



The direction of evolution: The rise of cooperative organization



John E. Stewart*

Evolution, Complexity and Cognition group, Center Leo Apostel, Vrije Universiteit Brussel, Krijgskundestraat 33, B-1160 Brussels, Belgium

ARTICLE INFO

Article history:

Received 16 January 2014
 Received in revised form 14 May 2014
 Accepted 16 May 2014
 Available online 2 June 2014

Keywords:

Direction of evolution
 Evolutionary progress
 Evolution of a global entity
 Evolution of cooperation
 Major evolutionary transitions

ABSTRACT

Two great trends are evident in the evolution of life on Earth: towards increasing diversification and towards increasing integration. Diversification has spread living processes across the planet, progressively increasing the range of environments and free energy sources exploited by life. Integration has proceeded through a stepwise process in which living entities at one level are integrated into cooperative groups that become larger-scale entities at the next level, and so on, producing cooperative organizations of increasing scale (for example, cooperative groups of simple cells gave rise to the more complex eukaryote cells, groups of these gave rise to multi-cellular organisms, and cooperative groups of these organisms produced animal societies). The trend towards increasing integration has continued during human evolution with the progressive increase in the scale of human groups and societies. The trends towards increasing diversification and integration are both driven by selection. An understanding of the trajectory and causal drivers of the trends suggests that they are likely to culminate in the emergence of a global entity. This entity would emerge from the integration of the living processes, matter, energy and technology of the planet into a global cooperative organization. Such an integration of the results of previous diversifications would enable the global entity to exploit the widest possible range of resources across the varied circumstances of the planet. This paper demonstrates that its case for directionality meets the tests and criticisms that have proven fatal to previous claims for directionality in evolution.

© 2014 The Authors. Published by Elsevier Ireland Ltd.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

1. Introduction

Is there an overall direction to evolution that is driven by selection? Do evolutionary processes drive the evolution of life in a particular direction? Is evolution headed somewhere?

Various attempts have been made to answer these questions by demonstrating the existence of large-scale directional patterns in evolution. None have yet attracted widespread acceptance.

The hypothesis that seems to have gained most support is that selection tends to drive increasing complexity as evolution proceeds (see [McShea, 1991](#) for an overview). Intuitively, this seems to be a plausible claim. There are many instances of increases in complexity during the evolution of life on Earth. However, strong arguments have been mounted against the claim that this apparent trend is driven by selection that directly favours increased complexity ([Gould, 1996](#); [McShea, 1991, 1994, 1996](#)).

In particular, it is obvious that complexity *per se* is not favoured by selection. There are numerous possible changes in organisms

that would increase complexity but are not advantageous in evolutionary terms. And changes that are less complex are not always inferior.

Compounding this difficulty, proponents of this claim have been unable to identify how known evolutionary processes would drive the supposed trend towards increasing complexity. This is a serious deficiency that also bedevils other attempts to demonstrate an overall, driven trend in evolution. To demonstrate such a trend, it is not sufficient to identify some supposed large-scale pattern in evolution and to marshal empirical evidence that substantiates the existence of the pattern. The pattern may be an artefact and not driven by selection that directly favours the pattern. It is therefore also necessary to provide the claimed directionality with micro-foundations at the level of natural selection that show how the pattern is driven by selection and related processes.

This has proven particularly challenging because it is not at all obvious how natural selection could drive a trajectory encompassing all living processes, given that it produces only local adaptation to local circumstances ([Gould, 1996](#); [Maynard Smith, 1988](#)).

This deficiency obviously cannot be overcome by the postulation of some new general 'force', 'tendency' or 'drive' that is

* Tel.: +61 3 96637959/422540984.

E-mail address: future.evolution@gmail.com (J.E. Stewart).

unsupported by appropriate micro-foundations. Nor can it be overcome by teleological explanations that rely on impermissible ‘pulls from the future’.

The weakness of the ‘complexity’ hypothesis is not only due to the absence of a convincing micro-foundational model that demonstrates how accepted evolutionary processes drive increased complexity. It has been further undermined by the demonstrations that show how the apparent pattern of increasing complexity could emerge passively – i.e. that show how complexity could increase without it being actively favoured by selection in any overall sense.

Gould (1996, 1988) and McShea (1994) have shown that complexity would be expected to increase as evolution unfolds merely as a consequence of the fact that the first living processes were necessarily simple. Life had nowhere else to go but to become more complex. Evolution necessarily began with the exploration of the possibility space which encompasses simple forms of life. It necessarily continued after that with the progressive exploration of the possibility space for living processes of greater and greater complexity. So as evolution proceeded, more complex forms of life emerged progressively, giving the appearance of a trend. Because life began as simply as possible, there was no countervailing exploration of more simple forms.

This model facilitates the identification of a key test for distinguishing an overall trend that is driven from one that is passive (in the sense that it is not driven directly by selection that favours the trend). If the trend is passive and there is no overall advantage to complexity, simpler organisms will not necessarily be replaced as evolution proceeds. In contrast, if the trend is universal and driven directly by selection, all organisms in all niches could be expected to evolve in the direction of the trajectory, or be replaced by others that do so. If complex organisms are favoured over the less complex overall, simpler organisms cannot be expected to continue to persist.

As Gould (1996) points out in detail, the complexity hypothesis fails this critical test. Relatively simple organisms such as bacteria are still abundant. For extended periods of time, they have not increased in complexity, nor have they been replaced by ones that have.

This test also proves fatal for other claims that evolution exhibits an overall trajectory that is driven by selection that directly favours the trend. These include claims that put humanity at the tip of the evolutionary trajectory (see Ruse, 1996 for examples). Such claims typically hypothesise that evolution favours attributes that are most highly developed in humans (e.g. intelligence, adaptability). The continued existence of huge numbers of organisms that exhibit less-developed forms of these attributes tends to suggest that they are not favoured in any overall or universal sense by selection.

This paper presents a case for directionality in evolution that does not suffer from the deficiencies that have undermined other claims.

It should be emphasized from the outset that the claim outlined here is made in relation to the evolution of all living processes on Earth, including humans and human organizations. As the paper will demonstrate, the trajectory of evolution can only be properly understood if the evolution of all living processes is taken into account, ultimately as a whole. In particular, the full nature of the trajectory cannot be identified and understood by focusing on, for example, only biological evolution. As we shall see, human cultural evolution and the evolution of human organisations and technology (including artificial intelligence) play a critically important role in driving the trajectory beyond a certain point. The nature of the evolutionary mechanisms that explore possibility space prove to be far less important in driving the trajectory than is the structure of the possibility space. In particular, the trajectory is shaped primarily by the nature and location of evolutionary attractors in possibility space.

Section 2 of the paper begins to outline the case for a particular form of overall directionality. It identifies a large-scale pattern that is evident in the evolution of life on Earth.

Section 3 provides the pattern with micro-foundations by presenting a model which demonstrates that this pattern is driven by natural selection and other accepted evolutionary processes.

