Teamspirit

People, money and Artificial Intelligence

How humans and robots can have a richer,
more real and rewarding relationship

50% of chief executives of financial services companies see AI as the most important technology for their organisation.*

65% of senior managers expect positive changes from its use.**

When it comes to the future of financial services, our money is on Artificial Intelligence (AI). But AI isn't just something to look forward to. It's already here, crunching data, retrieving documents, spotting patterns and making predictions to help with fraud detection, credit risk management for lending, and in automating previously tedious and time-consuming tasks. For example, JPMorgan Chase has been successfully using Robotic Process Automation (RPA) for a while now, to comply with Know Your Customer regulations.

In banking, AI powers the smart chatbots that provide customer support, while intelligent mobile apps help to manage our personal finances, tracking income and spending and giving budgeting tips, guiding us through the investment process and reminding us to pay our bills.

In the meantime, appreciation of the ethical considerations around the use of AI has been growing, culminating in the development of a set of 'Trustworthy AI' principles by the EU Commission, with similar principles adopted by the OECD's 36 member countries, along with Argentina, Brazil, Colombia, Costa Rica, Peru and Romania. Ethical issues are also being explored by the UK's newly created Centre for Data Ethics and Innovation (CDEI).

The seven central principles of Trustworthy AI comprise:

- Championing human agency and oversight
- Robustness and safety
- Privacy and data governance
- Transparency
- Non-discrimination
- Societal and environmental well-being
- Accountability

AI can write pop songs and play chess, but it can't yet provide a complete planning and advice service that's anything close to that provided by a human financial adviser. But we are entering a world of Open Banking, in which multiple providers can be aggregated into a single command-and-control interface and smart, connected devices that can talk to each other as well as to us. At the same time, new research into how our brains work at a structural level, is driving a whole new field of 'Deep Mind' innovation, to enable AI to learn as it grows. This means that in the future, we are likely to be able to manage our finances via a voice-activated, digital assistant, with not only new capabilities but more human qualities than we see today.

She (and she will invariably be a woman) won't just be there to carry out transactions and tasks, guide us through financial decisions, manage and monitor our spending and measure our progress towards our goals. Rather, she'll be our financial 'self', with whatever personality we ascribe to her and a calm, capable, reassuring presence in our lives.
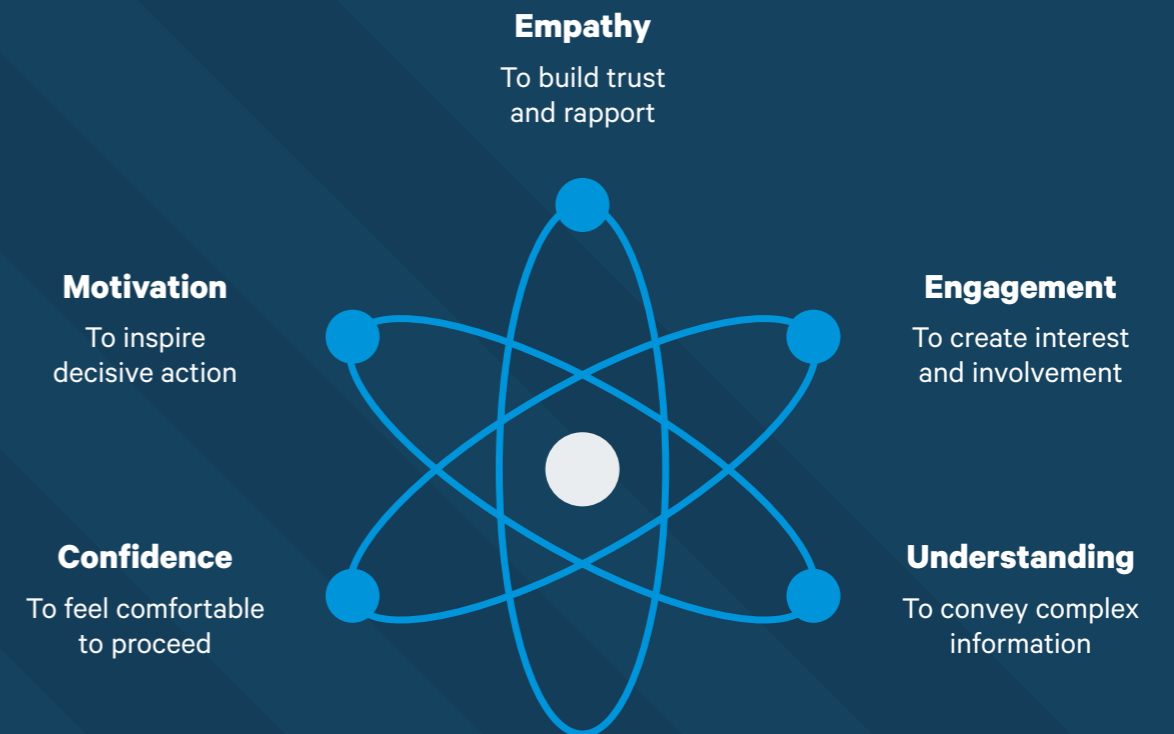
So where will that leave us, not just as financial consumers, but as individuals? And how will we feel and behave when one of our most trusted counsellors and confidantes is a machine?

