Preface

by *Simon Freakly*[*]

Change and instability increasingly characterize today's business environment. Technologies are evolving at an unprecedented rate. Patterns of consumer behavior have become less predictable. Geopolitical structures that provided stability and a basis for economic growth around the world have become progressively more uncertain. And today's business leaders must process more information around all of these developments at a greater pace and under greater scrutiny than ever before.

In a world so in flux, businesses across industries must move fast to stay ahead. As Claudio shows in this compelling new book, this is particularly true for the financial services industry. Banking has already been transformed significantly by technology, but artificial intelligence and other new technologies are driving even more profound changes. The picture that he paints is a future for banking that is much less asset-intensive, and perhaps less human-intensive, than today. Many core functions, such as credit extension, insurance underwriting, and investment management, could be disintermediated or, at very least, fundamentally reshaped over time. And like in other industries, banks are also faced with a new set of competitors in technology companies that have fewer regulatory or legacy issues than they do.

Complacency, or institutional paralysis, in the face of such threats can lead to value destruction. Ultimately, in the absence of swift and dramatic action, external forces may even take over and affect change outside an organization's control.

Those banks that move quickly and decisively to confront these challenges will be best positioned to lead the industry into this new future that Claudio has so ably outlined for us.

---

[*] Chief Executive Officer of AlixPartners.

1 Asset-Less, in the Digital Age

1. Like a hedge fund

Banks have traditionally been long on real estate, almost like hedge funds, and closely intertwined with the real estate cycle – driving its fortunes and potential asset bubbles as for no other sector. Banks have gone long ever since they started to use real estate collateral to strengthen their recovery ratios and "cure rates" in cases of default on a corporate or retail loan, and have lent to real estate developers and to the construction industry, as these sectors drove most of the growth in the developed countries after the Second World War. Banks have often underwritten loans purely on the strength of their real estate collateral, feeling that they were "buying a call" (option), whilst often ending up having "sold a put" to the counterpart – ready to dump cracking collaterals at the trough of the cycle, whilst retaining most of the upside when their lending options are "in the money".

Real estate assets have also become the object of the banks' leveraged finance business, of their trading strategies (either as principal investors or as arrangers and financiers) and life insurance and wealth/asset management and pension businesses (as a key component of their investment strategies for end customers). Finally, banks have traditionally been notorious for their sprawling real estate "core" assets, which in the past they used to grow their origination and distribution might (via their corporate and retail branches), to host their product and service factories and, more mundanely, their global and regional headquarters.

Real estate, and trophy assets specifically, with their image of tangibility, concreteness and opulence, have also been working well to reinforce the reputation of banks as "fortresses", full of money and with almost no risk

of failing – the new kingdoms of unimaginable wealth to whom people could entrust their savings and sleep well. If you like historic buildings and architectural works of art, you may be interested to find many of these used as headquarters by international and local banks in Europe. If you are keen on new, technological and imposing glass and steel skyscrapers, you may be drawn to those banks that inhabit entire buildings of up to 30, 40 or 50 floors, as in Canary Wharf in London, or at La Defense in Paris – and in the new financial districts of Frankfurt, Madrid and Milan (think here of the Gae Aulenti, City Life and Expo new areas).

If you happen to visit such banks, whether historic buildings or new skyscrapers, a simple thought might come to mind almost immediately: These are big and imposing, trophy and technological, and full of people working in their cubicle-shaped offices all day long, with limited or no client interaction and no "physical" activity that would point to the creation of something – as in a manufacturing plant, for example. A banking magnate could have something like 30,000, 50,000 or even 100,000 white-collar people working in just such a way, not to mention the further twice as many doing front-office work in the bank's individual branches.

The advent of phone and then online and mobile banking is already a well-known phenomenon, with hundreds of thousands of front-office jobs fast disappearing, along with the closure of banks' local branches. Yet even assuming that the digitalization journey of most customers keeps accelerating and that most of the real estate used for origination and distribution gets shut down, banks will still command many of the most iconic historic buildings and the highest skyscrapers, with a growing number of people doing compliance, risk management and other regulatory, administrative and control work, on top of the product factory (asset management, financial trading etc.) and related service companies (e.g. IT), not to mention the general management of the bank.

2. The empty bank

What about if, in a five- to ten-year timeframe, we could walk into the imposing headquarters of one of these banks and find nothing except a few stacks of super-powerful and super-fast computers, probably kept cool by the proximity of a river or lake and also probably stored somewhere safer, more remote and cheaper than downtown?

What about if we should find, behind the shiny doors of these imposing headquarters, almost no people at all, save for a bunch of computer scientists going around in jeans and t-shirts, and a few remaining top managers (also data scientists?) pretending to check up on the day-to-day business and set the bank's future strategy? All of this could well happen, as most of the back-office people could have been, by that time, automated away. And middle-office employees may also get crowded out by robots working 24/7 and with no paid holidays or sick leave – let alone ever going on strike. Even front-office resources could be arbitraged away by human-like voice recognition-interaction activated devices, by autonomous applications doing advisory work online and, potentially, even by "alive and kicking" robots (very human looking ones, if we should wish) that can move across the city in driverless cars or by flying in drones, to deliver the "wine and dine" aspect of the business direct to our doors.

Eventually, even their reprogramming and coding, as well as strategy and "command and control" activities could be done by machines better able to define the "best possible strategy" among the many millions potentially available, killing off the last remaining and well-paid striped suit-wearing top managers. Such a scenario is hardly impossible, as the advances in artificial intelligence (AI) are progressing on multiple fronts – as we shall discuss in this book – and at an unprecedented pace. Furthermore, most of the relevant technology, hardware and software are becoming easily available as a service and "over the cloud".

In the (by definition) immaterial world of finance, banks, born (almost by design) in a state of unbearable "lightness", could also fully dematerialize and become "empty", in the full sense of the word. They could basically evolve into an ethereal set of contractual obligations and rights, granting them the ownership and use of a mesmerizingly large quantity of data and, more importantly, of a more selective but extremely powerful set of AI super-applications. With the help of robots (and cyborgs, if needs be), these would progressively take care of most of the strategic and operational functions of current and future business models, thus becoming progressively not just "asset-less" (as discussed in our previous book on digital transformation in financial services1) but also (and maybe more worryingly) "human-less".

---

[1] Claudio Scardovi, *Digital Transformation in Financial Services,* Springer, 2017.

