The Developing Role of Software in Actuarial Consultancy

or, why all actuaries must have their own personal supermodel by 2018

Paul Hopkins, England

Preamble

Over recent years there has been a staggering increase in the amount of computing power and related technology to which individuals and organisations have immediate access on a day-to-day basis. The purpose of this paper is to set out a personal view as to a possible impact of the ongoing developments in information technology (IT) upon actuarial analysis and advice.

I am not an expert upon any of the specialisms which underlie or are peripheral to today's technological advances, whether in solid-state physics or chemistry, the mathematics of logic, the IT market, or studies into the social and economic consequences of IT. Such knowledge as I have is that which I have picked up during the course of my work in a life insurance consultancy. However I am indebted to my colleagues Graham Taplin and Mark White who have filled in some of the gaps for me in respect of the innovations likely to become available in the next few years.

My concern is that the profession may underestimate the impact of IT upon actuarial work in the coming years, and should be on its guard against encroachment by other professionals and technologists using IT more effectively. In particular, improvements in knowledge-based and learning systems and the ability to build increasing degrees of co-operation and communication between systems may undermine those elements of our work which give us an intellectual monopoly. As a result actuaries may no long be seen to add value and hence lose the 'right' to carry out certain pieces of work, and it will be necessary to find other ways to add value to the client. Clearly actuaries will have a headstart in terms of existing reputation, statutory responsibility, and client relationships; these are valuable advantages, but in isolation they are not enough to compete effectively in the longer term.

This viewpoint is very much a personal one; to an extent I am playing the Devil's Advocate and I certainly do not claim to have any sort of panacea. If the view seems alarmist, then so be it; I would certainly accept that it might be in the lower quartile or even decile of the distribution of possible outcomes. Nonetheless I believe that the scenario that I outline is plausible, and 'forewarned is forearmed'. I should also emphasise that the paper is not intending to damn any particular actuary, consultancy, or professional body as being short-sighted. Nor is it limited in scope to Consultants with a capital 'C' - all actuaries have at least one client in some shape or form, for example the insurance company they work for, and will have to meet the same challenges.

The time horizon I have chosen is 20 years. Since it is often quite difficult to understand what work we are likely to be carrying out in 2 to 5 years time, there is clearly a large element of speculation in what I write. Perhaps some cynics, having often found themselves sorting through reams of output on the network printer and wondering what happened to the promises of paperless offices, might go as far as to say that my view is guaranteed not to be fulfilled merely by being expressed. However, I thought that it was necessary to look 20 years ahead in order to have the right degree of vision and direction about the possibilities that might emerge. Having to be so speculative also has the advantage of reducing the reliance upon being an 'expert'.

I should add that expecting to retire on the back of the next 10 years' profits is not in itself good reason to ignore the 20-year view. Firstly, and obviously, the trends would emerge and have an increasing impact over the intervening period of time, not just after 19 years 11 months. Secondly, the rate of change anticipated could be too conservative.

The business need

Few, if any, of the problems which consultants are asked to address have a unique, 'correct' solution, although the building blocks of the possible solutions will often be questions with well-defined answers. As a simple example of what I mean, consider an individual choosing between different ways of financing a large purchase such as a house. To address the problem properly he may have to do a number of 'sums' involving discounting payments at various rates of interest. However none of these will in isolation tell him which option to take unless he is taking a very narrow view of the selection criteria. He is most likely to take into account a whole range of factors, some 'touchy-feely', such as his trust in the lender, the restrictions imposed by different methods, etc.

Fundamentally many problems become one of optimising a function such as

$$f(a_1, a_2, \ldots\ldots, a_m, w_1, w_2, \ldots\ldots, w_n)$$

where the a_i are variables and the w_j are weightings; in addition various constraints may arise on the values of the variables and upon the resulting function values. The weightings are used to describe the relative importance of factors in determining the outcome.

