the futures foundation
engaging all Australians in creating a better future...

Volume 24, No. 3, June 2024

FUTURE NEWS

TO CONNECT, TO INFORM AND TO INSPIRE

IN THIS EDITION

futures foundation Australia • Ross House, 247 Flinders Lane, Melbourne 3000 • Phone: 03 9029 5787

DON'T LET YOUR BIASES GUIDE YOUR DECISION-MAKING
WHAT BEHAVIOURAL ECONOMICS CAN TEACH US ABOUT PLANNING FOR THE FUTURE

by Martin Kruse


Pierre Wack

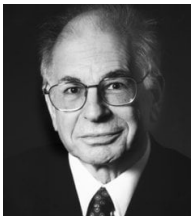
Daniel Kahneman

An icon within future studies, Pierre Wack, once said that the best and most difficult task of a futurist is to make people think of the world in a new way. He used the term 'reperception' to describe how people awaken to the possibility of the future being different from the past or from how you expect it to be. Wack believed that the greatest accomplishment, but also the most difficult for people working with foresight and scenarios, is to facilitate this transformation. But why is it so difficult for us to think of alternatives to the way things are?

Why are we predisposed to think in certain ways? This question has been on the mind of recently deceased psychologist Daniel Kahneman for years. His writings contain clues not just to why we think the way we do in general, but more specifically how we think and make decisions about the future.

What makes Kahneman's work interesting from a futurist's perspective is his focus on human errors in decision-making that arise from heuristics and biases. He thereby challenges the assumption held by economists for decades, that of the rational human acting based on objective self-interest. In so doing, he has provided strong arguments for looking at economics from a psychological angle, an area known as behavioural economics. Let's take a closer look at some of the biases drawn from the work of Kahneman that specifically relate to the challenge of imagining alternative futures – as well as some of the red flags to look for when these biases are applied to decision-making.

'CHANGE IS BAD FOR BUSINESS' – The status quo bias

Pierre Wack saw the act of reperceiving as crucial to opening the minds of executives and making them understand either the risks of disruption to their business or the possibilities that exist for them in alternative futures. Fundamentally, Wack was talking about how to overcome the so-called 'status quo bias'.

The fundamental problem with this kind of bias is that it does not permit change to be positive. Change will, for a number of reasons, be interpreted as a threat, especially for incumbent businesses that have lowered costs on core processes substantially to increase competitiveness. Businesses in this situation have invested a lot of money in organising their offerings to be efficient, and they are the kings of low cost in what is typically a red ocean market. For many such businesses, no change is preferred to constant change,

simply because the status quo (where the incumbent business is on top) is preferred to the available alternatives.

When doing strategic foresight however, you sometimes find yourself in a situation where an executive, from a logical point of view, agrees to all the driving forces causing a specific scenario, yet chooses to ignore the scenario presented to them, close their eyes, and hope for the best. This can be especially perplexing to a futurist because it is no longer a question of having the right arguments or the right data. Rather, it becomes a question of feelings. Some people simply choose to ignore the facts because they hope things will turn out differently in the future than what the most likely scenario suggests.

This problem is confounded by the fact that especially big corporations often need a sizeable revenue stream to replace their cash cow, and new business very often fails to deliver enough revenue right away to be of interest. Forecasts of future revenue are rarely very reliable, primarily because new products or new technologies create new markets, the size of which are naturally hard to predict. No wonder hoping for the best, even in the face of radical change, sometimes seems to be the best approach.

At its worst, the status quo bias can lead to what is known as 'persistence of discredited beliefs'. In a now-famous study undertaken in the 1950s and described in the book "When Prophecy Fails", psychologists studied a UFO cult that was convinced that the world would end on December 21, 1954. When in fact it did not, many of the members of the cult still clung to their beliefs, settling on alternative explanations for why the world had not ended yet. One might not be so surprised that this happens in a cult, but the fact of the matter is that something similar also often happens in large corporations, behind the walls in the boardrooms, and in governments as well.

This is a reason why it is one of the most important tasks of futurists to look for where opinions diverge between people within organisations and experts outside the organisation. When external experts have radically differing opinions about the state of the world than those inside the organisation, it is often a case of status quo bias, and that should raise a red flag.

'WHEN IN DOUBT, GO WITH WHAT YOU KNOW' – The confirmation bias

In many cases, the tendency to search for, interpret, and recall information that supports one's own beliefs actively stands in the way of choosing a better path forward. The 'confirmation bias' has been known for years, and rules to mitigate it are integrated into the scientific method and teachings of good scientific practice. However, it is very much a part of everyday media and politics, and it affects decision-making in many areas of society and business.

