CASPER JASPERS

TECHNOLOGY AND GOVERNANCE

MAKING TECHNOLOGY WORK FOR SOCIETY

2021 Prometheus Amsterdam

Chapter 1 | Introduction

1.1 Why we need technological progress and innovation

Like it or not, technology has a profound impact on our lives and our society. You can't hide from it even if you want to. Luckily, the relationship between humans and technology is largely a happy story of progress. By many metrics, our lives were made better by technology. We live longer than ever before, and the quality of life has improved significantly for most of the world's population. We have more options to choose from when it comes to education, travel and leisure. But technology has also had negative effects such as the abuse of the earth's resources and the impact on our mental health. Still, on balance we seem to have done well – for now.

Technological invention began long ago when we discovered using stone to tenderise food, which helped digestion and saved energy. The energy surplus allowed our brains to grow bigger.[1] More than a million years ago, we invented fire and started using it to grill food, which unlocked even more energy in the form of calories.[2] Fire also gave us much-needed protection from predators.[3] It is one of the explanations for how we could sleep safely on the ground instead of on branches in trees. As a result, the intensity of our sleep, especially the REM phases, increased. REM sleep is body-immobilising, making us defenceless, but getting more of it increased the complexity and connectivity of our brains.[4] These modern ancestors, *Homo erectus*,[5] had brains that were an astounding 50% larger than those of our habilis ancestors.[6] And

compared to *Homo erectus*, we have a 40-50% larger brain.[7]

Using our upgraded brain power, we made life-changing inventions such as wheels, combustion engines, electricity and the printing press. We also invented complex social structures. But with progress came technological anxiety. There is a general sense that we have lost control over technology and are becoming victims of its unanticipated consequences. For centuries, literature has portrayed technology as threatening and possibly uncontrollable. The ancient Greek myth of Prometheus is, if anything, a warning about technology.

Prometheus is a Titan and the protagonist of our story. He has the gift of foresight and embodies ideas about technology, science and civilisation's progress. One day, Zeus, the king of the gods, and Prometheus set out to create a new species, with the Titan modelling human figures from mud. Through the spittle of Zeus and the gentle breath of the goddess Athena, the mud men were given life. However, Zeus refused to give humans fire, as he believed that if men possessed it, they would no longer stand in reverence of the gods.[8] Prometheus defied Zeus and travelled to Mount Olympus. There, he stole fire from the gods and bestowed it upon mankind. Prometheus also gave men the skills of metal work which marked the beginning of civilisation. For these acts, which challenged Zeus' commands, Prometheus was punished in the most terrible manner. He was helplessly chained to a rock and an eagle would eat the liver of the immortal Titan every day into eternity. Zeus also wanted to teach men a lesson. He instructed Hephaestus, the god of fire, to create a woman named Pandora, an Ancient Greek version of the biblical Eve. Each Olympian god gave her a gift, but Zeus gave her a jar and told Pandora never to open it.[9] The jar was filled with every kind of misery and evil you can think of. As with Eve, curiosity got the better of Pandora and she removed the lid of this jar, thereby unleashing creatures that carried with them disease, poverty, misery, sadness, death, and all other evils of the world.

This Greek myth is a metaphor for human enlightenment and

the dangers that come from overreaching our limits. Prometheus furthered societal progress by giving men fire and tools, which they used to create and manipulate their world. Although Prometheus had the best intentions, he and the humans ended up paying a terrible price.

If we have been worrying about technological progress for such a long time and everything has gone well so far, you might ask why worry now and is it different this time? Throughout this book I'll show that we are more interconnected than ever before and that technology is increasingly powerful, seeping into every nook and cranny of our daily lives. While our economic might is growing, our technological capabilities also present unprecedented threats.[10] Our tightly coupled, globalised system makes it likely that any crisis will spread rapidly, like a row of dominos toppling over.[11] These complex systems increase the potential for unintended consequences.[12] So, as we continue to innovate, we are enhancing both our capabilities and the risk of systemic failures.[13]

On balance, resisting technological progress is futile and doesn't make sense given the enormous benefits societies have gained from technology. Many of the world's current problems, ranging from removing vast amounts of carbon dioxide from the atmosphere, to cleaning up plastic in oceans and rivers, to discovering new antibiotics that can kill drug-resistant bacteria can probably only be solved through development of new tech.[14] Too much fear of technology might stifle much-needed innovation, which is referred to as the *technology trap.*[15] Fear that machines will take over human jobs, for example, has led several experts, among whom is Bill Gates, to propose a robot tax. This would slow down automation and provide funds for a safety net. But it might well be that we need more robots and must speed up their innovation because developed countries struggle with an ageing and shrinking workforce.

