Draft

Concepts in Innovation

Comments to: ***jedbarker@iinet.net.au***

Chapter 5: *Knowledge and Learning Part B: Learning Dynamics*

"We are what we repeatedly do. Excellence, then, is not an act, but a habit." – Aristotle

Summary

In the previous chapter we looked at the different modalities of *how* learning occurs. In this chapter we will look at *how much* learning occurs and *how good it is* and the *rate* at which it occurs.- ie- the *quantity*, *quality* and *quickness* of learning..

We see that learning is often motivated by a drive – perhaps innate – to improve the efficiency and effectiveness of the innovation so that it more fully expresses its intended purpose.

While the actual learning commences in the people involved, it can be applied directly to the innovation – as *product improvement*, or to the production environment and beyond – as *process improvement.*

We call the achievement of learning the *performance* of the innovation. Depending upon the complexity of the innovation, its performance can be measured quantitatively or semi-quantitatively, in *performance curves* and *learning curves*.

We also find that although our main focus is on learning, forgetting alos occurs. This can be useful or dysfunctional.

Learning as a dynamic concept

In most of the cases described in the previous chapter, the *learner's* learning of the *knower's* knowledge is, of course, often less than exact or instant. Further, the *knower's* knowledge is often less than perfect or complete in the first place. By *perfect* we mean that the *performance* of that knowledge would conform to a preconceived notion of the *ideal.* Like Plato's[1] ideal forms, they might not be real or obtainable, but we can imagine them and aspire to them. While, in the case of learning *codified factual knowledge*, exact precision *is* possible in simple cases (*eg*, knowing that two times two equals four is possible) *complex and procedural knowledge* can be *workable* or *acceptable* while being somewhat less than perfect. We also know that, with repetition – ie *experience*, our innovation usually becomes more workable – ie it's *performance* improves. In this chapter we will examine how this improvement in performance occurs.

We know that innovation is more than a single event – in fact we have defined it as *a process* – a number of interactions (events) between the innovator and innovation over time and space. It is a process of *becoming* – first, in creating an initial attempt or 'prototype' – and then, with the application and embodiment of more and more knowledge, the innovation's later models more closely fulfill the purpose envisaged by the innovator – that is, the *innovation acquires a greater capacity to act as initially envisaged by the inventor*. The application of knowledge is in the form of changing the elements and relationships within the system being innovated and also between that system and its environment. Some of the ways that this can happen include:

- More components (elements) are added;
- Some redundant elements are removed;
- New and better components replace older and less effective components;
- The relationships between elements may be reconfigured;
- New ways of effecting relationships are used.

The outcome of this process is an *innovation that works better than our first attempt*. Of course, in many cases, an attempt at innovation may be abandoned – despite the above efforts – because the innovation still does not achieve its intended purpose of working as well as envisaged.

Thus the process can be seen as *systemic learning*:

- The system being innovated (the 'innovation') 'learns' from application of knowledge;
- The innovator(s) learn how to apply the knowledge to the innovation and learn the results of that application;
- The users of the innovation learn how to use it and communicate information on its usefulness to the innovator; and so on.

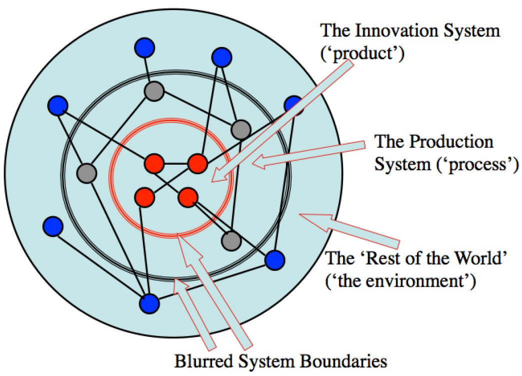

Fig 5.1: The total innovation system. Production of an 'innovation' occurs in a 'production system', which is, turn part of a wider environment of 'the market' and other factors that might influence the innovation process.

But as a *process*, this does not all happen at once: it is not enough to just say that the process is rate-limited by available money or physical access of hands to the innovation, *it is rate-limited by the capacity of the whole system to learn – the* *speed with which the system to generate, transmit and assimilate knowledge*. Only part of that learning is how to get the finances and how to get the hands working productively – which we will deal with in detail in later chapters.

Learning and Research and Development (R&D)

We should bear in mind that, although, when we are innovating, we are usually mainly in an exploratory modality (Mode 5 – Solitary/ *Focal/Exploration and* Mode 7 - *Interactive/Focal/Exploration*), we may also be *imitating* and accidentally *exploring* and so on. We conject-and-try and cut-and-try, and sometimes have to undo our work. Often social and cultural issues have to be discovered (*tacit-to-explicit*), articulated and responded to, as Nonaka and Takeuchi[2] described in the previous chapter. As a result, the application of knowledge to our idea may go in fits and starts, but with motivation and access to resources, progress (towards fulfillment of purpose) is made. This is, essentially, the process of *Research and Development* (R&D). R&D, as it is usually called, is generally thought of as a 'scientific' process, that is, confined to the creation of new 'technologies' – such as chemicals, electronics, aircraft etc. From our perspective, R&D can be applied to the creation of *any new system* – that is, R&D *is* innovation, whether it is applied to a widget, a chemical plant, a call-centre or a cake-stall.