As Kahneman points out in his book "Thinking Fast and Slow", confirmation bias tends to be strongest with emotionally charged issues and entrenched beliefs. The current media reality, increasingly defined by online echo chambers, tends to feed our confirmation biases by creating spaces where we can easily have our existing beliefs confirmed by likeminded individuals. The largest study ever done on the spread of falsehoods on Twitter was published in Science in 2018, and the results confirmed that the confirmation bias thrives in our fast-paced social media reality. The study, which was conducted by MIT researchers, tracked how news circulates and found that hoaxes, rumours, and falsehoods consistently dominated the conversation on Twitter. In fact, stories containing false information tended to reach people six times quicker than stories containing factually correct information.

For executives, the confirmation bias manifests itself most often when they choose to only listen to people who share their own opinions. This impulse can be so strong that it ends up being a defining trait of an organisational culture. This can lead to information contradicting the established truth not being circulated or taken seriously. In other words, a self-imposed censorship can take hold, which means that disruptive business models or technologies that are around the corner may be ignored at the detriment of the organisation. Other times, decision-makers will have put in so much effort into committing to a specific strategy that there is a sunk cost connected to switching lanes, and so, an executive may do their best to continuously seek out arguments that confirm that the chosen strategy is the right one. This can blind one to the possibility that other directions may be more beneficial in the long term.

'THIS IDEA IS SO GOOD IT COULDN'T POSSIBLY FAIL' – The optimism bias

One of the most commonly observed biases is called the 'optimism bias'. In our 2017 report Evaluating the Hype, we explored how this kind of bias often affects the assessment of what the impact of new technologies will be, and how fast they will reach maturity. Almost without exception, experts and media commentators alike tend to believe that things move faster than they actually do.

For this reason, when assessing the prospects for a technology's future breakthrough, it may be necessary to add two, five, ten, or even twenty years to that assessment (depending of course on the technology) if you think you may be suffering from optimism bias yourself. There are several reasons for this delay that may not immediately come to mind. For example, new technologies are often hemmed by standardisation issues, regulations impeding uptake, or high prices creating a tough transition between innovators and early adopters.

Optimism bias often makes an appearance whenever people try to envision how things may look in the future, both in regard to their

personal outlook and when assessing more general developments. Kahneman argues that there are several reasons for this, chief among which is that our judgment is affected by the goals or end-states that we aim for or desire. That is a fancy way of saying wishful thinking.

Optimism bias is often found going hand in hand with confirmation bias. The sense that one's own business is superior to the competitor is what happened to Martell, the producer of Barbie dolls, who found that despite having been able to fend off all the prior attacks on their core product, Bratz still managed to take a big market share to the big surprise of Martell's management.

Optimism bias is often present when new technology sees the light of day. Some readers may remember the hydrogen bubble in the early 2000s, during which President George W. Bush said fuel cell cars would be competitive with internal combustion engines by 2010 and would eliminate over 11 million barrels of oil demand per day in the US by 2040. Today, there are fewer than 20,000 heavily subsidised hydrogen fuel cell vehicles on the roads globally, nowhere close to the target.

Research has shown that this kind of bias is closely tied to mental well-being, with individuals suffering from depression showing less signs of optimism bias. The same study also made clear that even experts aren't free from optimism bias: 'Divorce lawyers underestimate the negative consequences of divorce, financial analysts expect improbably high profits, and medical doctors overestimate the effectiveness of their treatment', the researchers write.

AWARENESS IS THE FIRST STEP

The work of establishing what kind of biases are at play when we envision the future is of vital importance for how we plan for it. There are many other biases than the ones discussed here, and the work with identifying the ones that are specific to the field of futures studies and foresight is ongoing. The fundamental problem is that if we do not know what guides our decisions, we are not well equipped to make the right choices. This is especially true because more than ever, the problems we face in the future, be it climate change, loss of biodiversity, or pandemics, are shaped by the decisions we make today.

For some of these problems, we don't have the luxury of making the wrong decision. When it comes to climate change, time is running out. A big part of the explanation of why we have even gotten to this point is that we lack the imagination to see the future clearly because we have little or no past references to draw on. In order to get this point across, let us recall 9/11, a wild card event that permanently changed the global geopolitical landscape. It's not that no one could have seen it coming. Al-Qaeda's plans were known in advance by US intelligence since they had been disclosed

in an interrogation with captured members of the terror network, but still the information was never acted on. Why? One explanation, the one that was put forward in the 9/11 commission report, has to do with something known as 'availability heuristics'. This term explains how bits of information can be retrieved, generated, and combined from memory.

In the case of the terrorists' plans, there weren't many similar historical instances of giant skyscrapers being hit by airplanes to draw from. The fact that this information did not exist in the minds of the individuals in possession of the relevant intelligence was taken as evidence that it would not happen. As the report concluded, 'it was fundamentally 'a failure of imagination'.

