The final version of this paper was published in *Foundations of Science*, 15: 395-409 (2010)

THE MEANING OF LIFE IN A DEVELOPING UNIVERSE

John Stewart
Member of the Evolution, Complexity and Cognition Research Group, The Free University of Brussels
john.stewart@evolutionarymanifesto.com

Abstract: The evolution of life on Earth has produced an organism that is beginning to model and understand its own evolution and the possible future evolution of life in the universe. These models and associated evidence show that evolution on Earth has a trajectory. The scale over which living processes are organized cooperatively has increased progressively, as has its evolvability. Recent theoretical advances raise the possibility that this trajectory is itself part of a wider developmental process. According to these theories, the developmental process has been shaped by a yet larger evolutionary dynamic that involves the reproduction of universes. This evolutionary dynamic has tuned the key parameters of the universe to increase the likelihood that life will emerge and produce outcomes that are successful in the larger process (e.g. a key outcome may be to produce life and intelligence that intentionally reproduces the universe and tunes the parameters of 'offspring' universes). Theory suggests that when life emerges on a planet, it moves along this trajectory of its own accord. However, at a particular point evolution will continue to advance only if organisms emerge that decide to advance the developmental process intentionally. The organisms must be prepared to make this commitment even though the ultimate nature and destination of the process is uncertain, and may forever remain unknown. Organisms that complete this transition to intentional evolution will drive the further development of life and intelligence in the universe. Humanity's increasing understanding of the evolution of life in the universe is rapidly bringing it to the threshold of this major evolutionary transition.

Keywords: conscious evolution; development of the universe; evolution of the universe; intentional evolution.

1. Introduction

Until recently, a scientific understanding of the natural world has failed to provide humanity with a larger meaning and purpose for its existence. In fact, a scientific worldview has often been taken to imply that the emergence of humanity was an accident in a universe that is completely indifferent to human concerns, goals, and values (e.g. see Weinberg, 1993).

Humanity has had to supplement a naturalistic understanding with beliefs in supernatural beings and processes if it wanted a worldview that includes a meaningful role for humanity in a larger scheme of things.

But recent advances in evolutionary science are beginning to change this. In particular, we are rapidly improving our understanding of the evolutionary processes that have produced life on Earth and that will determine the future evolution of life in the universe. While it is far too early to tell with certainty, it is possible that the universe and the evolution of life within it have been shaped by yet larger evolutionary processes to perform particular functions that are relevant to these larger processes.

If this proves to be the case, then these functions have a purpose in the same sense that the functions performed by our eyes have a purpose in the larger evolutionary processes that have shaped humanity.

This paper explores some key implications for humanity of the larger-scale evolutionary and developmental processes that appear to operate in the universe and beyond. In particular, the paper shows that humanity has a role to play in these processes. It also shows that the success of the processes depends critically on humanity (and other organisms that reach a similar stage in evolution) understanding this role and embracing it intentionally.

We begin by briefly surveying some of the main theories of these larger-scale processes.

2. The Trajectory of Evolution

Many theorists have suggested that evolution exhibits large-scale trends (see Blitz, 1992; Ruse, 1996 for overviews).

In particular, it has often been noted that evolution has moved through a sequence of transitions in which smaller-scale entities are organized into larger-scale cooperatives. Self-replicating molecular processes were organized into the first simple cells, communities of simple cells formed the more complex eukaryote cell, organizations of these cells formed multi-cellular organisms, and organisms were organized into cooperative societies. A similar sequence appears to have unfolded in human evolution: from family groups, to bands, to tribes, to agricultural communities and city states, and so on (e.g. see de Chardin, 1965; Corning, 1983; Blitz, 1992)

As this sequence unfolded, a progressively higher proportion of living processes were organized into cooperative organizations of larger and larger scale. This long sequence also seems to have been associated with a series of improvements in evolvability (the capacity to discover effective adaptations).

However, evolutionary biologists have been very reluctant to accept that these apparent patterns represent actual evolutionary trajectories (e.g. see Gould, 1996; Ruse, 1996).

In large part this is because these hypotheses about directionality were not accompanied by explanations of how the claimed trajectories were produced by known evolutionary processes. This left them open to the criticism that they necessarily relied on impermissible teleological mechanisms.

The view that the evolutionary process is not directional eventually became widely accepted within the modern evolutionary synthesis (Gould, 1996). But this was not because any evidence or theory conclusively ruled out large-scale directionality. Instead, as demonstrated in detail by Ruse (1996), opposition to directionalism was given considerable impetus by the actions of the founders of the synthesis *who were in fact themselves directionalists*. As part of their intentional efforts to build the professional standing of evolution as a scientific discipline, the founders used their power as editors and peer reviewers to cleanse the synthesis of notions of direction, progress and purpose. Apparently they feared that to do otherwise would embroil their fledgling field in public controversy and attract criticism that it was unscientific. Ironically, the intentional and sustained efforts of directionalists paved the way for anti-directionalism to become mainstream dogma in evolutionary biology during the second half of the twentieth century.

