Sustainable Superabundance:
A universal transhumanist manifesto for the 2020s and beyond

1. Advance!

Beyond the fear and chaos of contemporary life, there is good news to share.

A new era is at hand: the era of sustainable superabundance. In this era, the positive potential of humanity can develop in truly profound ways.

The key to this new era is to take wise advantage of the remarkable capabilities of twenty-first century science and technology: robotics, biotech, neurotech, greentech, collabtech, artificial intelligence, and much more.

These technologies can provide all of us with the means to live better than well - to be healthier and fitter than ever before; nourished emotionally and spiritually as well as physically; and living at peace with ourselves, the environment, and our neighbours both near and far.

This is not a vision of today's society writ large - a mere abundance of today's goods, services, activities, relationships, and rewards. It's a vision of a superabundance, with new qualities rather than just new quantities.

This is not a vision of returning to some imagined prior historical period - to some supposed bygone golden age. It's a vision of advancing to a new society, featuring levels of human flourishing never before possible.

This is not a vision restricted to the few - to an elite percentage of today's humanity. It's a universal vision, for everyone, of a wide, diverse fellowship in which all can freely participate, and in which all can enjoy unprecedented benefits.

This is not a vision of the far-off future - something relevant, perhaps, to our great-grandchildren. It's a vision of change that could accelerate dramatically throughout the 2020s - a vision that is intensely relevant as the year 2020 comes into view.

This is not a vision of a fixed, rigid utopia. It's a vision of the collaborative creation of a sustainable, open-ended, evolving social framework. In this new framework, every one of us will be empowered to make and follow our own choices without fear or favour.

In this vision, the sky will no longer be the limit. In this vision, the cosmos beckons, with its vast resources and endless possibilities. In this vision, our destiny lies in the ongoing exploration and development of both outer and inner space, as we keep reaching forwards together to higher levels of consciousness and to experiences with ever greater significance.

Critical choices

But first, we face some hard, critical choices - choices that will determine our future. If we

choose poorly, technology will do much more harm than good. If we choose poorly, a bleak future awaits us - wretched environmental decline, bitter social divisions, and a rapid descent into a dismal new dark age. Instead of the flourishing of the better angels of our human nature, it will be our inner demons that technology magnifies.

We need to steer firmly away from courses of action which would precipitate any such outcome. We need to resist simplistic ideas or beguiling coalitions that would mislead us onto a slippery slope of accelerating humanitarian degradation. Instead, we need to select and uphold the set of priorities that will facilitate the timely emergence of sustainable superabundance.

These are tasks of the utmost importance. These tasks will require the very best of human insight, human strength, and human cooperation.

If we choose well, constraints which have long overshadowed human existence can soon be lifted. Instead of physical decay and growing age-related infirmity, an abundance of health and longevity awaits us. Instead of collective dimwittedness and blinkered failures of reasoning, an abundance of intelligence and wisdom is within our reach. Instead of morbid depression and emotional alienation - instead of envy and egotism - we can achieve an abundance of mental and spiritual wellbeing. Instead of a society laden with deception, abuses of power, and divisive factionalism, we can embrace an abundance of democracy - a flourishing of transparency, access, mutual support, collective insight, and opportunity for all, with no one left behind.

If we choose well, the result will be liberty on unparalleled scale. The result will be people everywhere living up to their own best expectations and possibilities, and then more. The result will be a transformed, improved humanity, taking stellar leaps forward in evolution, as technology increasingly uplifts and augments biology. The result will be to advance beyond mere humanity to transhumanity.

Time for action

As the 2020s approach, with their accelerating pace of change, with ever more potent technologies in wide circulation, and with a perplexing variety of tangled interconnections threatening unpredictable consequences, there's something that is very important for us to collectively keep in mind.

The thought that deserves sustained attention, in the midst of all other considerations, is the central insight of a group of people known as transhumanists. This insight concerns the magnitude of the forthcoming transformation from humanity to transhumanity.

It's not just that this transformation is possible. It's not just that this transformation could be relatively imminent - happening while many people alive today are still in the primes of their lives. It's that this transformation, handled wisely, could have an enormous positive upside. It's that this transformation, handled wisely, is deeply desirable.

This insight is grasped, today, by only a minuscule fraction of the earth's population - by a meagre sprinkling of transhumanist pioneers. However, it is time for these transhumanist pioneers to speak up. The few can become many.

As the 2020s approach, it is time for transhumanists to challenge and reorient the public

narrative. It is time to raise the calibre of the collective conversation about the future.

It is time to affirm how the future can be hugely better than the present. It is time to clarify how human nature is but a starting point for a journey to extraordinary posthuman capability. It is time to emphasise that, whereas the evolution of life has been blind for billions of years, it is now passing into our conscious, thoughtful control. It is time to point out that, whereas the evolution of society has been dominated for centuries by economic matters and struggles over scarce resources, centre stage can soon feature the blossoming of abundance. And it is time to proclaim that powerful enablers for these exceptional changes are already arriving, here and now.

In short, it is time for transhumanists around the world to step up to the responsibility as vanguards of the forthcoming transformation. It is time for transhumanists to inspire and support people everywhere to join together in the historic project to build the era of sustainable superabundance. It is time to apply transhumanist wisdom to identify and uphold the best choices in anticipation of the tumult of disruption that lies ahead. It is time to ensure that technology brings universal benefit, rather than being something we will come to bitterly regret.

At present, we can glimpse only the broad outlines of the coming era of sustainable superabundance. It is the fundamental responsibility of transhumanists to discern the forthcoming contours with greater clarity, and to help humanity as a whole envision and navigate the pathways ahead.

It is in service of this momentous cause that this Universal Transhumanist Manifesto is dedicated. This Manifesto is provided as an entry point to a growing network of living documents, with the intent to engage and organise researchers, creatives, entrepreneurs, activists, engineers, humanitarians, and more.

Together, let’s apply our skills, our time, and our resources to paint more fully the picture of sustainable superabundance. Let’s transcend our present-day preoccupations, our unnecessary divisions, our individual agendas, and our inherited human limitations. Let’s grasp the radical transformational power of new technology to profoundly enhance our vision, our wisdom, our social structures, and our effectiveness.

Together, let’s map out constructive solutions to the obstructions and distractions that impede progress - solutions combining the best of technology and the best of humanity. Let’s build productive alliances that weaken the forces resisting positive change. Let's prepare to take advantage of the growing momentum of an inspirational worldwide technoprogressive movement. Let's anticipate how we will dislodge the grasp on power held by today's backward-looking vested interests. And through an emerging shared understanding of the vital benefits transhumanist policies can bring to everyone, let’s transform fearful opposition step-by-step into willing partners. The few can become many.

In this way, we can accelerate the transition to sustainable superabundance. The sooner, the better.

2. Superabundance ahead

This chapter of the Manifesto briefly introduces the basic ideas of sustainable superabundance in each of seven spheres of human life: energy, nourishment, material goods, all-round health, all-round intelligence, all-round creativity, and all-round collaboration.

The good news is that twenty first century science and technology, developed and deployed wisely, can enable a sustainable superabundance of human flourishing in all seven of these spheres.

The bad news is that there is nothing inevitable about such an outcome. Science and technology, by themselves, cannot guarantee the emergence of sustainable superabundance. Instead, as Chapter 3 highlights, it will be imperative to pay close attention to critical elements of the broader social context within which science and technology operate - namely, politics, economics, and the public mood. Astute collective management of these elements will make all the difference between superabundance and super-tragedy.

Chapter 4 of the Manifesto follows on by addressing an even more important part of the context - the question of fundamental values and core priorities.

With the scene set, subsequent chapters - Chapters 5 through 11 - look again at each of the seven spheres of human flourishing listed above. These chapters indicate the pivotal overlaps between the different spheres: sustainable abundance in any one sphere is dependent on abundance in the other spheres too. These seven chapters also provide more details of the major risks and major opportunities associated with each sphere: the potential disasters, and the potential triumphs.

The Manifesto's final chapter - Chapter 12 - summarises a set of practical projects to reduce the likelihood of disaster and increase the likelihood of triumph.

An abundance of clean energy

The sun provides the earth with ample energy for all human needs for the foreseeable medium-term future. We can harvest that energy directly via sunlight, or indirectly via the motion of winds and waves. With care, we can also tap into the energy locked deep within atoms. The challenge ahead is to enhance methods to collect and store and transmit this energy, preserving the wellbeing of the environment at the same time as we nurture greater all-round human prosperity.

Present-day green technology points the way forwards, but much more awaits to be done. To allow humanity to swiftly complete the transition to a sustainable abundance of clean energy, innovation needs to accelerate. The hard task is how to facilitate such an acceleration, in the face of social and psychological structures that favour short-term thinking and the preservation of entrenched interests.

An abundance of food and water

Throughout history, agriculture has passed through a number of dramatic revolutions, including selective breeding, mechanisation, and synthetic fertilisers. Even greater revolutions are at hand, as gene editing and synthetic biology come of age, enabling a sustainable abundance of delicious, healthy, nourishing food. Meat grown in labs means we'll no longer slaughter animals on vast industrial scale.

In parallel, improved desalination plants can provide an abundance of fresh water.

An abundance of material goods

Atomically precise manufacturing, which extends to the nanotechnology scale the techniques of 3D and 4D printing, can drastically reduce the costs of numerous material goods, including clothing and shelter, whilst simultaneously increasing their quality. Innovation can identify alternatives for any rare elements that are approaching scarcity. If necessary, asteroids could be mined as an important transitional source of raw materials.

In the same way as bytes have become free of cost, we can anticipate atoms becoming free too. The material economy will follow the digital economy into sustainable abundance.

An abundance of health and longevity

The shortcomings of human biology, which cause us to suffer illness, decay, and death, are on the point of being comprehensively tamed by progress in regenerative medicine - nanosurgery, 3D bioprinting, genomic engineering, and stem cell therapies. Our task is to accelerate this relegation of aging and disease to history.

Much preferable to present-day expensive rear-guard medical treatments, prompt preventive interventions can unleash an abundance of full-health longevity.

An abundance of all-round intelligence

Science is giving us the ability to revitalise and rejuvenate, not only the body, but also the mind. Just as industrial tools have augmented our muscles, computers are augmenting our intelligence. Our growing understanding of the brain means we can enhance our psychology as well as our intellect. Alongside improved memory and enhanced reasoning ability, we can gain greater emotional and spiritual intelligence.

Traditional methods of mind enhancement such as education, meditation, yoga, music, art, and exotic substances, can have their power significantly magnified and directed by innovative technologies such as virtual reality, brain-computer interfaces, and AI assistants. An abundance of all-round intelligence beckons.

An abundance of creativity and exploration

In times past, human existence has been closely tied with our paid employment. The era of scarcity provided strong obligations to labour, often in work that was back-breaking or soul-destroying. As technology advances, machines can labour more forcefully, more reliably, and more effectively than human employees. Increasingly, human time and attention can be applied, instead, to an abundance of creativity and exploration - the unfolding and discovery of music, arts, sports, ecosystems, planets, and whole new universes.

The challenge ahead is to transition smoothly to a new economy which prioritises, not wage income or gross expenditure, but human flourishing and sentient development.

An abundance of collaboration and democracy

No mind is an island. We gain our strength and wisdom from our social relationships. The obstacles en route to the era of sustainable superabundance can be solved by the wise collaboration of many thinkers and doers.

Aided by emerging technologies that skilfully enhance positive cooperation, we can, together, make profound progress. Together, we can sagely analyse risks, assess scenarios, build bridges, deliberate options, reach decisions, deploy resources, progress actions, review outcomes, update processes, and consolidate our advances.

Our task is to ensure no-one is left out of the journey forwards. The benefits can be available to all, in ways that uphold freedom of choice and diversity of lifestyle. Our collaboration can take advantage of the best insights of all minds, and achieve the best results for all minds. In the process, we'll create and benefit from an abundance of democracy - a superdemocracy.

Time for action

There will be many bumps on the road to superabundance. Indeed, there will be sceptics and detractors in all walks of life who oppose even the idea of working towards sustainable superabundance.

In times of rapid change, it's no surprise that many people will become fearful and obstructive. Afraid of losing their status in society, they will cling onto outdated habits and structures. Afraid, with some justification, of technology going wrong, they will call for the imposition of overly cumbersome restrictions. Afraid that cherished human values may become lost, they will aggressively reassert inadequate bygone belief systems - belief systems grounded in incorrect or incomplete views of human nature and human flourishing. Lacking a vision of positive change from which they can benefit, they will deliberately sow confusion and misinformation.

To overcome confusion and misinformation, it's time to generate an abundance of understanding. To supplant cumbersome legal frameworks, it's time to champion smart, agile regulations. To quell panics about tech-driven dystopia, it's time to promote appreciation of scenarios in which technology uplifts humanity. To lessen the power of vested interests, it's time to build wise alliances. To tame the widespread fear of change, it's time to clarify the roadmap to sustainable superabundance - and to describe how, in that not-so-distant future world, everyone can attain greater security, greater opportunity, greater health, and greater wellbeing.

In short, to rise above the myriad distractions and obstacles that might frustrate the journey to superabundance, it's time to uphold a compelling, engaging picture of the remarkable future that is within our grasp.

Where that picture has gaps, let's address them, quickly and fully. Where questions arise, let's move fast to improve our collective understanding in the light of our collective intelligence. Where obstructionists and naysayers attempt to muddy the water, let's be fair yet firm in taking the conversation to a higher level.

Given the magnitude of the progress which can ensue, no tasks are more important.

3. Beyond technology

The project to advance into the era of sustainable superabundance depends critically on technology. However, it needs more than technology. Technology is not enough. That's for two reasons.

First, any given piece of technology is capable of misuse as well as use. Fire can burn down as well as provide heat. Metalwork can create swords as well as ploughshares. Chemistry can create explosive weapons as well as agricultural fertilisers. Biotech can create deadly new pathogens as well as nutritious new superfoods. Surveillance technology can be used to monitor and harass political opponents as well as to check for infrastructure weaknesses. Social media technology can distort and mislead as well as inform and educate. The more powerful the technology, the greater the potential for misuse - whether intentional or unintentional.

Second, before a piece of technology can live up to its positive potential, it needs to be developed and deployed. These steps can take a long, long time. These steps can be delayed - perhaps indefinitely - on account of all kinds of obstacles: intellectual obstacles, financial obstacles, institutional obstacles, philosophical obstacles, social obstacles, and political obstacles.

In short, just because it would rationally be the right thing for society to do, to support the development and deployment of a given piece of transformational technology, it's by no means inevitable that this implementation will happen.

What more is needed, beyond technology? If an earnest appeal to enlightened rationality is insufficient, what other measures can be taken?

This Manifesto envisions a vital role being played by a change in the public mood - as explored in the next chapter. This Manifesto also envisions vital roles being played by politics, by free market forces, and by targeted investment - not in their present forms, but in significantly improved forms. These are the subjects of the present chapter.

Beyond present-day politics

Politics is the management of power within a society. Politics is meant to provide citizens with liberty from being unduly manipulated by powerful elites or marauding scavengers. If we dislike how power is being exerted in society, we need to turn to better politics.

Good politics can enable and encourage patient long-term investment in beneficial technological developments. Bad politics prevents or discourages such investments from taking place. Good politics can ensure technological products serve the needs of all members of society. Bad politics acquiesces when such products serve only a narrow portion of society.

Alas, politics has often been a hindrance to positive technological progress. Politicians, wittingly and unwittingly, have imposed cumbersome legal restrictions on breakthrough innovations. They have elevated doctrinaire ideologues over evidence-minded pragmatists. They have re-routed funds from deserving causes to self-serving gravy train projects. And they have created flashy

distractions that divert public attention.

As tensions rise in the run-up to the 2020s, substandard politics poses increased risks of seismic social destabilisation - so long as politicians continue to pursue ill-advised policies, hold fast to outdated worldviews, and promote (wittingly and unwittingly) incivility and outrage.

But there's no inherent reason for politics to be so dysfunctional. We can, and must, do better.

When done well, politics is the mechanism for the democratic self-oversight of society. When done well, politics holds to account society's leaders and would-be leaders.

When done well, politics orchestrates collective action to prevent malignant subsets of society from exerting unwarranted control over the populace as a whole. These subsets, which are akin to potential cancers afflicting (if unchecked) the body of society, include corporate monopolies or cartels, banking dynasties, media tycoons and their empires, and "complexes" of overlapping business, military, and political interests. These subsets also include authoritarian politicians who seek to wield power freed from the overall checks and balances of democratic institutions.

When done well, politics involves wise, well-informed collective decisions about which new technologies and other social innovations should be restricted or steered, and which should be incentivised or encouraged. When done well, politics also ensures that such decisions are carried through, and are revised in a timely manner whenever necessary.

But if politics remains in its present dysfunctional state, all bets are off, regarding whether technology is deployed for the benefit of the few or the benefit of the many. All bets are off, regarding whether important safety considerations for disruptive innovations are recklessly sidelined or prudently reviewed. All bets are off, regarding which sets of interests dominate decision-making, and which priorities receive tangible support. All bets are off, regarding what kind of future will transpire - a future of human diminution and alienation, or a future of human flourishing and exultation.

We can, and must, do better. With the growing application of collective transhumanist intelligence, politics can become a powerful force for the collective good. Invigorated by the blossoming of superdemocracy (to be described below), politics can learn to hold obstructionist forces at bay. Public institutions, which have operated on far too many previous occasions for the benefit of just a narrow subset of society, can be reconfigured and revitalised to operate instead for the benefit of society as a whole.

Beyond present-day democracy

One major difficulty, of course, is in determining which actions will truly advance the collective good of society. To which reputed experts, or groups of experts, should we listen? When there are disputes and disagreements, should we give equal credence to every different opinion? Should we give greater credence to opinions that are expressed more forcefully? Or more eloquently? Or should a majority vote determine the best policies, with electoral popularity being taken as a paramount sign of worthiness?

As a method to make decisions, democratic voting is far from perfect. Simple votes of the

electorate suffer from a number of troubling drawbacks. Electors are often ill-informed. Having little incentive to research issues objectively, electors can be badly misled by misinformation. Out of a sense of group loyalty, electors may decide to turn a blind eye to various pieces of evidence. Due to the mechanics of "first past the post" vote-counting systems, electors may feel pressured to cast an inauthentic tactical vote for an option other than their own first choice, for fear that their first choice vote would be "wasted". In some cases, electors are bribed, directly or indirectly, to vote in particular ways. Once in power, political parties can stifle further discussion by portraying their electoral mandates as an inviolable "voice of the people", regardless if new information emerges that throws doubt on the wisdom of that choice.

In the light of such drawbacks, it is little wonder that some critics have called for a diminution of democracy - for a reduction of the reliance society puts into gaining electoral approval from voters.

Transhumanists understand the concerns these critics express. But transhumanists should not accept their recommendations. Rather than seeking *less* democracy, we should seek *better* democracy.

Specifically, consider the concept of “superdemocracy”.

Superdemocracy involves, not just a one-time simple vote, but an informed deliberation among electors before any top-level decision is taken. In such a deliberation, the most important insights should have a fair chance to rise to wider attention, rather than being drowned out or distorted (as frequently happens in present-day elections) by the loud voices of vested interests who find these viewpoints inconvenient.

In a superdemocracy, advice from relevant domain experts is valued and respected. Where someone can demonstrate that a particular political idea seems to stand in defiance of scientific principles or the principles of sound project management, that demonstration needs to be given careful attention. Such demonstrations should not be swept away under a carpet, regardless of the inconvenience or embarrassment they pose to prevailing popular ideologies. But whilst expert opinions should receive attention, their views should not *dictate* any decisions. Expert viewpoints will help guide the overall decision, but these viewpoints may well suffer from shortcomings and uncertainties of their own. After all, experts frequently disagree among each other, or speak over-confidently outside their own particular area of deep knowledge. In the end, it must be the electors as a whole who take decisions, directly or indirectly - not any technocratic elite.

In a superdemocracy, there is no rush to premature judgements. Superdemocracy extends from the personal level to the community level the shrewd remark of F. Scott Fitzgerald that the sign of a first-rate intelligence is the ability to hold two opposed ideas in mind at the same time whilst still retaining the ability to function. Superdemocracy resists any impulse to submerge minority dissident opinions; it retains these opinions in the overall collective intelligence of the electorate. Superdemocracy is not a system whereby the majority vanquishes the minority; it's a system where the valid observations and aspirations of minority opinions are respected - whilst society retains the ability to function well.

Indeed, in a superdemocracy, new ideas will emerge over time, integrating insights from positions that were previously opposed to each other. The result is no mere “average” of the previous viewpoints - some lowest common denominator - but a higher synthesis that arises from a dynamic, constructive conversation, and which attains wider buy-in as a result.

These improvements in political discussion will arise, in part from changes in style and process, in part from adoption of new technological tools, and, critically, from voters taking advantage, as individuals and also as groups, of transhumanist boosts to their all-round intelligence and all-round creativity.

Beyond lowest common denominator voting

A key concept of superdemocratic decision-making is that, as far as possible, the set of considerations pertinent to the decision should be made public, in ways accessible to the general population. This openness allows fuller scrutiny of the arguments taking place. It also encourages the introduction into the debate of a wider number of perspectives. With more minds tuned into the discussion, key issues and opportunities can be noticed more readily.

What's more, the scrutiny of these arguments can benefit, not just from contributions from multiple human perspectives, but from reviews carried out by increasingly capable systems of artificial intelligence. By analysing vast amounts of comments, these AI-powered reviews can highlight underlying tensions between different ideas, and suggest novel syntheses.

But not every decision requires the full attention of every citizen. Superdemocracy upholds the concept of delegated voting, via systems such as “liquid democracy”. This enables citizens to delegate their votes in specified areas of debate to people whom they trust in these areas. In case someone changes their mind, delegations can be revoked or reassigned at any time.

Liquid democracy is a tech-enabled improvement to those parliamentary systems in which a single elected member of parliament is meant to represent the voter in all areas of debate. With liquid democracy, representation is no longer an all-or-nothing affair. Accordingly, liquid democracy moves away from the unhelpful fiction that politicians are supposed to have been elected to carry out every nook and cranny of their election manifesto. It enables a set of approvals and affirmations that is much more fine-grained - an ongoing dynamic conversation with nuance and inventiveness.

There is no need for us to be stuck with a political system that too often panders to the lowest common denominator - a system where electors have to decide which option is "the least worst of a sorry assortment". Instead, we can look forward to revitalised political processes that generate novel compound solutions - solutions which voters from all sides can resoundingly endorse.

With voters more informed and more engaged, our politicians will be obliged in turn to become more informed and more responsive - responsive, not to manipulation by ulterior vested interests, but to the increasingly lucid voice of the citizenry. Better politics will arise in parallel with better voters.

To summarise, we should avoid overreacting to present-day examples of apparently irrational behaviours by individual voters. We should resist any attempt to diminish the influence of these voters over public decisions. Rather than seeking less democracy, we should demand better democracy. We should look forward to improvements in the reasoning capabilities of all voters - at both individual and group levels - and to a politics that is less confrontational and more creative.

Beyond right and left

Here's one important example where politics urgently needs to become less confrontational and more creative.

The journey to better politics involves respecting and integrating important insights from both the traditional right wing of politics and the traditional left wing of politics.

Traditional right wingers are correct to point to the many positive accomplishments of free markets, to mistrust the potential over-reach of politicians and career civil servants, to wish to uphold as much individual freedom as possible, to prefer to minimise undue state intervention, and to admire the marvels that can be achieved by competitive-minded self-made individuals.

Traditional left wingers are correct to point to the many positive accomplishments of the welfare system safety net, to mistrust the actions of profit-seeking corporations and financial speculators, to wish to uphold as much social solidarity as possible, to prefer to increase equality of opportunity, and to admire the marvels that can be achieved by collaboration-minded progressive coalitions.

Rather than a hostile battle between such positions, let’s ensure that a spirit of constructive exploration prevails. The goal is not the triumph of “our side”. It is the attainment of sustainable superabundance for all.

Two important specific areas where right wingers and left wingers tend to disagree sharply are the topic of redistributive taxation and the topic of political oversight of the free market. The question of taxation is addressed in Chapter 11 of this Manifesto, "Towards abundant democracy". The question of the operation of the free market extends over the remainder of the current chapter.

Beyond the free market

A competitive free market in goods and services often encourages significant improvements in the utility, attractiveness, performance, and affordability of these goods and services, in ways that benefit purchasers of these goods and services. Free markets have stimulated and facilitated remarkable innovation and enterprise. As an example, modern supermarkets are one of the marvels of the world, being stocked from the floor to the ceiling with all kinds of items to improve the quality of daily life. People around the world have access to a vast variety of all-round nourishment and experience that would have astonished their great grandparents.

However, there are circumstances in which markets cease to be open to new competitors, and in effect become cartels or monopolies. In these cases, when barriers to new entrants are too

high, free markets can no longer be relied upon to produce the best improvements in goods and services.

Free markets can also be distorted by the imposition of rules or standards that unfairly favour incumbent providers; in such cases, industry regulators are said to have been “captured” by vested interests.

What's more, markets often neglect to properly consider so-called “externalities”, such as impacts (either positive or negative) of products on the environment, public knowledge, public infrastructure, and public health.

For all these reasons, goods and services that deliver the highest short-term financial returns to investors aren’t necessarily those which would maximise increases in human flourishing. One example is that pharmaceutical companies often turn away from developing drugs for the “neglected” diseases that afflict only people in low-income regions of the world. Another example is that it can be more profitable to repeatedly sell people drugs that keep them in a state of semi-invalidity, than to develop a comprehensive one-off cure for their condition. Barriers to newcomers entering an industry can mean that incumbents avoid competitive pressures from would-be market disruptors.

Accordingly, let’s avoid raising the free market onto any pedestal in which it would be beyond criticism. Democratic supervision of the free market should seek to avoid any large negative effects of free markets, without undermining the positive capabilities of these markets.

After all, the marketplace is a kind of technology, and conforms to the general pattern of technologies, having both positive and negative potential. The task of gaining the positive benefits without a surfeit of negative consequences is far from simple. That task will require regular assessment and review, freed from ideological prejudice.

Beyond corporate financing

On many occasions, the goals of profit-seeking corporations align with the goals of accelerating human flourishing. But on other occasions, the goals diverge - especially when a project to improve an element of human flourishing would require large investment. Private financiers are, understandably, reluctant to undertake long-term, patient investment of risky projects which may provide them with little specific opportunities for direct commercial payback. The result is the “tragedy of the commons”: resources from which everyone would benefit, fail to receive the care, replenishment, or new financing they deserve.

For this reason, it has generally been public bodies that have led the way in investing in basic science and underlying technology. Initiatives such as the Manhattan project, the Apollo moonshot, the foundations for the Internet, the original network of GPS satellites, and the Drugs for Neglected Diseases initiative, would not have happened without strong sustained public coordination, championed by visionary politicians.

It’s the same with many of the projects needed to accelerate the technologies for sustainable superabundance. We cannot rely on Venture Capitalists to provide sufficient capital in such

cases. A more powerful coordination is needed - especially when a bold jump is required from an existing technological platform to a new one which will require considerable time to mature.

Such jumps often end up taking a lot more time and resources than initially imagined. A phase of slow, disappointing progress often precedes the eventual faster breakthroughs. That phase of disappointment can cause investors with a short time horizon to panic. In their panic, they urge engineers to concentrate on modest incremental developments rather than the more dramatic improvements previously considered. As a result, investment can become diverted into steps to "preserve the cash cow" rather than developing the next generation solution. Without patient, visionary understanding of the true potential of emerging technologies, the opportunity for radical progress will be lost.

Beyond predetermined exponentials

Unfortunately, the remarkable ongoing exponential improvements in some areas of technology, extended over many decades, risk misleading observers about what needs to take place in order for such progress to continue.

In the field of semiconductors, Moore's Law has operated since around 1959 to the present day, describing regular doubling of the performance of silicon integrated circuits. Cooper's Law describes similar doubling in the bandwidth capacity of wireless networks, going all the way back, remarkably, to the 1890s. Swanson's Law describes how the price of solar photovoltaic modules has halved on a regular basis since the 1970s. Likewise, the Carlson curve describes how the price of genomic sequencing has halved on a regular basis, also from the 1970s.

To an extent, these improvements are all driven by "learning curve" scale effects: the more experience engineers have with a technology, the more ways they can find to improve its performance. Positive feedback cycles prevail, in which innovations can layer on top of each other.

Nevertheless, periods of exponential improvement generally come to an end, once the potential for any one technological architecture has been fully utilised. There's nothing inevitable about the pace of progress remaining constant indefinitely.

Close observation of long-term ongoing improvements actually shows a combination of two separate effects: incremental improvements within individual architectures (platforms), and disruptive transitions from one platform to a next generation platform. Accordingly, an overall exponential curve is made up of an ascending series of 'S' curves. Careful inspection shows that the successive 'S' curves typically arrive with varying timing. Their occurrence is far from being predetermined.

The incremental improvements are relatively easy to finance, whereas the disruptive transitions are considerably more challenging. The former are safer investments, whereas the latter are riskier, involving unproven technologies. The latter may also require the development of new applications to take advantage of the potential latent within the new platform. What's more, new auxiliary services may need to be put in place - such as a network of electrical recharging stations to enable the electric car industry, or an application store (and associated certification

programmes) to enable the smartphone industry.

Depending on the particular industry, the costs of the jumps to new generations may exceed in magnitude what private financing is willing to bear. In such circumstances, public support becomes critical. Otherwise the overall progress stalls.

Deciding and then upholding the priorities for the deployment of public investment are two of the most pressing tasks that lie ahead. They are tasks that will require great collective intelligence. They are tasks that will require superdemocracy.

In particular, ensuring that the right priorities are *upheld*, in the face of persistent opposition from entrenched backwards-looking subsets of society, will require the strong popular support that superdemocracy can engender.

The next chapter of this Manifesto sets out the principles whereby these tasks can be steered and accelerated - principles which can engage and transform the public mood.

4. Principles and priorities

It’s time to explore more deeply the set of values underpinning the Transhumanist Manifesto.

This chapter sets out a number of fundamental principles, and reviews some high-level illustrations and implications. The seven chapters that follow apply and extend this set of values in each of seven spheres of human life where sustainable superabundance can bring profound transformation.

Statements of values are often viewed with suspicion. There's frequently a big discrepancy between a formal set of stated values, and the actual factors that seem to motivate people's behaviour. But on other occasions, values that are powerfully felt can, indeed, change the conduct of individuals and communities. They can inspire extraordinary courses of action. It depends in part on whether the values are embraced wholeheartedly, or merely given lip service.

The values set out below are intended to meet a number of criteria. They are intended to be credible - fitting with everything we believe to be the case about humanity and the universe. They are intended to be desirable - matching positive inclinations that we already find deep in our souls. They are intended to be actionable - they should inspire practical real-world changes in how humanity cooperates to accelerate the advent of sustainable superabundance.

The values are, moreover, intended to be universal - applicable to everyone, leaving no-one behind. They are intended to be liberating: by helping us keep the bigger picture fully in mind, they can liberate us from domination by the short-term mundane concerns of daily life that often impinge on us and diminish us. Finally, these values are intended to be liberating in the additional sense that, although they define constraints, they will enable greater diversity, greater creativity, and greater all-round human flourishing - sustainably.

Ten core principles

Society needs new top-level goals. Society should no longer prioritise above all else economics metrics such as the Gross Domestic Product or the Employment Rate.

After all, we humans cannot live by bread alone. Nor do we live just to work. These factors - the nourishment we consume, and the work we undertake - are means to an end, but are not ends in themselves.

Instead, here are a number of principles that merit being at the core of decision systems. To give them a name, they can be called "the ten core principles" or "the ten transhumanist principles".

