

Too much of a good thing can be a very bad thing: Technology's Future

Why we must act now to ensure that exponential technological progress remains a benefit for humanity and the planet



Gerd Leonhard, Futurist & Humanist. Zurich, Switzerland May 17 2021.

Also published on [Cognitive World](#).

Technology has undoubtedly been good to us. AI systems are now able to detect some forms of skin cancer at a higher rate of accuracy than human doctors. Robots perform ultrasound exams and surgery, sometimes with little to no human intervention. Autopilots fly and land planes in the most adverse conditions and may soon be able to steer personal air taxis. Sensors gather live data from machinery and their “digital twins” to predict upcoming failures or warn of crucial repairs. [3D printers](#) are capable of spitting out spare parts, improving maintenance options in remote locations.

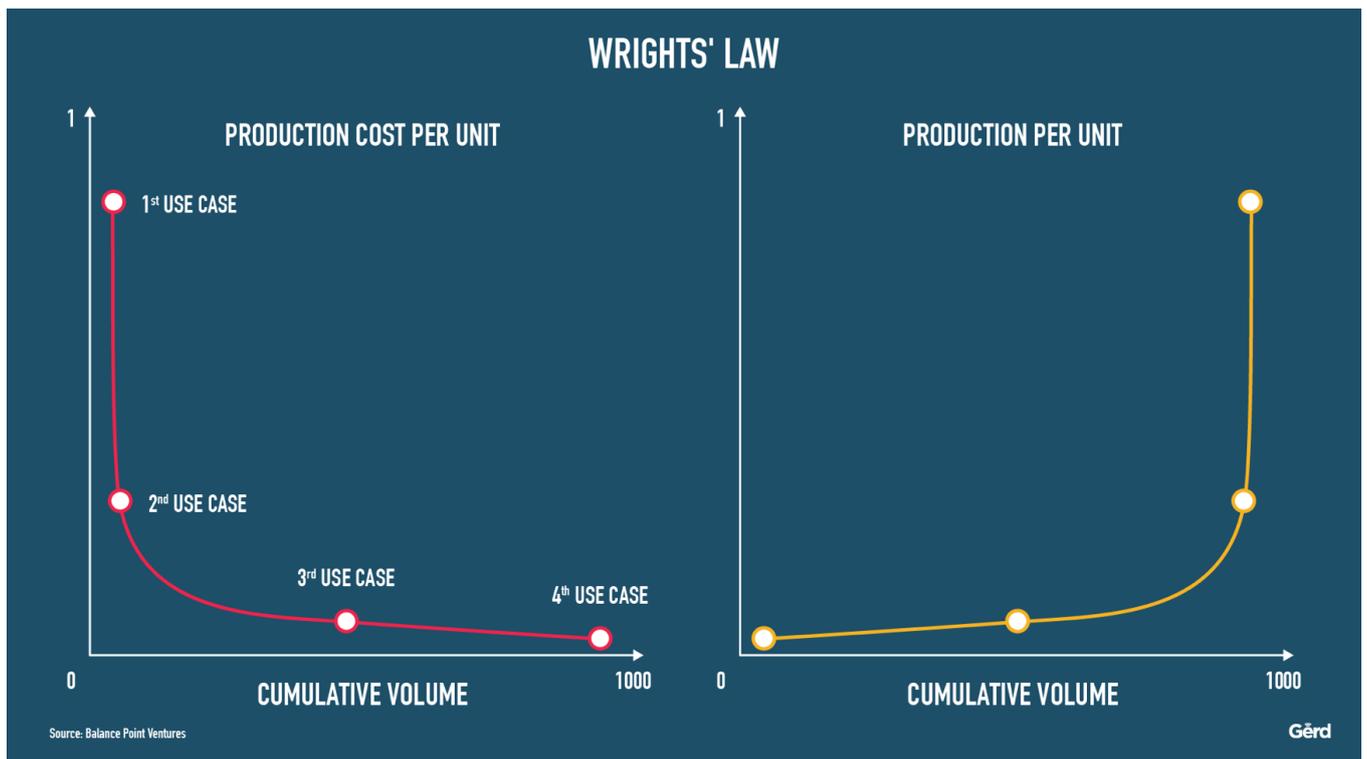
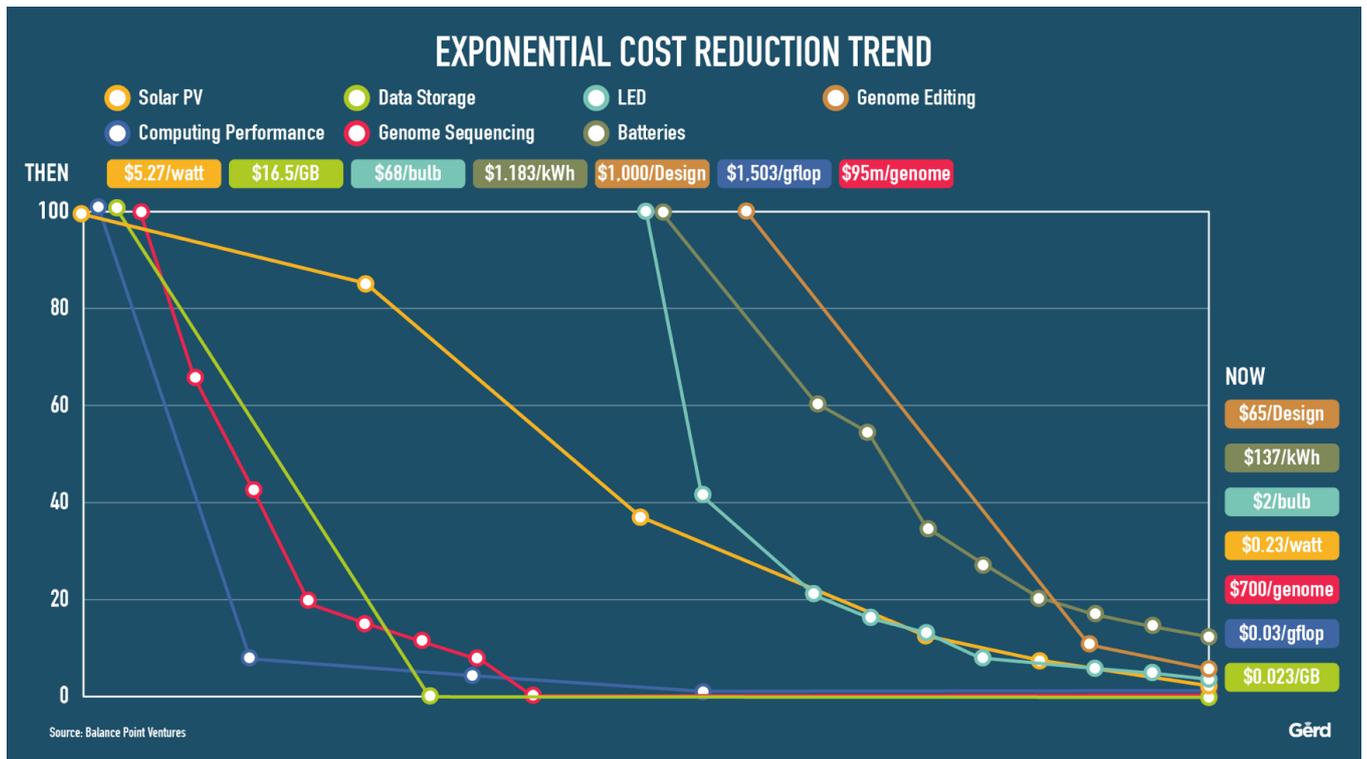
Yet as technology and [AI systems](#) in particular become ever more powerful, the downside of exponentially accelerating technological progress also comes into sharp focus. [Riffing off Marshall McLuhan](#) (57 years ago!), every time we extend our capabilities, we also amputate others. Today, social media promises to “extend us” by allowing all of us to share our thoughts online, yet it also ‘amputates’ us by making us easy targets for tracking and manipulation. We gain the reach but we lose the privacy.