2.1 The evolution of cooperation

Until near the end of the 20th century, the hypothesis that evolution moves in the direction of producing cooperative organizations of larger and larger scale gained little traction. In large part this was because mainstream biology held to the view that selfishness, rather than cooperation, is favored in evolution (e.g. see Williams, 1966; Dawkins 1976). This position notes that selection will act against entities that invest resources in cooperation but do not capture any of the benefits it produces. They will be outcompeted by 'selfish' entities that undermine cooperation by taking the benefits without contributing anything in return (e.g. cheats, free-riders and thieves).

According to this position, only limited forms of cooperative organization are likely to emerge at any level, and then only under restricted conditions. Cooperation will be restricted to those limited circumstances in which individual entities are somehow able to capture the benefits of cooperating. This can occur where the interactions between entities are 'collectively autocatalytic' (i.e. where the actions of each entity that participates in a cooperative process increases the fitness of one or more others, and its fitness is in turn increased by other entities). The simplest form of collective autocatalysis is where two entities engage in reciprocal exchanges of benefits.

Examples at various levels of organization of cooperation that is collectively autocatalytic are: autocatalytic sets of proteins (e.g. see Bagley and Farmer, 1991); RNA hypercycles (e.g. see Eigen and Schuster, 1979); autocatalytic cycles of indirect mutualism in ecosystems (e.g. see Ulanowicz, 1986); kin selection amongst multi-cellular organisms (e.g. see Hamilton, 1964); and reciprocal altruism (direct and indirect) amongst humans (e.g. see Trivers, 1971).

However these forms of collective autocatalysis fall far short of accounting for the major evolutionary transitions. They are unable to explain how evolution has apparently organized smaller-scale entities into complex larger-scale cooperatives that eventually become entities in their own right.

In part this is because free-riding, cheating and theft can be expected to undermine and disrupt autocatalytic processes. Furthermore, these processes will emerge only where interactions between entities just happen to form a closed autocatalytic system. There is no reason to expect that advantageous cooperative processes will be collectively autocatalytic. Those that are not

will fail to self-organize and will be undermined by individual selection. The complex forms of cooperative organization that are necessary if a group is to become an entity in its own right will not emerge.

2.2 Advances in understanding the evolution of complex cooperative organization

However, in the past two decades considerable progress has been made in understanding how evolution has repeatedly organized independent entities into larger-scale cooperatives. A number of researchers have contributed to the development of a thorough understanding of specific transitions, such as the transition to multi-cellularity (Buss, 1987; Michod, 1999). Others have attempted to develop more general models that are applicable to all the transitions to larger-scale cooperation (Stewart 1995, 1997a,b, 2000; Maynard Smith and Szathmary, 1995; Heylighen, 2006).

In general, this work has shown that evolution can organize complex cooperation amongst self-interested individuals once particular conditions are met.

Stewart (1995; 1997a,b, 2000) shows that evolution will favor the emergence of cooperation amongst self-interested entities when they are embedded in a particular form of organization that makes it in their interests to cooperate. In this form of organization, sets of evolvable constraints (managers) constrain the activities of the self-interested entities, preventing free riding and other actions that undermine cooperation. Furthermore, the evolvable constraints ensure that entities that contribute to effective cooperation are rewarded. As a result, entities capture the benefits of their cooperation and cooperation can be favored by individual selection.

An organization managed by a set of evolvable constraints therefore escapes the limitations that prevent collective autocatalysis from producing complex cooperative organization. The manager can ensure that any cooperative process that benefits the organization as a whole is sustainable, even if the process itself is not collectively autocatalytic. It achieves this by using its power to ensure that entities that contribute to the process benefit from doing so, as well as by restraining free-riding and cheating.

Examples of evolvable constraints include the RNA that managed early cells, the DNA that is reproduced in each cell of a multi-cellular organism (and that thereby manages the interactions between cells), and the governments that manage modern human societies (e.g. see Stewart, 2000).

Importantly, a manager is able to harvest some of the benefits that flow from any cooperation that it organizes. As well as using these resources to reward cooperators, the manager can also use some to enhance its own fitness. In this way it captures some of the benefits of organizing cooperation. Selection will therefore tend to favor managers that are able to organize effective cooperatives.

As a consequence, it is in the interests of the manager to organize cooperation, and in the interests of the smaller-scale entities to cooperate. In this form of organization, the interests of all the members of the organization (including the manger) are aligned with the interests of the organization as a whole.

Evolution will tend to favor cooperative organizations over independent entities because of the advantages that cooperation can provide. For example, cooperation enables the exploitation of synergies, including through specialization and division of labor (Corning, 1983; Stewart, 2000).

And the larger the scale of cooperative organization, the more resources commanded by the cooperative, the greater its power, the greater the impact and scale of its actions, and therefore the wider the range of environmental challenges that it can meet successfully. And the greater the evolvability, the greater the capacity to respond effectively to any adaptive needs and opportunities.

