

Thought Leadership

Shiny new toys. Does Wealth Management Need A.I.?



Greg B. Davies PhD

Founder, Centapse

GregBDavies@centapse.com

How to get ahead with A.I.

There are few things more enticing to an ambitious banking executive than a new technology that is poorly understood by the industry they're in. It offers the opportunity to appear innovative, cutting-edge, and entrepreneurial through nothing more than frequent jargon dropping. This is particularly true when this new technology is surrounded by popular media exposure and closely associated with the buzzy tech startup world. Current exemplars are *blockchain*, *Big Data*, *agile*, and, of course, *A.I.* and *machine learning*.

Not that A.I. is exactly new – Alan Turing proposed his eponymous test for A.I. in 1950, and it was nearly 20 years ago that IBM's Deep Blue finally beat chess grandmaster Gary Kasparov.

But advances in machine learning, and the rise of the data and tech unicorns has raised the tantalising promise of harnessing it commercially; something the burgeoning fintech bubble is increasingly forcing financial institutions to pay attention to... and that ambitious executives are keen to be associated with.

The problem with this enthusiasm is that it leads to *Shiny New Toy* syndrome. The recipe for career progression: find a hot new technology from outside your firm; persuade budget holders that this is the future and you need to invest in a pilot project or your firm will be left behind; spend innovation budget bringing outside technology in to solve, as yet unspecified,

problem; look for problem to solve; ignore existing internal solutions to this problem (they don't have your name attached to them); build expensive pilot solution to said non-problem; await promotion; rinse and repeat.

Picking the right problems

I exaggerate somewhat, but it seems almost self-evident that a sensible approach to innovation is to carefully define the problem that you need to solve and then pick the technology best suited to solve it. This may in some cases be A.I., but often the best (and usually more cost-effective) solution, is something more mundane, quite possibly some capability or design you already have in-house, lying half-developed for lack of senior support and, hence, budget. In wealth management there do seem to be a lot of shiny new A.I. hammers looking for nails.

Accepting that A.I. does offer novel solutions that traditional approaches won't be able to provide, where should you seek to apply it? Which opportunities have the potential to sufficiently enhance your proposition to warrant pouring your organisational resources into the A.I. drain? The important distinction is determining the *nature of the problem*.

A first observation is that many problems in wealth management are simply not complex enough, or do not require enough data processing to warrant the sledgehammer sophistication and expense of A.I. solutions.

Unfortunately, the fact that these problems are easier means that they are often the first choice for proof-of-concept problems to allow the firm to test out the shiny new toy they have just committed to.

Some of these are complicated: they have many moving parts which are difficult for humans to mentally process. This makes them seem ripe for A.I. tools. However, complicated is not the same as complex. If these moving parts can all be effectively handled in a big well-designed decision tree to arrive at the right answer, then there is no need for A.I. Just build the decision tree, which can be entirely deterministic. Indeed, A.I. might not only be unnecessary, but may actually deliver an inferior solution.

Suitability is the wrong problem

The paradigm example here is *suitability* – determining the best investment solution for each client based on his or her individual circumstances.

There are many details of each client that affect the solution, and this is an area where humans struggle to comprehend how these details interact, let alone integrate them all flawlessly: getting it wrong can have enormous regulatory and reputational costs.

But it is not a problem that needs A.I. Suitability merely requires mapping a (relatively small) set of information defining each client's circumstances and preferences, to a (relatively small) set of different investment solutions.

The suitability decision tree that connects the two, if well designed, is certainly complicated, but nonetheless just a mapping algorithm that can be easily built, run, and monitored. For each unique set of client characteristics there *should* be a clear, auditable, deterministic process that points to the space of suitable solutions, and

ideally, picks one. There is no need to over-engineer a solution to this problem, and indeed A.I. systems that pick solutions probabilistically, and which potentially evolve over time to provide unpredictably different solutions for the same client characteristics, would be a compliance nightmare, and potentially dangerous for the client.

So is Risk Profiling

Another example where there is frequently a misplaced desire to deploy A.I. is in risk profiling. This is over-engineering a (dangerous) solution to a problem that really doesn't exist. Risk Tolerance, correctly understood and measured, is a simple, stable psychometric trait⁴, and best assessed with simple psychometric questions. These could be presented in slick, tech-enabled formats – but this is UX, not A.I.

Attempts to assess Risk Tolerance with real time processing of investors' behaviour, social media footprint, or video monitoring of their facial expressions are misguided. Such 'revealed preferences' are extremely unstable and reflect all sorts of short-term behavioural biases and influences of context that should not be used as a foundation for the long-term Risk Profile of an investor's portfolio.

Opportunities for A.I.

Behavioural decision support

However, both these examples bring us to areas where A.I. *could* be useful. For example, the same revealed 'preferences' that are often wrongly confused with Risk Tolerance can be very useful in other ways: in particular they give us important signals of when individual clients might be *anxious* along the investment journey, or where they are in danger of making decisions that deviate from good practice.

▪ ⁴ Davies & Brooks (2014). Risk Tolerance: essential, behavioural, and misunderstood. Journal of Risk Management in Financial Institutions.