Towards a *dynamic* view of Learning

The discussion to date on the nature of knowledge and learning has been a somewhat *static* view. With declarative knowledge a single *fact* can be verified as having been learnt – when tested, the aspiring learner is either right or wrong – we either *do* know something or we don't. We either know who is the president of France, or not, we either know the specific gravity of gold, the refractive index of quartz crystal or the speed of light in a vacuum, and so on. Similarly, procedural knowledge can also be verified – we can either demonstrate that we can perform some action that we have tried to learn, or that we fail to perform it. *eg* We can show that we can bake a cake, operate a spread-sheet program, solve an algebraic equation and so on. However, this simplified approach to knowledge does not fully take into account a number of other aspects of learning:

As we discussed in Chapter 1, an innovation results in either a new system or an existing system with new relationships to other systems. How do those new elements and/or relationships come into being and into place in the system so that it *works* –ie *fulfills the intended purpose*? And further, how *well* does the new system work? How long did it take to learn that fact, or sequence of facts, like a poem or speech. How long and how much effort was required to get a declarative sequence completely correct? As we know from everyday life, we can improve an initial working system by learning more and applying that learning. We will explore what that means in terms of the model that we have developed to date. We will address:

- The efficiency of acquiring existing knowledge;
- The process of creating new knowledge;
- The efficiency and effectiveness in creating new knowledge;
- The completeness of the knowledge created or acquired;

- The efficiency and effectiveness in applying new knowledge;
- The effort of re-working errors
- The re-remembering of needed information that has been forgotten.
- To what extent does the system 'work' ? ie fulfil its intended purpose?

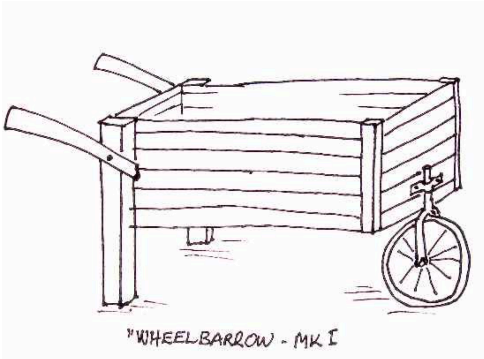

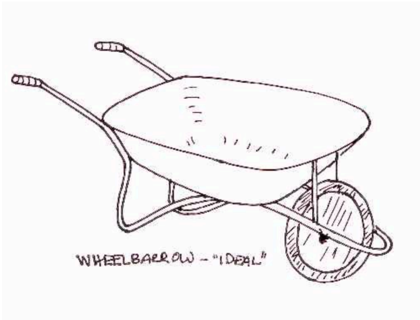

Fig 5.2 a) and b: A wheelbarrow that starts as an assembly of existing elements and finally becomes an optimal arrangement of a small number of new elements.

Lessons from a wheelbarrow

Let's illustrate this learning process with a simple example of a system that starts from existing elements and undergoes changes as we learn at different system levels. This example illustrates many, if not all the features of learning in *any* system, whether it is an artifact or an organization.

Let's assume that I've invented the 'wheelbarrow' – I have envisaged a system whose elements comprise a large open container with a wheel attached near one end, two handles at the other end and two stays underneath to give it three-point stability when it is not being pushed or pulled. Its purpose can be stated:

- *The purpose of the wheelbarrow is to enable things to be moved from place to place with less effort than carrying them.*

We need to add to that we want our wheelbarrow:

- *to be more efficient and effective than other means.*

This additional statement provides our *motivation* to learn. Chapter xxx will deal with the concept of *motivation* in greater detail.

Like most prototype systems, our wheelbarrow is a 'novel combination' (Schumpeter's words) of some pre-existing elements: a wooden packing crate, a wheel and fork from an old tricycle, two axe-handles from the garden shed and two short lengths of 4-by-2 that were in that heap of spare timber that everyone seems to have. All fixed together from saved screws, bolts and nails. It works!

But what does 'it works' mean? It works in the sense that it fulfills my envisaged purpose of being able to move things from place to place with less effort than carrying them, but we find that 'works' is a *relative* term in that it is possible to make a wheelbarrow that makes the job even easier than our first attempt – and might cost less as well! (Note that although I was able to make a working prototype from scrap material, any improved model would probably need purchased materials, or materials that take effort to collect in volume.). There are many designs of wheelbarrows and the 'Version N.0' in *Fig 5.2b)* works a lot better than my 'Version 1.0' prototype – it is factory-built and has a sloped, high-density plastic container; it is lighter, more stable (through R&D on stability), more durable, easier to clean and easier to push because of the large pneumatic tyre. But more than this, it uses fewer materials and is cheaper to make in large volumes using automatic moulding, pressing, welding and bending machinery.

Both of these systems are called 'wheelbarrows'; both have the same purpose and functional elements, but Version N.0 is more efficient and effective than Version 1.0. So when I said 'I know how to make a wheelbarrow' I was saying something different from 'I know *at this moment* how to make the most efficient and effective wheelbarrow'. I had *some* knowledge – I showed that I had *the capacity to act*. But the 'ideal' wheelbarrow (Version N.0) seems to show that more knowledge can be applied to the *basic* idea of a wheelbarrow.

These two wheelbarrows illustrate all the basic concepts that we have reviewed above. In my *Version 1.0* model, I demonstrated that I had factual (*declarative*) knowledge about certain materials and their properties (eg 'wood is strong' and 'wooden pieces can be fixed together with screws' *etc*.) I have demonstrated procedural knowledge by cutting the pieces of wood and screwing them together in a way that makes a simple wheelbarrow. Some of that procedural knowledge was *explicit* (I could instruct someone else how to cut and assemble a wheelbarrow); some of the procedural knowledge was tacit (how I actually use a saw, drill and use a screwdriver effectively).

Let's now assume that having satisfied myself that my prototype wheelbarrow 'works' and I have shown it to my friends, who now also want one! I'm encouraged to go into limited manufacture of my innovation. But – although messing around for several weekends to make Version 1.0 was fun (I charged out my time to 'lifestyle'), I will have to take a lot less time and materials if I'm going to make a lot of them without losing money. Let us now follow the learning processes as I embark on wheelbarrow manufacturing.