Larger scale and greater evolvability both have the potential to provide evolutionary advantage to living processes across a wide range of environments. This is because they are meta-adaptive capacities—they improve the ability to adapt in all circumstances, although they are not themselves an adaptation to any specific circumstance (it is also worth noting that both are deeply interrelated—increases in scale and power generally increase the range of possible adaptive responses, and hence enhance evolvability [Stewart, 1995])

As improvements in these capacities are discovered, life will tend to evolve directionally. Of course, this trajectory will often be masked by meandering, halting and back-tracking, particularly where the process that searches for improvements relies on blind trial and error. Furthermore, improvements in these capacities will be favored only when the advantages they provide outweigh their cost. As a consequence, directional change will often stall until evolution discovers a cost-effective way of enhancing the capacities.

Taken together, the research of the last two decades constitutes a very strong case that the apparent trajectory of evolution towards larger-scale cooperative organization and greater evolvability is 'real'. It is driven by processes that do not rely on teleology and are readily understandable within modern evolutionary theory.

3. Extrapolating the Trajectory

The extrapolation of this trajectory is reasonably straightforward, at least initially.

The next major transition on Earth would be the formation of global cooperative society (Stewart 1995, 1997a, 2000; Heylighen, 2007). Such a society would be enabled by a system of global constraints (governance) that organizes cooperation (including market processes) and that suppresses destructive competition (including war and international pollution). The evolvability of human society is also likely to increase rapidly through the continued development of artificial intelligence and other technology, and also through the fundamental transition in human psychology which will be discussed in Section 5 below.

Extrapolating this trajectory further would see the continued expansion of the scale of cooperative organization out into the solar system and beyond, and the further enhancement of evolvability (including through the intensification and compression of intelligence discussed by Smart, 2008). However, the expansion of the scale of cooperative organization might occur largely by linking with other organizations of living processes that originate elsewhere, rather than by 'empire building'. This linking up could be expected to greatly increase evolvability through the exchange of knowledge and technology (including artificial intelligence).

The possibility of life arising elsewhere seems high. There does not appear to be anything special about this planet, the emergence of life on it, or its evolutionary trajectory that make it unlikely to have occurred elsewhere. The details can be expected to differ, but it is likely that the general increase in evolvability and the step-wise increase in the scale of cooperative organization will be ubiquitous. And as we will discuss below, any other living processes that are expanding out into space can be expected to be cooperators, not hostile. (However, there is

no consensus within the scientific community about the likelihood of extraterrestrial life and the 'specialness' of Earth. For an overview of the debates, see Dick, 1996.)

If the trajectory continues in this way, the scale of cooperative organization would continue to expand throughout the universe, comprised of living processes from multiple origins. As it increased in evolvability and scale, its command over matter, energy and other resources would also expand, as would its ability to use these resources intelligently to pursue whatever objectives it chooses. The ability of cooperative life to influence large-scale events in the universe would increase, and it might even develop a capacity to impact on processes outside the universe we currently know.

The question of whether the trajectory is likely to continue in this way is discussed in Section 5.2.

4. The Developmental Possibility

4.1 Current developmental hypotheses

With our current level of knowledge and intelligence, we cannot say much about the nature of any larger-scale processes in which our universe is embedded. But as a number of theorists have noted, the following considerations raise some intriguing possibilities (see Davies, 2006 for a broad overview):

(1) The known universe exists (there is something rather than nothing), and it is reasonable to presume that it owes its existence to processes that exist outside it.

(2) The fundamental laws and parameters of the known universe seem extraordinarily finely tuned to the needs of life—slight changes would produce a universe in which life would seem highly unlikely.

(3) the evolution of life follows a trajectory that appears likely to eventually produce a universe that is controlled and managed in significant respects by intelligent life (including artificial intelligence).

A number of theorists have tried to account for these considerations by suggesting that our universe and the evolution of life within it is a developmental process that has itself been shaped by evolutionary processes of even wider scale (Crane, 1994; Harrison, 1995; Gardner, 2000, 2003, 2005, 2007; Smart, 2000, 2002, 2008; Vidal, 2008).

According to this hypothesis, the basis laws and parameters of the universe have been tuned so that it eventually develops into an entity that is managed by intelligence. This intelligence is 'developmentally destined' to organize the reproduction of the universe and to tune the parameters of the 'offspring universes' so that they are especially conducive to the development of life and intelligence. The effectiveness of the tuning of the developmental process is enhanced as the cycle repeats.

The developmental singularity hypothesis (Smart, 2002, 2008) includes much of this basic schema, but suggests that life will transcend the universe and initiate the reproduction cycle without linking up with other living processes on the scale of the universe. Smart builds on the idea that life on Earth will enter a post-biological technological singularity in the relatively near future, possibly this century (Adams, 1909; Good, 1965; Vinge, 1993; Broderick, 1997). This

accelerating trend towards higher evolvability will continue as intelligence rapidly increases in density and efficiency through compression in matter/energy and space/time. Eventually, this is postulated to produce local intelligences with black-hole-analogous features—a highly local, dense and maximally computability-efficient network of entities that the hypotheses terms a 'developmental singularity'. Smart goes on to suggest that this local (Earth originating) intelligence will interact with some other intelligences that originate elsewhere, and then begin the process of universe reproduction in the quantum domain of black holes.