We can find some of the answers to this, by looking to Behavioural Neuroscience and how our brain chemistry influences our emotions and behaviours. In fact, the early Artificial Neural Networks, the precursor to today's AI, were based on how our brains are wired and how we process information.

Here at Teamspirit, we have developed a five-part framework, to map out the key stages of an interaction and explore the role of Behavioural Neuroscience at each stage. This works equally for a single conversation or piece of communication, or a longer, phased, decision-making process. Using this framework, we looked at how effective communicators are more trustworthy and engaging, make complex information easier to understand, inspire confidence and motivate people to act.

But if the communicator is a machine and its responses have been pre-programmed (given that fully conscious AI is still some way off), does this framework still apply? The answer is yes – and it provides some useful guidance for how financial services consumers, as humans, can embrace AI as an opportunity, not a threat, and allow it to empower rather than diminish our sense of self and personal identity.

The Five Levers of Behavioural Neuroscience

Empathy
To build trust and rapport

Motivation
To inspire decisive action

Engagement
To create interest and involvement

Confidence
To feel comfortable to proceed

Understanding
To convey complex information

Empathy: to build trust and rapport

Bots can't have feelings for humans, but humans can have feelings for bots.

The first point to make, is that bots don't have feelings, they have no ability for what neuroscientists call 'affective empathy', by which we experience and share the emotions of others. But they do have the capacity for 'cognitive empathy', which is the ability to learn and recognise emotions. For example, facial, voice and text recognitions systems can be programmed to respond in an empathetic way, by expressing sympathy or acknowledgement of a person's feelings in what Google's Deep Mind team calls 'moments of attunement'.

Oxytocin, the hormone of trust, is reciprocal. We instinctively trust people who trust us. But that doesn't mean we can't trust bots. Humans are notorious for forming emotional connections with things that can't love us back, from dolls to cars. So this must mean we are prepared to assume our trust is reciprocated and we can feel safe, for example in providing personal data or being open to a conversation. Experiments with robot carers and pets have found that people develop genuine feelings for them, so there is nothing to suggest that a robo financial adviser, for example, can't bond with clients and be trusted.

The downside is that if trust can be chemically created, it can also be abused by fraudsters or unscrupulous providers offering what feels like a warm and friendly welcome in order to mislead. Likewise, artificial empathy could be manipulated to muddle our thinking, creating biases that lead to the wrong decisions.

Engagement:
to create interest and involvement

Given the deluge of information that bombards us every day, our brain avoids cognitive overload by only focusing on what matters. It protects us from potential risks, by ensuring we are only receptive to 'new news' and can safely ignore the rest.

Interestingly, this stage of an interaction relies not on emotions, but on facts. Research has found that when we are given choice, certainty and fairness, we have a more open and collaborative frame of mind. An AI bot can be programmed to provide these. Similarly, clear, direct communications, without fudge or dissembling, go down well with our human brain, and AI can easily deliver these. But this is a watch-out for developers, who in their attempts to establish trust and an emotional connection, might programme the messaging to be overly warm, friendly, and to our brains, fuzzy and irrelevant.

The issues of choice and fairness also create another potential pitfall for AI developers. For example, we might ask our digital assistant to set certain restrictions, such as putting limits on our spending in order to save money, or only allowing a certain risk appetite in our investing. It is our choice and, for that reason, fair, but if we struggle with these constraints, we will feel belittled and controlled rather than empowered by our assistant.
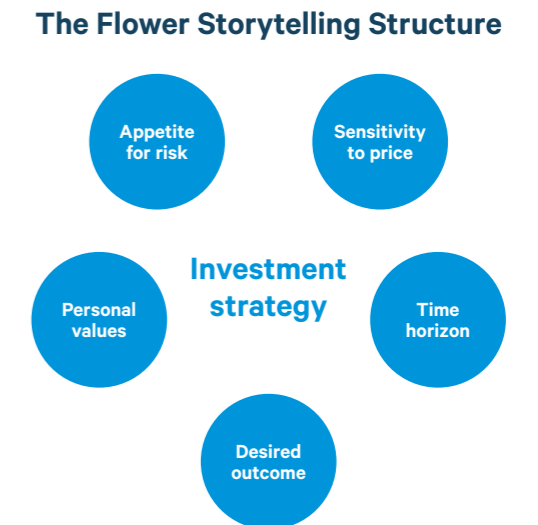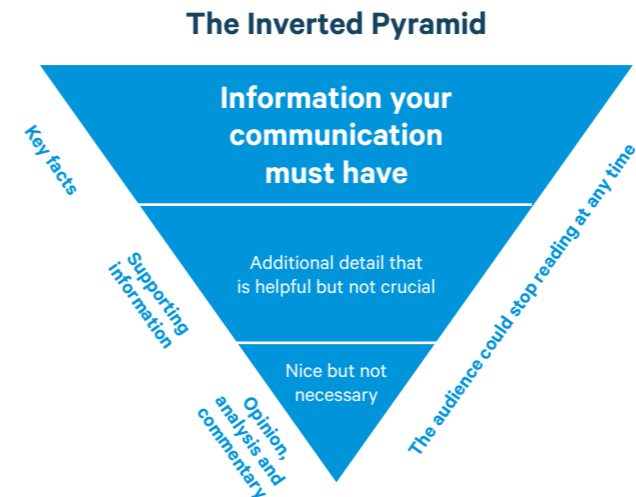
Choice, certainty and fairness make people more open and collaborative.

