Richard H Webb, CPNSS, LSE Session P4.6 at 12.15 on Saturday 13 September 2014

Innovating institutions

Organisations are inventing, developing, modifying and producing new stuff the whole time. Idealistic policy statements, new pharmaceuticals, fresh sales pitches, clever ways of staying within the law. Without this stream of innovations few organisations would survive for long, and none could track the changing environment.

Today I want to look at how they do it. I will focus on that class of institution (in its strict sense as a nexus of sustained sub-routines) that is dedicated to generating working variants and inventions intended to deliver benefit. Innovations, in short.

Innovation as an institutional process

But innovation is a slippery concept so we have to be careful. Let's restrict our use of the word here to mean those processes, operating within an institution, that apply novelty to deliver benefits. A set of processes, not their outputs.

Let's further limit ourselves to institutions of several or more people. That allows us to think in terms of agent-based systems. And it helps us dodge the conundrum of individual creativity.

This way we won't count as innovations: the patter of the salesman, the lies of the politician, or the solitary alchemist's potions (although of course they are). Rather, we address the working processes of: the marketing department that generates the sales pitch, the government ministry that crafts the policy document, and the research department that takes a new drug through to Phase II.

Managing innovation

That helps. We are left working with an area of unquestionable importance. One that drives much of our evolving technology and social policy. The effective management of innovation, in this sense of the word, is widely regarded as the key to improving competitiveness, creating wealth, driving social progress, building knowledge.

This is also the habitat of innovation managers and consultants. I am one myself. But if we are honest, we are engaged in a craft, not a science. We may deliver a valuable service, but there is not much theory. We work with rules-of-thumb, judgement and experience. We have a feel for what makes a difference, what can improve the innovation process. But it would be hard for anyone, even a business school professor, to claim that this was science.

The theory gap

So this is innovation: a process exploring the design space of possibilities, and churning out a huge proportion of the stuff around us. It supplies most of our technology, tunes, laws, foods and fashions. In some manner it generates a big chunk of the observable diversity of human culture.

So it is odd to have no settled scientific explanation of how it works. There isn't even consensus on what category of process we are dealing with.

Improbability and entropy

The sheer improbability of our artefacts and processes is seldom remarked upon. Look at: a new page of text, another molecular sequence, a revised operating procedure – all of them doing some sort of job, none of them previously known in the history of the universe. These are fresh configurations of matter and information – wildly, absurdly, astronomically unlikely configurations. We may not be clear they came about. But it certainly wasn't by chance.

The Second Law of Thermodynamics is not in question: all systems progress towards disorder. Entropy must rise, and nobody doubts that global entropy does increase. It takes a lot of hot air to create something beautiful; our islands of inspiration float in oceans of perspiration.

So, against this inexorable trend, how do we account for these examples of complexity and order, this appearance of design and purpose? Is it really good enough to say that we created them ourselves?

Life's diversity explained

If we look at living things around us, we know how all those improbable working systems came about. We have a robust explanation for organic diversity.

This explanation is usually simplified to the one word: evolution. Where we see function and effective working systems in nature we are confident that those designs evolved by a process called natural selection.

Natural selection operates over long periods of time, and it has been driving biological evolution since life began. Today's complexity evolved from previous complexity. Selective forces may change, but the driving process – natural selection – persists.

Not culture's diversity

By contrast, our explanations for the diversity of the non-living world of technology, language and other social constructs – culture in other words – are hazy. There has been some kind of evolution, evidently. But most people have concluded that, whatever is driving this cultural evolution, it is not natural selection.

To me, this is puzzling. After centuries of debate and controversy, we discover a self-sustaining process in our universe that can build order, complexity, structured information, functional networked memory and sustainable sub-routines, starting with little more than hot rocks and water. But we refuse to accept that it could have built our institutions.

So innovation is not evolution?

Innovation, as I have defined it today, is the process that generates functionality by the application of novelty, and that takes place inside (and across) institutions. Very few people would countenance natural selection as its organising principle. Not that they have any clear alternative.

In Darwin's day

To understand why natural selection seems so clearly unfit for the task, we have to step back to Darwin's day.

'The Origin of Species' published in 1859 was the compression of a much longer work that Darwin had been drafting for years, and in the end never completed. He was bounced into writing the single volume by his friends, after Wallace independently came up with a description of natural selection.