Smart points out that a particular strength of his hypothesis is its ability to parsimoniously explains the Fermi Paradox—the hypothesis suggests that the reason why we do not see evidence of intelligence or life elsewhere in our universe is because soon after life reaches our stage of development, it enters the developmental singularity, effectively disappearing from view in our space/time.

4.2 Other developmental possibilities

However, work on developmental models is only in its early infancy. Existing models do not explore all the broad classes of developmental possibilities and their implications.

In particular, the key models incorporate the assumption that the only source of inherited information provided to offspring universes is transmitted through the equivalent of the germ line—i.e. through the particular values of the fundamental laws and parameters that shape how the offspring universes will develop.

This assumption seems to be based largely on an analogy with the development and evolution of life on Earth prior to the emergence of cultural evolution. For most organisms on the planet, inherited information is transmitted primarily through the germ line. Very little is transmitted from parent to offspring during their lives, or between adults. This is a serious limitation in evolvability—all the adaptive information that is acquired by an organism during its life dies with it.

This limitation was overcome somewhat with the evolution of humanity and the emergence of cultural modes of transmission. Humans undergo a relatively short period of development in the womb where the information they inherit is largely restricted to the germ line. But this is followed by a much longer period in which they acquire cultural information that has been accumulated across the generations and is transmitted to them during their lives. The emergence of cultural transmission paved the way for the massive enhancement in evolvability that produced human science and technology.

The transmission mechanism postulated by current developmental models of the universe is not quite as limited as the mechanism that applied prior to the emergence of cultural transmission on Earth. Most models suggest that when an intelligence is tuning the parameters of offspring universes, it would draw on the knowledge it acquired during the life of the universe. But this is still an extremely limited information channel. Most of the science and knowledge acquired by the intelligence would be lost. This limitation would be even more serious if universes governed by intelligence engage in extra-universal interactions and activities that affect their evolutionary success. Most of what is learnt about those interactions would be lost.

We could therefore expect that extra-universal evolutionary processes would favor transmission between parent universes and their offspring, and between adult and young universes, once the young universes have developed sufficiently. It would also favor transmission between adult

universes. If these forms of transmission are achievable, they would significantly enhance the evolvability of offspring universes, including their capacity to engage in extra-universal interactions that affect their fitness. The germ line of a universe that fails to develop or to receive these forms of transmission could be expected to suffer a similar fate to a human germ line that fails to support cultural transmission.

Life and intelligence that is committed to contributing to the successful development of its universe can therefore be expected to search for every feasible way of opening up possibilities for such transmission. It will also seek to exploit any potential for other cooperative interactions between universes for whatever projects are relevant to evolution and development at that level.

The fact that we are not yet receiving such transmission does not rule out its existence. It might simply mean that like individual humans, life in the universe might have to achieve a particular level of development before this form of transmission is feasible and productive. This possibility is also consistent with the likelihood that intelligent life on a planet would not be contacted by life originating elsewhere until it reaches a particular level of development (Stewart, 2000; 2008a).

It is also far too early to rule out the possibility that transmission can occur between intelligent universes, or that intelligent universes can be involved in some extra-universal evolutionary dynamic that involves interaction between them and some larger environment. To attempt to decide these issues on the basis of current physics would be even less reliable than Lord Kelvin's impressive 1895 proof that heavier-than-air flight is impossible. The difficulty we face in trying to evaluate these possibilities at our current scale and intelligence would be similar to the challenge facing an intelligent bacterium in our gut that is trying to make sense of the social interactions that humans engage in.

We are in a similar position in relation to developmental hypotheses in general. We have insufficient data at our present state of knowledge and intelligence to adequately test these hypotheses. And there are a number of competing non-developmental hypotheses that can account for the same evidence. Some alternatives such as the multiverse hypothesis account for the apparent fine tuning of the universe for life on the basis of a kind of blind trial and error—it postulates a large number of universes each with different fundamental laws and parameters, with chance favoring the likelihood that some will suit the emergence of life (e.g. see Susskind, 1995). Other alternatives are similar to the developmental hypotheses in that they account for fine tuning through the operation of intelligence. For example, a class of models suggests that our universe could be a simulation designed and initiated by an unknown intelligence operating outside the universe (Bostrom, 2003; Martin, 2006).

To summarize to this point, there is a very strong case that the evolution of life on Earth has been directional. There is also good reason to suggest that this trajectory applies to life that originates elsewhere in the universe. If the trajectory continues, life from different origins will link up to form cooperative organizations of increasing scale and evolvability. This is consistent with the possibility that the evolution of life in the universe is itself part of a larger developmental and evolutionary process. Other evidence such as fine-tuning is consistent with this possibility. But it is far too early to treat developmental hypotheses as anything other than possibilities.

5. The Transition to Intentional Evolution and Development

This section focuses on a critical psychological transition that needs to occur if evolution on a planet is to advance beyond the emergence of a global society and go on to contribute to any larger-scale developmental and evolutionary processes.

Up until the emergence of global society, natural selection and cultural processes will tend to drive evolution along its trajectory towards greater evolvability and increased scale of cooperation. As we have seen, cooperative organizations that are larger in scale and more evolvable will out-compete others. But these forms of competition-driven selection will come to an end as a global society emerges. The global society will not be involved in direct or immediate competition with other planetary societies (Stewart, 2000).

