain from applying the model to real-world examples. Gell-Mann and Lloyd [9] have applied their analysis of informational measures of complexity to draw an interesting conclusion about the nature of scientific theory itself: formulating and describing a theoretical model that is representative of an ensemble comes at the expense of a reduction of the total information contained in the model. In other words, the more broadly applicable a model is, the less immediately useful it may be for application to a specific subject matter (such as a hierarchical domain of CAS); the coarse-graining of a representative description neglects the equations of motion that may be needed to describe the fine-grained resolution of a particular member of an ensemble. A model's simplicity is both its beauty and its burden, and that is the general approach I have taken in describing a generative model for the ensemble of CAS domains.

In this section I will try to bridge the gap between abstraction and the real-world with some brief commentary on applying the model to select case studies at opposite extremes of the CAS spectrum, and for which it may not be obvious that the model applies: 1) viruses, and 2) human social systems generally and economic systems in particular.

4.1 Viruses

Viruses are infectious agents that co-opt a host's cellular machinery for their own reproduction. They are not universally considered a form of life because they do not actively metabolize within their physical boundary. If they do not "actively" metabolize, then does the metabolic model of CAS still fit them? During their dormant, transmission stage it is true that they are not much more than some nucleic acid protected by a shell; they are consuming no energy and accumulating no information. However, the same is true for some species of bacteria, fungi, and even animals (the tardigrades) – during periods of environmental stress, they dessicate and form a thick outer protective shell, in which state they may remain dormant but reactivatable for surprisingly long periods even in extremely adverse conditions, until environmental conditions reactivate their metabolic process.

Viruses are often able to infiltrate by "tricking" the host's normal intrusion protections – a form of mutual information between the virus and its environs. Once a virus enters a cell (which it may do by releasing stored energy, such as the tobacco mosaic virus), we may then identify all the elements of the metabolic model. The virus' description $\langle {f_\mu}^{vir} \rangle$ contains sufficient mutual information with its host to reprogram the host's ${f_S}^{host}$ as a generator ${f_S}^{vir}$ for ${f_\mu}^{vir}$. The host's metabolic process now becomes the virus', and so the metabolic model certainly applies during the virus' active reproduction stage. From this perspective, the description $\langle {f_\mu}^{vir} \rangle$ "unpacks" into the totality of host cell machinery, ${f_\mu}^{host}$, and is thus perhaps one of the most compact information encodings that we can observe (it is of course, highly specialized in being reliant on select but critical mutual information with its host)!

4.2 Human social systems; economics

Some readers may argue that human social systems are inherently different, and question how this model could possibly apply to and account for the variety and complexity of interactions and structures in human society. Humans after all are purposeful – do we not shape our own destiny? Can our capacity for building cities and nation-states and all the great social institutions that have arisen across the globe and throughout history really be boiled down to notions as simple as extracting energy from the environment, and using it to reduce entropy? I submit that indeed they can, while at the same time acknowledging the awesome complexity of human social CAS – so awesome, and so seemingly qualitatively different, that it might appear to be governed by a fundamentally different model from the

CAS that are found in the biosphere and produced by the classical biological evolutionary paradigm.

I offer a simple explanation, within the metabolic model, for the astounding success of human social CAS – simple, but because of the recursive nature of CAS, having profound impact. That difference is this: the human capacity for conceptual modeling gives us the ability to generate new schema, new $\langle {f_\mu}^* \rangle$ – in a word, it is "imagination". We humans generate new ${f_\mu}^*$ at a much greatly accelerated rate, on the order of days, versus the years, centuries, and millenia of classical evolutionary adaptation. Furthermore, the rate at which we generate new ${f_\mu}^*$ is accelerating as we create technologies that assist us both with conceptualizing, and with instantiating, new ${f_\mu}^*$.

Let's examine this in a little more detail. We denote the metabolic entity of a single individual with ${f_I}^i$, where i is one of N individuals. In "flat" societies with little or no social organization or stratification (e.g., many hunter-gatherer societies), we may speak of a single social construct f_P that encapsulates the metabolic activities of its members, as in the single level of hierarchy schematically depicted in Figure 3. The operation of f_P is encoded in $\langle f_P \rangle$ – the customs, traditions, mores, and taboos that govern the relations between its members ${f_I}^i$. In small, flat societies, $\langle f_P \rangle$ is small enough to be internalized (as a mental schema) in its individuals, who use it to instantiate and regulate relations with others.