Today we make much of Darwin's personal struggle, during those years of prevarication, with the religious and social consequences of his insight. But it is also clear that he was wrestling with the challenge of articulating the set of problems to which evolution by natural selection was the answer. The questions, it turned out, were more difficult to state than the answers.

Darwin's explanation – evolution by natural selection – was not such a surprise. Evolution had been around for a while as an idea, although it was still entangled with notions of social and biological progress.

And natural selection, once it had been explained, using the domestication of animals as a worked example, was almost obvious to Darwin's contemporaries. His friend Huxley famously remarked, "How very stupid not to have thought of that!".

Challenges Darwin faced

But the problems that Darwin's solution – evolution by natural selection – was expected to address were legion. And they were toxic. They included the origins of life, the descent of man, the argument from design, the role of the Creator, and the inevitability of progress (with or without revolution).

Pity poor Darwin. As it happens, he wanted to tackle a different question, on the origin of species. That was the title of his book. He meant species as persistently distinct in-breeding populations – he was, after all, the inventor of ecology and a pioneer of population thinking.

But the species concept was not what the public wanted to discuss at the time, and it hasn't been since. So, Darwinian theory has evolved to tackle other questions.

The dogma of natural selection

150 years of vituperative argument has bred a fairly dogmatic theory. It had to be stark, to see off its challengers. So, today it is a central tenet of Darwinian thinking that natural selection does not permit any external directing agency, intentionality, foresight, or goal-directed progress.

There can be no room for intelligent design, because that would let the fundamentalists off the hook they have made for themselves. Time and again evolutionary theorists have put themselves on the line, declaring that a single piece of contrary evidence would negate their theory. And each time their theory has survived, and thrived.

As a consequence, natural selection today can really be applied only to biological life: a domain where there is nobody in charge. In their war with the metaphysical barbarians, the Darwinists have built a cage for themselves. No wonder it does not convince as a mechanism for institutional change.

Social Darwinism

It was not always thus. Even before Darwin, the history and evolution of social and technological constructs had been discussed. As soon as Darwin provided a plausible engine for evolution, others started applying it to social domains. Darwin himself was not necessarily unsympathetic to this.

When they moved away from biology, of course, those early Darwinists got it badly wrong. It was too difficult to detach themselves from prevailing assumptions about the human mind and better ways to structure society. There was no tradition of social sciences to help them. As a theory, what is today called Social Darwinism soon fizzled out. It didn't work; it was a blind alley.

But as a political idea it persisted. Notoriously it was co-opted 50 years after Darwin's death by certain fascist movements, to supply a bogus justification for their genocidal intentions. So noxious were their subsequent actions, that 70 years after the obliteration of the Nazis we still recoil from any attempt to associate natural selection with the social domain.

That is unfortunate. It means that today we are allowing the historical abuses of some ill-considered Victorian speculations to constrain our thinking. Long-dead men are preventing us from reviewing the scope of the greatest organising principle discovered in the universe. We should not be giving them that privilege.

Natural selection's essentials

Natural selection is simple enough in outline. It was first illustrated by a commonplace description of the domestication of animals. If we strip natural selection back to its essentials we can express it very succinctly.

Evolution by means of natural selection is an iterated process of variation and selection, operating on an information-based system.

It is possible to capture the underlying algorithm under 6 headings:

- Information
- Variation
- Recombination
- Selection
- Replication
- Iteration

Others have expressed it slightly differently, but this list is hardly controversial.

The General Theory of Evolution

This simple algorithm is a self-sufficient way to explore design space, discover novelty, and generate persistent examples of complexity and order. It does not contravene the Second Law of Thermodynamics of course, though it is sometimes described as a mechanism for generating local eddies in the global trend towards disorder.

Evolution by means of natural selection, in these terms, is context independent. It could work anywhere. I call it the General Theory of Evolution. It is a universal phenomenon, as fundamental to our universe as arithmetic is, and we could probably have discovered it in various ways. As it happens, we worked it out when asking questions about the origins and diversity of life on Earth.

Wherever it is at work, this algorithm can generate new configurations of information, and can operate independent of external actors, and without the need for goal-direction or intention.

The Special Theory of Evolution

As we all know, natural selection has been at work on the surface chemistry of our planet for billions of years. The outcome is a plethora of improbable designs that we call the biosphere. Life as we know it is an evolved and evolving system.

