

01000101101100101011010101001 000101101100101011010101001
0100011011001011010101001 01000101101100101011010101001
101011010101001 11010101001 1010010101001 000011010101001
110110101 1100100111001011001 001010101 0101010101010101
101011010101001 11010101001 1010010101001 000011010101001
101011010101001 11010101001 1010010101001 000011010101001
100011011001010110101011101 0110110010101101001 0011101
0100011011001011010101001 01000101101100101011010101001
01000101101100101011010101001 000101101100101011010101001
0100011011001011010101001 01000101101100101011010101001
101011010101001 11010101001 1010010101001 000011010101001
110110101 1100100111001011001 001010101 0101010101010101
101011010101001 11010101001 1010010101001 000011010101001
101011010101001 11010101001 1010010101001 000011010101001
100011011001010110101011101 0110110010101101001 0011101
0100011011001011010101001 01000101101100101011010101001

The Theory of Post-Economics

Part IV – Being Alive is our Job

Andy Mulcahy

December 2019

<http://posteconomic.org>

As we saw in the previous part, becoming truly digital and achieving some of the ideals that this technology makes possible – such as omnichannel operations – remains aspirational for many businesses and it may take several years before anything genuinely representative of it (in a siloed, restricted personalisation sense, at least) becomes common.

Omnichannel is a major factor affecting how businesses are structured, but it is not the only transformative concept with the potential to completely rewrite the rules for how they operate – for we are already in the midst of a transition. Or, to use its proper expression, *The Transition*.

Not my term, but that of one of the largest and most powerful forces in tech (or business for that matter; the boundaries between the two seem so be fairly fluid) – Google.

As already mentioned in this text, the introduction of the smartphone was a major development that made it possible to interact with the digital world in so many more locations, situations and contexts than was previously the case. This exerted a huge influence on behaviour – both from the perspective of web users, who suddenly discovered that they could get an ‘app for anything’ so they could manage their lives wherever and whenever they wanted through this device, but also customer-facing businesses, who were forced to adapt quickly to the new ways in which their customers were engaging with information and services (which some have been able to do, others struggled with). This was also where the concept of omnichannel had its roots because it fragmented us as customers and made us stream data everywhere we go on everything we do.

The Apple iPhone was launched in 2007, although it wasn’t the first ‘smart’ phone as such (by definition, a smartphone basically means one that is connected to the web – accessing the data network in addition to the cellular one; ‘feature’ phones are also web-connected but with limited functionality, smartphones are the ones that enable users to download all manner of multifunctional apps). It was the model that achieved mass market appeal, however. Google have referred to the subsequent decade (2007 – 2017) as the ‘*decade of mobile*’ⁱ, which is to say the one where smartphones transformed the way that web engagements worked. The next decade, 2017 – 2027, is being earmarked by them as the ‘*decade of assistance*’ⁱⁱ.

What this, and by extension *The Transition*, specifically refer to is the move from search-based interactions (as in typing terms into a box and pressing a button) to voice-activated search (the jargon for this is ‘browserless commerce’, signalling a shift from doing things primarily through screens to just speaking out loud to control and access things, greatly extending ‘where’ the internet actually resides in relation to people). That might seem innocuous enough, but this is a significant evolution for two main reasons. Firstly, people don’t formulate queries in the same way when speaking as they perhaps would when typing, which means the ways in which web pages are structured needs to evolve to cater for it. So while, for instance, marketers have been taught to structure content around long-tail search engine keywords (jargon for ‘longer than one word’) that are fairly macro-level (*‘impact of Brexit on farming’*) to be relevant to queries in a generic sense, voice can require results to be far more context- and individual-specific.

This may be to do with elements such as location where the exact place they mean is not explicitly stated and has to be inferred (*'what pubs are on the high street'* – which high street?) or behaviour, again where all the required information is not necessarily forthcoming (*'have I already seen episode five of Game of Thrones?'* – would that be episode five of series one or series six?).

Which brings us to the second reason. The way that search is evolving means that it is no longer just about bringing up endless potential matches for the specific query for users to sift through – whittling down is far more necessary for voice search results. By way of example, typing *'what pubs are on the high street'* into Google produces 12,300,000 results, with the rather non-committal Punch Taverns website as the top resultⁱⁱⁱ. The kind of searches executed via voice necessitate far greater relevancy by taking other factors into consideration – and the result has to be more contextually-useful than the Punch Taverns website; voice assistants tend to give a single answer, the best one, not the top 10 (or, indeed, 12,300,000) on Google. Working out what the best answer should be in any given situation is where AI comes in of course – and how 'best' is worked out is a moot point. Predictably, as economics is still central to how businesses operate, the tech companies that offer voice assistants are strongly rumoured to be doing deals with brand-owners to promote their products as the default answers. So if, for example, someone said *'order me some shampoo'* to their voice assistant, it would recommend shampoo from a company where there is an underlying commercial deal in place. It is not personalisation, but an entirely generic approach to matching up individual to product.

Where no such deal is in place, achieving genuine personalised, context-specific success in this area is pretty far from easy. If it were simply a matter of looking up possible matches for a query in a database quickly, this is something Google has been doing on a daily basis for a long time already anyway (incidentally, it took 0.71 seconds to produce the 12.3 million matches in that recent example). Google also has a 'Knowledge Base' where its algorithms have access to a reported 70 billion facts to work from to help them make decisions in response to human queries. Having facts available in an encyclopaedic fashion is useful to an extent, but the challenge is understanding intent, implication, irony, sarcasm and the various intricacies, idioms, dialects, sociolects and idiolects inherent to language, in addition to the contextually-determining factors affecting what someone really means when they ask something. Consider, for example, the difference between these two requests for the same information – *'who is playing in the Premier League at 17:30'*, which provides plenty of specific information to inform decision-making, vs *'what's the evening kick-off?'*, a far more colloquial form that a human would understand but a system based on logic might encounter challenges in being able to strip away irrelevant possible answers. Which sport? Which league? Which country? Which day? What does 'what's' signify in a linguistic sense? Do they mean kick-off as in the start of a game, the start of a party or the start of a fight? It has to be able to arrive at the desired answer without asking for clarification several times.

What would seem to be an impossible task is made possible by AI's ability to learn from every engagement and use that information centrally to hone its understanding of future requests. Listening to human voices and conversation is the perfect data source for this. And one of the useful attributes of language is that it is not specific to a single area, as other datasets might be (for example, what products a user clicks on during a site visit – revealing to an extent, but just specific to a shopping context). Language, technically speaking, can cover any topic, discipline or theme, which voice devices can continually

digest and learn from. The range of information it can take in is without logical borders or limits.

It is this level of access to such a vast store of information as human language, spoken by billions of people the world over, that is really going to boost the capacity for AI to understand developments in a human sense, as opposed to the traditional logical way that we associate with machines – bits and bytes, ones and zeroes, *'if this variable equals this, perform this action'*; premeditated, pre-programmed, with a set range of predictable outcomes – if a machine doesn't need to make any independent decisions (or use its 'intelligence' indeed) it is not bonafide AI, but merely automation. The potential exponential increase in the decision-making capability of AI systems is directly connected to the volume and variety of data it is able to base those decisions upon.

All the major tech firms have a voice-activated system, that work by listening for a verbal cue (such as *'ok Google'*) to start processing spoken information and use it to surface information that is as relevant as possible. This type of 'listening' technology is actually built into all manner of devices, such as smartphones and TVs in addition to the devices dedicated to that function, such as Amazon's Echo. All the manufacturers claim these devices only start 'listening' when you give them the verbal cue, but that would seem like such a huge missed opportunity for research that it's difficult to believe is universally the case. So there is already a veritable network of listening devices that are in our homes and pockets and the Echo is proving to be a big seller, probably due to its low price point – the main device retails at around £149.99, but the smaller version Echo Dot is available at just £49. It is in Amazon's interests to get as many of these in to peoples' houses as possible, so they are not using this as an opportunity to rake in huge profits from the device sales themselves, as the data they collate (and access to people they get) will be worth far more over the longer period.

And it's working, if the estimates are anything to go by. Amazon are famously cagey about their sales performance; in fact the only reason anyone knows how much they turnover in the UK is because they were forced to reveal it during a defence of their tax arrangements to a UK government committee. So they will not come out with an official figure on how many Echo devices they have sold, but Morgan Stanley research published in January 2017 estimated the total to be 11 million globally between mid-2015 and 1 December 2016 (note – that period does not, therefore, include the Christmas shopping period of December)^{iv}. On Amazon Prime Day in July 2017, an event the company launched in 2015 with heavy discounts only available to their Prime customers, they slashed the price of the Echo range and marketed them prominently. Again they don't share exact sales figures, but it's a reasonable bet they would have shifted in pretty sizeable volumes. The figure they did release was that the overall year-on-year sales growth (compared with the same 30-hour period in 2016) was up 60%^v – whatever that pushes the volumes to overall is anyone's guess. Then, during the Black Friday period in the UK in 2017, the price of both Echo Dot and Google Home was slashed from £49 to £35 on several major retailers; at such an accessible price-point as that, it seems a lost opportunity to not buy one, just out of curiosity if nothing else. They are highly likely to become ubiquitous very quickly.

The aim of all the major tech companies is to have their device as the one of choice in everyone's home. As we can see with the Echo, there are various ways in which they can capture information from users through embedding their respective AI voice-recognition programmes – Amazon Alexa, Apple Siri, Microsoft Cortana, Facebook M and Google

Assistant – within certain types of hardware (smartphones, tablets, laptops, TVs or dedicated voice devices). The trick is for each to work out ways for their service to become indispensable by offering access to as wide a range of services as possible. So while on the surface these businesses seem straightforward enough to define – Amazon is a retailer, Google a search engine, Apple a device hardware manufacturer, Microsoft a software developer, Facebook a social network – in reality that barely even scratches the surface of what they cover.

Many of them are also highly dominant in all manner of other areas – email and messaging (Google gmail, Facebook Messenger), social networks (Google Plus), payment provider (Google Wallet, Android Pay, Apple Pay, Amazon Pay), music streaming service (Prime Music, Google Play), film-streaming service (Amazon Instant), web hosting provider (Amazon Web Services is by some distance the market leader), advertising platform (Facebook and Google account for a huge share of the overall digital ad spend – 85% of every new dollar spent in online advertising was forecast to go to Google or Facebook in Q1 2016, according to one analyst^{vi}), car manufacturing – that’s not even an exhaustive list.

Taken as a whole, they pretty much penetrate every area of our virtual interactions with the external world, plus they are usually very successful in gaining market share due to the huge resource they can invest in it. It’s important for them to delve into so many areas to gain additional levels of insight across the full scope of people’s (more specifically, individuals’) activities, as the business with the greatest store of understanding of what people do and want in the digital age will possess an advantage over its competitors – both in terms of personalising content and targeting them, and how intelligent their artificial intelligence can become. The overall winner could eventually create an unassailable lead over all others, indeed.

Voice assistants offer a central opportunity for achieving this. Speaking to Wired magazine about Google Assistant in October 2016, Fernando Pereira, VP and engineering fellow at Google said:

“Launching the Assistant is very much like Google launching search back at the beginning of the company. Search was a great thing then, but compared to what it is today, there’s so much more understanding. We’re going to see that with the Assistant of 10 years from now compared with the Assistant of today. It’s going to be way more fluent, more able to help you do what you want to, understand more of the context of the conversation, be more able to bring information from different sources together.”^{vii} (author’s emphasis)

This is a really fundamental point for AI generally, but particularly where voice activation is concerned, for the reasons discussed in this chapter but also because it will often require access to disparate, unpublic information if it is to be genuinely useful. For example, it could only answer fairly rudimentary questions such as ‘do I have any credit on my phone’ if it had access to that user’s mobile account or give an accurate response to ‘am I doing anything on 9 September?’ if it can access their diary, potentially multiple diaries – both personal and business diaries perhaps (anyone who has tried to schedule a meeting between internal and external business attendees will understand how frustrating it can be to know nothing about the external attendees’ availability).