Section 4 subjects the model and its key predictions to appropriate tests, including those that have been failed by other claims for overall directionality.

Section 5 concludes the paper by providing an overview of the trajectory of evolution and discussing some of its key implications.

2. A large-scale pattern

If we stand back from the evolution of life on Earth and view it as a whole, a number of patterns are apparent.

2.1. Diversification

An obvious trend is that living processes have diversified as evolution proceeded. When life first began on Earth, it was limited to exploiting only a tiny proportion of available free energy sources under a very restricted range of environmental conditions. From there living processes have diversified progressively as evolution unfolded, spreading across the planet, adapting to an ever-widening range of environmental conditions and exploiting more and more sources of free energy. This trend towards increasing diversification has continued up until the present with the emergence of humans, albeit now mainly through the processes of cultural evolution, rather than through gene-based adaptation and speciation.

2.2. Integration into large-scale cooperative entities

But a less obvious trend that moves in a very different direction is also apparent. As well as the trajectory towards increasing diversification, there is also a trend towards increasing integration. As the evolution of life on Earth has unfolded, living processes have increasingly come to be integrated into cooperative organizations of larger and larger scale.

This has unfolded through a stepwise process. It began with the integration of self-reproducing molecular processes into cooperative organizations that became the first simple cells. These simple cells adaptively radiated, exploiting fresh sources of free energy and displacing proto-living processes from the sources they utilized.

But the process of integration did not end there. Cooperative organizations of simple cells eventually formed the more complex eukaryote cell (Margulis, 1981). Diversification of eukaryote cells by adaptive radiation was in turn followed by the integration of some of these cells into larger-scale cooperatives, eventually emerging as multi-cellular organisms. In a further repetition of this process, cooperative organizations of multi-cellular organisms produced animal societies.

This broad pattern in biological evolution has been described in one form or another by a number of researchers, including Corning (1983), Maynard Smith and Szathmáry (1995), Stewart (1995, 1997b, 2000), Turchin (1977) and Vidal (2014) (see Marcot and McShea, 2007; McShea, 2001 for more historical detail). Often different researchers have included different evolutionary events and transitions in the patterns they describe. This is because the causal factors that shape this pattern are as yet poorly understood. Researchers have therefore lacked a sound basis for deciding which particular evolutionary events fall within the pattern (e.g. see the criticisms by McShea and Simpson (2011) of the classification of major evolutionary transitions developed by Maynard Smith and Szathmáry (1995)).

This progressive integration of organisms into cooperative organizations of increasing scale is not limited to evolution driven by gene-based natural selection. The trend has continued in human evolution where cultural evolutionary processes now predominate: small kin groups were integrated into bands, bands were integrated into tribes, these formed the constituents of kingdoms and city states, and these in turn have been integrated into nation states (Stewart, 2000; Turchin, 1977).

At each step in this process of integration, smaller-scale entities are integrated into cooperative organizations that become larger-scale entities at the next level of organization. Typically the larger-scale entities undergo a relatively rapid diversification and adaptive radiation (e.g. see Knoll and Bambach, 2000). As evolution unfolds, this step-wise process repeats itself, producing cooperative organizations of living processes of greater and greater scale. At each step, a new level of nesting of entities within larger-scale entities arises. And as evolution proceeds, entities with greater levels of nestedness emerge. The result is the familiar nested-hierarchical structure of living processes. For example, a nation such as the US is an organization of states which are in turn organizations of local government entities, which are in turn organizations of families and other small groups, which are in turn organizations of individual humans, which are in turn organizations of eukaryote cells, which in turn are organizations of descendants of simpler prokaryote cells, and so on.

Because this process of integration has continued to repeat itself, and because its products have out-competed some earlier entities and also exploited fresh resources, an increasing proportion of life on Earth has come to be integrated into cooperative organizations as evolution has unfolded.

Like the trend towards diversification, the trend towards integration seems to be continuing apace at present. Humans are increasingly integrating other living processes into its organizations through activities such as farming, aquaculture and broader ecosystem management. And human organization itself seems likely to continue to increase in scale. Although rudimentary, the League of Nations and the United Nations were early attempts to build supra-national organizations on a global scale. Some forms of economic organization are already global, and regional cooperatives of nation states such as the European Union have emerged. Global crises such as human-induced climate change seem to be increasingly evoking coordinated responses across nation states. The idea that some form of global governance is essential for human survival and flourishing is now strongly supported by many leading international relations researchers (e.g. see Craig, 2008) and economists (e.g. see Walker et al., 2009).

Since life began on Earth, this process of integration has substantially increased the scale of cooperative organization. Three thousand million years ago, cooperation extended only between molecular processes that were separated by about a millionth of a meter, the scale of early cells. Now, some forms of cooperation extend between human organisms on the scale of the planet. Since life began on Earth, the scale of cooperation between living processes has increased by a factor of greater than 10^{13} .

It seems possible that humans and other living processes will increasingly be integrated into a cooperative entity on the scale of the planet (see Heylighen, 2007 for a history of ideas about a global superorganism). If this occurs it will unite the two great trends in the evolution of life on Earth, producing unity in diversity on the scale of the planet, and maintaining that unity as entities within the global entity continue to differentiate and diversify (Stewart, 1995, 1997a,b).

3. What produces the pattern?

What is the nature of the apparent trend towards increasing integration? Is it driven by selection that directly favours effective

integration, or is it largely passive? To what extent is the trend produced and shaped by normal processes of evolution, such as gene-based natural selection and cultural evolution?

I will set out to show that the apparent trend is shaped by two major factors:

- 1) Cooperation between living entities can produce significant fitness benefits.
- 2) Despite these potential benefits, cooperation does not evolve easily (“the cooperation problem”). If a group of evolving entities are to realize the potential benefits of cooperation, a complex and specific form of organization must first arise within the group.

3.1. The benefits of cooperation

That cooperation can produce significant fitness advantages is not controversial (e.g. see Corning, 1983; Dugatkin, 1999; Miller, 1978; Ridley, 1996; Stewart, 2000). Cooperative organizations have the potential to be more successful than isolated individuals. Whatever the evolutionary challenges, living processes can respond to them more effectively if they form cooperative organizations and if their actions are coordinated. In part this is because cooperation enables the exploitation of synergies, including through specialization and division of labour (Corning, 1983). Furthermore, the larger the scale of any cooperative organization, the more resources are commanded by it, the greater its power, the larger the impact and scale of its actions, the greater the potential for collective adaptation and intelligence, and therefore the wider the range of environmental challenges that it can meet. Cooperation also saves resources that would otherwise be dissipated by competition (e.g. it can eradicate much agonistic behaviour, including war amongst humans). Wherever impediments to the evolution of cooperation have been overcome to some extent, cooperation flourishes (e.g. amongst the cells of multicellular organisms, or amongst economic agents in an effective market for goods and services (Knoll and Bambach, 2000)). The potential benefits of cooperation are a powerful driver at all levels of organization, whether we are considering self-replicating molecular processes, complex cells, organisms, corporations or nations.