Once this stage is reached, the actions and objectives of the global society would be determined by the values and goals of its members (provided the global society is organized democratically). The society will do what is necessary to advance the evolutionary process only to the extent that this is consistent with the goals and motivations of its members.

However, it is highly unlikely that the desires and motivations of the members of the global society will be consistent with the needs of future evolution. Their desires and motivations will have been shaped by past evolution to be successful in previous environments, not for the future. In large part, members of the society will continue to pursue the proxies for evolutionary success implanted in them by their evolutionary past.

In the case of humanity, members of the global society will continue to spend their lives pursuing the positive feelings produced by experiences such as popularity, self-esteem, sex, friendship, romantic love, power, eating, and social status. And they will continue to strive to avoid the negative feelings that go with experiences such as stress, guilt, depression, loneliness, hunger, and shame.

The way in which these desires and motivations are satisfied will be influenced significantly by cultural factors, but the goals of behavior will be largely unchanged. And these goals will in turn determine the ends that will be served by technological advances, including the uses to which artificial intelligence and genetic engineering are put.

It is only by chance that the desires and motivations of global citizens will be consistent with the needs of future evolution and development. The selection processes that shaped these predispositions operated without foresight and were not influenced by the needs of future evolution.

In the much longer term, selection processes will operate to some extent between planetary societies. But these processes will not force immediate changes in values within the societies. For example, societies with values that lead them to vegetate on their planet of origin will have minimal impact and relevance in the future evolution of the universe. But unless new processes begin to operate, there is nothing to drive such planetary societies to change their values so as to align them with any longer-term developmental or evolutionary imperatives. There is nothing to stop them continuing indefinitely to shape their technology (including artificial intelligence) for the satisfaction of ‘stone age’ desires.

5.1 *The need for freedom from the constraints of past evolution*

For these reasons, a planetary society is unlikely to intentionally contribute to the success of any larger developmental and evolutionary processes until it realigns its values with the needs of those processes. The success of any wider processes would therefore seem to depend on the willingness of planetary societies to adopt pro-evolutionary goals and values (Stewart, 2000, 2001 and 2008a). Any wider developmental processes within the universe can succeed only to the extent that intelligent life is motivated to carry out the tasks that will advance the process and eventually reproduce the universe.

Here I will suggest that this fundamental change in values is likely to emerge as the members of the planetary society begin to realize the possibility that they are living in the midst of a developing universe.

Once evolution on a planet produces organisms that have the capacity to develop realistic models of their environment, they are likely to develop theories of the evolutionary process that produced them and the world about them. Eventually they are likely to begin to construct models which reveal the direction of evolution on their planet and how the trajectory is likely to continue throughout the cosmos in the future. They will begin to awaken to the possibility that this trajectory is part of a larger developmental process that has been shaped and tuned by wider evolutionary processes that may eventually reproduce the universe.

They will see from their models that selection will drive evolution along its trajectory to the point that they have reached. But they will also see that unless they now commit to intentionally advancing the evolutionary process, it will stall on their planet. Unless they intentionally align their values with those of the wider evolutionary process, life on their planet will not participate in any wider-scale developmental process in the universe. It is as if they are living in the midst of a developmental process that depends for its continued success on their commitment to intentionally advance the process.

Although their immediate desires and emotions will often clash with the demands of future evolution, it is likely that their most fundamental values will be consistent with making a commitment to pursue evolutionary and developmental goals. This is because the deepest and most fundamental values held by intelligent organisms that reach this stage in evolution are likely life-affirming and meaning-seeking.

These values will tend to motivate them to choose to act in ways that lead to the survival and thrival of the living systems of which they are part. They will see that it is only by assisting the advancement of any larger-scale developmental trajectories that they can contribute to something that has a chance of surviving indefinitely. Any other actions they could take would be futile—such actions would have only temporary effects, and in the long run, everything would be as if they had never lived. Setting out intentionally to contribute to the success of any wider-scale developmental and evolutionary trajectories is the only action they can take that keeps open the possibility that their life and actions can have meaning in a larger context.

Deriving pro-evolutionary goals from fundamental values in this way would not commit the naturalistic fallacy (Stewart, 2008a). The fallacy arises when values are derived from facts alone (i.e. when ‘oughts’ are derived from ‘is’s’ alone), not when they are derived from more general values (i.e. when ‘oughts’ are derived from other ‘oughts’ as well as ‘is’s’) [also see Wilson et al., 2003].

5.2 *Commitment before certainty*

It is of critical important to recognize that the reasons for the adoption of pro-evolutionary goals are valid irrespective of whether the organisms know with certainty that the universe is a developmental process. The reasoning applies even though they may be in the same position as humanity is at present. Humanity has discovered the trajectory of past evolution and can see how it is likely to continue into the future. But we cannot yet be certain that the trajectory is part of a wider developmental and evolutionary process that will reproduce the universe.