Availability heuristics, as well as our active biases, are of huge importance whenever we try to assess the likelihood of wild cards or black swan events. The Fukushima nuclear reactor disaster and the depth of the housing market crash in the US in 2008 leading to the financial crisis, are other examples of how wrong things can go if we are not mindful of this. Needless to say, not being able to foresee disasters or radical change has, in retrospect, often proven to be a case of biases rather than not being able to prepare for alternative futures.

For governments and businesses to make better decisions, we need to understand what drives this decision-making in the first place. Equipped with this knowledge, one of the main goals of futurists, that of facilitating reperception as Wack pointed out, should become easier.

This article is part of the Copenhagen Institute for Futures Studies "using the Future" compendium and is published with permission.

COPENHAGEN INSTITUTE FOR FUTURES STUDIES is an independent, global, and non-profit futures think tank. The Institute was founded in 1969 on initiative of former Finance Minister and OECD Secretary-General, Professor Thorkil Kristensen. The foundation was established in collaboration with visionary Danish public and private organisations to better qualify decision-making through futures studies for the betterment of our society. The Institute works to equip and inspire individuals and organisations, decision-makers and the public, to act on the future, today.

Read more at WWW.CIFS.DK

Martin Kruse
Senior Executive Advisor and Futurist, Chair of the Copenhagen Node of The Millennium Project, Employee Representative at the Board of the Copenhagen Institute for Futures Studies,

Book Review

by Charles Brass – Chair, Futures Foundation

Undelivered
The Never-Heard Speeches That Would Have Rewritten History
by Jeff Nussbaum

At first glance this is an unusual book to be reviewed in these pages – our focus is normally on foresight practice or principles. On closer examination, however, this is a novel, and highly important, exploration of one of the principles of strategic foresight – the creation of alternative futures – also known as What If? thinking.

Jeff Nussbaum has a long career as a speechwriter, and an equally long interest in speeches that were written (and seriously contemplated being given) but never actually given. As he explains in the introduction his interest in this area began when he was on the staff of Al Gore at the time of the 1996 US Presidential election. As is traditional (apparently) all the candidates (and their staff) prepare both acceptance and concession speeches. As those who remember this particular election (think hanging chad….), it presented some significant challenges for both candidates and their speechwriters. As it happens, Nussbaum has been unable to find copies of any of the speeches prepared for and by Gore that evening, but the experience set him on a journey to explore speeches written, and contemplated, but never delivered – and he found a good variety of them.

The book is divided into six sections, each comprising undelivered speeches of different kinds:

1. WORDS THAT ARE TOO HOT

There are times when an audience can be shaken too hard or challenged too much. What happens when someone other than the speaker decides that a speech goes too far?

Here Nussbaum has two examples, a speech proposed to be given alongside Martin Luther King during the march on Washington in 1963 and a speech planned to be given by a descendant of the American Indigenous population at the 350th anniversary of the landing at Plymouth Rock. For different reasons, neither was delivered, but records of both remain.

2. A CHANGE OF MIND OR A CHANGE OF HEART

Sometimes a speaker sees a speech in their hand and realises that they don't want to own those words and their consequences.

Of the five examples here, perhaps the most profound is speech written by and for

Richard Nixon in August 1974 in which he announces his intention NOT to resign.

3. CRISIS OR CRISIS AVERTED

When we look back at historical events, we do with a sense of certainty. but what we fail to realise is that there could easily have been a very different path.

Two examples are given – a speech in which Edward VII refuses to abdicate and one in which the mayor of New York declares bankruptcy in 1975 (a crisis which was averted at literally the last moment).

4. THE FOG OF WAR OR THE PATH TO PEACE

Important speeches often derive from times and places where outcomes are in doubt, and the stakes are high.

Three examples – Dwight Eisenhower apologising for the failure of the D-Day landings – Emperor Hirohito apologising for the Second World War and President Kennedy announcing the launch of missiles at Cuba are all chilling, and a reminder of how

5. THE PEOPLE CHOOSE

Elections provide no shortage of undelivered speeches, for nearly every delivered victory or concession speech, there exists its undelivered opposite.

the most powerful example here is the speech Hillary Clinton would have given had she won the 2016 US presidential election

6. EVENTS INTERVENE

It is not uncommon for a speech to go undelivered because of some unforeseen circumstance, including the death of the speaker.

Here Nussbaum looks at a speech US National Security Advisor Condoleezza Rice was scheduled to give on the afternoon of the 11th of September 2001, which was never delivered because of the attack on the Twin Towers in New York; and the speech that was supposed to be given when the controversial film “Moonlighting” won the 2017 Academy award for best picture, but was not delivered when the presenter announced the wrong film as winner and chaos ensued. He also explores the last words of Pope Pius XI, FDR, Albert Einstein and John F. Kennedy

Foresight practitioners are acutely aware of both the difficulty and the importance of having people deal with the inevitability of uncertainty, while at the same time navigating their way into their future. Books like this are both interesting to read and useful to gently remind us all not only that things don’t always turn out as planned, but that the world goes on even after things don’t go as planned.