It is evident that cooperation has the potential to produce advantages for a great diversity of living entities in varied local circumstances, irrespective of the nature of the mechanisms that evolve or adapt them. The ways in which cooperation is able to realize these advantages do not appear to be limited to any particular kinds of entities or any particular environments.

3.2. The cooperation problem

But cooperation does not evolve easily (e.g. see Boyd and Richerson, 2005; Buss, 1987; Olson, 1965; Williams, 1966). The reasons for this are well understood. Consider a population of living entities that compete for limited resources. Entities that invest resources in beneficial cooperation but fail to capture sufficient benefits of that cooperation, will tend to be out-competed. Other entities that take the benefits of cooperation without investing in the cooperation (free riders) will tend to do better than co-operators. Free riders undermine the capacity of co-operators to capture the benefits created by their cooperation.

In these circumstances cooperation will not persist in the population no matter how much its benefits exceed its costs. This is a very general impediment. It applies at all the levels of organization, under whatever selective processes govern adaptation at any particular level. It applies whether the interacting entities are members of the same species or of different species. And it also applies to entities that adapt through psychological or other adaptive processes that have been shaped by selection. So it

applies for example, to populations of self-reproducing molecular process, simple prokaryote cells, eukaryote cells, multicellular organisms, the members of ecosystems, human tribes, human economic entities (including corporations), and nation states.

If co-operators within a group of entities were able to capture all the benefits of their cooperation, cooperative organization would self-organize (in more general terms, the cooperation problem would be solved comprehensively if all the entities in a group capture the impacts of their actions on the group as a whole, whether the impacts are beneficial or harmful) (Stewart, 2000). Cooperation in which the benefits to the individual exceed the costs to the individual would be selected at the individual level (unless some alternative, more effective cooperation emerged). If this fundamental condition for cooperative self-organization were met, individual entities that engage in cooperation would out-compete non-co-operators. It would be in the evolutionary/adaptive interests of individuals to cooperate. As a consequence, the group would be able to explore the possibility space for cooperative organization, and any form of cooperation that was discovered and which was more advantageous would be able to persist in the population. Where it was beneficial, the kind of complex division of labour that is found amongst the cells within our bodies and amongst businesses in human economic systems would self-organize.

3.3. Limited emergence of cooperation

But complex cooperative organization will not self-organize in this way where the cooperation problem operates. Where the problem applies, co-operators do not capture all the benefits of their cooperation because some leak to free-riders (and free-riders do not capture the harmful effects they have on the group). In these circumstances complex cooperative organization will not self-organize, no matter what advantages it might provide to the group, and no matter how strong any selection that favours it. Selection cannot call cooperative organization into existence if that organization cannot persist within the population.

Nevertheless, are there circumstances in which co-operators can capture sufficient of the benefits of cooperation to enable some simpler forms of cooperation to persist?

Co-operators will capture proportionately more of the benefits of cooperation if they interact cooperatively with other co-operators more often than if all cooperative interactions are random. This will ensure that the benefits of cooperation are more likely to be shared amongst co-operators than leak to free-riders. If this condition is met, co-operators will capture a disproportionate share of the benefits of cooperation, and may capture sufficient to outweigh the costs of cooperation and the benefits that 'leak' to free riders. To the extent that this condition is met, co-operators will be collectively autocatalytic (they will collectively facilitate each other's success), and cooperative organization will be able to persist and be a target of selection (Ulanowicz, 2009).

It is conceivable that this condition could be met stochastically at times in a population. But it is likely to be met far more reliably if the cooperative interactions within the population are biased in some way. A huge literature has explored the specific kinds of circumstances in which some form of cooperation can arise in a population of competing entities (for overviews of particular parts of this literature, see Nowak, 2006; Perc and Szolnoki, 2010; Sigmund, 2010). In general, these circumstances can be understood as particular instances where constraints or other sources of bias increase the likelihood that co-operators interact cooperatively with other co-operators. This will enable co-operators to increase the proportion of cooperative benefits they capture beyond their representation in the population.

Two main ways in which this bias can occur are:

- 1) Population structure: cooperative interactions may be biased because the population of entities is structured in ways that increase the likelihood that co-operators interact with other co-operators. For example, low dispersal rates may concentrate co-operators and constrain the capacity of free riders to move freely throughout the population and exploit local concentrations of co-operators (e.g. entities might be fixed to a two dimensional surface for part of their life cycle (Nadell et al., 2010), or offspring may stay close to their parents (Griffin et al., 2004; Hamilton, 1964)). Or the population may be formed into groups that tend to concentrate co-operators and restrict invasion by free-riders (e.g. Okasha, 2006).
- 2) Active selection: interactions may also be biased because co-operators selectively choose to interact with other entities that are more likely to be co-operators (conversely, they may also selectively exclude or punish entities that are more likely to be non-co-operators). For example: entities may direct their cooperation towards others that carry particular markers that are closely correlated with a pre-disposition towards cooperation (the green beard effect, and some forms of kin selection in which individuals can discriminate in favour of kin by using cues that are correlated with relatedness (e.g. Keller and Ross, 1998; Mateo, 2003); entities that have the capacity to identify other individuals and assess their behaviour may use knowledge from repeated interactions to identify co-operators and to selectively direct their cooperation towards them (reciprocal altruism and some forms of indirect reciprocity (e.g. Roberts, 2008; Trivers, 1971); entities may also identify free-riders and decline to cooperate with them and/or punish them (strong reciprocity and policing (e.g. Boyd et al., 2005; Fehr and Gächter, 2002; Mouden et al., 2010)); and entities may be constrained in ways that enable them to cooperate only with other co-operators (autocatalytic sets of proteins and/or RNA in which some members are highly specific catalysts (e.g. Bagley and Farmer, 1991).

At every level of organization, living entities will tend to evolve some forms of cooperation through the operation of these kinds of processes. But if these are the only mechanisms that operate, co-operators will not generally be able to capture the full benefits of their cooperation. These mechanisms do not include processes that comprehensively prevent free-riders from undermining cooperation (including preventing cheats that imitate co-operators and therefore attract cooperation from selective co-operators, but do not themselves contribute to cooperation).

In summary, these processes can help cooperation get started at any level. But they are not responsible for the major steps in the trajectory of evolution towards increasing integration. They are incapable of producing the highly organized cooperative groups which exhibit a complex division of labour and that eventually become larger-scale entities at the next level of organization.

3.4. Comprehensive solutions to the cooperation problem

If the cooperation problem is to be overcome comprehensively, free-riding must be prevented, and as far as possible, the benefits of cooperation must go to the co-operators that create them. If this is to be achieved consistently and comprehensively in relation to a group of entities, special arrangements that have three key characteristics need to be in place (Stewart, 2000):

- 1) Power: the arrangements must have power over the entities in the group (including over co-operators and free riders), and the power to re-distribute the benefits of cooperation amongst

members of the group in favour of co-operators. Power means the ability to influence or constrain without being influenced in return. If the arrangements could be influenced in return by those they need to control, control would break down. For example, free riders would be able to escape effective suppression by the arrangements.