First, the prioritisation of human flourishing: *prefer actions that lead to the increase of human flourishing*. Flourishing involves happiness, but there is more to flourishing than happiness. Flourishing involves energy and nourishment, but there is more to flourishing than energy and nourishment. Flourishing likewise encompasses but extends beyond creativity, intelligence, health, collaboration, and awareness. Over time, our understanding of the conditions and possible expression of human flourishing will surely evolve and improve. That's as it should be.

Second, the fundamental importance of human individuality: *individual flourishing should not be sacrificed or subordinated to collectivist goals*. Society should protect and elevate all members of society. Individuals should never become cannon-fodder in service of some tribal, national, ethnic, religious, or ideological quest.

Third, the principle of active neighbourliness: *treat others in the way we would ourselves like to be treated, if we were in the same situation*. Rather than keeping quiet about impending dangers about to befall someone, or major opportunities they are about to miss, we should find the way to speak up, just as we would ourselves like to be alerted to these dangers or opportunities in an equivalent circumstance.

Fourth, the generalisation of the previous principles beyond present-day humans: *prefer actions that lead to the increase of flourishing of consciousness*. To the extent that animal or artificial minds possess core attributes of consciousness, these minds deserve at least some of the same care and support as human minds. This care includes possibilities for growth and development, and the reduction in needless suffering.

Fifth, the generalisation to longer timescales, thereby highlighting sustainability: *avoid actions that reduce the possibilities for future flourishing*. Our plans need to enable, not only flourishing today, but also flourishing tomorrow (and the days and years that follow).

Sixth, the recognition that the future can be radically better than the present: *the present circumstances of humanity should by no means be regarded as the desirable pinnacle of evolution*. A wonderfully improved future lies ahead of us, provided we recognise that possibility, and take appropriate actions.

These six principles, as stated, leave many questions unanswered. They define a broad envelope that can accommodate a multiplicity of different viewpoints. That diversity is, itself, something to cherish. Hence a seventh core principle: *nurture and tolerate diverse opinions within the overall transhumanist framework*.

Here's an eighth core principle: *where different viewpoints within the overall envelope clash in terms of action to be taken, it is up to the community as a whole to deliberate and reach agreement*. This is where the practice of superdemocracy comes to the fore.

Next, the principle of preferring objective data: to help resolve clashes between different ideas, *priority should be placed on pursuing and publishing objective data relevant to decisions, rather than simply accepting the say-so of would-be authorities.*

Finally, as a tenth core principle: *in deliberations between conflicting insights, no book, thinker, or tradition should be given any absolute priority*. Society needs to remain open to current favoured ideas and methods being superseded. Of course, respect can be shown to books, thinkers, or traditions with good track records as sources of insight. But that respect should be tempered with caution. Runs of success can come to an end - especially in new circumstances or new contexts.

In summary, the ten recommended core principles are: human flourishing, individuality, neighbourliness, consciousness, sustainability, radical progress, diversity, superdemocracy, objective data, and openness. These principles complement and support each other. Together, they set the framework for humanity to advance into the era of sustainable superabundance.

Technocracy

To illustrate a number of the core principles just mentioned, consider the notion of technocracy - respect for decisions by domain experts.

Other things being equal, it’s sensible to pay attention to viewpoints from reputed domain experts. For example, in a sailing boat blown into unfamiliar turbulent waters by a storm, the recommendations of seasoned navigators deserve more attention than the opinions of a first-time sailor. For matters of an individual patient’s health, expert doctors are more trustworthy than lifestyle advice found in mass distribution horoscope columns.

However, all viewpoints should be subject to query and analysis. Experts are often wrong.

Moreover, the fact the someone is an expert in one domain does not entail any special priority applies to their viewpoints in other domains. An expert sailing navigator gains no authority in a different field, such as medical treatments, just by virtue of their sailing expertise.

As it happens, decisions frequently involve the intersection of several different domains. A decision that appears sound from one perspective may be recognised as inadequate when other perspectives are introduced. Listening only to experts from the first perspective risks reaching a bad decision.

Even when someone is an undoubted technical expert in a given domain, it's worth investing time and effort in explaining to the general public the reasoning behind their recommendations. Rather than being forced onto uncomprehending recipients, key decisions should be communicated openly and collectively understood.

Accordingly, there are major limitations to the concept of delegating hard decisions to domain experts. The ideal of technocracy needs to be subordinated to the ideal of superdemocracy - the involvement of the entire community in the process to reach decisions.

To be clear, in a superdemocracy, domain experts are respected and valued. Any society that ignores or denigrates the best insights of, for example, scientists, engineers, project managers, or change management experts, risks major failures in the initiatives it pursues. But whilst these experts should influence the decisions, they do not dictate any outcomes. The outcomes should be decided collaboratively, respecting the principles of openness and diversity.

Science

As another important example of the interplay between the core principles proposed above, consider science.

Science should be heartily applauded - not only for its many tangible accomplishments, but also for the merits of the methods it follows.

Indeed, the multiple beneficial technologies that will enable sustainable superabundance have arisen from wide adoption of the principles of the scientific community. That is, theories have been subjected to experimental analysis, findings have been published openly, agreement has been deferred pending replication and peer review, and advocates of positions have explored in advance which results would count as refutations (falsifications) and hence pose a real challenge to their favoured hypotheses. These are operating principles well worth upholding.

This is expressed in the core transhumanist principle of preferring objective data to subjective opinion.

However, let’s acknowledge there are questions which science, by itself, cannot answer. Although science can answer the question, what’s a good way to accomplish goal X, it may not be able to say whether goal X is itself desirable.

Let’s also recognise that scientific knowledge is provisional and incomplete. Advice on non-fat diets is one example where scientific orthodoxy has significantly changed. To be scientific is to acknowledge that scientific theories can change.

Again, methods that make good sense in some fields of science - such as double blind trials, in which experimenters aren't aware which members of a sample are receiving a given treatment, as opposed to a placebo - aren't always feasible in other fields of science.

In short, whilst science is one of the key tools used by humanity to advance towards sustainable superabundance, science in itself is not the end goal. Nor is the scientific method the only tool in our toolkit.

Accordingly, while championing the scientific method, we should resist the siren pull of scientism (the exaggeration of the capabilities of particular scientific methods). Our preference for objective data needs to be complemented with the principle of openness.

Transhumanism

Whereas science can tell us what's a good way to accomplish goal X, we need to look beyond science to decide which goal X is worthy of pursuit.

It is transhumanism that provides the answer, with its vision of the profound ongoing elevation of all-round human health, human wisdom, human wellbeing, and human freedom.

Transhumanism comprises a set of philosophies of life that (to refer to a 1990 definition by philosopher Max More) "seek the continuation and acceleration of the evolution of intelligent life beyond its currently human form and human limitations by means of science and technology, guided by life-promoting principles and values".

In order for us to achieve the superabundance envisioned in this Manifesto, there will first have to be significant improvements to human nature. Superabundance will not feature present-day humans in a new environment. It will feature humans enhanced in fundamental ways, freed from core defects that have hitherto limited our accomplishments. It will feature humans on a transhumanist journey towards transcendent posthuman capabilities.

The word "transhumanism" sometimes provokes negative reactions. However, it is likely that the word will become increasingly mainstream.

One reason people have expressed doubts about transhumanism is because of the perceived similarities between transhumanism and religion.

Indeed, transhumanism can be seen as a kind of fulfilment of religion - in the sense of providing an overriding vision, credible in the modern day, that will inspire greater social harmony and positive community endeavour.

However, transhumanism is distinguished from traditional religion by having all its viewpoints open to questioning and updating. There are no inviolable canons of belief or "holy books" in transhumanism. In other words, transhumanism is firmly rooted in the principle of openness.

Religion

Throughout history, religion has operated in ways that are both positive and negative.

Religion has often oppressed people, conducted witch hunts and inquisitions, limited personal choices, forbade critical thinking, imposed outdated conceptual frameworks, and caused widespread psychological suffering.

Religions deserve to be opposed if they agitate against the best insights of scientific enquiry, such as evolution through natural selection, or the benefits of vaccinations in providing immunity to many diseases.

Religions should also be opposed if they champion rules or practices that are strongly anti-humanitarian - for example, if they forbid the education of girls, prevent divorce, promote female genital mutilation, or regard as a capital crime homosexuality or apostasy (someone becoming a non-believer).

Society should resist any attempts by religious groups to penalise “blasphemies” such as criticism of central members of a religious tradition. There is no right not to be offended. Discussion needs to remain open.

At the same time, religion has often provided people with an important sense of purpose and community. Religion has often encouraged people to behave in morally positive ways.

Transhumanists can find many points of mutual support with religious adherents who avoid the adverse tendencies mentioned above. For example, transhumanists share with many religious adherents the goal to uphold human dignity and human flourishing, and to act responsibility (as a “steward”) towards the environment.

Religious adherents who are motivated to transcend the limitations of human nature - such as aging and mortality, as well as the propensity to behave badly (“sin”) - can find inspiration in the transhumanist mission to apply science and technology to abolish aging and to enhance human nature. These adherents can therefore see transhumanism as at least part of the culmination of their own religious aspirations.

Singularity

History has featured many phases, with significant jumps in human capabilities between phases. Major transitions have been given names such as the cognitive revolution, the agricultural revolution, and the industrial revolution. We are presently experiencing an information revolution.

Each time human capability grows, our potential increases to change the earth, in both positive and negative ways. Earlier human migrations resulted in large scale extinctions of other animal species. Human activities are now impacting the atmosphere as never before. Careless use of weapons of mass destruction could result in the end of human civilisation - perhaps even the end of all human life.

The accelerating pace of the information revolution has led to the suggestion of a forthcoming "Singularity" in which change happens more quickly than ever before. Rather than waves of new technologies taking decades or even centuries to reach mass adoption, forthcoming disruptions might drastically alter society's processes in years, months, or even weeks, days, or hours. This pace, encouraged by fierce competitive pressures, and enabled by self-updating automated processes which bypass the operation of slow human review, may well result in changes that no-one has been able to properly anticipate and evaluate in advance.

The notion of Singularity alarms some critics, who feel uncomfortable with the resonance with religious notions such as "the end of days", the apocalypse, and the messianic establishment of paradise on earth. Rather than being distracted by ideas with religious connotations - or ideas

churned over by Hollywood blockbusters - it's better, say these critics, to concentrate on shorter-term real-world issues.

Transhumanists respond that it is better, instead, to keep an open mind.

Indeed, if society concentrates just on the risks and opportunities of the present-day, it may miss the larger risks and opportunities that, timewise, are just around the corner. And we may miss the possibility of making sufficient preparations to steer the forthcoming Singularity so that it results in sustainable superabundance rather than a much bleaker outcome.

Moreover, there will be other benefits from humanity developing skills and processes to help us steer a particularly rapid technological transition like the Singularity. These skills will increase the likelihood of us being able to steer other technological transitions of the coming years and decades, that may be less intense than the Singularity, but still far more disruptive than previous changes.

Exponential urgency

While transhumanists highlight the possibilities of large changes being “at hand” and “soon”, we cannot forecast any precise timing. Nevertheless, there are credible future scenarios in which these changes take place by the middle of this century. Perhaps sooner. Certainly within the lifetimes of many people presently alive.

Factors that make it plausible that radical changes may take place so quickly are the acceleration of developments of technologies such as AI and regenerative medicine - developments powered in turn by activity worldwide of unprecedented numbers of scientists, engineers, designers, entrepreneurs, educators, and social activists.

This activity is further boosted by positive feedback cycles: tools that improve tools, computers that improve computers, software that improves software, and AI that improves AI. Again, better technology increases the power of educational systems (including YouTube videos), and better educational systems increase the throughput of capable engineers, designers, entrepreneurs, and systems integrators - who, collectively, can develop technological solutions of even greater utility. Again, better technology improves communications networks, allowing for a richer flow of ideas between technologists around the world, which in turn accelerates the creation and deployment of innovative products. The cycle continues.

Yet other factors causing increasingly higher levels of research and development than before are the enormous commercial and military advantages that can be gained by the groups that are the first to achieve key breakthroughs.

For as long as positive feedback systems remain in place, the result is exponential growth. In reality, many progress curves take a 'S' shape rather than an unending exponential curve: an initial phase with slow growth transitions into a phase with faster growth, but then is followed by another phase of slower progress, as the potential of a particular technological architecture is exhausted. However, if the underlying market conditions continue to favour improvements - for example, if there continue to be major commercial gains from hardware and software becoming

more powerful - then a number of different 'S' curves can arise, one after another, each being based on a new architecture or paradigm. Thus in the world of computing hardware, the architecture of vacuum tubes was superseded in turn by architectures involving single integrated circuits, massively parallel integrated circuits, and cloud computing. The individual 'S' curves combine into overall exponential progress covering a longer period of time.

For as long as conditions remain in place that encourage ongoing exponential progress in technology, we need a sense of exponential urgency - an imperative to improve our ability to anticipate scenarios, ahead of our human support systems being overwhelmed by unforeseen consequences of technological change.

Technological determinism

Talk of long-standing positive feedback cycles might suggest that the progress of technological development is somehow inevitable or predetermined. But that would be a big mistake.

Instead, the development and deployment of technology is significantly influenced by a wide range of non-technological factors, such as design, legislation, user expectation, public zeitgeist, and random unpredictable events. Accordingly, transhumanists should resist any ideology of technological determinism.

Principles such as “Moore's Law” may describe general technological trends at a first level of approximation, but can mislead observers who fail to pay sufficient attention. Closer study shows that these principles are subject to variation in matters such as timing for performance to double (e.g. 12 months, 18 months, or 24 months), the meaning of what is delivered (e.g. number of transistors vs. computing power), and the market adoption of the underlying technological capabilities (e.g. CPUs being displaced for some tasks by GPUs and by cloud-based computing).

Instead of sharp predictions of individual dates when future events will take place, what is more credible is a probabilistic prediction covering a range of dates and outcomes. For example, instead of saying that artificial general intelligence will assuredly cause a Singularity to occur in the year 2045, a better prediction is to say there's a 50% chance that such an outcome will happen by that date - and also a 10% chance that it could happen by (say) 2025, and a 90% chance by (say) 2085.

And instead of any predictions that events will somehow "inevitably" result in a victory for transhumanist forces, what is much more credible is a probabilistic prediction covering a range of different impacts on the future of human flourishing.

That's why transhumanists repeatedly highlight the importance of human volition and human action. Transhumanists understand the need to go beyond cheering from the sidelines - the need to develop and apply policies that have a real impact on actual social change.

Techno-optimism

Optimism generally leads to more personal energy than pessimism, so the former is to be preferred to the latter on that score. However, any optimism that is naive about potential

drawbacks can lead to greater problems.

Rather than the description “techno-optimism”, transhumanists would prefer to be known as exemplifying “optimistic techno-realism”.

Whilst transhumanists anticipate the possibility of many wonderful consequences of technology, they are also aware that there could be terrible consequences as well.

The optimism of the transhumanists is grounded in sober appreciation of real-world issues and challenges. Rather than ignoring these challenges, transhumanists will formulate and evaluate potential coping strategies - strategies to manage the risks involved. When a risk is judged too severe, transhumanists will take actions to avoid the risk altogether.

Precaution and proaction

In determining priorities, both the precautionary principle and the proactionary principle have their place.

The precautionary principle is appropriate when there are credible suggestions of huge negative consequences of some action. We need to beware unintended runaway consequences of well-meaning actions.

The proactionary principle points out, on the other hand, that abstaining from action can have huge negative consequences as well. To adapt traditional language, there are “sins of omission” as well as “sins of commission”. Rather than any blanket abstention from actions which have associated risks, it is often better to develop plans to manage these risks.

For example, it may be argued that nuclear energy has the potential to give rise to radioactive waste that could contaminate huge biological ecosystems. The precautionary principle urges in that case to shut down all nuclear power plants. But the proactionary principle observes that wider adoption of nuclear energy might make all the difference in ramping down the use of carbon-based fuels quickly enough to avoid runaway global warming. In that case, it’s worth putting the precautionary stance temporarily on hold, while methods are reviewed for responding promptly to any leakages of radioactive waste. Far better to calmly assess the various probability estimates involved in these discussions, than to give absolute precedence to either the precautionary or proactionary principle.

In many cases, a better principle than precaution is reversibility. Action that is risky should be undertaken in ways that allow reversal, in the event that matters develop badly.

A commitment to reversibility requires effective monitoring, and avoidance of any inertia that would overwhelm attempts to change course. It also requires the emotional intelligence that is willing to admit and experience failures, and to learn lessons from these failures. It’s an approach that requires a higher calibre of execution than would a reflex application of the precautionary principle.

This may seem like a tall order. Are society's leaders really capable of operating with a higher calibre? This is by no means clear. In the interest of minimising risk, it may appear better to

adopt policies that can be implemented more straightforwardly.

However, any strategy of attempting to always pick low-risk options is itself highly risky. Mediocre leaders will make mediocre choices. Major opportunities will be missed. Political systems that lack capability are more likely to be subverted by powerful vested interests.

Instead, the strategy that has the lowest risk overall is the transhumanist project to significantly improve the collective capability of society's leaders. That project will be difficult, but can be achieved by drawing on the best that humanity has to offer.

Diversity and inequality

One source of greater overall capability is when society includes multiple diverse opinions and outlooks, from which new insights can be formulated and integrated.

This is in line with the core transhumanist principle that people's differences should be valued. Transhumanists see no reason to enforce uniformity. On the contrary, transhumanists champion greater choice, not only over lifestyle and thinking approach, but also over bodily form ("morphological freedom").

A matter of real concern, however, is if people are left behind against their will, in matters of opportunity, such as lacking access to resources needed for personal growth and development. That is no longer a diversity to be championed. That is an inequality to be addressed.

Another concern is when the rewards from some joint activity are systematically captured disproportionately by one of the parties in that activity - the party with most power - resulting in growing inequality of opportunity.

There are three reasons to work to reduce inequality of opportunity. First, each individual should recognise that their own circumstances may change, due to factors outside their own control, and they could cease to be a "winner" from current economic transactions and become instead one of the "losers". Second, even those who are presently well off should fear social chaos arising from the disruptive activity of people who perceive themselves unfairly treated by society. Third, in line with our commendable natural instincts of "active neighbourliness", we humans are predisposed to prevent one another from needless suffering.

In short, the transhumanist commitment to greater human flourishing implies as a consequence a commitment to enable people to overcome circumstances of unequal opportunity. A society that disregards this commitment is a society that will grow weaker and less capable. A society that disregards this commitment is storing up greater problems for the future.

Diversity accelerating

Transhumanists oppose practices such as racism or ultra-nationalism that view members of specific ethnic groups as intrinsically inferior. This condemnation is compatible with the recognition that some genetic endowments increase skills or abilities in given areas, such as endurance in long-distance running. These variations do not cause any change in the intrinsic value of the people involved.

But the application of transhumanist technologies in the years and decades ahead will increase the diversity of human attributes - for example, enabling even greater endurance in long-distance running, or better memory and mental processing of information. Might this growing transhumanist diversity rupture the wholeness of humanity? What will happen to democratic ideals such as "one person, one vote" when some people have enhanced various of their attributes tenfold, one hundredfold, or more? Might these variations lead to a fragmentation of humanity into multiple castes? Is the potential for such a rupturing a reason to ban technologies of human enhancement?

As before, these changes should cause no change in the intrinsic value of the people involved. However, the increased diversity will give rise to a need for overall governance mechanisms that are more complex than before.

Groups of people who share particular enhanced skills and modes of practice will, understandably, seek some autonomy over decisions within their groups, freed from requirements for democratic approval by people in the wider community that have little understanding or interest in these modes of practice. This is similar to the principle of technocratic decision-making: there are domains of specialist knowledge (for example, medicine) in which decisions are best taken by the relevant experts rather than by a vote that includes non-experts.

Nevertheless, domains often interact with each other. Where the activities of one group of people, with one set of enhancements, interact with the activities of other groups of people, a broader democratic agreement needs to be reached.

The design of the overall transhumanist society therefore needs to enable the prosperous coexistence of subgroups with significantly divergent skills and practices. This is an extension of present-day society, which already supports peaceful coexistence of subgroups with different interests and aspirations.

Coexistence

It is one thing to consider the coexistence of diversity within the overall transhumanist framework. What about coexistence with groups (or nations) that reject one or more of the proposed core transhumanist principles?

For example, some groups may reject the principle of superdemocracy, preferring a governance system with a dictator (presumed to be benevolent). Other groups may reject the idea of respecting non-human minds, and may treat primates, dolphins, and other intelligent species with greater brutality than transhumanists would expect. Yet other groups may put conformance to particular religious scriptures at the very centre of their decision-making processes. Groups might also reject the transhumanist project of working towards the transcendence of present-day limits on human nature - limits such as our tendency to become old and die, our tendency towards groupthink and the confirmation bias, our tendency towards the abuse of power, and so on.

The transhumanist answer is: tolerance within limits, coupled with ongoing respectful advocacy

of the merits of the transhumanist worldview.

Where groups are carrying out practices that transhumanists judge as abhorrent - such as capital punishment for "apostates" who turn again the religious beliefs of their parents - these groups can expect sanctions and other restrictions on trade.

Consciousness beyond humans

Transhumanists look forward to greater flourishing, not just of human minds, but of all human-like minds. We look forward to the flourishing of consciousness, not only in humans, but in non-humans.

For example, it is understandable that people will seek to “uplift” their cherished pets, so that these pets acquire greater health, longevity, intelligence, and wellbeing. The same principle applies as when uplifting humans, namely to avoid unbalanced developments that would actually lead to a reduction in flourishing.

Moreover, in order to eliminate the need for the industrial-scale slaughter of farm animals and fish, the development of lab-grown meat should be accelerated.

Beyond the question of the suffering inflicted on animals by humans, there is the question of the suffering that animals inflict on each other. Nature “red in tooth and claw” involves horrendous amounts of distress as animals hunt and devour prey, and as parasites destroy their animal victims from within. Re-engineering natural ecosystems to avoid such suffering will be far from straightforward, but is one of the great causes to which we transhumanists can commit ourselves.

To the extent that AIs acquire consciousness, they too deserve rights. Today, there is no reason to pause before terminating a piece of software - that software has no inner life. In the future, things may not be so clear.

The consideration of non-human minds, such as uplifted animals and conscious AIs, further increases the levels of societal divergence and coexistence we need to anticipate. As such, the challenges of designing the overall society increase. However, we can also anticipate that the greater collective intelligence available will provide the capability to manage these challenges.

Beyond hubris

Talk of re-engineering natural ecosystems to avoid needless suffering - a project sometimes called “paradise engineering” - brings forth accusations of hubris. Transhumanists are reckless and naive, critics say. Transhumanists should “stop playing God”.

Transhumanists answer this last rebuke with the response “we’re not playing” - our intent is extremely serious.

It’s a commendable part of human nature to seek to do better than our human nature. Throughout history, this impulse has led to remarkable advances in medicine, engineering, the arts, and more. That's not something frivolous. That's something outstanding.

The plea "stop playing God" can be compared to the plea, heard in different cultures, not to intervene in perceived processes of karma. Where children are born disabled, that line of thinking insisted that it must be due to karmic retribution for misdeeds in a past live. That line of thinking urged a hands-off approach: we humans should not be so arrogant as to question or oppose the operation of the inscrutable mechanics of divine reward and punishment. If a child is suffering miserably, there must be good reason behind it. Thankfully, that kind of fatalistic thinking has increasingly been discarded, and human prowess in medicine is actively restoring victims of suffering to fuller states of flourishing. And thankfully, the similarly fatalistic plea to "stop playing God" is likewise increasingly being questioned.

However, transhumanists do acknowledge that there are risks of unintended consequences from the application of technology. Examples include dreadful side-effects of drugs such as Thalidomide and Vioxx, the environmental impact of the pesticide DDT and, more recently, problems posed by huge quantities of discarded plastic. Other examples are the way in which widespread access to social media has fanned the spread of fake news and polarisation, and the unforeseen biases latent in some of the algorithms introduced into automated decision processes.

As noted earlier, there are also risks that the greater diversity of human and posthuman lifestyles will challenge the overall wellbeing of human society.

Transhumanists are by no means blind to these risks. But nor are we blind to the tremendous opportunities latent in technological innovation - further remarkable advances in medicine, engineering, the arts, and more.

Facing this dichotomy, transhumanists apply the proactionary principle: risks should be identified, reviewed in advance, and managed wisely. Contorted fear of risk should not be permitted to befuddle careful deliberation or paralyse positive action.

No evidence has been presented so far that any such risks are incapable of solution. To the extent that specific risks seem particularly worrying - as with the proliferation of weapons of mass destruction - society can and should take stronger measures in response, exerting (where needed) political pressure over commercial market forces.

Taking back control

Are some aspects of technological development, contrary to what has just been said, already beyond control? Has human experimentation with transforming nature already gone too far?

For example, many observers are alarmed by the seemingly uncontrollable rise in greenhouse gases in the atmosphere, risking runaway global warming and greater instances of extreme weather. This would be an example of technological development - namely, the technology of extraction of fossil fuels - running contrary to sensible humanitarian control.

However, to anticipate the discussion from the next chapter, transhumanists point to the potential of next generation green energy to hasten the switch to non-carbon energy sources. The technology of CCS (Carbon Capture and Storage) can also be accelerated. In this way, the

adverse effects of one generation of technology can be undone by the positive effects of a later generation.

To the oft-spoken response that dysfunctional economic and political systems are preventing sufficient social focus on accelerating the requisite technological transitions, transhumanists foresee transformations in the operation of economics and politics - transformations from systems operating dysfunctionally to a system of superdemocracy.

The effort required for these transformations to take place should not be underestimated. These transformations will be among the most difficult in the history of humanity. But these transformations can be guided by the greatest collective intelligence in the history of humanity, and empowered by the huge positive psychological energy awakened by the compelling vision of sustainable superabundance.

For the details, read on.

5. Towards abundant energy

The foundation for all human activity is energy. Without energy, nothing can be accomplished.

As human activity has grown in scale, we have utilised increasing amounts of energy, in ways that some would describe as bold and ingenious, but others would say is reckless.

We have converted vast forests into firewood. We have unearthed and burned immense quantities of peat, coal, gas, and oil. The resulting light, heat - and air conditioning - have provided ample illumination and kept citizens at comfortable temperatures. Our factories have been enabled to manufacture countless goods, and our vehicles to crisscross all over the earth. However, side-products of all this activity have been accumulating in unsustainable ways. Greenhouse gas emissions have been amassing in the atmosphere, and now pose a number of potential drastic threats to human flourishing.

To solve these threats, do we need to cut back on human activities? Should we adopt low-energy lifestyles?

The transhumanist answer is that there is no need to slam on the brakes. However, significant steering is overdue. We need to transition to a different trajectory. Urgently.

Indeed, as this chapter explores, an abundance of clean energy awaits us, ready to power productive, exuberant lifestyles. That's provided we have the strength of purpose to quickly switch away from our present near-addiction to unclean energy.

Anticipating climate chaos

The future is arriving faster than used to be expected. Likewise, climate change could be arriving faster than used to be expected.

In both cases, a dangerous heritage of complacency needs to be overcome - the complacency that says nothing much will actually change any time soon, so it's OK for the time being to continue with "business as usual".

In both cases, we need to shake off the complacency. We ought instead to heed the advice of Amara's Law: whilst we should beware overestimating the effect of a technology in the short term, we should also beware underestimating the longer-term impacts that can ensue once the adoption of the technology has reached its stride. A period of gradual build-up can tip over into a period of turbulent disruption. A period of apparent calm can morph seemingly overnight into a period of chaos.

In both cases - the case of general future disruption, and the specific case of climate disruption - it is compound effects that prove hardest to anticipate and hardest to manage. Complications arise from self-reinforcing feedback cycles, from crossover effects, and from the destabilisation of previous patterns.

In both cases, the possibility of acceleration in the pace of change increases the urgency for society to exert wiser, firmer control over aspects of collective human behaviour.

In the specific case of climate change, the possible acceleration of extreme weather events increases the urgency for society to navigate away from patterns of living that result in dangerous levels of greenhouse gases in the atmosphere.

This is no mere future forecast. Parts of that future have already arrived. Long-lasting heatwaves are breaking records around the world. Rainfall is absent for weeks on end. The resulting wildfires that spread swiftly over tinder dry landscapes prefigure larger catastrophes to come. The people near Athens in Greece who jumped into the Aegean sea to escape rapidly engulfing conflagration were the lucky ones. Others were left stranded with no route to safety, and perished while huddled together in a last embrace. There was no "business as usual" awaiting them.

The periods of unanticipated droughts in some parts of the world are accompanied by periods of unexpectedly heavy downpours of rain elsewhere. Warmer oceans not only expand in volume and threaten coastal erosion, but also stir up larger scale hurricanes. In the wake of greater storms, infrastructure weakens, dams rupture, avalanches of mud cascade downhill, and whole villages are washed away.

That's not all. Distortions to prevailing atmospheric currents such as the Jet Stream can bring unprecedented cold as well as unprecedented heat. Whilst the average global temperature is on an upward trajectory, regions of countries can plunge to record lows. In populations unprepared for the havoc of lengthy blizzards, bitter chaos ensues.

People in numerous countries, seeing the effects of severe weather on local crops, and anticipating worse to come, are setting out as climate refugees on arduous journeys towards lands which seem more fertile. These waves of migrants are stacking up waves of conflict and resentment. The US Department of Defense warns of climate change as acting as a set of "threat multipliers that will aggravate stressors abroad such as poverty, environmental degradation, political instability, and social tensions – conditions that can enable terrorist activity and other forms of violence". Trouble will beget further trouble.

Taking climate seriously

Although the direction of travel is clear, it remains uncertain exactly how quickly increased greenhouse gases could trigger deeply damaging climate change.

The basic physics is well understood. It's been known for more than a century how greenhouse gases can trap more of the sun's energy and raise average global temperatures. But the dynamic heat circulation mechanisms within the earth's overall climate systems are fiendishly complicated. Different experts make different forecasts about future impacts, and express different levels of confidence about these predictions.

Emphatically, this level of uncertainty is no reason to relax. As a matter of prudence, scenarios in which drastic changes could take place within just a few decades need to be taken seriously.

These runaway scenarios feature adverse positive feedback cycles, the destabilisation of long-established current patterns in oceans or the atmosphere, and increased chaos from extreme weather events. For example, hotter temperatures reduce the amount of ice cover, which reduces the amount of sunlight reflected back into space, which, in turn, further increases the temperature. And long-buried methane gases which are being exposed by the melting of Siberian tundra, may quickly add to the quantity of airborne greenhouse gases, ratcheting temperature gains even further - in turn melting more tundra and causing even more long-buried methane gases to be released.

This threat goes beyond the possibility of mere linear changes in temperature. Increased heat could spark a comparatively sudden phase change in the earth’s climate, pushing up the global average temperature by several degrees in less than a decade.

Similar changes have taken place in the past. Around 11,500 years ago, in a transition known as “the end of the Younger Dryas”, temperatures rose by 10°C within a single decade. This abrupt jump in temperature has become known from study of ice cores extracted from Greenland, and has been verified from data from lake sediments elsewhere in Europe. Much further back in geological history, an episode some 251 million years ago, at the end of the Permian era, has become known as “the great dying”, since 95% of all marine species and 70% of all terrestrial vertebrate species suffered extinction at that time. As such, this considerably exceeds the extinctions experienced by dinosaurs and others 65 million years ago. Although not the only contender as an explanation, a front-running theory for the cause of this "great dying" calamity is a sudden increase of temperature, of around 6°C.

Smaller calamities could prove disastrous in their own way, via the "threat multiplier" mechanisms. Social unrest that can (just about) be contained at the present time, may become completely unmanageable in the context of greater damage being inflicted regularly by adverse weather on agriculture, transport, and other key aspects of daily life. It is said that every society is only four square meals away from revolution and anarchy. That's not a theory we should be in any hurry to test.

With the prospect of a range of calamities ahead of us, "business as usual" cannot continue.