To take a purely arithmetic example, the problem might be to find an investment policy that maximises the expected returns without accepting an undue risk. In this case the function being maximised gives the level of profit arising from the values ascribed to the a_i, these being parameters describing the investment policy (not necessarily just in a fixed way but dynamically linked to circumstances). The weightings fall away since we assume that no investment policy is considered intrinsically worse than others. Constraints exist on the a_i, for example as a result of legislative requirements, and a constraint upon the resulting distribution of the function value limits in an absolute way the risk being taken. (If there was an additional, more onerous limit on risk that required investment return to be traded for less risk at any level where the 'price' was right, then weightings would be required to describe the relative values of risk and return.)

In practice it is rare for a problem to be such that all of the factors can be mathematically quantified; for example a variable might be related to the level of bad publicity. Even for those variables that are tractable to mathematical analyses the function may be highly complex, non-linear or discontinuous, and the solution sensitive to the chosen weightings. The latter point, in particular, usually precludes any notion of a 'right' answer.

In my view the tools we use to address business problems fall into three fundamental groups, namely reasoned, empirical, and analytical.

Reasoned arguments rely upon knowledge, experience, judgement, and intuition (the 'hunch'). The main advantage of them is that they can address variables not readily modelled mathematically, such as knowledge of law and human behaviour. However, the corollary is that they are very dependent on the processing environment, ie the individual carrying out the analysis. People are often blind to the prejudices and predispositions that might affect their reasoning. In addition it is not always easy to formalise a 'proof' of the result, especially where a large element of judgement or intuition is called upon.

Empirical tools are those which involve any process of trial and error, applied not to the actual problem being addressed but to models of it. This group includes techniques such as Monte Carlo simulation and the many optimisation routines that rely upon successively improving guesswork (of which Newton-Raphson iteration is probably the simplest and best known example). These tools can only be used where a problem is tractable to mathematical formulae. From a purist's viewpoint, they are deficient in that they may not establish a proof of the result or of its uniqueness, and can also be seen as the application of brute force (although individual techniques will sometimes rely upon the elegance of numerical analysis theorems in order to take shortcuts and avoid unnecessary work). However for many classes of problem they are the only viable approach and their power and range of application are being greatly enhanced in direct response to the increasing availability of computing power. The power has two effects; it enhances the ability of the techniques to solve complication optimisations in a useful period of time, and increases the chances of finding relationships within factors previously immune to numerical modelling.

Analytical tools establish a direct proof and often result in a 'plug-in' solution to which it is possible to provide a series of input values and 'instantly' obtain the result. The class of problem which can be addressed is restricted in that appropriate algebraic tools must exist to arrive at a solution, although over time the boundaries should be progressively pushed back. In fact the limits are often more restricted than appears at first glance. For example, the traditional way of deriving the relationship between the length of a pendulum and its period relies upon the approximation $x=\sin(x)$ for small x. Thus the resulting formula is not a solution to the original problem at all and begs a rather philosophical question - 'how reliable is the relationship between the precise solution to an approximate question and an approximate solution to the precise question?'.

Most tools will fit neatly into one of the three categories. However it is interesting to note that at least one technique of which I am aware, namely neural networks (does any reader use them?), is basically empirical but shares some elements of the reasoned and analytical approaches.

As knowledge progresses so individual classes of problem may become susceptible to tools in different categories. This has already been touched upon in the context of analytical techniques, a recent example being the recent analytical proof of Fermat's Last Theorem by Andrew Wiles. I author would suggest that even prior to this proof nearly all mathematicians **believed** the Theorem. This was not just a reflection of the esteem in which Fermat is held; empirical techniques, combined with some analysis, showed that the Theorem was true for all integers larger than 2 up to a very high limit. No **proof** existed, yet had the veracity of the theorem been a factor in a business decision then I suspect that many would have been willing to take a calculated risk that the Theorem was true on the basis of such an empirical approach.

For me this is a crucial point. With the explosion in availability and mobility of computing power, I foresee empirical techniques making large inroads into the area of reasoned argument, in particular those elements that require judgement. In many cases clients may be more impressed by a model which is able to demonstrate an approach of systematically evaluating a large range of possible outcomes than a 'gut feel'. Where such a model is coupled to a suitably codified database of relevant knowledge and other 'learning' software then it will be able to replicate many of the functions of reasoned argument.