The impact of exponential change is incredibly hard to imagine

We should take note that while today’s technological progress often seems very impressive and useful, mostly harmless and easily manageable, it would be very foolish indeed to assume that this will remain the same in the near future, as we’re literally leaping up the exponential curve from 4 to 8 to 16 and on. 30 steps linearly may get me across the street, 30 exponential steps equals 26 trips around the globe. The scope of magnitude is utterly different and whatever *may have* worked before the pivot point may prove disastrous later on.

Right here, right now we are at the take-off point for many of the foundational sciences and technologies of the 21st century, from artificial intelligence and deep learning to human genome editing to geo-engineering, longevity and human enhancement – our entire framework is

about to change – [not just the picture!](#)



30 LINEAR steps will take you
30 meters from the starting point

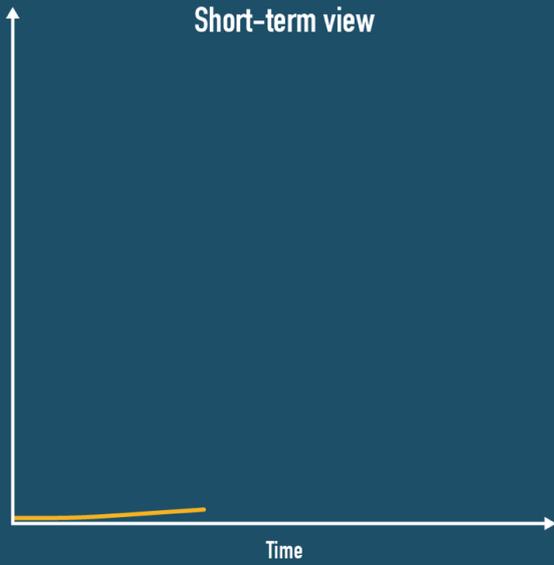
30 EXPONENTIAL steps will take you
26 trips around the world!