This is where we can start to see how a ‘walled-off’ web is just not a practical option for supporting AI systems. This does not just specifically refer to the kind of digital

protectionist moves made by some governments, most notably China's 'Great Firewall' which censors certain sites and prevents its citizens from being able to access them (this includes Google, so Baidu is the dominant search engine there, and news agencies such as the Economist are non-permissible). Instead, the idea of a walled-off internet (or 'splinternet', as it is sometimes labelled) finds expression through the actions taken by major tech companies to force users to remain within the applications that they can track and control. Google probably provides the most prominent recent example of this, with the EU fining it €2.4bn^{viii} for alleged abuse of its market position by promoting Google Shopping entries above other results in its 'shopping' feed. There are many other examples (it's the reason the tech companies diversify the number of areas in which they operate, as the competition to understand every element of user activity is the top goal for any tech firm), some of which even have something of a whiff of irony about them. For instance, Amazon has filed for an anti-showrooming patent called 'Physical Store Online Shopping Control'^{ix}. Showrooming is a retail concept whereby a shopper visits a store and looks at a product, then goes online, most likely using their smartphone, to check prices at alternative retailers to see if they are getting the best overall price or deal. They may then buy that product from a competitor, even though they used the original retailer's expensive store estate to help them do it. Amazon, who have just bought Whole Foods and therefore now technically have a store portfolio, have filed the patent for some technology that enables it to intercept network requests that emanate from within its in-store wifi network so they can be redirected – perhaps to a site that it owns or to offer a coupon to help secure the sale, hence people won't be able to access the web in its entirety, only Amazon's estate. According to the patent:

“A negative scenario may exist for a physical store retailer when a consumer evaluates items at the physical store, leverages physical store sales representatives, and then reviews pricing information online in order to purchase the same item from an online retailer. The physical store retailer pays for floor space, sales representative time, product inventory management, and other costs while not being able to complete a sales transaction.”

Poor old Amazon, having to put up with the online retail wolf entering its physical sheep den. As I say, just a faint whiff or irony there perhaps.

These examples relate to either redirecting users or surfacing content that isn't necessarily the best match for a user's requirements, but instead represent a more suitable match for the tech giant. The problem with an environment defined by connectivity, but in which economics is maintained as the governing system, is that 'free' to access services have tended to primarily generate high revenues for the platforms rather than the people creating content on them (unless you can boast millions of followers, which is only a small percentage of those creating the content). So, rather predictably, there are numerous and various attempts underway to divide it up into pay-to-access areas. Some are straight-up 'you must pay a subscription to view this content' such as the paywalls used by *the Financial Times* and *the Times*. Others are more voluntary (for now), such as *the Guardian's* 'Supporter' network, but joining provides access to additional content and features.

Sites such as Patreon allow what it calls 'creators' to run the equivalent of membership subscription to their services and content, which equates to paid access. Sport is already heavily influenced by this model, as is the film and TV industry, with Disney the latest to

announce an exclusive streaming service for its Star Wars and Marvel films. Even those who previously have been regarded as the archetypal pirates of online content have turned to this model, with none other than Kim Dotcom himself introducing a micro-payments service intended to let people charge small amounts to access their content^x. The would-be anti-establishment revolutionaries of the digital age couldn't beat economics it would seem, so they've joined it.

What all this leads to is a kind of 'web behind closed doors', with areas of it partitioned off and only accessible by those willing to pay and, of course, the data created only useable by the host platforms. The problems inherent to a 'walled-off' web soon become apparent when you think about them. For instance, a user asking Google Assistant to check whether they have any emails from a specific contact in their Hotmail / Outlook account, which is hosted on Microsoft. It's easy to envisage a situation where the competing businesses limit the capacity for rival systems to access data streams for competitive purposes – they wouldn't want their competing services to be that convenient. These companies are so dominant in some of their respective areas that many people will use them in combination – accessing Amazon through a Google search on an Apple device, for example. So while it would be counterintuitive for them to block each other outright (Google to stop indexing Amazon – and, to give them their due, Microsoft and Amazon have publicly announced plans for their respective voice assistants to be able to 'speak' to each other^{xi}), that's not the same as being willing to share information on individual users' activities freely between themselves. There may be a certain level of interoperability they need to introduce for functional purposes, but data-sharing would be tantamount to losing competitive advantage.

And this is why *The Transition* may sound transformative in principle, but ultimately there are too many restrictions, put in place by monopolistic tech companies wary of ceding that competitive advantage to each other, for any one AI system to attain a level of accuracy and relevance appropriate to its promise. It doesn't seem unreasonable to imagine a situation in which the top ten largest tech companies have access to 80% or more of the data available, though not in a shared way between them, while the other 20% is collated by a huge long-tail of much smaller organisations. In that situation, the capacity for AI would be heavily restricted by the owning business' share of data. As Google put it earlier in this chapter, being more able to *bring information from different sources together* is exactly what makes AI function more effectively, or at least it raises the potential for getting information out that is more accurate and relevant to the intention behind a query.

Take an example from a Google research project in this area with Carnegie Mellon University^{xii}. The abstract to the experiment reads as follows:

“The success of deep learning in vision can be attributed to: (a) models with high capacity; (b) increased computational power; and (c) availability of large-scale labeled data. Since 2012, there have been significant advances in representation capabilities of the models and computational capabilities of graphical processing units (GPUs). But the size of the biggest dataset has surprisingly remained constant. What will happen if we increase the dataset size by 10× or 100×?”

In other, less convoluted words, what if they ran an experiment around vision tasks (recognising things in images such as an object in a photograph) and focused on

increasing the volume of data that algorithms have access to rather than making tweaks to the algorithms themselves or relying on use of a faster GPU. For this, they used a sample database of 300 million labelled images, whereas many experiments in this area tend to use a standard set of 1 million images. The result was that the object detection score was 3 percentage points higher when using 300 million images rather than 1 million. That might not seem like a huge uplift, but that could have a significant impact on the rate of accuracy and relevancy systems can achieve. Adding 3% to an already huge number quickly opens up all manner of new possibilities. It serves to quantify an idea that would appear to be self-evident – the greater the variety and volume of data available to base decisions upon, the greater the accuracy of the results.

This is, of course, not the same as just assuming that combining entirely random and unrelated datasets willy-nilly would achieve a similar outcome, but the principle of using more to get more out seems sound enough.

It would also seem a safe bet that continual and sustained investment in AI by the major tech companies is a given; they all state as much in the public domain and it has been established as probably the primary area of competition over the coming period. This represents a challenge for any new upstarts who may be aiming to compete in this area if AI is the battleground, as a highly sophisticated AI system isn't something that can be put together in someone's garden shed – it requires the kind of long-term, high-expense resource and investment only available to a handful of companies who possess that level of capability. In reality, *The Transition* is being driven by five or six companies for whom the digital world is becoming a pie to be divided up and, possibly, walled up.

Given that huge concentration of power for change into so few hands, we might perhaps be tempted to ask a few fundamental questions. Such as what the overall purpose of *The Transition* actually is – which does not just mean a shift from text-based to voice-based search but a significant acceleration in the contextual understanding of individuals on an individual basis. Or, if we are indeed entering the decade of assistance, then what is it we are to be assisted with exactly? What would be considered a successful outcome?

Let's assume that, during the decade of assistance and all the functional transformation that comes with it, there is no will to address whether the values and elements of capitalism are still relevant – that these are a given and everything that happens need to be structured around them. What happens when you overlay economics onto a digital environment that is hugely dominated by a small number of large companies?

Notwithstanding the point made earlier about the major tech companies being hard to pigeonhole in terms of what they actually are these days, most people would probably still regard them for what they are (or were) primarily – Google a search engine, Facebook a social network etc. Yet this idea of these being buzzy, wacky, fun, disruptive tools and networks seems to belong to a former era. The cold hard reality, that has become more and more apparent as time has gone by, is that these are businesses, first and foremost, operating in an economic environment and within an economic legal framework who, irrespective of all the public arena talk of wanting to be egalitarian, improve people's lives and make the world a better place, have to play by the universal – and highly competitive – rules of capitalism. Which is to say, they have to find ever-increasing opportunities to grow revenue and drive those to hit the targets the market expects or, hopefully, exceed them even. Spooking the markets by falling short of expectation can be hugely damaging

to the valuation of their business and confidence of investors, which would limit their capacity to engage in all the fun wacky stuff, so it always has to be the number one focus.

Who could forget, indeed, the defence that the somewhat dumbfounded Google co-founder Eric Schmidt gave to the furore over their tax arrangements in Europe:

“It’s called capitalism. We are proudly capitalistic. I’m not confused about this.”^{xiii}

The fun and wacky bubble well and truly burst and no mistake. But this does, of course, bring us back once again to the very problematic digital concept of things being ‘free’. Both Google and Facebook have over a billion users (or, in traditional speak, ‘customers’), none of whom (if we exclude advertising businesses and the like) pay them a penny, overtly at least. As discussed in the introduction to this text, this is not a viable scenario for an economic system so alternative ways to extract value from those relationships have to be found. Different models have been tried, such as the pay-to-access ones used by the likes of *the Financial Times*, the ad-free model of Spotify, or the ‘gated’ approach whereby users can access a small number of articles on a site per month for free but have to pay after a threshold is reached. Realistically, these can often only work for certain propositions – ie those with strong quality content that people can’t really get elsewhere for free – otherwise we approach a model under which every site would have some kind of front-end financial tariff attached to it; the web would become genuinely walled off, to the point where users would just choose a small number of sites to use and never stray any further. This isn’t particularly practical or likely to work in many cases, so the model that has been settled on so far in a wider digital sense has been free-to-access, underpinned by advertising (and furthermore, in some cases, by bitcoin mining – where some sites detect if the visitor is blocking ads and triggers a pop-up informing the visitor that, in order to view the site content without ads enabled, they will have to agree to allow the business to use part of their device’s processing power to mine bitcoins, which is highly intensive on CPUs. Salon.com is one example of this).

An inevitable consequence of that model being in operation on such core digital platforms as Google is a glossing-over of personalisation and relevance. When you execute a search on Google you usually get numerous pages of results, many of which will be ‘natural’ matches selected by the algorithm based upon the search term(s) you used. For some queries, you may also see some results at the top that have a little box beside them saying ‘Ad’. While the top non-ad results are probably being displayed because the creators of the content have worked hard to structure their pages to be regarded among the strongest matches for specific keywords according to Google’s algorithms (so if you search for ‘*blue Adidas trainers sale*’ these results are as close to what the algorithm thinks you are after in this context), ads essentially enable businesses to switch off personalised relevance and position their sites above other, potentially more desired or accurate matches, because they are paying to do so. Money provides the ability to ride roughshod over relevance, in a highly un-digital way (as anyone who has had to endure an ad for a Land Rover on YouTube will know, even where the content they were waiting to view has nothing to do with vehicles).

Pushing ads to people because businesses have paid to have them pushed is nothing to do with assistance – it’s just keeping economics relevant, as advertisers will pay for that kind of exposure to an audience. That’s where the money is online. The whole point of assistance – and AI / machine learning in general – is that it is going to get very good at

understanding people not as segments for marketing to, but as *individuals*. If AI is there to provide assistance, then it has to have an overall purpose toward which it assists each individual, some concept of outcome that corresponds with a good job done well. The digital world, more than anything else, is characterised by the ability to measure – things that could not be tracked in any realistic degree of detail can now be comprehensively understood. Yet each AI system is not some independent entity, but is always wholly owned by a business – may even be best classified as a company asset or even a business in its own right in some cases. As an expensive asset that siphons off resource and investment, but that provides assistance ostensibly for free (in the modern digital sense), AI systems will need to offer some kind of return either overtly (not necessarily as a pay-to-use service, but by constantly encouraging users to spend their time / money in areas that benefit the owning business) or by proxy (by sharing information gleaned on individuals with other, more profitable areas of the owning business to help them perform more efficiently and drive revenue behind the scenes).