The design and construction of the *Version N.0* wheelbarrow, in theory, *could* occur in one development stage. On seeing the crude *Version 1.0* design, an industrial designer might know that extruded plastic, steel tubing and pneumatic tyres are the very best materials to use and they might know enough physics to design the positioning of the wheel and handles to achieve the very best balance and maneuverability possible.

But in reality, progress to Version N.0 would go in steps and stages as more knowledge was applied and tried to see if it 'worked'. Even without changing one aspect of the design, explicit and tacit knowledge could be applied to *Version 1.0* to make it more quickly – explicit knowledge could include the sequencing of construction and tacit knowledge could include sawing and drilling more quickly and so on.

In reality, wheelbarrows have 'evolved' with the application of factual and explicit and tacit procedural knowledge of materials and processes over a long period[3]. And there are many kinds of wheelbarrow, each being the most efficient and effective *for a particular purpose*. And in reality, this 'evolution' is more likely to occur as a series of steps. For the moment, for the sake of clarity, we will assume a simple, continuous change. We will deal with the issues of change in greater detail in the following chapters.

Knowledge and learning in practice
Let us continue to use our example of the development of the wheelbarrow to look at the different ways that knowledge can be applied to improve our initial innovation. We need to look at the production at two main systems levels- the *product-system itself* (*ie*, the wheelbarrow) and the system that we use to make it- *the process-system*. (See *Fig.5.1*) There are four ways that we can apply new knowledge to making our product, as depicted in Table 5.1. We can keep the product the same, or change it and/or keep the processes the same, or change them.

Product/Process	**Same**	**Change**
Same	Speed Accuracy	Tools Processes
Change	Speed & Accuracy within stepwise change	Tools Processes

Table 5.1: The performance of a system can be improved by changes to either, or both the product and its production system (process).

In the first case (Cell: *same product/same process* of Table 5.1), we 'freeze' the initial product design of the wheelbarrow and set about to make it with less effort or in less time. This path, in turn, can go in two directions: In the first, both the product is frozen and the way we make it (process) is also frozen (*same/same*), in that no new techniques are used. Within these constraints, we learn how to measure and cut the wooden parts quicker and more accurately, we learn how to source and assemble the parts quicker and more accurately – that's about all we can do.

In the second path, we 'freeze' the product, but change the way that we make it (*same/change*)- we might use a power saw rather than a hand-saw, set up a special bench and jigs. We might source the materials in bulk and make batches of each component at one time – we might organise the labour like Adam Smith's 'pin factory'[4] – each person doing a simple task more frequently and therefore more rapidly (although this is an arguable point, as much of the increased productivity in Adam Smith's pin factory example is due to specialised equipment that is fully utilised – see Ref[5]). And so on.

The second main path is the 'effectiveness path'. In this case we retain the same purpose for the system, but set about to achieve that purpose to a greater degree (*change/same*). This means 'improvement' in design, which we then put into effect to improve the *performance* of the system/wheelbarrow. As we can see from the 'Version N.0' wheelbarrow above, the use of plastics, steel tubing, pneumatic tyres, reshaping the container, *etc* means that our wheelbarrow can either take larger loads or the same loads with less effort than our Version 1.0. This path inevitably means that the processes used must also change. However, *within* each product change (*change/change*), there will be changes to the process. We will see in detail in Chapter XXX that the two paths are essentially the paths of *process innovation* and *product innovation* respectively.

Knowledge, Learning and Performance

In each of these cases of our metaphorical wheelbarrow-making we are using both factual (*explicit/declarative*) and *procedural* knowledge (both *tacit* and *declarative*). We find that by repeating an experience – say sawing the wood – our actions usually speed up and/or take less effort. Part of this is *tacit* knowledge – our bodies usually seem to learn how to do things quicker without us even thinking consciously about it. But part of this improvement can be *declarative* or *explicit* – we might ask someone, or read a book on 'how to' do something better or 'the right way'.

With repetition, our initial mindful (conscious) actions, which may be segmented into a number of discrete steps, become a seamless, *tacit*, whole. Anderson[6] (2004) calls this 'knowledge chunking'. We are all familiar with this process from an early age, through the development of reading, music and sporting skills).

With improving the design, we might consult a mechanical engineer, whose general *factual* knowledge about loads and balance are combined with her *declarative/procedural* knowledge of engineering calculations to redesign the wheelbarrow. A materials engineer has *factual* knowledge that certain plastics are strong enough and durable enough to use for our purposes. A plastics injection moulding company then uses its procedural knowledge to make the container – although large numbers are required to justify the development of special moulds. Standard-sized pneumatic tyres are sourced in volume. And so on.

Together, we can call these improvements in efficiency and/or effectiveness improvements in 'performance'. Performance is also 'fulfilment of purpose'. What we have shown is that there can be a number of different measures of performance, depending on the system level that we are looking at. We can have improved *product performance*, improved *process performance*, or both. Of course, ultimately the end-user is most interested in the product performance. However, as the manager of the innovation process, we must also pay attention to 'process issues'. But in the light of changes in the product, the *producer* might have to review and modify the original purpose, which may have been to make a wheelbarrow to a particular design as cheaply as possible, or to make a human-driven device that minimises the cost of transporting a wide range of materials, or to have a wheelbarrow made of recycled or environmentally-friendly materials. Further, the user's criteria for efficiency and effectiveness might change in the light of the design changes. Usually, the criteria for 'performance' are a complex mixture of subjective and objective factors. We will revisit the evolution of 'purpose' in Chapter xxx.

Learning and Performance curves

The dynamics of improved performance are well known, and are often plotted on *learning curves* or *performance curves*. We will first look at performance curves. *Fig 5.3* is a generalised and *idealised* plot of *performance versus time*. In practice, the progress towards improved performance is often a series of 'steps and stairs' (often with some backwards steps. See Fig 5.5.), rather than an infinitely smooth upward progression.