Gërd

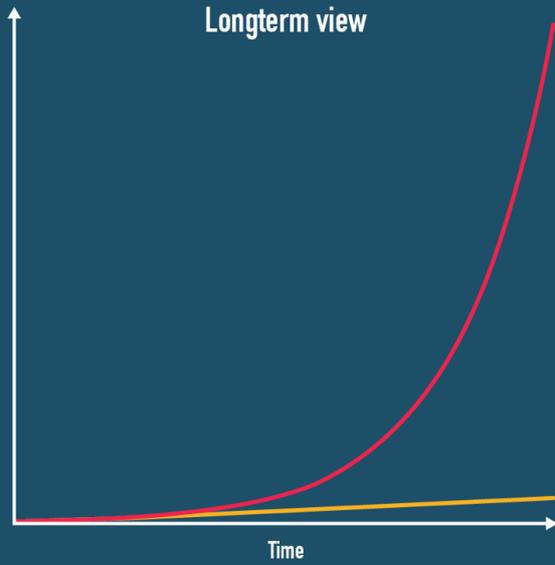
LINEAR VS EXPONENTIAL

Short-term view

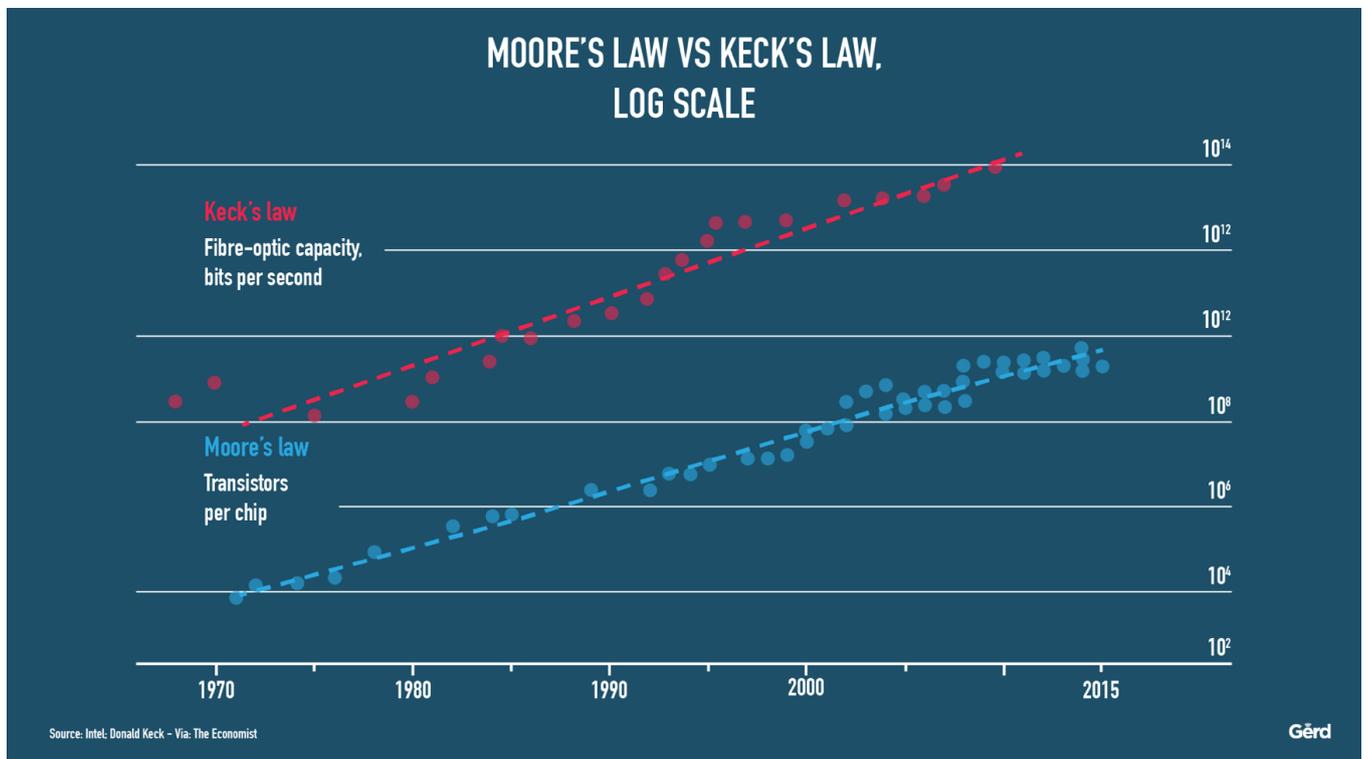


Source: Michael Simmons

Longterm view



Gërd



During the next decade, every single one of the domains mentioned above (what I often call the [Game-Changers](#)) has more change in store than we witnessed in the entire previous century, in total. [Science fiction is increasingly becoming science fact](#) – and since humans are biological ([for now](#)) and don't progress exponentially, we often have a very hard time understanding where all of this is going. What may be really amazing, magical and "good" for us today, may soon become "too much of a good thing", and go from being a tool to being the purpose, from being a god-send to being a doomsday machine.

Whether it's enhancing humans, designing chimeras or robots going to war: What will happen [during the next 10 years](#) will sweep away most of our notions about gradual progress and upend our views about how technological innovation usually and maybe even inadvertently furthers the common good.