What we should do is mitigate the dangers of rapidly advancing technologies. But what exactly are people afraid of? Firstly, there are worries about *techno-determinism,* meaning that technology deter-

mines the development of social structures and cultural values, and that there's nothing we can do about it. One concern, for example, is that automation will lead to humans being replaced by machines, which will result in mass unemployment and further widening of the inequality gap. Secondly, people worry about the moral implications of technology taking over our lives. How will it affect our privacy, safety and security, freedom of thought and freedom of expression?[16] Video games and online platforms are designed to trigger dopamine responses to increase *user engagement,* and they employ similar tactics to the ones used in the gambling industry. We don't let kids gamble in casinos, so why do we let them freely play these games and participate on social media platforms without limitations? And what if we are constantly fed misinformation by malicious actors trying to break down the fundamentals of democratic societies? Should people be allowed to upgrade their brains by connecting them to the cloud? And should we let companies offer artificial intelligence resurrection services, allowing you to bring back to life a deceased loved one based on data from social media platforms, videos, messages and voice recordings? Thirdly, as China is getting more powerful and the country's technologies are finding their way into Western infrastructure, a global contest between Chinese autocracy and Western liberal values will take place. Finally, there's the overarching question of whether we should make the world a better place for future generations.

If we want to anticipate and tackle these problems, we first need a basic understanding of emerging technologies and their foreseeable consequences. We must start taking time to think ahead and ensure that technological progress is progress for society as a whole.

1.2 WHY WE MUST THINK AHEAD

What our society will look like in the future is a matter of scientific and ethical significance. More importantly, the future will

depend on choices made in the present. Over time, humanity has undergone numerous significant societal changes, for example due to the development of agriculture and industrialisation. Research indicates that the rapid pace of technological and social change suggests that some sort of major societal transformation is likely to occur relatively soon.[17] But much of modern society – in particular government and politics – is suffering from *temporal exhaustion*: we spend so much time and energy dealing with the present that there is little left for anticipating the future.[18]

Because we are able to create and believe fiction, we can dream up different kinds of scenarios.[19] We perform *mental time travel*, which allows us to go back in time and foresee and shape times ahead. The human brain uses the same area – that of episodic memory – for both activities; to remember is simply to imagine the past.[20] Thinking about the future means finding a probabilistic answer to *what would happen if*. This counterfactual thinking makes you investigate how causal relationships change given some kind of intervention.[21] We can imagine situations such as what we will do tomorrow, next week, where we will go on holiday, what career path to pursue, and with whom we have relationships. We can also imagine alternative versions of those events, which we can evaluate in terms of their likelihood and desirability. This ability to imagine different types of scenarios appears to be unique to humans.[22] What is perhaps even more remarkable is our capacity to describe these imaginary scenes to other people so they understand them, allowing others to co-create the future.

There is good reason to think ahead because the longevity of civilisation depends on it. Most of us are familiar with the examples of the rapid collapse of cultures. But more often the process is drawn-out and mild, and leaves societies floundering for many decades.[23] This phenomenon is called *creeping normalcy*, where every year things get just a little bit worse, but never bad enough for anyone to really notice. There is a metaphor that is often used: legend has it that if you put a frog in a pan of boiling water it will instantly jump

out. But if you put that frog in a pan with lukewarm water and gradually turn the heat up, it will stay in the water until it boils to death.[24] The hypothesis is that certain problems accrue and materialise in a gradual manner that we are likely to overlook, until crystallisation is well underway. Luckily for the frog, it is just a legend, and the frog's biological system has a kind of foresight that makes it jump out before the water boils, thus ensuring survival.[25]

The collapse of civilisation can be defined as a rapid and enduring loss of population, identity and socio-economic complexity. Public services crumble and disorder ensues as governments lose control of their monopoly on violence.[26] There's an ongoing debate about why societies fail, with causes ranging from external factors (environmental damage, barbarian invasions), to population growth leading to strife, to economic and political inequality leading to revolution,[27] to implosion resulting from complexity and bureaucracy,[28] to the capture and use of energy (fossil fuels and nuclear energy).[29] Complete destruction of a society is rare, however, and often it rebuilds itself on the remains. Perhaps societies innovate through failure, and we should therefore view collapse as a feature instead of a bug in our software.[30]

Whether we can completely avert crisis, failure or collapse of societies is debatable. But we can at least try to soften the blow if and when it happens, or make systems anti-fragile. We must plan for failure. There are four main reasons communities fail to prepare for change:[31]

1 Failure to anticipate a problem before the problem actually arises.
2 When the problem does arise, the group may fail to perceive it.
3 Then, if it perceives the problem, it may fail even to try to solve it.
4 Finally, it may try to solve it but may not succeed.

Let's take the example of a machine that is capable of understanding or learning any intellectual task that a human can perform. After its inception, we also give our robot an upgraded human-looking body equipped with super strength and speed.

First, we could fail to anticipate the problems we could face if such a robot were to exist. That could be because we have no prior experience with the situation and simply do not know what to expect. But, even worse, we might draw false analogies. To come to grips with a new situation humans have a tendency to compare it to an old, familiar situation. A prime example is the discussion around disruption of the labour market as a result of automation: people tend to point out that in the past technological progress has always led to the creation of more jobs overall, meaning the number of new jobs is greater than the jobs lost. But this is not a law of physics we can assume will apply again.

