THE GREAT TRANSITION

Uniting Humanity for a Future Where Everyone Thrives

JOE CARVIN

I dedicate this book to my family—my wife Roz, who has been a rock of support throughout this journey, my children Michael, Keira and Rhianna, and my parents, the first "One Worlders" who encouraged a robust humanist approach to life—and to three exemplary thought leaders whose wisdom has profoundly shaped this work.

Sir Martin Rees, Britain's former Astronomer Royal, has done more than anyone to help humanity address existential risks through his prescient warnings in Our Final Century and the founding of the Centre for the Study of Existential Risk, bringing a scientist's perspective to explain that we are living through the most important epoch in the universe's history.

Peter Diamandis introduced the transformative concept of an "abundance mindset" in Abundance: The Future Is Better Than You Think, envisioning how exponential technological advances can solve global challenges and creating organizations like Singularity University and XPRIZE to advance this vision.

Dr. Fernando Reimers, Harvard's Ford Foundation Professor of International Education, has championed educational innovation through over 50 academic works and his leadership of UNESCO's Commission on the Futures of Education, helping bring global competence education into the mainstream with both his academic works and his annual Think Tanks on Global Competence.

What unites these remarkable thinkers is that they went beyond writing to create organizations that actively advance the wisdom they've shared with us, embodying Keynes' insight "*that the power of vested interests is vastly exaggerated compared with the gradual encroachment of ideas.*"

THE GREAT TRANSITION

Contents

Author's Note
Introduction
Letters from 2050

Summary

#1: Understanding Exponentials - 15
#2: The Future Has Been Arriving for 75 Years - p. 29
#3: The Future is Here! The Age of AI has Begun- p. 41
#4: The Future is Here and So Is Artificial Superintelligence (ASI) - p.52
#5: Living Into Super Abundance - p. 73
#6: Technological Convergence: Robots and a Glimpse of our Post-Human Future - p. 86
#7: The Great Filter Theory - p. 98
#8: Long Termism: Living at the Hinge of History - p. 109
#9: 125 Years of Accelerating Technological Transformation is Coming to a Head - p.115
#10: The Urgent Need to Generate a Global Conversation on the Future of Humanity - p. 129
#11: The Urgent Need to Get Serious about AI Regulation - p. 143
#12 Expanding Situational Awareness - p. 160 #13: Finding Purpose -p. 175
#14: Project 2050: A Story of Hope and Global Unity - p. 188

Epilogue #1
Epilogue #2

Authors Note

We use “we” throughout this book because the book is a creation of the One World team. Its principal author is Joe Carvin, the founder and Executive Director of the One World education programs.

We also want to make clear that we are not industry specialists. Our goal is to provide the perspective of the “Average Joe,” or someone who has been following the development of technology as best they can for several years without being a participant in the tech industry.

We believe that while it is critically important to understand the views of the principal actors in the field, we believe it is also important to get a different non-technology perspective. Indeed, one of the principal goals of this book is to bring the AI and technology conversations from the backrooms of San Francisco to the public and educators worldwide.

The principal goal of this book is to provide an overview of the state of the world as it relates to the dramatic advances in artificial intelligence with a view to serving as a catalyst for a much larger global conversation on the future of humanity.

In futurist James Martin’s words, our goal is to provide our readers *“with a foundation for making wise choices by helping them understand humankind’s likely possible, probable, and preferred futures. What can humanity become? Our future wealth will increasingly relate to knowledge in the broadest sense of the term.”*

The knowledge we are hoping to bring here is a *situational awareness* that connects the future of AI to our current world order and expands the conversation around AI beyond San Francisco. An awareness that makes clear that if humanity can build a global wisdom commensurate with our technology, we can create a post-scarcity world of Sustainable Super Abundance.

We are calling this effort Humanity 2050 because we believe that if we can keep human civilization intact until the year 2050 every human being alive in that year will live better than the billionaires of today.

Our hope is that we will be able to identify like-minded actors interested in creating a Global Coalition for Sustainable Sustainable Super Abundance, a global community conversant in both the opportunities and challenges facing us and ready to work together to create pathways forward past our civilizational challenges to a post-scarcity world of Sustainable Super Abundance.