In that time, life has also evolved a high level of evolvability. This sounds recursive, but it only means that living systems today possess attributes that make further evolution more likely. Competition has ensured that those systems that were less conducive to continuing evolution have died out.

Today we can see encoded information, instantiated in macromolecules. We know that these codes themselves evolved. The macromolecules build large metabolic machines, that we call organisms. We can discern lineages of these phenotypes, with parents and offspring, generations and family trees. But all of this took uncounted millions of years to settle down.

So, life is a very special domain, with a very special history, within and throughout which the General Theory of Evolution has been at work. It is a special case. So, with a wry nod to Einstein, I call the evolutionary theory that forms the bedrock of biology: the Special Theory of Evolution.

Context dependency

The Special Theory of Evolution is how evolution works in the highly evolved domain of life on Earth. By contrast, the General Theory of Evolution is context-independent. It applies in any domain. It is just an algorithm: a way of doing things. Albeit one that has shown it can self-start using some fairly unpromising materials, and thereafter keep on rolling, without requiring anybody to be in charge.

The algorithm could work in agent-based systems, to generate new configurations. It could be deployed to march through design space in search of inventions that solve pre-determined problems. There is no theoretical reason why a creator could not co-opt it.

That sounds a bit like innovation...

What about innovation?

Innovation in human institutions is an iterative process for exploring design space. Variant options are put through a series of screens to evaluate their performance. The most promising designs are further varied, recombined and tested against increasingly stringent criteria. This may be followed by the controlled release of prototypes (drafts, beta-versions, pilots). At the end of that round of selection the winning designs are mass produced and released into the world.

They may then enter a further evolutionary cycle in a marketplace of products or ideas. Some designs will disappear almost without trace, because they do not match the requirements of the competitive environment. Others may be modified, changing aspects of their design, or perhaps copying elements from other more successful designs.

Classic examples of this include so-called high-throughput screening.

High-throughput screening

Many of today’s standard pharmaceuticals were designed by exploring vast libraries of unrelated chemicals for their performance as drugs. These were often chemicals from other industries: photography perhaps, or explosives.

Researchers use a battery of simple pharmacological tests to search for activity, hunting blind, without knowing the structure of the chemical compounds. Only where they find a promising effect will they go to the trouble of looking at its chemical structure.

This structure is then deliberately varied, synthesising many versions, and testing them in turn. Promising leads are further varied, and sometimes recombined, for testing again. At the end of a process in which the screens became increasingly stringent, and the chemists run out of potential variations to make, the drugs are either abandoned as unsatisfactory, or are brought to market.

Living examples

This is a powerful problem-solving methodology. It is particularly effective at tackling unknown problems with the minimum of preconceptions. So we should not be surprised that there are plenty of examples to be found in the living world of evolution in this sense – obeying the General Theory of Evolution, but not the Special Theory.

Adaptive immune system

The mammalian immune system is one. Evolvable information is encoded in amino-acid sequences. Variations in the critical sequences are generated in vast numbers, combined into Y-shaped antibodies, and sent into the bloodstream in search of foreign bodies. Those that successfully bind to and kill antigens send back their information. Their sequences are mass-produced and copies of these antibodies are released into the bloodstream. These in turn bind to antigens, and by iteration the best killer molecules come to dominate.

The adaptive immune system's responsiveness and power to generate new designs make it one of the best ways that a mammal can defend itself against the capacity of its pathogens to evolve offensively.

Foraging ants

A different example can be found in the foraging strategies of certain ants. They operate as an agent-based system. Each ant follows simple rules. It has a home (the nest) and food that is clumped unpredictably and temporarily across an unknown environment. The ant has limited information-processing powers, and some basic onboard sensing and navigation equipment.

Ants set off on an increasingly random walk, generating a variety of paths but leaving no trails. Those that encounter food pick it up and head home, leaving a directional chemical trail that fades over time. Unladen ants that encounter such trails head towards the food, but may wander a bit. Laden ants use trails they encounter to modulate their own choice of homeward route.

The result is a pattern of foraging trails that evolve over the course of a day to exploit patchy food resources efficiently. It is a robust and efficient way for a crowd of cooperating creatures to exploit a transient and patchy environment, without the need for much memory or computing power.