Defining 'The Good Future'

In my view, what I recently started calling "[the good future](#)" (i.e. a utopia or better yet a [protopia](#) as opposed to a dystopia) is now at stake, and we cannot afford to shy away from the hard work such as collectively defining what a "good future" actually means for us, collectively and as a species,

and how we can best shape it, together. As arduous as a consensus on the question of “good” may seem, the future must be created collectively or there will be none at all – at least not one that involves us humans (see the nuclear treaty / NPT analogies below).

Right now, as we see some semblance of a light at the end of the pandemic tunnel, is a good time to start an honest and open discussion about our shared goals and our *telos* (i.e. the purpose), and about the policies we will need in order to create that '[Good Future](#)'. One thing is certain: The free-market doctrine has recently failed us a least twice (during the 2007 financial crisis, and during the Covid crisis), and the role of government as a market-shaper is bound to be much more prominent in the coming decade. The Good Future is not likely to come about with good old-fashioned capitalism, and neither will it be ushered in by technology. Rather, it is a question of ethics, values – and the policies we craft based on them.

Time to Rehumanize!

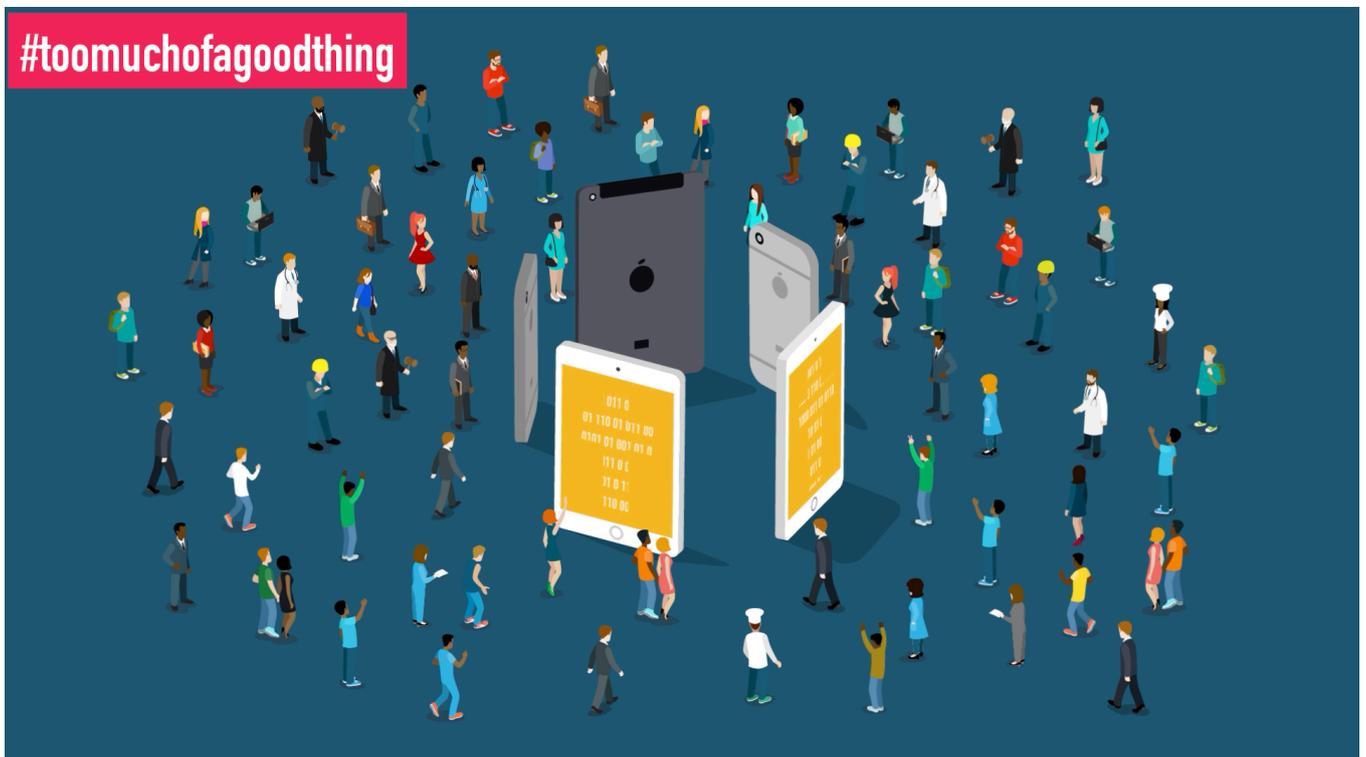
Now is the time to think about and act on “civilizing” or “re[humanizing](#)” and yes, regulating (but not strangulating) the global giants of technology – because in this age of advanced exponential change, today's amazing breakthroughs may well turn out to be major problems not too far down the road. What looks just fine at “1” may no longer be ok when we reach “16” on the exponential curve (i.e. after 5 doublings). We urgently need to look much further ahead as we are moving into the future at warp-drive speed. Foresight is simply mission-critical now and no longer “nice to have.” ‘The Future’ is becoming a daily agenda item for every leader and public official.