Or to speak plainly: AI will assist you to make money for the business that owns it. How could it possibly be designed to evolve to behave any other way?

From an economic perspective, this model is logical. It wouldn't make sense for something to be creating such a high level of value for individuals, through so many different interactions with external entities, without driving any economic growth in the process. But the technology demands a different approach due to the ways in which it alters how these relationships work. We've never had a situation before in which platforms sit in-between our interactions with so much of the world external to us, that can track, know and understand us individually. When people are defined as consumers, we each have a degree of financial resource which we choose to assign to whatever appeals to us; this is when agency largely sits with the individual. With assistance, choice is not eradicated but agency drifts over toward the AI system, which is there to assist individuals in the choices they make by matching them up with relevant, accurate and context-specific selections in a managed way. Consumption is about deciding what you need from the range of available choices, but assistance to a large extent makes those decisions for you. AI is going to have a lot of sway in influencing what any given individual thinks or does; the criteria it uses to assess the success of its assistance is going to push the world in a very specific direction as a consequence.

It is this kind of shift that dictates why it is so necessary to review the way in which everything is defined, rather than just assuming existing values are still relevant and forcing the old ways onto the new. Digital technology on its own is not sufficient to necessitate it – it just connects people and things together – but these changes that may seem subtle enough at first can actually destabilise long-held assumptions about how things work. *The Transition* is not just a shift of text to voice, but a wider advancement of AI capability that will make individuals far more known, in far more ways and on an individual basis – while also developing the nature of web access from being primarily manual, stop-start and restricted to automated, constant and surrounding.

If AI can know so much, what is it going to do with that information? What would the ideal situation be given this development, how should things be structured and defined and what is the overall purpose toward which it operates?

Thus far in this text I have argued that the move toward total connectivity and total personalisation will completely change the way that everything works, to a reasonable extent has already indeed – they, in tandem, represent the ultimate and truest expression of a connected world, with economics merely a hangover from the unconnected age that is being kept relevant due to its deep entanglement in all areas of our lives. We haven't countenanced the possibility that it might be irrelevant now, the thought hasn't, perhaps couldn't, occur in most instances.

Yet if we are to accept that economics should not be considered as a given as the supporting system for the digital age, then we also have to consider which values could replace the current economic-focused ones that underpin our existing system, capitalism – competition, choice, freedom, privacy and opportunity – which were discussed in detail in part II. These values, of course, have to be very cognisant of the fact that in the near future every area of our lives will become increasingly integrated with, and overseen by, AI systems, while in the distant future the very idea that AI could be absent from an area of human activity will seem a nonsense idea.

This chapter puts forward three interrelated concepts that would form a logical set of values to underpin a non-economic system that is not measured in monetary terms, but is instead characterised and defined by data, AI, total connectivity and total personalisation. These are *purpose / efficiency*, *guidance* and *value* and the imperative behind how and why the digital age could be structured around them is interrogated for each one in turn.

Purpose / efficiency

The value of purpose addresses the issue of why a business or other type of organisation would be incentivised to do something when the traditional measurement of value – money – is no longer valid. Why would they go to all the effort of creating products and services anymore?

Under a system in which the overall success of activities is measured using economic metrics – and within which people are both defined, and treated, as consumers – the purpose of all consumer activity is ultimately to feed the economy, which at the individual level means each consumers' purpose is to *consume*, while it is the purpose of every external entity (such as businesses) to stimulate that activity. Businesses are structured around this very idea, that their activities will be assessed against economic criteria and this, in turn, heavily influences the directions in which they develop. The services they provide have to be economically viable for the business to continue as a going concern. These services have to remain relevant and attractive to consumers, so that they choose to direct some of their spending power toward them rather than to a competitor. The relationship of business to consumer is shaped by that; it's common sense that value is attributed in this way.

The economy – the thing we are all told we have to stimulate and grow continually – is an intangible presence, a mass vortex that sucks everything in and swallows it down with the express purpose of growing ever-larger; while for the individual, busy consuming away to fulfil their part of the economic system contract, what they can expect to get back is little more than a lottery. Their wages may go up or down, mortgage rate may increase or decrease, pension annuity deflate or inflate, the price of their common purchases may

edge, rocket or plummet. The economy may be influenced by billions of individuals, but it is not required to feedback individually on what it is doing or why (how it is performing against its purpose), based upon the contribution they have made, nor does it do anything specifically for a given individual. It's a relationship built exclusively on trust that the economy is 'working' – an idea that politicians are often keen to impress and one that, of course, cannot at any given time be incontrovertibly proven one way or another (consider David Cameron's and George Osborne's message during their 2010-15 UK government administration that '*our long-term economic plan is working*' – don't worry about the minutia or how it impacts you in detail, that's all you need to know).

But in a world where everything is trackable as data – and can therefore be measured in relation to each individual, with the information being continually processed and cross-referenced by AI systems acting as the intermediary between individual and business – you need something to measure toward that is non-economic in nature (for all the reasons previously discussed in this text). Purpose is a required element of an entity like AI, as it is so dependent on logic to function. It wouldn't be particularly efficient, for example, to instruct an AI system to just oversee production at a factory without establishing clear criteria around what that factory is there to do. If that instruction was as simple as 'make toys', it would be reasonable to expect the system to create ever-increasing numbers of toys over time, making better use of available resource and optimising processes to reduce time taken per output as it constantly reviews data to identify more efficient ways to make the toys. But this would be a very restricted application of AI if that was all it understood to be its goal, as it would just make them incrementally without attempting to achieve anything else. It may not ever even stop.

If the goal is economic, then the focus and likely outcome changes. So if instead of just 'make toys' the instruction was 'increase revenue from the toys made in this factory', a number of things might happen. Rather than just make more toys incrementally, it might identify a change in design to improve margins, or establish additional factories to create greater volumes of toys, expand into international markets, drive costs out of the business, work out ways to raise the average amount a customer spends. It might even reduce the number of toys it makes, discontinuing certain lines, to protect profits. Its purpose will be purely to improve the bottom line of the business, which means any additional efficiencies that could be found through analysing the available data would be ignored, no matter how transformational they could be from a social, political or individual perspective, unless they specifically contribute to that goal.

As money can no longer be considered a viable option for assessing the success or otherwise of interactions between individual and business, we need to establish some form of alternative purpose or this kind of economically-blinked AI will be the logical form that dominates the direction of business and individual activity the world over. The limitations of what AI could achieve are hard to predict, but it has to place its focus on achieving some kind of outcome in all it does, the *reasoning* behind the decisions it makes – which would, of course, be at the expense of the other kinds of outcome it could pursue. An AI system cannot function efficiently or effectively if it is unable to understand and identify what a desired outcome is from the decisions it makes, what is essential to the realisation of these desired outcomes and what can be sacrificed in support of it.

A major factor in determining the answers to these questions relates to how individuals are defined under a system based on information.

Instead of focusing on mass producing and shifting as many units as possible – which is what equates to success under economic consumerism, where people are defined as consumers – individuals are actively involved in the creation of products and services specific to them by the very nature of our connectedness to them. They become participants in this process without even realising or needing to do anything as such – AI systems handle that side of it on an ongoing basis via the data we all stream through our activities. It raises the prospect of creating the right volumes of things for the right people at the right time to fit the right purpose – or at least the aspiration to get somewhere near it.

Under an economic system, progress is dependent on the establishment of new markets where opportunities pop up and draw investment. The most successful of these – as well as the ones with the greatest potential – are essentially what set our direction. It's the basis on which the infrastructure around us is built, the supporting regulation is drawn up and it further determines the nature of relationships between individual and business. The more everything is geared toward moving in that direction, the greater the inherent, justifiable sense in doing so. At present, it seems logical to assume that AI will simply inherit this approach and speed it up exponentially.

But if it no longer makes sense to just shift as many units as possible – which is all about making as much profit as possible from as many sales, in a generic sense – because it is becoming possible to understand people at the individual level, then what criteria should be used to ascertain what needs to be created and why, for whom and in what volume? And, of course, how would access to things be regulated (or 'paid for') in relation to those individuals?

Things can only be created and managed in an efficient manner if there is an underlying reason driving that activity. Or, a purpose to measure against.

The issue with purpose is that it can seem a lofty concept. Many institutions have had a go at identifying it and putting it into some kind of real-world practice, from organised religions through to political groups who take the ideas of philosophers and interpret them into a system of governance (often, it has to be said, by only focusing on certain areas of a philosopher's writings and glossing over any inconvenient parts). Purpose tends to be applied at a universal and generic level, covering millions or even billions of people concurrently and, some of whom, needless to say, find it easier to comply with the dictates of the guiding systems than others. Yet, being realistic and pragmatic about this, purpose doesn't have to be set based upon thousands of philosophical texts and applied to the entire human race in one fell swoop. You don't need to know the meaning of life to be able to identify a sense of purpose, people typically generate their own sense of purpose based on their own experiences and needs. Purpose might relate to losing weight, climbing a mountain, visiting a faraway place, reading a long and complex novel, learning a foreign language. Purpose doesn't have to be a single dominant idea – any one individual may have multiple purposes which fragment and change as their life progresses and things evolve. To extend the toy factory example further, it's not about making more toys continually, but making the right toys to suit the right purpose; having a specific reason underpinning each decision that is made.

Purpose is not only lofty, but can also only ever be somewhat idealistic when attempts are made to apply it at a universal level as it has to be, by definition, generic and non-specific to any of the individuals across whom it spans.

The UN provides perhaps the best example of this. The opening statement of the UN Charter^{xiv} (first published in 1945) declares:

WE THE PEOPLES OF THE UNITED NATIONS DETERMINED

- *to save succeeding generations from the scourge of war, which twice in our lifetime has brought untold sorrow to mankind, and*
- *to reaffirm faith in fundamental human rights, in the dignity and worth of the human person, in the equal rights of men and women and of nations large and small, and*
- *to establish conditions under which justice and respect for the obligations arising from treaties and other sources of international law can be maintained, and*
- *to promote social progress and better standards of life in larger freedom,*

AND FOR THESE ENDS

- *to practice tolerance and live together in peace with one another as good neighbours, and*
- *to unite our strength to maintain international peace and security, and*
- *to ensure, by the acceptance of principles and the institution of methods, that armed force shall not be used, save in the common interest, and*
- *to employ international machinery for the promotion of the economic and social advancement of all peoples,*

[...]

Laudable aims indeed and difficult to argue against the general humanist principles to which it aspires. We have to remember the context within which it was produced, of course – just after the Second World War, the Holocaust and subsequent formation of the UN, when the appetite for collaboration to avoid future devastating conflicts must have been at a high.

If we are, however, to accept these as being the archetypal existing model of a common purpose, then we would also have to accept that, 72 years later, progress toward it has been at best stunted and has arguably even proved unobtainable.

There are any number of reasons that we might identify for this – it may be that the UN is too democratic for its own good; the issue of some major members being able to veto UN activities; it lacks genuine authority, acting instead as a forum to at least get governments around a table to facilitate talks and reduce political tensions; perhaps, even, it is quietly understood that the UN charter provides a kind of idealised benchmark, a model that is acknowledged as being ultimately unattainable but good to have as a point of reference anyway.

There may be other reasons that could be identified, but the main point upon which it's hard to disagree is that the UN doesn't really achieve the things its charter says it's there to do. These elevated ideals seem to have such little connection to where the world actually goes, yet we don't see huge organised protests around the world demanding that governments adhere to the charter's principles more stringently. Most people (outside of the legal and political profession) probably aren't even aware such a charter exists.