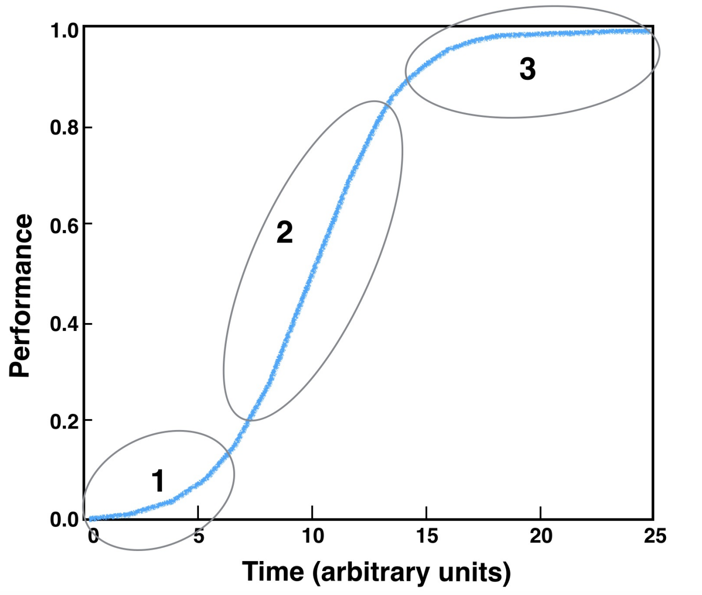

Fig 5.3: An idealized *Performance Curve.* The vertical axis depicts the fraction of the ultimate maximum performance achieved at a particular time.

We notice that the curve comprises three main regions: First, a rather flat region (1 in *Fig 5.3*) is followed by a very rapid rise (possibly exponential) in performance in the second region (2) and then in the third (3) region performance seems to slow down and flatten out, perhaps to a maximum level. Let us look at these three regions of *Fig 5.3* in more detail:

Region 1: Improvements in performance are relatively slow. At first, the innovator is 'getting to know the ropes', often doing some things for the first time, taking care and thinking things through – such as 'measuring twice- cutting once'.

Region 2: The project starts to 'click together' (the knowledge 'chunks' are becoming a seamless whole) and the time-rate of performance improvement is rapid. Relationships between different parts of the system are identified and dealt with synergistically. The trends of previous experiences are realised and extrapolated. In general, there is a positive feedback loop between the last experience and the next action.

Region 3: The rate of performance improvement slows as ways of improving become less obvious and trade-offs in improvement in one area are made at the expense of decline in other areas. In general, the purpose – and therefore the scope – of the system, as conceived may be limited.

Readers may be familiar with *Fig 5.3* as the 'generalised logistics curve'. It occurs in many places in both natural and invented systems. *Fig 5.4* shows four different curves, illustrating that learning towards a particular goal can be faster or slower, depending on the goal and the circumstances.

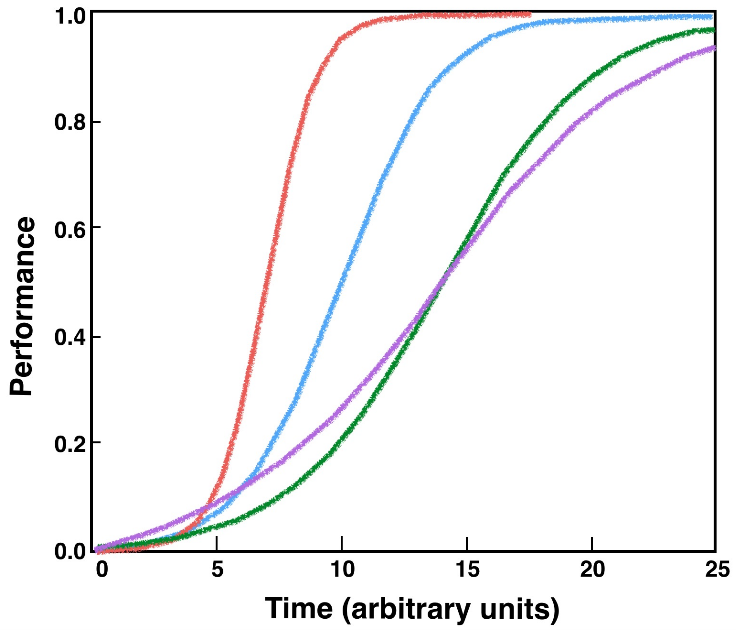

Fig 5.4: Showing four different performance curves.

The composition of the performance curve: Clarifying the 'Irishman's axe' issue

What do we mean by 'the wheelbarrow has evolved'? Clearly, our wheelbarrow is inanimate and our original Version1.0 is gathering dust in the back of the shed. What has evolved? This is why we need to be careful with our definitions – of course it is our "*idea of a system that fulfills a particular purpose as well as we can make it*" that has evolved– the idea of the functions of the wheelbarrow remain (fairly) constant while we make new versions. This may seem a trivial point, but one which, as we will see later, becomes quite complex. The issue of 'same' is a popular one in undergraduate philosophy classes and is exemplified by the 'Irishman' who claims that he has worked all his life with one wonderful axe – although, of course it has had 20 new handles and five new heads in that time!

Although *Fig. 5.3* may be essentially correct for, say, the progress of an individual's performance of some set task, many systems of interest will look more like *Fig. 5.5*, where the overall 'envelope' of performance of the stated purpose comprises a number (in this case five) separate performance curves. Each of these curves could be a 'version' of the wheelbarrow, where improvement in performance of each 'version' occurs within the constraints of keeping the main attributes of the basic system constant.

In *Fig 5.5* we can see that the end of one curve isn't continuous with the beginning of the next curve. We will discus this in greater detail in Chapter xxx. In summary, there are often 'transition costs' in moving to a significantly changed version. The 'bedding in' process of the new version may reduce its initial performance, with the prospect of its later performance being greater than the version that it replaced. As Fig. 5 also shows, there is sometimes an immediate improvement in performance with the new version – for example the transition from the first to second versions of the wheelbarrow and the third to fourth version.