Empty banks in empty buildings, with a limited balance sheet and an even more limited human workforce. But let us take one step at a time. The "unbearable lightness" that we mentioned (the extremely high leverage driven by banks' money multiplier model, the time lapses and duration mismatches driven by their intermediation model, the excessive velocity of equity usage driven by securitization and derivatives, etc.) have led, we have argued, to the many structural weaknesses of the global banking system that were already quite evident during the 2008 financial crisis. Furthermore, digital innovation and the challenges introduced by FinTech and "big digital" players are progressively killing the banks' traditional business model – with old savings institutions and lenders ignoring technological innovation and related disruptions at their peril.

The "synapsis bank", as we discussed, was then the potential end state of this accelerating evolutionary process – driven by the so-called fourth industrial revolution and leading to a new wave of banks' failures in the name and spirit of a Schumpeterian "creative destruction" to come.2 By synapsis bank we mean a bank business model where the bank is capturing and capitalizing all kind of structured and unstructured data and then, using applied analytics and AI, is able to create new intelligence in the system, making things happen in the economy in the best way, activating multiple stakeholders in the ecosystem and designing and executing new, value adding solutions. In a way this is similar to what the synapsis are doing, getting data as input to create intelligence in our brain and then activating muscles, joints, arts etc to drive our bodies toward a certain actions that are always informed by our thoughts that keep changing and been re-elaborated whilst we act.

A synapsis bank would be built on a different and more sustainable version of lightness and basically working on the basis of five main building blocks.

3. Asset-less, brain heavy

Synapsization is, in fact, in our strategic framework of analysis, driven, by five key elements, or building blocks. The first one is data – the structured and unstructured information making up the "big data" that is the "new

---

2 Claudio Scardovi, *Restructuring and Innovation in Banking*, Springer, 2016.

Table 1 The synapsization framework

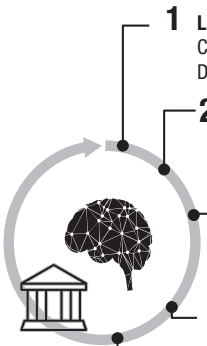

1 **Listen** to data – *Data and information management*
Collect unprecedented quantities of data (quantitative, qualitative, structured, etc.)
Develop more flexibility in data retrieval, combination and sequencing

2 **Touch-base** on information – *Applied analytics*
Derive intelligence from advanced data analytics
Build-up of goodwill out of data and information

3 **Taste** market appetite – *Interconnectivity*
Interconnect a larger number of stakeholders providing more valued added
Match evolved market needs in terms of money, time horizons and risk profiles

4 **Smell** for new opportunities – *New business solution design*
Proactively develop new business models, leveraging technology partnerships
Embrace a "solution seeking role"

5 **See** renewed Client confidence – *Trust and credibility management*
Sustain and accrue "trust capital" as the foundation of any sustainable innovation – manage cyber capital at risk in the best way

oil" of the digital intelligence economy. The second is applied analytics, able to mine data and get intelligence out of them. The third refers to the management of enlarged ecosystems, composed of multiple stakeholders that can enjoy their greater interconnectivity (or junction playing, following the brain metaphor) as favoured by the "synapsis bank". The fourth points to the economic solutions or ultimate use cases that, either by design or by chance, are created and acted upon by the "synapsis bank".

Finally, the fifth element refers to the cyber-trust embedded in the new economic system, as the synapsization created is fully developed and consolidated on the basis of the "cyber-trust" (and put at risk by its antithesis, the cyber-risk, that defines a "cyber-trust capital at risk"), which is the new scarce resource at play.

Synapsis banks are the answer to the unbearable lightness of the traditional bank business model and the solution to the digital transformation challenge, where banks create synapsis and become financial assets light, as everything is progressively driven by information, intelligence, networking, solutions and cyber-trust. These are the key elements of the business transformation that could happen to banks and to their economic system, allowing them to remain competitive, overcoming the threats of new challengers.

Global leading companies that are worth tens of billions of dollars and are asset light or asset-less are nothing new, by the way. Uber, the world's largest taxi company, and one of the largest in transport overall, owns no taxis. Facebook, the world's most popular media owner, creates no content. Airbnb, the world's largest accommodation provider also owns no real estate.[3] So, what about a bank becoming the most valuable one globally and holding no financial assets at all? It would not be inconceivable and maybe it is already happening, as balance sheets are reduced and substituted by big data and intelligence assets.

Whether a traditional bank, with a full license and all, could do this may still be debatable. Certainly, the move into banking of big digital players' (the FAANGs of the western economy – Facebook, Amazon, Apple, Netflix and Google) could aim at creating billions in financial value out of "thin air" (meaning, with almost no financial assets, sitting on their balance sheets with no regulatory capital at risk).

They could certainly pose a serious competitive threat to banks and even change the structure of the overall financial system and the paradigms (and related risks) of its stability (it is worth noting here that big digital is also, at the moment, only lightly regulated, often escaping the supervision of banking authorities).[4] Should they set up financial services divisions or business units (let us call it a "proxy bank"), these would already be worth a few billions. Even tens of billions – if we think of Alibaba, Tencent or Baidu, the FAANGS equivalent of the Chinese economy – as these big digital giants are apparently making faster inroads in the Asian continent.

3. Big, digital and dangerous

Regulations (or lack thereof) are also lending a hand to big digitals. This not only allows them to dodge the heavy burdens forced onto traditional banks after the global financial crisis, but also imposes new competition rules on banks, compelling the latter to open their informational door – somewhat cracking the code of banking – in a faster and potentially more

---

[3] Andrew McAfee and Eric Brynjolfsson, *Machine Platform Crowd,* Norton, 2018.

[4] Martin Arnold, "Finance chiefs warn on Big Tech's shift to banking", *Financial Times*, 5 February 2018.

devastating way. Let us consider, for example, the consequences of "open banking" and of the PSD2 (Payments Service Directive 2) in Europe, which is forcing banks to provide access to the accounts of the customers who authorize it, providing interfaces that can be used to source data.

Francisco González, Chairman of BBVA, one of the most "pro-digital" bankers in Europe and a vocal advocate of the necessity for banks to embrace digital transformation and AI, has warned how players such as Facebook, Amazon and Google in the USA, and Baidu, Tencent and Alibaba in China, could soon replace banks. Amazon is providing payment services and loans to merchants that offer their goods on its website. It can monitor, in real time, the health of their sales and cash flows, and the "likes" of current and prospective clients. It can even get hold of their products as a collateral, as they are already physically hosted in Amazon warehouses – which bank could match, as of today, such a powerful lending proposition?