Just an analogy?

But these examples from the realm of biology do not always persuade. They look like just more ways in which evolution works in biology. They are not evidence that such processes might apply to for example institutional innovation.

Similarities between the domains are widely recognised. But they are generally viewed as analogical. Time and again people have looked to see if natural selection could be the mechanism behind observed evolution in non-living systems, such as institutions. Time and again they have concluded that it cannot be.

There is no consensus about what kind of evolution this is. But most people would find it perverse to claim that it really was evolution by natural selection. Such a conclusion does not seem to survive even cursory investigation.

If hard pressed, objectors might concede that something that looks a bit like natural selection may be part of the design process. But it is not true natural selection.

Reasons for failure

There are many reasons why people fail to find natural selection at work. (That is interesting in itself.) Some of these are good reasons. Others are clearly based on mistaken analysis.

Sometimes people misunderstand what is evolving.

The theory of the firm

The ‘theory of the firm’ is an established field of economics that looks at the dynamical change over time of institutions like firms. But the firm is not a helpful starting point for examining evolutionary mechanisms. That is because firms are insufficiently coherent entities. An experienced evolutionist would not expect them to be tight enough units to be candidates for natural selection. To misunderstand this is a classic ‘group selection’ error.

Evolution of markets

Markets evolve, but not by natural selection. They are too broad a category. This is a ‘level of selection’ error, like a biologist looking for the evolution of an ecosystem. Ecosystems do evolve, in the sense of changing over time. But their dynamics is the consequence of interactions of many species that are themselves evolving. Ecosystem changes are driven by evolution working at the gene level via living organisms.

In the same way, markets change over time, so you might say that they evolve. But the study of changing markets will give little clue of underlying organising processes, if these are happening at a lower level. Natural selection may be operating within markets, but it does not operate on markets. Natural selection can design a bird’s wing, and, who knows, perhaps a plane’s wing. But it could never design an ecosystem or a market.

Anthropocentric

Sometimes people overstate the role of humans in the evolving mix. This mistaken perspective has a long history. Time and again our understanding of the universe has only advanced once we moved the human away from the centre of the picture.

So you often see fitness (of technological innovations, for example) measured in terms of benefits provided to humans. As if an invention (perhaps a piece of useful computer code) could be considered successful only if it helped humanity. A better measure of fitness would be frequency and persistence (thus computer viruses may be considered highly successful).

It can be difficult to let go of the illusion that we (or our gods) are in charge. It took astronomy a long time, and biology longer. Some corners of the social sciences still have some way to go.

Lamarckian inheritance

Sometimes people fail to understand that they are looking for a mechanism that is able on its own to explore design space. Natural selection needs no 'black box' to give a guiding hand or set the standards for its design choices. The external environment is enough.

So-called 'Lamarckian evolution', or the inheritance of acquired characteristics, has often been presented as if it were a potential mechanism. But (if it is anything at all) Lamarckian evolution is a kind of learning. To enable it to drive change would require the presence and regular maintenance of various cognitive or other information-processing tools. It cannot be a candidate mechanism.

The famous example of the giraffe that stretches its neck and thereby leaves longer-necked offspring is not conceptually impossible. But that giraffe would need some means to detect salient signals in its environment, and this would have to be linked to another system for generating and prioritising desirable attributes for future fitness, before it could even start to engineer its offspring's DNA.

In addressing how the giraffe might evolve, Lamarckian approaches beg the far harder question of where those imagined information-processing tools come from, and how they might be evolving.

Applying the Special Theory to culture

Above I have described some classes of error people make who do not properly understand evolution. But plenty of people with a sound grasp of natural selection still reach a similar conclusion.

If you believe the text-books, the Special Theory is The Theory of evolution. Naturally, as the Special Theory applies only within the highly evolved domain of life on Earth, it is unlikely to pass muster as a plausible mechanism for evolution in non-living systems (cultural evolution, in short).

This is not fair. Unlike life, culture has not had billions of years to evolve digital information in a standard code. So there is no reason to expect to find the equivalent of genes, or even to find a separation of genotype and phenotype. That acid test for the biologist, the distinction between replicator and interactor, is a feature of a highly evolved system, which most of culture is not.