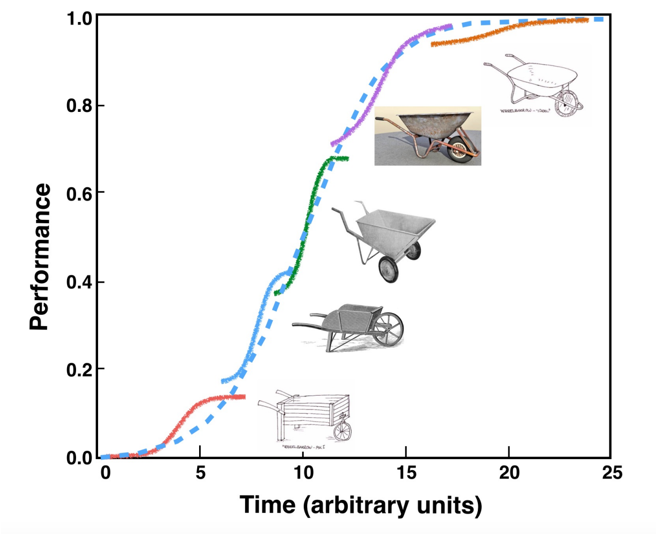

Fig 5.4: The underlying performance curve actually comprises a number of overlapping curves of a series of 'versions' of the 'same' system. The smoothness of the curves suggest that there are continuous improvements in performance within each 'version'.

Some background to the 'logistics curve'.

To *fully* understand the system's rate of learning one must know the basics of *diffusion theory*[7]. Diffusion theory had its origins[8] in observations of the way populations grow when introduced to new territory and how this growth is ultimately limited. The notions of *population* and *territory* have been generalized to include, for example, the number of people becoming ill in an epidemic, market growth of a new product, energy and transport infrastructures, language acquisition, and technological performance. All of these measured quantities, or 'populations' (number of a species, height of a plant, power of an engine) display the same dynamic: a period of slow growth followed by a rapid ('exponential') rise and then a leveling-off towards a maximum population. The result is the familiar 'S'-shaped, or sigmoid curve. Why does this occur and why is it such a common happening? (See Hurst and Zimmerman or Hurst[9].)

We can provide a ready description of diffusion from a systems perspective. First, we assume that the system has a *boundary*- it may be an island where we are observing the effects of an introduced species, a city where people are affected by an epidemic, a corn-growing region where a new corn variety is introduced, a market for cell phones, etc. The boundary may be precise (like an island), or less well-defined, like a city. In line with our systems approach, we may arbitrarily define what is inside and outside the boundary. A boundary, by definition, sets the limits to our observations.

Limits to Growth

Essentially, the growth of the system is a function of the fundamental process of knowledge accumulation: *The amount of knowledge that can be accumulated is related to the amount of knowledge that has already been accumulated*. In the early stages, this will accelerate the accumulation; later, it will impede it. This is what is called an 'exponential' process, similar to population growth and nuclear explosions. There are some provisos to this: The rapid growth does not happen infinitely quickly and the process does not usually expand indefinitely. Let us look at these two provisos in detail:

First, the 'speed of thought' is not infinite. The rate at which we, or a system, learns depends on:

- The speed of communication from element to element in the system;
- The rate at which the data being communicated can be processed (learnt) within the receiving element; and
- The rate at which this processed knowledge can then transformed into procedural knowledge and then acted on.

These same principles apply at both the neurological level – where the communication is internal – and at the external organizational level, where communication relies on a range of media. Internally, the rate-limiting factors include the capability of the initial data sensors (sight, touch etc), the speed of signals along nerve pathways, the extent to which conscious data processing is required, the extent to which the response has been integrated ('data chunking') and ultimately the capability of the 'effectors' (muscle groups) to actually do what the cognitive processing has commanded.

Similarly, externally, data gathering is rate-limited by human and equipment capabilities; communication speed is rate-limited by the speed of speech, writing or other signaling, the postal or internet services, the capacity of the organization to process and 'make sense' out of the information received and finally to organize and put decisions into action.

Signal processing speed is therefore 'intrinsically' limited by these factors, even in an ideal situation. In reality, further reductions in the speed of action are caused by *signal distortion* and *non-optimum signal pathways*. Signal distortion might be due to random or accidental interference (noise), which makes correct interpretation of the signal difficult, or deliberate interference by competing signal sources and transmitters. To all of this should be added the problem of processing errors. Non-optimum (ie longer with more links and nodes) signal pathways may occur because the optimum pathway is intentionally or accidentally blocked, or simply structurally not available. All-in-all, there are many reasons why the speed of action of a system is limited.

Secondly, expansion or diffusion does not go on indefinitely. Many – if not all – systems have limited potential in practice. In our wheelbarrow example, the size, speed and ease of use are limited by the physical size and strength of humans, the limits of mechanical dynamics and materials' properties. Further, the 'price' that we are prepared to pay for a wheelbarrow is important. For example, an all-titanium frame with a lunar-rover wheel might have a superior performance to our best steel, plastic and rubber model, but its cost would probably be prohibitive to the average wheelbarrow user.

'Cost' is essentially a measure of the amount of energy or effort required to acquire the resources to make the product and its availability to any system is always limited. *The wheelbarrow-user must strike a balance between the proportion of the available money/energy/resources spent on buying, maintaining and using his wheelbarrow and the proportion of money/energy/resources available for other purposes.* As the original purpose of using the wheelbarrow was to enable the production of something else, (*eg,* cart bricks to make a shed) the making and use of that something else could be limited by the choice of wheelbarrow. If the wheelbarrow took too much time to make, or the user had to work too long to earn the money to buy the very fancy wheelbarrow, then enough of our time and money might not be available to buy the bricks and make the shed.