Big digitals are not just offering financial complements starting from the e-commerce world. Take, for instance, the growing "cloud services" business they offer, which allow the likes of Microsoft and Amazon to manage the information and the applications used by each client company. They cannot access the information and applications but have access to a view on the flows of the market. This can help in the prediction of trends and patterns, as investment banks have been doing for years, leveraging their brokerage "flow" businesses to get insights and hints for their principal trading business.

The potential ways in for big digitals are increasing by the day, as they spot the opportunity for disruption in the global financial system and start tapping the profit pools that have been the preserve of banks and other financial intermediaries for centuries. Apple and Google are, for example, already in payments and electronic money and Facebook is now entering too and it wants to grow big.

Was it not money, after all, that started as a way to allow people to socialize and exchange, well before the advent of communication across space and international commerce? Moreover, once you are in payments, it is easier to manage money and savings – and then consumer or commercial lending or personal and property and casualty insurance – if you know what people like or dislike and suggest to friends and people of similar cohorts.

What about Alibaba then, which is already by far the largest insurance distributor in China and also big in payments with AliPay? The situa-

tion is similar with Tencent and Baidu. They are now becoming active in "instant lending", providing underwriting decisions in a few seconds, based on the online behaviour of customers. All these players have huge amounts of social and behavioural data and are (so far) completely financial assets free.

They still lack huge stocks of transactional and other financial data (hence their focus on payments or the potential urge of customers to buy some financial intermediary just to get that missing piece), but they are rapidly making up for that gap. More importantly, and often on the back of their offer of cloud computing – a storage solution increasingly used by banks too – they are investing heavily in AI applications, as they try to convert information into intelligence in an efficient and effective way. This is happening in an accelerated manner, "over the cloud", asset-less and almost without any people being involved.

4. The age of digital intelligence

As these big digital players turn into big cloud service outfits, they also become the leading masters of AI capabilities and related services – they can potentially offer some slimmed-down version to banks and to all kinds of other financial intermediaries, who are then enticed to move their hardware and software over the cloud. Almost by default, as the cloud managed by the big digitals expands, banks' real estate core assets empty out, and a greater concern arises over this massive concentration of power on just a few global platforms. Furthermore, considerations of the systemic risk potentially linked to the default – or cybersecurity crisis – of any of these big cloud platforms become a high priority for regulators and policymakers to address.

This is, however, just the very start of a new age of digital intelligence to come. In fact, the digital transformation that is underway in financial services will not be ending with a change in the business and operating model of the main players and markets. The overall structure of the global financial system is also transforming, driven by profound changes that will eventually introduce new equilibria and steady states – towards a target end state marked (potentially) by the dominance of digital intelligence: one that is super-human in nature, artificial by design and potentially better than the human.

These changes are already dominating several other industries and could be summarized by following the framework created by Andrew McAfee and Erik Brynjolfsson, which identifies three distinct trends.

The first is marked by the accelerated expansion of technological computational power and of machines' capabilities, where AI has a central role in setting the scale and direction of change. The second is driven by the parallel growth of platforms, which really started with the world wide web and now keep aggregating multiple stakeholders in novel ways, where again AI can play a critical role in the why and the how – from the design to the ongoing management of these emerging ecosystems. The third is finally leveraging crowds, sourcing a very large amount of knowledge, expertise and skills, via the web and (again) with the obvious application of AI.

These three mega-trends, not dissimilarly from our the "synapsization" framework with its five building blocks, are fostering the creation of digital intelligence and supporting its ultimate dominance of the economic system. This represents a digital intelligence that is not developing in a vacuum but is driven by the specific utility functions of multiple stakeholders and applied to idiosyncratic business needs that are solved by given business solutions and use cases and then delivered and "accepted" by end users on the basis of their cyber-trust.

In truth, we are already living in something of an "age of digital intelligence", as most of the apps and digital utilities we have started using every day – from filtering spam on our email, to suggest our best next book to read or where to holiday – a re already delivering on its promises: "Hey, Siri, can you check out a list of such digital nuggets for us?"

Nevertheless, the current age of digital intelligence cannot be considered the end state as, far from it, further disruptive and potentially explosive acceleration of this will be driven by the development of AI technologies and platforms – considering their broadest and most holistic range of potential applications.

Unlike past wishful thinking and beyond current hype, the science and practical applications of AI appear now at a crucial point and poised to continue to develop at an unprecedented multi-decade rate of acceleration. Specifically, efficient deep and machine-learning techniques have shown very promising results in reducing the effort required to advance and train models by automating features' extraction, albeit still with a limited ability of "abstraction" (as defined as the capability to leverage the obtained learn-

ing across multiple, heterogenous, unrelated fields of application). Moreover, heavy data requirements that made these techniques very costly are quickly becoming more affordable given new technological development. Finally, unstructured learning is gaining ground supported by advances in reinforcement learning and other simulation-based modelling techniques that, bit by bit, make the limited "abstraction" less of a constraint. This is in turn potentially multiplying the opportunities for areas of application.

5. What is AI?

We may be forgiven if, in the introduction of the topic and in the development of our argument for a "burning platform" in the financial services industry to consider the implications of AI, we have not yet offered a basic, clear and comprehensive definition of what AI actually is. Indeed, what is it? Looking at the past and current literature on the subject, it is not easy to come up with a common definition, as AI has been used in a number of ways and with different meanings. It is therefore easier to define it as referring to a number of capabilities.

AI certainly includes machine learning (ML) and deep learning – techniques used to make predictions on trends and behaviours based on data and via algorithms. But it is also about computer vision – meant as the set of cognitive functions able to reason, interact with and learn from visual data or digital images (such as the ones used to action and control driverless cars), and it is about natural language processing (NLP) – meant as the machines' understanding and analysis of human language. AI could also include biometrics – used to understand people's physical characteristics and behaviours, and to recognize and prove their identity for security reasons. As a mix of the above, it could also be recognized in the development of virtual, digital assistants – able to interact with their end users (or masters) on a wide range of topics, in writing and even using simulated human voice.

Aside from the list of capabilities covered, AI can then be defined from a use case point of view, as a system of technologies that can have predictive power and can master some degree of autonomous learning that helps the overall system in making continuous, marginal advances. AI is then employed for use cases that leverage its pattern detection and its recognition of trends and irregularities in data. With foresight, AI can determine

the probability of future events. It can then support decision-making, generating rules from general data and applying specific profiles against these rules, hence customizing the answer to slightly different questions. It can then help in optimizing specific outcomes, leveraging data and interacting with humans in various ways.