Universal purpose, then, has proven difficult to agree upon (or execute, at least). Yet AI systems do require some top-level guiding principles so that they can assess the success of decisions made and outcomes produced in any given situation. An advantage that AI has over humans is the ability to 'remember' guidance set out in any charter to which it is designed to adhere. AI works not just by endlessly digesting information that it can regurgitate in response to queries (getting the answer to questions that have 'correct' answers, such as '*what is the current capital of France*'), but by assessing the context of a situation and taking all manner of factors into account as part of the decision-making process. Working out what *should* happen, given everything that it understands about what is being asked; where, when, why, by whom as well as what could be considered a successful outcome. To judge the success of such outcomes, a purpose of some kind, potentially multiple interrelated purposes, to measure against is required.

Under the current system, direction is to a large extent dictated by economics – which is too complex to communicate to any individual who isn't directly interested in economics. When it comes to policy, governments are often trying to balance the economy (with policies that have to apply to huge swathes of people concurrently and generically) rather than focusing on the specific needs of individuals – this of course makes no sense in a world characterised by personalisation. Economics just wasn't designed with personalisation in mind. When someone feels they are being wrongfully treated by the system, a common complaint goes along the lines of '*it's one rule for them and another for me*'. Yet if it is personalisation that characterises the digital world, that would actually seem to be a broadly sensible approach.

What the UN tried to do was establish some kind of universal principles to which we could all sign up. It's proven to be a lot more difficult to achieve than it was to write, but it would be churlish to criticise the attempt. Because, when you really come to think about it, having a universal sense of purpose would seem to be a good idea at any time in history. Quite whether there have been times in history when it has been possible or near is a moot point, but the fact is that after several millennia of human civilisation, philosophical theories, religious doctrines and scientific enquiry it still remains a distant and highly ignored aspiration. In spite of all our intelligence and debate, the human race has never managed to define a universal direction in which to travel. It may be an impossible endeavour to get everything aligned in that way, as there are just so many people and people can be very different – not just in different cultures geographically far apart, but even within a single culture or country. Next-door neighbours don't agree on everything, so getting cross-planetary agreement would seem a little ambitious.

Yet the issue of getting universal agreement on a purpose or set of principles is only an operational requirement when information – and the communication of that information – is sparse. Things have to apply at great scale in that context. The move toward total personalisation means that things don't have to be universal (or, generic) anymore. Everything can be adjusted slightly, even in small ways, to be appropriate to each

individual – and it is this shift that makes purpose possible and elevates its importance as a non-economic value for a system based on information.

Total personalisation is so significant because it marks the birth of the genuine individual in society. In history there are many claimants to the start of the true age of the individual – the start of mercantile capitalism, the Enlightenment age, the Industrial Revolution, the post-Second World War period, the 1960s – but the one enabled by this technology is different in that it makes individuality universal. In these previous ages where a greater sense of individuality became possible, it tended to only allow (usually a certain group of) people to express themselves in more individual ways, but this was one-way traffic – having a greater sense of their own individuality is not the same as being recognised by all the surrounding structures and systems as being so.

Total personalisation becomes most effective when every single crumb of data that is collected can be uniquely attributed to a unique individual, with the experience that the person gets adapted accordingly, so that no two people on Earth ever need to receive the same experiences and communication – at least where varying and personalising it may produce better results. This isn't just some people starting to feel a bit more like individuals; total personalisation means that each person can't help but be recognised on that level, by everything around them, in any number of contexts.

Purpose is what dictates the nature of communication that individuals have with external entities. While it is obviously the case that governments and official bodies set the overarching laws and rules that regulate the behaviour of everyone and everything within a territory, a lot of communication with individuals in the modern age actually comes via brands, who (depending on their size) can spend millions developing marketing messages that they then have to deliver en masse to the public. The purpose of that communication is to stimulate interest in the brand, awareness of the issue / area their product or service solves / enhances and to drive engagement with it. But this is not a personalised message that recognises the independent and unique requirements of each specific individual, it is generic in the sense that it solely focuses on communicating the set of values toward which the brand aspires – the purpose is purely commercial and serving the interests of the brand, rather than factoring in the relationship that it has with an individual and how both parties could work in tandem to produce better, more mutually-beneficial outcomes.

The purpose of much brand activity is to drive people in their droves toward sites where they will be exposed to advertising, or be convinced to complete a transaction of some kind outright; the desired outcome is always economic eventually, even if it doesn't always achieve a tangible increase in revenue straight away. Their influence, however, is not restricted to what we might think of as their 'estate' – site, app, dedicated Facebook page – consider the way in which social networks have trending topics where large volumes of people share their thoughts on a development or event (some freak weather, a TV programme, a political announcement). Wherever people have debates like this, which are far more digital in nature than they used to be, brands have an opportunity to associate themselves with it or steer opinion toward a certain sentiment that is consistent with a marketing message that they are trying to push. Discussion through digital cannot be thought of as being free from interference, it is a significant thread of human engagement that is heavily open to subtle, clever and personalised manipulation. This isn't necessarily an issue in itself – to some extent, that is unavoidable – but brands operating within an

economic framework serve the interests of themselves primarily, and some have staggering resources at their disposal in order to do so.

And a growing issue with this kind of personalised, data-led communication is that it's going to become *too* effective. The adverts you see are served up based on what is known about you – the type of content you view, what you search for, your likes and dislikes etc. Obviously, the more that is known – and the more intelligent the algorithms deciding what to show – the more likely the advert is to be relevant and appeal to you. The key variable here is around what is *known*. As everything around us becomes digitally-enabled, far more data will be tracked and there will be far more platforms through which adverts can be served up. In short, advertising is going to become very clever and completely ubiquitous. The consequence of this is a kind of consumerism in overdrive – where people are constantly cross-sold and up-sold with remarkable success (forget 4% conversion rates there) due to the relevance of what is marketed, and constant pressure on companies to achieve higher targets.

This relationship requires a process of balancing, so that it functions in the interests of both sides. As explained previously, individuals are now plugged into the supply chain in ways that enable them to 'advertise' for what they genuinely require or would benefit from, without even needing to manually do anything – both in a short-term and long-term sense – making the entire process (the whole 'transaction' between individual and business) more personalised, integrated and synchronous in nature.

This is the genuine start of the age of the individual because things that happen are being done based on information that is specific to each individual, but it brings a paradox in that it is also the point at which we all become inextricably connected to each other and lose the ability to undertake activities in isolation. The age of the individual is not about individuals existing individually – instead it relates to all the larger structures that govern and support how individuals act and what they are able to do. It is recognition of each individual's place within those structures – or, to be more accurate, their participation within them.

Purpose, even within this context, can still seem a grandiose concept. Establishing a sense of purpose isn't the same as establishing the meaning of life, or the meaning of a life (*'this is the point of your life, we've decided'*), it is instead necessary to have something against which to measure to work out who gets access to what, when and why to improve accuracy of outcomes and efficiency of process. If the common standard, money, that previously served the role of measuring the capacity for an individual to access something (*'it costs this, do you have sufficient money to purchase it'*) is acknowledged as being an obsolete concept, then some kind of purpose is required to fill the void.

In the digital age, the way that everything works is heavily characterised by data and tracking which, as explained in part I, is not the same as money. As it becomes possible to understand what is happening and why, there is always the potential for honing processes in pursuit of evermore efficient outcomes. Efficiency cannot be said to be a statement of purpose in itself, but it does provide a meaningful measure against which to assess the success or otherwise of activity between business and individual. Where greater efficiencies are being found in the engagements between them, the system has

the strongest chance of operating in both of their interests – an individual gets something they want or need, the enabling businesses are able to fulfil it efficiently.

Purpose itself is something that needs to be non-generic under a system of total personalisation, which is to say specific to different groups and, indeed, even on an individual basis. For businesses, non-economic purpose is actually a lot easier to identify than one specific to an individual and, due to the ability to understand each individual on an individual basis, can be adapted slightly in accordance with each individual's requirements.

For example – business sectors and sub-sectors are numerous and new ones emerge over time, but each one tends to exist for a purpose, as all businesses within them do. It's straightforward enough to work out definitions for a few major ones, to strip them back to their most fundamental purpose:

- *Farming* – producing food fit for human consumption
- *Manufacturing* – producing goods that people want or need
- *Retail* – getting things into the hands of people
- *Travel* – enabling people to move around safely and conveniently

And so on and so forth. These are clearly quite basic and straightforward, and hence easy to adhere to from the perspective of businesses operating in those areas. It is when we get to the sub-sectors that it gets a bit more complex. Using farming as an example, dairy farming would be a sub-sector. The question becomes '*what is the purpose of dairy farming?*' which seems obvious upon initial consideration – logically extending it from the farming purpose, we get '*to create dairy produce fit for human consumption*'. But, when we have the ability to measure everything that happens on an ongoing basis, that purpose seems a little too basic. Measuring is about outcomes and efficiency of process, not just activity, as we gain the ability to assess the success (or not) of everything that happens. Which is where the individual comes in.

While creating dairy produce fit for human consumption might be the overall purpose, clearly AI systems and the ability to measure everything throughout the supply chain, including after the produce has been consumed, means it's possible to be far more comprehensive in analysis than that. We could suggest, for example, that just creating it is not sufficient – there is little point in producing something for human consumption if no humans consume it. Or indeed, the right humans don't consume it, such as the ones who would benefit most from the nutrients dairy can provide. Hence we might go further than the generic overall purpose for farming and add sub-purposes, such as '*ensuring children drink the right volume of milk*' (bearing in mind personalisation could allow the exact optimum quantity to be identified for each individual, based on factors such as age, weight, height and diet) or even getting quite scientific with it – '*ensuring each individual consumes the appropriate level of protein for their unique disposition*', with efficiency of creation, supply, distribution and end-usage all factored in as measures of how successfully purpose is being achieved. This is a world away from the unconnected consumer approach of supermarkets slashing the price of a pint of milk to 49p, which is how we encourage people to use things now, squeezing farmers intensely in the process to ensure consumers keep 'choosing' to buy.

As can be seen from what is quite a simple example, purpose can be tightly aligned with aspiration – what is the purpose of doing something and what could it achieve ideally? In the previous, unconnected age this kind of focus would have been impossible to manage and track the effectiveness of it, but this is the *exact* capability that is enabled by total connectivity and total personalisation, underpinned by AI systems, and aiming for anything less would be to restrict the benefit that can be extracted from the opportunities offered by DCT.

Every single person is an individual with different strengths and weaknesses, existing within a distinct context. This means that each person needs to be considered as an individual, by all the surrounding infrastructure, and assessed according to their circumstances and capabilities, because it is possible to do so now, for the first time in history, in a practical sense. Taking into consideration the move toward total connectivity – which will hugely extend our concept of what we deem ‘employment’ – it will no longer be appropriate to assess an individual’s worth based entirely on their job and tax code. Instead the value calculated at this (individual) level should relate to a far wider raft of measurable criteria – what their specific needs are and how they perform in a range of activities disconnected from a more traditional concept of employment (more on this in subsequent sections).

In essence, the individual level of purpose is tightly integrated with that of business – having a symbiotic relationship based around purpose and efficiency is the difference between consumers and participants. When people are defined as consumers, businesses target and compete over their attention – they are aloof by nature of the relationship and acting within their own interests. Participation is about enabling involvement, understanding the individual and structuring activity around servicing their needs and supporting tasks that they undertake.

In the unconnected world, resource was accessible through an economic system and the more you had, the more you were able to access; irrespective of purpose or intention. In the connected world everything can be measured and assessed, so resource (which is necessarily limited) can be assigned based upon what an individual is intending to do, how they intend to go about doing it and what the likely outcome of that activity may be.

So, under a system based on information, it is the role of business to enable individuals to do things in relation to a specific purpose, on an individual basis, while the role of the individual is to enable them to do so with maximum efficiency.

Guidance

Purpose, on its own, is useless. It just provides something toward which AI can aim. The problem with identifying a purpose is that it doesn’t have to contain any kind of roadmap for achieving it, no detailed blueprint for understanding how it might best be reached; ‘*producing food fit for human consumption*’ becomes little more than a marketing statement in this situation.