- 2) **Evolvability:** the arrangements must be evolvable/adaptable. This enables the arrangements to explore the space of possibilities for suppressing free-riders and for supporting beneficial cooperation. It gives the arrangements the capacity to optimize their use of power over entities, and to adapt their control as free-riders and other non-co-operators evolve and adapt to escape their control.
- 3) **Alignment of interests:** the evolutionary/adaptive interests of the arrangements must be aligned with the overall evolutionary/adaptive interests of the group of entities that it manages. Evolvability/adaptability *per se* is not enough. Unless interests are aligned in this way, the arrangements will not necessarily evolve/adapt in the direction needed to solve the cooperation problem. They will not necessarily use their power and evolvability to suppress free-riders and to support cooperation.

When we consider cooperative groups at a scale we are familiar with (i.e. at the scale of human organization), it is not difficult to identify arrangements of this kind that organize cooperative groups. Three examples are:

- 1) The CEO of a corporation uses power over employees to suppress free-riding and to reward actions that contribute positively to the corporation.
- 2) The government of a modern state uses its power to establish and operate a police force and legal system that punishes theft and fraud, enforces contracts between economic actors, and collects taxes that can be used to reward the provision of public goods and other cooperative actions (e.g. national defence and education). A democratic method of choosing the government can go some way towards aligning the adaptive interests of the government with the interests of the society as a whole. Strong competition between societies can have a similar effect.
- 3) A farmer uses his power over plants and animals on his farm to create conditions favourable to the reproduction and growth of those species that contribute to the economic success of the farm. And the farmer excludes or destroys species that consume useful resources but have no economic value themselves.

An understanding of these examples helps to identify cases in which similar kinds of management arrangements emerged during biological evolution to organize cooperative groups, including with the origin of life itself. For example:

- 1) Self-reproducing RNA developed the power to manage a proto-metabolism of simpler organic and inorganic molecules (Dyson, 1985; Stewart, 1995, 1997a, 2000). The emergence of proto-metabolisms began with cooperation amongst molecular species in the form of collective autocatalysis (Kauffman, 1993). Molecules cooperated with each other in autocatalytic cycles and autocatalytic sets to promote each other's formation and to thereby reproduce collectively. But the evolutionary possibility space that could be explored by these cooperative arrangements was very limited. In large part this was because molecular species that could make a significant contribution to the effectiveness of a proto-metabolism were often not themselves reproduced by any autocatalytic cycle or autocatalytic set in the proto-metabolism. These molecules could not

therefore capture the benefits of their contribution. So despite their usefulness to the proto-metabolism, they could not persist in the proto-metabolism and could not be the target of selection.

The emergence of RNA managers changed this radically. RNA could use its catalytic capacity to support molecular species and processes in the proto-metabolism that increased the fitness of the RNA and of the proto-cell as a whole, and to suppress processes that undermined it, including free-riders (Stewart, 1997a, 2000).

The emergence of RNA as an evolvable manager whose evolutionary interests were aligned with those of the proto-metabolism it managed was arguably the most important step in the origin of life. By overcoming the cooperation problem within the proto-metabolism, it enabled the realization of an enormous range of cooperative relationships within the metabolism that could never have been explored in the absence of an evolvable manager. It massively expanded the possibility space that could be accessed, paving the way for open-ended evolution, and moving significantly beyond what was possible through spontaneous chemical reactions and processes. Without management, proto-metabolisms could not generate the range of sustainable variation that is the essential raw material of evolution. Without sustainable variation to operate on, selection is incapable of producing evolution. Unmanaged autocatalytic cycles and sets have very limited capacity to evolve (Hordijk et al., 2012). Their possibilities are circumscribed by the dictates of chemistry. Although they can be self-reproducing, they fall on the abiotic side of the divide between living and non-living. The subsequent emergence of DNA as the manager of both RNA and the proto-cell as a whole was a further important step in the transition from chemistry to life.

- 2) A similar form of management was also critical in the emergence of the eukaryote cell. The precursors of cellular organelles such as mitochondria and chloroplasts were bacteria that were engulfed by a larger cell (Margulis, 1981). Because the larger cell had control over the environment of the engulfed bacteria, it had the power to evolvably manage the bacteria. It could constrain the bacteria to suppress competition amongst them that was against the interests of the cell, promote actions that advantaged the cell as a whole, and suppress free riders (Maynard Smith and Szathmáry, 1995). It could also constrain the evolvability of the engulfed bacteria to reduce the possibility that they would escape control by the larger cell (for example by transferring some of their DNA to the chromosomes of the larger cell).

It is useful to classify the constraints applied by management processes into two categories, although the categories represent extremes on a continuum:

- 1) **Prescriptive constraints:** these specify more or less precisely the particular outcomes that occur in the managed group. For example, DNA determines the specific proteins that are produced in a cell, including the quantities. And in a human command economy, the central authority prescribes specific economic outcomes, such as the nature and volume of the consumer goods that are to be produced. Where constraints are prescriptive, evolvability resides primarily in the manager, not in the other entities in the group.
- 2) **Enabling constraints:** these achieve outcomes that are best for the group without specifying what those outcomes are. They accomplish this by aligning the interests of group members with the interests of the group as a whole, and then letting the entities adapt freely in pursuit of those aligned interests. So

when managed entities pursue their own local self-interest, they adapt in the interests of the group. For example in a human economic market, the system of governance constrains theft and cheating on contracts, thereby ensuring that the only way economic agents can pursue their own interests is by engaging in fair market exchanges. In this way, enabling constraints do not suppress or override the evolvability of members of the group, and in fact harness it for the benefit of the group.

Of course, enabling constraints became more effective once evolution produced entities that were highly evolvable and capable of pursuing their own interests adaptively. The potential advantages of having evolvable entities managed by enabling rather than prescriptive constraints were demonstrated by the competitive superiority of free market economies over command economies in the 20th century. As this example illustrates, a key advantage of enabling constraints is that they make use of the superior local knowledge and the diverse perspectives of the members of the group, rather than relying on the relatively limited knowledge and perspective of the central manager (Hayek, 1948).

3.5. How management emerges – vertical self-organization

How do powerful managers emerge in populations in which entities begin as equals, and do not have power over each other? How do the evolutionary interests of managers come to be aligned with those of the entities they manage, so that selection favours the use of power to manage a group of entities in the interests of the group as a whole?

It is not difficult to identify scenarios in which selection would favour entities that develop power over others. For example, their power may enable them to monopolize resources (including access to reproductive opportunities), or to predate others. But exercising power in these ways does not necessarily align their evolutionary/adaptive interests with those they have power over.