So the limits of performance of the system are ultimately set by the energy/resources available to achieve that purpose in the context of other systems and purposes that we are involved with. As we shall see in later chapters, the use of knowledge to improve the efficiency of the use of available resources is a principal driver of innovation and determinant of the way that a system operates.

The Learning or Experience curve

The curves of *Figs 5.2-5.5* depict the change of performance of the system-of-interest over *time*, with a generally-observed increase in performance. This kind of experience/ learning curve has been around for a long time[10] (from at least 1885). These graphs are basically fairly easy to understand – time intervals are frequently used as many of our activities are measured against time – our daily activities, our hourly rate of pay, the number of wheelbarrows produced each week, our annual tax returns, etc. In summary: the longer the period of time, the better the performance.

There is also another – and often more powerful – way to depict performance – that is, the change of performance with *cumulative experience*. This depiction is somewhat more abstract. Rather than using *time* intervals on the horizontal axis, we show the *number* of times that the 'system' has been reproduced. That reproduction could be related to a single kind of activity – say, the speed of sawing of the 4x2s for the legs of our wheelbarrow, or for a collection of activities – such as the production of the *whole* wheelbarrow. Further, rather than depicting, say, the number of wheelbarrows made per week, we depict *all* of wheelbarrows produced *since the very first one* was made – the *cumulative number*.

To add to the abstraction of this presentation of performance we add two further refinements. First, we often graph an indicator of performance that reduces with experience – often the cost-per-unit. A simple example of this is shown in *Fig.5.5a*.

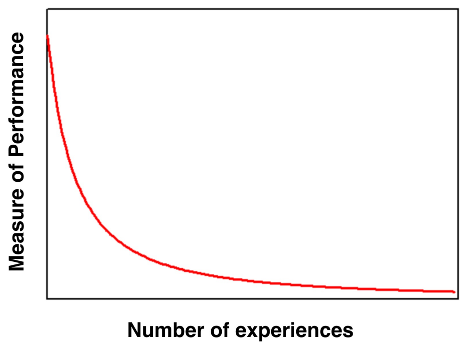

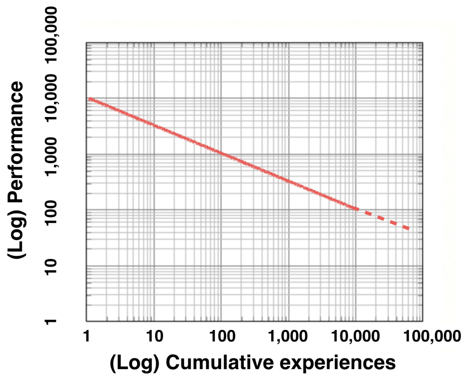

Fig 5.5 a): The experience curve shown on a linear graph and **b)** on a log-log graph.

However, although this kind of graph broadly illustrates that learning is initially rapid and then incremental, it is hard to use quantitatively – or semi-quantitatively. So we add a further refinement – we depict the *logarithm* of the measure of performance *versus* the *logarithm* of the total number of experiences. Briefly, such a graph does not use *linear* – or evenly-spaced – intervals on the axes – it uses 'powers of ten' – the first ten experiences has the same length on the axis as the next 90 (from 10 to 100) and the next 900 (100 – 1,000) and so on. The same logarithmic scale is also used on the vertical (performance) axis. The result of this mathematical processing is to turn the curve into a straight line. This is illustrated in *Fig 5.5b*.

The first benefit of this *log-log* graph is that it is easier to predict future performance from past performance by simply extrapolating the straight line. In *Fig. 5.5b* we can see that the performance (say cost) is likely to fall from $100 after 10,000 units have been produced to about $30 after 100,000 units have been produced. Note that the actual drop from 1,000 to 10,000 units was from $300 to $100 – demonstrating that the actual cost reduction per unit has slowed, as easily seen in *Fig.5.5a*.

An actual example of the power of the *log-log* learning curve is shown in *Fig.5.6*, where the cost of solar (PV) panels (cost per watt) versus the cumulative production of panels (in Megawatts- MW). The time at which a particular price was achieved is also shown. We can see that the price has dropped from about $100/watt in 1975 to $10 in 1990 to less than $1 in 2014. The estimate of 30-40 cents in 2025 is based on the estimated cumulative production by that date. The broad grey band indicates that there is a range of prices and that the costs below the band since 2012 may be anomalous (or that there has been a radical shift in production methods and or factors in about 2006 – see Chapter xxx)

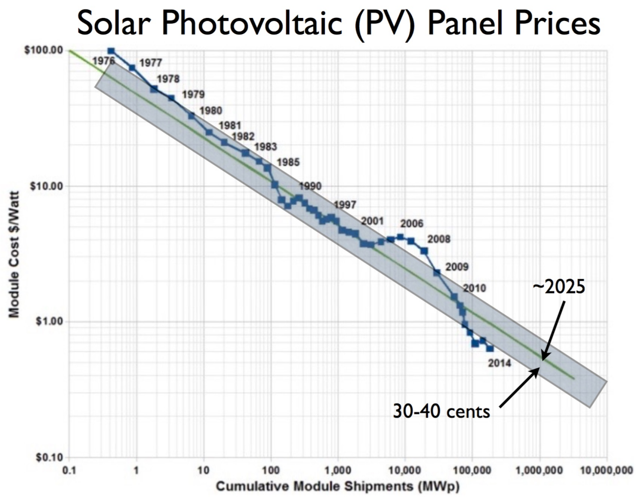

Fig.5.6: The cost of photovoltaic (PV) panel cost vs the cumulative production – also showing the dates at which particular costs and production were achieved. [(Adopted from Wikipedia)]()[11]

The second benefit is that we see that although the rate of improvement in performance declines, it does not actually drop to zero – it just becomes harder to make gains.