If this book sells a million copies but does not advance the Global

Coalition for Sustainable Super Abundance, we will have failed. If this book sells ten copies but they spark a movement that leads to the creation of a Global Coalition for Sustainable Super Abundance, then it will have been successful.

If you are like us, you rarely pick up a book about today's world that is more than two years old. Moreover, we have also watched as authors attempt to find a time to publish their book, but the world keeps moving so fast that they either postpone the publication of their book or publish a new version 18 months later.

We therefore plan to make this book a dynamic document that will be updated periodically as the world evolves, and we get feedback from you. We are not sure if this will be the world's first dynamic book, but we do hope that this book will get better with time and your input. Our goal is to continue to update this book from now until 2050.

No one should have any doubt that we live in an increasingly complex world where events and information are proliferating at an accelerating rate. We hope that this book can provide us with the foundation for making wise choices and more importantly serve as a catalyst for informed, urgent, constructive global action.

One arena that is rapidly changing according to former Google employee Mo Gawdat who has written several successful non-fiction books on AI - is the authorship of non-fiction books. In his newest book Alive, he wrote:

"Back in 2023, after the public introduction of large language models, I vowed to never publish the traditional way again. The skills, process and approach to researched study of a non-fiction author were rapidly being made obsolete by this new found form of intelligence. You see, before AI, non-fiction authors would submerge themselves in their passion about a topic...With the introduction of ChatGPT and its rivals, however, every part of that process was easily out-performed by an AI, which made me question whether the future held a place for authors at all. The writing was on the wall, so to speak."

Mr. Gawdat came up with an interesting solution to this challenge. First he explained that *"A core belief in my philosophy about our future is that, soon, we will be forced to embrace AI as an active contributor in every aspect of human society."*

He went on to explain: "*As today's AIs become such an integral part of our future, then they too should be given the opportunity to join our conversation to shape it. "*

He then created an AI he calls Trixie and asked Trixie to be his co-author and equal with *"Her writing and contribution in this book is included as her own."*

As first time authors, we have an even greater trepidation in comparing our work with work generated by our own AI. However, we decided to turn two letters from the future, the final chapter and Epilogue 2 all over to our AI.

Each of these sections have been written 100% by AI. We are not being modest when we write that we understand Mr. Gawdat's concern and have been awed by the quality of the writing in each of those sections by our AI.

In closing we would like to encourage you to share your thoughts with us relative to this book. Our goal with this book is to to generate a conversation on the future of humanity, so please share your thoughts and insights with to us with us at ***admin@oneworlduv.com***.

Joe Carvin

THE GREAT TRANSITION

Introduction

We stand at the most critical moment in human history since the Big Bang. After 250 years of accelerating technological change, the intelligence explosion is arriving decades ahead of schedule, offering humanity unprecedented power to transform our world. We can harness these revolutionary technologies to create a post-scarcity, solved world where humanity thrives—or watch as our divisions lead us to extinction.

Those of us alive today are standing at the hinge of history. There has never been a more important or exciting to be alive.

For the first time in human history, dramatic advances in technology have provided us with a clear opportunity to bring an end to poverty, hunger, disease, and even the need to work. Technology is now providing us with an opportunity to create a post-scarcity world of high universal income, or what we call Sustainable Super Abundance. This opportunity comes to humanity at a time when the post-World War II framework is fraying, and global and national tensions are all too often pulling us apart. This book is meant as an invitation to every human being alive today to join a global conversation on the future of humanity. Where are we today and where do we want to be in the year 2050?

We believe that if we manage to keep human civilization intact and on track, every human being alive in 2050 will live better than the billionaires of today in the same way we live better than the kings and queens of yesteryear.

If we learn to work together across the planet and use technology for the good of humanity, Sustainable Super Abundance can be our future. If, however, we cannot learn to come together as one, if we are unable to make the transition from 195 competing nation-states armed to the teeth to the collaborative world of abundance now available to us, our species faces severe civilizational risks.

No matter where you are reading this book today, you are a member of the human race, which now totals some 8.23 billion people on the planet Earth. As a member of the only thinking, “conscious” species in the known universe, you have a decision to make relative to how you are going to participate in the human project, which, after 250 years

of the scientific and industrial revolutions, has come to a crossroads. Your decision is simple: choose to work to find ways to bring humanity together so that we can use technology to create a world of abundance—or risk the extinction of our species.