Since we also know that many of our members regularly give presentations, the book also provides a number of useful tips for developing and delivering memorable speeches.

FUTURISTS IN ACTION

THINK LIKE A FUTURIST FOR IMPACT AND FUTURE GROWTH

by James Clampett

The 2020s are years of upheaval and uncertainty. The pandemic has upended our lives and left many of us feeling anxious about the future. It's natural to wonder what the next year will bring and whether we are headed towards a future of abundance or scarcity.

One way to make sense of what's happening is to think like a futurist. Futurists use foresight to examine trends and identify potential scenarios for the future. This helps them to see both the opportunities and the challenges that lie ahead. While futurists can't predict the future with certainty, they can help us to understand the different pathways that might lie ahead. This is valuable knowledge for anyone who wants to make the most of what happens next.

WHAT'S GOING TO HAPPEN NEXT?

It's human nature to want to know what's going to happen next. We want to be able to plan for upcoming events and avoid any potential pitfalls. However, the future is inherently unpredictable. So why bother trying to think about it? According to futurists, foresight is a way of thinking that helps us to make better decisions in the present. By considering a range of potential future scenarios, we can start to prepare for them. This helps us to minimize the impact of negative events, and capitalize on positive ones. In other words, foresight gives us a form of insurance against the unknown.

Whilst thinking about futures and all the different potential pathways that could unfold is interesting. It's even more fascinating to consider how our actions in the present can shape futures. This is what futures intelligence is all about: understanding the trends and dynamics at play in the world today so that we can make informed decisions about futures. This process of imagining different futures and mapping out potential pathways is essential for making progress on big challenges like impact and future growth. By using futures intelligence, we can make sure that futures aren't just something that happens to us, but something that we actively create.

FIVE TO TEN YEARS AHEAD.

It's interesting to think about how far into the future we can reasonably predict. Obviously, the further out we try to look, the more variables there are that could impact what happens. But foresight is more accurate when considering events that are closer in time. At insight & foresight we like a five to ten years timeframe for people to think in because the futures intelligence available is stronger and people can generally agree on possible, probable, and preferred futures within that timespan. And finally, this gives us just enough time to take action on our findings and make a real impact. So while it's important to monitor longer-term global trends and changes, the five to ten-year horizon is our focal point when it comes to foresight and using futures intelligence.

FUTURES INTELLIGENCE.

Foresight is all about understanding and preparing for change. Futures intelligence is the data we use to support our thinking. Futures intelligence can come from anywhere. They could be a technology, a product or service, a social media topic, a research project, a news story, or just a piece of

data that shows something is different or changing.

And right now, there are plenty of signs that change is on the horizon. From new technologies to social trends, we are seeing the world around us shifting in some pretty big ways.

Of course, change is always happening. But sometimes things happen slowly and gradually. Other times, there are inflection points – moments when change happens more rapidly or dramatically. We seem to be at one of those inflection points right now, with signals of the future all around us.

So what exactly are these trends, signals, and emerging changes? They can take many forms, but usually, they are things or developments that are on the margins – things that look strange or weird or out-of-the-ordinary at first glance. But upon closer inspection, we can see that they are forerunners of much larger trends.

Some examples of trends, signals, and emerging changes include:

New technologies like artificial intelligence, blockchain, Web3 and 6G wifi

Social trends like the rise of the gig economy or the popularity of veganism

Political movements like populism or nationalism

Economic shifts like the rise of China or the slowing of global growth

These are just a few examples, but you get the idea. By paying attention to futures intelligence, we can start to see patterns and trends that give us a glimpse into what lies ahead.

Of ourse, not all futures intelligence is positive. Some signal potential problems or challenges that we will need to address. For example, the rise of populism or nationalism could signal a future of increased conflict and division. The slowing of global growth could signal a future of economic insecurity and hardship. So while it's important to be aware of both positive and negative signals, we also need to be thoughtful about how we respond to them.

This is where foresight comes in. By understanding the trends and dynamics at play in the world today, we can make informed decisions about the future. We can use foresight to see both the opportunities and the challenges that lie ahead. And we can use our insights from the past to help us navigate the present and shape the future.

FORESIGHT TOOLS.

Futurists often use a tool called scenario planning to explore different futures. This involves imagining different potential outcomes and mapping out possible pathways. It helps us to understand the different factors that could impact the future and identify areas of opportunity or concern.

Scenario planning is a useful tool, but it's not the only way to think about the future.