However, interests begin to be aligned to some extent if the powerful entities discover ways in which they can harvest an on-going stream of benefits from those they control. Once this occurs, they may do better if they use their power to help the group survive and thrive, and thereby produce a larger stream of harvestable benefits, not just a once-off dividend. As we have seen, proto-managers can substantially boost the productivity of a group by solving the cooperation problem. More specifically, they can increase the stream of harvestable benefits by supporting co-operators who would not otherwise persist in the group. Proto-managers can also suppress free-riders who would otherwise take some of the benefits of cooperation without cooperating in return.

In some circumstances, proto-managers that solve the cooperation problem and harvest an on-going stream of benefits may be able to do better than if they move between groups, exploiting them as they go. Where this is the case, selection operating at the level of individual proto-managers will tend to favour those that remain with a group and use their power to increase the stream of benefits that they harvest from it (Stewart, 1995, 1997a, 2000).

To solve the cooperation problem comprehensively, a proto-manager must have power over the group. It must be able to operate on a sufficiently large scale to act across the group, applying constraints to all its members. It must be able to stand outside the dynamical interaction between group members and not be dependent on involvement in those interactions. An entity that is a normal member of a group will not have the capacity to apply constraints across the group or to harvest resources from it. This is readily seen in human groups: the only way a normal member of a human group can take control of the group is to

acquire power by, for example, using weapons or by increasing the scale over which it operates by forming alliances with others.

Once such a proto-manager is able to harvest benefits from its group on an on-going basis, its interests will tend to be aligned with the success of the group. Selection acting on managers will also favour managers that use their power to align the evolutionary interests of the members of the group with their own interests. If a population of managed groups are in competition, selection will favour managers that are best able to govern their group in ways that enhance the success of the group as a whole.

The process of vertical self-organization embodied in this scenario can account for the evolution of key steps in the trajectory towards increasing integration outlined above (i.e. steps in which individual entities that initially compete with each other have come to be organized into cooperatives that eventually form larger-scale entities at a higher level of organization). For example:

- 1) Some self-reproducing RNA molecules shifted from plundering autocatalytic cycles and sets to governing metabolisms on an on-going basis (Dyson, 1985; Stewart, 1995, 1997a, 2000).
- 2) Some large cells shifted from engulfing and feeding on some forms of bacteria to managing and cultivating them within themselves.
- 3) Some marauding bands of humans shifted from raping and pillaging agricultural villages and communities to establishing kingdoms that governed those communities.

However in modern human societies, the path to managed organizations is much more straightforward. Human intelligence is able to envisage somewhat the benefits of managed organizations and see how they can be organized effectively. This enables humans to construct managed organizations without the intermediary sequence of steps.

3.6. Management and constraints

But not all of the major steps in which cooperative organizations arise and form higher-level entities can be accounted for by the emergence of this form of managed organization. In the kind of managed organization that we have considered to this point, the processes that manage the organization are physically located outside the entities that are managed, and are larger in scale than them. For example, DNA and RNA are external to the entities and processes they manage, as is the governance of modern human societies. But this does appear to be the case for the first multicellular organisms (there is no powerful entity that manages the cells and is external to them), or for early human tribal societies (again there is no powerful chief or other ruler that manages simple tribes (Boehm, 1999)). Yet without external management, multicellular organisms and early human tribes were able to achieve complex cooperative organization, including division of labour. Apparently the cooperation problem was solved in these instances to some significant extent. How was this realized?

Salthe (1985) demonstrates that constraints that can control a dynamic of interacting entities may arise in either of two ways:

- 1) Upper-level constraints: these arise external to the dynamic of entities. They can influence the dynamic without being influenced in return. This is often because they are larger in scale than the entities they constrain, and are constituted by processes that operate significantly more slowly than the interactions in the dynamic. Examples of abiotic upper-level constraints that act on a population of entities include features of the environment that are relatively unchanging from the perspective of the interacting entities, such as large-scale physical structure in the environment. The external managers

referred to in this paper are evolvable systems of upper-level constraints. It is worth emphasizing here that they are often constituted by processes rather than entities.

- 2) Lower-level constraints: these arise within the entities themselves. These constraints are relatively fixed, internal features of the interacting entities that can influence how entities behave in interactions, but are not influenced in return. In effect, they hardwire entities. Examples of lower level constraints in living entities include genetic elements, and internalized norms, customs and beliefs.

3.7. Internal distributed management

However it is not immediately obvious how lower-level constraints that are internal to each entity could manage a group of entities to overcome the cooperation problem. These constraints have the power to constrain an entity that contains them, but how can this power operate across a group of entities?

It can do so if a suitable cluster of constraints are reproduced in each of the entities in the group. Such a cluster of constraints would constitute an internal distributed manager that is capable of acting and coordinating across the group, just as an external manager can. The cluster of constraints could, for example, predispose entities that contain it to: reciprocate in cooperative exchanges; refrain from free-riding; collectively punish or kill any non-co-operator members of the group who do not contain the cluster of constraints (and who therefore are not constrained to cooperate); and prevent any free-riders from joining the group and exploiting it (Stewart, 1995, 1997a, 2000).

Importantly, the evolutionary interests of a distributed internal manager tend to be aligned with the interests of the group that it manages. The manager will capture the benefits of any cooperative activities within the group, because it exists in each of the members of the group. Selection will therefore favour any variant internal manager that constrains the members of the group in ways that promote beneficial cooperation. As a manager evolves, it can be expected to develop a co-adapted assemblage of constraints that organizes the group into a cooperative whole. The manager itself will evolve and operate as a whole. In principle, the evolution of such a manager in gene-based organisms can be reduced to kin selection and related processes. However, the emergence of the manager and the system of evolvable constraints that constitute it can only be properly understood as a co-evolving whole. A fully-developed manager will solve the cooperation problem, and complex cooperative organization will self-organize.

Examples include:

- 1) The transition from independent eukaryote cells to multicellular organisms was organized by the emergence of internal, distributed genetic managers that were reproduced in each cell.
- 2) Early human tribes were organized in the main by internal managers comprising co-adapted sets of norms, mores and inculcated beliefs (to an extent some current human groups such as the hutterites continue to be organized in this way (Sober and Wilson, 1998)).
- 3) Insect societies are organized by genetic internal managers reproduced in members of the society.

3.8. How the global trajectory is driven and shaped

Now that we have identified the particular forms of organization that can comprehensively overcome the cooperation problem, we are in a position to explain the global trajectory towards increasing integration and cooperation. We can identify the causal processes that drive and shape it.

The trajectory is propelled by the potential fitness and utility of cooperative organization. Cooperative groups have the potential to out-compete isolated individuals. Selection will favour effective cooperation, provided it is available for selection.

Importantly, the potential advantages of cooperation are universal—they apply for all living processes in all environments. Whatever the circumstances, co-operators are potentially superior. The potential benefits of cooperation therefore have the capacity to drive a global trend in the evolution of life.