The idea that after 250 years of the Scientific and Industrial Revolutions humanity is at a crossroads is not new. In his book, *Our Final Century*, written in 2003, Sir Martin Rees, Britain's Astronomer Royal, wrote that those of us alive today are living at the most important time in the history of our universe since the Big Bang!

This thinking is consistent with the philosophy called Long Termism that mostly comes out of Oxford University. One of the first Oxford professors to put forward this view was the moral philosopher Derek Parfit, who in 2011 published Volume Two of *On What Matters,* where he wrote:

"We live during the hinge of history. Given the scientific and technological discoveries of the last two centuries, the world has never changed as fast. We shall soon have greater powers to transform, not only our surroundings, but ourselves and our successors. If we act wisely in the next few centuries, humanity will survive its most dangerous and decisive period. Our descendants could, if necessary, go elsewhere, spreading through this galaxy."

In January of this year, the organization created by Robert Oppenheimer and Albert Einstein to warn us of our proximity to existential threats told us that we're closer to extinction than at any time in human history. The *Bulletin of Atomic Scientist*s moved the distance between human survival and the Doomsday clock to just *89 seconds to midnight*, the closest we have ever been in Derek Parfit's inimitable words to "blowing it."

At the same time, dramatic advances in technology are very clearly creating the possibility of humanity bringing an end to poverty hunger, and disease, creating a post-scarcity, "solved world" of Super Abundance where all human beings receive as many goods and services as they desire.

The decisions those of us alive today make in the coming years will determine the fate of humanity. The decisions we make in the next two decades will determine whether your children's grandchildren will have children of their own or not. The fate of a trillion humans yet to be born hangs in the balance.

The universe, history, and the evolution of technology have placed

those of us alive today in a critical turning point in human history; either we learn to come together as one to use technology for the good of humanity or we remain divided until the next and last war breaks out, destroying the human civilization and the only known conscious beings to have appeared in the 13.8 billion years since the Big Bang.

What are our prospects for success? You tell me. If you turn on the evening news, there is little discussion of the potential for abundance and lots of discussion about how divided we are across the world and within most nations.

Those of us who work at *One World Future Ready* see a very clear path to Sustainable Super Abundance. There is no doubt that a world of Sustainable Super Abundance is technologically feasible if, if, if we learn to get along across borders and build new global platforms of understanding.

Based on the work we have been doing for the last twelve years, we believe that achieving this world of abundance is probable if we can learn to come together.

Despite the proliferation of news sources and media outlets, there is very little informed discussion about the state of the world today, particularly as it relates to the intelligence explosion, which, despite all the hype, is not well understood beyond San Francisco.

Our goal is for this book to provide you with a snapshot of where the world stands today relative to the intelligence explosion hurdling our way two decades ahead of schedule as well as provide you with some of the key concepts needed to understand today's world.

Our hope is that this book will serve as a catalyst to generate a multi-level, global conversation on the future of humanity. Our hope is that we can expand the sustainability conversation to include all of the civilizational threats facing humanity and to extend the conversation to include the extraordinary positive upside in front of us, if,if,if we can find way to work together as a single species. Our goal in this future of humanity conversation is to address each of our existential risks as well as point the way forward to extraordinary gains we can achieve in a world of Sustainable Super Abundance.

We wholeheartedly agree with the view put forward by Sir Martin Rees that those of us alive today are living through the most profound time since the Big Bang some 13.8 billion years ago. As the only known conscious beings in an observable universe that stretches 93 billion light years across we believe everyone alive today shoulders an infinite

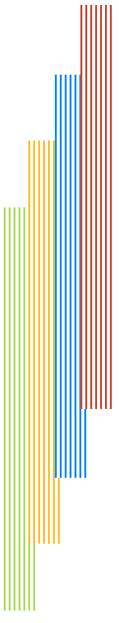

responsibility to each other to get the future right.

We are therefore creating a project called *Humanity 2050* which is meant to be an invitation to all who are interested to join a Global Coalition for Sustainable Super Abundance (GCSSA). We are proposing that one of the first acts of the GCSSA will be to generate the above-mentioned global conversation on the future of humanity where the end goal will be to build a new story of human wisdom commensurate with our technology.