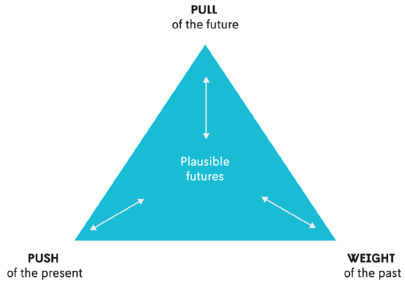

Futurists also use futures triangle, emerging issues analysis, three tomorrows framework, cause layered analysis, visioning and backcasting as well as tools like the Futures Platform and other methods to explore different possible futures. The important thing is to start thinking about the future more intentionally.

The world is changing rapidly, and it can be hard to keep up. But by thinking like a futurist, we can gain a better understanding of the trends and changes that are happening in the world today. We can use this knowledge to make more informed decisions about the future. And we can take action to make a positive impact on the world around us.

insight & foresight supports for-purpose CEOs and Boards to use foresight to plan and make decisions for impact and future growth with more confidence and great reliability. James Clampett is principal at insight and foresight based on Gadigal land in Sydney.

This post is reproduced with permission.

Signals in the Noise

FIVE AI PREDICTIONS FOR THE YEAR 2030

by Rob Toews

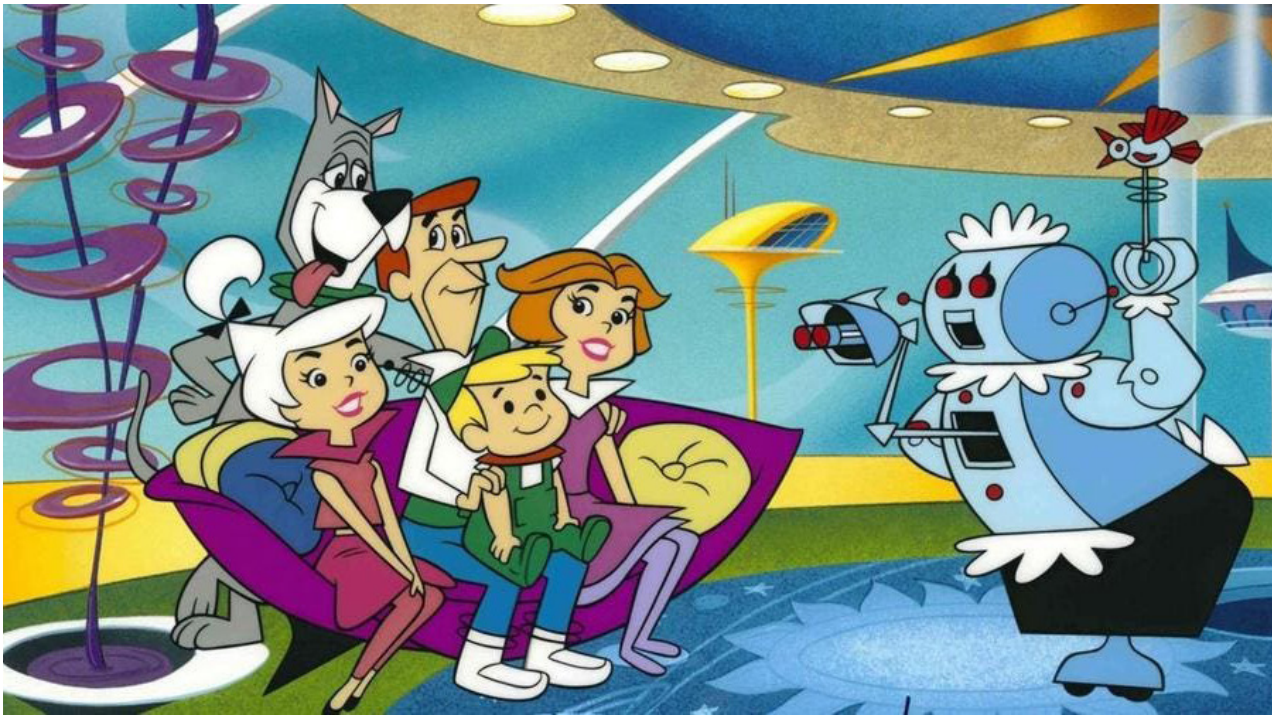

How will the field of artificial intelligence, and the world, have changed by 2030? THE JETSONS

Making predictions for the year ahead is difficult enough. What if we tried to predict the future not one year out, but half a decade out?

The further into the future we attempt to peer, the hazier things look and the more speculative our thinking must become. If one thing is certain in technology, it is that no one can actually predict the future—and that we are all going to be surprised by how things play out.

But putting a stake in the ground about how things will unfold is nonetheless an informative and fun thought experiment.

Below are five bold predictions about what the world of artificial intelligence will look like in the year 2030. Whether you agree or disagree with these predictions, we hope they get you thinking.