So let's stop asking “what will the future bring” – it can soon, quite literally, bring almost anything we can imagine. Instead, [we need to ask: “what kind of future do we WANT?”](#)

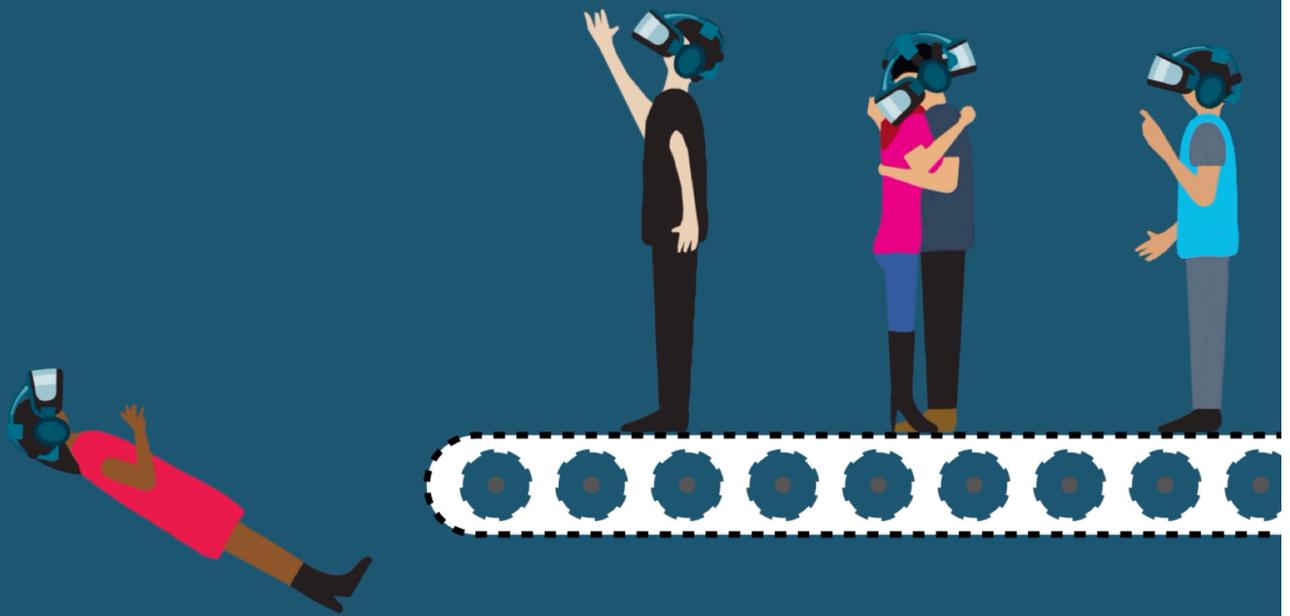
The amazing feat of developing not just one but several efficient vaccines against SARS-CoV-2 is a case in point as far as our accelerated

exponential status goes. What used to take a decade and more, now took a mere 12 months, thanks in large part to the global capabilities of cloud computing and AI. Within a decade, that development cycle is likely to shrink to six weeks or less. Just [imagine](#).

Yet this accomplishment should not distract from the basic truth that technology always has dual usages; it can be both good and not so good at the same time: It's "morally neutral until we use it" (W. Gibson). "[Too much of a good thing](#)" describes this conundrum very well: something quite useful can quickly become something deeply harmful and corrosive to society – as recently evidenced with social media, which suffers from excessive monetization obsession and utter lack of accountability.