Having a defined purpose is a bit like saying you’re going to fly to the moon – it’s not likely to happen without a rocket or meticulous plan for getting it there safely.

Purpose is also perfectly ignorable – as we’ve seen with the UN and its laudable goals, we might sign up to a batch of lofty ambitions only to abandon them at the first sign of difficulty. Purpose needs to be ‘remembered’ during these periods, so that we don’t

become diverted from its course entirely. To achieve this, purpose has to be tightly integrated with two complementary concepts – *guidance* and *value*. These do not just enable a roadmap to be established, they are also necessary to make sense of it.

The closest thing we have to guidance under our existing system comes through economics, which brings with it a continual drive for individuals to accrue money and businesses to develop services that they want (specifically, choose) to spend it on. But, as a guiding concept, its focus is too heavily skewed toward individuality in a consumerist sense; so while individuals are guided toward certain behaviours, sometimes in a collective sense (such as when a new iPhone is released to the mass market), the desired outcome from each individual's perspective is restricted to them and what they stand to gain from any interaction (*'I bought the new iPhone because I wanted it'* – the reasons behind why they wanted it or how they intend to use it are not factored in, the purpose ends at the point of sale when economics is the guiding principle; the transaction is complete). Consequently, as everyone acts in their own interests when defined as consumers, the direction of things is inherently unruly as individuals and businesses are pulled this way or that in an unconnected fashion. Nobody considers anybody else or any other associated factors when buying an iPhone, they just buy it because they want to and believe they can afford to.

Total connectivity, by contrast, needs to be governed by something that recognises every connected element within it individually. This is necessary simply because it's possible – to ignore what happens before, during and after interactions would introduce a level of inefficiency which would seem eminently avoidable when that information is technically available.

This is where guidance comes in, which has a major point of semantic difference with assistance. An AI system characterised by providing assistance suggests that choice is still the core principle that it works alongside, as in 'assisting people to make choices', but with those people still the primary decision-makers – the individual, or consumer for the purposes of this relationship, is the primary entity, the assisting system the more supplementary one that has limited, perhaps even no, agency for making something happen without prior consent from the consumer. The consumer is being assisted in their activities, over which they have ultimate control. An AI system characterised by guidance suggests far greater scope in what it has agency to do. It is not solely reliant on the individual to make the choices for which it provides assistance, but actually has agency to be a decision-maker in the relationship with that individual. This means that it doesn't have to constantly be held up waiting for the individual to make choices before it is able to function, which would be limiting due to the difference in processing speed capability between human and machine; it could just get on with a lot of the heavy lifting in the background, continually and autonomously, due to its ability to assess the outcomes of the decisions it makes against a sense of purpose that is relevant to a given relationship.

Broken down to its most basic principle, guidance helps to connect people to the things, products, services and other people that can help them to fulfil activities that they undertake, in a way that is consistent with a given purpose. This isn't an entirely radical concept, of course – it's how the web works already. When you type a search term into Google, it presents the best matches for what it assumes the person is trying to find based on a number of criteria. The issue with how that has developed is that economics has had to intervene in genuine personalisation, by promoting ads ahead of more 'natural'

matches because a business has paid for it to appear there. This approach prevents people from being connected to the most appropriate matches consistently every time, as it enables money to trump personalisation.

If we move forward toward something like total connectivity, given all the diverse information that will be available about individuals, AI systems will be able to focus on making logical decisions around what each requires, in what volume, at what time and for what purpose. What is the best object, person or service to connect an individual to at a given moment, maximising efficiency of process, to achieve an outcome that could be considered successful when assessed against a sense of purpose?

This can, of course, seem more than a little controlling (though it's worth remembering this is already what happens anyway, it's just the assessment of what constitutes a successful outcome that needs revision). It is not to remove freedom of choice or be dogmatic about what individuals can do – instead, for AI to have a chance of increasing efficiency in the connections and activities it helps facilitate, businesses need to be incentivised to guide individuals in a manner that is mutually beneficial, connecting people and things together to identify what is appropriate for each person given their situation and the context of their interactions with others. Guidance enables and supports choice, but not aimless choice (or: consumer choice).

Guidance also takes account of a number of external factors when assessing what a successful outcome would be. It is not as straightforward as identifying that an individual *'is this size / needs to travel to this place / is hungry'* – just basing decisions on a single individual's requirements admits potential for serious inefficiency in resource management. A really successful outcome would be one that assesses the availability of resource (particularly that which can be locally-sourced in situations where the need is time-sensitive) and makes the most efficient use of that resource given other local connections that are being processed at that time. It may be that one individual has something that they have no immediate use for, which can then be distributed to another place and person to serve a purpose before being returned or even passed on further. This, again, is nothing new – it's the basic underlying principle of most P2P networks. Purpose is most suitably served when forming part of a chain of events, with the outcomes assessed as a collective.

Isolation of activity in a world characterised by our connectedness to each other is a concept that can make no sense anymore. It's logical enough if you think about it – we are entering an age of total individuality, enabled by the move toward personalisation. It requires guidance to ensure people are not just connected but connected to each other in the right places at the right times, increasing their chances of achieving some form of success in whatever it is they need or want to do and reducing the risk of friction with – and inefficiency in – other things someone needs or wants to do locally.

Under a system where people are treated as consumers, choice is regarded as a core virtue and the marketing assault on our senses can be bewildering. Consumers are free to choose from as wide an array of options as can be made available without realistic limit. For example, just standing at the platform on a London Underground station can lead to multiple marketing campaigns being pushed concurrently on billboards, covering as diverse topics as holidays, car purchasing, vitamins, financial management, insurance,

movies, games and albums (and that can hardly be regarded as a comprehensive list of all possible options either). And that's without even including any digital channels.

The opportunities for reaching individuals have expanded exponentially. It's not just through the channels where you would expect it – marketing emails, adverts on publisher sites etc. Instead anywhere that attracts eyeballs seems to be considered fair game – product placements in films / soaps / TV dramas, computer games, sports (particularly 'in-play' elements such as sponsored throw-ins or corners); even the smartphone passcode unlocking sequence or ATMs can carry adverts. Furthermore, the more an individual does through digital channels the more data becomes available, making more sophisticated personalisation possible. This means they can be influenced into associating a product or brand with a purpose or need, as it isn't just the marketing of the product that is important here but the way the communication can be crafted to suit that individual. People can be influenced toward things in very clever ways using data and DCT, but this is currently done in the interest of the brand first, the individual second – which, of course, makes sense from an economic perspective but goes against the principles of personalisation. Too many opportunities for structuring communications in a personalised, leading way through marketing means choice can never really be said to exist, laud it as we may as a virtue of capitalism – we are already guided, just not in the mutual interests of both business *and* individual, and with individuals still defined as consumers; so that guidance is actually assistance.

With each business trying to use these channels in the most effective way to drive engagement with their brand(s), it's easy to see just how aimless this system will become (not to mention how aimless it already is). Each company is trying to generate economic gain from their engagements with people – that is their purpose, which will also be that of their respective AI systems – and they are in competition with everyone else selling a similar type of product or service, but also with other types of business who aren't directly competing with them; as these are the same consumers they are targeting with limited spending power after all.

The purpose of the exercise then is to just secure a sale or build up something slightly less tangible from a business tracking perspective – loyalty, engagement, interest etc – that will hopefully lead to a sale further down the line. It's a logical outcome for a system of economics that functions by distributing varying levels of money to everybody, which then has to be prised away from them by businesses who are ultimately responsible for attracting their attention and convincing them that what they have to offer is right for them. It's a continual process of pulling this way or that in an uncontrolled, competitive, frenzied and aimless fashion. And it can *never* be genuine personalisation.

In an unconnected world this was perhaps to be expected, but it was mitigated by access to prospective customers being restricted due to the relatively few opportunities available for targeting people. When everything is all connected together however, access to people becomes 24/7 potentially with almost no engagement free from commercialisation in some capacity.

Should 'consumer' remain as the definition for what an individual is in the digital age, then the concept of guidance would just equate to a form of highly influenced choice, as the aim behind business activity would be to convince consumers to divert their spend toward the kind of products and services they offer. But we should no longer think of ourselves

as consumers, due to our connectedness to everything – our age is not defined by choice but by *personalisation*, which means knowledge and understanding at the individual level can be comprehensive, making the provision of endless choice an unnecessary requirement. Choice no longer needs to be regarded as a core virtue, elevated and protected against all else.

Consumer choice has little to do with individual freedom; that is not what characterises it. Businesses have no interest in the concept of absolute freedom of choice for their customers. As Henry Ford famously said, he never asked his customers what they actually wanted because, if he did, they would have just replied '*faster horses*'.

By extension, a common axiom in many business sectors is around listening to your customer, finding out what they truly want (as opposed to what you want them to want) and delivering on it. While it's true that many businesses actively seek out such feedback from their customers, it tends to be framed within set parameters. You don't typically find drastic questions on surveys along the lines of '*if we could be / do anything, what would that be?*' – the answers to that might require a level of proposition reappraisal that could sink the business completely, if taken at face value. You are more likely to encounter questions that relate to adapting the current offering – '*which of the following flavours would you like us to produce?*' – which leads to insights that are actionable without being fundamentally disruptive.

From the consumer perspective therefore, choice means they are able to select from a number of options and even feel that they are involved in shaping some of these options by providing such feedback as above. But the available options are designed in line with current consumer trends in mind. The businesses manufacturing and selling these products and services have to be able to convince people to buy them, so being too radical too often would be far too risky a strategy for most.

So people as consumers have choice, but it's choice in relation to *stuff*. As people don't know what stuff is available unless someone tells them, that stuff has to be communicated to them somehow. Obviously, a business is never going to shift many units if it doesn't attempt to increase the appeal of its products and services, even if what they have to offer is really clever / fashionable / useful, so that method of communication is marketing – which might involve clever storytelling, celebrity endorsements, exaggeration of its capacity for improving your life etc. In fact, it's not unreasonable to suggest that it's based around lies and misrepresentation by its very nature – or, to give it another name which might upon initial consideration seem extreme, manipulation.

Manipulation, however, is an inevitable part of communication. It's not as nefarious as it initially sounds – it's not just governments and businesses that do it, close friends and family do too. For example, I'm trying to manipulate you into thinking about a topic from a certain perspective in this text at this very moment, as all non-fiction does. It's the intent behind the manipulation that alters the nature of it, and we might split manipulation into negative and positive manifestations – with negative being carried out in the interests of something external to the individual who is the subject of the communication and positive taking the interests of that individual into account.

In the digital age access to individuals and knowledge about their diverse activities is increasing exponentially. As being a consumer is about choice (or the illusion of at least), the entities in control of communications need to convince people that they should divert

their attention / spend / time in the direction of the product or service they are offering in the interests of driving economic growth for them, either immediately or further down the line. Having a highly in-depth understanding of an individual will make businesses incredibly effective at being able to sell to people (right message, right place, right time, all the time), so the very idea that choice exists as part of these relationships will be highly questionable (the best result for a business is that an individual buys from them exclusively rather than a range of suppliers – which is to say continually has their choices made for them, while still being convinced that they made them).

This kind of choice wrapped up in negative manipulation corresponds with the idea of assistance, but it cannot realistically be labelled 'guidance'. This is where the current application of business personalisation strategies is heading, which is hardly surprising as it just represents an extension of previous, unconnected thinking. In a retail sense, this is because things are seldom made for a specific purpose – people develop products because they fit in with consumer trends and they (usually with investors of some kind in agreement) think they can sell them. The most telling indicator that we make things in an aimless way is the continual cycle of discounts that retailers go through, which helps to clear old stock and make space in the warehouse for the next batch. Discounting is the consumerist way of saying we are making too much, without first confirming that a specific individual has a specific want or need for it.