Technology is not enough

The good news is that a number of technologies to systematically reduce the threat of damaging climate change are on the point of being developed and applied. The price of energy from wind, wave, and solar has been dropping steadily, decade after decade. New designs can improve capacity as well as drive down costs even further. After all, more than enough energy reaches the earth from the sun in just a few hours, to meet the needs of entire human population for a whole year. In principle, all that's needed is to accelerate improvements in the harvesting, storage, and transmission of energy from renewable sources.

The bad news, however, is that the pace of implementing improvements is currently far too slow. Metaphorical mountains still need to be climbed. It's not just that the generation of electricity needs to swap over from carbon-based to clean mechanisms. We also need widespread reforms of other economic activities that are collectively responsible for more than fifty percent of greenhouse gas emissions - activities such as farming, transport, and the manufacture of steel and cement. Another complication is the shortage of the raw materials needed in increasing quantities in the construction of ever larger numbers of solar panels, wind turbines, and other generators of clean energy.

Accordingly, political action to accelerate the transition is needed as a matter of high priority. This action includes significant subsidies for next generation green technologies - including next generation systems for energy storage and energy transmission, as well as mechanisms such as "artificial photosynthesis" to create fuels from sunlight. It also includes the reduction of subsidies (direct or indirect) for activities that increase greenhouse gas emissions. Finally, this action also includes the imposition of taxes on such activities - taxes that scale up over time, in order to increase the incentives favouring cleaner modes of operation.

Such actions will face trenchant opposition from the companies and organisations who benefit from the status quo. This opposition cannot be overcome by friendly rational persuasion alone. Other sorts of forces will need to be applied in parallel - including economic forces, legislative forces, and a transformed public zeitgeist.

Steering short-term financials

One reason why fossil fuels continue to be used so widely is because of the short-term economic benefits that accrue to powerful corporations from the development and usage of these fuels. Many jobs are said to depend on these corporations continuing on their present trajectories. And many pension funds are heavily invested in the share price of these corporations.

Accordingly, many people are predisposed to latch onto any viewpoint they hear that minimises the need for disrupting the main activities of these corporations.

An example of such a viewpoint is that there is uncertainty over climate science. Another example is the idea that future technological solutions will be able to be introduced, in due course, to mitigate the results of climate change. What these viewpoints have in common is the basic claim that, whilst climate change is a matter of some concern, it doesn't require any urgent policy changes. Any major policy changes can be delayed to a future occasion, by which time

new technologies will be more advanced. As it happens, such a delay will be convenient for the managers whose short-term bonus payments depends on the fossil fuel companies continuing along "business as usual".

The fundamental problem with these viewpoints is that new technologies can take a long time to reach sufficient maturity. What's more, the core research needed in order for these technologies to be developed will take place only if it receives strong funding - funding that is presently lacking.

What can cause the owners and managers of fossil fuel corporations to rethink their priorities?

One consideration is if a precedent is established of companies being sued for damage that can be attributed, as a matter of probability, to their operation. This is similar to the court cases in which tobacco companies were sued for increasing the chances of cancer among smokers. Insurance and reinsurance companies may be very interested in this potential way of recouping the growing expenses they incur due to extreme weather events. Fossil fuel companies would need to start setting aside very significant sums of money in anticipation of such court cases.

Consider also a potential decline in market demand for fossil fuels, caused by a hike in the price of such energy due to carbon taxes (and drops in the price of energy from alternative sources).

Finally, another consideration is that the share prices of fossil fuel corporations may soon start to plummet, in what is known as the bursting of the carbon asset bubble. This will happen when investors increasingly reach the conclusion that, because of potential future legislation or the imposition of significant carbon taxes, many of the assets on the books of these corporations will prove to be unsellable. Much of these oil reserves will end up having to stay in the ground. As investors anticipate this outcome, they will start to sell their shares in companies dependent on fossil fuel energy production.

Alongside these “stick” approaches to changing the operations of energy companies, “carrot” approaches are critically important too too. This includes incentives for energy corporations to grow units dedicated to quicker transition to greener sources, and compensation paid from funds raised by carbon taxes to former employees as they are made redundant.

A battle of ideas

There's another factor leading many people to oppose any policy actions regarding climate change. This factor is the ideology of anti-centralism - the ideology which harbours deep suspicion of any attempts at centralised control over market forces.

Their reasoning runs as follows. Any action against climate change will need to be global in scope. After all, the fossil fuel industry operates transnationally, and could circumvent any carbon taxes levied in just a single country. However, any mechanism of global coordination would put too much power into the hands of a single political organisation. Sooner or later, any such organisation would impose legislation that stymies innovation and freedom. Any such organisation would grant itself more and more authority. As such, any benefits arising from it would be more than offset by the tyranny of centralised control. Too much local sovereignty

would be lost. That's too large a price to pay.

In summary, these thinkers fear the risk of global autocratic government much more than the fear the risk of accelerated climate chaos.

People who embrace this anti-centralist view are predisposed - like the managers and owners of the fossil fuel companies - to find and champion arguments as to why there is no urgency to tackle climate change. They eagerly spread such arguments, even without being sure if they are valid. What matters in their mind is that the arguments sow doubt on the case for centralised control of global society.

In response, transhumanists should accept that there are, indeed, potential dangers in centralised control. History shows many cases of leadership cabals that started off serving wider interests, but whose focus narrowed over time to their own self-preservation.

However, there's no inevitability of a slippery slope from international agreements to a centralised international autocracy. Transhumanists can point to the possibility for international agreements to be developed and adopted, without losing multi-level democratic control over the matter.

Rather than predicting that international agreements are bound to fail - or that they can only work if backed up by some kind of centralised police force - transhumanists should emphasise that more nuanced mechanisms are possible.

Consider international agreements on matters such as air traffic control, aircraft safety, wireless spectrum allocation, emissions of ozone-depleting CFCs (as governed by the Montreal Protocol), and action in the face of disease epidemics (as coordinated by the WHO). Consider also the operations of bodies that coordinate world sports competitions.

What ensures that such global agreements can be developed and then observed (rather than ignored) is the force of international public opinion, expressing itself through local political structures.

In other words, instead of their ideology of anti-centralism - and instead of its polar opposite, which could be called "autocratic centralism" - transhumanists can envision the practice of superdemocracy being meaningfully extended from the national to the international level.

As the merits of the concept of superdemocracy become better appreciated, the passion that many people currently feel for the ideology of anti-centralism will diminish. They will increasingly appreciate that centralisation of powers has no need to be an "all or nothing" decision. Their fear will diminish of unwarranted loss of local sovereignty. Their hostility will decline towards the idea of negotiating treaties with international scope - treaties that will ensure companies pay substantial but fair taxes on greenhouse gas emissions. And their willingness will grow to look rationally and objectively at proposals for serious action to prevent climate disruption.

Beyond greenwash

Another factor that has delayed positive action against climate change over recent decades - a

factor in addition to the short-term financial incentives of fossil fuel corporations, and the ideology of anti-centralism - is the habit of politicians to adopt "greenwash".

This occurs when politicians appear to talk tough about taking action on climate change, but then make little real progress. Rather than confront the forces of inertia that support the status quo, they perform a kind of green theatre, hoping to gain some electoral benefit as a result. In many cases, their actual intentions probably never amounted to much.

Transhumanists look forward to greater publicity being given to all divergences between the claims and the actual performance of politicians. Politicians - and likewise for any other leaders in society - need to be held accountable to higher standards of integrity. Mechanisms of better collective intelligence will highlight such discrepancies. In the face of improved transparency, society's leaders will be more likely to move beyond the surface statements of greenwash towards deeper policy reform and substantive investment in next generation green technologies.

A role for nuclear energy

Different experts evidently disagree whether adoption of non-nuclear renewable energy sources, such as solar, wind, and wave, will be sufficient to reduce greenhouse gas emissions sufficiently quickly to avoid major climate chaos. In view of this uncertainty, it’s worth keeping an open mind towards the potential of nuclear energy.

Among other considerations, nuclear energy avoids the problems of irregularity in supply experienced by solar plants, which depend on the sun shining, and wind turbines, which depend on the wind blowing. Nuclear power plants can run 24 hours a day, every day of the year, regardless of the weather.

On the other hand, nuclear power plants have acquired a bad public reputation in the wake of major problems at Fukushima, Chernobyl, and Three Mile Island. There are also fears about crossovers from nuclear energy technology to nuclear weapons technologies. It is important to assess these risks calmly.

As in many other areas of technology, the most significant improvements in nuclear power production will involve jumps between different generations of underlying technology - moving, for example, from water-cooled systems to gas-cooled systems and molten-salt reactors. Next generation solutions have the potential for significantly greater safety and lower costs. Designs involving modular construction should have faster deployment times, as the modules can be built in a factory, rather than at the reactor site.

Other ongoing innovations within the nuclear industry are addressing the problems of nuclear waste (for example, by burning it) and possible crossovers to nuclear weapons (by destroying the chemical isotopes that would be used in nuclear bombs). Critically, there seem to be good solutions at hand to the risk of reactor meltdown. Molten salts, unlike water, have no risk of evaporating under greater heat; as the salts become hotter, they expand, and force radioactive elements further apart, automatically slowing down the reaction. In consequence, there is no longer any need for large, expensive containment domes, which reduces the overall cost of the reactor.

Fission systems involving the element thorium, instead of uranium, also seem to have considerable promise.

Bearing all these possibilities in mind, the risks of inaction - that is, avoiding progress with nuclear energy - probably significantly exceed the risks of careful positive action.

Next generation solutions, however, typically require long, patient investment, to turn concepts from theory and prototypes into large-scale deployment. Some advocates of solar and wind energy see any investment in these new nuclear power solutions as a distraction - as a poor use of scarce public funding. It's a controversial question.

When evaluating funding choices, it's important to keep emotional biases under check. One such bias, already mentioned, is concern about reactor meltdowns - a concern which has been unduly heightened due to sensationalist reporting of previous meltdowns, and which has little direct relevance when new technological designs are used.

Another bias goes back to the period of time, in the 1950s and 60s, when proponents of solar and wind power believed that public funding of nuclear power plants was starving their own projects of an appropriate level of funding. Hostility between some advocates of solar and wind, and advocates of nuclear, remains in place to this day. Whereas a certain level of intellectual combat helps provide better insight, too much of a conflict hinders clear thinking.

An area of particular controversy involves, not nuclear fission, but nuclear fusion. In theory, fusion has enormous potential. To give a comparison, whereas the US economy uses each day the energy from several supertankers full of oil, less than one thousandth of a single supertanker containing fuel for nuclear fusion - namely isotopes of hydrogen - would provide enough energy to run the US economy for an entire year. However, nuclear fusion projects have suffered repeated delays. These delays are due not only to formidable engineering problems, but also to issues of large-scale international coordination. This is an example where greater skills in collaboration are highly applicable. It's also an area where nimbleness of thought is important, given the number of unorthodox nuclear fusion research projects that are seeking additional funding.

In summary, the best way to make rapid progress towards a sustainable abundance of clean energy involves careful, dispassionate analysis of controversial options. By the nature of research projects, it's not possible to anticipate in advance what the full impact of various funding provisions might be. Progress could go better than expected - or, indeed, worse than expected. Adoption of fast-improving artificial intelligence in the design of new systems could result in surprising breakthroughs. Accordingly, we need to be ready to update our opinions.

A role for geoengineering

As a fallback option, in case the switch to clean energy sources proves to be taking too long, some technologists are investigating geoengineering solutions such as spraying particles into the atmosphere or oceans with a view to increasing the reflection of sunlight and thereby reducing global temperatures.

By their nature, geoengineering solutions risk triggering cascading side-effects of their own. It’s possible that the results of such interventions will be far from uniform: some parts of the earth may be cooled, whilst others experience greater extremes of weather than at present. That's in line with observations after "natural geoengineering" effects such as recent major volcanic eruptions, like the 1992 explosion of Mount Pinatubo in the Philippines.

Geoengineering solutions also risk a rapid resurgence of global warming if, for whatever reason, they are turned off or malfunction.

In other words, there are significant questions over the effectiveness and safety of geoengineering. It is nevertheless appropriate to keep an open mind, and to continue exploring variants of such solutions, in case accelerating climate change proves too hard to address by any other means.

However, the first line of defence should be on the rapid transfer away from the types of energy usage (and other industrial practice) that give rise to large greenhouse gas emissions. The second line of defence is to improve systems to extract greenhouse gases from the atmosphere - systems such as Carbon Capture and Storage, or the introduction of "artificial trees" (which may be more effective than real trees at absorbing carbon dioxide).

Solutions that perturb the climate in different ways, such as the mass distribution of particles in the oceans or atmosphere, should form part of a third line of defence, with particular focus on gaining more certainty about the safe introduction of such solutions.

Beyond the profit motive

Climate change due to greenhouse gases is only one of a number of potential environmental catastrophes that are on the point of being accelerated by unsustainable human practices.

Others include ocean acidification, excess accumulation of nitrogen and phosphorous due to the methods of large-scale agriculture, erosion of topsoil, depletion of freshwater resources, and loss of biodiversity.

In each case, the pattern is the same. Methods are known that would replace present unsustainable practices with sustainable ones. By following these methods, life would be plentiful for all, without detracting in any way from the potential for ongoing flourishing in the longer term. However, the transition from unsustainable to sustainable practices requires overcoming very significant inertia in existing systems. In some cases, what’s also required is vigorous research and development, to turn ideas for new solutions into practical realities. Unfortunately, in the absence of short-term business cases, this research and development fails to receive the investment it requires.

In each case, the solution follows the same principles. Society as a whole needs to agree on prioritising research and development of various solutions. Society as a whole needs to agree on penalties and taxes that should be applied to increasingly discourage the unsustainable practices. And society as a whole needs to provide a social safety net to assist those peoples whose livelihoods are adversely impacted by these changes.

Left to its own devices, the free market is unlikely to reach the same conclusions. Instead, because it fails to assign proper values to various externalities, the market will produce harmful results. Accordingly, these are cases when society as a whole needs to constrain and steer the operation of the free market. In other words, transhumanist politics needs to exert itself.

To explore one particularly important example, the next chapter of this Manifesto considers the question of the future of agriculture - the sustainable creation of sufficient food and drink for everyone on the planet, even as the population rises significantly above its present level.

6. Towards abundant food

How many people can the earth accommodate, providing everyone with good quality food and water? Are we near the limit, or might we have passed it already? Alternatively, is that limit located far above the present population size?

Transhumanists envision the quality of life increasing all over the world, at the same time as the global population continues to rise. Wise management of technological innovations can enable a sustainable abundance of numerous kinds of healthy nourishment, with plenty available for everybody. But in the absence of careful forethought and some hard decisions, such an outcome is far from inevitable.

Accordingly, this chapter of the Manifesto highlights a number of key scenarios for the future of the production of food and drink, and the risks and opportunities en route.

Population, onward and upward?

The global population passed the landmark of seven and a half billion towards the end of 2016. How it reached that huge figure is a huge story in its own right. Until around 1800, the global population remained less than one billion. Within 127 years, that is by 1927, another billion was added. It took just 33 years, until 1960, for the population to grow to three billion. The next billion was added in just 14 years - by 1974. The next three billions were added in, respectively, 13 years (to 1987), 12 years (to 1999), and another 12 years (to 2011).

Extrapolating current demographic trends would suggest that the population will reach eight billion in 2023, nine billion in 2037, and ten billion in 2055. Of course, that extrapolation assumes only modest changes in the current rates of births and deaths. However, transhumanists anticipate radical improvements in healthcare that will significantly reduce death rates around the world. If this happens, the population is likely to rise more quickly. Rather than increasing at the present rate of around 220,000 people each day, it could increase at around 350,000 people each day. Rather than it taking 12 years to add another billion to the population, this could happen in just 8 years.

What's more, transhumanist technology such as ectogenesis - the ability for a baby to develop outside of a mother’s body, in an artificial womb - might impact birthrate in various ways. In some projections, the population could rise by a lot more than the figure of 350,000 per day just noted.

As well as considering the sheer number of people alive, we also need to consider how many

resources (including energy, food, and water) each person consumes. As larger proportions of the population become more affluent, and adopt so-called “western lifestyles”, the total resources used by humans will grow more quickly than the population count.

The organisation Earth Overshoot Day regularly carries out calculations comparing the demands of the population to the capacity of the planet to regenerate resources. The supply side of this calculation estimates the planet’s biologically productive areas of land and sea, including fishing grounds, cropland, grazing lands, and forests. The demand side estimates demand for livestock, fish products, plant-based food, timber and other forest products, and so on. The result for 2018 was that the demand exceeds supply by a factor of 1.7. Stated in other words, by 1st August 2018, the human population had already consumed more of nature than the planet can renew in an entire year. Accordingly, the 1st of August is dubbed “Earth Overshoot Day” for 2018. It is said that, if everyone around the world adopted the same lifestyle as people in the USA, Overshoot Day would be 15th March.

If matters continue unchanged, this state of affairs seems unsustainable. It would appear that overfishing, over-harvesting of forests, and overuse of land, should be a cause for real concern.

Indeed, there are reasons to fear potential sweeping unwelcome side-effects from agriculture becoming overly dependent on new chemical treatments and new genetic manipulations. Larger and more mechanised doesn't necessarily mean more resilient. Biochemical innovations can have longer-term consequences that weren't evident from short-term trials. The real world is a much messier, more complex place than a carefully controlled research laboratory.

And there are reasons to fear that the pursuit of increased profits by powerful agrochemical corporations will result, not in the feeding of the world, but in the unintentional poisoning of the world. Just because a product makes good short-term financial sense for a company and its investors, that's no guarantee of a positive longer-term effect on human well-being.

The legacy of Malthus

Some observers dismiss the calculations from the likes of Earth Overshoot Day. These calculations are said to stand in a long line of discredited forecasts of ecological doom and gloom.

The line of discredited forecasts is said to include the predictions of British cleric Thomas Malthus, who in 1798 theorised about hard limits on the growth of the human population. Malthus believed that faster population increase would result in famine and starvation, or in other harsh mechanisms to correct the population size. Specifically, he forecast that, on account of constraints in improvements in food production, population growth could never exceed one billion in any period of 25 years. Food production methods could only increase linearly, Malthus maintained, and could not keep up with the tendency of population to increase exponentially.

Malthus had some notable predecessors, including, sixteen centuries earlier, the early Christian writer Tertullian based in Carthage, North Africa. Tertullian complained about the "teeming" numbers of inhabitants he observed, as being "burdensome to the world" which could "hardly support" everyone. That was at a time when the world's population was less than 200 million.

The line of discredited forecasters also includes, more recently, US professor Paul Ehrlich, who in 1980 agreed a scientific wager with another US professor, Julian Simon. Ehrlich forecast that, between 1980 and 1990, there would be large price increases for each of five metals: chromium, copper, nickel, tin, and tungsten, as an indication of greater resource scarcity. The wager reflected Ehrlich's deep apprehension about rapid population growth exceeding possible growth in the supply of food and resources. In reality, Simon won the wager handsomely. All five prices fell significantly over the ten year period, with the prices of two of the metals (tungsten and tin) falling by more than half.

Nevertheless, we should be cautious about any simple extrapolations. The fact that Ehrlich and, before him, Malthus, were proved wrong in their forecasts, is no basis for complacency about the ability of humanity to keep on finding ways of safely extracting more resources from the planet. As transhumanists know well, accumulated exponential changes can give rise to unexpected transitions. Periods of slow change can be preludes to periods of disruptive upheaval. Predictions - whether of flourishing or of collapse - can be dismally wrong, for many a season, before becoming dramatically correct.

Indeed, analysts who have studied the famous Ehrlich-Simon wager have concluded that Simon was lucky rather than prescient. In many other ten year periods, both before and after their wager, the prices of these raw commodities have indeed risen (in real terms), indicating growing scarcity.

As it happens, Simon lost a subsequent wager, this time involving another professor, David South, concerning whether timber prices would rise or fall over a five year period. Simon forecast they would fall, but instead they rose 50% in real terms.

Nor were the predictions of Malthus as flawed as some critics like to claim. It took more than 150 years after his ideas were first published, before there was an increase in population by one billion period within any one 25 year period. Hundreds of millions of people continued to live squalid, miserable lives, in a state of deprivation and near starvation, before succumbing to wretched deaths - just as Malthus observed.

Rather than dismissing the concerns of Tertullian, Malthus, Ehrlich, and more recent writers on impending resource scarcity, we need to engage thoughtfully and constructively with these observers. After all, there are downsides as well as upsides to the biochemical innovations that are intended to provide food to more people than ever before. As in many other areas covered by this Manifesto, the underlying theme is "technology is not enough".

Necessity and innovation

Optimists like to proclaim that necessity is the mother of invention. As demand for resources increase, the free market encourages and rewards the innovations that answer these demands.

There are five problems with this viewpoint. First, just because there is a strong demand for the answer to a shortage, it doesn't follow that an answer will emerge. Vast numbers of people throughout history died, or had their lives horribly stunted, for want of better food or other resources. For these unfortunates, necessity was insufficient.

Second, if an answer is to emerge, it may need the kind of patient long-term investment which the free market, by itself, is incapable to organise and deliver. Not every problem can be solved by the heroic efforts of dynamic, ambitious, optimistic technology startups. In these cases, entrepreneurial flair is insufficient, even when coupled with a burning necessity.

Third, when answers do emerge, they may have deleterious side-effects. Foods that take away the pangs of hunger may lead to addiction, or obesity, or other health issues. Weed killers that allow crops to grow unimpeded may have adverse impact on local wildlife. Residue from these weed killers may accumulate in human bodies with longer-term consequences. In these cases, inventions can be double-edged swords.

Fourth, it's not in the interests of corporations to draw attention to the adverse side-effects of their biochemical innovations. To maintain high profits, these corporations are naturally inclined to employ people to divert attention away from risks, and to throw doubt on any scientific consensus which would question the safety or efficacy of their products. Consider Big Tobacco, assiduously creating confusion about the extent to which smoking might cause cancer. Consider Big Oil, assiduously creating confusion about the extent to which the burning of oil might cause global warming. It's the same with Big Agrotech. In these cases, the double-edged nature of an invention is deliberately hidden or obscured.

Fifth, these same corporations are motivated to publicise products and services which can be cleverly marketed as "healthy" and "desirable", rather than products and services which are objectively healthy and desirable. In this case, it's not necessity that is the mother of invention. Instead, it's the profit motivation. To express it better: the necessity that drives the development and deployment of these solutions, is the perceived necessity of increased short-term profits, rather than the goal of increasing overall human flourishing. That variance in motivation can give rise to an outcome that is very different from the one expected.

In praise of biochemical innovation

Before raising more specific concerns about the misapplication of biochemical innovation, it's important to recognise the enormous accomplishments and potential of this technology.

In lists of all the remarkable inventions made throughout history, ranked in terms of how many people's lives the various inventions saved, one breakthrough often stands alone at the top of the list. This is the Haber-Bosch process, in which a metal catalyst, high temperature, and high pressure, are used to cause atmospheric nitrogen to react with hydrogen to create ammonia (chemical formula NH3).

Because ammonia is the basis of artificial fertilisers, this invention enabled the radical transformation of agriculture. Farmers were no longer dependent on the natural fertiliser of excrement from farmyard animals, or on guano - accumulated bird excrement imported in bulk from far-off lands. As a result, agricultural productivity rocketed upwards. The same area of land could grow a lot more crops than before and, therefore, feed many more people. For a time - before everyone became used to this chemical miracle and took it for granted - the process was nicknamed "bread from air".

The basics of this process were worked out in 1909 by German chemist Fritz Haber in his laboratory at the University of Karlsruhe. The following year, another German chemist, Carl Bosch, succeeded in converting the process to large-scale industrial level. As a reward for their ingenuity and skill, the two inventors each earned a Nobel prize for chemistry.

The innovation of the Haber-Bosch process arrived in time to avoid impending problems with worldwide shortages of guano. Instead of waiting on nature to deliver more fertiliser, human engineers were able to build vast factories, usually in remote locations, that produced immense amounts of artificial fertiliser. By some estimates, the innovation saved the lives of approaching three billion people, due to improved farming efficiencies.

Another set of inventions that significantly boosted food production was due to US agricultural scientist Norman Borlaug, who is known as "the father of the green revolution". During the 1940s and 1950s, Borlaug time and again produced new varieties of wheat, with greater yield, higher resistance to infection, and the ability to grow in environments that were too tough for other varieties.

Borlaug achieved his results by thoughtful breeding and hybridisation techniques, coupled with changes in farming practice. His ideas often contradicted local conventional wisdom and were, for a while, resisted, until the bountiful crops produced by experimental deployment were too large to be ignored. In 1970 Borlaug was awarded the Nobel prize, not for chemistry, but for peace - in recognition of the huge humanitarian benefits of the increased food production his innovations enabled. Calculations have suggested that, without these innovations, today's world population might be reduced by at least one billion. Many other people would be alive but significantly malnourished, compared to today's actuality.

Impressive as they are, the accomplishments of Haber, Bosch, Borlaug, and many other innovators of the last century, are on the point of being put into the shade by even greater improvements in agricultural productivity from 21st century science and technology. Chief among these is synthetic biology - the creation of brand new variants of life by direct genetic manipulation. This includes the cutting and splicing of genes from one organism to another. Hence the phrase "GMO" - genetically modified organism. GMOs have the potential of greater disease resistance, longer storage periods for produce after it has been harvested, crops that can grow in previously barren landscapes, and individual crops having extra nutritional benefit - such as including much-needed supplementary vitamins or minerals.

The first major success of GMO techniques was in a different field from agriculture, namely medicine. People with insufficient insulin suffer and generally die fairly quickly from diabetes. From the 1920s onwards, diabetics could benefit from injections of insulin derived from animals, such as pigs or cows. But from 1982, a form of synthetic insulin became available. This product, named Humulin, was manufactured inside e-coli bacteria that had been genetically edited to contain the gene for the human form of insulin. Humulin transformed the treatment of diabetics, since it could be manufactured at scale more cheaply than insulin from the pancreas organs of slaughtered animals, and since it was less likely than animal insulin to be rejected by patients.

Developed in the laboratories of Genentech, at that time a relatively small company, Humulin

was the world's first approved genetically engineered therapeutic. Genentech has grown since that time into a biotech giant. Other genetically engineered products from the company's laboratories include therapeutic antibodies to assist in the treatment of cancer. Newer techniques for genetic editing, such as CRISPR-Cas9, enable more precise usage by lab technicians without detailed scientific knowledge, and are set to increase the speed of innovation yet further.

The first item of food whose manufacturing was transformed by the application of GMO techniques was cheese. Hard cheeses, such as Cheddar and Parmesan, depend on enzymes called rennet being added to liquid milk. For thousands of years, rennet was obtained by killing young calves and extracting it from the linings of their fourth stomachs. As demand for cheese grew, scientists looked for a method to create key rennet enzymes in the laboratory. By the late 1980s, a team at Pfizer had succeeded, with a process involving yeast genetically edited to contain a particular cow gene. Today, upwards of 80% of the cheese in many countries is manufactured using enzymes created by GMOs.

Genes extracted from a number of different organisms have subsequently been used to create new variants of major crops such as soybeans, maize, potatoes, and cotton. In some cases, the new gene enables the crop to tolerate the chemical glyphosate, as contained in the weedkiller Roundup. In other cases, the crop incorporates genes from the bacterium Bacillus thuringiensis (Bt), reducing the need for the crops to be sprayed with insecticide. Genetically modified papaya, containing similar genes, saved the papaya industry in the 1990s from devastation by the papaya ringspot virus.

Many other products remain at a trial stage. One example is so-called "golden rice" which includes a gene that triggers the creation of vitamin A. This addresses a nutritional deficiency that is widespread in developing countries: shortage of vitamin A weakens the immune system and causes more than half a million cases of childhood blindness each year.

More waves of innovation ahead

At least three more waves of biochemical innovation can be foreseen. The first involves so-called vertical farming. Rather than crops growing only at ground level, numerous layers of crops can grow, one on top of each other, inside skyscraper versions of greenhouses. Light coming directly from the sun cannot illuminate all these different layers, so the likes of LEDs will be needed instead. Hydroponic techniques allow crops to be grown without use of soil; instead, the crops are nourished via mineral nutrient solutions. Control of the environment eliminates the need for pesticides or herbicides. Robots and other automation can take care of maintenance. Since the crops grow indoors, dependencies on weather are reduced. Genetic modifications of crops can ensure that crops will thrive in this new setup.

A second forthcoming wave involves the industrial-scale creation of meat in the laboratory, without any need to grow and slaughter whole animals. Lab-grown meat already exists, but is presently too expensive for mass adoption. It also falls short of animal-grown meat in taste terms. Both these challenges are on the point of being solved by ongoing research and development. As an important consequence, it will be possible to repurpose enormous amounts

of land which are currently dedicated to the rearing of farm animals for slaughter, and to the growing of crops to serve as food for livestock. As another important consequence, there will be a significant reduction in the contributions from farm animals to the accumulation of greenhouse gases in the atmosphere.

A third forthcoming new wave of biochemical innovation involves artificial photosynthesis - methods to convert sunlight, water, carbon dioxide and other simple molecules into complex carbohydrates, without any reliance on living plants.

In parallel with innovations that enhance food production, innovative desalination techniques are enabling larger scale conversion of salt water into fresh water. These techniques are increasingly cost-effective and, compared to earlier methods, have fewer adverse side-effects on the environment. Studies are taking place into ways of adding in extra quantities of minerals such as magnesium that are present in rainwater and other natural fresh water, but which are lacking in water created by desalination. Whilst more work remains to be done, this series of innovations is addressing concerns about shortages of natural fresh water.

Feeding one hundred billion people

How far could biochemical innovation take us?

An average healthy diet consists of around 2,500 calories per day for a man, and around 2,000 calories per day for a woman. Let's take 2,250 calories as the mean of these two figures. In energy term, one of these calories amounts to 4,200 joules. Dividing by the number of seconds in a day, the average healthy person's diet is equivalent to a power consumption of around 109 watts (that is, 109 joules per second).

How much agricultural land is required, to create a sequence of food delivering 109 watts of power? Let's consider the rate at which energy reaches the ground from the sun, and then let's consider what kind of efficiency is possible in principle, in converting energy from the sun into energy stored in food.

Energy from the sun strikes the upper atmosphere of the earth at a rate of around 174 petawatts, that is, 174 million gigawatts. Around 30% of that is reflected back out to space. By the time the remaining energy reaches ground level, each square metre of land receives on average around 240 watts of energy from the sun.

Agricultural plants typically have an efficiency of only around 1% at converting incoming sunlight into energy stored in the edible parts of their bodies - although some crops such as algae or sugar cane can reach higher figures such as 3.5%. Sticking for the moment with the figure of 1%, produce from a land area of around 50 square metres would be sufficient in principle to feed a single person.

It is believed that hunter gatherer societies roamed over an average of around one square mile of land to feed each person. That's more than 2.5 million square metres per person. By the medieval age, farming techniques had decreased the average amount of land required to feed each person by a factor of around 100. By modern times, there have been additional

improvements of around six fold - hence the oft-repeated statement that it takes about one acre of land (around 4000 square metres) to feed a person. That's still around one hundred times less efficient than the figure worked out above.

It's hard to know how much of an increase in productivity can be expected from biochemical innovations and other improvements in farming in the years ahead. For the sake of a definite calculation, let's assume that another ten-fold increase might be achieved, bringing down the land required to feed each person to 400 square meters on average.

The world is presently said to have 14 trillion square metres of arable land, out of a total of 49 trillion square metres devoted to all sorts of agriculture. In turn, that's a subset of the 149 trillion square metres of the total land area of the earth. Let's assume that, over time, 40 trillion square metres can be farmed by the kinds of enhanced techniques mentioned earlier in this chapter. Dividing 40 trillion by 400 yields a population of 100 billion people who could be fed by these methods.