Understanding: to convey complex information

In the real world, there are plenty of ways to make complex messages and data more digestible, and therefore less likely to be interpreted by our brains as a threat. For example, the inverted Pyramid format, which condenses the most important message into the headline and puts the supporting information at the end, is often used in news pieces, while the various storytelling structures such as The Bombshell and The Flower are used to take the reader on a journey. AI developers can use these to good effect and, over time, bots will be able to deliver more engaging media and content, such as video and interactive infographics, tailored to the recipient's needs, in real time.

However, looking at storytelling from the perspective of the AI, thought needs to be given to how complex information is captured, as well as presented. While humans don't always like giving away information, especially private, personal data, they just love telling stories about themselves. According to one study, talking about ourselves activates the same part of the brain as eating our favourite food. The trick for the AI developer, is to create a question 'journey' that feels natural to the storyteller, with a structure that echoes the context of the questioning. For example, an investment planning factfind could use the Flower structure, using questions on appetite for risk, sensitivity to price and the investor's personal values and interests, to connect into and build towards the central concept of the overall investment strategy and approach.

The Inverted Pyramid

Key facts

Supporting information

Opinion, analysis and commentary

The audience could stop reading at any time

Information your communication must have

Additional detail that is helpful but not crucial

Nice but not necessary

The Flower Storytelling Structure

Appetite for risk

Sensitivity to price

Personal values

Investment strategy

Time horizon

Desired outcome

Confidence: to feel comfortable to proceed

In the future, AI will be able to bring our choices and decisions to life in rich and realistic ways.

The next stage in the journey, is when we humans need to feel sure that we are doing the right thing as we approach the decision itself. To do this, our brain actually wants to jump to the end of the story to see how it turns out, and then post-rationalise backwards from there. When our brain is able to visualise the outcome and effectively 'rehearse it' to experience the benefit, we feel more positive and optimistic.

The learning here for AI developers, is in creating visualisations of scenarios with the customer at the heart of the story. For example, while today's retirement cashflow modelling tools can present data back to us graphically, in the future, our AI will be able to bring potential scenarios to life for us in a much richer, more immersive and compelling way. This is likely to have significant consequences in our attitude and behaviour towards events that are set in the distant future, such as saving for retirement or making a will.

Another aspect of Behavioural Neuroscience that comes into play at this stage, is a technique that works in the real world – acknowledging people's hopes and fears about a decision or a situation and openly addressing them. But if an AI, having analysed our expression, voice or language, says, "I can tell you're feeling nervous about committing to this mortgage," for example, how might we react to that? Many people will feel that their privacy has been violated and they are being exploited, thereby losing confidence and disengaging from the process.

Motivation: to inspire decisive action

At this final stage in the journey, our brain just wants an easy life, to be able to proceed without over-thinking. In the real world, effective communicators will present the call-to-action as a series of small, simple steps rather than a giant leap and an AI would do the same. The key is to ask for actions, rather than decisions.

However, by the same token, this desire for an easy life is the same button that fraudsters and unscrupulous providers would press. Time-limited offers or restricted availability which create a false sense of urgency, confusing options and reams of small print can panic or overwhelm the victim into making the wrong decisions.

In conclusion

While the opportunities for AI at the 'back end' of financial services firms are certainly positive, in terms of business efficiency and lower costs, improved risk management and better planning and strategy, it's the ability to create better customer experiences that will have the biggest impact on the category in future.

According to Gartner, by 2020, 85% of our engagement with an enterprise will occur without interacting with another human, and consumers will be perfectly happy with that, as long as the experience and the outcome work for them. Of course, there is the possibility that human financial advisers become disintermediated and their clients set adrift in a lonely world of bots and automation.

Likewise, wherever there is money to be made, there will be fraudsters and unscrupulous providers. But advances in fraud detection, powered by AI and blockchain technologies, able to spot suspicious behaviour and trigger security alerts, are becoming ever-more sophisticated and robust, to keep consumers safe, while regulations around ethics, transparency and fairness will continue to evolve to keep pace with the market.

At Teamspirit, our own research has found that service and support don't need to be human, as long as they feel human. By using Behavioural Neuroscience to understand how real people think and feel, and how we make decisions, financial brands can create customer experiences that help us enjoy a more positive relationship with our money and our financial selves.

By 2020, 85% of our interactions with a company will be non-human.

Teamspirit

Here at Teamspirit, we work with some of the most exciting, forward-looking financial services companies in the world, to bring their vision of the future to life. We're helping them to transform how they connect and communicate with their clients in an age of digital transformation. If you would like to talk to us about what we could do for you, please contact us at **hello@teamspirit.uk.com**