Memes

Memes were postulated as the cultural equivalent of genes. Clearly, there are some memes to be found – words are a good example. You can even watch them evolve. But memes are probably not very common. Memetics itself is too close to an analogy of biological evolution to be very helpful or persuasive. Its followers hoped to make it the science of cultural evolution. They have not made much progress so far.

Applying the General Theory

Let us return to that category of institution that is the subject of this talk: those that are there to generate the effective application of novelty. The big question remains unanswered: is natural selection, in the sense of The General Theory, a plausible mechanism?

First, let's be clear, do we mean: the only mechanism, the primary mechanism, or just one mechanism of many? For biologists that distinction is critical, but it matters much less to us. Any understanding we can gain on what is driving innovation would be useful. We don't need a single mechanism to the exclusion of all others.

But what would it mean to say that a system was partly designed by natural selection? How could we partition the effects of more than one algorithm – say natural selection and something intentional and forward looking? Would it imply a different exploration pathway, or a different outcome? Would these processes be competitive or additive? How could you tell?

Domesticated species

The domestication of animals and plants gives us a classic example of a mixed or hybrid process. Let's look at the canids.

Wolves are wild animals that evolved by natural selection. In understanding their design and diversity we do not need to look for any other forces.

Dogs are domesticated wolves. To understand them, we need to know that human choices shaped the selection pressures on them and enabled otherwise improbable hybridisation between lineages. As a consequence of human actions, wolf genes have been rearranged in fresh ways across the dog breeds.

Some of this has been deliberate and intentional. There are rule books of required features for dog breeds, and even genetic tests to help identify them. That is an explicit and intentional selection pressure.

Most selection pressure has been less intentional. Cute puppies, with expressive faces and friendly demeanours were more likely to be kept. Brave and loyal dogs were more highly valued and ended up with more surviving offspring.

And much of the selective environment has been nothing to do with human decisions. Dogs able to live off human waste materials, and to cope with human-borne diseases, survived. Today disease resistance and dietary requirements differ significantly between wolves and dogs.

And let us not forget the dingos. These were introduced to Australia as domesticated dogs, but went feral. For thousands of years they have run wild, scavenging around, being killed by, and occasionally preying upon humans. How do we want to describe their evolution?

Human influences

I have discussed dogs, but it might have been show roses, or insecticide-resistant cockroaches, or Monsanto's RoundUp Ready corn. In all these examples, humans have intervened in various ways, intentional or otherwise, to change the design of organisms, to shape the course of their evolution.

Let's refer to our checklist for the algorithm that defines the General Theory of Evolution, to see how this fits these new examples:

- *Information*
  In these cases, the information system being evolved is DNA
- *Variation*

Most of the variation is generated in the usual way by random mutation. Sometimes (as in roses) this is speeded up to give breeders more to work with. Occasionally as in GMO corn the variation has been deliberately chosen and added to the breeding mix.

- *Recombination*
  Most of this happens in the standard biological way, by means of sex. Some of the cross-breeding is deliberate, looking for features in two parents to combine in one offspring. In transgenic crops that cross-breeding is quite exotic, picking up a gene from one species and introducing it to another.
- *Selection*
  In all these cases humans are shaping selection. But much of the selection is either not intentional, like insecticide resistance, or applied at a fairly high level, like the show roses (which varieties do best when grafted onto root stock and grown in flower beds?). Some of the selection is delivered by the market place, an agent-based system under nobody's control.
- *Replication*
  Often this is beyond human control, like the cockroaches whose populations flourish under the same conditions we enjoy ourselves. Sometimes it is engineered, like the breeder who propagates thousands of roses to distribute to retailers. And the market plays an important part, where success determines future production schedules.
- *Iteration*
  All these are iterated examples.

Robust organising principle

I find it telling that the General Theory of Evolution can be used in this way to partition the forces that shape domestication of plants and animals. It suggests to me that, as a guiding principle, the General Theory is robust to many human design processes.

Notice how we can be explicit that intentionality is at work. It is expressed as shaping the selective environment, or structuring the breeding structures. It makes very little difference to the model that variation has been deliberately enhanced, or recombination engineered.

Notice too how much of the design remains outside human control. I will illustrate this by expanding on the example of genetically engineered maize.

RoundUp Ready (RR) corn

Monsanto took a herbicide tolerant (HT) gene from another species and inserted it into the genome of a high-yielding maize variety, then bred it up in vast numbers to be sold to farmers all over the world keen to reduce their operating costs. This looks like engineering, not evolution.