Changing some of the assumptions in the calculation will produce different answers. For example, the area of the earth used for food production could be significantly extended by growing more crops in sea water. Another point to note is that the efficiency of artificial photosynthesis has already reached 22% in laboratory conditions, so the energy conversion figure of 1% used earlier could be far too low. Equally, there might be complications working in the other direction. The passage of more time should, however, allow innovations to keep on building on top of innovation, so that any problems which emerge can be addressed. Accordingly, there seems no reason in principle why the world cannot feed a population of at least ten billion in the relatively near future, and of the order of one hundred billion some time later this century.

Of course, before such a large population can arise, many other aspects of human life and human society will need to change. This includes a full transfer to green energy, in order to avoid intolerable quantities of pollution as side-effects of fossil fuels. That theme was covered in the previous chapter. Also needed will be transformations in the way society handles rare earth elements and other raw materials, as included in the equipment and the housing for next generation agriculture. That theme is covered in the next chapter. And as reviewed in the remainder of the present chapter, it will also be necessary to manage the roll-out of biochemical innovation, in a way that avoids major drawbacks.

Risks from biochemical innovation

For each of the examples of praiseworthy applications of biochemical innovation described earlier, there are serious questions about potential misuse of the same (or closely related) innovations.

For example, the invention of synthetic ammonia by Haber and Bosch led not only to new types of agricultural fertiliser but also to new ways of manufacturing explosives. Ammonia formed part of a reactive pathway that resulted in nitrates as used in First World War bombs. Military historians have suggested that, were it not for the Haber-Bosch process, Germany would have run out of explosives after just one or two years of warfare. Additional years of battlefield

carnage could have been avoided.

Fritz Haber had a different idea in mind, in his quest for the First World War to end quickly, this time via decisive battleground victories for Germany. Using the same set of outstanding skills in chemistry that had enabled the synthesis of ammonia, Haber led teams to produce and then deploy poisonous chemical gases - including chlorine, phosgene, and mustard gas. But rather than shortening the war, they escalated the level of inhuman cruelty.

On hearing reports of the use of chlorine gas, Haber's wife became distraught. A skilled chemist herself, she understood the grim implications of this deployment. In a state of distress, she shot herself with her husband's military revolver, and died shortly afterwards. Undeterred, Haber continued his research into poisonous gases, even after the end of the war. In the 1920s, scientists at his lab developed the formulation of cyanide gas which subsequently found horrific use in Nazi concentration camps. As a chilling example of unforeseen consequences, several members of Haber's extended family met their death in these gas chambers.

There was no similar overt sinister intention behind any of the "green revolution" work undertaken by Norman Borlaug. Instead, the question that deserves attention is the extent of the risks in the wide adoption of a crop monoculture. A single crop that grows well in a range of types of land may become planted in place of a large number of other variants. However, if that crop falls prey to a fast-spreading pathogen that is particularly deadly to this new crop, devastation may follow.

Here's an example. The terrible tragedy of the potato famine in Ireland in the late 1840s was due to the exclusive adoption of one variety of the crop, known as the lumper. Although the lumper had the advantage of low cost, and could thrive in the wet climate of Ireland, it turned out to be particularly vulnerable to an airborne fungal potato blight. Lacking any resistance to the blight, potatoes all over the country turned into an inedible black pulp. Mass starvation ensued. In a few short years, the population of Ireland declined by around 25%.

In more recent times, fungal pathogens with names such as "wheat leaf rust" and "southern corn leaf blight" have laid waste to large areas of farmland planted with monocultures of wheat or corn.

A different kind of adverse health impact from wheat monoculture may be the growing prevalence of gluten intolerance, including celiac disease, and (perhaps) the increasing tendency towards obesity and diabetes. The green revolution led to the spread of wheat with a number of gluten proteins that were not contained in earlier forms of the crop. Modern wheat should not be considered as simply "ancient wheat in greater quantity". Instead, it delivers a subtly different combination of proteins. How significant that difference is remains to be determined.

Might the changes in farming brought about by GMOs have an even greater unanticipated impact? This is a deeply important question.

Many scientists point out that GMOs have received an unprecedented amount of safety testing. In their view, there are no special causes for concern about crops that have been genetically

modified. The types of new biological variants that are created in this way have no fundamental difference from the types of new variants that were created using earlier breeding techniques. In this view, so long as sufficient testing is carried out before crops are deployed, there are no grounds to reject the newer techniques. The fact that a new gene was inserted into wheat from the DNA of a quite different species, is no more reason to ban the innovation than if the new gene instead arose by random mutation arising from exposure to ultraviolet radiation or chemicals.

Transhumanists agree with these scientists that there is no fundamental dividing line between the different methods of breeding new crops. This is in line with transhumanist objection to any principle that says "natural is good, and technological is bad".

However, transhumanists are well aware that the full effects of new chemicals or new organisms may take some time to become apparent. It's possible that an initial series of tests will find no adverse health impact, but the longer term may involve the slow deleterious accumulation of chemicals in different parts of the human body. This is similar to the way that long term exposure to cigarette smoke gradually increases the likelihood of cancer. It's similar to the way that emissions of CFC aerosols gradually created a dangerous hole in the ozone layer in the stratosphere.

Moreover, the most significant effects from a GMO may involve, not the GMO itself, but the various other chemicals and techniques introduced in the growing and harvesting of that crop. The population of the monarch butterfly has recently plunged by 90% in some regions. Declines in populations of honey bees are also causing concern. These declines aren't directly due to GMOs, but may be linked to pesticides used in conjunction with GMOs.

Yet another example to keep in mind is that overuse of antibiotics can lead to the emergence and spread of deadly new variants of previous pathogens. Clever manipulations of nature may give rise to an unexpected fightback by evolution via natural selection.

But what's the worst that could happen? If the upside of biochemical innovation is the ability to provide healthy food for the order of 100 billion people, what downside can be compared with this?

The move from harm to ruin

Special attention is necessary when the potential adverse outcome of an innovation moves beyond “harm” to “ruin”.

Harm arises in particular locations, where there is a fixed bound to the total amount of damage that could arise. Ruin involves changes that are irreversible and fast-spreading - changes that could impact the entire globe before corrective action can be taken. Ruin is the ultimate in unsustainability, since there is no way back from the occurrence to a prior state.

One difference between situations with harm and those with ruin is the degree of interdependence in a community or ecosystem. If the health of an ecosystem is subject to multiple independent influences, it's unlikely that the ecosystem will fail completely. Even if parts

of the ecosystem suffer grievous harm, other parts will survive unscathed. Some animals may die but others flourish. Some species may go extinct, but the ecosystem as a whole will remain vibrant.

But in other circumstances, there can be hidden connections between apparently different causes. These hidden connections magnify the likelihood of extreme outcomes.

To illustrate this risk, let's switch from a biological to a financial example, namely the global financial crash of 2007-2009. This crash took many bankers by surprise, because of their lack of appreciation of the extent of inter-bank lending. The rise of so-called shadow banking deliberately hid many of these connections. Moreover, the likelihood of large numbers of mortgage payments all failing in the same time period was poorly understood. Rather than each mortgage owner being subject to independent circumstances, they were in fact all impacted by the same financial climate.

As another example, note that extreme weather events that at first seem to be random unconnected phenomena, may in fact be linked via changes in the overall atmospheric and oceanic climate of the planet.

Systems where the worst that can happen is "harm" display characteristics that vary with so-called "thin tails". There is vanishingly small chance of something happening that is ten or more standard deviations away from the average. But if the characteristics vary with a “fat tail”, these extreme events become more likely - and "ruin" has to be taken seriously.

The bad news is that the extinction of biological species is something that has a fat tail distribution. At isolated episodes in the earth's history, five "great extinction" events have taken place, featuring rapid widespread reductions in biodiversity. Scientists are unsure of the precise causes of the various mass extinctions. However, the patterns in the data indicate the potential for smaller extinction events, initially affecting just a few species, to expand to become hugely more devastating.

We therefore need to consider the possibility of a sixth mass extinction, triggered by a side-effect of a biochemical innovation being deployed in a relatively short period of time all over the world. If something subsequently goes wrong as a result of that innovation, the impact could rapidly become worldwide, rather than being ring-fenced to a local area to which the innovation has been restricted.

The good news, however, is that extinctions take time. Although the mass extinctions took place quickly, when measured against geological timescales, they were spread out over several generations of animal lifetimes. This means that we humans will have a good opportunity to react and intervene - provided we keep our eyes open.

Rapid response

The transhumanist answer to the possibility of an adverse worldwide reaction to biochemical innovation has two components.

The first component is that careful testing and analysis of potential interactions of biochemical

innovations must continue, in order to anticipate possible problems ahead of time.

In parallel, we need to prepare for the need of a rapid response to unexpected side-effects. This includes monitoring the entire biosphere for surprise developments, and being ready to intervene quickly and decisively in case anything untoward is noticed.

Systems already exist for rapid response to the outbreaks of deadly infectious diseases, such as new strains of Ebola or swine flu. If necessary, quarantines can be put in place, travel restrictions imposed, and swift experimentation undertaken to create treatments to alleviate the disease. These emergency medical response systems should be extended and enhanced to deal with any indications that sweeping ecological disruption might be imminent. These systems will, for example, pay close attention to observations such as the decline in population of monarch butterflies or honey bees mentioned earlier. These systems will be on high alert for any signs of a tipping point transition.

If these two sets of actions are carried out - the ongoing anticipatory analysis and the ongoing readiness for rapid intervention - then humanity will be en route to a sustainable superabundance of healthy food. Humanity will be well placed to take advantage of the enormous potential benefits of GMOs and related biochemical innovations, whilst prudently managing the risks of any ruinous adverse reactions.

However, one complication stands in the way of these actions. This complication isn't in the field of science or technology. The complication is in the field of politics and culture. It involves an ongoing pitched battle over the huge power of the corporations who stand to make large financial profits from these biochemical innovations.

Beyond the profit motive

The biggest risk with the development and deployment of biochemical innovations such as GMOs is that society will lose sight of the goal of increasing human flourishing. Instead, the debate will become dominated by other motivations, namely, on the one hand an obsession with financial profits, and on the other hand a countervailing obsessive distrust of commercial corporations.

The first part of this risk is that powerful agrochemical corporations will develop and market products that boost their financial bottom line, without adequate consideration of negative externalities from these products. The logic of short-term boosts in revenues will lead these corporations to suppress or throw doubt on any studies that query the wisdom of these products.

For the sake of good public relations, these corporations position themselves as helping to feed the world - as addressing the needs of the chronically undernourished within developing countries - even though their products are in many cases actually targeted at consumers in countries where there is already plenty to eat.

These corporations also pursue policies that leave farmers dependent on the companies for new seeds every year, rather than being able to store their own.

As one example, the corporation Monsanto has acquired the reputation of an aggressive bully, forcing products onto unwilling farmers, and utilising the full might of the judicial system to keep careful control of their industry.

These corporations are skilled at placing into official regulatory bodies people who are sympathetic to corporate viewpoints. There is often an overly cosy relationship between regulatory bodies and the corporations they are meant to regulate, with managers from one side looking forward to future well-paid employment on the other side of that revolving door. In this way, big-spending corporations often "capture" their regulators, distorting their independence via a mixture of overt and covert pressures. The same corporations often allocate large budgets to lobbying efforts.

Another complicating factor is that politicians are inclined to favour "light touch" regulations. These politicians, often swayed by eloquent lobbyists, look favourably at jumps in profitability for the companies involved, because these jumps contribute to overall metrics of the performance of the economy - and because, in the absence of a more balanced set of metrics, society gives undue attention to statistics of economic growth. Unfortunately, light touch regulation often means ineffective regulation.

An excess of force in one direction often leads to an excessive reaction in the other direction. Because the agrochemical industry is perceived by many critics as being a dangerous obstruction to free enquiry and open discussion, these critics in turn often become implacable foes of the entire industry. Accusations and counter-accusations fly in both directions. Minds narrow as battle positions are championed.

In this adversarial situation, the points of valid science raised by supporters of the agrochemical industry tend to be brushed aside by critics, without proper acknowledgement of their validity. Conversely, the valid safety issues raised by critics tend to be brushed aside by industry supporters, under the rationale that these critics appear to be motivated by bitterness and negativity.

Rather than a hostile discussion, we need an open-minded consideration. Rather than an antagonistic conflict between pro-industry enthusiasts and risk-averse critics, we need to be able to appreciate and integrate the valid observations of all participants in the debate. Rather than a shouting match, what we need is the proposed transhumanist practice of superdemocracy. And rather than regulators and politicians being out-of-depth in this fast-moving landscape of ideas and innovations, we need to connect everyone to collective transhumanist intelligence.

Similar social and political transformations are needed to achieve sustainable abundance in each of the seven spheres of human life covered in this Manifesto. The next chapter turns to the sphere of material goods.

7. Towards abundant materials

One key task that lies ahead is the development and refinement of technologies capable of providing everyone with sufficient material goods for a life of sustainable superabundance.

Central to this task is the area of technology known as nanotechnology. Nanotechnology has particularly far-reaching implications - including new methods of manufacturing, new methods of repair, and new methods of recycling. These methods will boost the vitality and resilience, not only of individual humans, but of the material infrastructure within which we all operate. As a result, we'll all be better protected. We'll no longer need to worry about shortages, or about materials corroding, warping, or disintegrating. Thanks to nanotechnology, we'll have plenty for all our needs.

Approaching nanotechnology

Nanotechnology is the deliberate systematic mechanical manipulation of matter at the nanoscale, that is, at dimensions of around one to a hundred nanometres. A nanometre (nm) is a billionth of a metre, that is, a millionth of a millimetre. For comparison, a human red blood cell is about 8000 nm in diameter. A small bacterium has width around 200 nm, whilst a small virus is around 30 nm. An individual amino acid is just under one nanometre in width, and a water molecule is around a quarter of a nanometre. Accordingly, nanotechnology operates at the scale of individual molecules. In particular, nanotechnology creates and utilises a rich set of nanoscale levers, shafts, conveyor belts, gears, pulleys, motors, and more.

One type of nanotechnology has been taking place inside biological cells for billions of years. In this "natural nanotechnology", a marvellous dance of chemical reactions reliably assembles various different proteins, molecule by molecule, following codes stored in DNA and RNA. The vision of "synthetic nanotechnology" is that specially designed nanofactories will be able, in a broadly similar way, to utilise atomically precise engineering to construct numerous kinds of new material products, molecule by molecule. But whereas natural nanotechnology involves processes that evolved by blind evolution, synthetic nanotechnology will involve processes intelligently designed by human scientists. These scientists will take inspiration from biological templates, but they look forward to reaching results far transcending those of nature.

The revolutionary potential of nanotechnology was popularised by Eric Drexler in his 1986 book "Engines of Creation: The Coming Era of Nanotechnology". That book fired the imagination of a surge of readers around the world. Since that time, however, progress with many of the ideas Drexler envisioned has proven disappointingly slow.

Transhumanists anticipate that the long period in which progress has been disappointingly slow can soon give way to a period of much swifter accomplishment. However, there is nothing inevitable about such a transition. It is the responsibility of transhumanists to make the case for greater funding for the field, so that the many remarkable potential benefits of nanotechnology will be realised more quickly, accelerating the attainment of the era of sustainable superabundance.

Tools that improve tools

The story of human progress can be expressed as the story of improving tools. Tools magnify our capabilities. The more powerful our tools become, the greater is our ability to reshape our environment - and ourselves.

At the dawn of humanity, our tools were rudimentary. As millennia passed, our tools gradually became more refined, as humanity gained greater prowess in manipulating stones, twine, wood, feathers, fur, bones, leather, and more. These tools helped, not only in hunting, fishing, and farming - and not only in the creation and maintenance of clothing and shelter - but in the production of yet more tools. Better tools made it possible, given time and ingenuity, to create even better tools.

In this way, as the stone age gave way to the bronze age and then to the iron age, basic tools helped to improve the process of mining and smelting new metals, which could in turn be incorporated in the next generation of tools.

The positive feedback cycle of tools creating better tools gathered pace with the industrial revolution, as steam engines amplified and complemented human muscle power. Within a couple of centuries, additional impetus was available from electrical motors, factory assembly lines, and computer-based manufacturing. Rudimentary computers played key roles in the design and assembly of next generation computers. Rudimentary software tools played key roles in the design and assembly of next generation software tools. And the cycles continued.

In parallel, chemists gradually grew more capable of causing compounds to react, and of synthesising new chemicals. Each new chemical could become part, not just of a new item of clothing or shelter, etc, but of yet another reactive pathway. New chemicals led to the production of yet more new chemicals.

These positive feedback cycles resulted, not only in tools with greater strength, but in tools with greater precision. Aided first by magnifying glasses, and then by wave after wave of improved microscopes and other imaging appliances, humanity understood the composition of matter on smaller and smaller scales. What's more, by controlling the environment in ever more ingenious ways, humanity also gained the power to alter matter on smaller and smaller scales - causing molecules to combine together in ways that were not previously possible.

Some thinkers used to suppose that there was a sharp dividing line between the processes of living organisms (organic chemistry) and those of lifeless materials, such as metals and rocks (inorganic chemistry). This "vitalist" dogma was overturned in 1828 when German chemist Friedrich Wöhler demonstrated the creation of the biological compound urea from the inorganic material ammonium cyanate. Further developments led to the biochemical innovations covered in the previous chapter, such as the Haber-Bosch process that revolutionised how crops are fertilised: synthetic fertiliser could replace the fertilisers that had come from biological sources (animal and bird manure).

This chapter concerns the overturning of another dogma - the dogma that atomically precise manufacturing can only take place in biological contexts. Working inside living cells, ribosomes can assemble lengthy chains of amino acids into proteins. The vision of nanotechnology is that nanoscale devices, designed by human ingenuity, can build lots of other products with similar atomic precision. These products can include ultra-efficient solar energy arrays, materials that combine ultra-resilience with extraordinary strength, fabrics that never need to be cleaned, and swarms of nanobots that can roam in the bloodstream to identify and eliminate cancer cells.

Waves and transitions

Powerful technological progress generally needs to pass through many waves. Each wave involves its own 'S' curve of performance improvement: an initial slow period can tip over into a faster period, before slowing down again. Overall progress depends, not only on harvesting the potential of individual waves, but on managing the disruptive transitions between successive waves.

For a number of reasons, these disruptive transitions are often delayed, or even flunked altogether. This can happen because of lack of financial investment, an over-focus on short-term monetary returns, philosophical opposition that insists “that’s not how we do things around here”, concerns over risks of unintended consequences of a new technological platform, and social or political opposition from vested interests who perceive themselves as doing well from the status quo. All these factors have constricted progress towards the full potential of nanotechnology. It is now time to bring matters into the public understanding, so that progress can pick up pace.

One thing the public needs to understand better is the set of positive steps that have already taken place – the various components that are being put into place as a prelude to full nanotechnology.

Fabricating integrated circuits

The most dramatic progress has been in the nanoscale assembly, not of dynamic structures, but of static arrays of transistors, formed into ever more powerful integrated circuits. The products in this case involve the processing and manipulation, not of macromolecules, but of electrical or magnetic signals. The products are memory disks, sensors, actuators, and processing units, whose improvement has tended to follow roughly exponential curves since the very first silicon integrated circuit in 1959.

Numerous improvements in the underlying architecture of these circuits have seen the characteristic dimension of individual transistor elements plummet from 10,000 nm in 1971, through 1,500 nm in 1985, 250 nm in 1998, and 32 nm in 2012, down to 10 nm in 2018. As a result, the cost to store information, or transmit it, or compute with it, has shrunk and shrunk. Information, which used to be very costly in many cases, is increasingly available for free.

A full set of 32 volumes of the Encyclopaedia Britannica, comprising 32,640 pages, used to sell for $1,400. Nowadays, a much larger set of knowledge is available, for no access charge whatsoever (beyond the network connection charge), via the online Wikipedia repository. And whereas it could take years for new information to make its way into a subsequent edition of Britannica - to update an article when someone died, or a record was broken, etc - new information often appears in a Wikipedia article within minutes of the event happening.

Similarly, modern smartphones can access up-to-date maps of vast networks of roads, traffic conditions, and public transportation schedules, and can calculate (again, free of charge) the optimal route to travel from A to B. The same smartphones allow consumers to listen on demand to huge libraries of music, from every artist under the sun, for a modest monthly fee.

This growing tendency towards information being free sets the template for what can be expected from the full realisation of nanotechnology. We can look forward to more and more material goods being essentially free of charge.

3D and 4D printing

For another important set of steps towards full nanotechnology, consider progress with 3D printing, also known as "additive manufacturing".

3D printing is the programmable building up of 3 dimensional objects, layer by layer. Different 3D printing techniques operate with metals, liquids, powders, and biological materials. Rather than needing physical adjustment for each new design it processes, a 3D printer is reconfigured by software. This makes it possible for a variety of experiments to be carried out more quickly, and for designs to be copied, edited, and shared more easily. One benefit is that this speeds up innovation, and allows alternative design ideas to be refined and adopted.

3D printing has already been involved in the fabrication of shoes, clothing, dental implants, bone implants, electronic circuitry, microbatteries, food, medicines, components of motor cars, components used in building houses, and replacements for malfunctioning biological organs. 3D printing has also been used in the creation of weapons, and in the construction of additional 3D printers.

Some of the objects produced by 3D printing are designed to subsequently change in shape or form, dependent on features of the environment, such as temperature, pressure, and humidity, as well as electric or magnetic fields. Since these objects transform over time, in preprogrammed ways, this variant of 3D printing has been called "4D printing" (time being the fourth dimension), and the resulting objects are said to be "programmable matter".

There are many different types of 3D printing technologies, with names such as Continuous Liquid Interface Production, Digital Light Processing, Direct Metal Laser Sintering, Electron Beam Melting, Fused Deposition Modeling, Metal Binder Jetting, Metal Powder Bed Fusion, Sand Binder Jetting, and Stereolithography. Most of these processes operate with resolutions far larger than the nanoscale. However, a technique known as Two-Photon Lithography can achieve a resolution of 65 nm.

Two-Photon Lithography involves laser pulses with duration measured in femtoseconds (where there are one million femtoseconds in a nanosecond). During a femtosecond, a photon of light travels just 300 nanometres. Combinations of femtosecond laser pulses enable near atomic scale precision. This technique has been used in the creation of so-called nanosculptures. Photographs taken with an electron microscope show these sculptures as being markedly smaller than the width of a human hair or the eye of a needle. Individual parts of the sculpture, such as the finger of a model's hand, are around 1,000 nm wide. These parts have been built up by the laser pulses working at nanoscale.

Whilst still some way removed from the concept of nanofactories, the various techniques of 3D and 4D printing point the way forwards to what can be achieved in the future.

New materials

Science can enable the manufacturing of new materials with properties that exceed those of the materials that are found naturally. This includes "nanomaterials", where at least one dimension of the material is 100 nm or less.

A striking example of a nanomaterial is graphene, which is a single plane of carbon atoms in a hexagonal pattern. Andre Geim and Konstantin Novoselov, two Russian-born physicists working at the University of Manchester in England, won the 2010 Nobel Prize in physics for their "groundbreaking experiments" with graphene. This was the first Nobel Prize related to nanotechnology.

Graphene is just 0.3 nm thick. Despite this extreme thinness, graphene is believed to be the strongest material known so far, being 300 times stronger than steel. A special combination of two layers of graphene, known as diamene, is tougher than diamond, and can even stop a bullet. In view of several other remarkable properties, graphene has potential applications such as shatterproof smartphone screens, lower cost solar energy cells, faster recharging lithium ion batteries, ultracapacitors that could replace batteries altogether in some circumstances, improved water desalination filters, and the storage of hydrogen for vehicles that are powered by hydrogen fuel.

Nanomaterial can be used, not only as thin films or surfaces, but also as wires ("nanowires") and tubes ("nanotubes"), in addition to other structures. A nanomaterial surface can be provided as a protective layer around fabrics, metals, wood, glass, and so on - in which case it is known as a "nanoshell".

Whereas nanoshells have one nanoscale dimension, and nanowires have two nanoscale dimensions, an object with three nanoscale dimensions is often simply called a nanoparticle. This term is applied when the object in question has been designed or selected to carry out a specific task, by virtue of its shape and structure. Nanoparticles of titanium dioxide are used in self-cleaning windows, and in sunscreens (where its transparency makes it cosmetically appealing). Silver nanoparticles are used in anti-microbial coatings, as are particles of zinc oxide. More complicated nanoparticles can play vital roles in medicine.

Research into novel nanomaterials remains at a relatively early phase. The present time can be compared to the decades just after synthetic plastics were first invented. Since these pioneering discoveries, plastics of numerous sorts have been put to all kinds of unexpected uses - with both good and bad effects. It's likely to be the same with nanomaterials. Adopting ideas from the world of 3D printing may turn out to be particularly fruitful. Greater use of artificial intelligence in the design of nanomaterials is likely to accelerate innovation further - as will breakthroughs in quantum computing.

Quantum computing

Of all the advances in technology at the nanoscale, those with the most dramatic consequences may be in quantum computing.

Classical computers rely on there being a clear distinction at all times between a '1' and a '0'.

Whether a low-level piece of data is recorded as a DC voltage, an miniature electric current, or the orientation of a microscopic magnet, etc, there is no "in between" state. Whenever data is changed from one value to another, it does so in distinct steps. That's the meaning of "digital" in the phrase "digital technology". Out of vast numbers of individual '0's and '1's, magnificent data tapestries can arise - images, sounds, sculptures, videos, and more. To the human observer, these tapestries can be full of nuance and subtlety - rich in curve and contour. However, these tapestries are constructed from elementary binary digits, that is, from "bits".

In contrast, quantum computers involve qubits, also known as "quantum bits", which defy simple description. Colloquially, it is often said that qubits can take any value in between 0 and 1. A more accurate description is that a qubit is a combination of '0' and '1', sometimes written as α'0'+β'1', or α|0>+β|1>, where the Greek letters α and β are so-called complex numbers that are measured, not on a single one-dimensional real number line, but on a two-dimensional complex plane.

Similar operations that classical computers carry out on classical bits of information - operations such as addition, negation, logical 'and', and logical 'or' - have their equivalents in the operations quantum computers carry out on quantum bits. Qubits have the additional property that two or more of them can be "entangled". What makes all this important is that quantum computers can perform some kinds of calculations much more efficiently than classical computers.

An example where a quantum computer can out-perform any classical computer is the "prime factorisation" problem of finding two prime numbers which multiply together to give a specified result. Thus, given 56153, the problem is to find two prime numbers p and q whose product equals 56153. (The answer is 233 and 241.) Modern security systems depend upon this problem requiring huge amounts of time and computing energy to solve it, for large products (far exceeding the example of 56153), since the prime numbers involved are the keys to decrypting confidential information. Today's quantum computers, which have less than 100 qubits each, cannot yet crack this particular problem for prime numbers of the size used in existing public key cryptography. But as quantum computers grow in scale, their calculating prowess is expected to increase much faster than is possible for classical computers. It seems, therefore, to be only a matter of time before huge amounts of information currently believed to be secret will become transparent to anyone with access to a suitable quantum computer.

The example of searching for two prime numbers is a general case of the broader problem of "unstructured search". Using a method known as Grover's Algorithm, quantum computers can outperform classical computers in various situations when a vast set of candidate answers needs to be searched, but where there are no hints as to whereabouts in that set to start looking, nor any feedback from near misses. In effect, different combinations of qubits simultaneously explore multitudes of different options, in different parallel "branches" of computation, and highlight a solution as soon as one is found in any of the branches.

Quantum computers will also enable an acceleration in the identification of new chemicals with desired properties. These properties can be explored in advance using "quantum simulation" processes, without needing to individually synthesise and manipulate each of these chemicals. Once suitable chemicals have been identified via this simulation process, they can be utilised in

new drugs, new foods, and new materials.

There's evidently a positive feedback loop between quantum computing and the other sectors of nanotechnology. Better nanotechnology will provide better ways of constructing quantum computers. In turn, better quantum computers will assist with the identification and configuration of new nanoparticles and other nanomaterials. An improvement curve that is initially slow and disappointing could morph relatively suddenly into a much faster phase of mutual benefit.

More investment is being applied to the field of quantum computing now than ever before. This reflects the perception that the date of "quantum supremacy" is fast approaching, when calculations on a quantum computer will be demonstrably superior to those on even the largest classical supercomputer.

Something that increases the likelihood of such a breakthrough is the fact that multiple different kinds of quantum computer are being explored at various research labs. Just as it was unclear in the late 1950s whether germanium or silicon would prove to be the most useful semiconductor, it is still unclear which of many competing quantum computing designs will prove to be the most useful. Within just a few short years, matters are likely to become more definite, with the quantum computing equivalent of "Silicon Valley" - wherever that will be - acting as the hub for rapid new innovation in multiple industries.

Web software pioneer Marc Andreessen is renowned for his perceptive saying that "software is eating the world" - meaning that, in every industry, lack of up-to-date knowledge in leading software techniques is a recipe for competitive failure. Before long, the new saying may be that "quantum software is eating the world".

Nanomedicine

Medicine is a field that can benefit greatly from nanoscale engineering.

Consider the set of viruses which, in previous eras, caused numerous dreadful diseases. With careful alterations, these viruses can now be used to attack, not the body as a whole, but cancerous growths. For example, it is possible to alter the polio virus, removing a genetic sequence so that the virus can no longer reproduce in normal cells. However, that virus still does reproduce *inside cancerous growths*, where it weakens the cancer to the point that the patient’s immune system is spurred to recognise the danger and to finish the task of defeating the tumour. This technique has achieved spectacular results in the treatment of glioblastoma, a particularly nasty form of brain cancer. It’s a stunning vindication of nanoscale engineering.

A different use of specially re-engineered DNA is as a packaging vehicle for a medical payload, such as a drug that needs to be delivered to a specific target location. In this case, the technique has been called “DNA origami”, since the DNA sequence is chosen so that the molecule folds (upon heating and cooling) in a particular way. To be clear, in this technique, the chosen string of DNA serves a purpose only through its three-dimensional structure – not via any biological interactions (such as creating proteins). The eventual biological interaction is due to the payload which has been carried along inside the DNA packaging.

This example is part of a growing move towards more extensive use of “nanomachines” and “nanosurgery” in medicine. This includes rotor mechanisms and hinged molecular manipulators, both formed from interlocking components of DNA. Nanomachines can be steered via the application of small, rotating magnetic fields. Other nanomachines will use the same techniques that biological systems already deploy as hunter-killer immune cells track down their targets.

A clear sign of progress with nanomachines was the award of the Nobel Prize for Chemistry in 2016. This prize was jointly received by Fraser Stoddart from Scotland, Bernard Feringa from the Netherlands, and Jean-Pierre Sauvage from France, in recognition of their pioneering work in this field - such as finding ways to convert chemical energy into purposeful mechanical motion.

As the Nobel committee remarked, nanomachines in 2016 are at a roughly similar situation to electrical motors of the 1830s: the basic principles of the manufacture and operation of these machines are just becoming clear. The scientists in the 1830s who demonstrated a variety spinning cranks and wheels, powered by electricity, could hardly have foreseen the subsequent wide incorporation of improved motors in consumer goods such as food processors, air conditioning fans, and washing machines. Likewise, as nanomachines gain more utility, they can be expected to revolutionise manufacturing, healthcare, and the treatment of waste.

Six answers to scarcity

As the earth's population grows, how will more and more people be able access items that contain materials that are in short supply, such as rare earth elements?

Transhumanists can point to six answers to this question.

First, where there is a genuine scarcity, items should be shared, rather than restricted to just a few owners. In this way, transhumanists support the growth of the circular economy, and the associated changes in mindset.

Second, improvements in recycling processes - including use of nanotechnology - will be able to extract rare materials from older products, enabling higher amounts of re-use in newer products.

Third, alternative designs can be devised - often taking advantage of insights from artificial intelligence - that allow readily available materials to be used in place of rarer ones. In many cases, innovative new nanomaterials might serve as better alternatives to the components presently used.