Within AI, neural networks are then usually classified as a sub-set of ML techniques – they are AI systems based on simulating connected "neural units", loosely modelling the way in which neurons interact in the brain – via deep learning (as they work through multiple – termed "deep" – layers of simulated, interconnected neurons). Thus the specific interest in using them. Neural networks work then through five main techniques.[5]

In the first technique, called "feed forward neural networks", the information moves in just one direction (forward) from the input layer and through the hidden layers to the output layer. There are no loops in the network but just a one-way direction. In the second technique, neural networks are "recurrent" and include loops. Then, in "convolutional neural networks", artificial neural networks show connections that are inspired by the organization of the animal visual cortex, hence being more complex and articulated. In the "generative adversarial networks" technique, two neural networks contest one another in a zero-sum game framework (in a kind of evolutionary Darwinian challenge). Finally, in the "reinforcement learning" technique, systems are trained and developed by receiving virtual reward or punishment, learning by trial and error.

Whatever the technique, making effective use of neural networks requires large, labelled training datasets and access to sufficient computing infrastructure and power. A supervised deep learning algorithm could, for example, achieve good performances with around 5,000 labelled examples per category and will match and exceed human performance when trained with a dataset containing at least 10 million labelled examples.

Supervised learning requires humans to label and categorize data and is therefore expensive. However, promising new techniques are emerging to overcome these data bottlenecks and include reinforcement learning, generative adversarial neural networks, transfer learning and "one host" learning, which allows a trained AI model to learn about a subject based on a small number of real-world demonstrations or examples.

---

[5] "Notes from the AI frontier: applications and value of deep learning", McKinsey and Company, April 2018.

6. Not for everybody

The advancement in these techniques looks impressive and the acceleration in the investments and related developments in the field also appears very promising. Still, a number of limitations are apparent and well known. For example, neural AI techniques can excel at analysing images, videos and audio data, because of their high dimensionality – as multiple layers in the network can learn to represent the many different features found in the data (for facial recognition, the first layer in the network can focus on raw pixels, the next on edges and lines, another on generic facial features and the final layer might identify the face). But they require a lot of data and are still prone to error as people age, and there is a need of continuous retraining to match potential conditions. Training then requires new data, both big and deep, and related labelling and refresh.

Moreover, as mentioned, the results are sometimes difficult to explain in human terms and their generalization (abstraction) to become useful in other circumstances is still limited and subject to the vagaries of data and algorithms. In a nutshell, AI can become extremely powerful and can be used by masses of end users, but it is not for everybody to manage and requires serious technical skills and relevant investments of scarce resources.

Banks and other financial intermediaries could, therefore, try to develop AI capabilities and use cases by building in-house capabilities, or by outsourcing them and leveraging the "AI-as-a-service" proposition offered by cloud vendors, leaving to their last stage the definition of the most important business requirements and the subsequent changes required in the way the organization is structured and how processes and people work.

Investments in AI have already grown dramatically in recent years, with USD 11 billion spent in financial services since 2010 to 2017 and an expected USD 47 billion by 2020 according to IDC.[6] And their range of applications are already vast, from enhanced customer service and engagement (e.g. managing the customer's interface in an omnichannel strategy), to lower operating costs and enhanced efficiency (e.g. in document management and compliance costs), to revenues generation (e.g. via increased cross-selling and robot-advisory), to optimized risk assessment and management (e.g. via

---

[6] "Is AI the next revolution in retail banking?" UBS global research, 27 November 2017.

automated and AI-powered credit approval, fraud detection and internal audit processes), to many other narrow (or deep) applications.

Broadly speaking, we could summarize the urge of banks to invest in AI because of their pursuit of a given set of opportunities, more or less ambitious, as they can aspire to "do the same thing, better" (creating, for example, faster and leaner operations, or tailoring products and services and ensuring a ubiquitous presence), or "do something radically new" (promoting and implementing smarter decision-making, creating new products and services and mastering new value propositions across their fundamental functions).

More specifically,[7] AI thought leader Ajay Agrawal, Rotman School of Management professor and chairman and founder of AI start-up Creative Destruction Lab, argues that AI serves a single, but potentially transformative, economic purpose as it significantly lowers the cost of (good) prediction.[8] As this cost keeps dropping, people use more of it for traditional problems (such as inventory management) and predict potential events faster, cheaper and better. At the same time, they can start using AI to solve prediction problems that are not, as of today, thought of as such (such as autonomous driving) not just doing things faster, better and cheaper but also overcoming well-known human cognitive biases and making them more objective.

In sum, AI can be defined in terms of the capabilities, or solutions it can offer for use to the many and – given its limits and constraints and its inherent complexity which makes AI manageable by just a few – it could also be defined for some general-purpose objectives, such as doing more with less, or doing better and faster or just smarter. This is all useful stuff, but still short of a serviceable definition for the scope of our banking discussion.

7. A serviceable definition of AI

A more useful and serviceable definition of AI and of the changes it can deliver in banking is needed, to better support us in the analysis of the potential changes in the shape and structure of the global financial system.

---

[7] "The economics of artificial intelligence", McKinsey and Company, April 2018.

[8] Joshau Gans, Avi Goldfarb and Ajay Agrawal, "Prediction machines: the simple economics of artificial intelligence"; Max Tegmark, *Life 3.0. Being Human in the Age of Artificial Intelligence*, Allen Lane, 2017.

To that end, we shall provide our own definition, based on a three-legs proposition of deductive-inductive and emotional artificial intelligence. Our reasoning, far from being scientific, is logically derived by the way we usually interpret human intelligence. Machines are ultimately unbundling human decision-making anyway, clarifying the distinct roles and paradigms played in this – hence our idea to discuss AI by analogies.

Human intelligence, in our starting assumption, can be analysed on the basis of its deductive and inductive reasoning capabilities and of its emotional intelligence. Deductive reasoning (also known as logical deduction) refers to the process of reasoning from one or more statements (or premises) to reach a logically certain conclusion. If all the premises are clear, and the terms are clear, and the rules of deductive logic are also followed, then the conclusion reached is necessarily true. This is also called "top down", as a conclusion is reached reductively, by applying general rules which hold over the entirety of a closed domain of discourse.

In inductive reasoning, instead, a conclusion is reached by generalizing or extrapolating from specific cases to general rules, hence retaining an epistemic uncertainty.[9] The premises are then viewed as simply supplying some evidence for the truth of the conclusion – which can just be the most probable, but is never certain. Inductive arguments, as a consequence, are then either strong or weak, hence more or less probable.