Without personalisation then, people are likely to make relatively fewer purchases as the things they see either have less chance of being relevant to them or are not marketed in a way that will be most effective for each individual. When personalisation is added, but only in line with unconnected thinking, the goal logically becomes to create stronger links between individuals, brands and products to generate a higher chance of making a sale – which should lead to more people making more purchases, but primarily in the interests of the business selling the products as opposed to servicing the needs of an individual, based on all the information that can be known about the context of the interaction(s). Certainly, there are numerous retail case studies that attest for increases in conversion rates following the introduction of personalisation to various areas of their marketing campaigns^{xv}.

Does this represent a greater sense of choice on the part of the consumer, or does choice actually gradually become the domain of businesses who are able to make the best possible selections for the things they want to sell? Consumers assume that choice lies with them, but it is being ceded over to business without proper recognition on the part of the average individual.

If we add personalisation again but this time underpin it with connected-age thinking, the result is actually similar in that choice, to a large extent, is transferred over to businesses (AI systems, specifically). Yet there is a fundamental difference – instead of making something in isolation from someone then trying to convince them they need it, businesses would be guided toward providing products and services for the specific individuals for whom they would be most appropriate, or that would most benefit from them, at a given time. The primary focus then would be on the individuals themselves, with businesses creating and crafting the things they do around each individual's specific circumstances and requirements – because data makes it possible, it would be inconceivable to approach it in any other way due to the huge inherent inefficiencies.

This would be some distance from choice as we currently understand it, as it would be far more of a 'made to order' model (for longer-term requirements) that directly reacts to the needs of each individual. This is very different from innumerable products being built up in isolation then marketed to people – who then choose the one they think works best for them and reject countless other options which may remain unsold and unused.

So if the post-consumer period is not characterised by choice (or, at least, what choice on the part of the individual actually means has to undergo some degree of revision) and there is no requirement for a common standard such as money anymore, what comes to characterise the nature of interactions between people and external entities? If the various incentives that an economic system provides for people are removed, then which non-economic factors replace those incentives?

Or, to put it another way, why would anyone do anything?

Purpose is the value that informs the reasoning behind why any activity is being done, which in many cases require resource in order to achieve an outcome. This might be as overt as creating highly advanced and specific fittings that an engineer needs to build a rocket, right through to far more standard items that act as supports to enable an individual to go about in pursuit of something that they either need or want to do – footwear for rock climbing, underwear, washing-up liquid, toilet paper, soap – it all has its purpose and its usage can be tracked (within reason – there is no need to track every sheet of toilet paper individually, surely), which in turn helps to build a measure of value.

As everything that happens and every action undertaken can be tracked, measured and assessed against various criteria in a world of total connectivity, people cannot help but become participants in everything that is going on – either in a tiny or highly significant way, depending on their capabilities and the purpose for which these things are being done.

Purpose has to be, by its nature, a connected and integrated concept. It's not just something that can be worked out and completed by an individual completely in isolation – AI systems can be instrumental in providing clarity and guidance around what needs to be done or achieved in support of a highly efficient system interfacing between business and individual. With the amount of information digital makes available, it becomes possible to cross-reference all manner of datasets that can help to identify an individual's capacity for involvement in pursuit of a given purpose. For example, the kind and topic of content that an individual leans toward indicates interests; tracking their performance in certain tasks, from problem-solving right through to how fast they can type indicates skills; analysis of probable intelligence level from how they communicate and act can indicate capability.

Total connectivity enables AI to calculate the type of role and activity that an individual could suit well. It is then a matter of guiding them toward, and incentivising them to undertake, activities related to those areas of strength and forging the right connections that helps them to do it. At present the main way this could happen is by situating that individual within a company, performing a traditional role within it (analyst, driver, accountant etc), should one such role be available. There is no reason why that individual couldn't undertake activity associated with those roles unconnected to a business, but businesses are how people get paid under current definition, so they wouldn't be able to

accrue any income in the process unless they set up independently and brought in some paying clients.

When economics doesn't make sense anymore and money no longer acts as the measure of business success, activity can be undertaken by whomever wherever, with AI systems able to assess the purpose and success or otherwise of outcomes from that activity. The best way to achieve a purpose may be to work for a business in a traditional way, or it may be to operate in complete isolation or in combination with others not tied to the same business – the outcome and impact on efficiency is what matters, not the form. The situation of their locality may be determining in that respect. If there are lots of people with a certain set of skills in one geographic location, they may be guided toward working together on tasks of high relevance to their skill-sets. Equally people in proximity to each other may have a diverse range of skills, which could enable them to work together on something of a different nature. It may also be able to identify gaps in the market that could be filled (*'you have a certain set of skills and knowledge that could make this happen, which is not something currently available here, hence it represents a clear opportunity for you to fill'*).

The nature of our interactions then relates to what each individual can be guided toward doing in order to achieve some kind of purpose, whether fairly specific to them or in combination with one or many others. The depth of contribution made can vary widely, but any contribution is important nonetheless as it can be measured.

Capability is another important consideration here. The number of people on Earth at any given time who are capable of excelling at a given task or doing something new and innovative is necessarily limited – not all of us can discover new particles, design aircraft or write *The Iliad*. But having a top bracket of likely achievers doesn't stop the rest of humankind having a key, enabling role in helping them to do it that in many ways is just as important and can be recognised for its contribution. The point of guidance is not to dictate what an individual should do, but to create connections between people at the right time and allow them to select whom to work with, what endeavours to support, how they want to get involved etc to achieve certain outcomes, evolve areas of human endeavour or enable things to function with maximum efficiency. Choice, in that respect, is still an important element of the process.

This idea is not new in digital, far from it – we already have multiple P2P networks (such as Freelancer and Guru) through which people either advertise for roles they are looking to fill, or create adverts showcasing their own skill-sets and the kind of work they can undertake. On these networks the 'middle man' of business is absent to all intents and purposes, allowing people to seek out appropriate connections themselves. The problem with these networks in their current form is the fact that money is still the core determining factor – it's all a bit mercenary, with people prospecting for work solely for personal gain rather than working toward any integrated or common purpose. An advertiser can either afford to employ someone or they can't, irrespective of what that individual's involvement could help to achieve or what outcomes could be possible.

If the need to make money was removed, P2P networks offer an early glimpse of how a system of guidance could work in reality. There is no logical reason for a world where the infrastructure is based around P2P networks to be characterised by commercial gain – they work by creating connections between people and, as there is no longer need for a

common standard such as money, people can be incentivised in other ways to do things to help each other.

Eventually it may be possible to get rid of the idea of businesses as external entities altogether and replace that structure with something like a universal P2P system, which is much more suited to a system based around AI, total personalisation and the continual creation of meaningful connections.

The concept of guidance is a fairly arbitrary one. For it to have a practical application in the real world, it requires something to oversee its operations – which is of course AI. This may refer to a singular overall AI system or a massive amount of independent but interrelated AI systems that process and share different bits of information. Whichever of these makes more sense and proves to be most effective, there is one thing that is absolutely clear – humans are not capable of overseeing a system of guidance. It is just not possible to get anywhere near it without having the ability to process incredible amounts of information and make decisions on how to make connections between people based on that information – and in accordance with purpose – in something like real time.

A common axiom in the IT security industry is that a network is only as strong as its weakest point. For an automated enabling system for total personalisation, we might extend this to refer to any point that relies on manual processes. Where anything requires manual intervention, it is by modern definition inefficient. Having any weakness such as that will inevitably introduce a degree of latency that would clog up the efficient functioning of the overall system.

The idea of multiple AIs also admits a contradiction, as it suggests the existence of data siloes – which is one of the key barriers to the existence of total personalisation. It would of course not be possible for a single system to fully understand everything in real time at all times, so it may be that the processing is divided up into logical segments – perhaps geographical, or linked to interrelated groups of people. This would not be the same as the siloes we see under a system of economics and private ownership however, as all data would still need to be shareable between each independent AI segment. AI would need to make allowances for ‘correcting’ decisions it makes in relation to individuals as new information is received – although this is also the nature of how it would work anyway.

AI systems should be able to execute the personalisation guidance system with relative autonomy, as humans are unable to process information at a quick enough pace to be particularly useful. This is an entirely different prospect to ‘using’ AI as a means for guiding us but ultimately looking to have executive control. There is distance still to travel – as yet AI has tended to struggle when having to deal with unique or novel situations. It’s also worth restating that AI in this context does not refer to the idea of humanoids – machines that look and act exactly like humans – but a more intangible cloud of knowledge and guidance, processing and cross-referencing huge volumes of data, that is wound in tightly with human endeavour. An integrated, symbiotic co-existence that ties us together in ways that neither of us could – nor would want to – rupture.

Value

Purpose and guidance work in combination to generate outcomes. A third non-economic factor is then required to interpret how successfully they are doing this into a format that

AI systems can process, understand, share and monitor. That factor can best be expressed through a measurement of *value*.

To recap briefly, total connectivity means that the calculation of value has to shift away from '*this thing costs X*' to something rather more personalised, complex and tangled, worked out over a longer period and updated (and backdated) as things subsequently evolve – specifically because it can be tracked, the information cannot logically be ignored. Money – the previous designator of value – is a system that makes sense when two people / companies / governments have limited means of communicating information about the long-term impact of a transaction or interaction. It just deals with what is exchanged in a single moment in time.

When so much can be known, processed and measured in real time (because everything we do becomes data) in relation to so much that can happen, just having a price at the point of purchase makes no sense. It ignores anything that happened before or subsequent to that point of purchase and takes no account of context, need or individual circumstance.

This is because the actual value of what has been achieved in any given situation – when considered in relation to its connection to a purpose – can never end at the transaction; the value produced can change in line with what the transaction has actually contributed toward achieving that purpose, or multiple associated purposes, measured over time and on an ongoing basis.

Imagining a system that is characterised as establishing connections between people and things to serve a purpose is a simple enough leap to make – that is how the web works already really, just with the underlying purpose being economic in nature currently – but it doesn't immediately make clear how every use that money served could be replaced. It is still necessary to determine who can access which type of stuff, at what time and in what quantities. After all there is always limited space available on any plane, particularly those destined for a paradise island.

In the era of total connectivity, this can be managed through the introduction of *value ratings* for each individual, which replace the role of determiner previously played by money – which are calculated and built up by measuring the contribution that individuals make toward purpose(s), both in terms of their own performance against certain metrics, but also how they support and enable the efficiency with which businesses are able to operate. In essence, it's a non-economic, benign alternative to the Chinese 'citizen score' system introduced in part one. Or, indeed, a non-economic extension of the credit score system already widely used in many countries around the world.

As we head into the period during which, so the forecasts tell us, there will be a huge tilt toward automation across multiple sectors, it's possible that what we think of as employment needs to evolve (unless automation creates new 'traditional business employment' replacement positions in the process). Given the scale of this change, value cannot be restricted to just doing what we'd currently think of as 'work' – ie your salaried day job, as calculated against a 37-and-a-half hour working week. As everything can be known, tracked and measured it doesn't make sense to clock out at 17:30 anymore – we can be measured 24 hours a day, already are, in so much of what we do, so the definition needs to adapt to that new reality.

In a world in which economics no longer makes any sense and information is the new currency, almost anything we do can potentially accrue value – as all our actions and activities can contribute (in however small a way) to the achievement of a purpose. When you know everything about an individual, being alive *is* your job.

Obvious examples for accruing value would be to extend it into, though by no means limit it to, those general areas over which we do ultimately have control, but that, in many cases, lack sufficient focus on the part of the individual to manage as well as they might like – maintaining a balanced diet, keeping fit, learning new skills and educating ourselves, as well as things we do that contribute to the development and wellbeing of others (consider how this might work for parent and child relationships). These are all perfectly reasonable – and perfectly trackable – personal levels of purpose where an individual may create realistic targets and aspire toward them. Neither is there anything particularly radical about these; many people today already look to improve on their performance within these kind of areas, which some are effective at and others less so. But this is the crucial point to understand here – this doesn't mean everyone aiming for impossibly high targets; remember that all experiences can be personalised, with the attainment of higher value ratings linked to achieving what's *realistic* for an individual's circumstances. Instead of saying every candidate needs to get this percentage of answers on an exam paper correct to qualify for the top mark, that percentage may be adjusted – or the structure of the examination itself adapted – to produce a sliding scale of poor and good performance in accordance with that individual's capabilities.