Let's look at some of the detail:

Monsanto had a number of promising genes (traits) to choose from. They also had a number of germ-lines (new maize breeds) as candidates for hosting the new traits. Those choices were made by selecting the ones that performed best in trials. [ie variation-recombination-selection]

That HT gene was blasted under high pressure into the cells of maize, and those cells were bred up to plantlets and tested for tolerance to herbicide. Only where the gene had inserted itself into the genome in a place where it would be expressed would the plantlet survive. [ie recombination-selection]

These promising plantlets were then bred up to full plants and tested throughout their life-cycle to see how the engineered crop would perform on yield, nutrition value, etc. The plants differed in the expression of some of their other genes, according to where the new HT trait had inserted itself. After extensive field trials a small number of germ-lines were chosen. [ie replication-selection-replication-selection]

The new HT strains were bred up in very large quantities to generate the seed stock to sell to the marketplace. Distributors chose the RR strain that was available in the necessary quantities and that they felt was best suited to their needs. Individual farmers decided whether they would sign up to Monsanto's deal (rather than buying their seed and inputs from elsewhere). The following year those farmers again made the choice. [ie replication-selection-selection-replication-et seq]

Non-living systems

Some might suggest this is special pleading: that these examples fit so well to the General Theory only because I have chosen living systems. Also that I am understating the importance of non-random variation and human intentions. Let me deal with these objections in turn.

True, these are living systems. We know the basis and coding of the information that defines them: it is DNA. In other cases (the design of turbine blades perhaps) it is something else (spatial mathematics perhaps).

My response is, so what? Most innovation explores design space using other information systems. But there is nothing magical about our DNA triplet code. It just happens to be the one that evolved here by natural selection, some billions of years ago, probably in a marine thermal vent.

I have used the example of domesticated species (inspired by Darwin) to start the reader off on familiar ground. Having accepted that this process describes what is happening in biotechnology innovation, we can see that it may be applied to other areas of innovation.

Intentionality

So, what about human intentions? In particular, what about pre-determined goal-setting, and the description of a solution ahead of discovering a route towards it? How do they affect the nature of the underlying process?

I think this is an important challenge. I have two lines of response to it.

Screening for potential future performance

The first, as described above, is that intentionality can often be modelled as selection. In other words, intentions can be expressed in the evolutionary algorithm as selection pressures. For example, if it has been decided that a new widget has to be manufactured at below a certain cost, then the institution responsible for innovation sets up various tests that are designed to screen out options that, were they to be scaled up in the future, would fail the low-cost criteria.

Is this incompatible with the General Theory of Evolution? I don't see why. Even under the Special Theory, living creatures routinely use proxy tests to evaluate potential partners for their ability to deliver future performance (as mates with good genes, or as hard-working fathers, or as symbionts that will do their share of metabolic work). Intentions can shape the direction of evolution without changing the mechanism.

Post-hoc intentions

The second response is more uncompromising. I believe that most intention leads to nothing, and that most success stories are post-hoc, self-serving and not to be trusted.

Most inventions fail. Among those that succeed, most deviate significantly from their original intentions. Very few original conceptions make it through unchanged. Even among those that do, most are quite different in execution from what was originally intended. Many products (drugs are a good example) are used for purposes quite different from their original intentions.

People with success stories tend to reverse the arrow of causality. Few admit that they dreamed up a host of wildly implausible designs and waited to see which might work. Instead they claim that they designed a successful product, conceiving from the beginning how it would work.

Every successful design, whether it be an advertising jingle, a hi-tech glass coating, or the speech given by a government minister is the survivor of a very large number of antecedent variants that fell victim to Darwin's 'struggle for existence'. If each of these failed variants is also to be considered a purposeful design, then perhaps it would be better to describe intention as a means of generating variation for natural selection to work on.

If so, would this be considered as directed variation? And is such direction compatible with the General Theory of Evolution?

Directed variation

Directed variation is often presented as a fatal negative for innovation as a Darwinian process. Every evolution textbook asserts the necessity that variations (mutations) be random with respect to their future purpose. Evolution is a blind watchmaker, remember.

Well, not quite. Genetic mutations are not random, neither in their structure nor in the likelihood that they will give rise to viable design variants. Mutation is itself a highly evolved attribute.