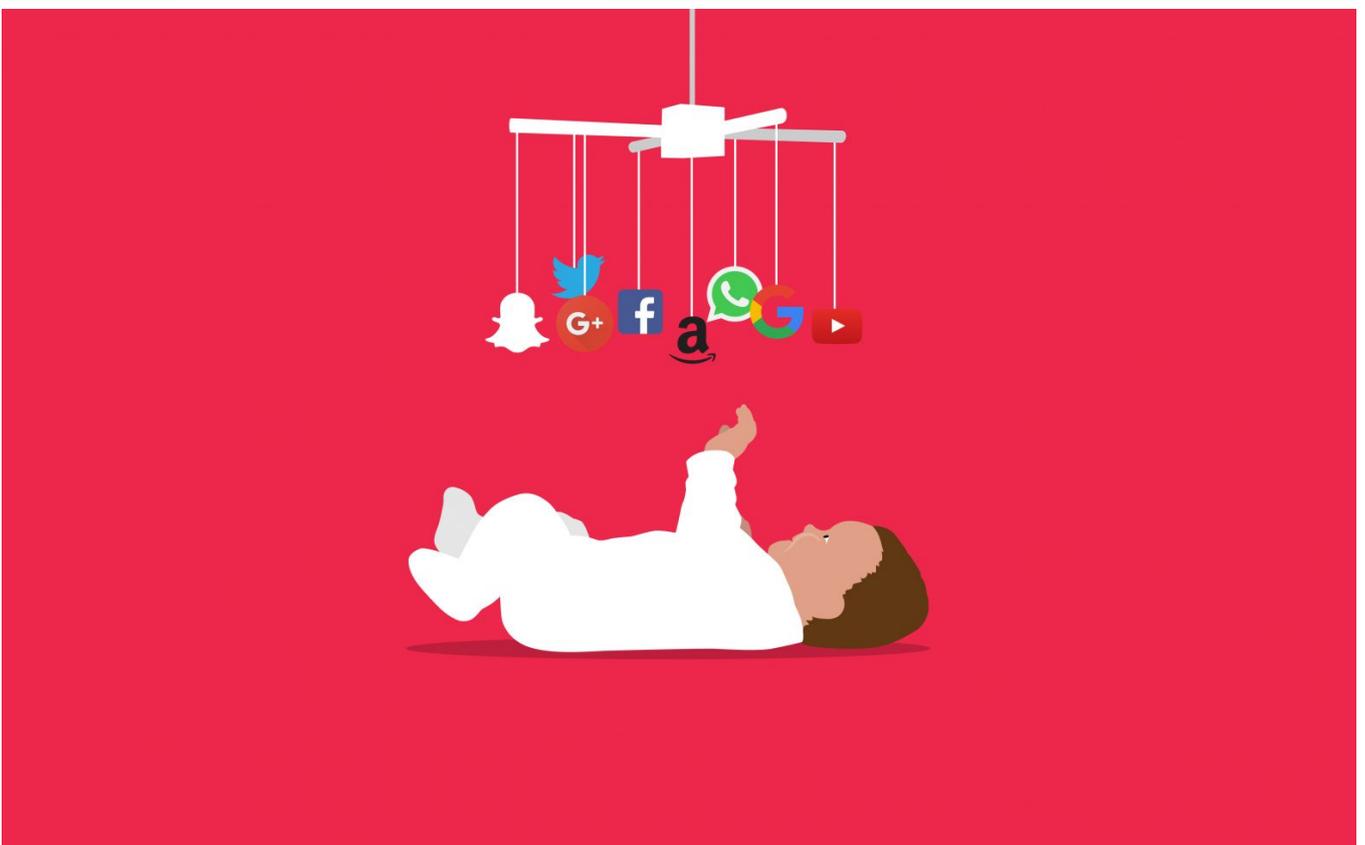
#toomuchofagoodthing



#toomuchofagoodthing



#toomuchofagoodthing





This is why I think the European Commission is on the right track with its recent [AI and technology regulation efforts](#). Yes, it's complicated, it's painful and may sometimes be overbearing – but somehow, we must get started to look beyond economic factors and growth at all costs, and to move towards what I've been calling the quadruple bottom line of "[people, planet, purpose and prosperity](#)" (extending the work done by Elkington and others on the triple bottom line).

“Too much of a good thing” – sound familiar?

It is not a coincidence that the overuse and dependency on devices and online services are often considered an addiction. Many of us use drugs in some form, from alcohol and tobacco to even controlled substances. Societies have learned how to deal with those addictions by way of laws, regulations, and a slew of social contracts designed to limit their reach and negative impact. Technology is not dissimilar in that excessive use or over-exposure also ruins lives. Many studies have documented that [social network power users frequently rank at the top of suicide-prone lists](#), and that FOMO (far of missing out) can lead to serious anxiety disorders. But just wait until we have 5G and pervasive virtual reality applications and “smart glasses”.

There is a palpable danger that this rapid and leaping technological progress may not be governed wisely before its constantly expanding capabilities overwhelm us, and presents us with a de-facto future most of us would not wish for their children. We therefore need a



a mix of “new norms” and social contracts, demonstrated responsibility and accountability in business and industry, and real safeguards such as regulation, laws and treaties. Our response must be broad and flexible, fair and democratic yet decisive when it comes down to acting on the side of precaution.

Imagine, for example, what will soon become possible once most of us have our [healthcare data stored in the cloud](#) and regularly interact with digital assistants and chatbots for things like mental health support. Pooling massive datasets derived from the human genome, biome and from behavioural observations (for instance when using a mobile device or a smart watch or when wearing a VR set) can give us tools to detect degenerative diseases and possibly even [pre-diagnose cancer](#). Yet the

consequences of privacy breaches in this domain could also be a thousand times worse than today's pitfalls e.g. when people's location data is illicitly captured. Again, as the song goes: *"You ain't seen nothing yet"*.

Regulation and the challenge of a global consensus: the nuclear analogy

A useful example exists in the way we dealt with the threat of nuclear proliferation. Unfortunately, it took two nuclear bombs and almost 25 years of hard work before the global [NPTs went into force in 1970](#). Today, given how fast and wide technology is advancing, we won't have the luxury of such a long runway anymore – and it's a lot harder to build a nuclear bomb than to write code. Or – taking a dimmer view – could it be that in a not-too-distant future, we must first experience catastrophic incidents caused by over-relying on unfit and unsafe "intelligent machines," or face even more ecosystem collapse due to hasty experiments with geo-engineering?

Locally, nationally, regionally and inevitably on a global scale, we need to agree on the bottom lines of what we can agree on, and set realistic 'lowest common denominator' thresholds beyond which certain technological applications cannot be pursued just "because we can" or because enormous economic benefits may be achievable. One such example of attainable consensus is a [ban of autonomous weapons systems](#) that kill without human supervision, or to ban the use of face recognition and more recently affect recognition for dubious commercial purposes. It's another issue the EU Commission has started to address.