Secondly, we could fail to perceive the problem. It could be that we do not have the technology to measure or monitor what the robot is doing, making the problem literally imperceptible. It could also be that there is too much distance between the decision-makers and the problem. The people in charge simply don't notice what is going on because the problem is not brought to their attention. Or we might fail to perceive the problem because progress takes place in the form of a slow trend with upward and downward fluctuations, like climate change. Related to that is another cause for failing to perceive problems called *landscape amnesia*. We tend to forget what the world looked like a decade ago and therefore don't notice change. Just try to imagine that we used to live in a world without smartphones.

Thirdly, even if we perceive the problem, we could fail at trying to solve it. Sometimes this is the result of a dislike of the people who do perceive the problem; we think they are just complaining. Other times, we dismiss warnings because of past warnings that proved to be false alarms. It could also be that rational behaviour prevents us from trying to solve an issue. It is not uncommon for

people to advance their own interest by behaviour harmful to others, many times with the excuse: if I don't do it, someone else will. United States foreign policymakers love this justification, for instance when it comes to selling arms to foreign nations.[32]

Fourthly and finally, we could simply fail to solve the problem. Maybe the robot we created is so intelligent that it is beyond our capacities to shut it down. Or the solution exists but is too expensive to implement. It could also be that our efforts are too little and too late.

That is not to say that there aren't several governments and organisations working on mapping and solving potential problems. Wales, for example, has a special commissioner for future generations to ensure that policy decisions are made with long-term effects of at least thirty years in mind.[33] Finland and Sweden have installed parliamentary advisory groups to promote longer-term planning, and Hungary has appointed an ombudsman for future generations. Then there are organisations such as the Future of Life Institute, a non-profit research institute, working to mitigate existential risks facing humanity, especially dangers from advanced artificial intelligence (AI). Beyond the existential risks threatening humanity such as nuclear war, climate change, anthropogenic risks, asteroid impacts and pandemics, the risk of a catastrophe due to AI occurring within the next 100 years is the most pressing. There is an estimated 10% chance a rogue AI will become an existential threat.[34]

Today, technology is by and large a market-driven industry. Private companies, and their shareholders, determine the direction of technological development. We are letting them, and their technologies, drive our choices. Where, when and how governments should intervene remain open questions. Your view on government intervention likely also depends on whether you think firms owe society anything at all. Do private companies have moral

obligations towards the community in which they exist? I'll argue that governments have a responsibility to guide technological progress in cooperation with the private sector. In many places around the world, policymakers need to regain trust by starting to tackle the big issues that will dominate tomorrow. Technological disruption is not a leading item on the political agenda, but it should be. If the direction of technological progress can be steered, which remains to be seen, we will be forced to make decisions along the way. Therefore, we need a deeper understanding of the policy and governance tools available to lawmakers and regulators, and we must better understand the emerging technologies. Taking decisions on the direction of technology also requires an investigation into moral concepts such as justice and fairness. In this book, I'll argue that in democratic societies governments and politics are primarily responsible for providing a forum where this discussion can take place. While we might not be able to agree on everything, we'll see that it is worth having a public debate if we want to design a governance system for emerging technology that has any chance of being successful. We must set the outcomes that we want technology to help achieve.

Technology appears to be agnostic, merely providing a function with a certain utility. There are gun owners who tend to use the following argument: *guns don't kill people, people kill people*. Guns are used for good purposes (protecting people) and bad ones (hurting people). But there is an aspect other than utility, namely the design, which leads to the conclusion that technology is not agnostic. The choices made during the design of technology impact how we use it. This is increasingly true as technology gets more advanced and more capable of influencing our behaviour. So there are two fronts upon which to make ethical choices: use and design. When moral issues are solved by the technological activities of designers instead of democratic activities of politicians, it will be technology that is in control, and not society.[35] We will discuss the importance of finding a democratic way to moralise

technology, and this book offers a possible solution.

While it's notoriously hard to make predictions about the future, we can extrapolate trends from current technological developments. But as a general rule, we tend to overestimate the effect of technology in the short run and underestimate its effect in the long run.[36] And much harder to predict, if not impossible, are breakthrough discoveries. Imagining the future will leave us with a wide range of possible paths, which are constantly changing as a result of technological developments and scientific research. That is why we should shift from old-fashioned regulatory infrastructure to new, more adaptive and more flexible governance models that are co-created by the public and private sectors. If we want technological progress to benefit society, we must regain our first-mover advantage and proactively work together with technology companies.

1.3 What you'll learn from reading this book

I firmly believe that executives, professionals and practitioners in various areas will find this book useful. In 2017, I was working on a book on corporate governance.[37] I interviewed more than forty prominent executives, non-executives and policymakers, several of whom worked at reputable listed companies such as Shell and ING Bank. Technology was a topic that came up every time, but there appeared to be a gap between what people thought they knew about technology and what they actually knew. My suspicions were confirmed in a report that came out in 2019, which concluded that many directors, especially the ones in more supervisory roles, lack digital literacy.[38] The same applies to policymakers: they also have catching up to do.[39] Perhaps this is unsurprising as, like many of us, these directors and policymakers are incredibly busy and have trouble finding the time to catch up with the high volume and relenting pace of technological progress.