This shouldn't be used to change the Risk Profile, but *should* be used to help guide the client to better decisions in line with their stable Risk Profile. This is definitely a problem that could benefit from extensive data processing.

A.I. approaches could be deployed to look for the costly behavioural patterns in each

individual investor, and connect with behaviourally designed systems of targeted communication, news feeds, and nudges to help them control these detrimental impulses; or integrate with CRM systems, sending real

“Decision prosthetics can help improve human decision-making”

time signals to advisors to help them navigate exactly what to say, to whom, and when.

A.I. systems thus should not be used to determine suitability, but could be used to dynamically guide advisor and client behaviour *towards* the suitable solution. This problem, unlike risk profiling, has characteristics that A.I. is suited to: rapid processing of large quantities of disparate data; complex pattern recognition; and probabilistic responses, offering guidance where there is not necessarily a right answer, but a range of more or less likely responses (much as IBM's Watson approached the questions in *Jeopardy*).

Dynamic suitability

A further opportunity is to use A.I. to move towards *dynamic suitability* – that is, constantly updating suitability in response to changing client circumstances and preferences. This is not about changing the suitability framework used to determine the right answer for each client, but instead using A.I. to constantly update the data used as *inputs* to the suitability framework. If the client's balance sheet, circumstances, and goals and preferences are

continually in flux, then so should be the suitable solution.

Applied well, such technology could increasingly blur the distinction between a) engagement, b) profiling, and c) suitability. Processing real-time data will allow A.I. systems to continually refine and update the client's behavioural profile, use this to prompt appropriate engagement and interaction, which in turn provides profiling information for the suitability inputs, and permits constantly updated solutions that serve the client's constantly changing needs.

There are, of course, many related use-cases that could draw on the same systems and technology: *just-in-time education* providing targeted information and education nuggets to clients when they are most receptive, enhancing understanding; client prospecting using data processes to identify life-events; finding appropriate cross-sell opportunities; emotional recognition using biometric data; and finding patterns of bank-client interaction that the organisation cannot itself discern because of the omnipresent barrier of organisational siloes.

However, it is very likely that these solutions would need be developed within an organisation, building the intelligence to suit the specifics of the proposition, client base, sales processes, and culture. This may require external technology, but a simple 'lift and drop' of technology from external vendors is unlikely to be successful.

Centaur design

Note that in all of these examples machines are not being used to make the decision, but as *decision prosthetics* to help improve human decision-making in pursuit of suitable solutions that are driven by much more traditional, stable models. This is another general feature of problems in wealth management that are most suited to A.I. solutions: they are more about decision-support than automation.

After Gary Kasparov was defeated by Deep Blue in 1997, he did something very novel – he started competing *with* machines, rather than against them, creating ‘centaurs’ of man and machine. And, despite the dramatic advances in A.I. in the last 20 years, the best chess player in the world today is not a machine, it is a man-machine centaur.

This is because humans and machines think very differently: there are things that machines are good at (data processing; pattern recognition; consistency; low error rates), and there are things humans are good at (empathy; coping with unstructured problems; creativity; generating insights from association across completely different problems; coping with dynamic environments and multiple objectives).

Because these skill sets are different there is value to combining the two – our ideal should be *centaur design*: applying A.I. to the right parts of the right problems, using them as decision prosthetics to help with what humans are bad at, but recognising that ultimately humans are required to make the decisions and navigate the deep uncertainty of a changing environment, and their own unstable preferences. Only when systems are designed to enable both parts of the centaur to work together will A.I. change financial decision making for the better.

About the author: *Greg B. Davies is a specialist in applied decision science and behavioural finance, turning academic insight into practical applications. In April 2016 he founded Centapse, a firm dedicated to applying sophisticated behavioural insight to design, develop, and deploy solutions across industry to help people (and organisations) make better decisions.*

Greg started, and for a decade built and led, the banking world’s first behavioural finance team as Head of Behavioural-Quant Finance at Barclays. He was the architect of Barclays’ behavioural profiling tools and holistic Wealth Philosophy, delivering solutions tailored to both financial and emotional investment needs; and he designed Barclays’ innovative behavioural approach to impact investing and philanthropy.

He holds a PhD in Behavioural Decision Theory from University of Cambridge; is an Associate Fellow at Oxford’s Saïd Business School; a lecturer at Imperial College London; and author of ‘Behavioral Investment Management’. Greg is also the creator of Open Outcry, a ‘reality opera’ premiered in London in 2012, creating live performance from a functioning trading floor.

There are parts of these systems, like suitability frameworks, where A.I. is not the right tool to provide the right answer and we need something more intentionally designed, engineered, and stable. However, A.I. might help to guide behaviour towards the suitable answer, or ensure the inputs to the stable system are much more current and error free, or even serve as a natural language (and inexpensive) assistant to help human users navigate the process and user interface.

But fundamental to using A.I. well is to first define the problem, and then look for an appropriate technology to solve it, rather than buying expensive shiny new toys and throwing them hopefully at every problem.

And those ambitious executives looking to associate themselves with a rising technology? They would do well to consider this serenity prayer of A.I.:

Grant me:

The serenity to accept what computers do better than people,

The courage to let people do what they do better than computers, and...

The wisdom to know the difference.

From Andreas Weigend