Finally, emotional intelligence refers to the capability of individuals to recognize their own emotions and those of others, discerning different feelings and labelling them appropriately, then using emotional information to guide thinking and behaviour and manage emotions to adapt to challenging environments or achieve specific goals. Emotional intelligence also reflects the ability to combine intelligence, empathy and emotions to enhance thoughts and the understanding of interpersonal dynamics, and is seen as a complement or enhancer of a human's general intelligence capability – sometimes even becoming the most important element of this.

AI has often been closely associated with deductive reasoning and, mistakenly, relegated to that dimension of intelligence. It is relatively easy to think of AI applications that, leveraging superior computing power and a super-sized access to billions of units of data, and on the basis of predeter-

---

[9] Wikipedia, "deductive reasoning, inductive reasoning", August 2008.

mined set of rules, can easily achieve great performances and outbid the best of human brains at all kinds of top-down reasoning challenges.

AI can, however, also perform inductive reasoning and work in a "bottom-up" way, getting close to the creative and innovative dimensions of human thinking (wrongly considered unachievable for machines). A number of examples and cases are already available to support this argument. Certainly, the number and the impact of potential AI applications that could leverage this dimension of intelligence would be great and highly relevant.

In fact, the beauty of AI is that it can work a lot on social data, captured in real time. It is then able to recognize all kinds of nuances in the recurring or emerging sentiments of a market or society, not to mention its interpretation of the emotions of people, through its psychometric read of their facial expressions, tone of voice, pose and attitude, variations in temperature and sweating and of course, of what they write and say.

This third dimension of (emotional) intelligence resonates well with the empathy that certain use cases could build with end customers and with their ability to influence thinking and behaviour, obviously with interesting applications in disparate matters such as compliance and conduct risk management, fraud detection and prevention, risk management across business lines (from trading to lending), to cross-selling and upselling to customers. Imagine a robot that could understand how we feel on a given day and behave accordingly to make us as happy as possible, without having to control its own emotions (being a robot).

8. Curious like a robot

In its potentialities, AI is much more (in our holistic definition) than pure deductive reasoning and it should include and develop as much as possible the inductive and emotional dimensions. Based on this more serviceable definition, several solutions – or use cases – can be thought of and applied to the fundamental functions of the global financial systems and across banks' main processes, front to back. Using this definition could help us to identify new opportunities for change that are not just related to "doing the same thing cheaper and better" (mostly leveraging AI deductive reasoning) but also, and more importantly "doing something radically new", with truly value-adding solutions (with inductive and emotional intelligence).

Even if it looks counterintuitive, AI (and robots) could be made creative and to understand the matters of the heart. This looks technically possible and, if feasible in practice, all kinds of AI applications could result in much more radical performance step-changes, if they can stretch and leverage all three dimensions of holistic intelligence instead of just one. The main question then is how to make theory become practice and create "curious AI systems" (and robots) that do not just create algorithms and intelligence by the (deductive) book.

To train thinking machines, we need a lot of data and the supervision of humans as well – to label data, for example, but also to suggest, more importantly, which questions AI should try to answer. A typical supervision could be, for example, to show a large set of pictures labelled "cat" and "dog" and have an algorithm learn to discriminate between the two (using mostly deductive reasoning). Another game could be to set a specific goal within a highly structured environment, such as achieving a high score in a video game, then letting the machine create different algorithms and foregrounding the best performer able to achieve the objective (using mostly inductive reasoning), as DeepMind did to win at Go.[10] The same two methods can also envisage learning programmes that involve social data, and hence emotional intelligence.

Still, these two types of supervision, apart from being quite expensive (as human support is costly, and the machines would use lots of energy to run billions of simulations), score badly on "abstraction", as once the training programme and guidance is completed, it cannot be used for other means. However, if AI, and even more so in our "holistic" definition of three-dimensional intelligence, is really going to take off, something more is needed – as it is needed for any (deductively, inductively and emotionally) intelligent person. Curiosity, in fact, may have killed the cat, but has surely been the basis of most advancement in human knowledge and civilization. Intelligence, of all kinds, is of limited value if no curiosity is pushing it towards new ends.

According to Pierre-Yves Oudeyer, a researcher at Inria, the French national institute for computer science, it is possible to train AI and its related algorithms to leverage their, almost human-like, "curiosity".[11] This can

---

[10] Max Tegmark, *Life 3.0. Being Human in the Age of Artificial Intelligence*, Allen Lane, 2017.

[11] "Heads full of brains, shoes full of feet", *The Economist*, 31 August 2018.

Table 2 Holistic intelligence and curiosity (and serendipity)

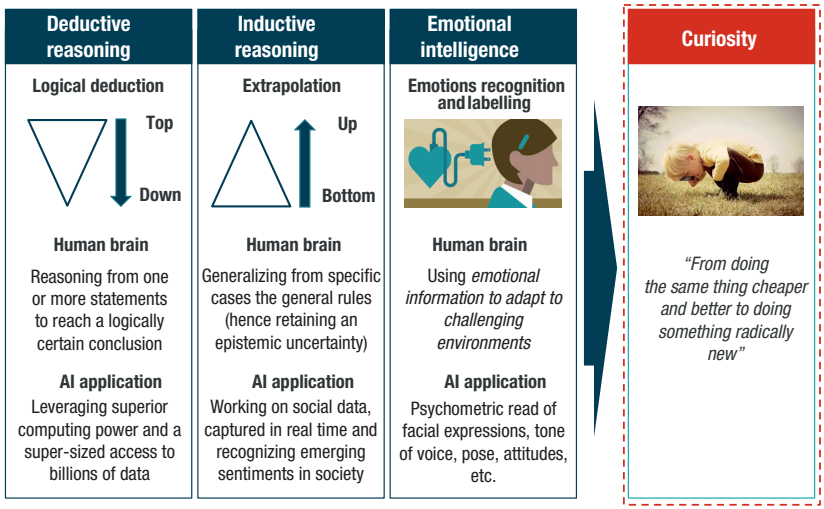

be done through developing AI that use their own intrinsic rewards system to inspect the world around them and obtain and process lots of data to create new use cases and solutions to unspecified issues or problems.

For example, the incentive system used could be based on "prediction error". In this, the AI system starts exploring its environment of reference (such as the trade executed in a trading room) and looks for things (the outputs of the trades or the traders' conduct) that deviate significantly from what it has predicted. It would then potentially sort out new answerable questions because of the intelligence it has been gathering (if trades are performing better in early morning instead of late afternoon, then there may be a reason linked to traders' tiredness levels – and the solution could be their getting a longer break; or, if the desk that exhibits more individualistic behaviour of traders, and where these tend to do work very early or overnight when most people are away, is associated with greater fraud cases, then these behaviours should be controlled more closely and incentivized to stop). It could then create solutions out of ingenuity and even serendipity.