1. NVIDIA'S MARKET CAPITALIZATION WILL BE MEANINGFULLY LOWER THAN IT IS TODAY. INTEL'S WILL BE MEANINGFULLY HIGHER THAN IT IS TODAY.

Nvidia is the hottest company in the world right now. It has been the biggest beneficiary of today's generative AI boom, with its market cap skyrocketing from under $300 billion in late 2022 to over $2 trillion today.

But Nvidia's position as the single dominant provider of chips for AI cannot and will not last.

What Nvidia has built is difficult, but not impossible, to replicate. A resurgent AMD is emerging as a credible alternative provider of advanced GPUs, with its cutting-edge new MI300 chip about to become widely available. The big tech companies—Amazon, Microsoft, Alphabet, Meta—are all investing heavily to develop their own AI chips in order to lessen their dependence on Nvidia. OpenAI's Sam Altman is seeking up to trillions of dollars of capital to build a new chip company in order to diversify the world's supply of AI hardware.

As demand for AI chips continues to grow in the years ahead, relentless market forces will ensure that more competitors will enter, supply will increase, prices will drop, margins will tighten and Nvidia's market share will fall.

In addition, as the market matures in the years ahead, the primary type of AI computing workload will shift from training to inference: that is, from building AI models to deploying those models in real-world settings. Nvidia's highly specialized chips are unrivaled when it comes to training models. But inference can be done with cheaper and more commoditized chips, which may undermine Nvidia's advantage in the market and create an opening for competitors.

None of this is to say that Nvidia will not still be an important part of the AI ecosystem in 2030. But the current stratospheric runup in its stock price—which has made it the third most valuable company in the world as of this writing, larger than Amazon or Alphabet—will in retrospect look like irrational exuberance.

Meanwhile: what is the one thing that sets Intel apart from virtually every other chip company in the world?

It manufactures its own chips.

Nvidia, AMD, Qualcomm, Broadcom, Alphabet, Microsoft, Amazon, Tesla, Cerebras, SambaNova, Groq: none of these companies build their own chips. Instead, they design chips and then they rely on other companies—most importantly, the Taiwan Semiconductor Manufacturing Company (TSMC)—to produce those chips for them.

Intel alone owns and operates its own chip fabrication facilities.

The ability to manufacture chips has become a vital geopolitical asset. Case in point: China's utter dependence on foreign semiconductor suppliers has enabled the U.S. to handicap China's domestic AI industry by banning the import of AI chips to China.

U.S. policymakers are acutely aware of the vulnerabilities posed by the extreme concentration of chip manufacturing in Taiwan today, especially as China adopts an increasingly hawkish stance toward the island. Promoting advanced semiconductor manufacturing on U.S. soil has become a top policy priority for the U.S. government. U.S. lawmakers are taking decisive action to advance this goal, including committing a whopping $280 billion to the effort under the 2022 CHIPS Act.

It is no secret that Intel has fallen behind TSMC over the past decade in its ability to manufacture cutting-edge chips. Yet it still remains one of the few companies in the world capable of fabricating advanced semiconductors. Under CEO Pat Gelsinger, who took the helm in 2021, Intel has reprioritized chip fabrication and undertaken an ambitious strategy to reclaim its former position as the world's preeminent chip manufacturer. There are recent indications that the company is making progress toward that goal.

And perhaps most importantly: there is simply no other option to serve as America's homegrown chip manufacturing leader.

U.S. Commerce Secretary Gina Raimundo, who leads the Biden administration's efforts on AI and chips, acknowledged this directly in a recent speech: "Intel is the country's champion chip company."

Put simply, America needs Intel. And that bodes well for Intel's commercial prospects.

Nvidia's market cap today is $2.2 trillion. Intel's, at $186 billion, is more than an order of magnitude smaller. We predict that this gap will have shrunk significantly by 2030.

2. WE WILL INTERACT WITH A WIDE RANGE OF AIS IN OUR DAILY LIVES AS NATURALLY AS WE INTERACT WITH OTHER HUMANS TODAY.

Even though the entire world is buzzing about artificial intelligence right now, the number of touchpoints that the average person actually has with cutting-edge AI systems today is limited: the occasional query to ChatGPT or Google Bard/Gemini, perhaps.

By the year 2030, this will have changed in dramatic fashion.

We will use AIs as our personal assistants, our tutors, our career counselors, our therapists, our accountants, our lawyers.

They will be ubiquitous in our work lives: conducting analyses, writing code, building products, selling products, supporting customers, coordinating across teams and organizations, making strategic decisions.

And yes—by 2030, it will be commonplace for humans to have AIs as significant others.

As with any new technology, there will be an adoption curve. Some portions of the population will more readily adjust to interacting with their new AI peers; others will resist for longer. The proliferation of AIs throughout our society will unfold like the famous Ernest Hemingway line about how people go bankrupt: "Gradually, then suddenly."