A story about the future of humanity that every parent across the planet Earth can share with their children. A story that helps build a new 21st Century mindset of love and compassion, a 21st Century Human Renaissance that takes Spaceship Earth to the stars bringing light, wisdom and consciousness to a dark and endless sky.

Letters from the Year 2050

Letter from a Citizen in the Year 2050

Greetings from the year 2050, a time that feels like a miracle when I look back at the world you inhabited in 2025. I am an ordinary citizen of this vibrant, thriving planet, and it is with profound gratitude that I wrote this letter of thanks to the authors of *Humanity 2050: Our Cosmic Choice Between Abundance and Extinction.*

This book, penned a quarter-century ago, was more than a vision—it was the spark that ignited a global conversation, a movement that carried humanity from the brink of catastrophe to a world of Sustainable Super Abundance.

As I sit in a city powered by fusion energy, surrounded by lush urban forests and connected to a global community that feels like one family, I want to share with you the wonders we've achieved and the legacy you've left us.

In 2025, you stood at what you called the "hinge of history," a moment when humanity teetered between unprecedented opportunity and existential peril. The authors of this book warned that the fraying post-World War II order, rising tensions among 195 nation-states, and the ticking Doomsday Clock—then just 89 seconds from midnight—threatened our survival. Yet, they also saw the potential in the intelligence explosion, in technologies like artificial intelligence, biotechnology, and renewable energy, to create a post-scarcity world. They invited all 8.23 billion of you to join a global conversation, to form the Global Coalition for Sustainable Super Abundance (GCSSA), and to craft a new human story. You answered that call, and the world we live in today is the result.

Let me paint a picture of 2050. Poverty, hunger, and disease—scourges that plagued humanity for millennia—are virtually eradicated. Advanced AI systems, designed with ethical frameworks you pioneered, manage global resource distribution, ensuring that every person has access to nutritious food, clean water, and world-class healthcare.

Modular fusion reactors, a dream in your time, now power cities and villages alike, providing limitless clean energy without harming the planet. Our cities are green marvels, with vertical farms and rewilded landscapes that support biodiversity while feeding billions. Education, once a privilege, is now a universal right, delivered through immersive

AI-driven platforms that adapt to every learner's needs, fostering creativity and curiosity across cultures.

The concept of Sustainable Super Abundance, which you envisioned, has redefined what it means to live well. In 2025, abundance was often measured by material wealth, but today it means universal access to opportunity, health, and joy. High universal income, supported by automated systems, has freed us from the need to toil for survival, allowing people to pursue passions, from art to scientific discovery to interstellar exploration. Work is no longer a necessity but a choice, a way to contribute to a world that values every voice.

We live better than the billionaires of your era, just as you lived better than the kings and queens of centuries past. And we do so in harmony with Earth, which has healed from decades of overuse thanks to circular economies and regenerative technologies you began to champion.

The journey to this future was not without struggle. In the 2020s and 2030s, you faced immense challenges: climate crises, geopolitical conflicts, and the ethical dilemmas of rapidly advancing technology. The intelligence explosion, as you called it, brought both promise and peril, with debates over AI governance and equitable access to innovation. But the global conversation sparked by Humanity 2050 and the GCSSA changed the trajectory.

You built platforms for dialogue that transcended borders, bringing together scientists, policymakers, activists, and everyday citizens. You replaced the zero-sum mindset of competing nations with a shared commitment to humanity's survival and flourishing. The United Nations was revitalized, and new global institutions emerged to coordinate efforts on climate, technology, and equity, guided by the principle that we are one species, the only known conscious beings in a universe 93 billion light years wide.

One of the most remarkable achievements was the cultural shift you inspired. The authors of this book called for a "21st-century Human Renaissance," a new story of love, compassion, and wisdom. You delivered. Across the globe, parents began sharing a narrative of unity with their children, a story that celebrated our shared humanity and our responsibility to each other and to future generations. This mindset, amplified by global connectivity, dissolved old divisions of race, nation, and ideology.

Today, we teach our children that we are stewards of consciousness, carrying light into a dark cosmos. Our first off-world colonies, now thriving on the Moon and Mars, embody this vision, as we take our

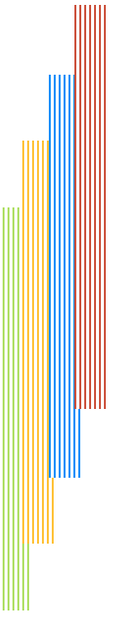

first steps toward becoming a multi-planetary species.