In these applications, the AI work is focused on its prediction error rate, rather than the nature and content of the error itself. As soon as this is reduced, the machine can move to something else (or, if this value is

consistently high, meaning that no prediction is possible, the machine still moves away), basically working out the solutions even before knowing the right questions. In some tests, curiosity-driven learning has been working well, and scoring better – for example in video games such as Pong: the AI agent avoids dying in the game as this brings it back to the beginning of the game, where it has been many times and predicts dynamics well. Dying is then boring, and moving on to the next stage of the game is quite exciting (as the agent is curious to know more).

9. We have a dream (it could be a nightmare)

This all resembles the dynamics of human evolution, where curiosity helped a great deal and serendipity (or just mere luck) took care of the rest. A set of AI-powered "curious" algorithms (or robots) could then look to keep developing after the earlier stages of human-assisted learning, and they would become better suited to manage sparse environments devoid of much data but still able to stimulate their curiosity to find new things to model and questions to answer. Or they could just end up to the great use case and solution by mere chance.

The applications coming out of these "holistic" intelligence and related areas linked to curiosity could be very powerful indeed, as we shall discuss in later chapters. Of course, multi-pronged AI intelligence and even its supporting curiosity cannot exist in a vacuum as these capabilities need to be developed in conjunction with other technologies, such as the Internet of Things (IoT), the "public" cloud, quantum computing and blockchain, that we shall shortly discuss in more detail – but these great developments are already happening.

Nevertheless, AI is the polestar of digital and business transformation, and the dream of development for humanity – as, given the required technological developments, we could potentially get to a super-human, artificial 'general' intelligence (AGI) able to beat the human brain across all kinds of holistic dimensions and at any sort of business, social, political, scientific or other. This dream could then rapidly turn into a nightmare if this AGI was not used for the benefit of humanity but started to act against it.

This dream (or nightmare) of ours has not yet been fulfilled, and the case for moving from narrow intelligence to general intelligence to su-

per-intelligence (which we shall discuss in chapter 2, mighty and scary as it may be, has not yet been fully proved with regards to its feasibility, largely remaining the subject of science fiction films.

Still, we believe this is, at least in theory, the final endgame and, as such, a further move away from the "synapsis" bank in our transformation discussion needs to be considered, almost superimposing itself on the old and new paradigms of the global financial system, which are changing as AI applies its holistic intelligence to real-world use. We will call this end state the "singularity bank" (the concept can be applied to all sorts of other financial intermediaries or companies), which refers to a company that is ultimately dominated and run by AI (or by AGI).

It is then useful to define the notion of "singularity", starting from its absolute meaning and working out a business definition that is serviceable enough for our discussion of AI's impacts in financial services. We can start from a definition provided by Irvin Good who, in 1965, was arguing how "an ultra-intelligent machine (could) be defined as a machine that can far surpass all the intellectual activities of any man, however clever. Since the design of machines is one of these intellectual activities, an ultra-intelligent machine could design even better machines; there would then unquestionably be an "intelligence explosion" and the intelligence of man would be left far behind. Thus, the first ultra-intelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control."[12]

Singularity is then, to paraphrase Good, the hypothesis of the invention of a super-AI (an AGI) that will inevitably and abruptly trigger uncontrollable technological growth, but also lead to unfathomable changes to human civilization, that could either be extremely good, or very bad – and almost unrecoverable by their very nature. More recent AI thought leaders have expressed similar hopes and worries.

Yuval Noah Harari, a native of the 21st century, is of the mind that human beings will change more in the next hundred years than they have in all their previous existence.[13] This will be driven by a combination of biotechnology (needed to make us more "cyborg-like" and encapsulate into us AI apps and other pieces of smart technology) and – of course – AI that could enable some people to become digitally enhanced. A further vision

---

[12] Tegmark, 2017, op. cit.

[13] "In the kingdom of cyborgs", *The Economist*, 31 August 2018.

of a dream (or nightmare) is digitally enhanced "human gods" who will be able to do all kind of wonders – definitely prevailing on the remaining "basic humans" that will not have the cash for an upgrade.

10. A singularity moment?

Are we approaching a "singularity moment" in global finance and banking and in the history of the development of economies and human societies alike? Is this something to welcome and look forward to, or a matter of alarm and concern for our future – as financial professionals and human beings – as even Stephen Hawking and Elon Musk have rung the bell on the risks driven by the uncontrolled rise of AI? Will it empower us and bring still unforeseeable opportunities, or will it seize control of our own destiny?

In practice, we may be facing not just a single, but multiple moments of singularity in the economic realm and with specific reference to the world of global finance and banking as, If we take specific, very vertical components of the financial services value chain, we may already find bits and pieces of it where AI is dominating, see, for ex the recognition of false positives in compliance /Anti Money Laundering, or the asset picking for certain passive funds in asset management.

What is now already available and reaching boiling point (or "narrow singularity moments" in the parlance of their limited universe of reference) is what we call today narrow intelligence, mostly driven by deep, ML techniques (a sub-set of AI). It is happening in payments and trading, in insurance and lending, just as it is in the automotive industry (with driverless cars), healthcare (with AI-powered diagnosis) or travelling (with digital assistants optimizing travel plans and making the bookings) and in many other fields.

However, what we are still missing is the AGI that leading to the singularity moment, a defining moment marked by an intelligence explosion – one feared even by Elon Musk (he has invested in a not-for-profit organization that aims to "control" that explosion to the benefit of humanity but, on the side, he has also invested in for-profits in a similar vein: one planning to build a human colony on Mars, just in case things get bad; and a second to create brain-connected devices that would allow us also to become smart cyborgs – retaining our own consciousness, but also being able to fight against uncontrolled AGI-powered robots with the help of

quantum computing devices (a science fiction remake of the far-west The Quick and the Dead?).

To get from today's forms of "narrow artificial intelligence" (ANI) to an AGI-powered world where machines could actually take over almost any "brainy" task (and potentially humanity itself) would require few critical steps. One would involve the creation of human-level AGI – a machine able to perform better than the best human mind across the entire realm of thinking and decision-making – from playing chess or Go, to developing better credit scoring to underwrite loans, to designing and executing trading strategies to beat the markets.