This conclusion is not negated by the fact that gene-based natural selection can drive only local adaptation. Since the benefits of cooperation apply both locally and universally, an organism that exploits the benefits of cooperation can achieve local evolutionary advantage while also participating in a global trend.

But the potential of cooperation can be realized only to the extent that the cooperation problem is overcome. Unless the cooperation problem is solved, complex cooperative organization will not arise. It will not come into existence while individual entities fail to capture sufficient of the benefits of their cooperation. Selection at the group level, no matter how powerful, cannot call it into existence.

We have seen that evolvable management, whether external or internal, can overcome the cooperation problem. It can manage a group of entities to ensure that beneficial cooperation can be sustained within the group and can therefore be the target of selection. This massively expands the possibility space that can be explored as the group evolves.

Where a powerful manager is able to harvest an on-going stream of benefits from the group it manages, it will be able to capture the benefits of any management that increases the productivity of the group. It will therefore be able to advance its own evolutionary/adaptive interests by promoting cooperation within the managed group. If managed groups compete with each other, and if a manager is less successful if it lives independently of its group, the most effective way in which it will be able to advance its interests will be to advance those of the managed group. In these circumstances the manager's evolutionary/adaptive interests will tend to be aligned with those of the managed group as a whole (Stewart, 1997a, 2000).

Because the manager's evolutionary/adaptive interests will tend to be aligned with those of the group, selection acting on the manager will favour management which aligns the interests of the entities it manages with the interests of the group. As a result, all the members of the group, manager and managed entities alike, will adapt cooperatively to serve the group as a whole. Members of the group will be favoured by selection only insofar as they serve the adaptive interests of the group as a whole. As a consequence, the group will increasingly come to be organized and adapted to function as an entity in its own right. Selection operating at the level of the manager and at the level of individual entities will favour entification of the group. The emerging entity will come to be comprised of parts that function interdependently to ensure the survival and reproduction of the whole. Entity-level adaptations will be favoured and shaped by selection at the individual level. The entity will increasingly operate as an integrated functional unit, and will come to act on the world as a coherent whole, as do for example, bacteria, eukaryote cells, and multicellular organisms (Stewart, 1997a, 2000).

But even though managers may solve the cooperation problem for the groups they organize, this simply exports the cooperation problem to a higher level and to a larger scale. Managers can only solve the cooperation problem at the scale over which they can exercise power. At this scale they can suppress free-riding and support co-operators within the group they manage. But the cooperative groups that they organize at this scale will compete with each other, and the cooperation problem will therefore apply

between groups. Complex cooperation will not be able to persist at this larger scale, and will not be available for selection to operate on. The full benefits of cooperation between groups will be unrealized. These unrealized benefits will drive the emergence and evolution of new managers that can capture sufficient of the benefits by organizing groups into larger-scale cooperatives. In this way, the trajectory of evolution will be driven towards further cooperative integration that is able to realize some of the benefits of cooperation at larger-scales (Stewart, 1997b, 2000).

This process repeats itself at larger scales and higher levels, giving rise to the step-wise nature of the trajectory towards increasing integration. At each step in this trajectory, the scale of cooperative organization increases. As the trajectory unfolds, the potential benefits of cooperation are progressively realized through the actual integration of living processes into larger-scale organizations. At each step, the emergence of management greatly expands evolutionary possibility space.

As we have seen, the benefits of cooperation exist interspecifically as well as intraspecifically. The cooperation problem stands in the way of the realization of the benefits. These unrealized benefits provide opportunities for suitable managers to emerge and manage the integration of multi-specie systems into larger-scale cooperative entities. Integration is not limited to members of the same species and will progressively include the integration of all living processes. Examples of steps in this process include the integration of prokaryotes into the eukaryote cell, and the agriculture practised by various species of ants, termites and beetles (Mueller et al., 2005).

The potential benefits of cooperation can be expected to continue to drive the step-wise process of integration until it reaches a natural limit. This limit is the scale of the planet. Once a global cooperative organization emerges that integrates the living processes of the globe, it will not be in competition with other entities on the planet. It will not be subject to a cooperation problem at the level of the globe.

However, while ever cooperative organizations of the largest-scale are smaller than the planet, they will constitute a population of organizations that compete with each other. This will produce the cooperation problem. Organizations that act cooperatively towards others will tend to be out-competed. This dynamic is currently evident at the level of nation states. Global warming and international war are both manifestations of the cooperation problem. The existence of the cooperation problem at this level means that potential benefits can be realized through the emergence of global management (including global governance). Global management has the potential to overcome the cooperation problem at the level of nation states.

Effective global management would not only suppress destructive competition between nations. Ultimately it would also have the potential to realize the benefits of integrating all lower level entities into the global organization, including by supporting entities that contribute positively to the global organization. A global manager that is sufficiently powerful and evolvable has the potential to control a hierarchy of management that integrates the living and non-living processes of the planet into a cooperative and unified global entity (for more detailed discussion, see Stewart, 2000).

4. Testing the directionality hypothesis

4.1. The apparent absence of directional change in bacteria and other lineages

As discussed in Section 1, the central challenge confronting any claim for a driven, global trend in evolution is that many lineages do not appear to show any evidence of such a trend for long periods of evolutionary time. How is this challenge met by the claim that

living processes tend to be progressively integrated into cooperative organizations of larger and larger scale as evolution unfolds?

The apparent absence of cooperative integration amongst many lineages of organisms does not appear to be due to a lack of potential benefits that could flow from their integration. The benefits of cooperation are just as available to these lineages as they are to the lineages that have cooperated successfully. Potentially, any given population of organisms could exploit its resources more effectively if the population was coordinated and organized cooperatively. But these potentials would not be able to be realized until the emergence of appropriate forms of management that can exploit the benefits of cooperation cost/effectively. And this might not occur for long periods – effective management requires complex arrangements that cannot be expected to evolve easily. Nor would we expect lineages that fail to exploit the potential benefits of cooperation to be always replaced by those that have. This is because they may be adapted to their niches by specializations that outweigh any fitness advantages that might accrue to co-operators that are not specialized. For example, the small size of many species of bacteria gives them a significant advantage over eukaryote cells and multicellular organisms in many niches.

But the case for an overall driven trend towards increasing integration does not have to rely on these largely theoretical arguments. It is evident that many previously 'trend-less' lineages are now beginning to be integrated into larger-scale cooperatives, and this process can be expected to continue. This integration has been made possible by the emergence of organisms that are highly evolvable (intelligent), powerful and able to coordinate actions over much larger scales than bacteria and other organisms—i.e. due to the emergence of humans and cooperative organizations of humans. Humans are beginning to develop the technology and know-how to manage other organisms and integrate them into larger-scale human-dominated organizations for human ends. Humanity is starting to exploit the potential benefits of cooperation that so far have been unrealized for extended periods by many other lineages.