It may be the case that environmental stressors (or opportunities) warrant a transition from f_P to ${f_P}^*$ – a social reorganization. In such cases, the population must generate a new schema $\langle {f_P}^* \rangle$ and instantiate it. These reorganizations can be quite drastic and rapid (see Flannery and Marcus [7] for a survey of such cases). Since individuals are limited in time, energy, and access to resources, it may be the case that ${f_P}^*$ requires dedicated subprocesses ${f_{P'}}^{*i}$ to accomplish, as in the arbitrary social organization depicted in Figure 4.

In addition to new ways of organizing people and their labor activities (what Nelson [14] and Beinhocker [5] refer to as "social technologies") to increase the total information-accumulating activity of a society, humans also create physical technologies ([14], [5]) – the physical artifacts and tools that we typically associate with the word "technology". Invention is the process of conceptualizing and describing some machine – a $\langle f_T \rangle$ – and turning it into a physically realized product f_T. These technologies, if they are useful, extend our ability to collect energy from the environment or reduce the total entropy. It is to the engineering disciplines that we look to optimize a technology's function to its intended environment – its mutual information with the environment. Key technologies are occasionally invented that enable acceleration of new f_P: agriculture (meeting the energy requirements of a larger population), heat engines (using the stored energy of fossil fuels to more quickly produce physical technologies), universal computation (enabling the more rapid creation of much more complicated $\langle f_T \rangle$ and even testing of their operation in a simulated environment) are just a small set of examples.

Based on this metabolic model then, we may perceive that economic value is primarily derived from entropy-reducing functions instantiated in society. There is wide consensus that economic growth is inherently related to technological progress; in the metabolic model, economic growth through technical progress is internally consistent with – one might say, equivalent to – an entropy reduction theory of value. A full treatment of this concept is not practical here, and I will explore it at greater length in future work.