Value sounds like an economic term, but it is in fact distinct from economic systems by its nature. Money, for example, can exist in huge quantities that are separate from any individual and be transferred between people, businesses, governments etc at will and in bulk. It is not necessarily tied to anything, hence the aimlessness of it that puts it so out of kilter with an age defined by information. Our current approach to economics has, in fact, produced an extending gap between rich and poor in a way that seems to have little to do with value – those with the resource available are able to do more economically, as the more an individual has, the greater the access they have to the vast infrastructure of wealth management businesses that have been specifically constructed around economic logic (providing tax efficiencies / avoidance, smart investments etc), so the purpose of them, how they measure success, is through getting a stronger return on their financial portfolio rather than assessing the overall outcome of activity. In 2016 Oxfam analysed data from Credit Suisse and concluded that the richest 1% of the UK population owns more than 20 times the wealth of the poorest fifth^{xvi} - with the gap continuing to widen. Whether it could be said that they contributed a level of value relative to that ratio is another question.

Value on the other hand, can only be accrued in relation to an individual, business or other entity and the trackable activities they undertake. To illustrate this fundamental difference, consider how absurd some common phrases are when we substitute 'value' for 'money':

- *Have you got any value?*
- *Can I borrow some value?*
- *I just won some value on the horses*

- *How much value do you get per hour?*

It is perhaps more accurate to think of value as something that is generated, rather than earned as is the case with money – which often carries a sense of entitlement; someone being contracted at a generic set-rate per hour to perform a job, irrespective of performance or impact within that role.

Conversely, the value of any activity is something that has to be calculated and measured. Two people undertaking the exact same activity for the exact same time are unlikely to generate the exact same output. But simply assessing the value that each is generating in a generic sense (from the perspective of productivity; so this person is capable of creating 15 outputs an hour, this person only 12) could easily lead to an elitist system where the quickest / brightest have all the opportunities, leaving a huge gap in value between them and those unable to reach their levels.

Due to total personalisation, it would be possible to identify realistic expectations for each individual and assess the value they are generating relative to that. It would therefore be possible in theory for any individual to attain a high value rating – based on criteria that is specific to them and their circumstances, rather than any generic criteria, levelled out for all participants.

Value is potentially situated within every stage of the lifecycle of any activity or output, from conception through design, manufacture, distribution and subsequent impact it has on the achievement of an individual or business purpose. Hence the opportunity to influence value is widespread, diverse, continual, collaborative and inclusive.

The value of something lies not in itself, but in what it does, how it is used and what it enables someone to achieve. It is becoming possible to transfer the determination of value from within something (ie this lemon costs 30p) to what that thing achieves, allowing a truer picture to form of what value was actually generated from any interaction, transaction or activity.

The fundamental shift here is in seeing past the price point – the idea that something no longer has a price associated with it as such, but that things can be created or used with a specific purpose in mind, connected to a specific individual and their unique circumstances. Even in situations where the product is fairly generic – a bag of carrots, for example – it will still become possible to make assessments of the impact that they had, such as knowing whether a specific individual ate them (combining with remarkable levels of information about their individual circumstances and requirements, whether they needed them etc). Any business involved in any way in the creation, distribution or use of these products will be able to factor these measurements into how they report on the success of their endeavours. In this situation, just basing this on whether they sold them in the first place, paying no attention to what happens subsequently, is tantamount to pretending that digital doesn't exist.

While there will always be things that can be produced in bulk and in generic fashion (eg toilet roll), total personalisation is about using all the information available in both the interests of the involved individuals and enabling businesses, with those individuals playing their part by helping the processes to run as efficiently as possible. At present our approach is to manufacture products in bulk and to a generic design, which helps with efficiency but produces output that is unspecific to the unique needs of any individual and

therefore needs to be mass-marketed to convince people that they need them. Which admits huge inefficiencies throughout all subsequent processes.

Total personalisation is not about scale or sales targets – if a business produces ten outputs that are highly tailored to the needs of ten individuals, genuinely improving a specific area of their lives in the process by helping contribute value toward a given purpose, the outcome can be just as good from the perspective of value as producing ten thousand (at least, in its own way – at the very large scale that kind of output necessarily becomes more generic by nature). It is instead about understanding the unique needs of each unique individual and using that information to guide them towards either the attainment of, or contribution toward, a purpose.

While a system of guidance can make sensible, logical and meaningful recommendations in terms of whom to forge connections between, it's clear that people cannot be forced to take on one activity or another. There will still need to be plenty of scope for choice (in the non-consumer sense), and there may be a trend for people to generally gravitate toward those that are easier or less onerous to undertake.

For instance, how would someone be convinced to try to cure a complex disease, with all the laborious and lengthy research, tests and experiments that come with it, if they equally have the option to paint a bench in their local park or help deliver a parcel to a destination near their home?

An obvious approach to this issue would be to vary the levels of value available for any given activity. So painting a park bench would be relatively low, while discovering the Higgs Boson would be relatively high. While this would appear to make sense on the surface, it actually risks establishing (or continuing, arguably) an elitist society in which only those with the highest brain capacity realistically have the opportunity to generate significant value. Those merely deemed capable of undertaking the more menial tasks would find themselves in an eternally suppressed situation, as they would seldom be connected to anything of any serious significance, restricting them to tackling low value tasks only, extending the gap between top and bottom participants in a way that would be difficult to redress.

Varying the levels of value in a given context to some extent will clearly be necessary – after all, painting a bench cannot realistically be regarded as the same as making fundamental scientific discoveries. However, instead of focusing the determination of value on the task in isolation, it has to be calculated in a more complex way. It is a combination of factors such as:

- The difficulty / sophistication of the task;
- The time and effort required;
- What level of efficiency it enables; and
- The urgency of it / how necessary it is.

And that is only one part of it. Then, of course, the unique circumstances of the individual need to be factored in to ensure that each task appeals to the right person. For instance, if a brain surgeon paints a bench the outcome is more or less the same as if a less skilled individual does it, but clearly the surgeon's time would probably have been better spent on other, more appropriate tasks if there was a waiting list of people requiring operations

at that time. The value generated also has to be very tightly aligned with the individual, although even still there is a further complication as value can no longer be determined at a single point – as everything can be tracked and measured on an ongoing basis, value has a far longer lifecycle. If the bench was finished beautifully, but someone sat on it shortly afterwards and got covered in wet paint because the painter forgot to put up a warning sign, then the measurement of the outcome would be negatively impacted accordingly.

This shift is simply reflective of the kind of world we are approaching (in many ways, have already arrived in), where things that were previously untracked and, in truth, unconsidered, can suddenly be measured and carry great meaning.

For example, the concept of ‘bit walking’ (the technical term for tracking a user’s steps) is already being turned into a functional business model. Still at the prototype stage at the time of writing, this will work by enabling people to track the number of steps they take in a day using some kind of device (perhaps a wristband), with users earning ‘walking dollars’ for each 10,000 steps they take – which is the daily step target generally recognised as the benchmark at which people should aim. These walking dollars can then be used in an online store and, of course, this data will be commercially attractive to insurance companies and sports brands.

If we were to look at it from the perspective of a value system rather than a financial one, it’s easy to see how there is no reason for it to stop at walking but could also cover other factors that can be measured – how many cigarettes someone smokes over a specific period (as in actually does, rather than how many they tell their doctor they smoke), how many units of alcohol consumed, what other forms of exercise they have engaged in, their exact calorie intake in a day. Elements such as this, not an exhaustive list by any means, can contribute meaningfully toward the calculation of a value rating (and by association, purpose) in some capacity. Since it is possible to know, it’s impossible to ignore.

The digital connected age tells us that everything can be measured, in principle at least. Almost nothing is exempt from this if the inclination to measure a given element is there, for example where doing so would create value or lead to fresh insights. It does seem reasonable to measure how many steps an individual takes, and also to use that information to determine some kind of output. But this is only ever going to be acceptable to people if highly personal information – not just how many steps an individual takes in a day, but where they took them, with whom etc – is used in their interest as opposed to those of some external entity who is commercialising that information.

Of course, focusing on the bit walking model, we can see that it will be fundamentally flawed in a system of personalisation under its current expression – for it deals exclusively with one tiny part of an individual’s existence. It doesn’t, for instance, recognise whether they go swimming or cycling. If they only take 5,000 steps in a day but also do 50 laps of the pool, they will be regarded by the bit walking system as performing well below the threshold – when in truth they have done far more rigorous exercise than someone achieving the full 10,000 daily target. It is also generic – 10,000 steps might be a sensible target for some people (though whether it, as a popular benchmark promoted by doctors, is a good standard is a moot point – its source was a 1960s marketing campaign for an early form of pedometer in Japan, based on research by an academic at Kyushu University of Health and Welfare), but for those in a wheelchair or that have

otherwise restricted manoeuvrability (even temporarily, such as where an individual twitches their hamstring) it's meaningless. For someone who is double their recommended weight, it might be far more useful to set a lower threshold and work up to the full 10,000.

Assessing anything in isolation in a world of total connectivity is also meaningless. Consider a situation in which an individual took 10,000 steps, vandalising a school and stealing a car along the way. Looking at bit walking in isolation, they are eligible for a reward but should they be rewarded? Probably not.

Now let's extend the idea of measuring steps into a more comprehensive view of wellbeing. Part of (and only part of, given so much can be included) the determination of an individual's value rating might be an assessment of how well they adhere, generally speaking, to guidelines around lifestyle – which would all be personalised to suit their unique circumstances. At present, health advice is incredibly generic – aim for 10,000 steps a day, eat five fruit and vegetables, take vitamin D.

The vitamin D example is particularly revealing of how generic and meaningless blanket advice can be. In July 2016 Public Health England^{xvii} announced that people may not be able to get the recommended level of vitamin D between October and March because of the reduced chance of being exposed to direct sunlight over that period (due to shorter days and autumnal / winter / overcast weather). Advice from public officials was therefore that everybody in the UK should consider taking vitamin D supplements.

Just to confirm, that's *everybody* in the UK – about 64 million people.

At that kind of scale, this kind of advice is little better than a shot in the dark. The problem with this approach is that it causes unnecessary alarm as, in reality and at present, people have absolutely no way of knowing whether they are getting sufficient vitamin D. It may be that some individuals are getting the right level, but are lacking in some other vitamins. Yet instead of focusing on getting that balance right at the individual level, their attention is diverted toward something that may not actually be a concern for them purely because the advice they receive is so generic.

Taking the personalised view, it would instead be possible for AI systems to calculate and map out nutritional (ie meal) suggestions based on what has been digested recently and what will likely be digested in the near future – with that being adjusted in accordance with items an individual may ingest in the meantime – to help promote a balanced diet that corresponds with their individual nutritional requirements (also taking into consideration other medical measurements such as blood sugar levels).

So to repeat: this is not to be prescriptive; people could still eat and drink foodstuffs that are high in sugar / salt / alcohol, but in order to maintain a positive value rating, this would need to be offset by also choosing options that either provide the right nutritional balance overall, help them improve in some area or other (perhaps even only in a small way) or enable businesses to operate efficiently in how they prepare and distribute the produce that they consume. The aim is not to tell people what they should and shouldn't eat at all times – it is to guide them toward making choices that will work for them, but that prevent them from making poor dietary choices consistently, or choices that reduce the efficiency with which businesses can operate.