But organisms that reach this stage in the trajectory of evolution will realize that they cannot wait until certainty is achieved before they commit to advancing the evolutionary process. If any developmental process is to have the greatest chance of success, they need to begin immediately to invest their resources and intelligence in advancing the process. It may be a very long time until their science and intelligence is developed to a point where they can know for sure whether they are in the midst of a developmental process. It is likely that the true nature of any larger-scale evolutionary processes (particularly any extra-universal aspects) will only be discovered gradually, after significant improvements in evolvability.

In fact, it seems likely that absolute certainty may never be reached. Any large-scale processes that have shaped the development of the universe may in turn be shaped by even larger-scale processes, and so on. And it would seem that intelligence could never know if it had discovered the processes of the largest-scale, even if there are such processes—it would seem impossible to ever rule out the possibility of new discoveries and advances, or the existence of processes of even wider scale. The development and evolution of life and intelligence seems likely to prove to be in the nature of a journey without a final destination.

Given this context of fundamental uncertainty, and given that the success of any developmental process depends on their willingness to advance the process, organisms would appear to have only one option to ensure that they can participate in larger-scale processes if they exist. They will have to act as if they are in the midst of a developing universe and intentionally take the action needed to advance the process. Strategically, this means they will have to begin immediately to build the capacities they will need to participate in the wider processes if they prove to exist. They need to get into the game and to stay in it long before its final nature is clear. Only by acting as if the universe is a developmental process can they ensure that they will be able to contribute positively to any developmental and evolutionary processes that are eventually proven to exist. A key element of this strategy will be to invest significant resources in attempting to discover all they can about the larger-scale evolutionary processes within the universe and beyond.

For as long as their science leaves open the possibility that they can participate in meaningful, larger-scale developmental and evolutionary processes, they will need to continue to build capacity so they can take advantage of any possibilities that arise. Even if their universe happens to be the first to arise in which life emerges, such a strategy would maximize their chances of developing the means to reproduce the universe and initiate an on-going process.

Advancing the local trajectory of evolution by linking up with life that originates elsewhere would appear an important and productive way to build this capacity. It would enable the formation of cooperatives of larger and larger scale that are more powerful and able to influence larger-scale events in the universe (this would be of critical importance if mature universes prove to participate in extra-universal events and associated evolutionary dynamics). As mentioned earlier, it would also enable evolvability to increase through the sharing of knowledge,

intelligence and different perspectives. Life that wants to position itself to contribute to the success of any larger-scale developmental processes in the universe would appear to have good reason to build capacity in this way.

However, some theorists have questioned whether the continued expansion of the scale of cooperative organization is likely. Smart (2008) argues that the key trend will be towards the increasing efficiency and density of intelligent computation which will remain local. Although he acknowledges that the linking up of intelligences from different origins to share knowledge and intelligence would be advantageous, he postulates that this will occur in localized areas, such as black holes. Cirkovic (2008) critiques some of the arguments that have been advanced previously to support the view that the scale of cooperative organization will continue to expand, but not the main argument outlined in this paper.

It is not possible to decide conclusively between these competing hypotheses at our current level of knowledge and intelligence. However, we can be more certain that these issues will be of great interest to any living processes that decide to contribute to the success of developmental and evolutionary processes in the universe. A key priority for pro-developmental life will be to identify the capacities it should build and the actions it should take to best advance these processes.

5.3 *The inadequacy of an intellectual commitment*

However, a mere intellectual decision to align their goals with the trajectory of evolution will not free the organisms from the desires, motivations and emotions inherited from their biological and cultural past (Stewart, 2000, 2001). Their thought processes are unlikely to be able to easily modify their inherited predispositions. When a capacity for thought first arises, it would be unable to understand why motivations and emotions influence behavior in the particular ways that they do. It would not understand their adaptive purposes. Selection would therefore be likely to act against the emergence of any capacity for thought to override these predispositions, since it is likely the outcome would be maladaptive.

Thus humans do not use their thought processes to choose their desires and emotional responses, or their likes and dislikes. Humans generally use thinking and rationality to devise better ways to achieve their ends, not to determine their ends.

As a result, in order to fully align their values with the trajectory of evolution, the organisms would first have to free themselves from the predispositions inherited from their biological and cultural past. Since these changes will not be driven by natural selection, the organisms will have to develop this new psychological capacity intentionally, through changes to their 'software' rather that their 'hardware' (Stewart, 2001).

Living processes that have completed this transition will not be hostile as they move out into space. They will be motivated by pro-developmental goals, and know that in order to achieve these goals, they will need to work cooperatively with other living processes.

5.4 *The significance of the transition as a major enhancement of evolvability*

The transition to intentional evolution produces a fundamental change in evolvability and in the way evolution unfolds (Stewart, 2000 and 2001). Prior to the transition, evolution relies largely on natural selection to advance the evolutionary process. But natural selection operates mainly

by trial and error, and has no foresight—it is blind to future possibilities. Evolution's search for innovation and adaptation is driven by a particularly unintelligent process.

After the transition, intelligent organisms intentionally seek to advance evolution. When deciding how to adapt, they use foresight and modeling to take into account the future effects of alternative adaptations, including the long-term evolutionary effects. As a result, the evolutionary process itself will begin to advance intelligently and with foresight.