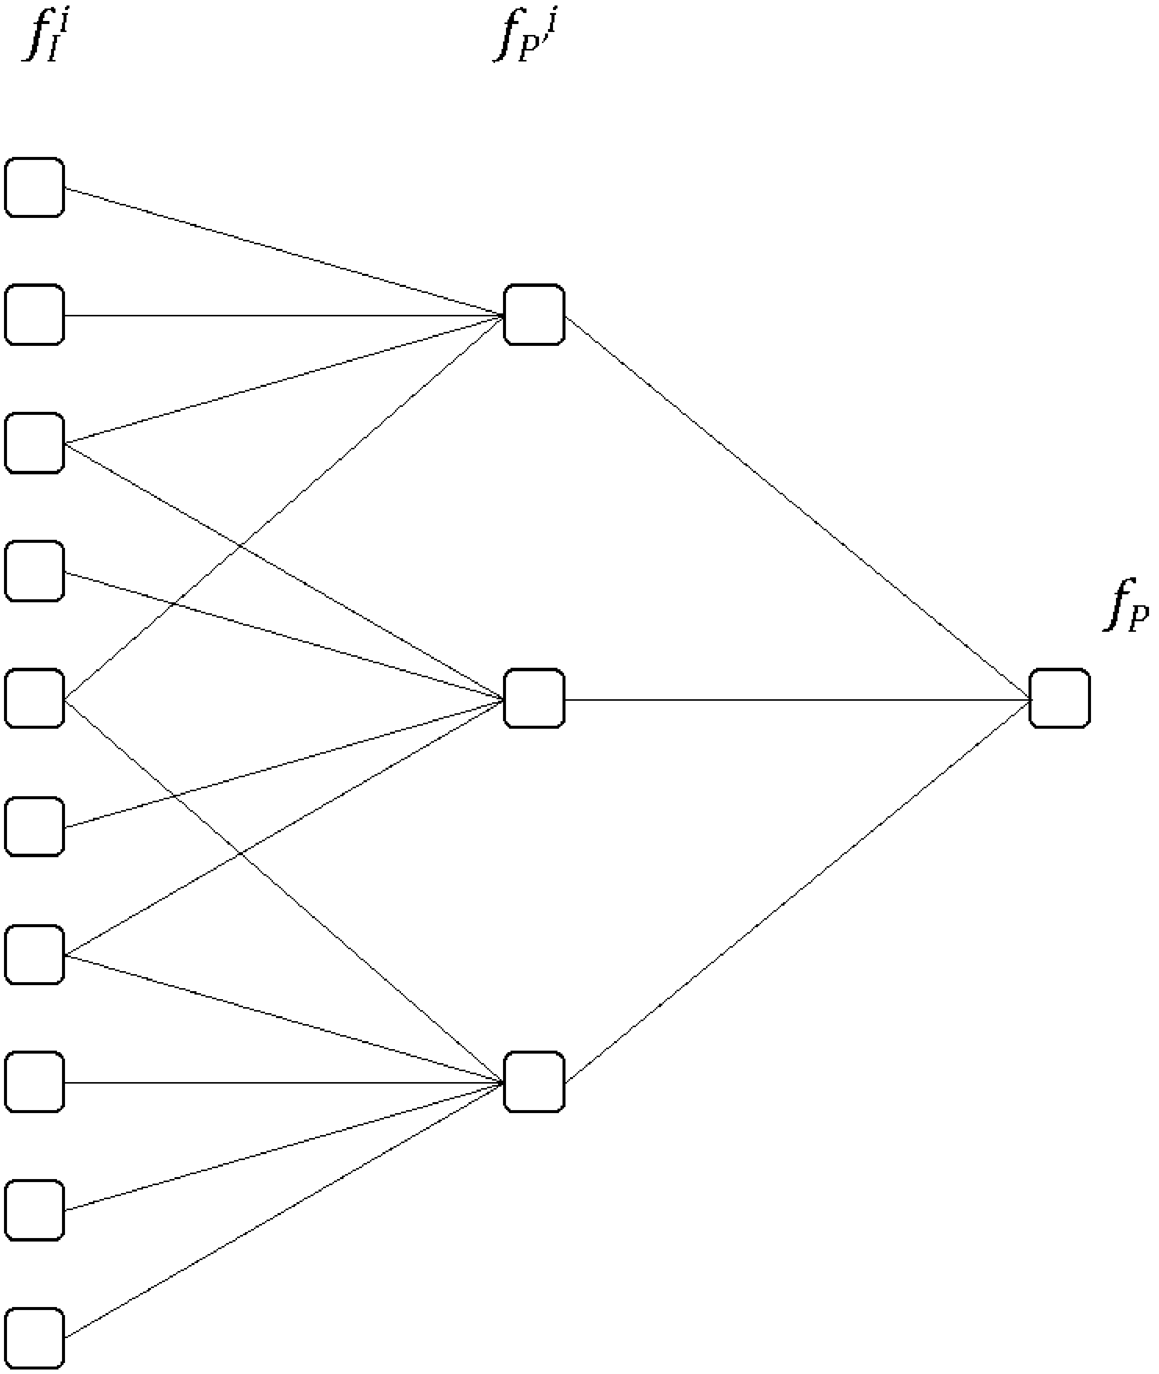

Figure 4: Multilevel Social CAS: Individuals ${f_I}^i$ form metabolic entities ${f_{P'}}^j$ (individuals may participate in multiple functions). These intermediate entities ${f_{P'}}^j$ compose the total social metabolic function f_P.

5 Conclusion

Metabolic activity, when described generically as a composite of thermodynamic and information theoretic functions, serves as a generative model for features common to all natural CAS when thermodynamic constraints are applied. Entropy is the most generic natural selector. Replication is required to repair degeneration due to entropy, and growth can simply result from “over correction” by repair mechanisms; replication also introduces recursion over the entire process. Variation (with entropy as possible cause) results in divergence of predecessor functions. Mutual information accumulates in metabolic entities, and characterizes the diversity, and – inherently – the adaptation associated with CAS. While conceptually a relatively simple model, its depth comes from recursion, and its breadth from variation.

Stated generically in the context of physical systems that must obey the laws of thermodynamics, the model applies not only to natural CAS found in the wild, but also to the “artificial” systems that humans build; i.e., our technologies may be seen as extensions (albeit directed and non-random) of the general metabolic model. Thus the metabolic model of CAS provides an internally consistent unification of economic growth and a theory of value based on entropy-reduction. In future work I will more fully explore application of the metabolic model to economic systems.

Finally, the reader who has studied complex systems may note that there is a phenomenological regularity of complex systems that I did not address: scale-invariance (i.e. power-law relationships: $x = y^a$, where a is a constant). This is a rich subject for further investigation. The discovery of a power-law relationship in empirical data is a tell-tale sign that a system is governed by non-equilibrium dynamics – that one is indeed observing complex phenomena. Having been observed in domain-specific empirical data, the power-law relationship is perforce also domain-specific, as are the theoretical equations of motion offered as generative mechanisms to explain them.