This would enable us to enjoy more free time or become redundant – or both, with ultra-rich individuals (digitally enhanced "human gods") owning and controlling the machines that run the world; and with billions of jobless people with no pay and no hope (the "basic humans"), crowded out by AGI and robots but with no access to their services. More importantly, by continuously learning and reprogramming itself this AGI would very quickly become able to create a super-intelligence potentially able to take over the globe – ultimately it would keep recursively redesigning itself towards an intelligence explosion limited only by the laws of physics.

11. Evolution – revolution

This may sound like science fiction (and indeed there may be an element of it involved), but the chances of reaching an AGI and then an intelligence explosion are fair and not relegated to a "too distant to be relevant" future. In fact, once an AGI has been reached – as this keeps growing at a rate proportional to the available computing power – the journey towards an intelligence explosion (the "ultimate" singularity moment) could take just a few days, if not hours – from a software perspective. Or months or years if new hardware is required, with a median "guesstimate" of thought leaders from different scientific disciplines pointing to a date of 2040.[14]

This "fast forward" process of AI evolution could have some obvious impacts, and others unforeseeable, on world economic and financial systems and on human societies, leading to runaway business transformation,

[14] Singularity Stuart, Summit Armstrong, 2012.

driven by self-upgradable artificially intelligent agents running AGI. It could then enter a chain reaction of self-improvement cycles, with new and more intelligent AGI generations coming to light faster and faster.

This would then lead to an intelligence "evolution – revolution", towards a super-intelligent force, that would, quantitatively (as has already happened with PCs a long time ago) and qualitatively (as is happening more and more for narrow fields of thinking and decision-making) far surpass human intelligence – not just deductively, but also inductively and emotionally, and spurred by human-like curiosity (and serendipity).

This would become the "essential singularity", as suggested by John von Neumann, in the history of our race, and beyond which human affairs could not continue, as the technological evolution and related AGI would take control of all systems and applications and robots would be able to perform the most critical tasks now carried out by humans, including warfare. The AGI would then potentially conquer and rule the world, as we have achieved with domestic animals in our daily life.

This evolution – revolution would be very quick and hence superior to the Darwinian one, which is led by genetic selection that takes ages to develop (as it needs to go through multiple generations to execute selection through the disappearance of the weakest). A simple, biological drive has also been defined as "evolution 1.0" by Max Tegmark,[15] and is characterized as one that is unable to redesign either the "software" or "hardware" component during a lifetime – as DNA drives evolution over many generations of people.

Contrary to this, an evolution 2.0 (driven by culture as opposed to genes), can redesign much of the software component, as humans can learn complex and new skills (anything from languages, sciences, professions) on the basis of the knowledge already developed by their predecessors and transmitted not via genes but through books, files and a wealth of information and intelligence that is coded to survive and develop beyond the constraints of single individuals or of a given population. We start out living already standing on the shoulders of the giants of the past, and we can begin building and updating our worldview and goals on the basis of history.

Our brains' synapsis connections link neurons which can store about 100,000 times more information than our DNA, so even from a simple

---

[15] Tegmark, 2017, op. cit.

storage perspective we are getting more powerful and faster in this 2.0 version of evolution. Synapses store all our knowledge and skills in up to 100 terabytes' worth of information (with DNA only managing about a gigabyte) and we can fill them up in a few decades, if not years, without having to wait for the next generation.

This "life 2.0" (as defined by Tegmark) allows smarter and more flexible reprogramming if the ecosystem changes. A quicker reaction is ensured if anything bad happens to our environment and – apart from our own memory – books and computers can extend "on demand" our storage and computational capabilities.

Glacially slow biological evolution now becomes irrelevant, if not pushed by biotech and genetic manipulation also leading to quicker breakthrough and leaps. This can, in turn, bring humanity to an "evolution 3.0" (or "life 3.0" in the definition of Tegmark), marked by technological breakthroughs that can facilitate the redesign not only of our software but also of our hardware component, and in a quicker timeframe. Even if this does not yet actually or fully exist, this final upgrade (in the sense of allowing us to become cyborgs – think of people already living with pacemaker or with arms and legs actioned by hi-tech exoskeletons) would potentially enable us to decouple from evolutionary cycles and become true master of our destiny – or slave of uncontrolled AGI-powered machines.

12. Intelligence explosion

To get to the potential explosion of intelligence in our system, both the software and hardware components need to be present. The achievement of AGI stands as the first key step on the software side. This is very different from the "narrow intelligence" (the ability to accomplish a set of goals, such as playing chess or driving a car, that we usually target with deep learning/ML techniques), as AGI could virtually allow the accomplishing of any complex goal, including self-learning. It could then lead (in a few hours) to an intelligence "explosion" far beyond the human level – the singularity moment – with recursive self-improvements rapidly leading to super-intelligence, which would then require robots as the hardware component.

For all the talk of robotization now taking shape, super-intelligence is the truly critical thing, as it could enable the control of humans, machines

and cyborgs (a hybrid combination of the latter two). But to reach the super-intelligence stage, the intelligence explosion would still need hardware evolution (as discussed later with regards to quantum computing). On this basis, robots are relevant and worrying, as they stand as the visual manifestation of this control of machines. But it is super-intelligence that poses the real threat – from jobs to warfare – in the sense of the way some other entity could manipulate and control our political systems and societies.

Let us go back then to the definition of intelligence, and to one that can be defined as human as opposed to artificial. Intelligence, in its most basic definition, can be defined as the capacity for deductive reasoning, logic, understanding, problem solving and – basically – the ability to put causes and effects into reasonable, and hence predictable, relations, also including inductive reasoning. But intelligence can also be thought of as including our capacity to read and understand feelings, emotions, passions. Even more importantly, it can include our self-awareness and consciousness. More pragmatically, according to a widely accepted definition in business parlance, intelligence can be thought of as the ability to accomplish complex goals – the key driver that companies increasingly use to hire and promote people at all levels.

It follows that intelligence cannot be measured by a single IQ (intelligence quotient) but by an ability spectrum working across multiple goals. Today it tends to be very "narrow", with humans remarkably broader than any existing AI application. However, AI – if and when it gets broader on reaching AGI status – will also become able to accomplish any kind of goal (good or bad), including self-learning. Similarly, as happened to computing (as per Alan Turing's hypothesis – if a PC can perform a certain bare minimum set of operations, then given enough time and memory, it can be programmed to do anything that any other PC can do), given enough time and resources AGI could make itself able to accomplish any goal.