6. Implications for Humanity

Humanity is beginning to enter the early stages of the transition to intentional evolution and development.

As outlined earlier, our evolutionary science has established a strong case that the evolution of life on Earth is directional. We have good reasons to believe that this trajectory applies to life wherever it originates, and continues after life reaches our stage. Some of the most recent developments in evolutionary science suggest the possibility that this wider trajectory is itself part of a large-scale developmental process that has been shaped and tuned to reproduce the universe.

Until now, evolution on Earth has moved along this trajectory of its own accord. However, it is becoming increasingly clear to evolutionary science that evolution will continue to advance only if certain conditions are met: humanity must awaken to the possibility they we are living in the midst of a developmental process; we must realize that the continued success of the process depends on our actions; and humanity must commit to intentionally moving the process forward (Stewart, 2008a).

As yet there is no certainty that such meaningful larger-scale developmental and evolutionary processes exist. However, it is clear that if humanity is to put itself in a position to contribute to these processes if they do exist, it must commit to the pursuit of evolutionary goals now, long before certainty is achieved. Robust strategizing does not have to await robust science. It is possible to identify courses of action that are strategically optimal despite radical long-term uncertainty.

As humanity begins to enter the transition to intentional evolution, we are seeing events emerge that are of great evolutionary significance. Similar events are likely to occur on any planet that moves through the transition. In particular, a key milestone on each planet will be its first global scientific conference that is convened to discuss the large-scale evolution of life in the universe and beyond. On this planet, this evolutionary milestone occurred in 2008 (The conference on the Evolution and Development of the Universe, held in Paris in October 2008).

If the transition is to be completed successfully, humanity will have to free itself from the dictates of its biological and cultural past. Humanity will have to align its goals and motivations with the needs of any larger-scale evolutionary and developmental processes (Stewart, 2008a). The imminent possibility of a technological singularity lends great urgency to the development of these new psychological capacities and to the adoption of pro-evolutionary goals. Unless we develop these capacities before hand, the artificial intelligence and other technologies that enter any singularity will have been engineered to serve humanity's 'stone age' desires, as they are at present. They will not be engineered to serve pro-developmental goals.

There is abundant evidence that the psychological organization of humans is such that we have the potential to free our behavior from the dictates of our evolutionary past (Stewart 2000, 2001, 2007). The world's spiritual and contemplative traditions have discovered a variety of practices and techniques that can be adapted to develop the new psychological software that is needed. These practices can also enhance the ability of humans to understand and manage complex systems (Stewart, 2007, 2008a). Work has begun on using the tools of scientific inquiry to model and understand how these capacities can be developed and how the practices that train them can be enhanced (Stewart, 2007). The discoveries of the contemplative traditions about the human potential for enhanced modes of consciousness are being shorn of their mystical associations and are being integrated into mainstream science.

As humanity begins to enter the transition we also see the first attempts to initiate a political program that explicitly seeks to advance the evolutionary process (Stewart 2008a and 2008b).

If humanity goes on to complete this great evolutionary transition, we will have embraced a role that provides meaning and purpose for our existence.

7. Acknowledgments

Thanks to Clement Vidal and John Smart for helpful comments on an earlier version of the paper.

8. References

H. B. Adams: 1909 (1919), The Rule of Phase Applied to History. In: The Degradation of the Democratic Dogma, Macmillan, New York, pp. 267—311.

R. J. Bagley and J. D. Farmer: 1991, Spontaneous emergence of a metabolism. In: C. Langton et al (Eds): Artificial Life II. Addison and Wesley, NewYork.

D. Blitz: 1992, Emergent Evolution: Qualitative Novelty and the Levels of Reality. Kluwer Academic Publishers, Dordrecht.

N. Bostrom: 2003, Are you living in a simulation? Philosophical Quarterly 53: pp. 243--255.

D. Broderick: 1997, The Spike: Accelerating into the Unimaginable Future. Reed Books, Kew, Victoria.

L. W. Buss: 1987, The Evolution of Individuality. Princeton University Press, Princeton.

M. Cirkovic: 2008, Against the Empire. Journal of the British Interplanetary Society 61: (In Press).

P. Corning: 1983, The Synergism Hypothesis: A Theory of Progressive Evolution. McGraw Hill, New York.

L. Crane: 1994, Possible Implications of the Quantum Theory of Gravity: An Introduction to the Meduso-Anthropic Principle. arXiv:hep-th/9402104v1

P. Davies: 2007, The Goldilocks Enigma. Penguin Books, New York.

T. de Chardin: 1965, The Phenomenon of Man. Harper and Row, New York.

R. Dawkins: 1976, The Selfish Gene. Oxford University Press, New York.

S. J. Dick: 1996, The Biological Universe: The Twentieth Century Extraterrestrial Life Debate and the Limits of Science. Cambridge University Press, New York.