The author is not aware of a theory of generic mechanics that predicts how and where scale-invariance arises in specific domains; and it may be that such a theory for all complex systems is unachievable because of the “coarse-grained” nature of a universal theory [9]. If, however, the metabolic model is valid for all CAS, a strong validation for it would be the analytical or computational discovery of a power-law relationship generated by the basic model, and the ability to to use that relationship as an indicator for power-law relationships among variables in specific CAS domains. While not necessarily applicable to all complex systems, this would be an important advancement for the study and understanding of CAS, and should be the subject of more research.

References

[1] W Brian Arthur. “Complexity and the economy.” In: *science* 284.5411 (1999), pp. 107–109.

[2] Robert U Ayres. “Eco-thermodynamics: economics and the second law.” In: *Ecological economics* 26.2 (1998), pp. 189–209.

[3] Robert U Ayres, Leslie W Ayres, and Benjamin Warr. “Exergy, power and work in the US economy, 1900–1998.” In: *Energy* 28.3 (2003), pp. 219–273.

[4] Robert U Ayres and Jeroen CJM van den Bergh. "A theory of economic growth with material/energy resources and dematerialization: Interaction of three growth mechanisms." In: *Ecological Economics* 55.1 (2005), pp. 96–118.

[5] Eric D Beinhocker. *The origin of wealth: Evolution, complexity, and the radical remaking of economics*. Harvard Business Press, 2006.

[6] Roderick C Dewar et al. "Beyond the second law: an overview." In: *Beyond the Second Law*. Springer, 2014, pp. 3–27.

[7] Kent Flannery and Joyce Marcus. "The creation of inequality." In: *How our prehistoric ancestors set the stage for monarchy, slavery and empire* (2012).

[8] Murray Gell-Mann. "Complex adaptive systems." In: *SANTA FE INSTITUTE STUDIES IN THE SCIENCES OF COMPLEXITY-PROCEEDINGS VOLUME-*. Vol. 22. ADDISON-WESLEY PUBLISHING CO. 1995, pp. 11–11.

[9] Murray Gell-Mann and Seth Lloyd. "Information measures, effective complexity, and total information." In: *Complexity* 2.1 (1996), pp. 44–52.

[10] John Henry Holland. *Hidden order: How adaptation builds complexity*. Basic Books, 1995.

[11] Andrei N Kolmogorov. "Combinatorial foundations of information theory and the calculus of probabilities." In: *Russian Mathematical Surveys* 38.4 (1983), pp. 29–40.

[12] David C Krakauer. "Darwinian demons, evolutionary complexity, and information maximization." In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 21.3 (2011), p. 037110.

[13] Reiner Kümmel, Robert U Ayres, and Dietmar Lindenberger. "Thermodynamic laws, economic methods and the productive power of energy." In: *Journal of Non-Equilibrium Thermodynamics* 35.2 (2010), pp. 145–179.

[14] Richard R Nelson et al. "Physical and social technologies and their evolution." In: *Économie appliquée* 56.3 (2003), pp. 13–32.

[15] Georgescu-Roegen Nicholas. "The entropy law and the economic process." In: *Harvard Un* (1971).

[16] Eric D Schneider and James J Kay. "Complexity and thermodynamics: towards a new ecology." In: *Futures* 26.6 (1994), pp. 626–647.

[17] Eric D Schneider and James J Kay. "Life as a manifestation of the second law of thermodynamics." In: *Mathematical and computer modelling* 19.6-8 (1994), pp. 25–48.

[18] Erwin Schrodinger. *What is life? and, mind and matter, by erwin schrodinger (p)*. Cambridge University Press, 1967.

[19] Michael Sipser. *Introduction to the Theory of Computation*. Vol. 2. Thomson Course Technology Boston, 2006.

[20] Eric Smith. "Thermodynamics of natural selection I: Energy flow and the limits on organization." In: *Journal of theoretical biology* 252.2 (2008), pp. 185–197.

[21] Eric Smith. "Thermodynamics of natural selection II: Chemical Carnot cycles." In: *Journal of theoretical biology* 252.2 (2008), pp. 198–212.

[22] Eric Smith. "Thermodynamics of natural selection III: Landauer's principle in computation and chemistry." In: *Journal of theoretical biology* 252.2 (2008), pp. 213–220.

[23] Eric Smith and Duncan K Foley. "Classical thermodynamics and economic general equilibrium theory." In: *Journal of economic dynamics and control* 32.1 (2008), pp. 7–65.

[24] James V Stone. *Information theory: a tutorial introduction.* Sebtel Press, 2015.

[25] Myron Tribus and Edward C McIrvine. “Energy and information.” In: *Scientific American* 225.3 (1971), pp. 179–188.