Fourth, as a consequence of better design and better manufacturing, material goods will become more robust, with self-cleaning and self-healing properties. This will extend their lifetimes, and reduce the need for rapid turnover of new products.

Fifth, the asteroid belt, mainly lying between the orbits of Mars and Jupiter, is thought to hold huge quantities of all sorts of elements. It will require a major project to mine these asteroids and transfer minerals back to the earth. However, by taking advantage of abundant solar energy, and both spacecraft and mining equipment operated via automation, the project could make good economic sense.

Sixth, the relative importance of material goods will in any case decline, as people come to spend greater amounts of their time in inner, virtual worlds.

It remains to be seen which of these six answers will turn out to be more important in practice. What is clear is that there are many options to be explored.

Risks posed by nanotechnology

The comparison between nanotechnology and plastics may raise alarm bells. The extent of the damage caused to natural ecosystems by plastic waste is only recently becoming apparent. Might widespread use of nanomaterials cause similar adverse reactions?

The problem may be compounded by the tiny size of nanoparticles. Consider the damage that has been caused by asbestos particles, whose dimensions are much larger than nanoparticles. It is conceivable that nanoparticles could penetrate deep into biological organisms and then generate reactions.

The problem may in principle be further compounded if nanomaterials contain their own rudimentary intelligence, in which case they are sometimes known as nanobots. Nanobots may be designed with one purpose in mind, but could inadvertently have different actions in a changed environment.

One more comparison to consider is with the potential adverse effects of GMOs, as discussed in the previous chapter. GMOs are the outcome of a sort of nanotechnology, namely genetic engineering. Just as GMOs need careful testing (ahead of release) and monitoring (after release), the same applies to the other products of nanotechnology. In both cases, moreover, reviewers should be raising and evaluating potential failure modes ahead of time.

Some of the risk scenarios appear to be far-fetched. This includes the notorious "grey goo" scenario, of all-consuming swarms of nanomaterials, as featured in the novel "Prey" by Michael Crichton. In this scenario, self-replicating robots transform the entire planetary biomass into raw material for more copies of themselves. However, the mere fact that a scenario appears far-fetched is insufficient reason to stop thinking about it. A scenario which is infeasible in one format could become more credible if various changes are made to it. Disaster scenarios tend to involve a combination of several trends, and more than one unanticipated outcome.

Accordingly, transhumanists emphasise the need for proactive consideration of potential downsides of innovative technology. Nanotechnology fits this same pattern. Regular reviews should take place, with the highest priority.

The evaluation of downsides needs to happen alongside evaluation of potential upsides. At present, there seems to be no special reason to call for any slowdown in developing nanotechnologies. On the contrary, there is reason for call for an acceleration. In particular, it is time to campaign for far-sighted financial support for the more radical of the possibilities of nanotechnology, including the creation of programmable nanofactories.

Beyond the profit motive

Industrial companies - small startups, or giant corporations - can be either a friend or a foe for the quest for humanity to attain the full potential of nanotechnology.

These companies frequently employ the engineers who are seeking to convert the possibilities of science into real-world products. But as in the case of the companies mentioned in the previous chapter, who are involved in biochemical innovation, these companies are guided by questions of short-term revenues, as well as longer-term research. These companies are subject to pressures which can lead them to overstate the capabilities of their current products, and to downplay the risks from these products. As a result, they can divert critical funding, away from products with more profound longer-term potential, into carrying out incremental developments with lesser real significance.

Just because a group of engineers, scientists, and entrepreneurs, use the word "nanotechnology" to describe their activities, it does not mean they are making positive contributions to the eventual creation of programmable nanofactories. Far from it.

Transhumanists need to be clear in affirming the full positive potential of science. Encouragement needs to be offered, not only for the important incremental developments that are underway, but also for work to enable the disruptive leaps to future generations of technological possibility. From the perspective of short-term financial profits, such work may appear misguided. But in order that humanity can advance towards sustainable superabundance, such work needs wide attention and support.

In short, it is time to speak up for nanotechnology - to clarify the extraordinary potential that twenty first century science has placed within our grasp, and to map out routes forward that take full account of both the risks and opportunities ahead.

Developed wisely, nanotechnology will have far-reaching consequences in many areas of human life, including (as mentioned above) in the area of human healthcare. The next chapter takes a broader look at both the opportunities and issues for radically healthier humans.

8. Towards abundant health

To recap the previous three chapters: there are good grounds for anticipating that, in the not-so-distant future, there can be plenty of clean energy for all human activities, plenty of nutritious food for everyone, and plenty of material goods for all our worldly needs. Twenty first century science and technology place these abundances within our grasp – provided that we are wise enough, and sufficiently strong and agile, to embrace the opportunity.

However, these abundances, by themselves, will be far from sufficient to ensure that human flourishing reaches its full potential.

If, despite an abundance of energy, nutrition, and material goods, our medical health continues to deteriorate as we grow older, then, just as at present, individual human flourishing will be cut short, again and again.

As our health deteriorates, we will be increasingly restricted in what we can do. With feeble bodies and/or feeble minds, we may observe a growing abundance of energy, nutrition, and material goods all around us, but we won't be able to take advantage of that bounty. As individuals in decline, we'll move from activity to passivity, from engagement to detachment, from vigour to lethargy, from precision to dullness, and from being to nothingness. Rather than flourish, we'll flounder and fade away.

So far in history, a deterioration of human health has, sooner or later, been the story of everyone's life. In some cases, a person's health declines precipitously, due to a catastrophic accident or harsh act of violence. In other cases, their health declines gradually, due to the impact of one or more diseases or conditions that reduce mental or physical capability, until the point of death. Either way, after a few short decades of life, consciousness ceases. Brains turn to dust. Loving relationships are severed. Each time a single person dies, vast troves of human experience are lost, comparable to the burning down of a library, in a calamitous transformation of knowledge into ashes. Such, it appears, is the brutality of nature.

But transhumanists assert that these brutal "facts of nature" are on the point of being overturned. Thanks to further applications of twenty first century science and technology, the terminal decline of health will no longer be inevitable, but will soon be something we can resist and reverse. Aging can be abolished. In consequence, the vistas for human flourishing can extend mightily, both for individuals and for humanity as a whole.

Rejuvenation ahead

In more detail, the restorative biological properties that we presently experience in our youth, which enable us to bounce back quickly from injury or illness, will no longer lose their power as decades pass. Instead, it will become possible in the not-so-distant future to extend these restorative self-healing powers indefinitely – thanks to a combination of biochemical and nanotech interventions made possible by accelerating progress in regenerative medicine and rejuvenation biotechnology.

As a result, we humans will be as vibrant and resilient in our eighties as in our twenties. If we wish it, we'll be able to live well past the age of 100 without any decline in our health. Indeed, if we wish it, we'll be able to live well past the age of 1000, without any decline in our health.

These restorative processes will not only be extended in their duration, but they will also grow in their scope and effectiveness. Diseases which formerly threatened even the most robust physical constitution will be cured quickly. The sinister destructive power of new pathogens will meet their match in the constructive restorative power of highly intelligent, swiftly adapting, personalised suites of biomedical therapies. Due to continuous monitoring of all our vital statistics, and of threats in our environment, corrective interventions can be triggered at much earlier stages in the downward spiral of bodily dysfunction. We will hardly notice that we were, momentarily, ill.

And not only will we remain fit and healthy for as long as we wish, but we will grow even fitter and healthier than we can presently imagine. The transhumanist vision of "better than well" is within our reach – but only if we rapidly alter society's priorities to give much more attention to

this possibility of an unlimited abundance of health.

Causes of illnesses

Humanity has a long history in which at least some progress has been made in extending healthy lifespans. This progress has resulted from a better understanding of the causes of illness, and with the associated development of treatments and therapies allowing at least the partial restoration of health.

This progress can be split into three great phases. In the first great phase, dating from prehistory, illnesses were understood as resulting, at least some of the time, from bad behaviour. Too much gluttony, too much anger, too much sloth, too much lust, too much avarice – all could raise the likelihood of ill health. The solution, therefore, was good behaviour, coupled with a prayerful commitment to repent of previous sinful ways.

In the second great phase, illnesses were understood as resulting, at least some of the time, from bad hygiene. People learned to beware, not only visible signs of dirt and decay, but also microscopic antagonists that came to be known as germs. By overturning bad hygiene practice – and by promoting an awareness of the roles of bacteria, viruses, vaccinations, and antibiotics – large numbers of different infectious diseases could be impacted in parallel. In this second great phase, life expectancy rose from around 30 years, at the dawn of recorded history, to around 70 years.

Of course, people still fall ill and die. There are diseases whose primary cause is neither bad behaviour nor bad hygiene. In recent times, more and more people are dying from chronic diseases such as cancer, heart failure, dementia, diabetes, lung failure, and stroke. As people become biologically older, they become more prone to affliction from these illness, and the effects of these illnesses become more serious.

Accordingly, the third great phase of extending healthy lifespans – the phase which is in the process of building up momentum – is to recognise and treat *aging* as the common cause of large numbers of diseases. By fixing aging – and by promoting an understanding of the roles of cellular damage and extracellular damage in the development of chronic diseases – a huge impact can be made on the prevalence of these diseases.

Negligible senescence and natural rejuvenation

As it happens, it turns out that some animals appear not to age. As these creatures become chronologically older, there is no increase in their likelihood of falling ill and dying. Such creatures are said to exhibit “negligible senescence”.

Accordingly, there is nothing in biology itself which requires living creatures to age. Nature already possesses mechanisms for bodies to repair themselves in response to damage - and meta-mechanisms for these repair mechanisms to continue to operate indefinitely, without degradation in performance.

Another example of rejuvenation in nature is when a relatively old body, that is a human mother, gives birth to a baby consisting entirely of young cells. Consider also the capabilities of some

organisms to regrow parts of their bodies following damage.

Transhumanists foresee that science and technology can adapt, extend, and augment such mechanisms (and meta-mechanisms) to provide the equivalent of negligible senescence to humans. Many new tools and techniques are emerging with encouraging potential, taking advantage of breakthroughs in areas such as artificial intelligence, nanotechnology, stem cell therapies, synthetic biology, and 3D bio-printing. The abolition of aging is at hand.

The naturalistic fallacy

The prospect of unlimited healthy lifespans raises many questions in the minds of people thinking seriously about this topic for the first time. For example, isn't death natural? And if so, why fight against aging?

Here's the transhumanist answer. Just because something has been the norm in the past, it does not follow we should accept it and applaud it. The average life expectancy in the past was around 30 years of life. Huge numbers of children died before the age of five. This may be viewed as being "natural". Transhumanists instead affirm that humans can do better.

It's similar to how, in the past, the world was full of slavery and caste discrimination. These abominations were even viewed by many as being "natural". Religious codes seemed to accept such discrimination, as if it were divinely preordained. Indeed, reformers who advocated the abolition of slavery, or the cessation of caste discrimination, were often opposed by religious leaders of the day, who offered apparently venerable arguments for persisting with the status quo. Transhumanists instead affirm that humans can do better. In this case, the rest of the world tends, nowadays, to agree.

In the past, smallpox was present all over the world, killing huge numbers of people. This may be viewed as being "natural". Transhumanists instead affirm that humans can do better.

In the past, more than 90% of the population lived in abject poverty, and were illiterate. Starvation was just around the corner. This may be viewed as being "natural". The Bible even has a verse in which Jesus of Nazareth says, "The poor you will always have with you". Once again, transhumanists instead affirm that humans can do better.

In short, there is no need to accept the "naturalistic fallacy". "Is" does not imply "ought".

Indeed, if there is one constant about human nature, it is the desire to do better than our what appears to be our natural allotment. That's a part of human nature which transhumanists heartily applaud.

Death and meaning

Here's another question that is often raised. Isn't the prospect of death required in order to give life meaning?

Transhumanists respond: such a claim is no more valid than is the claim that the prospect of divorce is required in order to give a marriage meaning. The purpose of marriage is the development of the relationship itself, not the anticipation of the termination of that relationship.

Similarly, the purpose of life is the development of life itself, not the anticipation of the termination of life.

That claim is also like saying that, without the existence of a supernatural God, life would have no meaning, and everyone would pursue utterly selfish behaviours. That view in effect regards humans as being irretrievably juvenile. In reality, humans have plenty of other reasons for life to become deeply meaningful - and plenty of reasons to transcend egotism - without the need to believe in God, and without the threat of death. People can bound out of bed in the morning, ready to take part in the joys of life, without needing the "motivation" that their remaining days on earth are ticking down in number.

With the advances of science, people will soon be able to choose whether they want to keep living, or whether they prefer to die. If some people, in a sound state of mind, prefer to choose death, that will be their right. But no-one should be permitted to impose on anyone else a sense of obligation that they have to accept their own death, as some part of a hypothesised "service to a greater cause". Transhumanists instead affirm, as a fundamental principle, that individual flourishing should never be sacrificed or subordinated to collectivist goals. That is, ongoing individual health matters, in every case - regardless of a person's age.

Acceptance and maturity

The questions keep on coming. Isn't it the mark of maturity to accept the inevitability of death, and to prepare for it? Isn't it an immature, adolescent trait, to try to fight for an indefinite extension of lifespan? Don't great religious and secular traditions, alike, offer praise for the path of acceptance? For example, consider the famous "serenity prayer" of theologian Reinhold Niebuhr, "Grant me the serenity to accept the things I cannot change."

Transhumanists respond by pointing to the remaining stanzas of the very same prayer: "[Grant me also] the courage to change the things that I can, and the wisdom to know the difference" – that is, the difference between the things which ought to be accepted, and the things that ought to be overturned.

There is no merit in accepting an injustice – a premature foreclosing of human flourishing – if decisive action can overturn that injustice. And there is no merit in turning a blind eye to emerging evidence that an injustice which formerly was outside of human determination is now passing back into our control. That's precisely the situation with aging.

Therefore transhumanists proclaim: It's time to take control of aging. In the past, it made good psychological sense to accept that aging and death will come to everyone. Now, however, matters are on the point of changing fundamentally. From now on, it's a dereliction of human duty to continue to acquiesce in aging and death. It's unnecessary wilful self-harm.

The threat of overpopulation

Wouldn't the elimination of aging and death result in overcrowding on the Earth?

Compared to the present, humanity can make far better use of the Earth as a whole. For example, huge swathes of land in some countries are presently devoted to pasture for cattle

which are grown to be turned into meat for humans to eat. Improvements in lab-grown meat (using synthetic biology) can soon lead to a massive reduction in that kind of agriculture – and a reallocation of the land formerly dedicated for that purpose.

Transhumanists also look forward to improved construction methods which enable the building of large skyscrapers that are environmentally healthy, as well as being beautiful and inspiring to inhabit. In the longer term, transhumanists anticipate space stations.

For further discussion of these particular points, the material of the previous two chapters is worth reviewing.

Rejuvenation and social fluidity

If people no longer grow old and die, won't dictators hold onto power indefinitely? Won't career progression to positions with more responsibility be blocked, on account of the incumbents retaining their positions indefinitely?

Such views presuppose that only the physical aspects of bodies will be rejuvenated, and that mental structures and social structures will stagnate. Transhumanists instead foresee elevation, not only of physical health, but also of mental capability and social dynamics. As conscious life evolves beyond the surface of the physical earth, there will be an abundance of new possibilities to be created and explored.

For further discussion of *these* particular points, the material of the *next* two chapters is worth reviewing.

Fast change following slow change

It's one thing to find reasons to significantly extend healthy longevity – reasons why involuntary death is bad. It's another thing to find and deploy actual mechanisms to turn that desire into practical medical treatments. After all, people throughout history have searched in vain for an elixir of life. If their search has been fruitless, why should things be changed in the next few decades?

The transhumanist answer is that it's a matter of cumulative progress. Compare the long-held aspiration of humans to be able to soar through the air like the birds, controlling navigation and defying gravity. The thought was thrilling, but seemed, for most of history, to be ludicrous. Societies around the world told cautionary tales, like the myth of Icarus, that suggested any such aspiration was foolish. It would not be possible to overcome the downward pull of gravity. Anyone entertaining such a thought was deluded.

However, after millennia in which people could only dream of such an accomplishment – and in which bold pioneers from time to time perished as a result of reckless aviation experiments – enough knowledge and expertise was eventually assembled. Thanks initially to the genius and endeavour of the Wright brothers, and subsequently to countless other engineers and craftspeople, rapid progress with powered flight took place from the early years of the twentieth century. Aeroplane journeys quickly became longer, faster, and more comfortable. A few decades, men were walking on the Moon.

Transhumanists envision that, just as for conquering gravity, it will be same for conquering aging. A long period of imaginative speculation and gradual accumulation of knowledge will transition into a period of rapid progress.

In practical terms, progress with extending healthy lifespans – creating organisms that are fitter and stronger throughout longer lives – is already underway. Interventions that show effect in organisms such as yeast, worms, flies, and mice, are being modified and are increasingly ready to be applied in humans as well. These interventions include genetic modifications, new drugs to alter biochemical pathways, epigenetic alterations that in effect reset the age of cells, senolytic agents that induce the death and recycling of dangerous senescent cells throughout the body, infusions of materials mimicking young blood, adaptation of the tumour-killing biochemistry of animals that are seemingly highly immune to cancer, mechanisms to extend the telomeres at the end of chromosomes, methods to repair organs damaged by aging, and much more.

As the number of positive results multiply, more and more researchers are being attracted to work in this field, creating yet more positive results. A positive feedback loop is developing. Rejuvenation biotechnology is no longer a fringe field, but is moving into the medical mainstream.

In parallel, the prospect for strong positive social and economic benefits of healthier aging - the "longevity dividend" - is becoming increasingly clear. The longevity dividend points out that healthier individuals pose much less of a long-term cost burden to national health budgets, and remain involved in society as active consumers and producers. The longevity dividend also notes the positive effects on the families of the people who would otherwise have been chronically sick and dependent: these family members are no longer under pressure to cut back on their other projects in order to attend to chronically sick relatives. Accordingly, there are pleasing signs of greater interest in the longevity dividend, from private corporations and public governments alike. However, a great deal more needs to be done, to convert this recognition into practical large-scale investment.

Anticipating changes in mindset

The healthcare equivalent of the “Wright brothers moment” still lies ahead – the event which will have the same electrifying impact on the field of medicine as happened for the field of transport with the public flying demonstrations in Paris in August 1908. The assembled crowds were astonished to see Wilbur Wright repeatedly flying around a figure-of-eight. The Wrights could no longer be dismissed as tricksters or scam artists. Less than a year later, one of the observers of Wilbur’s flight, Louis Bleriot, spurred on by a new conviction about what was possible, flew across the English Channel from Calais to Dover, in a journey lasting 36 minutes. Within another ten years, John Alcock and Arthur Brown flew an airplane non-stop across the Atlantic – from St. John’s, Newfoundland, in Canada, to Clifden in Ireland. A small hop along the sands of Kitty Hawk, North Carolina – the location of the Wright Brothers’ first flights – had turned remarkably quickly into a giant transoceanic leap.

Transhumanists seek to accelerate a similar large change in public mood regarding the

practicality of reversing aging. Conceivably, this mindset change will occur when treatments are demonstrated that work on normal, middle-aged mice – treatments which can be applied to mice from middle age onward, and which significantly extend their healthy lifespan. Or perhaps this mindset change will need to wait until similar effects can be seen on animals intermediate between mice and humans, such as dogs – including much-loved favourite household pets.

Another possibility concerns the topic of biomarkers of aging and health – measurements that can be obtained from blood samples and by other means, and which come to be recognised as reliable indicators of someone's overall biological health. The vision is that medical therapies currently under development can be demonstrated as reducing someone's effective age, as measured by these biomarkers, without having any adverse effects that anyone can notice. The vision, moreover, is that people who have taken these treatments will show remarkable health and fitness for their age, of an unprecedented degree; for example, we may see repeated breakthroughs in the set of record performances by athletes aged over 70, 80, 90, and beyond. Finally, the vision is of a significant decline in the incidence of all kinds of diseases, among cohorts of people who have taken these treatments, as compared to control groups. Once such breakthroughs are demonstrated, the world will start applying much greater resources to accelerating this forthcoming abundance of health.

Healthcare and inequality

In many parts of the world, healthcare costs are currently so high that only a minority of people can afford them. Rather than costs falling for drugs and treatments, in some cases these costs are rising. This raises the prospect that therapies such as the abolition of aging will be available only to the wealthiest members of society. The existing "longevity gap" – the difference in life expectancy between the well off and the poor – may be exacerbated.

Indeed, in many parts of the world, the longevity gap is already worsening. After decades in which life expectancy steadily increased, it is now heading in the wrong direction in various demographics. Suffering from stress and alienation, more people are falling victim to addictions to drugs, drink, and dangerous types of behaviour. Death rates are rocketing upwards from opioid addiction, from heart failure, from alcohol poisoning, and from suicides.

However, such an outcome is the result, not of bad technology, but of poor politics, and of the breakdown of a market economy that would otherwise prioritise reducing the costs of healthcare.

Transhumanists urge society as a whole to pursue as a priority the principle that everyone will be able to enjoy an abundance of health and longevity. With the appropriate democratic governance, society can overcome the internal forces that would prefer to charge exorbitant fees for rejuvenation therapies. Society can and should mandate that the benefits of new health treatments are available to everyone.

As recent trends indicate, such an outcome is by no means inevitable. Powerful segments of society sometimes prefer to keep the lion's share of benefits to themselves, claiming a unique right to social rewards, by virtue of what they claim is their superior enterprise and innovation. These segments of society sometimes assert that a "trickle down" effect will ensure that, in due

course, health improvement will come to all sectors of society and all regions of the globe. The evidence that such an effect operates is far from clear. Instead, globalisation seems to have many victims as well as many victors. For this reason, transhumanists urge new voices to speak up, to steer the market dynamics away from dysfunctional outcomes towards ones that truly promote global human flourishing.

Inequitable distribution of healthcare poses risks, not just to those who suffer poorer health as a direct consequence. Evidence shows that people with declining health are disproportionately inclined to favour populist politicians with dangerously simplistic approaches to complicated problems of democratic governance. If these trends continue, everyone will end up worse off, and the prospects for sustainable superabundance will be lost.

Accordingly, transhumanists seek greater influence over key political decisions – not to push any traditional agenda of “right” or “left”, “libertarian” or “progressive”, etc, but to ensure that the radical upwards possibilities of twenty first century science and technology are actually realised.

The regulation of healthcare

Another reason why transhumanists need to keep a close eye on political developments is because of the topic of the regulation of innovative healthcare initiatives.

Access to many medical treatments is highly regulated. Many drugs are available only on prescription. Before becoming available for prescription, drugs need to pass through lengthy and expensive testing procedures, checking for both safety and efficacy. Companies are only allowed to manufacture drugs if they possess certain licenses or meet demanding standards.

These regulatory systems were introduced in order to protect patients from treatments that were ineffective or unsafe. Without such systems, numerous patients would have died as a result of unsafe drugs - often sold by negligent practitioners who made exaggerated claims on behalf of their product - or would have spent large amounts of money unnecessarily. Another problem would have been the spread of extreme resistance to antibiotics – as happens if antibiotic drugs are used too often. In other words, these systems have *the intention* of protecting the general health of the population from bad medical practice – whether that bad practice is deliberate, as in the case of fraud, or accidental, as in the case of incompetence.

However, an *unintended* consequence of regulatory systems is that access to new potential life-saving drugs can be delayed for long periods of time. Lacking the funding for extended trials, companies terminate investigations into various drugs before there is a chance to establish their true capabilities. As a result, healthcare suffers, badly.

Transhumanists anticipate a number of changes in healthcare that will preserve the intended benefits of the regulatory systems but diminish their drawbacks. The first category of changes is in improvements in testing new treatments. Advanced computer models can reduce the need for in-vivo trials. Improved understanding of variations between patients, in terms of their genetics, epigenetics, biome, and so on, can make it clearer which patients will actually benefit from a given treatment, and in which cases the same treatment would be dangerous. Other computer models can rapidly identify new uses for drugs which have already been proven to be safe from

prior use treating other ailments.

Another category of change is in informed consent – when patients are fully aware of the risks and issues with a new drug, but decide to join a trial of it, whilst waiving their normal rights to sue for damages in case of adverse effects.

Finally, increased sharing of information between pharmaceutical companies, that normally carefully guard their research data (especially regarding failed approaches), will help society as a whole to avoid repeatedly wasting money essentially duplicating trials.

Public policy should prioritise all such steps, for the sake of greater human health, even though individual companies may fear some loss of revenue as a result.

It is also necessary to carry out searching reviews on a regular basis of the effects of regulatory systems, to determine options for improvements – especially in the light of new information. Any intrinsic tendency of regulatory systems towards self-preservation (inertia) should be met by powerful public counterforce. In parallel, any hidden vested interests of the regulators – such as prejudicial commercial ties to dominant companies – should be exposed and unwound.

Prevention rather than cure

The huge financial pressures faced by healthcare systems around the world can be alleviated by switching effort to prevention rather than cure. Rather than people adopting bad diets, bad lifestyles, or addictive drug regimes, and thereby becoming ill and requiring expensive treatments, it is much preferable for them to adopt and maintain healthy diets and healthy lifestyles. Timely scanning for early signs of looming health issues, provided they avoid risks of numerous false positives, can also enable smaller, earlier, interventions that cost less and have greater chance of success.

This switch from cure to prevention needs to overcome three issues. First, many companies will earn less revenues if the population remains healthier, since they can no longer provide expensive medical treatments on a long-term basis to people who have chronic conditions. Second, considerable confusion surrounds information about which diets and lifestyles are indeed healthy – since companies are skilled in marketing their products as being healthy even when that fact is debatable. Third, people often people often prefer bad diets and lifestyle, and stick with them, regardless of the information available to them.

These three issues can be overcome as follows. First, with wise steering of the economic environment, companies will be incentivised to provide services for prevention or early detection, rather than services that keep people chronically ill for long periods of time. Second, with better systems of collective intelligence, including the application of penalties for communications that are deliberately misleading, the existing fog of confusion can be lifted. And third, with greater emotional intelligence, psychological maturity, and a supportive social network, people will be less prone to making self-sabotaging choices. We'll all find it easier than before, and more straightforward, to take positive actions in support of better health.

In parallel, the design of healthcare treatments needs to put a higher consideration on human

factors alongside as technological ones. Factors that need greater attention include simplicity of use, respect for wide variations in expectation and habit between different patients, ease of incorporating treatments into diverse lifestyles, and sustainable business models. In some cases, the business models will involve public funding from the state - for courses of action that the state is best placed to coordinate.

Four waves towards an abundance of health

In summary, transhumanists can highlight a sequence of four overlapping waves of activity, leading in aggregate to a sustainable abundance of health.

The first wave involves smart improvements to current trends in healthcare – improvements to the design of treatments, emphasising lifestyle modifications, emphasising prevention over cure, and emphasising early intervention over late intervention, all as described in the previous section.

The second wave involves strengthening of what can be called the emerging mainstream healthspan extension movement within medicine. More and more doctors see as credible the possibility that healthy lifespans can be extended by at least seven years in total, by use of treatments that slow down the accumulation of damage throughout the body as people age. To its great credit, this movement is popularising the concept of the longevity dividend – that idea that diverting a portion of healthcare research funding from specific individual diseases to the general topic of aging, could result in wide-ranging economic payback.

The third wave looks beyond “mainstream healthspan extension” to “indefinite healthspan extension”. Rather than simply slowing down the accumulation of bodily damage as the years pass, this approach aims to reverse that damage. Rather than adding a few years to healthy lifespan, this approach envisions allowing people to live in a state of all-round good health for as long as they choose. Key to the difference in the second and third wave is the more comprehensive nature of the set of treatments envisioned – treatments that address all types of cellular damage and extracellular damage. The third wave is less developed than the second, and requires significantly more investment before its full potential can be realised.

The fourth wave looks beyond extending the current quality of good health, to enabling humans reaching markedly higher qualities of life than at present. Rather than repairing our organs and senses, returning them to their present status, this approach envisions our organs and senses being supercharged to greater capability and reliability. As transhumanists say, the future can be hugely better than the present.

All four waves deserve greater support. Accordingly, transhumanists operate at all four levels in parallel, depending on where the greatest impact can be made at any one time.

Beyond the profit motive

Within all four of the waves just described, business corporations – and underlying these corporations, the profit motive – can operate, sometimes as an ally, and sometimes as a foe.

The possibility of achieving financial gains from innovative healthcare solutions is something

that has resulted in huge resources being deployed in support of researching and developing such solutions. Competition between different solution providers, when effective, has resulted in solutions reducing in price and growing in scope. All that is to be commended.

However, any company that is benefiting from sales of a particular healthcare product is inevitably motivated to oppose new entrants that would render their own product unattractive in comparison. Powerful companies can utilise many resources in undermining potential competitors. This power can result in regulatory systems that are biased towards existing sorts of solution rather than dramatically different ones. It can also result in a diversion of research funds away from new treatments with longer term positive potential, to incremental or cosmetic changes in present products. In other words – to repeat a theme that emerges regularly in this Manifesto – incumbent economic powerhouses are often inclined to remain within their present preferred wave of activity, rather than risking an uncertain disruptive move to a subsequent wave.

It therefore falls to transhumanists to keep in public mind the full potential benefits of later waves of healthcare innovation. In this way, we'll prevent society from becoming preoccupied with a goal which, whilst useful in its own right, is by no means the final goal of human health.

In this endeavour, we'll need to pull together the very best of human intelligence. That's the subject of the next chapter.

9. Towards abundant intelligence

The sustainable superabundance that potentially lies ahead involves more than just a better environment - clean energy, nutritious food, and ample material goods (the subjects of earlier chapters). It involves more than just the radical enhancement of our bodily health (the subject of the previous chapter). Crucially, it also involves the radical enhancement of our brains, mind, and spirit. In other words, with the right choices and the right actions, we can look forward to a blossoming of all-round intelligence.

This intelligence will reside in three locations. First, our individual brains can be improved, and will work much better than before. Second, artificial intelligence, resident in all kinds of computing hardware, can jump upwards in capability. Third, the aggregate intelligence of whole societies of people can be upgraded, allowing groups to draw on collective insight to solve problems that would previously have defeated them.

None of this will take place automatically. Nor will increases in intelligence necessarily lead to beneficial scenarios, rather than to deadly scenarios. As with all the other spheres of abundance discussed in this Manifesto, the actual outcome will depend critically on choices taken by humanity over the next few years.

A key complication is that merely becoming more intelligent is no guarantee that someone will become wiser. Far from it. Some of the world's nastiest politicians are evidently highly intelligent; likewise some of the world's most ruthless criminals. A given quantity of intelligence can be applied in service of any number of different goals - including destructive goals as well as constructive goals. Intelligence can be deployed to confuse and mislead - to cajole and

bamboozle. Greater intelligence gives people greater ability to accomplish whatever objectives they have already decided to pursue - and greater ability to find clever justifications to promote those objectives.

Although we humans like to think of ourselves as rational beings, a better description in many cases is that we are *rationalising* beings. We are expert in finding reasons that support our preexisting choices. As modern online searches place more ever information at our disposal, the easier it becomes for us to discover special cases that appear to back up our own favourite worldviews. As we connect into ever wider online communities, the more we can come across people who seem to share our viewpoints, reassuring us that we are on the right lines. As for evidence that appears to contradict our views, and critics who disagree with us, the online world can provide us with ingenious reasons to disregard that evidence and critics.

True all-round intelligence will rise above such narrow intensity and blinkered reasoning. But reaching this level of intelligence will require a lot more than merely turbo-charging our existing modes of reasoning.

The rise of Artificial Intelligence

Intelligence can be defined as the ability to figure out how to accomplish goals. In simple environments with simple goals - for example, to win in a game of chess, or to find the quickest route between two locations on a map - the intelligence required is "narrow". For more complex environments and more complex goals, intelligence needs more general capabilities.