As well as managing particular species for human purposes, humanity is also beginning to actively manage entire ecosystems and the large-scale ecological cycles and processes that are essential for human survival and flourishing. Bacteria and other organisms play a critically important role in these processes. Human disruption of ecological systems will increasingly demand active management of ecosystems and the organisms that comprise them. Current attempts to regulate the amount of carbon dioxide in the atmosphere are just a small, first step in this direction.

These trends toward the integration of bacteria and other organisms into large-scale cooperative organization can be expected to continue apace. The potential benefits of cooperation at the global level will incentivize human systems to increasingly manage the energy, matter, living processes and technology of the planet into a cooperative organization on the scale of the planet. From the perspective of this global organization, the planet's ecosystems and bio-geo-chemical cycles will constitute important parts of its metabolism that are managed by human systems in the interests of the organization (just as DNA manages the metabolic processes within a cell). The benefits of cooperation at the global level will drive entification of the global organization, including the management of all living processes into a symbiotic community that will be like an organism on the scale of the planet. As it has at every other level of organization, entification can be expected to eventually produce an entity that is able to act coherently as a whole, with goals and intelligence of its own. Heylighen (2012) argues that the intelligence of this emerging global entity will be constituted by humans and artificial

intelligence in the form of a global brain. To the extent that the global entity becomes aware of the nature of the large-scale evolutionary processes that have formed it and that will determine its future success, it will manage the living and non-living processes of the planet to create the best-possible platform for its further development and evolution (Stewart, 2010).

4.2. Predictions

The ultimate test for a theory is whether it can account for phenomena that are hitherto unexplained, and whether it can make surprising predictions that prove to be accurate. Many detailed predictions and explanations can be derived from the model developed here. Perhaps the most significant of these is that the theory:

- 1) Accounts for the nested-hierarchical structure of living processes, including in relation to human organization.
- 2) Hypothesises that each level in these nested hierarchies arose through the emergence of external management or internal distributed management (or a combination), and predicts that this will be evident in the way each of these levels is organized (i.e. each level will continue to include forms of organization that originally emerged as internal or external managers).
- 3) Predicts and explains the current rapid integration of living processes into human organizations, and predicts that this will continue and accelerate as part of the eventual emergence of a global entity that integrates the living processes, matter, energy and technology of the planet, and undergoes an entification process in its own right.
- 4) Predicts that the further integration of humanity into a global entity will be driven by the potential benefits of overcoming the cooperation problem on a planetary scale. Humanity will increasingly encounter challenges such as climate change and war that demand global coordination because they cannot be adequately solved by nation states or other smaller-scale entities acting alone (of course, the forces that drive the trajectory of evolution do not provide any absolute guarantee that human civilization will respond effectively to these challenges and survive them).

5. Conclusion

If we stand back from the history of life on Earth so that we can see past, present and future evolution as a whole, and if we take into account the large-scale processes that have begun to emerge recently and that are likely to continue strongly into the future, an overall direction is unmistakable. Evolution has been heading towards the emergence of a coordinated and integrated global entity.

It is difficult to discern the overall direction of evolution without this wider perspective. It is like looking at a developing chicken embryo, and focusing closely on some of the specific processes within the embryo for short periods of time. In many cases, it would not be evident that the particular processes are headed anywhere in particular. It is only when all the processes of the developing embryo are considered as an integrated whole over a longer time frame, that it becomes abundantly clear that there is an overall trajectory to the developing embryo.

When we stand back from the evolutionary process on this planet and consider it as a coherent whole, we see that there are two great trends within evolution. One is towards diversification. The other is towards integration and cooperation. As we have seen in some detail, both trends are driven by selection processes that are consistent with mainstream evolutionary theory.

When living processes first emerged, they were very small in scale and restricted to specific environmental conditions. From there life spread, specializing and diversifying into other environments, exploiting an increasing variety of free energy and other resources. Some living processes discovered ways to cooperate and to coordinate their activities, significantly enhancing their effectiveness in particular environments. These larger-scale cooperatives in turn spread and diversified. But because of limits to their ability to exploit the benefits of cooperation, they were not able to exploit all niches more effectively than their predecessors, and therefore did not replace them entirely.

Some of these larger-scale cooperatives in turn discovered ways to coordinate and cooperate, thereby further increasing their effectiveness. This enabled them to spread and diversify further across the face of the planet. Some of these larger-scale entities developed the power and evolvability to integrate other organisms into their organizations, overcoming impediments that prevented the organisms from exploiting the benefits of cooperation on their own. Repetitions of these evolutionary integrations produced cooperative organizations of greater and greater scale, and now seem highly likely to lead to the emergence of a cooperative organization on the scale of the planet. Such a global organization seems likely to eventually integrate the diversified and specialized living processes of the planet into a coherent whole, enabling the emerging global entity to exploit energy and other resources from diverse environments.

It can be seen at this point that the diversification and specialization that occurred at lower levels is an integral part of this directional process. This diversification enabled the widest range of resources to be exploited by life. The subsequent integration of these diverse smaller-scale living processes into an emerging global entity will enable the entity to exploit the widest possible range of resources across the varied circumstances of the planet.

As the global entity emerges, it can be expected to increasingly manage the living processes, energy, matter and technology of the entire planet into a coordinated whole. As it develops, it will optimize all the processes on which it depends (including large-scale ecological systems) in order to create the most effective platform for its future evolution. Of course, this is not likely to be the end of the evolutionary trajectory. The trajectory is likely to have unfolded elsewhere (although the details are likely to be different), and similar diversifications and integrations are therefore likely to be repeated at larger and larger scales across space and time. The further unfolding of this trajectory and its profound implications for humanity are considered by Stewart (2010).

Acknowledgements

I gratefully acknowledge the benefit of useful discussions and comments from David Richards, Wilson Kenell and Losa Bodie, and useful comments from two anonymous reviewers.

References

- Bagley, R.J., Farmer, J.D., 1991. Spontaneous emergence of a metabolism. In: Langton, C., Taylor, J., Rasmussen, S. (Eds.), *Artificial Life II*. Addison and Wesley, New York, pp. 93–141.
- Boehm, C., 1999. *Hierarchy in the Forest: The Evolution of Egalitarian Behavior*. Harvard University Press, Cambridge, MA.
- Boyd, R., Gintis, H., Bowles, S., Richerson, P.J., 2005. The evolution of altruistic punishment. In: Gintis, H., Bowles, S., Boyd, R., Fehr, E. (Eds.), *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press, Cambridge, pp. 215–227.
- Boyd, R., Richerson, P., 2005. Solving the puzzle of human cooperation. In: Levinson, E., Stephen, C., Jaisson, P. (Eds.), *Evolution and Culture*. MIT Press, Cambridge, pp. 105–132.
- Buss, L.W., 1987. *The Evolution of Individuality*. Princeton University Press, Princeton.