M. Eigen and P. Schuster: 1979, The Hypercycle: A Principle of Natural Self-Organization. Springer-Verlag, New York.

J. N. Gardner: 2000, The selfish biocosm: complexity as cosmology. Complexity 5: pp. 34--45.

_____: 2003, Biocosm: The New Scientific Theory of Evolution – Intelligent Life is the Architect of the Universe. Inner Ocean Publishing, Makawao, Maui, Hawaii.

_____: 2005, Coevolution of the Cosmic Past and Future: The Selfish Biocosm as a Closed Time like Curve. Complexity 10: pp. 14--21.

_____: 2007, The Intelligent Universe: AI, ET, and the Emerging Mind of the Cosmos. New Page Books, Franklin Lakes, NJ.

I. J. Good: 1965, Speculations Concerning the First Ultra Intelligent Machine. In F.L. Alt and M. Rubinoff (eds): Advances in Computers, Vol.6, Academic Press, pp. 31--88.

S. J. Gould: 1996, Full House: the Spread of Excellence from Plato to Darwin. Harmony Books, New York.

W. Hamilton: 1964, The Genetical Evolution of Social Behaviour. Journal of Theoretical Biology. 7: pp. 1--52.

E. R. Harrison: 1995, The Natural Selection of Universes Containing Intelligent Life. Quarterly Journal of the Royal Astronomical Society 36: p.193.

F. Heylighen: 2006, Mediator Evolution: a general scenario for the origin of dynamical hierarchies. In: D. Aerts, B. D'Hooghe & N. Note (eds): Worldviews, Science and Us. World Scientific, Singapore.

_____: 2007, The Global Superorganism: an evolutionary-cybernetic model of the emerging network society. Social Evolution and History 6: pp. 58--119.

M. Martin: 2006, The Mocking Memes: A Basis for Automated Intelligence. AuthorHouse.

J. Maynard Smith and E. Szathmary: 1995, The Major Transitions in Evolution. W. H. Freeman, Oxford.

R. E. Michod: 1999, Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality. Princeton Univ. Press, Princeton.

M. Ruse: 1996, Monad to Man. Harvard University Press, Cambridge, MA.

J. Smart: 2000 (2007), Intro to the Developmental Singularity Hypothesis. http://www.accelerationwatch.com/developmentalsinghypothesis.html

_____: 2002 (2007), Answering the Fermi Paradox: Exploring the Mechanisms of Universal Transcension. Journal of Evolution and Technology, June 2002. http://accelerating.org/articles/answeringfermiparadox.html

_____: 2008, Evo Devo Universe? A Framework for Speculations on Cosmic Culture. In S. J. Dick (ed): Cosmos and Culture, NASA Press (forthcoming).

J. E. Stewart: 1995, Metaevolution. Journal of Social and Evolutionary Systems 18: pp. 113--47

_____: 1997b, Evolutionary transitions and artificial life. Artificial Life 3: pp.101--120

_____: 1997a, Evolutionary Progress. Journal of Social and Evolutionary Systems 20: pp. 335--362.

_____: 2000, Evolution's Arrow: the direction of evolution and the future of humanity. Chapman Press, Canberra.

_____: 2001, Future psychological evolution. Dynamical Psychology retrieved October 6 2006 from http://www.goertzel.org/dynapsyc/2001/PsyEv.htm

_____: 2007, The future evolution of consciousness. Journal of Consciousness Studies 14**:** pp. 58--92.

_____: 2008a, The Evolutionary Manifesto. retrieved March 2008 from http://evolutionarymanifesto.com

_____: 2008b Evolution toward global cooperation. Kosmos Journal VII: pp. 56--8

L. Susskind: 2005, The Cosmic Landscape: string theory and the illusion of intelligent design. Little Brown, New York.

R. L. Trivers: 1971, The evolution of reciprocal altruism. Quarterly Review of Biology 46: pp 35--57.

R. E. Ulanowicz: 1986, Growth and Development: Ecosystems Phenomenology. Springer, New York.

C. Vidal: 2008, The Future of Scientific Simulations: from Artificial Life to Artificial Cosmogenesis. arXiv 0803.1087.

V. Vinge: 1993, The Coming Technological Singularity. [Presented at the VISION-21 Symposium, NASA Lewis Research Center, 30 Mar, 1993].

S. Weinberg: 1993, The first three minutes. Basic Books, New York.

G. C. Williams: 1966, Adaptation and Natural Selection. Princeton University Press, Princeton.

D. S. Wilson, E. Dietrich, and A. B. Clark: 2003, On the inappropriate use of the naturalistic fallacy in evolutionary psychology. Biology and Philosophy 18: pp. 669--682.

Short Biography: John Stewart is an Australian-based member of the Evolution of Complexity and Cognition (ECCO) Research Group of the Free University of Brussels. His main interest is in the development of an evolutionary worldview that reveals to us who we are and what we should be doing with our lives. His work on the directionality of evolution and its implications for humanity has been published in key papers in international science journals. A number of his recent papers have focused on psychological development, including the future evolution of consciousness. He is the author of the book 'Evolution's Arrow: the direction of evolution and the future of humanity'. More recently he has finalized 'The Evolutionary Manifesto'. The Manifesto outlines an evolutionary worldview and explores its relevance to humanity. It is at www.evolutionarymanifesto.com.