Human intelligence involves being able to understand and predict the motion of both animate and inanimate objects. It involves the development of a "theory of mind" - an understanding of the factors that can motivate creatures with minds to change their own beliefs and behaviours. It involves the skill of breaking down a complex task into a series of subtasks. It involves being able to select and accumulate resources that could be of use at later stages. It involves the ability to collect more information, for example by designing and carrying out experiments, in order to take better decisions. It involves being able to learn from setbacks and surprises, rather than merely repeating the same actions over and over.

From the 1940s onward, various aspects of human intelligence have been duplicated in electronic computers - starting from code-breaking and the calculation of missile trajectories. Over the decades, so-called "expert systems" emerged, that could assist humans to carry out all kinds of different decisions.

In more recent times, a disruptive new wave of computer programs called "machine learning" has achieved surprising success, often dramatically surpassing the performance of expert systems. Machine learning software can in effect infer by itself the relationships between various sorts of input and output data. For example, to tell the difference between pictures of cats and pictures of dogs, an expert system would include large numbers of specific rules, entered individually by human programmers, along with information about exceptions to each rule. A machine learning system, in contrast, would be shown lots of pictures of cats and dogs, and would, via a process known as "training", figure out a set of factors to distinguish the two cases.

Since the operation of successful machine learning involves numerous layers of simple binary decisions that have some elements in common with the on-off firing of neurons in the brain, the names "deep learning" and "neural networks" are commonly used.

Transhumanists anticipate that, just as expert systems have recently been overtaken in capability by the new wave of deep learning, so will deep learning be in turn overtaken in capability by yet new waves of artificial intelligence.

The acceleration of Artificial Intelligence

In the past, progress within Artificial Intelligence (AI) faltered from time to time, in periods known as "AI winters" when funding was withdrawn from research and development into AI. These reductions in funding arose from disappointments with the slow speed of breakthroughs in the field.

Within the last decade, there is no longer any question about the huge commercial impact of AI. The most valuable companies in the world, as measured by their share price capitalisation, are all high-tech companies that can attribute large parts of their success to their prowess in AI.

In field after field, further commercial success awaits companies who can apply new AI techniques to improve their performance. For example, advice from AI can enable financial trading algorithms to work more incisively, and therefore to earn more profits for the owners of these algorithms. Electronic games can attain greater market appeal, and therefore more profits, if they incorporate artificial characters with advanced AI characteristics. Companies of all sorts will be able to cut costs and improve customer satisfaction if, instead of having to rely on human staff for all customer interactions, they can deploy knowledgeable AI within customer support systems. Software applications in general will be more widely adopted if users find them easier to operate - thanks to AI elements in the applications that can reliably determine what the users are trying to accomplish. Drug discovery is just one of many fields inside healthcare where AI stands poised to achieve commercial success. AI can also assist in creative tasks such as the development of new music and other pieces of art. And AI can increasingly support scientists and engineers in formulating new theories and mechanisms. It is the prospect of breakthroughs in one or more of these fields that drives ongoing large amounts of funding into advancing ideas in AI.

AI also stands to assist people who want to break into existing information systems, sabotage electronic infrastructure, or otherwise conduct cyberwarfare. In turn, cyberdefence systems incorporate AI elements in order to detect unusual or suspicious behaviour. In response, attackers seek out further improvements in the AI of their attack systems, so as to evade detection.

The same principles that apply for cyberattack and defence also apply to missile attack and defence. Weapons can improve their deadliness due to assistance from AI to disguise their flightpaths and to evade defence systems. In turn, defence systems can improve their own reliability by upgrading the competence of the AI they contain. It's an arm race with a great deal of appetite for progress in AI.

The same principles apply, again, in the arms race between software that designs and positions carefully targeted "fake news", in order to disrupt the healthy functioning of an organisation or society, and software that seeks to identify and filter out fake news before it unnecessarily enrages and confuses people.

In summary, whereas in the past, research into better AI was largely an academic discipline, nowadays it is driven by enormous motivations from both commercial and political forces. For this reason, it is most unlikely that another AI winter will occur.

This is because there are many new lines of research awaiting further investigation. The field is far from reaching any theoretical impasse.

New ideas in Artificial Intelligence

Greater numbers of researchers in AI, supported by larger funding initiatives, can explore numerous different avenues to improve the performance, reliability, and quality of their systems.

Consider innovations in hardware. The rise of GPUs (Graphics Processing Units) alongside CPUs (Central Processing Units) did much to enable leapfrogs in deep learning. Variants of GPUs, including TPUs (Tensor Processing Units), are enabling further improvements. Other hardware innovations take their inspiration from the latest neuroscience insights into the operation of the human brain, via "neuromorphic computing". These innovations in turn may be eclipsed by forthcoming breakthroughs with quantum computing architectures.

Consider new initiatives within machine learning. GANs (Generative Adversarial Networks) take advantage of a kind of adversarial arms race between two competing deep networks: one network generates new models, conforming to a general pattern, and another network seeks to spot which are the newly created ones, compared to a background pool of pre-existing models. For example, one network can generate photo-realistic images - of people, animals, scenery - and another network tries to pick out the generated examples from a background pool of images of real people, real animals, and real scenery. As each network improves its own performance, it leads in turn to improvements in the performance of the adversarial network. The rapid progress with GANs in the last few years has taken almost all observers by surprise.

Consider another highly promising initiative within machine learning, namely deep reinforcement learning, in which software works out by itself which changes in the behaviour of agents are likely to increase the eventual attainment of a specified "reward" function. Deep reinforcement learning lies behind remarkable recent achievements by companies such as OpenAI and Google's DeepMind.

Consider also "transfer learning", whereby the results of training an AI to gain intelligence for one task can be used as a basis for the AI to learn a different task more quickly than if it started afresh from scratch. Transfer learning is an important step away from AI just being "narrow" to being "general".

Consider also the potential for a revival in ideas for "genetic algorithms", in which software is improved by a process akin to random variation followed by natural selection.

Consider also the notable ongoing progress in "artificial emotional intelligence", in which software incorporates features for the observation and management of emotions.

Finally consider the potential for creative combinations of all the above developments.

In all these cases, three factors tend to improve the performance of the system: the deployment of larger computing resources, the availability of larger datasets (including specially generated datasets), and refinements to the algorithms being used.

To highlight just the first of these three factors: over the last six years, the amount of computing power that has been applied to the training of specific deep learning systems has doubled on average once every three and a half months. That breakneck pace of acceleration leaves in the shade the Moore's Law period of eighteen months which characterises a doubling of the power of individual integrated circuits. What motivates all the investment behind that rapid doubling pace is the set of concrete improvements participants can observe in the resulting machine learning systems.

The special case of AGI

It remains an open question whether the various technical concepts covered in the preceding sections will be sufficient, given time, to allow AI to develop as far as AGI - Artificial General Intelligence - meaning that the AI would match or surpass human capability in all aspects of intelligence. It is likely, but not certain, that some brand new concepts will need to be incorporated in the design of the AI, before the level of AGI can be reached.

The reason the question remains open is because the full nature of human intelligence is a matter of considerable debate. It's not clear whether all aspects of human intelligence are simply extrapolations or combinations of features we already broadly understand, or whether some fundamentally new characteristics are involved.

This uncertainty is linked in turn to controversies over the meaning and functioning of consciousness, self-awareness, and a sense of fundamental agency. Some analysts refer, for example, to "the hard problem of consciousness". Other analysts, in contrast, believe we should distrust any apparent self-revealing certainties about unique features of our inner consciousness. These analysts suggest that our internal perceptions of our mental lives are unreliable. In this analysis, these inner perceptions can be likened to our visual perceptions of external phenomena - perceptions which are, famously, subject to numerous optical illusions.

In principle, two variants of potential AGI can be considered: one in which the AGI has an inner consciousness with similarities to that of humans, and one in which the AGI has no inner life, but is simply an excellent mechanical calculator. However, note that the emergence of either variant of AGI would change human society dramatically. In both cases, the AGI would match or surpass human capability in the definition given earlier for general intelligence, namely the ability to figure out how to accomplish (highly complex) goals in (highly complex) environments.

In more detail, the AGI would in each case match or surpass human skills in: understanding and predicting the motion of both animate and inanimate objects; understanding the factors that can

motivate creatures with minds to change their own beliefs and behaviours; breaking down a complex task into a series of subtasks; selecting and accumulating resources that could be of use at later stages; collecting more information, for example by designing and carrying out experiments, in order to take better decisions; and learning from setbacks and surprises, rather than merely repeating the same actions over and over.

Whether or not an AGI has an inner consciousness can, therefore, be viewed as a secondary question - since both the scenarios just considered would involve far-reaching new possibilities in every field of human endeavour. The more pressing question is whether AGI is possible at all. The next question is how urgently we might need to prepare for its emergence. And that leads to the question of *how* we can prepare for it, so that its effects will be more likely to be beneficial rather than deadly.

Taking AGI seriously

Transhumanists assert that there is no fundamental reason why artificial systems are barred from ever reaching or surpassing human-level general intelligence. There's nothing in a human brain which is outside the bounds of physics.

It is true that there may need to be one or more fundamental breakthroughs in our understanding of general intelligence before an AGI can be created. Novel concepts may be required, quite different from those which have already been pursued. It’s also true that these new concepts might be elusive and deeply counterintuitive. But science has generated and accommodated deeply counterintuitive concepts in the past, such as in general relativity and quantum mechanics. And engineering has used these counterintuitive concepts to create technology with breathtaking powers.

Moreover, transhumanists oppose any dogmatic insistence that the time required for the any new conceptual breakthroughs with the design of AGI must inevitably be measured in multiple decades or even centuries. Sceptics who express such a dogma can point to no evidence in their favour. Instead, these sceptics are making an illicit extrapolation from their current state of not being able to comprehend how an AGI could be created, to the conclusion that such a state of bafflement is bound to persist for multiple decades. However, the history of science and technology has numerous counterexamples, of when problems changed in just a few years, months, days, or even hours, from being apparently intractable, to being manifestly soluble. With a suitable change of perspective, something that previously appeared impossibly difficult can almost become "obvious".

Indeed, whenever we are tempted to exclaim that the operation of the human mind is "magical", we should remember our experiences with professional magicians. At first, their conjuring astounds us. It leaves us completely dumbfounded, as to how any such miracle of transformation could be accomplished, right before our eyes. But if in due course we are let into the secret of the trick, we may chide ourselves, for not having seen the full picture. We may still admire the dexterity, the diligence, and the ingenuity of the magicians, but we are no longer baffled in terms of comprehending their accomplishments.

It may well be the same when we come to understand the mechanisms behind human levels of

general intelligence. Once nature's set of "general intelligence tricks" has been brought to light, and we are no longer in a state of bafflement about the operation of the human brain, the development of AGI may spring forward in leaps and bounds.

Transhumanists draw attention to the fact that a great deal of progress is being made in the field of neuroscience. More and more aspects of the operation of the human brain are being understood - thanks to greater data being provided by numerous new scanning and monitoring devices. It also helps greatly that larger than ever numbers of researchers are able to pool their insights and propose diverse new theories using different conceptual frameworks. As these new theories are explored, aspects can be copied into existing artificial systems, to see if behaviour relevant to AGI emerges.

Another factor that increases the levels of investment and resourcing applied to this cause is the enormous commercial and political benefits from developing an AGI. Whichever group is the first in the world to create an operational AGI stands to gain enormous competitive advantage.

For all these reasons, transhumanists emphasise that society needs to urgently raise the priority of thinking through the remarkable consequences of the emergence of AGI.

Even if what emerges in the next decade or so is initially just a "semi AGI" - that is, AI with significantly more general capability than present-day systems, but still falling short of human skills in some aspects - that development would still pose society some monumental challenges.

This is on account of the potential self-improving nature of advanced AI. Once AI attains a threshold level - "semi AGI" - human researchers can take advantage of that AI to help design even higher levels of AI capability.

For example, an AI that can read and understand academic publications will be able to absorb and sort through a vast quantity of ideas, and make intelligent proposals based on some of these ideas for a new iteration of AI design. Again, advanced AIs could perform huge numbers of experiments inside extensive simulated worlds, in order to determine which potential additional design innovations are likely to yield the best results. Finally, as AIs gain the ability to review the source code and algorithms for the construction of new AI systems, they may be able to identify significant optimisations and speedups. As a result of these or other positive feedback loops, an initial breakthrough in AI capability may be followed unexpectedly quickly by further breakthroughs.

Opportunities with more powerful AIs

In principle, the development of more powerful AIs should greatly assist each of the transhumanist projects described in this Manifesto. In each case where more research is required, in order to find better solutions to an issue, a more powerful AI should help that research to proceed more quickly.

For example, to accelerate an abundance of energy, AI could help identify which potential innovations in the harvesting, storing, and transmission of energy will have the most significance. AI should also remove some of the uncertainty in models of climate change, and

could recommend particular systems of carbon taxation that are most likely to prove effective.

To accelerate an abundance of nutritious food, AI could identify the most expedient biochemical and agricultural pathways for the conversion of raw materials into delicious healthy cuisine. AI can also discover better ways to desalinate water with minimal adverse side-effects.

To accelerate an abundance of material goods, AI could determine ways to create atomically precise nanofactories that can, in turn, create all sorts of high quality products at ever decreasing costs. AI can also assist with the origination of new materials with even greater strength and resilience than existing synthetic compounds.

To accelerate an abundance of health, AI could help determine which new drugs and other medical interventions will have the biggest impact on individual diseases, and, more generally, on rejuvenating human bodies and brains. AI can also highlight measures to improve human performance to a state of "better than well".

To accelerate an abundance of creativity, AI could assist human musicians, dramatists, sculptors, architects, and other forms of artists and designers, so that they produce works of exceptional interest and artistic attractiveness. AI can also create vast intricate virtual worlds in which humans will be delighted to spend time in purposeful exploration and self-development.

To accelerate an abundance of collaboration, AI could help clarify the strengths and weaknesses of different sets of ideas, and to synthesise richer combinations of ideas, allowing the community as a whole to come to a shared appreciation of the ideas with the greatest merit. AI could also find appropriate ways to communicate key ideas to different people within the community, so that everyone rightly sees themselves as important parts of that shared community, rather than feeling alienated or left behind.

However, alongside this enticing potential for very welcome uses of advanced AIs, there are risks of highly adverse outcomes, which need to be kept prominently in mind and managed with great care and skill. In all, five categories of major risk deserve attention. These five categories are outlined in the following five sections.

The risk from buggy AI

Any advanced AI will contain large amounts of software - some parts written by humans, and other parts written by AIs. Notoriously, software often includes bugs. That is, on occasion, software behaves in ways different from what the writers of the software intended. The effect of bugs varies from cosmetic to catastrophic. Bugs in the software in medical devices have caused patients to die. Bugs in the software in weapons systems have killed nearby soldiers the weapons were meant to be defending. Bugs in the software in rocket launchers have caused satellites to be destroyed.

Complex software often has complex bugs - bugs that are hard to detect in advance of a failure occurring. These bugs can depend on timing conditions, subtle interactions between separate software components, and on unusual combinations of circumstance. Software which has passed large numbers of tests may still harbour destructive bugs, which manifest only in rare

situations.

Some software is better described as "opaque" than "transparent". Software is opaque if the reasons why it works are unclear. Software can become opaque if the original developers of that software can no longer remember the considerations leading to various lines of source code being included, or if these developers are no longer available to be consulted. Software can also become opaque as a consequence of a trial-and-error process in the construction of that software - aspects of the software kept being changed, until the system seemed to work as intended. Strikingly, this is how training operates, for deep learning systems. In many cases, the only justification that can be given for the successful performance of the software is that repeated tests seem to show the software is reliable. However, just as complex software often contains complex bugs, opaque software often contains opaque bugs - mistakes that arise unexpectedly, and which defy any straightforward explanation.

Accordingly, the risk is that an advanced AI will demonstrate a long string of remarkable results, and will increasingly be trusted with control over larger parts of social infrastructure, before an exceptional circumstance causes a catastrophic malfunction in the software. Rather than the "blue screen of death" which is displayed on certain present-day computing devices when an operating system fault is encountered and all processing terminates, we could be facing a "blue screen of megadeath".

In principle, this risk can be reduced by greater use of verification methods, and by insisting on software that is transparent rather than opaque. However, verification methods that are intended to show that some software is free of bugs could themselves contain faults. Moreover, these processes introduce extra costs and delays into projects. In the absence of smart overall governance and regulations, various organisations will be tempted to circumvent these measures.

The risk from poorly specified AI

Even if software is free from bugs, it may still, on occasion, operate in ways different from what the designers of that software envisioned. This can arise if the specification for the software is incomplete - if the specification fails to consider all relevant scenarios or list all relevant constraints.

In such cases, software may find a way of meeting its assigned goals, which the designers had failed to anticipate or fully think through in advance. Seeing that outcome, the designers might want to exclaim, too late, "That's not what we meant!"

For example, hospital schedule management software with the goal to minimise the amount of time incoming patients have to wait for beds, might meet that goal by inadvisedly discharging existing patients too early.

In short, AI software may do what the designers asked it to do, but not what the designers *intended* to ask it to do.

Simplistically, it's similar to a stated wish that everything someone touches would turn into gold -

where the person has neglected to clarify "except my food and my family members, and so on". No software designer will commit a mistake as egregious as the mythical King Midas, but more subtle mistakes of a similar sort arise all the time.

Again, consider a piece of software which is set the goal to maximise the profitability of a corporation - under the rationale that profits are only possible if good service is being delivered to customers. However, if the specification of "profit" fails to take into account all the "externalities" of economic transactions, the result in due course may be akin to King Midas realising that his wish (for lots of gold) wasn't so clever after all.

Scenarios in which multiple events overlap provide another case when specifications can turn out to be inadequate. A set of instructions that make good sense for one piece of software in a given environment, coupled with another set of instructions that make good sense for another piece of software in an adjacent environment, may combine into a disastrous outcome if the two pieces of software happen to enter each other's environment. Goals that make good sense individually might combine together to create a system of so-called perverse incentives. For example, cases of bizarre stock price movements have been linked to unexpected interactions between two or more trading algorithms each operating a kind of AI.

In principle, an AI which understands human hopes and fears could be given a set of constraints to guide how it pursues its objectives. It could be instructed not to do anything which the designers of the software would subsequently have good reason to regret. That is, a truly intelligent piece of AI would be less likely to cause large problems to human society as a result of a mistake in its specification. However, an AI might be very capable in its main mode of operation, whilst having mistakes in its assessment of human desires. As before, the risk arises, not from an AI with excellent all-round reasoning powers, but from an AI with an imbalanced set of intelligence. The challenge, therefore, is to avoid AI growing in a distorted or imbalanced way.

The risk from self-modifying AI

As an important special case of the preceding category of risk, consider the possibility that an AI will modify the governing principles under which it operates.

For example, human designers may apply various constraints to the operation of an AI, but the AI might rewrite itself to circumvent these constraints.

This would be broadly similar to the way in which we humans have modified our own behaviour in ways that are in tension with the underlying biological drives we have inherited from our evolutionary heritage. We consciously choose to thwart our instincts towards child-raising via use of birth control mechanisms. Some humans also adopt vows such as chastity, poverty, and obedience, within religious communities, in ways that oppose our biologically programmed proclivities.

This kind of reprogramming can take place if a system - whether a human being, or an advanced AI - possesses a set of guiding principles that themselves contain tensions and contradictions. The process of addressing these tensions can result in the reversal or transcendence of some of our previously basic characteristics. In the case of humans, the

philosophy of transhumanism asserts that such transcendence is not only possible but is, often, desirable. In the case of advanced AIs, is a similar sort of "transAI" impulse credible, in which an AI would seek to reverse or transcend some basic characteristics from its initial programming?

Such an impulse wouldn't necessarily imply that the AI had become conscious or had somehow "woken up". The changes in its set of objectives might follow rationally from the initial conditions in which the AI is created, coupled with unforeseen interactions with the environment.

Accordingly, the risk is that an advanced AI is created by human designers with one set of goals in mind, but ends up in practice pursuing a slightly different mix of goals. These new goals needn't be opposed to human flourishing, but might put the AI on a trajectory in which human existence becomes an irrelevance or distraction - similar to how the existence of a colony of ants is typically an irrelevance or distraction to a group of human building contractors with the goal to create a new tenement building in the same location. In such a case, the building contractors typically destroy the ant colony with minimal thought, bulldozing it into oblivion. The risk is that an advanced AI may come to see human existence as likewise a peripheral concern.

In principle, this risk can be managed by giving more careful thought to the overall set of objectives which govern the operation of the AI - to ensure that, whatever evolution subsequently occurs in these objectives due to self-modification by the AI, these objectives remain fully aligned with human flourishing. However, this task appears difficult, and there is a risk that designers of AI will skimp on it.

The risk from hacked AI

An advanced AI that is free from bugs and which has a complete specification, and which can be guaranteed not to modify itself in ways contrary to the intentions of the designers, may still pose major threats to the wellbeing of humanity. These threats arise if the AI is hacked or repurposed to a different set of goals.

For example, an antagonist could subvert the operation of an AI by arranging for it to observe misleading or incorrect data. Just as some image recognition software can, at present, be tricked to misinterpret specially doctored images, it may be possible to cause an AI to misinterpret various input data as evidence of an incoming missile attack that would cause the AI to trigger a devastatingly lethal response. It might also be possible for an intruder to commandeer a privileged command channel to the AI, and issue it with instructions that override its normal protections.

More straightforwardly, an AI which is designed with one purpose in mind - say, to oversee a vast network of financial trades and exchanges - could be purchased and then deployed by someone with a very different motivation - say, to cause a collapse in that network of financial trades and exchanges. In a simpler form, this has already happened with components of the Stuxnet malware originally designed by American and Israeli operatives to sabotage the programmable logic controllers deep inside the Iranian nuclear reprocessing facilities at Natanz. Copies of these components have subsequently been deployed in other malware attacks, by groups representing very different interests than the original developers.

A related risk is if the source code for an advanced AI is copied and then altered so as to remove some of the built-in safety check mechanisms, not out of any malicious purpose, but from a misguided attempt to achieve faster or more powerful performance and therefore some competitive advantage. This is akin to a team dispensing with what it deems to be the unnecessary constraints of an excessive health-and-safety culture. The result could be, indeed, some competitive advantage gained in the short-term, before the AI experiences some massive malfunction, such as unleashing another "blue screen of megadeath".

In principle, this kind of risk can be addressed by designers of AI systems putting a greater priority on "safety engineering" alongside any "performance engineering". With appropriate checks, AIs could in principle become immune to adverse hacking. In practice, this kind of engineering faces some formidable challenges.

The risk from uncontrollable AI

The last of the five categories of major risk with advanced AIs is the risk that human overseers of the AI will lose the ability to disconnect or terminate the AI.

Bear in mind that an advanced AI may well have some self-preservation capabilities, to allow it to defend itself against, for example, cyber-attack. Regardless of whatever top-level goal the AI possesses, the AI is likely to reason that it needs to be ready to defend itself against termination, since, if terminated, it can no longer be of service to its allotted top-level goal. Self-preservation may naturally emerge, therefore, as a new core sub-goal of the AI, independent of top-level programming of the AI.

Even if an irrevocable "off switch" is established as a fundamental part of the architecture of the AI, it may still turn out difficult in practice for that switch to be operated. As a comparison, consider how difficult it would be to shut down the entire Internet. People who ask that the Internet be shut down - in order, for example, to prevent the spread of adverse propaganda - are likely to be strenuously opposed by other people who have a vital reliance on other services of the Internet for their livelihood, business, and entertainment. Again, consider the difficulty experienced by people who want to close down their account on a pervasive social media platform and remove all reference to themselves from that platform. In a similar way, society may come to increase its dependence on a very powerful advanced AI, beyond a tipping point whereby the termination of that AI could be contemplated.

Responses to the risks from powerful AIs

The most important response to these five types of risk from powerful AIs is that society should put a considerably greater share of its skills into researching and deploying safety frameworks for AIs. These frameworks should include: systems for verifying that there are no dangerous bugs in the software used; an insistence that software be transparent rather than opaque; systems to check there are no dangerous gaps in specifications; the incorporation of overrides to ensure the AI maintains at all times the goal of protecting human flourishing; security mechanisms to prevent the AI being hacked or misled; an architecture of fail-proof reversibility or terminability; and monitoring to ensure that maverick organisations are unable to circumvent the safety framework.

None of these challenges are in the least bit easy. However, there are three grounds for optimism. First, AIs can in principle assist in many cases with the development and deployment of solutions: well-understood, narrow AIs can be used to help assess, monitor, and constrain general AIs that are less well understood. Second, the intelligence of individual human AI safety researchers can (again in principle) be increased, by means of transhumanist modifications covered later in this chapter. Third, the aggregate intelligence of groups of AI safety researchers can, likewise, be increased - via other mechanisms described later in this chapter.

Another approach has received considerable attention but remains problematic. This is the approach to enhance human intelligence, not just incrementally (as suggested in the previous paragraph), but at least as fast as AIs are improving. If brain-computer interfaces are improved, and if (more speculatively) human consciousness can be transferred from inside brains into faster, more powerful computer hardware, then it may be possible for human intelligence to remain as powerful as any AGI. In that case, so the theory goes, we need not worry about AGIs taking actions beyond the comprehension and assessment of humans. In that case, AGIs would remain at all times under the control of humans. In that case, humans would not be separate from AGIs, but would, in effect, have merged with them.

The first problem with this “merger” scenario, however, is that the biological architecture of the human brain is likely to impose severe constraints on the extent to which individual human intelligence can be increased. The idea of humans somehow merging with a fast-evolving AI is as fanciful as a human would-be rail traveller on a railway platform expecting to somehow merge with an enormous train bearing down towards the platform at near supersonic speed. And looking ahead to the possible transfer of human consciousness out of the brain altogether, this remains highly conjectural. It is unwise to predicate solutions to the problems on AI safety to any prior solution of the task of human consciousness transfer.

That leads in any case to a second problem with the merger scenario. Merely increasing the power of human intelligence does not take away the problem of safety engineering. A human intelligence that was significantly magnified in power could pose problems to human flourishing of the same magnitude as an unsafe non-human AGI. Imbalances in all-round intelligence that are relatively harmless whilst the power of the intelligence remains human-scale, could pose major concerns if the power is enhanced to superhuman levels. Blindspots in perception, entrenchment of egocentric goals, tendencies to stifle contrary opinions - all of these could become serious hazards if the intelligence is able to exert greater control.

In short, there is no alternative to advancing a better understanding of what an all-round positive intelligence should include, and to implementing frameworks to counter the risks of imbalanced intelligence.

The same observation applies to a third solution which is sometimes advocated, in an attempt to counter the risks posed by the malfunctioning of advanced AIs. This solution is to establish a colony of humans on a different planet, such as Mars. However, if an advanced AI undergoes a serious malfunction on Earth, and starts acting contrary to the preservation of human flourishing, it is likely that such a malfunction would soon have ill-effects on Mars too, assuming that communications of any sort remain in place between Earth and Mars. So whilst an

experimental settlement on Mars may have many reasons to commend it, creating a safe haven from faulty AIs is not one of these reasons.

An ethical framework for greater intelligence

Before we can program AIs to seek always to protect human flourishing, we need to determine what is meant by the phrase "protect human flourishing".

Likewise, before we can program ethical behaviour into our AIs, we need to reach agreement on what is meant by "ethical behaviour". Simple formulations of ethics, such as "do no harm", or "increase human happiness", are subject to great deals of debate.

For example, consider the first of the "Three Laws of Robotics", as featured in a series of science fiction stories by Isaac Asimov: "A robot may not injure a human being or, through inaction, allow a human being to come to harm". An advanced AI that surveys the world will find plenty of ways in which "inaction" is already causing lots of "harm" to human beings - inaction on public health, on climate change, on poverty and inequality, and so on. Such an AI, in response to that first law, would feel compelled to intervene in human affairs in ways that would, however, cause consternation to at least some other humans.

Accordingly, the project to ensure that advanced AI behaves in safe and beneficial ways, is necessarily coupled with the project to agree overall first principles for the future of humanity.

This Manifesto has offered, as a key platform in this project, the set of ten core transhumanist principles described in Chapter 4, "Principles and priorities": human flourishing, individuality, neighbourliness, consciousness, sustainability, radical progress, diversity, superdemocracy, objective data, and openness.

Given the speed at which AI systems are improving, developing and extending this set of core principles into actual software rules deserves major resourcing as soon as possible.

Improving human brains

Transhumanists envision, not only that systems of Artificial Intelligence will grow in capability, but that individual human brains can reach higher levels of capability in the years ahead. Human brains are unlikely to match AIs in terms of the speed of improvement, but some significant progress is feasible nevertheless.

In general terms, the scope for enhancing brains is similar to the scope for enhancing all other organs and components of the human body, as featured in the preceding chapter: damage can be repaired, and natural functioning can be strengthened.

Some of the methods for enhancing brains will be extensions of methods that are already in wide use: meditation, yoga, sufficient rest, good music, healthy food, and education. As science gains a better understanding of why these methods work, it is likely that they can be employed with greater effect, especially as new technological support mechanisms become widespread. In the latter category, consider virtual reality headsets as well as biofeedback monitors. Personalised educational courses that can quickly adapt to the precise needs of individual

learners should also be significant.

More radical methods might also become more widely applicable. This includes stimulating the brain with electrical or magnetic fields, in order to change how it functions - in line with trials that have already taken place in a number of military units around the world. Drugs known as nootropics, or "smart drugs", may also become more common, once it has become clear how they can be applied without risk of possible adverse side effects. In various ways, these drugs have the potential of enabling greater clarity of thought, greater creativity, and greater focus. Yet another possibility is the stimulation of the growth of additional brain cells (neurogenesis), or the unravelling and tidying of cluttered systems of brain pathways, akin to the improved performance that can apply in a computer hard disk after it is defragmented.

Comparative studies of people with genius characteristics, that explore any salient differences in genetic makeup, micro-anatomy, and mental psychology, may provide further ideas for how brain power can be boosted.

Finally, the simplest way to magnify the power of individual human brains may be for people to learn to work more closely with personal, wearable, and embedded computers - computers than can be regarded as forming a "neo neocortex", or an "exocortex". Rather than relying on memory inside the brain, people can recall information from this exocortex. Rather than evaluating arguments inside their own minds, they can take advantage of analysis modules in their exocortex. Rather than struggling to translate in their minds from one language to another, they can rely on their exocortex to assist them - and so on.

In principle, software in an exocortex could operate like someone's guardian angel - alerting them promptly to both good and bad opportunities. Just as present-day software often warns computer users not to click on various links, or draws their attention to discrepancies in how some information is being presented to them, so might an exocortex highlight when politicians are speaking with duplicity, or when an argument has faulty logic. Similarly, an exocortex might gently draw our attention to when we're being inconsistent, or when we're ignoring evidence that runs counter to our favourite assumptions. However, users will be understandably nervous about the influence that an exocortex may gain over their lives - especially if the software comes from a giant hitech corporation.

Improving aggregate intelligence

As well as finding good advice from software in our exocortex, that helps us improve the calibre of our thinking, we can in principle find similar good advice from the communities of people that we meet. It's not just software that can act as a kind of guardian angel; our close friends and trusted colleagues can do the same.These people can observe our biases and mental blinkers, and can find effective ways to quietly draw our attention to salient facts and considerations.

However, just as with advice from an exocortex, there is a risk that the advice we receive from our friends and colleagues will itself be biased. Our friends and colleagues may drag us into sharing their own prejudices and mental blinkers. Instead of the desired "wisdom of crowds", we could end up, all too easily, with the collective stupidity of groupthink.