However, in order to believe that this is a possible scenario it is necessary to understand how much computing power might improve over the coming years.

20 years to date

Before looking forward it is salutary to look briefly at the advances that have occurred over the last 20 years and where we have reached now. In 1978 I was still at school and many years away from his working career. I therefore cannot give a firsthand account of the working environment at that time, although even over the last 12 years there have been some fundamental changes.

If my memory serves correctly, in 1978 the pocket calculator had not long ceased to be a novelty. Personal computers were still very much in the province of enthusiasts able (and willing!) to write in machine code, although the next couple of years would see the launch of a number of personal computers for the domestic and business market. Presumably Commodore already had the PET and IBM its PC on their respective drawing boards; in England Clive Sinclair had not yet launched his ZX80 which, although affordable and programmable in a rudimentary BASIC language, was really not much more than a toy. One William Gates was still nurturing his embryonic ambitions in a lockup garage. In short, the mainframe ruled supreme.

Communication has also been transformed so that it is now rightfully regarded as an arm of IT. However in 1978 the closest a telephone came to being mobile was to be mounted in a car; the cellular concept had not been invented, or at least was not a practical proposition. Our idea of high fidelity music reproduction was to run a needle around a plastic disk in an undulating valley full of corrupting dust and other blemishes. (It should be added that for some people this is still as good as hi-fi gets, on the basis that digitising the signal irredeemably corrupts it.) Video recorders were large and unwieldy with poor colour reproduction; for personal movies silent 8mm film was the consumer standard, at least in the UK.

I set out the above not as a reason for self-congratulation but in order to emphasise the possible changes at current rates of innovation. We have now probably reached a point where in the case of most office and domestic functions, at least, people have access to more computing power than they need (although this would not be true for the more advanced scientific and financial analyses). For the first time, software in many areas of application is attempting to catch up with the power of the hardware rather than vice versa.

It is interesting to ask why computers are suddenly so necessary to business, and to actuarial work in particular, when civilisation survived for so long without them, since this casts light on why it is important not to be left behind. I would suggest three principal reasons:

- Using computers in one part of a business means that other parts need to be able to keep in step. For example, computer administration enables companies to offer life insurance products of a complexity that would not be practical using manual methods. As a corollary, computers are required for financial reporting of these products in order to manage the calculations arising from the greater complexity.

- Where computers bring a competitive edge, then it only takes one company to exploit that in order for the others to be forced to follow. Such an edge may arise from sources as diverse as more efficient administration or making better use of capital by a greater understanding of the underlying financial processes.

- There is almost certainly a genuine advantage in that the more thorough analysis of a financial institution results in a lower risk of disaster, **other things being equal**.

20 years' time

I have given below a few ideas for the technologies that might available on a commodity basis and some of the resulting innovations in daily life; some of these are not directly relevant to the point of the paper but are interesting corollaries. Some of the ideas are already almost with us, while others are based upon existing extrapolations of current trends or upon my own flights of fancy. Of course a basic assumption behind all of them is that progress is not impeded by some sort of disaster or by shortages of the required energy and resources.

Computing power

At current rates of advance the real cost of computing will have fallen so that we all have ready access to power 10 thousand to 1 million times greater than now. Clearly, a number of heroic assumptions lie behind such a statement. At one extreme, even if current technology eventually hits a development limit and no replacement technology is found, that limit is not yet very close. Systems currently under test in laboratories, coupled with parallel processing techniques, seem to ensure a likely improvement factor of a thousand or so over the coming years. At the other extreme, scientists are talking in serious terms of chips costing the equivalent of 1 penny, and festooning literally every aspect of our life from birth to grave.

Either way, combining the widespread distribution of computing power with enhanced communication will make more efficient use of the available power, since it will be possible to distribute processes according to the availability and suitability of processors.

Communication

The third generation of cellular telephones, already planned for introduction early in the next century, will have introduced mobile internet access and live video. By 2018 the cellular telephone will have mutated into a personal transactor/note-pad, incorporating many of the functions personal organiser and computer. However it will be extremely small and light, having only a bare minimum of functions locally, and communicating with a 'basecamp' computer at home or in the office to provide additional processing power and data access.