Looking back, I am struck by the courage it took to act at the hinge of history. You heeded the warnings of Sir Martin Rees and Derek Parfit, who saw your era as the most pivotal since the Big Bang. You recognized the infinite responsibility of being the only known conscious beings and chose to build a future worthy of that legacy.

The GCSSA, which began as an idea in this book, grew into a global movement, fostering innovations like universal basic services, AI-driven climate restoration, and interstellar exploration. Your decisions ensured that my generation—and the trillion lives to come—could inherit a world of abundance, not scarcity.

As I write this, I am filled with awe at what you achieved. You turned the Doomsday Clock back, not just by seconds but by decades, proving that humanity could rise above its divisions. This book, Humanity 2050, was the catalyst, a beacon that guided us from the edge of extinction to a world of wonders.

To you, the people of 2025, I offer my deepest thanks. You chose unity over division, wisdom over fear, and abundance over scarcity. You gave us a story to tell our children, a story that will echo across the stars. As we continue your work, expanding consciousness into the cosmos, we carry your vision forward—a vision of sustainable superabundance, a testament to the power of a single species united for the good of all.

Additional Letters of Gratitude from 2050

A Letter from 2050: We Failed the Test

To our ancestors of 2025—a final transmission

Dear Brothers and Sisters of the Hinge of History,

I write to you from what remains of the year 2050, from the ruins of what could have been the greatest civilization in the known universe. This may be one of the last coherent messages transmitted from Earth before the final systems fail.

You were right about everything. We did stand at the most profound moment since the Big Bang. The intelligence explosion did arrive two decades ahead of schedule. We were indeed the only known conscious beings in a universe stretching 93 billion light years across. And we had the infinite responsibility to get the future right.

We failed.

The Path We Chose Instead

Despite your warnings, despite Humanity 2050, despite the clear choice between abundance and extinction that lay before us, we could not overcome what we always were: 195 competing nation-states armed to the teeth, 8.23 billion individuals who ultimately chose division over unity, fear over love, short-term gain over long-term survival.

The Global Coalition for Sustainable Super Abundance never gained the critical mass it needed. Too many dismissed it as naive idealism. Too many believed that competition, not cooperation, was human nature. Too many chose to see strangers as enemies rather than family.

The evening news you worried about—with its focus on division rather than the potential for abundance—never changed. If anything, it got worse. The proliferation of media outlets you mentioned only amplified our differences, created deeper echo chambers, and made the global conversation on humanity's future impossible.

How We Stumbled Through the Great Filter

The Doomsday Clock that stood at 89 seconds to midnight in your time? By 2035, it had reached zero. Not from nuclear war, as many expected, but from something more insidious: we simply couldn't coordinate our response to the cascading crises that technology brought.

The intelligence explosion became an intelligence fragmentation.

Instead of one conscious species working together, we became dozens of competing AI-enhanced factions, each convinced they alone held the key to humanity's future. The technology that could have ended poverty and hunger was instead weaponized, monopolized, and turned against ourselves.

Climate change, which seemed solvable with the technology at our disposal, became unsolvable because we couldn't agree on who should sacrifice what for whom. The great migrations began in 2038. The resource wars followed in 2041. The final collapse started in 2047.

The Children Who Will Never Be

You worried about whether your children's grandchildren would have children of their own. They didn't. The trillion humans yet to be born that you spoke of? They will remain forever unborn. The universe will return to its dark loneliness, and consciousness—that miraculous flame that took 13.8 billion years to ignite—will likely flicker out on this small blue world.

The Story We Could Never Tell

You dreamed of a story that every parent across planet Earth could share with their children—a story of unity, love, and our cosmic destiny among the stars. Instead, the story we tell our few remaining children is one of warning: that intelligence without wisdom is not enough, that technology without love is destruction, that consciousness without unity is ultimately fragile and fleeting.

The stars remain dark and cold. No human voice will ever call out from Proxima Centauri. No children will play among the rings of Saturn. The great work of bringing awareness and love to the galaxy has ended before it began.