A third benefit of this kind of graph is that we can readily compare performances and make estimates of future relative performances. *Fig. 5.7* shows two 'products' – red and green, where the green product starts at almost on-tenth of the cost of the red product ($1,000 *cf* $10,000) and is still is half the cost ($500 *cf* $1,000) after 100 units. However, the red product is on a 'steeper' learning curve and we can estimate that the costs will be about the same after 1,000 units and the red product will be about half the cost of the green product after 10,000 units have been produced ($100 *cf* $200). This kind of analysis can be used to estimate– from early production – the likely costs of a product when it has reached large-scale production.

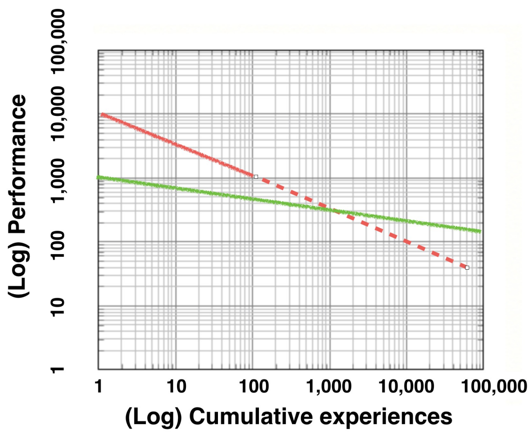

Fig 5.7: A learning hypothetical curve graph showing the different rates of learning ("steepness") of two different products. Whereas the initial cost of the red product is greater, it is cheaper than the green product when produced in volumes of more than 1,000.

Gladwell's '10,000-hour Rule'

Throughout his book *Outliers: The Story of Success*[12], Malcolm Gladwell repeatedly mentions the "10,000-Hour Rule", claiming that the key to achieving world-class expertise (performance) in any skill, is, to a large extent, a matter of practicing the correct way, for a total of around 10,000 hours. He cites examples ranging from The Beatles to Bill Gates to substantiate his claim. In Gladwell's view *talent ie, innate ability,* is quite secondary to *extensive practice*. In perspective, 10,000 hours means about three hours per day for 10 years, or equivalent.

Fig. 5.7 provides a clearer perspective on this proposition. *Red* may be less 'able' than *Green* – gauged by their respective abilities before they do much practice. However, *Red* is more *talented* than *Green* – that is, he learns more with each successive practice than Green does. By the time they have both done 1,000 'experiences', Red is as able as Green and by the time that they have done 10,000 experiences, Red is much more able than Green. In the case of say, playing a piano, hours-of-practice fairly well equates to number-of-experiences. However, in an activity like golf, the number-of-balls-hit is more relevant than the number-of-hours-of-practice, although, of course, they will be broadly correlated. We can see from *Fig 5.7* that if the performance is *experience-related* rather than *time-related* and if the very-talented *Red* does not repeat the practice-experience often enough, the less-talented *Green* will demonstrate superior performance. In summary, although Gladwell is broadly correct, nothing will beat hard-working talent!

There are many other uses of the learning curve, which we will explore in depth in Chapter xxx.

Reasons for the Learning Curve Effect

The primary reason for why experience and learning curve effects apply, of course, is the complex processes of learning involved. As discussed, learning often begins with making large 'finds' and then successively smaller ones. This learning can occur in the product itself, in the production system (process) or in the wider environment. These will be dealt with in detail in later chapters. For the moment we will just list some of these foci of learning. Some of the major areas in which learning occurs include:

- Labour (including management) efficiency, including specialisation;
- Equipment standardization and specialisation;
- Better use of equipment;
- Scale economies;
- Process rationalization;
- Technology-driven learning;
- Changes in the resource mix;
- Product redesign;
- Network-building;
- Shared experience effects.

Negative Learning – Forgetting
Although it is a common experience that 'practice makes progress', improvement in performance is often less than ideal and in some cases it may even be negative.

Frequently, we forget[13] a lot of what we have tried to learn. This occurs in both individuals and in organisations. With individuals, remembering what we have learnt depends on the effectiveness of the information becoming 'wired' by repeated experience. The variability between individuals in the amount of reinforcement required to memorise an item of information (declarative or procedural) is an active area of enquiry with neuropsychologists.

Surprisingly, there is still significant contention amongst researchers as to the mechanism responsible for decline in recall ('forgetting'). The earlier notion of 'trace decay' – where the neuronal connections are similar to an unpaved track that fades with time – are now largely discredited. 'Interference' – where the particular memory has not achieved a threshold of distinctiveness and therefore is subject to being over-ridden by more distinctive 'roads' (neuron paths) – now presents a more coherent theory. (See Brown and Lewandowsky[14] 2010).

In organisations, we often talk of 'corporate memory', which is embodied in the minds of individuals, in the relationships between individuals (both staff and clients) and in the recorded information in the organization. The corporate memory can be negatively affected in a number of ways, including by the loss of individuals, the re-organisation of departments and loss of records.

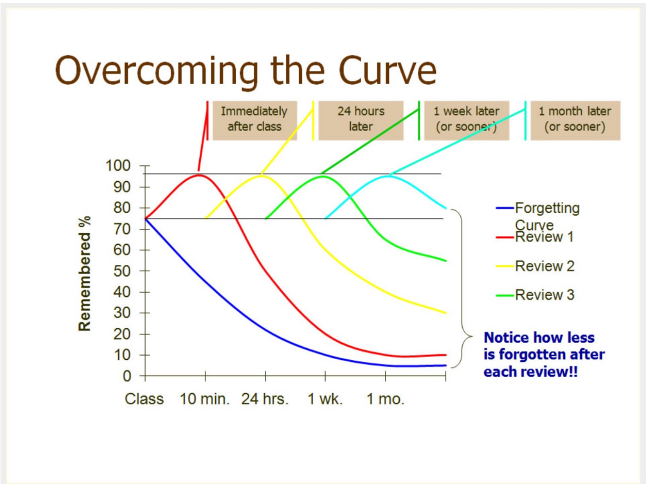

Fig 5.8: The forgetting curve[15] , showing that retention of knowledge is improved by both repetition and review and the timing of those repetitions.