Available radio bandwidth will be full, with information theory demonstrating that not much more can be squeezed in without going up to unmanageably high frequencies. Some increase in capacity will be realised by technical improvements such as using smaller cells. There will also be an increasing emphasis on using fixed links wherever possible

Virtual reality (VR) will be highly realistic, but will not have totally usurped the direct human contact that most people will still consider to be highly desirable. For much the same reason the office will still exist, but will operate much more flexibly. On arrival at the office employees will go to a convenient workstation, perhaps to be near others with whom they are currently working on a project. Since there the principal method of permanent storage of information will be electronic, there will not be any requirement to remain close to their own files. There will also be no need to provide enough space for all employees since at any one moment in time a proportion will be tele-working.

Voice control of machines will be normal, although direct brain links will still elude us except for cruder functions such as the control of artificial limbs - we will still not really understand very much about the thought process. However we will have very reliable security devices able to detect characteristic brain waves that are not capable of being fraudulently imitated. They will have the additional advantage for highly sensitive areas of being to detect individuals who, for whatever reason, have become mentally disturbed.

Data

Storage will be photosensitive substances with a highly regular crystalline structure storing data in three-dimensions. All distinction between permanent storage and RAM-type memory will have ceased, since the devices will need no power to retain data, will offer fully random access, and will be able to be rewritten indefinitely. One cube of dimension 10 mm will hold 1000 gigabyte of data.

Multimedia

Data-handling technologies will be totally merged, so that audio and video are stored and manipulated in the same format as any computer data. The same cube above can hold your entire music collection or 50 full length movies.

Software and applications

Artificial intelligence and other learning software will be widespread, but it will be quasi-intelligence. As stated above we will still have an inadequate understanding of the mind. Although computers will appear to imitate much of the functionality of intelligence, it will not be truly intelligent and will never have a truly original or creative thought.

Many important banks of knowledge will have been substantially codified and readily available for direct interrogation by other systems; key data banks might include (non-personal) medical data and legislation.

Despite the lack of true intelligence, expert systems will be able to perform the many tasks carried out by professionals and other specialists that are fundamentally mechanical. As an example of this, much legal investigation will be carried out directly by computers able to assess the merits of a case directly from the available evidence, legislation and precedent. For many court cases there will no longer be the need for human adversaries and the arguments will be considered by one computer able to weigh the evidence for both sides more thoroughly than teams of lawyers dedicated to a single point of view. Factors requiring weighing might include public opinion and the prevailing political climate, as well as straightforward moral and legal prohibitions. The computer would give a summary to a judge, together with an optimal recommendation for the final decision. (Much to the relief of the judiciary, the general public will still be suspicious enough of computers to want the final decision to be taken by a man - even a semi-senile one.)

Actuarial work

So what will this all mean for the actuary. A few initial pointers:

- Firstly a simple one - the commutation function, so long acting as a convenient shorthand for valuing cashflows, will go the way of the slide-rule. The actuary will have ready access to enough computing power to make such a short cut more trouble than it is worth.

- Much bread and butter modelling, for example as typically implemented by students, will no longer underpin our work, since a lot of the current requirement to involve actuarial experience in modelling is based fundamentally in a lack of adequate computing power. For example, it would in principle be possible to model a life company by accelerating a clone of it's administration system through the daily processing cycles over the required projection period, under the specified assumptions with respect to mortality, economic conditions, etc. It is only the requirement for this to occur in a reasonable period of time that makes actuarial involvement necessary, by understanding how to introduce simplifying approximations and the effects of them. To the extent that computing power makes this understanding redundant, the operation becomes one of basic competent programming.

- Direct links with, or incorporation into, administration systems will enable models to keep on top of product, policy and asset changes.

- A single logical framework will support the full range of studies such as product pricing, capital planning, statutory solvency testing, valuation, asset liability modelling, and sales monitoring on a consistent basis.