But make no mistake: this transition will be inevitable. It will be inevitable because AIs will be able to do so much of what humans do today, except cheaper, faster, and more reliably.

3. OVER ONE HUNDRED THOUSAND HUMANOID ROBOTS WILL BE DEPLOYED IN THE REAL WORLD.

Today's AI boom has unfolded almost entirely in the digital realm.

Generative models that can converse knowledgeably on any topic, or produce high-quality videos on demand, or write complex code represent important advances in artificial intelligence. But these advances all occur in the world of software, the world of bits.

There is a whole other domain that is waiting to be transformed by today's cutting-edge AI: the physical world, the world of atoms.

The field of robotics has been around for decades, of course. There are millions of robots in operation around the world today that automate different types of physical activity.

But today's robots have narrowly defined capabilities and limited intelligence. They are typically purpose-built for a particular task—say, moving boxes around a warehouse, or completing a specific step in a manufacturing process, or vacuuming a floor. They possess nowhere near the fluid adaptability and generalized understanding of large language models like ChatGPT.

This is going to change in the years ahead. Generative AI is going to conquer the world of atoms—and it will make everything that has happened to date in AI seem modest by comparison.

Dating back to the dawn of digital computing, a recurring theme in technology has been to make hardware platforms as general as possible and to preserve as much flexibility as possible for the software layer.

This principle was championed by Alan Turing himself, the intellectual godfather of computers and artificial intelligence, who immortalized it in his concept of a "Turing machine": a machine capable of executing any possible algorithm.

The early evolution of the digital computer validated Turing's foundational insight. In the 1940s, different physical computers were built for different tasks: one to calculate the trajectories of missiles, say, and another to decipher enemy messages. But by the 1950s, general-purpose, fully programmable computers had emerged as the dominant computing architecture. Their versatility and adaptability across use cases proved a decisive advantage: they could be continuously updated and used for any new application simply by writing new software.

In more recent history, consider how many different physical devices were collapsed into a single product, the iPhone, thanks to the genius of Steve Jobs and others: phone, camera, video recorder, tape recorder, MP3 player, GPS navigator, e-book reader, gaming device, flashlight, compass.

(An analogous pattern can even be traced out in the recent trajectory of AI models, though in this example everything is software. Narrow, function-specific models—one model for language translation, another for sentiment analysis, and so on—have over the past few years given way to general-purpose "foundation models" capable of carrying out the full range of downstream tasks.)

We will see this same shift play out in robotics over the coming years: away from specialized machines with narrowly defined use cases and toward a more general-purpose, flexible, adaptable, universal hardware platform.

What will this general-purpose hardware platform look like? What form factor will it need to have in order to flexibly act in a wide range of different physical settings?

The answer is clear: it will need to look like a human.

Our entire civilization has been designed and built by humans, for humans. Our physical infrastructure, our tools, our products, the size of our buildings, the size of our rooms, the size of our doors: all are optimized for human bodies. If we want to develop a generalist robot capable of operating in factories, and in warehouses, and in hospitals, and in stores, and in schools, and in hotels, and in our homes—that robot will need to be shaped like us. No other form factor would work nearly as well.

This is why the opportunity for humanoid robots is so vast. Bringing cutting-edge AI into the real world is the next great frontier for artificial intelligence.

Large language models will automate vast swaths of cognitive work in the years ahead. In parallel, humanoid robots will automate vast swaths of physical work.

And these robots are no longer a distant science fiction dream. Though most people don't yet realize it, humanoids are on the verge of being deployed in the real world.

Signals in the Noise

FIVE AI PREDICTIONS FOR THE YEAR 2030

Humanoid robots are on the verge of being deployed in real-world settings. 1X TECHNOLOGIES

Tesla is investing heavily to develop a humanoid robot, named Optimus. The company aims to begin shipping the robots to customers in 2025.

Tesla CEO Elon Musk has stated in no uncertain terms how important he expects this technology to be for the company and the world: "I am surprised that people do not realize the magnitude of the Optimus robot program. The importance of Optimus will become apparent in the coming years. Those who are insightful or looking, listening carefully, will understand that Optimus will ultimately be worth more than Tesla's car business, worth more than [full self-driving]."

A handful of younger startups are likewise making rapid progress here.

Just last week, Bay Area-based Figure announced a $675 million funding round from investors including Nvidia, Microsoft, OpenAI and Jeff Bezos. A couple months ago, the company released an impressive video of its humanoid robot making a cup of coffee.

Another leading humanoid startup, 1X Technologies, announced a $100 million financing in January. 1X already offers one version of its humanoid robot (with wheels) for sale, and plans to release its next generation (with two legs) soon.