The only question that really matters: What kind of world do we want to leave to our children?

We must realize that within this coming decade it will no longer be about what future we *can* build, but [what future we want to build](#). It will no longer be just about great engineering or clever economics, it will be about values, ethics and purpose. Scientific and technological limitations will

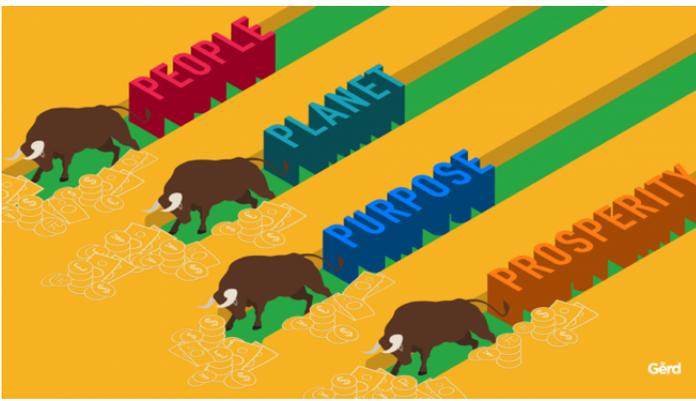
continue to fall away at a rapid pace, while the ethical and societal implications will become the real focal point. As E.O. Wilson famously said: *"The real problem of humanity is the following: we have paleolithic emotions; medieval institutions; and god-like technology"*

Further to this, Winston Churchill proclaimed: *"You can always count on Americans to do the right thing – after they've tried everything else,"* and I often adapt this quote to read "humans" instead of Americans. But there will inevitably be things that should never be tried because they may have irreversible and existential consequences for humans and our planet, and that [list is getting longer every day](#). Starting with the pursuit of artificial general intelligence, geo-engineering and human genome editing – all of these technologies could potentially be [heaven, or they could be hell](#). And who will be ["mission control"](#) for humanity?

In particular, Silicon Valley's well-documented approach to experiment with almost anything and ask for permission later ("Move fast and break things," in [Zuckerberg's pre-IPO parlance](#)) will take us down a very dangerous road. This notion may have been tolerable or even exciting at the beginning of the exponential curve (where it looks almost the same as a linear curve), but going forward this paradigm is decidedly ill-suited to help bring about "the good future." Asking for forgiveness simply doesn't work when technology can swiftly bring on existential threats that may be far beyond the challenge we face(d) with nuclear weapons.

Take the idea to somehow make humans "superhuman" by chasing extreme longevity or by [connecting our brains to the internet](#) (yes, Elon Musk seems to be clamouring to become leading purveyor of these ideas). I fear we risk ending up with half-baked solutions introduced by for-profit corporations while the public has to pick up the pieces and fix the externalities – all for the sake of progress.

We are at a fork in the road: The future is our choice, by action or by inaction

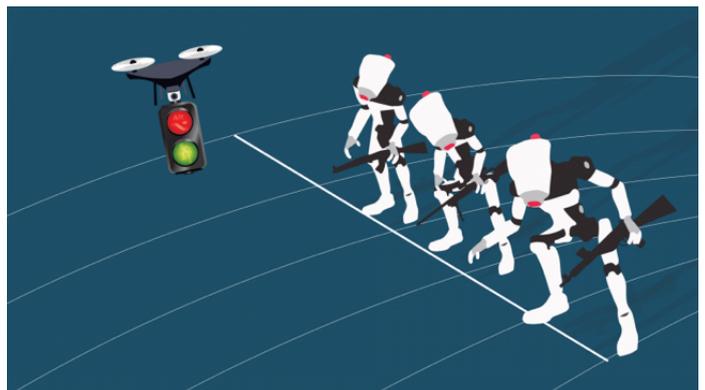


[The operating paradigm for the future](#)

Time is of the essence. [We have ten short years](#) to debate, agree on and implement new, global (or at least territorial) frameworks and rules as far as governing exponential technologies are concerned.

If we fail to agree on at least the lowest common denominators (such as clear-cut bans on lethal autonomous weapons systems, or bans on using genetic engineering of humans for military purposes) I worry that the [externalities of unrestrained and unsupervised technological growth](#) in AI, human genome editing and geo-engineering will prove to be just as harmful as those that were created by the almost unrestricted growth of the fossil fuel economy. Instead of climate change, it will be 'human change' that will present us with very wicked challenges if we don't apply some more ['precautionary thinking'](#) here.

Going forward, no one should be allowed to outsource the negative externalities of their business, whether it is in energy, advertising, search, social media, cloud computing, the IoT, AI or with Virtual Reality.