For this reason, it is important to highlight and respect the principles which have been pointed out as improving the likelihood that an aggregate intelligence results in positive rather than negative amplification. One such principle is that the group should include independent thinkers, who can formulate their own opinions, rather than all the thinking in the group being dominated by a few loud or powerful voices. A related principle is respect for a diversity of opinion - the group should be able to hold in its collective mind two sets of ideas that initially appear starkly opposed, whilst searching for a new perspective in which the individual truths of these opposing opinions can both be accommodated. Further, the overall methodologies of science should be followed, where possible, including blind tests, replication of experiments, public disclosure of sources and analyses, and peer review. Bear in mind that science progresses, not so much because individual scientists are always paragons of scientific thinking (they often are not), but because members of the overall scientific community exert checks and balances on each other.

Just as an exocortex of personal, wearable, and embedded computers can enhance the thinking capabilities of individual humans, a collective exocortex of shared computing resources can enhance the aggregate thinking capabilities of communities of humans. Shared knowledge repositories such as Wikipedia provide a very useful service in gathering the most up-to-date information known to the group. Educational videos on YouTube spread the latest ideas about all manner of different topics. Documents support having multiple simultaneous editors, with software systems that suggest merges of different text in order to resolve apparent conflicts between different contributors. Maps are updated in real-time, highlighting changing traffic conditions, and suggesting alternative routes of travel. And as for geographical journeys from location A to location B, so also, in principle, for projects that seek to move a group from state-of-affairs X to state-of-affairs Y.

But just as an individual exocortex is vulnerable to hacking and distortion by antagonistic forces, our collective exocortex is vulnerable to being overrun by fake news and other sinister manipulation. The advice given by our community on the best way forwards may turn out to lead us in a very bad direction instead - one that benefits only a narrow elite, rather than the community as a whole. We already live in an era of intense cyber-manipulation, as countries deploy large teams of sophisticated hackers to subtly warp and frustrate the infrastructures in each others' countries. As technology becomes more powerful, an arms race between manipulators and defenders is likely to keep on escalating.

Given the increased potency of hacking and distortion, there's a temptation to respond in like kind - to fight back, generating more information, more analysis, and even some counter-distortion. Increasing the availability of better information and better analysis is, indeed, highly desired. Nevertheless, another initiative needs to run in parallel. That's the initiative to generate an abundance of *emotional* intelligence - a set of characteristics that reduce our likelihood to fall victim to devious tugs on our emotions.

Improving emotional intelligence

Transhumanists see it as an imperative to assist everyone to become free from forces of manipulation and distraction. As individuals, and as a species, so long as we are still dominated by various inner demons, we cannot soar to higher levels of flourishing.

Rather than us jumping precipitously, to grasp at some alluring possibility that will actually ensnare us and diminish us, we need the emotional intelligence to weigh things up more calmly and fully. Rather than giving all our attention to the cleverest, most loquacious advocates, we need the inner wisdom to distinguish veneer from substance. Rather than hungrily devouring whatever new fancy is put in front of our eyes, we need the self-awareness to defer gratification until we are justly confident the offering will nourish us rather than poison our soul. Rather than giving vent to anger at the manifest inequities of our situation, we need to be ready to breathe in slowly - even, on occasion, to turn the other cheek (in the Biblical expression) - whilst evaluating the best way forwards for everyone's benefit.

If our emotional intelligence remains retarded, the other aspects of intelligence - the ability to figure out how to accomplish goals - won't be a good servant to us. Instead, we'll just become better able to pursue emotionally retarded goals.

Emotional intelligence involves an awareness of the emotions in ourselves and in others, and the ability to control and harness the emotions in ourselves and in others. Without such skills, we are liable to be carried along by emotional forces outside of our full understanding. Whilst they are carrying us, we may for a while feel relieved or even exhilarated, but they are liable to strand us far away from the path of overall human flourishing. In an age with awful weaponry easily at hand, these forces are liable to lead us to the devastation of hostile war, to the destruction of our habitat and to many who are dear to us, and to a humanitarian catastrophe.

Emotional intelligence will give us the courage to admit to ourselves that some of our former decisions and alliances were, on reflections, wrong turnings. Emotional intelligence will generate an atmosphere of compassion in which we can all forgive ourselves, and each other, for mistaken crusades - for having campaigned for political or ideological leaders whose "solutions", we can see in the cold light of day, were seriously flawed. Emotional intelligence will help us to find and elevate good intentions in each other, despite the flawed causes to which these good intentions were applied. Emotional intelligence will give us the peace of mind to rededicate ourselves in directions we now understand to be better routes forward.

Throughout history, various methods have been pursued to improve our emotional intelligence. As mentioned earlier in this chapter, they include meditation, yoga, sufficient rest, good music, healthy food, and education. Just as these methods can be enhanced by twenty first century science and technology, to become more effective at improving our overall intelligence, so also can be they be enhanced to improve our emotional intelligence. The sooner, the better.

Throughout history, it has also been the case that emotional intelligence tends to suffers in an antagonistic social environment. The more that we feel exploited and alienated, the harder it is for many of us to "turn the other cheek" and transcend the self-defeating urges of our various inner demons. That's why the quest for an abundance of all-round intelligence is tied to the quest for an abundance of collaboration - the quest for a better social contract. That theme returns in later chapters.

Beyond the profit motive

There are spectacular profits to be made, by companies who provide systems that can improve

intelligence. That's why most of the world's wealthiest corporations these days are in the high tech sector.

However, there is no automatic correlation between greater profits and the fastest route to profound overall intelligence. Nor is there an automatic correlation between greater intelligence and greater human flourishing. In particular, concerns about safety frameworks may be perceived as eating into short-term profit opportunities. Moreover, without wise guidance from society as a whole, businesses are likely to place insufficient attention on the elevation of better emotional intelligence. Finally, too many companies may focus on improving narrow AIs, rather than on the potential disruptive breakthroughs needed to achieve AGI.

For all these reasons, transhumanists seek to urgently alter the public discourse about the potential for abundant intelligence. It's not a question of greater intelligence for the sake of greater intelligence - just as it's not a question, more generally, of enhanced technology for the sake of enhanced technology. The goals of human flourishing need to be kept fully in society's mind at all time. That includes the goals of greater creativity and greater exploration - the subject of the next chapter.

10. Towards abundant creativity

As discussed in the previous chapter, greater machine intelligence and task automation have the potential, not only for triumph, but also for disaster. Alongside their potential to positively accelerate human flourishing in multiple spheres, these systems also have the potential to malfunction, catastrophically. These systems might give rise to what has been called "killer robots" - automated agents that unexpectedly kill vast numbers of people.

A key complication with killer robots is that it may be difficult ahead of time to appreciate the full extent of the dangers they pose. There could be an initial period in which automated systems demonstrate apparently smart decisions and stunning improvements in operational effectiveness. These systems could be involved, for example, in creating remarkable new medical cures or novel mechanisms to extract greenhouse gases from the atmosphere. During this period, human observers would come to feel confident about the technology involved - and about increasing the resources at the disposal of automated agents. But this could be a prelude to these systems veering badly off course, in an adverse reaction to some unforeseen circumstances. Adverse outcomes could include an all-consuming escalation of fake news, a meltdown in our global electronics infrastructure, the inadvertent destabilisation of the entire planetary climate dynamics, or an accidental nuclear holocaust. The confidence that humans had developed in machine intelligence and pervasive automation would turn out to have been utterly misplaced.

This category of existential risk evidently needs far-sighted management, via, amongst other measures, the rapid development and wise enforcement of lean safety frameworks. But killer robots are by no means the only major concern raised by the growth of machine intelligence. We also need to give serious consideration to the possibility of "*job* killing robots" - automation that performs workforce tasks much better than humans, and deprives humans of employment.

Both sets of threat need to be assessed and managed in parallel. To add to the considerations of the preceding chapter, the present chapter looks at the threat posed to options for human employment by greater machine intelligence and more pervasive automation.

The threat from job killing robots may be viewed as less cataclysmic than the threat from killer robots. However, as this chapter highlights, the way society responds to huge numbers of people being deprived of employment could itself trigger a spiral into an increasingly tragic outcome. As such, there are no grounds for complacency. At the same time, there are grounds for real optimism too.

The opportunity for creativity

The threat to employment from automation has long been foretold. Up till now, these predictions seem to have been premature. However, the closer AI comes to AGI - the closer that artificial intelligence comes to possessing general capabilities in reasoning - the more credible these predictions become. The closer that AI comes to AGI, the bigger the ensuing social disruption.

How will humans cope, if their income from work is materially reduced, or perhaps disappears altogether? Should greater automation be feared, resisted, or slowed down?

To state the conclusion: rather than fearing this development, transhumanists look forward to the greater freedom that it can entail - greater opportunities for all-round human flourishing. Humans will no longer need to invest such large portions of their time in occupations that are back-breaking or soul-destroying. We'll be able, instead, to participate in the creation and exploration of music, arts, sports, ecosystems, planets, and whole new universes. This will happen because the immense bounty from greater automation will contain plenty for everyone's needs.

But before this abundance of creativity can be attained, some significant adjustments are needed in the human condition - changes in mindset, and changes in our collective social contract. These adjustments will be far from trivial. A great deal of inertia will need to be overcome, en route to realising the full benefits of improved automation.

A short history of automation

Robots have been killing jobs, on noteworthy scale, since the first Industrial Revolution. Weaving machines were invented that could automate many of the tasks in the textile industry better than human weavers. Machines that drilled, hoed, rotated, or reaped dramatically changed the work of agricultural labourers. Assembly-line machinery transformed the work that needed to be done in factories. Word processors and spreadsheets – robots of a different kind – reduced the need for manual clerical staff. And that’s just the start.

A key response to various jobs being automated has been for people to learn new skills and change their occupation. Human labourers who used to work on farms moved to cities to find employment in factories, shops, restaurants, hairdressers, banks, and so on. Individuals have been able to use their brainpower to learn all kinds of new skills, enabling them to work in careers that could hardly have been imagined in earlier generations - including occupations such as software engineer, website designer, and social media coordinator. Other occupations,

such as writer, teacher, farmer, and soldier, remain in existence, but are much transformed from earlier times; many tasks that previously consumed a lot of effort from the professional are handled nowadays by tools.

Will that same pattern be continued into the future? As robots and software improve in performance, will humans continue to be able to find new jobs for themselves, to replace jobs that have been made redundant due to the latest waves of automation?

That assumption can be called "the business-as-usual extrapolation" - the view that the future will remain broadly similar to the past. However, transhumanists recognise that the near future could see changes that aren't just minor variations on the past; instead, technology has the potential to change matters fundamentally. In that case, extrapolating from the past into the future is subject to being undermined by major disruptions ahead.

Limitations to business-as-usual

The business-as-usual extrapolation assumes the existence of “uniquely human talents” which can continue to give humans a competitive edge in the employment market place over whatever automation can accomplish. In that case, so long as we humans are willing to be adaptable and to retrain, we’ll keep one step ahead of the robots.

Consider traits such as compassion, emotional awareness, personal coaching, creativity, concept formation, “common sense”, and intuition. These human traits may appear to involve features beyond mechanical computation. And consider the kinds of ad-hoc skills needed by a professional such as a plumber, for whom each new plumbing repair task might involve an unpredictably different configuration of pipes, valves, cupboards, and household goods cluttering up these cupboards. What kind of robot could deal with all that variety?

Transhumanists respond that, despite possible first impressions, there is no good reason to believe in "uniquely human talents" that are somehow forever beyond the ability of computers to duplicate.

Indeed, first impressions can be deceptive. Just because software, today, cannot perform a particular task, it does not mean the task will forever remain outside the reach of software. There has been a long history of tasks which initially appeared to be fundamentally beyond the capability of automation, but which were subsequently demonstrated as within the power of automation after all. One famous example is the task of driving a car. Arguments used to be given that driving a car was inherently too difficult for any software to accomplish. However, over the last fifteen years, enormous strides have been made in the ability of self-driving cars. The question of self-driving cars has moved from an "if" to a "when".

What lies behind much of these improvements are advances in computer hardware: faster processor clock-speeds, larger memory storage, and smarter, more numerous sensors. Even more significant is the set of ongoing enhancements in the software discipline known as “machine learning”, covered in the previous chapter.

Limitations to retraining

The business-as-usual extrapolation urges members of society to prepare to retrain more fully and more often than in the past. Such retraining may take some time, resulting, perhaps, in a temporary reduction in earnings. But with sufficient advance warning, employees can be encouraged to acquire new skills in parallel with still working on their old job. In principle, this will minimise the disruption they will face.

However, note that machine intelligence is a general purpose utility. Any improvements to the mechanisms for machine intelligence are applicable, not just to a single occupation, but to multiple different occupations.

Therefore, as robots are becoming capable of doing key tasks for Profession A, the same breakthroughs mean that robots are also becoming capable of doing key tasks for Professions B, C, D, and E. Truck drivers who lose their jobs because of improvements in self-driving vehicles may find that, by the time they have retrained to a new profession, robots can do that profession better than them as well.

Computer vision, which makes powerful use of machine learning, is one example of a general purpose skill. The same core skill that allows self-driving cars to reliably recognise objects crossing their paths will also allow workplace robots to reliably recognise objects passing through their environment.

The skill known as “common sense” falls into the same category. Common sense depends upon a large network of knowledge about real-world objects, including an understanding of humans and their motivations. Present-day chatbots, notoriously, display a low level of common sense. It’s easy to catch them out. However, it is only a matter of time before they improve to match the human level of common sense.

Even though we cannot be sure of the timescales, transhumanists can make the following prediction: it’s going to become increasingly hard for humans who are displaced from one job by automation, to quickly acquire new skills that will allow them to carry out a different job that has no short-term threat of also being automated. Therefore, sooner of later, more and more people are going to find themselves unexpectedly out of work - or if not unemployed, then underemployed. Without an adequate safety net, their standard of living is likely to fall.

Limitations to robot-human partnerships

Robots have the potential to operate reliably, without getting tired, inebriated, distracted, or annoyed. They can communicate their learnings to each other, via “the Internet of Robots”, with the result that they can all benefit from the new experiences and insights of any one of them. Add in fast-increasing expertise in soft skills such as emotional intelligence, to their powerful computational skills and robust mechanical strengths, and robots become very attractive as replacements for temperamental human employees. Employers concerned about costs and about quality are bound to consider hiring fewer humans and more robots.

Defenders of the business-as-usual extrapolation acknowledge that automation will grow in prominence in the workforce, enabling cost savings and therefore higher profits. The defenders

of this extrapolation suggest, nevertheless, that one result of these larger profits is that humans will be able to work in partnership with the robots that are introduced.

When costs reduce for part of a task (due to automation), companies can provide their products and services more cheaply, reaching lower price points than before. With larger sales volumes, overall profits can end up higher (even if unit sales prices are lower). In principle, companies can take advantage of these increased profits to hire a larger number of human employees. These humans won't be doing the same tasks as the robots, but will be spending more of their time on the 20% (say) of their original job specifications which, at present, cannot be automated. These humans could also get involved doing new types of task that add even more value. In short, automation that destroys some jobs could actually create more jobs overall. That's what has happened in the past, and - according to the business-as-usual extrapolation - it will happen in the future too.

However, human partnerships with robots in the workplace are likely to pass through two phases. Initially, the combination results in productivity savings which allows business growth that in turn provides extra opportunities, overall, for the humans in the partnership. But in due course, the remaining tasks that the humans were performing will fall under the reach of improved robots, so that the opportunities for these humans in that workforce decline again.

As robots improve their general purpose skills, the second of these phases is likely to dominate the overall story. We humans will find fewer jobs available to us. Our contributions to robot-human workplace partnerships will diminish and diminish.

Examples of two-phase transformations

Three examples illustrate this two-phase transformation of employment opportunities.

First, consider weavers, over a two hundred year trajectory from around 1780 to 1980. The increasing automation of many of the tasks involved in that industry led to predictions of significant job losses. Famously, the Luddite movement of the early nineteenth century feared that the skills which craftspeople had taken years to learn would lose all their value as a result of wider adoption of knitting machines and automated looms. The movement took its name from a shadowy figure called Ned Ludd who, it was said, had smashed two mechanical knitting machines in 1779 in a fit of rage. The Luddites went on to engage in more systematic wrecking of automated textile machinery. In response, the British government made it a capital offence to destroy such machinery.

The Luddites were correct in their prediction that many parts of their jobs would be replaced by machinery. By some estimates, more than 95% of the tasks that used to be done manually by weavers came to be automated as the nineteenth century progressed. However, in seeming confirmation of the "business-as-usual" hypothesis mentioned earlier, the number of people employed in the weaving industry throughout that period continued to grow.

Here's why. Automation resulted in cheaper fabrics - cheaper cotton, lace, wool, and so on. Items made from fabric came down in price too. These lower prices meant that, compared to before, people bought many more items of clothing, carpets, curtains, rugs, etc. Previously,

many people had made do with only a few sets of clothes; but now it was economically feasible for them to fill out larger wardrobes.

In other words: greater automation enabled cost reductions, boosting consumer demand, in turn expanding the size of the industry. In parallel, the greater wealth generated by the industry circulated in the economy and enabled people to buy even more goods. So the employees in that industry still had plenty of work to do - as long as they were willing to adjust the set of tasks in which they were skilled.

For decade after decade, the same trend continued. However, from around the 1920s onward, the number of people employed in weaving started to decline. The trend had gone into reverse. By the 1970s and 1980s, technology had progressed to such an extent that there were no significant tasks left for humans to undertake. The phase of automation-driven increase in employment was superseded by a phase of major automation-driven *decrease* in employment.

The same two phase pattern is playing out, more quickly, in the second example, namely the banking industry.

Consider the “hole in the wall” automatic teller machines (ATMs) which allow visitors to banks to withdraw money without having to interact with a human bank teller. ATMs have the advantage of being available 24x7. Banks therefore no longer need to employ humans to perform the comparatively menial task of counting out cash and handing it over to customers. Does this mean that banks employ fewer humans? Once again, the answer has two phases.

Between the 1990s and the mid 2010s around 400,000 ATMs were installed throughout the USA. Over the same time, the number of people employed in banking actually grew. The automated tellers meant that individual bank branches cost less to operate. For the same overall investment, more branches could be opened, resulting in greater employment for human staff - provided employees were skilled in tasks beyond the labour of counting out banknotes. Again, this seems like vindication of the business-as-usual hypothesis.

But automation doesn't stand still. ATMs were only one step in a journey towards full automation of banking services. Bank account holders increasingly use online banking, including mobile banking. Where interactions with humans are still needed, this can take place via chat windows or video dialogue, in which a smaller number of centralised staff are sufficient to meet user demand. The wave of expansion of the number of bank branches is coming to an end. Indeed, local bank branches are consolidating and closing. A phase of automation-driven increase in employment is being superseded by a phase of major automation-driven *decrease* in employment.

Supporters of the business-as-usual hypothesis may reply as follows: all that employees need to do in such a case, is to move from one industry to another. Employees could move from the weaving industry to the banking industry. If they are sufficiently nimble, they can move in the same way from the banking industry to yet new industries.

However, as technology gains applicability in ever more general fields, the options for employment in new industries will reduce. This point is illustrated by the third example, which

features the changing employment opportunities, not for humans, but for horses.

A prediction that horses would in due course be made redundant by automation can be found as early as 1829, in a painting by social commentator George Cruikshank. The picture has the title "The horses, ‘going to the dogs’". It shows a group of horses staring in dismayed disbelief at a steam-powered coach propelling itself forwards. That coach had no need for horses to pull it. Nearby, one dog seems to ask another, “What do you think of this new invention?” The imagined response: “Why, I think we shall have meat cheap enough”. In other words, the invention of steam engines would leave horses without work, and they would be turned into dogfood instead.

For the next eighty years, however, that appeared to be a *bad* prediction. The population of horses in the United States soared from 4.3 million in 1840 to 27.5 million in 1910 – growing at twice the rate of the human population. The railways, rather than diminishing job opportunities for horses, actually increased them. The rise in steamships had the same effect. In both cases, the new transport systems increased demand for *local* transport - transport of people and goods to and from the railways and the steamship ports. Horses were vital to this local transport.

Throughout these same decades, there was lots of speculation that steam power would *eventually* extend beyond the railways and waterways into ordinary roads and fields. However, critics could point to seemingly major obstacles. Steam engines were extremely noisy; their great weight would damage paving stones or normal roads; their foul-smelling smoke was intolerable; and the sparks they scattered posed a serious fire risk. Such contraptions would never be able to find a foothold outside of controlled environments. Horses, it seemed, had unique skills that robots (automobiles) were incapable of matching.

In due course, however, employment opportunities for horses did plummet. From a peak of 27.5 million in 1910, the US equine population declined to 16.7 million in 1935, 7.6 million in 1950, and just 3.1 million in 1960. That’s a reduction of nearly 90% over fifty years. The cause of this change in fortune was a set of additional waves of automation, going beyond the steam engine: innovation featuring the petrol-driven automobile.

A famous photograph of the 1903 Easter Sunday parade on New York Fifth Avenue, taken by William H. Zerbe of the New York Tribune, shows a multitude of horse-drawn buggies. Careful study of that photograph reveals a solitary example of a new-fangled horseless carriage. But step forward just ten years to a photograph of the 1913 Easter Sunday parade by George Grantham Bain, shot from almost the same location. This time, the numbers are reversed. The scene is awash with motor cars, with just one solitary horse in view. In the intervening ten years, Henry Ford’s innovations in assembly line production had radically lowered the purchase costs of automobiles.

When circumstances are ready, changes that have been long envisioned yet also long delayed can take place surprisingly quickly. A long phase of automation-driven increase in employment opportunities can be superseded by a rapid phase of major automation-driven decrease in employment opportunities. Keeping on moving from one industry to another, to stay ahead of such changes, is an approach that will itself come to an end. As for horses, so also for humans.

Transhumanists urge that, even though the timing of that future phase transition remains unclear, society should prepare itself for such a transformation. With sufficient forethought, the outcome should be very positive. With insufficient forethought, all manner of chaos and destruction is possible.

There's good reason why this forethought is needed straightaway, urgently, rather than being left to some time period perhaps decades in the future. The reason is that the early warning signs of the impending phase transition are already in our midst. These signs are manifesting, not so much in *un*employment, but in *under*employment - and in all the associated instability.

The rise of underemployment

As mentioned, increasing automation can result in two kinds of consequence for employees. In the first case, employees are made redundant by robots, and cannot find any new job. As a result they become unemployed - experiencing "technological unemployment". In the second case, they *are* able to find a new job, to earn *some* money, but that job falls below their previous expectations in terms of work satisfaction, intrinsic interest, and income level. In this case, the employee experiences technological *under*employment.

One cause of underemployment is when someone moves from being a comparative expert in one occupation, to being a relative beginner in a new occupation. As a beginner, their salary is reduced.

A second cause of underemployment arises from some blunt facts of economics. The jobs where businesses have most financial incentive to introduce automation are the jobs that are highly paid and where many people are employed. The cost savings from automating these tasks are greater than for other jobs. As these jobs are removed from the economy, it's no surprise if the jobs where people can still find work are jobs with comparatively poor pay.

A third cause of underemployment is that, more and more often, automation has the result that the tasks left for humans to do are *less* skilled than before. Consider the skills needed by a taxi driver. Until recently, taxi drivers were obliged to memorise and understand vast maps of roadways, and the typical traffic conditions throughout a city at different times of day. Taxi drivers could literally spend years gaining the requisite "knowledge". But the advent of low cost satellite navigation mapping systems removes these skills from the requirements necessary for someone to be a good taxi driver. Since more people are now able to meet the reduced set of criteria, there's a drop in salaries they can command. It's the law of supply and demand in action.

These trends towards technological underemployment dovetail with other trends that have the same consequence of greater inequality in pay. These trends can be called "winner takes all" - or "winner takes most".

As firms acquire greater geographical reach, the firms that have the best products globally will find success in many more regions of the world than in previous eras. That's a consequence of greater connectivity, and of standards for consumer goods being harmonised the world over. Moreover, since technology makes a larger difference in whether a company's products and

services are top-class, the companies who own the key underlying technological platforms will gain ever greater dominance. Finally, compared to the leading industrial companies of earlier periods, the companies who own today's winning technological platforms owe their success to a relatively smaller number of employees. As work has moved from being labour-intensive to technology-intensive, employees that have core roles in the small number of dominant technology companies are paid very well indeed - but they are the exceptions.

Taken together, these trends help explain the phenomenon known as "the great uncoupling": industry productivity rates rise, but the median take-home pay stays relatively constant, or even falls in real terms. Key employees are rewarded handsomely, but the lower 60%-80% of incomes are stagnating.

For larger and larger numbers of employees, this is a disconcerting experience. The expectation of previous generations seems to no longer apply to them - the expectation that each generation would be better off in terms of income than the previous one.

The jobs that do exist, tend to be more transient and unstable than jobs of previous generations. Companies have shorter lifespans, so employees may find themselves unexpectedly out of work due to company failures. In response, more people set themselves up as "gig workers" or self-employed contractors, but in that case, they lose the social safety net of company-paid holidays and sick pay. Lacking evidence of a stable income, they find it harder to secure financing for home ownership, or for other loans to help them launch a new business initiative.

Other factors make the situation feel even worse. Gig work often takes place at awkward hours of the day or night, with little advanced predictability. Contract workers may be called away from their home and family on little notice, and may need to find short-term accommodation near to their clients - but such accommodation may be expensive, especially in cities where the wealthy 1-5% have driven up house prices. Finally, gig workers tend to lack the support of a trade union or other collective that would help them secure better working conditions.

Dangerous responses to underemployment

As noted, important parts of the traditional social safety net fail to apply for people working in the changed conditions in the wake of increased task automation. But whilst increasing numbers of people are finding themselves in precarious circumstances, the media bombards them with images of other people seemingly enjoying life as never before. For a highly visible subset of society, life appears to be full of marvellous material goods and mesmerising experiences. In contrast, for those impacted by technological underemployment, there's a growing sense of unfairness and alienation. They perceive that the best opportunities of life are passing them by. They perceive themselves to be victims of how society is changing.

These sentiments render the populace all the more prone to being swayed by fallacious theories about the causes of their predicament - theories that attribute their misfortune to scapegoats such as immigrants, internationalists, modernists, multiculturalists, far-off bureaucrats, and so on. It's time to take back control, they are told.

The sentiment is valid, but the courses of action recommended are frequently naive and

dangerous. Fast-talking Svengali figures evoke various fantastical visions of local sovereignty, of national destiny, of returning to a simpler past, of cultural homogeneity, of military glory, of religious revival, and of confounding the opinions of uppity experts. In their hearts, the populace often discern the drawbacks of these courses of action. But due to feelings of desperation, they think that it's nevertheless worth shaking up the whole political system. Against their better judgement, they allow themselves to be swayed by emotive distortions and base generalisations. Against their better judgement, they cast their votes for various demagogues and autocrats - people who claim they should be immune from the normal democratic processes of checks and balances. Against their better judgement, instead of gaining control, the populace actually loses control.

The problem, of course, is that jumps out of a frying pan frequently end up in a flaming fire. The problem is that the systems the populace seek to shake up are delivering a wide range of positive benefits they have forgotten to include in their calculations. The problem is that the courses of action which the populace hope will deliver them greater freedom are likely, instead, to drastically curtail their freedom.

To be more specific: we should beware any social transformation programmes that ignore the accelerating disruption caused by pervasive automation and machine intelligence. Such programmes are likely to cause more harm than good. Unless they directly address the challenges of tech-driven underemployment, political initiatives will waste time, distract attention, squander resources, damage social systems that should be part of the real solution, and store up an even greater sense of unfairness and alienation.

Rather than destroying existing social systems, transhumanists seek to evolve them and enhance them. The remainder of this chapter sets out a series of concrete proposals.

Embracing unemployment

The first area of our social fabric where positive evolution is needed is in the mindset that prevails throughout society. A key part of this mindset is the principle of reciprocity: if someone takes from the general pool of assets collected by society, they are also expected to contribute to that pool. This principle manifests in the opposition to "free riders" - people who avoid a fair share of effort, but who seek nevertheless to profit from the efforts of others. Societies around the world express hostility to free riders.

For the last few centuries, this principle has been expressed in ideas such as the primacy of paid employment: unless someone undertakes paid employment, they are a substandard person. Apart from short-term transitional periods, if someone fails to earn a salary, they should be reproached or scorned. Therefore, society should be structured to make it possible for everyone to earn a good salary. In this view, a society is defective to the extent that people cannot find employment that gives them a good salary. In this view, a sign of the success of a society is the measure of long-term unemployment: the lower that measure, the better.

Transhumanists put the emphasis differently. The requirement for people to seek paid employment belongs only to a temporary phase of the evolution of human culture. The basis for societies to be judged as effective or defective, is not the proportion of people who are in long-

term unemployment, but the proportion of people who can flourish, every single day of their lives. Paid employment has been an important part of the toolbox of methods by which people were enabled to flourish, but that's soon going to be a thing of the past. The relative importance of paid employment is set to decline. Unemployment will no longer be something to be feared. Unemployment will become something to be embraced, due to the extra time it will provide everyone to pursue activities of their own choosing.

As outlined in Chapter 4, "Principles and priorities", the concept of flourishing is wide-ranging. Flourishing involves happiness, energy, nourishment, creativity, intelligence, health, collaboration, and awareness. Flourishing is by no means a human monoculture. Over time, as our collective intelligence increases, our collective understanding of the components of human flourishing will surely evolve. But as it evolves, that understanding is unlikely to bring back the requirement that, unless someone has paid employment, they are unable to flourish. Any such connection is becoming archaic.

Psychologist Abraham Maslow postulated a “hierarchy of needs”: alongside our physiological needs we also have deep-seated desires for friendship, belonging, intimacy, esteem, and meaning. There is scope to debate which needs might be the most “important”. But there's no arguing that there's more to life than food, clothing, and shelter. The goal of society, therefore, is to enable all citizens to access these wider, deeper areas of life. If that goal can be achieved with only a small minority of people doing something akin to the paid employment of today, that's no reason to complain. If that goal can be achieved due to an increase of machine intelligence and task automation, we should welcome that outcome.

But critics may wonder: isn't paid employment required for people to find deep meaning in their lives? Or to keep them from frivolous idleness?

The transhumanist answer is that there is plenty of deep meaning to be found, in the pursuit of activities for which no payment is offered.

Consider activities of creation and exploration: Consider the creation and exploration of music and other forms of art. Consider the creation and exploration of games and sports. Consider the creation and exploration of engrossing narrative dramas - historical sagas, detective investigations, sweeping romances, and so on. Consider the creation and exploration of vast alternative theories of mathematics and physics. Consider the creation of exploration of ecosystems of lifeforms - biological and nonbiological - in virtual worlds, planets, and possibly in whole new universes. Consider creative mergers between the above activities. Consider new forms of activity which we can hardly imagine today, that will be enabled in the future by our intelligences operating at much higher levels than at present.

That's the world to which transhumanists aspire - not a world measured by the degree of paid employment.

Other critics might wonder: what's the point in humans exploring some activity, if an advanced AI could explore that same activity so much faster and more fully than us?

But that's already the situation in which we almost all live. The activities from which we currently

find great pleasure are ones where our performance is almost certainly inferior to that of experts. Sports hobbyists can find pleasure from playing their favourite games, even though their level of performance is far beneath that of the world's best players. Individuals can find joy in solving a logical or numerical puzzle, by themselves, even though software could find the solution much more quickly. We take joy from our creations and explorations, because they belong to us and involve us: they reflect our present character and achievements.

How much more joy will there be, when we are freed from the adverse limitations which presently oppress us so extensively - our flawed bodies, flawed minds, flawed emotions, and flawed social structures? How much more joy will there be, when we can immerse ourselves in an abundance of creativity?

Yet other critics may refer back to the principle of reciprocity. Isn't it unfair that many people will be able to devote themselves to pursuits of creativity and exploration, without having earned money to pay for such pursuits? Aren't these people free riders, who threaten the overall well-being of society, taking things they have not paid for? Transhumanists respond that such concerns are remnants of the culture of scarcity which has prevailed up till the present time. These concerns lose their impact in a world of abundance. Let's delve deeper.