Over the next few years, these companies will ramp from small-scale customer pilots to mass production. By the decade's end, expect to see hundreds of thousands (if not millions) of humanoid robots deployed in real-world settings.

4. "AGENTS" AND "AGI" WILL BE OUTDATED TERMS THAT ARE NO LONGER WIDELY USED.

Two of the hottest topics in AI today are agents and artificial general intelligence (AGI).

Agents are AI systems that can complete loosely defined tasks: say, planning and booking your upcoming trip. AGI refers to an artificial intelligence system that meets or exceeds human capabilities on every dimension.

When people envision the state of AI in 2030, agents and/or AGI are often front and center.

Yet we predict that these two terms won't even be widely used by 2030. Why? Because they will cease to be relevant as independent concepts.

Let's start with "agents".

By 2030, agentic behavior will have become a fundamental, essential element of any advanced AI system.

What we today refer to using the umbrella term "agents" is actually just a core set of capabilities that any truly intelligent entity possesses: the ability to think long-term, plan, and take action in pursuit of open-ended goals. Becoming "agentic" is the natural and inevitable end state for today's artificial intelligence. Cutting-edge AI systems in 2030 will not just generate output when prompted; they will get stuff done.

In other words, "agents" will no longer be one intriguing subfield within AI research, as they are today. AI will be agents, and agents will be AI. There will thus be no use for the term "agent" as a standalone concept.

What about the term "AGI"?

Artificial intelligence is fundamentally unlike human intelligence, a basic truth that people often fail to grasp.

AI will become mind-bogglingly more powerful in the years ahead. But we will stop conceptualizing its trajectory as heading toward some "generalized" end state, especially one whose contours are defined by human capabilities.

AI great Yann LeCun summed it up well: "There is no such thing as AGI....Even humans are specialized."

Using human intelligence as the ultimate anchor and yardstick for the development of artificial intelligence fails to recognize the full range of powerful, profound, unexpected, societally beneficial, utterly non-human abilities that machine intelligence might be capable of.

By 2030, AI will be unfathomably more powerful than humans in ways that will transform our world. It will also continue to lag human capabilities in other ways. If an artificial intelligence can, say, understand and explain every detail of human biology down to the atomic level, who cares if it is "general" in the sense of matching human capabilities across the board?

The concept of artificial general intelligence is not particularly coherent. As AI races forward in the years ahead, the term will become increasingly unhelpful and irrelevant.

5. AI-DRIVEN JOB LOSS WILL BE ONE OF THE MOST WIDELY DISCUSSED POLITICAL AND SOCIAL ISSUES.

Concerns about technology-driven job loss are a familiar theme in modern society, dating back to the Industrial Revolution and the Luddites. The AI era is no exception.

But to this point, discussions about the impact of AI on job markets have been mostly theoretical and long-term-oriented, confined to academic research and think tank whitepapers.

This is going to change much more abruptly than most people appreciate. Before the decade is out, AI-driven job loss will be a concrete and pressing reality in everyday citizens' lives.

We are already beginning to see canaries in the coalmine here. Last month, fintech giant Klarna announced that its new customer service AI system is handling the work of 700 full-time human agents. Plagiarism detection company Turnitin recently projected that it would reduce its workforce by 20% over the next 18 months thanks to advances in AI.

In the years ahead, organizations will find that they can boost profitability and productivity by using AI to complete more and more work that previously required humans. This will happen across industries and pay grades: from customer service agents to accountants, from data scientists to cashiers, from lawyers to security guards, from court reporters to pathologists, from taxi drivers to management consultants, from journalists to musicians.

This is not a distant possibility. The technology is in many cases already good enough today.

If we are honest with ourselves, a major reason why we are all so excited about AI in the first place—a major reason why AI offers such transformative economic opportunity—is that it will be able to do things more cheaply, more quickly and more accurately than humans can do them today. Once AI can deliver on this promise, there will be less need and less economic justification to employ as many humans as today in most fields. Almost by definition, in order for AI to have an impact on society and the economy, it will take people's jobs. Of course, new jobs will also be created—but not as quickly and not as many, at least at first.

This job loss will bring with it tremendous near-term pain and dislocation. Political movements and leaders will arise in fierce opposition to this trend. Other segments of society will just as vocally champion the benefits of technology and AI. Civil unrest and protests will be inevitable; they will no doubt turn violent at times.

Citizens will clamor for their elected officials to take action, in one direction or another. Creative policy proposals like universal basic income will go from fringe theories to adopted legislation.

There will be no easy solutions or clear-cut ethical choices. Political affiliations and social identities will increasingly be determined by one's opinions on how society should navigate the spread of AI throughout the economy.

If you think the political moment in 2024 is tumultuous: buckle up.

Rob Toews is a venture capitalist at Radical Ventures.

Future News is published by the Futures Foundation six times a year for its members.