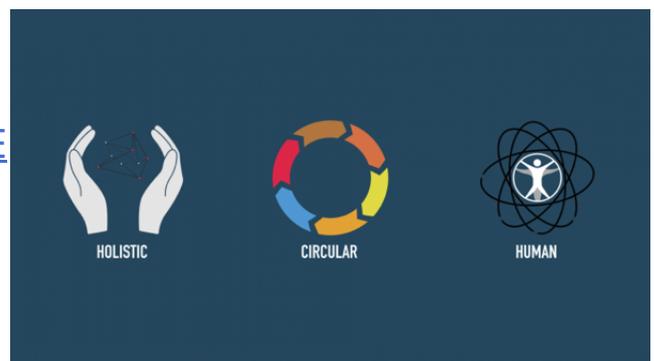


When a company such as Alphabet's Google wants to turn an entire Toronto neighbourhood into its smart city lab ([Side Walks Lab](#)), it is the initiating business — not the community and civil society — that has to be held accountable for the ensuing societal changes that pervasive connectivity and tracking bring about – both good and bad. Companies

like Google must realize that in this exponential era (i.e. “beyond [the pivot point](#)”) they are indeed responsible for what they invent and how they roll it out. Every technological issue is now an ethical issue, as well, and success in technology is no longer just about “great science, amazing engineering and scalable business models” (and maybe it never was).

Instead, the responsibilities for addressing potentially serious (and increasingly [existential](#)) downsides need to be baked into each business model, from the very beginning, and consideration must be given to all possible side-effects during the entire business design process and all subsequent roll-out phases. We need to discard the default ‘[dinner first, then morals](#)’ – doctrine of Silicon Valley or its Chinese equivalents. Holistic approaches (i.e. models that benefit society, industry, business and humanity at large) must replace “[move fast and break things](#)” attitudes, and [stakeholder logic](#) must replace shareholder primacy (*watch Paul Polman, former Unilever CEO, speak about this, [here](#))

Within the next five years, I think we are going to see new stock markets emerge (such as the still-nascent [LTSE](#) in San Francisco) where companies that operate on the premise of “People, Planet, Purpose and Prosperity” will be listed. [EUROPE should seize this opportunity](#) and take the lead!



3 of my 6 Future Principles

Call me an optimist or a utopian, but I am convinced that we can still define and shape a “good future” for all of mankind as well as for our planet.

Here are some suggestions to guide this debate:

- 1. The ethical issues of exponential technologies must be deemed as important as the economic gains they may generate.**

Technology is not the answer to every problem, and we would be ill-

advised to look to technology to “solve everything.” If there is one thing I have learned in my 20 years as a Futurist, it is this: Technology will not solve social, cultural or human problems such as (in)equality or injustice. On the contrary, it often make matters worse because technology drives efficiency regardless of its moral consequences. [Facebook](#) is the best example for this effect: a really powerful concept, very cleverly engineered, working incredibly well (just as designed) – and yet deeply corrosive to our society (watch my talks on this topic, [here](#)). Apple’s CEO Tim Cook nails it when he says: *“Technology can do great things, but it doesn't want to do great things. It doesn't want anything.”* The “want” is our job, and we shouldn’t shirk that responsibility just because it is tedious or may slow down monetization.

2. **Humans must remain in-the-loop (HITL):** For the foreseeable future and until we are much further along with understanding *ourselves* (and not just our brains) and capable of controlling the technologies that could potentially “be like us” (aka [AGI or ASI](#)), I strongly believe that we need to keep humans in the loop, even if it’s less efficient, slower or more costly. Otherwise, we may be on the path towards dehumanisation a lot faster than we think. [Pro-action and pre-caution](#) must be carefully considered and constantly balanced.
3. **We need to ban short-term thinking** and get used to taking both an exponential and a long-term view – while respecting the fact that [humans are organic](#) (i.e. the opposite of exponential). As mentioned above, we are just now crossing that magic line where the pace of change rockets skywards: 1-2-3 may look a lot like 1-2-4, but leaping from 4 to 8 to 16 is very different from going from 3 to 4 to 5. We must learn and practice how to live linearly but imagine exponentially; one leg in the present and one leg in the future.
4. **We need new “non-proliferation treaties.”** Artificial general intelligence (AGI), geo-engineering or human genome editing are both magical opportunities as well as significant existential threats (just like nuclear power, some would argue), but we will probably not

have the [luxury of dabbling in AGI](#) and then recover from failure. An “[intelligence explosion](#)” in AI may well be an irreversible event, and so might germ-line human genome modifications. We will therefore need forward-looking, binding and collective standards, starting with an international memorandum on AGI.

5. I have been speaking about the need for a “**Digital Ethics Council**” a lot during the past five years, starting with my 2016 book “Technology vs Humanity.” Many countries, states and cities have started establishing similar bodies lately, such as the EU’s European Artificial Intelligence Board. We need more of these, regionally, nationally, internationally, and globally!



So here we are in 2021, standing at this fateful and also very hopeful fork in the road. It is upon all of us to choose the right path and start building the “good future.”

It is not too late – have hope, have courage and pursue wisdom!

What kind of Future do we want? What's your PREFERRED FUTURE? ...

#Business as usual is dead or dying. An odd kind of 'Golden Age' is ...

No country would be safe from fully autonomous weapons

The Future of #Europe: Technology, Society & Policy. Riveting Keyno...