- Corning, P.A., 1983. *The Synergism Hypothesis: A Theory of Progressive Evolution*. McGraw-Hill, New York.
- Craig, C., 2008. The resurgent idea of world government. *Ethics Int. Aff.* 22, 133–142.
- Dugatkin, L., 1999. *Cheating Monkeys and Citizen Bees: The Nature of Cooperation in Animals and Humans*. Free Press, New York.
- Dyson, F., 1985. *Origins of Life*. Cambridge University Press, London.
- Fehr, E., Gächter, S., 2002. Altruistic punishment in humans. *Nature* 415, 137–140.
- Gould, S.J., 1988. Trends as changes in variance: a new slant on progress and directionality in evolution. *J. Paleontol.* 62, 319–329.
- Gould, S.J., 1996. *Full House: The Spread of Excellence From Plato to Darwin*. Harmony Books, New York.
- Griffin, A.S., West, S.A., Buckling, A., 2004. Cooperation and competition in pathogenic bacteria. *Nature* 430, 1024–1027.
- Hamilton, W.D., 1964. The genetical evolution of social behaviour, I & II. *J. Theor. Biol.* 7, 1–52.
- Hayek, F.A., 1948. *Individualism and Economic Order*. University of Chicago Press, Chicago.
- Heylighen, F., 2007. The global superorganism: an evolutionary-cybernetic model of the emerging network society. *Soc. Evol. Hist.* 6, 58–119.
- Heylighen, F., 2012. Conceptions of a global brain: an historical review. In: Rodrigue, B., Grinin, B., Korotayev, A. (Eds.), *From Big Bang to Global Civilization: A Big History Anthology*. University of California Press, Berkeley.
- Hordijk, W., Steel, M., Kauffman, S., 2012. The structure of autocatalytic sets: evolvability, enablement, and emergence. *Acta Biotheor.* 60, 379–392.
- Kauffman, S.A., 1993. *The Origins of Order: Self-organisation and Selection in Evolution*. Oxford University Press, New York.
- Keller, L., Ross, K., 1998. Selfish genes: a green beard in the red fire ant. *Nature* 394, 573–575.
- Knoll, A., Bambach, R., 2000. Directionality in the history of life: diffusion from the left wall or repeated scaling of the right? *Paleobiology* 26 (Suppl.), 1–14.
- Marcot, J., McShea, D., 2007. Increasing hierarchical complexity throughout the history of life: phylogenetic tests of trend mechanisms. *Paleobiology* 33, 182–200.
- Margulis, L., 1981. *Symbiosis in Cell Evolution*. W.H. Freeman and Company, San Francisco.
- Mateo, J.M., 2003. Kin recognition in ground squirrels and other rodents. *J. Mammal.* 84, 1163–1181.
- Maynard Smith, J., 1988. Levels of selection. In: Nitecki, M.H. (Ed.), *Evolutionary Progress*. University of Chicago Press, Chicago, pp. 219–236.
- Maynard Smith, J., Szathmáry, E., 1995. *The Major Transitions in Evolution*. Freeman, Oxford.
- McShea, D., 1991. Complexity and evolution: what everyone knows. *Biol. Philos.* 6, 303–324.
- McShea, D., 1994. Mechanisms of large-scale evolutionary trends. *Evolution* 48, 1747–1763.
- McShea, D., 1996. Metazoan complexity and evolution: is there a trend? *Evolution* 50, 477–492.
- McShea, D., 2001. The “minor transitions” in hierarchical evolution and the question of directional bias. *J. Evol. Biol.* 14, 502–518.
- McShea, D., Simpson, C., 2011. The miscellaneous transitions in evolution. In: Sterelny, K., Calcott, B. (Eds.), *The Major Transitions in Evolution Revisited*. The MIT Press, Cambridge, pp. 19–33.
- Miller, J.G., 1978. *Living Systems*. McGraw-Hill, New York.
- Mouden, C.E., West, S.A., Gardner, A., 2010. The enforcement of cooperation by policing. *Evolution* 64, 2139–2152.
- Mueller, U.G., Gerardo, N.M., Aanen, D.K., Six, D.L., Schultz, T.R., 2005. The evolution of agriculture in insects. *Annu. Rev. Ecol. Evol. Syst.* 41, 563–595.
- Nadell, C.D., Foster, K.R., Xavier, J.B., 2010. Emergence of spatial structure in cell groups and the evolution of cooperation. *PLoS Comput. Biol.* 6, e1000716.
- Nowak, M.A., 2006. *Evolutionary Dynamics*. Harvard University Press, Cambridge, MA.
- Okasha, O., 2006. *Evolution and the Levels of Selection*. Clarendon Press, Oxford.
- Olson, M., 1965. *The Logic of Collective Action*. Harvard University Press, Cambridge, MA.
- Perc, M., Szolnoki, A., 2010. Coevolutionary games – a mini review. *BioSystems* 99, 109–125.
- Roberts, G., 2008. Evolution of direct and indirect reciprocity. *Proc. Royal Soc. B* 275, 173–179.
- Ridley, M., 1996. *The Origins of Virtue*. Viking, London.
- Ruse, M., 1996. *Monad to Man: the Concept of Progress in Evolutionary Biology*. Harvard University Press, Cambridge, MA.
- Salthe, S., 1985. *Evolving Hierarchical Systems*. Columbia University Press, Columbia.
- Sigmund, K., 2010. *The Calculus of Selfishness*. Princeton Univ. Press, Princeton, MA.
- Sober, E., Wilson, D.S., 1998. *Unto Others: The Evolution and Psychology of Unselfish Behaviour*. Harvard University Press, Cambridge.
- Stewart, J.E., 1995. Metaevolution. *J. Soc. Evol. Syst.* 18, 113–114.
- Stewart, J.E., 1997a. Evolutionary transitions and artificial life. *Artif. Life* 3, 101–120.
- Stewart, J.E., 1997b. Evolutionary Progress. *J. Soc. Evol. Syst.* 20, 335–362.
- Stewart, J.E., 2000. *Evolution's Arrow: the Direction of Evolution and the Future of Humanity*. Chapman Press, Canberra.
- Stewart, J.E., 2010. The meaning of life in a developing universe. *Found. Sci.* 15, 395–409.
- Trivers, R.L., 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46, 35–57.
- Turchin, V.F., 1977. *The Phenomenon of Science: A Cybernetic Approach to Human Evolution*. Columbia University Press, New York.
- Ulanowicz, R.E., 2009. *A Third Window: Natural Life Beyond Newton and Darwin*. Templeton Foundation Press, West Conshohocken, PA.
- Vidal, C., 2014. *The Beginning and the End: The Meaning of Life in a Cosmological Perspective*. Springer, Dordrecht.
- Walker, B., Barrett, S., Polasky, S., et al., 2009. Looming global-scale failures and missing institutions. *Science* 325, 1345–1346.
- Williams, G.C., 1966. *Adaptation and Natural Selection*. Princeton University Press, Princeton.