- Asset/liability modelling, usually stochastic but as a minimum involving multiple scenarios, will be normal procedure. It is possibly useful to clarify what is meant by such modelling. Any valuation is an asset/liability model in a rudimentary sense since an assumption is made about investment returns earned on assets. What I am driving at is a holistic approach in which all factors which affect a company are incorporated within the model including the interactions, or supposed interactions, between them. Where possible this would include modelling of the predispositions of human elements.

- Expert systems will be able to learn from completed or current projects and use them as examples or templates for future work.

- Today's problems typically involve answering questions on the consequences of decisions, or supposed decisions, already taken by humans. The increased power of empirical models will enable systems to **provide** suggested optimal decisions in a practical period of time.

So let us consider a typical meeting for an actuary, assuming that the profession has succeeded in using technology to retain its position as the first port of call for actuarial advice and has implemented appropriate models.

An A1 client has asked you in to discuss changing investment manager; such meetings are still typically conducted face to face. After discussion you discover that they are looking at three new managers, A, B and C, as well as making fundamental changes to investment options offered to policyholders. You are asked to consider what effect each might have on the office finances, assuming new business according to the latest business plan, and using the 7th generation Wilkie model seeded with the client's own parameters.

Speaking to your telephone/personal transactor, you request that the model be run stochastically for all three managers. This request, together with the changes, is fed back to the base system which immediately starts pulling in the required elements - the economic model, the profiles on the three managers, the office model which it will adjust for the new policy options. The manager profiles are regularly updated by your systems; these look at performance, investment strategy etc. to derive parameter sets comprehensive describing each manager's behaviour. It also requests an update on the policy and asset database from the client's own computer system; since this is an unusual request the client's computer telephones the general manager, who is in the meeting with you, for authorisation.

Examining the product changes, your computer 'decides' it needs clarification. Knowing that with this client you do not wish to be interrupted unless really necessary, it opts initially to interrogate a member of staff who might be able to assist. Your message is replayed to the staff who, despite not being an actuary, provides the required information; so far this has been the single longest stage in the process. The system then interrogates the insurance regulations to verify the impact upon valuation of the new product feature. Now having all the information to hand, the system starts work on the calculations. So far three minutes have elapsed; each of the 1000-simulation, seriatim policy projections will take a further 90 seconds.

In the meantime, you chat amiably to the client. You discuss the forthcoming wedding of your daughter and invite all three representatives of the client in the meeting to the wedding. Your daughter won't like it but, what the heck, clients like this one don't grow on trees and this is one piece of account management that damn computer can't initiate. The investment manager then mentions that he really didn't like client C very much; he seemed over-confident and a bit out on a limb in some of his views. Just at that point the results come back – notwithstanding the investment manager's views the order of preference is C>B>A. Somehow this seems to stiffen the investment manager's resolve, and after discussion it is agreed that the choice will be B. At the same time your system advises a small change to the proposed product modifications that will reduce valuation strain, but offer most of the benefits; after discussion the client decides to stick with the original design.

Having successfully concluded the meeting to the satisfaction of all, you stand up to leave. You then have a horrible thought - very recently you have learned that manager B runs neural networks that emulate the processes of advisers such as yourself in order to optimise his presentations. You quickly confirm that B was told of your involvement, although in any case he would have been aware that you were a regular adviser. Somewhat embarrassed by your oversight you decide to 'come clean', and advise that the result was close enough that it might be worth investigating if B could have swung the decision by such means. Getting on the telephone to a senior student actuary, you find that the work which last week you instigated to emulate manager B emulating you has made little progress. However he feels that a couple of hours work will be enough to obtain a rough initial result.

> You inform him that if it takes a minute longer than this his next attendance at an IACA conference will be via VR; he starts work promptly.
>
> You are due back in the office for a lunch-time internal meeting, but would rather stay until the results arrive. It is agreed that you can use the client's own VR facilities to take part in your meeting. You don't like VR – there is always a disconcerting disconnected feeling; however there is little choice on this occasion. You are taken to one of the client's own one-man VR cubicles, a miserable black hole that with a single command can be transformed to almost anywhere on the planet or beyond. Today it transports you to a meeting room at your office, although in fact the room is also a fiction of your office's VR system - you note that the other six attendees failed to book one of the prized 'client' meeting rooms with real wood tables, internal décor and windows. Since virtual food is virtually useless, the system blanks out your office sandwiches and wraps the rest of the hologram around food placed in front of you.