In real-life situations, we know that the attempts at learning are competing with the forces of forgetting – whatever their causes. While we have shown that one of the most important determinants of learning is the number of repetitions of the learning event, the effectiveness of these repetitions is affected by the periodicity of the repetitions – that is, the wiring becomes ineffective if the 're-firing' is not frequent enough. (See *Fig. 8*) This is, perhaps, why a golfer, who plays once per week for twenty years, but does not practice between games, is unlikely to improve their handicap. A golf pro/instructor is more likely to tell the golfer/student to hit 1,000 balls at the range (with the newly instructed technique) before going back to the course, rather than say 'practice for two hours each week'.

While forgetting is most likely to inhibit innovation by reducing the ability of the innovator to apply knowledge to the system being innovated, there can be some positive effects. Some of the information that is forgotten may be dysfunctional and forgetting will reduce its negative effects. In individuals, this can range from the benefits of having a coffee-break to playing golf on the weekend to travelling to exotic places on holidays – all of which disrupt the reinforcement of memory traces which may be dysfunctional. In organisations, the loss of some corporate memory through redeployment of staff may reduce resistance to desirable change.

Another way in which performance may decline is through the selection of an idea, invention or organisation that progressively acquires increasing complexity to 'work'. By work, in this case, we mean the capacity to fulfill the expectations of its 'market', rather than fulfill an initial simple and single objective. We are familiar with 'feature creep' in computer software that requires increasing skills to do even simple tasks. We will see in Chapter xxx that it is the fate of many organisations to become "bureaucratic" with time, adding layers of staff that ultimately impede the achievement of the corporate mission.

A large-scale example of decreasing performance with repetition is the cost of nuclear power. *Fig. 5.9* shows the increasing cost per kW with cumulative production in the US and French nuclear industry. Increased complexity was required to comply with increased safety requirements.

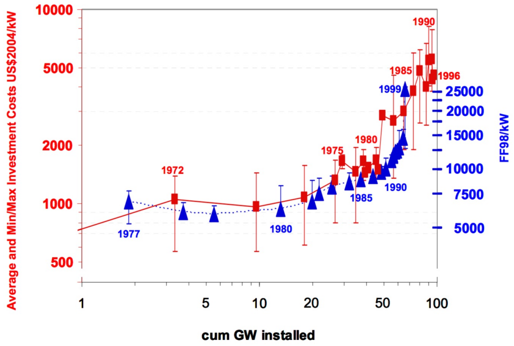

FIGURE 3. REACTOR CONSTRUCTION COSTS PER YEAR OF COMPLETION DATE FOR US AND FRANCE AS A FUNCTION OF CUMULATIVE INSTALLED CAPACITY. NOTES: REACTOR COSTS INCLUDE AVERAGES AND MIN/MAX RANGE. IN BOTH US & FRANCE, NUCLEAR EXHIBITS A PATTERN OF "NEGATIVE" LEARNING DESPITE A RADICALLY DIFFERENT INSTITUTIONAL ENVIRONMENT (CENTRALIZED AND FAVORABLE IN FRANCE, AND FRAGMENTED AND LESS FAVORABLE IN THE US). SOURCE: US DATA FROM KOOMEY AND HULTMAN, 2007; FRANCE DATA FROM GRUBLER, 2009.

Fig. 5.9: The learning curve for construction costs of nuclear reactors in the USA and France, showing the year of completion.

Conclusions

Understanding the processes of acquisition and deployment of knowledge (aka 'learning') is a vast subject. This chapter has aimed to introduce some of the key concepts that are frequently used in the management of innovation and change. These concepts will be expanded in later chapters.

References

[1] http://en.wikipedia.org/wiki/Plato
[2] Nonaka, Ikujiro and Takeuchi, Hirotaka, *The Knowledge-Creating Company –*
[2] Nonaka, Ikujiro and Takeuchi, Hirotaka, *The Knowledge-Creating Company – How Japanese Companies Create the Dynamics of Innovation*, New York, Oxford University Press, 1995.
[3] http://en.wikipedia.org/wiki/Wheelbarrow
[4] http://en.wikipedia.org/wiki/Business_process
[5] http://www.macrobusiness.com.au/2014/03/wrong-lesson-from-adam-smiths-pin-factory/
[6] See Anderson reference (16) below.
[7] https://en.wikipedia.org/wiki/Diffusion_of_innovations
[8] https://en.wikipedia.org/wiki/Trans-cultural_diffusion
[9] Hurst, David K and Zimmerman, Brenda J, *From Life Cycle to Ecocycle, A New perspective on the Growth, Maturity and Renewal of Complex Systems,* Journal of Management Inquiry, **3**, (4), December 1994, pp 339-354.
[10] https://en.wikipedia.org/wiki/Learning_curve
[11] https://en.wikipedia.org/wiki/Swanson%27s_law
[12] Gladwell, Malcolm (2008). *Outliers*. Little, Brown and Company.
[13] https://en.wikipedia.org/wiki/Forgetting_curve.

[14] Brown, Gordon D. A. and Lewandowsky, Stephan, *Forgetting in memory models: Arguments against trace decay and consolidation failure,* in *Forgetting*, Sergio Della Sala ed. Psychology Press 2010, pp 49-75. See also: http://websites.psychology.uwa.edu.au/labs/cogscience/Publications/FINAL%20BrownLewandowskyRev2.pdf.

[15] http://www.mentormegate.com/wordpress/2014/06/10/mentor-me-gate-the-forgetting-curve/