The lowest mutation rates are found in genes that code for proteins with least scope for improvement (such as the oxidation enzyme cytochrome-c). The highest rates are found in genes coding for proteins where it might be useful to have an abundance of varieties (such as odour receptors in mammals).

Also, mutation rates can be modulated as a response to the environment. Move a colony of bacteria to a new food source and they will mutate the genetic region that codes for digestive enzymes, so some at least of their children can eat the food their parents couldn't.

Directed recombination

Recombination is a special kind of variation, taking elements from two or more places and bringing them together. In the living world this is mostly done by sex (in eukaryotes) or the horizontal transfer of genetic material (in prokaryotes).

Recombination rates are highly responsive to the environment, and to rules of thumb about future prospects and the future environment. Biologists don't consider this to be foresight or purposive design. It doesn't threaten their Special Theory.

So: stressed trees flower splendidly, as if knowing they will soon die. Aphids, after generations of abstinence, grow sex organs at the first signs of winter. Stunted male trout pretend to be small females and sneak in among the large spawning males to do their stuff.

Furthermore, in nature the choice of what traits to recombine is anything but random. Everywhere organisms go to enormous trouble to select genetic material that shows promise. Much of life's richest show, from the flowers to the brightly coloured coral fish to the human dance floor is about using proxy tests to choose what genes to put into the mix for the next generation.

So, it seems that directed recombination is compatible with evolution by natural selection.

Literally unpredictable

The future is unpredictable - literally. The environment is a chaotic system whose non-linear interactions mean it cannot be modelled with any accuracy more than a very short time ahead. The wisest of us cannot tell what tomorrow will bring.

Markets, too, are literally unpredictable. People strive to appeal to them, or struggle to control them, but even at the macro level we are only placing bets on an unpredictable future.

The four-dimensional space of all possible environments is close to infinitely large. The most sophisticated design system in the world is not very reliable at exploring even neighbouring design space. Institutions that innovate are exploring design spaces that cannot be predicted or modelled.

Building innovation capability

So what is the use of all out human creativity and ingenuity, all our institutional capabilities and knowledge, if it is not the motor of innovation?

We can use these skills and processes to structure our innovating institutions in such a way as to improve their chances of success, raise their efficiency, and increase the impact and quality of their outputs. In other words, we can raise the innovation capability of our institutions.

This is where the innovation managers come into play:

- They can cast more widely for variation, and co-opt the best minds to spot opportunities for recombination.
- They can set up their trials more intelligently, linking these to the most relevant parameters.
- They can raise the system's responsiveness to change in the external competitive environment, especially by improving screens and modulating selection criteria.
- They can reduce future uncertainty by identifying and monitoring leading indicators, and linking these to the design of their tests.
- They can ride promising trends by tracking other fields, and borrowing skills and solutions from elsewhere.
- They can iterate faster and cheaper, and link discovery, commercialisation and production into a seamless pipeline.

In other words, they can do what they have long been doing without realising, which is to exploit the underlying driving force of innovation: evolution by natural selection.

Only perhaps, they can now be told that this is what they are doing.

Conclusions

I am left with the conclusion that evolution by natural selection is the primary organising principle involved in innovation, or the process of applying novelty to deliver benefit, as conducted by specialist institutions.

However, we have to move away from the textbook descriptions of the evolutionary process, because most of the rules laid out there refer to the very special case that has generated the complexity and diversity of life on Earth.

The simple evolutionary algorithm embodied in what I call the General Theory of Evolution is a powerful way to understand the processes that drive innovation. This is especially clear when we study institutional agent-based systems operating over long periods of time and exploring unpredictable design spaces.

It can be applied with advantage as a heuristic to assist in the direction and management of innovation. It is not an analogy, but a mechanism. Not necessarily the only mechanism, but a dominant one.

It may be influenced but not invalidated by the various protocols that people use, both rational and irrational, to manage innovation. These, like directed variation, engineered recombination, goal-setting and future visioning, can be overlaid on the evolutionary process.

Thus, human activities add new criteria, set the parameters of the selection screens, limit the variation and recombination options, and govern the allocation of resources to the working system. But the exploration of design space that is the heart of the innovation process remains an evolutionary process.

It is our job as innovators to harness the power of evolution.