Although elements of the above cameo not directly relevant to the point, such a scenario does not seem an unreasonable proposition to me. Note that I place human contact as still being an essential precursor to good business. However I do not know if it is the **actuary** that has to be the human at the centre of such a scenario. It is interesting to ask what added value is brought to the above proceedings by the consultant in the capacity of an actuary rather of any intelligent, business-wise creative consultant. Of course, to an extent the answer is that I have deliberately left out any suggestion that the actuary added much actuarial value, partly to make a point & partly because I can only guess how this might be achieved in the future. Nonetheless I believe that where much of the expert knowledge can be packaged into systems it is at least plausible that other specialists, who may have their own problems with competition from expert systems and need to diversify, will be able to compete effectively in this way.

For example, it already seems to me that many investment banks and the like perform much more sophisticated asset analysis and modelling than do most actuaries. We perhaps have more of a monopoly of knowledge in respect of the liability side. However, giving the increasing availability of suitably powerful tools, it is a moot point as to the relative difficulties of a bank understanding the liabilities and of actuaries understanding the assets. What is to stop an investment manager offering an insurance company a holistic analysis by adding the liability side of the equation to his existing work?

We also cannot forget the accountants. Most competent members of that profession understand the concepts of discounting and projected cashflows. Given access to increasingly powerful empirical models, even ones that have to be seeded by a tame actuary's knowledge, it is not clear to me that they could not learn to provide most or all of our services as part of a broader package.

For the record, I accept that that a number of factors may work counter to such trends, such as the value to the client of independence and reduced conflicts by spreading work around.

Conclusions

So what conclusions are to be drawn from the above, and how will it affect us all? I am not about to prescribe policies for matters such as systems development and for recruitment, since it would be highly presumptuous for me to do so. I would suggest however some pointers:

- More serious research is required into the potential of systems and analytical techniques; we need not be ashamed to take advantage of academic research in these areas.

- We should investigate more thoroughly what colleagues/colleagues in other disciplines are doing. Part of this involves inviting them to professional meetings, but they are perhaps likely to learn more than they teach at such events. We should therefore be more willing to invest time in proactive research.

- Keep alert for trends in financial analysis and reporting and be prepared to go out on the limb a little when extrapolating where they might lead.

In addition it is quite possible encroachment by other disciplines may in some cases be premature with an incomplete or incompetent understanding of the issues, leading to inept advice. Actuaries should be alert to ensure that this does not also damage the profession.

As indicated, my own area of specialism is life insurance, and this paper is based on observation of trends in that area. However I think that the fundamental point translates well to most areas of actuarial specialism.

I am not necessarily forecasting waves of redundancy across the profession; a succinct summary might be 'be prepared to adapt or die'. To date IT innovations have made redundant many of the clerical rôles that traditionally underpinned businesses. However new jobs have been created, partly from psychological necessity and partly as an economic inevitability (although there may be some accompanying change in the 'natural' level of unemployment). The functional successors of yesterdays' clerks may be spreadsheets and PCs, but ranks of their spiritual successors are to be found in the halls of tele-sales operations. It is my contention that such changes will inevitably work their way up organisations as analytical tools emulate more specialised staff and skills. I do not think that we can say where a ceiling might be or even if it exists.

And the 'paperless office'?

I actually think it will happen in the next 20 years, or at least nearly. Paper has two key uses, firstly as a permanent repository of information and secondly as a 'scratchpad' to assist the user with thinking. Although it will probably be possible to develop very thin, lightweight and pen-sensitive display screens looking, feeling and acting much like paper, I think that the value of being able to rifle through sheaves of paper will remain irresistible. I believe that we develop artificial materials that can be printed or written on and then totally erased at will. When working on a document, the computer will print the latest draft on 'paper' that can be annotated and amended at will by hand. A scanner will then read in and implement your changes in the electronic copy, and the 'paper' then erased for reuse.

In the current climate, it is perhaps fitting to end on such a 'green' note.