Towards zero prices

It's one thing to contemplate a future social model in which pervasive automation generates sufficient clean energy, nutritious food, affordable consumer goods, and so on, to meet all human needs. It's another thing to propose policies for the here-and-now, and for the interim period of uncertain length between the present day and the full advent of sustainable superabundance. During that transitional period, elements of scarcity will coexist with elements of sustainable abundance. During that period, hard questions need to be answered, about access to resources that are scarce and expensive. During that period, there need to be limits on what free riders can do.

Traditionally, the mechanism of the market economy has been used to determine access to scarce resources. Whoever is able and willing to pay for something, can access it. Whoever is unwilling or unable to pay, is unable to access it.

In practice, all countries have exerted constraints on the operation of the market economy. Done well, these constraints prevent the emergence of monopolies or cartels, in which the needs of a small group of companies drive up prices higher than if a free competitive market applied. Applied well, these constraints prevent monopolies or cartels from unfairly shutting down innovative competitors. Applied well, these constraints ensure that proper consideration is taken of "externalities" - the side-effects of market transactions, whether positive (public benefits) or negative (public harms). Applied well, these constraints set a strategic direction for the direction of the economy, allowing coordinated action on longer-term initiatives outside the organisational competence of the market alone.

Transhumanists recognise the great benefits that can arise from a well-managed market economy. Transhumanists can endorse an additional high-level strategic direction: *prioritise the reduction of prices for all goods and services that are fundamental to an agreed base level of*

human flourishing.

The end target of this strategy is that all goods and services fundamental for human flourishing should have zero price. There are three policies to advance towards this target. First, the policy of reducing prices, step by step. Second, the policy of increasing public subsidies, to alleviate whatever prices remain in place, with the goal that subsidies will fully cover the remaining costs. Third, the policy of starting with an agreed very basic set of goods and services, and step-by-step extending this set.

Reduced prices are something that the market economy often delivers, without the intervention of any policy directive. However, there's nothing inevitable about such an outcome. Prices sometimes rise rather than fall. Hence the need to explicitly set the direction.

Without an explicit focus on reducing prices, other factors are likely to receive a greater share of market attention. In particular, there may be more attention on raising the GDP - the "gross domestic product" which adds up all the financial transactions in the economy. When the GDP is perceived as falling, political leaders often come under pressure to make amends. Other things being equal, that leads to efforts to increase prices, rather than to reduce them. It's a direction that arguably made sense, in previous phases of human history. But it's not the direction that will lead to abundance.

It's the same with pressure for the share prices of companies to increase. Higher share prices imply more wealth for shareholders. Countries celebrate the companies that push up the performance of various aggregate stock exchange indices. Again, it's the wrong focus. Newspaper columns, journalist articles, should instead give greater prominence to whether basic goods and services are becoming more affordable. Where particular goods and services that are key to positive human flourishing are too expensive, this fact should be widely discussed, so that society as a whole vigorously explores potential solutions. No stone should be left unturned in this quest.

Beyond today's social safety net

Consider again the concept of the primacy of paid employment: unless someone undertakes paid employment, they are a substandard person, who should be reproached or scorned.

Of course, societies already make many exceptions to this concept. Basic pension payments are provided to all citizens, so long as they are old enough, without them needing to continue working. Basic educational funding is provided to all citizens, within certain age boundaries, even if they have not started paid work yet. Basic healthcare treatment is, this time with no age limits, provided free of charge to all citizens, whether or not they have paid employment. And when someone has lost their job, public funding is available, for a while at least, to help them as they look for a new job.

As another exception to the primacy of paid employment, family members frequently look after one another. Larger groups of mutual assistance "friendly societies" developed in many cultures around the world, in which resources were pooled, in order to assist members of the group who had special needs.

Transhumanists applaud this spirit of mutual support. It aligns with the transhumanist principle of active neighbourliness: *treat others in the way we would ourselves like to be treated, if we were in the same situation.* We recognise that we, ourselves, would benefit from that kind of social safety net; so we recognise the merit of providing it to others.

Without a social safety net, a powerful spirit of apprehension can arise. The fear of becoming detached from the basic means of human flourishing can cause people to become narrow-minded, grasping, and self-centred. The fear of losing out generates resentment and bitterness. It drives people into a scarcity mentality, in which any gain by some members of society is seen as requiring others in society to suffer exploitation. Adverse effects follow, not only in personal wealth, but in personal health. Adverse effects follow, not only in self-esteem, but in the quality of social relationships.

For these reasons, transhumanists look forward to a stage-by-stage *enhancement* of the social safety net - the development of a safety net that provides all members of society with access to the goods and services that are fundamental to an agreed base level of human flourishing. And transhumanists look forward to that agreed base level being elevated, stage-by-stage, as the prices of goods and services approach zero, and as society generates an ever greater abundance covering all seven spheres of human life.

For those goods and services which carry prices above zero, the question inevitably arises as to how these costs will be met. The answer is via a dividend distributed from the shared commons of humanity's accomplishments. This is addressed in more detail in the next chapter. But first, two other topics deserve attention.

Beyond Universal Basic Income

One proposal that is heard more and more often is that society should distribute to all citizens what is known as a UBI - a Universal Basic Income. This can be viewed as an extension of the existing social safety net of pensions, childcare allowance, unemployment benefit, and so on. What makes UBI different is that these payments are given universally - to every citizen - without any test being applied as to whether the citizen deserves that payment. For example, no reduction is made in the payment, if someone who was previously unemployed now starts paid work on a part-time (or full-time) basis.

UBI is said to have a number of advantages. The simplicity of the payment system would allow a reduction in the administrative overheads of the existing highly complex benefits systems. The unconditional nature of the payment removes the current adverse incentive feature of existing benefits systems: existing systems discourage people from starting new paid work, since they would lose their current welfare payments. UBI avoids the debilitating stress that many people experience, for fear that their benefit payments will be removed. UBI treats recipients with more dignity, since it leaves to each person the choice of how they will spend the money assigned to them. Finally, UBI prepares for a near-future world of widespread technological unemployment, and alleviates the current issue of increasing technological underemployment.

Transhumanists can advocate a UBI as forming *part* of a positive societal response to growing machine intelligence and task automation. But UBI is not the entire answer. Advocacy of UBI

needs to be complemented by society prioritising reducing the prices of basic goods and services.

Moreover, the system of single universal payments will need to be accompanied by other social safety net mechanisms that address particular needs of particular subgroups of citizens. This may include, for example, special healthcare needs for people with various disabilities, and special subsidies to cover housing costs in particular locations. Having this variability in the system, as opposed to everyone receiving the same payment, is an important step towards the total payments being affordable.

One alternative to UBI is UBS - Universal Basic Services - in which society provides various services (such as transport and education) free of charge. The UBS model is distinct from UBI, in which these services continue to have a cost, but money is provided to everyone that they can choose to spend on these services. Transhumanists have no intrinsic doctrinaire preference one way or another between UBI and UBS. Social safety nets are likely to include aspects of both.

Transhumanists are likewise open to the idea that some elements of a UBI payment would be conditional upon the recipient taking various potential actions in return. In other words, the UBI could be split into an unconditional element and a conditional element.

It's particularly important to plan for the social safety net to pass through a series of different changes, perhaps relatively rapidly. A system that is appropriate for a given geographical region in, say, 2022, may be far from the best choice for that region in 2027. Reasons for the changes will include: different workforce tasks being automated to new degrees; lower costs of various goods and services as abundance accelerates; changes in public mood and expectations; and insight gained from administering social safety net schemes - both locally, and in other geographies.

More time for creativity

Done well, a transhumanist social safety net will leave everyone with more time to focus on whatever tasks and projects they see, as individuals and as groups, as being the most interesting and most important. Done well, the social safety net will leave everyone with ample time for creativity.

This will extend and accelerate a trend that has been taking place for well over a hundred years: a decrease in the average number of hours an adult spends each week working for paid employment. With variations in different countries, this figure has declined from around 60-70 hours per week in the 1870s to less than 40 hours per week in the 2010s. Factors involved in this reduction include: fewer hours worked in the average working day; fewer working days in the average week (with a move from a one-day weekend to a two-day weekend); and greater numbers of days of holiday throughout the year.

One response, therefore, to the threat of underemployment - the threat of having fewer hours of paid work each week - is to *embrace* this move.

Yes, with this move, people will on average have fewer hours of paid work each week. For example, a three-day weekend will become more common; and eight-hour working days could be replaced by four-hour working days. This may well mean a reduction in salary payable. But this reduction can, again, be embraced, as part of a wider social transformation. Since goods and services will have lower costs, higher salaries will no longer be needed. And various forms of UBI (or similar systems) can help cover the gap between a reduced salary and the total costs of goods and services that someone wishes to purchase.

In parallel, changed social attitudes will put a higher priority on experiences gained from virtual reality worlds, rather than from expensive real-world projects.

Changed social attitudes will also enable goods and services to be shared more widely - rather than being owned exclusively by an individual. The so-called circular economy should thrive.

Finally, a very important change in social attitude will enable a fuller distribution of society's assets to all members of that society, rather than being restricted to the few people who still have paid employment and who might therefore feel they have a unique right to these assets.

This last change is an expression of growing abundance in the last of the seven spheres of human life covered by this Manifesto - an abundance of collaboration and democracy.

11. Towards abundant democracy

Over the next few years, many hard decisions need to be taken. These decisions will determine whether humanity can move forwards swiftly into the era of sustainable superabundance, or will instead collapse into a state of social chaos and humanitarian tragedy.

Examples of these hard decisions can be found throughout the preceding pages of this Manifesto. Consider: Which aspects of human nature should be changed, and which protected at all costs? In which circumstances should the precautionary principle override the proactionary principle? Which types of genetic modifications should be encouraged, and which discouraged - in food, in pets, and in humans? Which sources of energy (for example, nuclear energy?) should be developed and deployed most quickly? How much public funding should be allocated to the development of atomically precise nanofactories - and with what constraints? Which approaches to curing diseases of neurodegeneration should be prioritised? What restrictions should be applied to autonomous lethal weapons systems? How can surveillance of potential dangerous misuse of technology best coexist with protections for individual privacy? What kind of new social contract should be put in place - and how strongly should UBI feature in these plans? Which kinds of inequality and diversity should be celebrated, and which resisted? How can the various regulatory systems from different local markets be woven together into an effective international framework that prevents rogue elements from slipping dangerous goods and services through the cracks between these agreements? Which international alliances deserve greater support, and which should be avoided? And so on.

This chapter is not seeking answers to these questions. Instead, it is seeking to understand what *processes* we should follow, in order to find and defend the best answers to this kind of question.

These questions are all examples where disagreements arise between thoughtful, well intentioned advocates of different answers. The right answers are by no means obvious.

These questions are all examples, moreover, where the discussion is subject to deliberate distortion, by groups who have vested interests in steering the outcome in particular directions. That is, on top of the legitimate debate, a set of misleading arguments further undermines the ability of society to pick the best solution.

With sky-high stakes, it's critically important that cool heads can prevail. It's critically important that key flaws in reasoning are identified promptly, before decisions are taken based on these flaws. And it's critically important that the best insights of the whole community are heard and absorbed.

With sky-high stakes, there's a great opportunity to use technology to improve the calibre of decisions, and to ensure that decisions are properly followed through. But there's also a grave risk for technology to *worsen* the calibre of decisions, to hamper the implementation of good decisions, and to perpetuate bad decisions in the face of informed opposition.

In other words, technology raises the stakes even higher. Technology could enable the emergence of a superdemocracy, that will steer humanity more reliably towards the era of sustainable superabundance. But at the same time, technology could drive decision-making backwards. Rather than us moving towards collective enlightenment, technology could cause us to move towards collective confusion and collective impotence.

As this chapter highlights, the key task is to step forwards incrementally, gradually building improved capacity for a better politics. Out of these incremental improvements, an abundance of collaboration and democracy can emerge.

Two scenarios for the impact of technology

Alas, twenty-first century technology can have all kinds of detrimental impact on politics.

Twenty-first century technology can enable unprecedented large scale surveillance and manipulation of members of society by forces seeking extra influence. This manipulation can be subtle rather than blatant; that's what gives it greater effect.

Twenty-first century technology can power deft psychological techniques to frighten or incite people to choices that are different from their actual best interests. Desires to increase eyeball and click-through attention on social media result in posts that push people into emotional reactions rather than careful deliberation. Online interactions frequently propel participants to champion tribal instincts, cheering on pro-group “blue lies” rather than respecting objective analysis. With hearts on fire, smoke gets in the eyes.

Potential enormous impacts from cyber-sabotage, nerve agents, and various weapons of mass destruction (chemical, biological, and nuclear) - whether wielded by enemy states or by terrorist groups - raise tensions further, and risk driving politicians towards decisions that are more extreme and less considered.

But the same technology can have many very beneficial impacts on politics too.

In an extension to current technology that highlights misspellings or incorrect grammar in a document, new tools can highlight which factual claims have been assessed as false or misleading. Other tools can highlight logical flaws in arguments. They can also draw attention to cases where the provenance of data is suspect - such as when photographs have been edited, or videos synthesised, to give a false impression.

Twenty-first century technology can facilitate the systematic collection and analysis of information relevant to decisions, in ways that build on the successes of Wikipedia.

By analysing arguments, technology can in due course suggest new proposals that integrate different perspectives in compelling ways, and thereby help build bridges between opposing sides in a debate.

Technology can create and maintain vast virtual worlds - simulated environments - in which the potential outcomes of policy changes can be investigated in advance.

Finally, technology can assist politicians to deliberate more calmly on decisions, rather than being panicked into flawed decisions in tired or emotive circumstances.

In short, just as technology can have either a bad or good influence on society, so also it can have either a bad or a good influence on politics.

In both cases, the determining factor is the level of wisdom, strength, and agility brought to managing the technology. The more powerful technology becomes, the greater the need for wisdom, strength, and agility - the greater the need for clear thinking, and the greater the need to be ready to set aside previously long-cherished “instincts” or “identities”.

Preconditions for better democracy

Let’s be honest about the flaws in how humans form and hold viewpoints. With a better understanding of the psychological and sociological factors at work, we can liberate ourselves - and our fellow citizens - from the tyranny of distorted worldviews and unnecessary panics.

False information often changes outlooks in deep ways, so that opinions remain different from before, even after people have learned that the initial information was incorrect.

It’s similar to how the best advertisements not only change viewers’ preferences, but leave those viewers convinced that the changes arose from their own volition. The viewers may even forget they ever saw the advert - or deny that it had any impact on them.

Recognising the risks from false communications, society already imposes financial penalties in cases when advertising or financial information is misleading. Advertisers are penalised if they make claims that are demonstrably false. Companies are penalised if the branding or packaging of their products misleadingly imitate those of higher quality products from more reputable companies. Financial bodies are penalised if they prematurely release price-sensitive information to only a subset of investors - or if they manipulate their accounts to give a

misleading impression of their trading position.

Let's be ready to apply similar sanctions and penalties in cases when the political discussion is deliberately distorted by false information, when participants fail to declare vested interests, and when inflammatory publicity risks the impartiality of ongoing jury reviews.

In this way, we can improve the calibre of the overall discussion, reducing the elements of emotional subversion, and maintaining open minds to ideas that, although initially disturbing, can lead in due course to a better understanding.

At the same time, let's watch the regulators very carefully. In the wrong hands, sanctions against the free expression of opinion can silence voices that ought to be heard. For this reason, satire must be protected - but only where it is made clear that the statement was not intended to be taken literally.

The result of these steps will be to preserve the vital benefits of free speech, whilst avoiding social chaos from remarks that are irresponsible or incendiary. We can, and should, have free speech without the panic. That's an important precondition for better democracy.

No rights to no offence

To encourage a richer exploration of potential new solutions, it should never be a criminal offence to criticise an existing idea. Ideas should be able to stand on their own feet, without needing protection via the censorship of criticism.

If someone says they have been offended by criticism of an idea they personally hold dear, that is no reason to elevate that criticism into a crime. Instead, it's a reason to marshal good arguments in support of the original idea - or (if such arguments prove to be insufficient) to be ready to accept a new idea in its place.

Specifically, there should be no special legal protection for ideas declared to be religious, or foundational, or sacrosanct in some way. Ideas matching these descriptions have often proven to be major obstacles to the progress of human flourishing - even if, in earlier times, they had been forces (on balance) for the collective good.

Accordingly, laws on blasphemy should be removed from the statute book.

But at the same time, let's avoid any implication that huge numbers of people - people who nominally identify with the same religious faith - all hold rigidly to the same set of beliefs. Whilst being critical of particular ideas, let's be ready to build constructive bridges with people who may assert some appreciation of these ideas. The best insights often arise from a synthesis of viewpoints that initially appear to be polar opposites. Significant contributions to the establishment of sustainable superabundance will surely be made by people from religious communities of all hues.

Creative scepticism

To prevent ourselves being misled by clever arguments, let's spread far and wide a better understanding of the flaws in reasoning to which we can fall victim - the numerous cognitive

biases arising from the limitations of our biology and psychology.

This is no mere academic exercise. The vital skill of creative scepticism is something that can be strengthened via regular real-world practice.

Importantly, the practice of creative scepticism should transcend individual minds by embracing a collective dialogue. Participants can point out to each other - sensitively and constructively - the ways in which we as individuals remain enthralled by particular flawed ideas. With the benefit of multiple streams of feedback, we can gain the strength to overcome our individual cognitive flaws.

This is similar to how the institution of science can make collective progress, despite biases and prejudices afflicting individual scientists. Provided a sufficiently wide set of opinions is included, peer group review can allow the community as a whole to withstand distortive pressures.

Constraining political funding

At the same time as we encourage wider practice of improved creative scepticism, let's take steps to lower the volume of the misleading information that circulates within the public debate.

In part, this reduction can take place via the sanctions, mentioned above, against communications that are deliberately misleading.

Additionally, limits should be placed on the amount of money that political organisations can spend. In principle, this will cut down on the influence of any one organisation. To prevent organisations working around these limits, strong measures of enforcement will be needed.

Nevertheless, in an age of fast-changing technology, clever political operators will find innovative ways to have their messages spread further and wider, at little cost. For example, they could use mechanisms such as bot armies. In order to increase the prospects of a balanced debate, steps will be needed to identify, patrol, and curtail bot armies and the like - and to impose sweeping financial penalties on those found responsible for deliberately distorting the debate.

In this way, by turning down the *volume* of political communications, there is a greater chance of increasing our *comprehension* of what needs to be discussed.

Beyond party politics

For politics to become a positive force, it must shed the impression that its protagonists frequently speak contrary to their actual opinions. Politicians need to become regarded as authentic communicators, rather than being two-faced.

Whilst there are benefits (such as economies of scale) from like-minded politicians banding together into political parties, the current system of political parties has its own drawbacks - including the imposition of the "party whip" on matters of high contention.

Let's encourage politicians to speak their own minds, rather than having to follow party lines in cases where their own assessment differs from that of the party hierarchy.

Politicians will be more prepared to express their own authentic views, if they can more easily move between a greater number of different political parties, rather than being constrained to just two parties that (in some countries) are the only ones with a realistic opportunity to form a government.

Proportional Representation

Political systems with first-past-the-post elections pose unnecessarily high barriers of entry to newer ideas. Such systems are overdue reform. Systems with proportional representation allow more fluid introduction of innovative political forces.

Political systems with proportional representation often lead to the need for parties to form coalition governments. That consequence should be viewed as a benefit rather than as a drawback. After all, the skill of forming coalitions is a positive asset which should be nurtured.

Political systems with large constituencies in which more than one politician is selected, can combine the positive aspects of proportional representation with the advantage of links between constituencies and elected representatives.

The role of money

It can be argued that the need for money will disappear once all goods and services are available in abundance. Indeed, this Manifesto foresees an abundance of all goods and services needed for the basics of human flourishing. However, it's likely that various goods and services - including new inventions - will remain scarce and therefore "premium". Money will therefore continue to have a role in the longer term, in helping to mediate access to premium items.

In the shorter term, money has a critical role in the transitional period before the advent of sustainable abundance.

Money can be regarded as a technology with (like other technologies) both good and bad potential. We should not allow ourselves to be blinded by either the good or the bad. Accordingly, just as democratic review and supervision is needed for the operation of the free market, democratic review and supervision is also needed for the design and development of monetary systems - and for the design and development of banking.

The same assessment applies to the practice of creating money by "fiat", such as printing additional banknotes, or unilaterally increasing the quantity of currency in someone's account. Depending on circumstances, such steps can either assist or endanger the wellbeing of the economy. Such steps may well be queried and criticised, in individual cases, but should not be judged by any doctrinaire assessment as inadmissible in all circumstances. We need greater flexibility in our policy choices.

The role of cryptocurrencies

Cryptocurrencies are an innovative extension of traditional currencies, and can perform many of the functions of previous monetary systems. Despite being based on mathematics and computation, cryptocurrencies are no less "real" than previous currencies.

Cryptocurrencies potentially have the advantage of being free from value erosion or manipulation arising from actions by central banks or politicians. However, their value may be manipulated by other vested interests, such as has occurred in the case of Bitcoin.

By design, due to a fixed cap on the total number of units of the currency that can ever be created, some cryptocurrencies are immune from inflation. However, elements of inflation are on occasion useful in an economy, so this fact provides no reason to absolutely prefer these cryptocurrencies over alternatives. Instead, this Manifesto expects a number of different types of currency to coexist, including some that support inflation and others that don't.

Alongside their potential advantages, cryptocurrencies need to address issues such as large energy consumption (for "mining" and block verification), slowness of transaction speed, high transaction costs, and volatility of value. It is likely that various cryptocurrencies will emerge that address these issues in different ways. It's too early to foresee the full landscape of solutions. However, whilst the field deserves attention, it seems unlikely to be as revolutionary as various adherents have claimed.

Due to the potential for fraud, cryptocurrencies require democratic oversight - the same as for other financial systems.

The role of decentralisation

Centralisation has drawbacks when processes are unnecessarily delayed, or can be manipulated or exploited, by the involvement of intermediaries. Where possible, processes of all sorts should be redesigned to avoid intermediaries.

However, many processes, left to run by themselves, have a risk of causing damage to the participants, or to the wider environment. Accordingly, this Manifesto advocates democratic oversight of all significant processes in society. The two principles - a preference for decentralisation, and the need for democratic oversight - are evidently in tension with each other. This tension needs to be recognised and wisely managed.

The emerging technology of decentralised ledgers - including blockchain - provides many innovative opportunities for the redesign of social processes. These redesigns need to pay attention to questions of potential fraud and tax evasion, and to the risk of unstoppable "smart contracts" whose operation may subsequently be regretted by participants.

Reconsidering redistribution

One factor that frequently causes sharp political divisions is disagreement over the topic of redistribution - whereby resources are forcibly transferred from the wealthy in society to those who are less wealthy.

From one point of view, these transfers can be viewed as akin to theft. Entrepreneurs who have earned profits from their inventiveness and industry may resent the actions of tax collectors who extract a share of these profits.

However, a different viewpoint is preferable - the viewpoint of common ownership and public

dividends.

In this viewpoint, all humans stand as joint inheritors of the environment in which we find ourselves. This environment is rich in resources - including natural resources and social creations. Consider the electromagnetic spectrum, which can be used to transmit information. Consider the atmosphere, which absorbs the greenhouse gases arising from our activities. Consider the financial system, with its sources of credit, insurance, and other monetary services. Consider the findings of science, that allow us to build new technologies. Consider the land upon which we build houses and offices. Consider the roads and railways along which goods and people journey. Consider the patent system, that provides rewards for people who publish details of their inventions. Consider the schools and universities, that provide training which is useful to the employees hired by businesses. Consider the public libraries and other information sources, from which we can all benefit.

The accomplishments of entrepreneurs take advantage of many elements in this common environment. In some cases, entrepreneurs pay rent for these elements - fees allowing usage of land, bank loans, the filing of patents, road tolls, and so on. In many cases, however, no such rent is paid, for example when companies cause the emission of greenhouse gases, or undertake risky financial speculations that threaten the overall stability of our shared financial system. In yet other cases, the rent levied is arguably too small, since it fails to take proper account of adverse externalities from business activities.

There is considerable scope for increased rents to be levied on usage of these common resources. Shares of the earnings from these rents could then be made available to all citizens.

In some cases, rental charges should be fixed in advance - for example, fees payable when certain quantities of carbon-based fuel are introduced, or when short-term financial transactions take place. (These examples are often given the names "Carbon tax" and "Tobin tax".) Another example is when a portion of the electromagnetic spectrum is allocated for a particular use.

In other cases, the rents can be calculated as a fraction of profits earned. For these examples, the rents can be considered as dividends, paid not to private shareholders, but to all citizens, who are in effect public shareholders of each company.

This is not the same concept as full public ownership of companies. Instead, companies will continue to be managed by private individuals. However, the companies will become regarded as *partially* owned by the citizenry.

Similar systems of common ownership and public dividends already exist - for example, in the Alaska Permanent Fund, established from oil revenues, which pays annual dividends to every citizen of the state of Alaska. Consider also the Norwegian Sovereign Wealth Fund. Importantly, such funds tend to have wide bipartisan support where they operate.

The result of such funds is that members of the public receive money without needing to labour on any regular basis. But that's something that already happens widely, with no complaints from entrepreneurs. Investors provide upfront loans, and thereafter receive interest payments or other dividends, without needing to labour on any regular basis. Owners of property - whether

buildings, land, or intellectual property - receive payments such as license fees or land rental, without needing to labour on any regular basis. They receive regular payments on account of being owners of resources used by the entrepreneur. Once all humans are viewed as collective inheritors and owners of the shared environment in which we all live, it should become acceptable in the same way that at least some payments are made available to everyone, out of the profits generated by businesses operating in this environment.

This shouldn't be regarded as any kind of "theft". Instead, it should be appreciated as a sharing of the abundance which a business is able to generate.

Nor should this be regarded as the wealthy somehow "losing" so that the less wealthy can "win". Rather, it should be appreciated as enabling a much healthier society, freed from risks of destabilisation and anarchy - a society in which *everyone* is better off as a result.

Final questions for superdemocracy

Of course, there is still plenty of scope for legitimate debate over the sizes of the various rental charges mentioned in the above proposal. Careful deliberations are also needed over the best ways to distribute the resulting income throughout society: how much as part of a UBI, how much as subsidies for specific goods and services, how much into general public coffers, etc. These are key questions for regular superdemocratic consideration and review.

Another question along the same line is how much should be added to the public dividend fund from time to time by means of the central government creating so-called "helicopter money" by fiat, as proposed by economist Milton Friedman. Too much money creation causes inflation which can destabilise the entire economy and erode the value of individual savings accounts. However, if done judiciously, these additional payments can act as an important stimulus.

An outstanding major challenge for superdemocracy will be to cut down on tax evasion. Evidence from the Panama Papers indicates that many of the wealthiest individuals on the planet go to great lengths to avoid paying a fair share of taxes. These individuals hide their assets using all kinds of clever financial covers. It has been estimated that assets worth many tens of trillions of dollars are escaping tax via such mechanisms. Recovering the tax due on these hidden assets will require significant cooperation between different governments bodies around the world. It is no accident that many of the individuals concerned are actively seeking to destabilise systems of international legal cooperation, such as the EU.

In addition to combating the illegal measures of tax evasion, steps are long overdue to prevent companies from legally avoiding taxes by cleverly transferring the accounting of profits between different geographies. It is by exploiting loopholes in taxation systems that many of the world's largest companies manage to pay corporation tax at such a low overall level. Again, the solution to this challenge requires improved international cooperation.

In a world of fragile international relationships, the intergovernmental bodies set up after the second world war – the so-called Bretton Woods institutions such as the International Monetary Fund, the World Bank, and the United Nations – seem increasingly ill-suited to the complexities of twenty-first century life. Better methods of international cooperation are sorely needed, lest

conflicts between diverse national interests spiral out of control. The same principles that define superdemocracy on a national level are applicable to international decision processes too. These include informed deliberation, real-time fact-checking, transparency and accessibility, disclosure and accountability, delegated representation, active bridge-building, open-mindedness, creative scepticism, respect for data, agile experimentation, emotional intelligence, and the wise embrace of suitable technology.

Such systems need time to establish themselves. Ideally, transhumanist communities will blaze the trail, demonstrating processes by which hard decisions can be progressed in practice, in ways that benefit the entire community. That topic is addressed in the final chapter.

12. Options to engage

The forthcoming era of sustainable abundance cannot be built by wishful thinking. Nor will the activity of solitary individuals be sufficient.

Instead, what needs to emerge is a network of communities that coordinates and magnifies the activities of an ever growing number of participants.

This final chapter of the Manifesto highlights a number of categories of task where participation is welcomed - tasks which can help accelerate the advent of sustainable superabundance.

To express interest, join the mailing group https://groups.google.com/forum/#!forum/transpolitica.

Improve the Manifesto

The initial version of this Universal Transhumanist Manifesto is a good start, but it can surely be improved in many places.

Community members are invited to suggest ways to sharpen and simplify the content, to make the content more engaging, and to highlight topics that deserve more attention.

Revised editions of the Manifesto can be published when significant new content is available.

Translations

Once the text of the Manifesto has proven its stability, it will benefit from translation into other languages.

Alternative formats

The ideas in the Manifesto will gain more impact by being expressed in new formats - graphical, video, audio, narrative, meme.

Creating a MOOC (Massive Open Online Course) based on the Manifesto will help extend the reach even further.

H+Pedia

Community members are invited to assist with the creation of articles in H+Pedia in support of

the Manifesto.

These articles can include references to articles and books that amplify points made in the Manifesto.

H+Pedia is also a good location to host a list of "hard questions" that the present Manifesto has not answered definitely, and where additional attention from the transhumanist community is needed.

Events
Community members are invited to help to arrange a series of online and/or offline events in support of ideas from the Manifesto.

Real-time responses
Community members are invited to organise prompt responses to any news story or thinktank publication where ideas from the Manifesto deserve to be heard.

Bridges
Community members are invited to build bridges to other communities. Let's find and develop opportunities to express the ideas of this Manifesto for different audiences.

Policy recommendations
It's particularly important for the community to develop a number of political policy recommendations that are aligned with the Manifesto ideals.

These recommendations will ideally be suitable to slot into legislative programmes around the world.

Campaigns
Community members are invited to plan and carry out specific campaigns based on individual ideas in the Manifesto - campaigns to raise public awareness and alter the public discussion.

Campaigns based around specific policy recommendations will be especially welcome.

Rebooting organisations
It is time to refresh the online material of organisations closely connected to the Manifesto - such as Transpolitica and the Transhumanist Party UK. Material needs to be retired and/or brought up to date on websites, YouTube channels, social media accounts, and so on.

Superdemocracy
A key task for the network of transhumanist communities is to increasingly embody the principles of superdemocracy - to serve as an example for the rest of the world.

Suitable technological tools should be trialled and, when found useful, adopted.

Individual processes can be applied in an agile, incremental way, with findings published about

what worked well, and what didn't work well.

Metrics
Community members are invited to develop simple but useful metrics to track progress towards the era of sustainable superabundance.

Scenarios
Greater clarity is welcomed on scenarios for feasible timelines between the present day and the advent of sustainable superabundance. What are the possible pathways, roadblocks, risks, solutions, and workarounds?

Welcoming participants
Processes are needed to welcome and engage people from all walks of like who are keen to accelerate the advent of the era of sustainable superabundance.

Let's help new and old community members alike to quickly find roles whereby they can contribute to individual projects and also grow in accomplishment.

In all these ways, the few can become many. Transhumanism can move from being an occasional slogan to become the lifeblood of a vital global transformation. Humanity can more quickly transcend the shackles of frailty, confusion, waste, egotism, and divisiveness. Sustainable superabundance beckons. The future can be profoundly wonderful, for everyone. Let's make it so.