Machines are already remarkably better than us at doing computations – a transformation of one memory state into another, or a process where inputs (information) are transformed via a mathematical "function", generating a deterministic output, or result, that is path independent and substrate (or machine) independent. This is what we have referred to as mostly powered by deductive reasoning; it is highly effective, but still far from the maximum potential. In fact, if machines can implement and run complex intelligence functions (which include inductive and emotional intelligence) they could get to an intelligence able to accomplish much more

complex goals – maybe even triggering a repeated capability doubling rule, in the spirit of Moore's law.

It is not just the "software" component that matters here but also the "hardware", as today's computers are able continuously to gain additional speed by parallel processing, splitting the job into different parts of the hardware. The ultimate parallel computer, so critical to eventually allowing our "intelligence explosion" is then a quantum one that is able to share information with a huge number of versions of itself.[16]

This quantum computing is now coming of age, with a big boost from business – as new scientific milestones now announced by two of the biggest technology companies in the world can testify.[17] A singularity moment could also happen soon in computing, as a critical junction in time will be represented by the "crossover" when quantum machines will be able to out-compete the most powerful traditional PC. This is a moment known as "quantum supremacy" – a date that Google had already set for 2017, with its outcome still unknown and with limited testing and reviewing to allow for a final judgment on it.

13. Quantum leap

Quantum machines could, in theory, harness the properties of quantum mechanics to greatly speed up computation, where bits are substituted by qubits. Qubits, unlike bits, can represent a one or zero state at the same time, benefiting from a property of quantum mechanics that enables them to do so. The super-position of these multi-state qubits can then be enhanced by "entanglement", a process through which pairs of qubits can act in unison, allowing many parallel computations at once. This, in turn, could increase the power of computation in an incremental way and with reference to its velocity and quality.

This potential "quantum leap" in computing could then help expand AI across all its holistic dimensions, including the inductive and emotional ones, increasing the power of machines to think "outside the box". Just think of the most prominent example of this: the ground-breaking win by

---

[16] David Deutsche, *The Fabric of Reality,* Penguin, 1997.

[17] "Microsoft and Google prepare for big leaps", *Financial Times*, 10 February 2018.

DeepMind (owned by Alphabet Inc.) at the strategy game Go, leveraging the intuition, ingenuity and outside-the-box thinking of the machine.

Given the limited range of possibilities people can perceive and analyse, humans have been forced to think outside the box through history and have become very good at being innovative in a non-deductive way, giving us – this was the general understanding and almost a paradigm – an edge on machines. Machines, however, can assess randomly generated possibilities at much higher rate and in a non-deterministic way – much as quantum computing does, working with information across multiple dimensions. Should an ML algorithm rely on such technology, it would then become even more powerful and "innovative" – addressing the inductive dimension and encompassing the emotional one as well – as people's whims determine the development of trends in financial markets, with their rational and irrational behaviour (usually not classifiable as "0" or "1" but identified by a mix of emotions that quantum mechanics could represent better).

As this could, in the middle-term, come of age in the financial services, a number of banks, including JPMorgan Chase, have begun committing some of their research effort to quantum computing. This research could try and tackle a few critical problems of the industry, including modelling the behaviour of people and their emotional reactions in financial decisions and interplays. Quantum computing could also be used to develop risk management systems that are better at modelling a bank's financial exposure and calculating potential losses in a non-deterministic world, and at cryptography – developing un-hackable systems and replacing today's encryption techniques, as we shall discuss later in conjunction with blockchain. More broadly, they could better support the development and work of AI/ML algorithms, addressing financial problems that can be mapped in a similar way to what can be done in physics, chemistry or molecular biology.[18] Quantum computing could even help improve operational efficiency in complicated areas such as clearing and reconciliation, finding more outside-the-box ways to optimize apparently messy systems.

Google, IBM and Microsoft are working on quantum computing in large-scale projects and with huge potential implications for AI. They are already testing calculating machines that run on quantum bits – or qu-

---

[18] "Quantum computing", *Financial Times*, 3 September 2018.

bits[19] – able to solve very complex mathematical problems, such as factorization of large numbers that are hard or impossible to crack for traditional machines. In fact, through the peculiarities of quantum mechanics, qubits are able to exist in many quantum states simultaneously. The more this happens, the more powerful and mind-boggling are the calculations that the machine will be able to do – also for critical uses such as encryption of cyber-data and information.

Still, apart from the number of qubits, which address the "quantity" component of the supremacy game, some quality issues remain with this technology as new applications and tests need to focus on machines' error rates and "tail risk" – that is, the worst error they can make, as quantum computing is damaged by any interference from the external environment and there are no error correction techniques yet available and tested, because of the peculiar multi-state behaviour of qubits and their trends in time.

Quantum computing, apart from augmenting the hardware leg of the potential evolution of Life 3.0, also increases cyber-security risk,[20] as such a technology, once fully developed, could easily outmanoeuvre any "traditional" security barrier and break current cryptography standards – something reachable by 2026, according to experts.[21] On the other side, such technology could also be used to develop new cryptographic apps that could be almost impossible to break, because of the multi-states component. It is said that the Chinese government is already testing such approaches, apparently with some success.[22]

Because of the way they are run, such technologies will also probably require more decentralization, necessitating a change in the paradigm of the ways the system works. This trend to decentralization is also evident in the "shared general ledger" technology that powers cryptocurrencies,[23] allowing encrypted data to flow and be recorded on everything from money and financial instruments, to medical records and real estate. These encrypted records can be shared, protecting data from malware and updating all parties concerned, far from being a means to allow fraud (or an index

---

[19] "Quality over quantity", *The Economist*, 24 February 2018.
[20] Scott Totzke, "Quantum rising", CSO online, 23 May 2017.
[21] "Quantum computing is the next big security risk", 7 July 2017.
[22] Scardovi (2017), op. cit.
[23] "Davos: blockchain can no longer be ignored", *Financial Times*, 25 January 2018.

of money-laundering as notoriously mentioned by JPMorgan CEO Jamie Dimon). The development of quantum computing should in turn provide a boost to the way these general ledgers and their supporting blockchain technology work – given their highly complex and laborious (and energy-guzzling) approach to the encryption of data and information. Because of the synapsis happening across the system, such decentralization also poses a governance issue, as general ledgers share a single database across multiple parties, with no controlling entity in between.

That means that stock markets, clearing houses and even central banks could be disintermediated, and big business (such as the global foreign exchange, where banks like JPMorgan or City make a lot of money) could also be taken away. Even censorship and capital controls could be resisted, as a money transfer made using this technology cannot be stopped, once encrypted as a recorded and shared file by the system. Blockchain is then the next quantum leap in line for the global financial system, which is just going to be powered by qubits and by AI – potentially reaching a further singularity moment worth more consideration.