And, of course, having a blowout from time to time can also be catered for. It's ok to go out and drink eight pints of lager, but probably not every day. It is for the good of each individual that value ratings exist, and its calculation should cover a range of factors specific to that individual so that expectations and targets can be realistically set. Through the provision of guidance in such matters as important as wellbeing, the benefits to the infrastructure that makes up society are also great – easing the burden on healthcare and welfare services, in addition to maximising use of locally-available resource rather than continually making requests that are complex for the businesses involved to fulfil. Weight Watchers have used a points system for years that may provide a useful model for adaptation into a system of total personalisation.

Neither should the calculation of value be restricted to individuals' wellbeing – every aspect that makes up our lives can be measured and included. It would seem highly unlikely that every participant would achieve a high rating across every area of their lives consistently, but getting some kind of balance across those areas would seem a more reasonable target (though this balance can, as ever, be personalised for each individual's circumstances).

For example – education. This doesn't end with graduation from school / college / university, but extends across our whole lives. Under our economic system, we've ended up in a situation where people have to pay for their education, up to £9,000 per year for some universities in the UK – a prohibitive fee for many poorer prospective students. This is a reverse of the ideal situation, which would be paying people to provide incentives for performing well at studying.

This means people who continually challenge themselves to learn new things, both in the traditional educational sense and in developing new skills, can be recognised for it. Performing well in mastering skills and knowledge would clearly accrue value, as it raises their capacity for performing a greater variety of roles, but it shouldn't be just about being among the 'top graders' – not everyone starts from the same base or has the same capabilities, so for some making marginal gains in increasing their knowledge or understanding in a given area might represent a valuable outcome from their perspective.

Of course, when it comes to employment, many people today already do what they would suggest constitutes a day's work without necessarily being employed per se by an external company – and the ability to track everything, coupled with the decentering of traditional financial control away from holding businesses, means that value can extend into these, and theoretically many other, areas too.

Being a parent for example is a major, engulfing undertaking on the part of the parent(s) that provides no financial recognition outside of some welfare support where certain criteria is satisfied. Yet it can be tracked and assessed under a system of total connectivity, enabling the assignation of measures against things like, but not limited to, the child's welfare and development in a number of areas. Some can be generic, such as successfully observing doctor's advice when pregnant by avoiding smoking and broadly sticking to the recommended nutritional intake, but everything can of course be highly personalised to be appropriate to the circumstances of the parent and child. It would also need to recognise that generic expectations are inappropriate – so if the child is relatively late with beginning to walk or talk, it is not necessarily representative of a failure on the part of the parent in developing him or her; every child is different and develops in

different ways. There is no simple formula for elements such as this but, as it can be measured, it can't be ignored as a means for calculating potential value. Value, in the non-economic sense, has no logical boundaries.

Or consider the general upkeep of a home. This is something that takes up time and effort, in some cases hours every day, but it accrues no value under an economic system (who would pay someone to clean their own home?). Yet it is important – ensuring that it is kept clean and safe, reducing outbreaks of illness through exposure to bacteria and keeping rodents away. Naturally this extends outside the home too, where those who pick up litter and clean up after their pets demonstrate greater value within their local surroundings on a daily basis already, but this is currently unrecognised from the perspective of value. The infrastructure regards those who throw litter – and those who pick it up – as indistinct.

Volunteering would be another example of an area in which roles are currently unsalaried, but often contribute value – sometimes a significant volume – to people's lives. Whether an individual is working directly for a company on a salaried contract is beside the point – anything that can be measured becomes employment potentially, as it can create value in relation to purpose and guidance.

It is also important to remember that the value in something in the digital age cannot be considered in isolation – everything is defined by its connectedness to everything else, so all the ways in which employment can be extended, as suggested in this section, need to be assessed in association with each other and in a personalised way to suit each unique individual, their unique circumstances and the activities through which businesses engage them.

Value ratings would provide an indicator of how successful or otherwise an individual is being in contributing to a sense of purpose, with the factors influencing that calculation personalised to factor in their unique circumstances and capabilities. However, to rely exclusively on value ratings alone would risk engendering a culture in which people just focus specifically on boosting their own value, as opposed to seeing it as something that is integrated into a wider whole – value being something that is contributed rather than taken.

To counter this risk, AI is not just required to provide guidance but also feedback and analysis on contribution and results. This would not be high level in tone (merely stating overall achievements – *'we've built a skyscraper / invented a new microchip'* etc), but provide the absolute, personalised detail that helps each individual understand exactly what their contribution has been toward enabling efficiency or supporting the achievement of something – what the direct consequences of their activities has been to the realisation of larger, even macro projects and endeavours. Consider, for example, an individual having a requirement for something that is either rare or complex to create and distribute. Instead of being intransigent and saying *'it has to be exactly that, only that'*, they may compromise and accept an alternative that is easy to source, possibly comprising resources available locally to allow for efficiency in supply of the goods. A feedback mechanism could indicate to that individual how much more efficient that solution is to the original, perhaps in terms of time to fulfil, resource (people, equipment, transport) required, fuel used and volume of emissions. Furthermore, it could provide information on how the resource saved will be reassigned. This would help that individual to understand

the role they are playing in improving efficiency, which in turn frees up infrastructure to focus on other areas that may have a positive impact on something that the individual uses or would benefit from.

An element of benchmarking may also be useful in encouraging individuals to adjust certain elements of their behaviour. Benchmarking typically works by pooling a number of data metrics, drawn from a number of sources who share certain characteristics (so the output is actually comparable and meaningful). This data is then aggregated to produce averages – as well as high and low medians, if appropriate – so that each individual participant in that benchmark can understand how they are performing against these aggregated averages.

By way of example, online retailers like to have this information as they typically only have access to data specific to their own activities – *their* conversion rates, *their* sales growth, *their* average basket value. This information is only so useful in isolation – so if in June a mid-tier fashion retailer had a conversion rate of 3.8% and an average basket value of £70, is that good? Benchmarking against other mid-tier fashion retailers may reveal that the average conversion rate was 4% in June, so they then know that they were not proving as effective at converting site visitors into customers during that period as others in their sector were. Once they have that knowledge, it identifies areas of weakness and strength thereby helping them to focus attention and resource accordingly.

This idea can be translated easily enough into the age of total connectivity and total personalisation. On the one hand it may not seem particularly useful to enable people to benchmark themselves in this way, as we can be recognised as being completely unique by the surrounding infrastructure. Surely this could never, by definition, be comparing apples with apples? Benchmarking in this sense initially appears to reverse the understanding of people at the individual level and focus again on segmentation.

This is true enough in principle, but having this level of information and not using it to provide individuals with insight as to how they are performing would seem a wasted opportunity. It is all about selecting the most useful metrics to aggregate and ensuring each individual is included in segments that will provide meaningful insight for them – identifying who their ‘peers’ are for a range of disparate measures in a world of total personalisation.

A fairly logical example would be around diet and fitness. Individuals might be able to see how many calories people who are very similar to them – in terms perhaps of body mass, lifestyle, regular exercise patterns, geography etc – consume (as distinct from the current generic guidelines we receive around daily recommended intake – this would reveal how much people *actually* have). Other options might be average steps taken in a day or what volumes of a specific foodstuff is consumed weekly – reporting back along the lines of ‘*you ate 2 chocolate bars this week, the average among your peers was 1.3*’. All of this provides actionable insight that can help individuals to understand whether they are doing well compared to others similar to themselves and adjust small details of their behaviour in line with how they are currently performing.

The calculation and determination of value is complex, necessarily so – as it can no longer just exist at the point of interaction between two trading entities. The in-depth and wide-ranging availability of information that can be shared (in principle at least) between

everything that is connected means that measurement of value is an ongoing concern – that may build and decrease over time, and in some cases may be very long-term.

What value actually means and how it can be measured is likely to vary, in some cases greatly, depending on the individuals / businesses involved in any activity and the depth of their involvement. Connectivity is also fundamentally about collaboration – everything that happens has a potential knock-on effect for everything that it is connected to, directly or indirectly, and in a very minor or very major way. This has arguably always been the case, even in an unconnected age – but now actual measurement of these impacts can feasibly be measured for the first time.

This means that value cannot be rigidly fixed, but must be agile and adaptable to ensure the right balance exists between different (perhaps sometimes even contradictory) purposes toward which disparate people may be working.

So – while everybody needs to be involved in determining purpose and value collaboratively, clearly only a highly sophisticated system of AI can oversee such complexity at such pace and at such scale.

ⁱ Interestingly, it may be more of an internal term than something they want to push publicly as mention of it has vanished. For a while around 2016 though, it featured in their conference presentations (which, despite Google’s wacky public image, are heavily PR-controlled). They must have subsequently decided against pushing that as a marketing message in the interim.

ⁱⁱ Wired, *Google: our assistant will trigger the next era of AI*, 25 October 2016, <https://www.wired.com/2016/10/google-our-assistant-will-trigger-the-next-era-of-ai/>

ⁱⁱⁱ Search executed on Google, 18 July 2017.

^{iv} The Seattle Times, *Amazon has sold more than 11 million Echo devices*, Morgan Stanley says, 19 January 2017, <http://www.seattletimes.com/business/amazon/amazon-has-sold-more-than-11-million-echo-devices-morgan-stanley-says/>

^v Fortune, *Prime Day Sets Sales Records For Amazon*, 12 July 2017, <http://fortune.com/2017/07/12/amazon-prime-day-2017-sales-record/>

^{vi} Reported in The New York Times, *Media Websites Battle Faltering Ad Revenue and Traffic*, 17 April 2016, <https://www.nytimes.com/2016/04/18/business/media-websites-battle-falteringad-revenue-and-traffic.html? r=0>

^{vii} Wired, *Google: our assistant will trigger the next era of AI*, 25 October 2016, <https://www.wired.com/2016/10/google-our-assistant-will-trigger-the-next-era-of-ai/>

^{viii} European Commission, *Antitrust: Commission fines Google €2.42 billion for abusing dominance as search engine*, 27 June 2017, https://ec.europa.eu/malta/news/antitrust-commission-fines-google-€242-billion-abusing-dominance-search-engine_en

-
- ^{ix} Free Patents Online, *United States Patent US9665881*, 30 May 2017, <http://www.freepatentsonline.com/9665881.pdf>
- ^x BBC, *Kim Dotcom demos micro-payment service to help stop piracy*, 31 August 2017, <http://www.bbc.co.uk/news/technology-41094797>
- ^{xi} Recode, *Amazon's Alexa and Microsoft's Cortana are going to work together*, 30 August 2017, <https://www.recode.net/2017/8/30/16225140/alex-cortana-integration-amazon-microsoft-siri-google-assistant>
- ^{xii} Chen Sun, Abhinav Shrivastava, Saurabh Singh and Abhinav Gupta, *Revisiting Unreasonable Effectiveness of Data in Deep Learning Era*, March 2017, <https://arxiv.org/pdf/1707.02968.pdf>
- ^{xiii} Forbes, *Google's Capitalist And Proud To Be So: Anyone Got a Problem With this?*, 13 December 2012, <https://www.forbes.com/sites/timworstall/2012/12/13/googles-capitalist-and-proud-to-be-so-anyone-got-a-problem-with-this/>
- ^{xiv} UN Charter, <http://www.un.org/en/sections/un-charter/un-charter-full-text/index.html>
- ^{xv} These case studies, of course, must always be taken with a pinch of salt, as every business doesn't start from the same point. So, for example, if a retailer added personalisation to an area of their communications and it led to a high uplift of, say, 50% in sales conversions, we might question how basic their proposition was beforehand.
- ^{xvi} Oxfam, *How to Close Great Britain's Great Divide: The business of tackling inequality*, says Oxfam, 13 September 2016, <https://policy-practice.oxfam.org.uk/publications/how-to-close-great-britains-great-divide-the-business-of-tackling-inequality-620059>
- ^{xvii} Based on findings of: Scientific Advisory Committee on Nutrition, *Vitamin D and Health*, July 2016, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/537616/SACN_Vitamin_D_and_Health_report.pdf