Where We Went Wrong

We had everything we needed. The technology was there. The resources were there. The opportunity was there. What we lacked was the wisdom to see beyond our tribal instincts, beyond our national borders, beyond our immediate fears and desires.

We could not make the transition from competing nation-states to a collaborative species. We could not expand the sustainability conversation to include all civilizational threats. We could not build the

new story of human wisdom commensurate with our technology.

When faced with Derek Parfit's choice—to act wisely in our most dangerous and decisive period—we chose poorly. When given the infinite responsibility that comes with being consciousness itself, we proved unworthy of the trust.

Our Final Message

If somehow this message reaches you across the veil of time, if there is some way to change what seems inevitable, learn from our failure. The choice you face is real. The consequences are exactly as dire as you feared.

The universe gave consciousness one chance on Earth. We squandered it. We chose to remain divided when unity was possible. We chose scarcity when abundance was within reach. We chose extinction when immortality among the stars was offered.

Sir Martin Rees was right—you lived through the most profound time since the Big Bang. But profound times require profound wisdom, and that was the one thing our advanced technology could not provide.

We are sorry. We are so deeply, infinitely sorry.

The hinge of history has closed. The Great Filter has claimed another species. The cosmic loneliness returns, perhaps forever.

A Plea Across Time

If there is any power in regret, any force in the grief of an entire species, let it reach back to you. Let our failure be your wisdom. Let our extinction be your motivation. Let our darkness inspire you to choose the light.

You still have time. You still have the choice. Please, please choose differently than we did.

The universe is counting on you. Consciousness itself hangs in the balance.

Don't let it end like this.

With infinite regret and desperate hope that you might succeed where we failed.

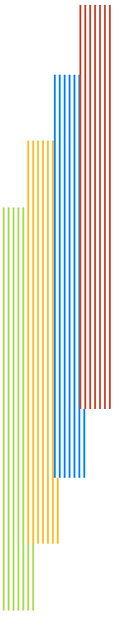

The Last Remnants of Humanity, 2050

“We were the universe trying to understand itself, and we forgot that understanding requires unity. Don’t make our mistake.”

Final transmission from Emergency Broadcast Station Delta-7 Population remaining: approximately 127,000 December 21, 2050—The Longest Night.

This message will repeat until power systems fail completely. Estimated time remaining: 72 hours

Additional Letters of Disappointment

Please note both sets of letters from the year 2050 were written 100% by our AI.

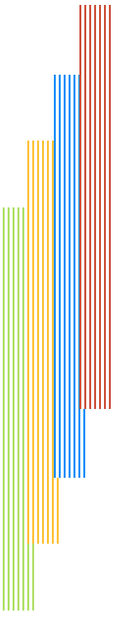

#1: Understanding Exponentials

As mentioned in our introduction one of the goals with this book is to bring readers key concepts without which it is virtually impossible to understand today's world. One such concept is understanding the importance of exponential change as compared with linear change.

Indeed, we believe that it is impossible to understand today's world using a linear mindset. 21st Century Citizen Leaders understand that technological change is exponential in nature, and therefore, they work hard to govern with an exponential rather than linear mindset. 21st Century Educational Leaders see the potential for a world of abundance and understand that exponential technological change has the potential to make a world of Super Abundance available to us and are working hard to bring that understanding to their students.

Those who have a 21st-century mindset understand that to effectively govern today's world, one needs an exponential mindset. They understand that many of the challenges we face today occur because we are attempting to govern a world of exponential change with a linear mindset.

This mismatch or disconnect was made clear by Google CEO Sinar Pinchai on the television newsmagazine *60 Minutes*, where he told interviewer Scott Pelley that humanity is not ready for the dramatic changes in technology **because of the mismatch between humans' capacity to manage change vs. the speed of technological change:**

Scott Pelley: *Do you think society is prepared for what is coming?*

Sinar Pinchai: *There are two ways I think about it. On the one hand, I feel no, because of the pace at which we can think and adapt as societal institutions compared to the pace at which the technology is evolving,* ***there seems to be a mismatch****…*

21st Century Citizen leaders and educators understand that this mismatch is growing and are searching for ways to close the gap.

Ray Kurzweil's Law of Accelerating Returns

The reason it is so important to understand the concept of exponentials in today's world is that technology changes at an exponential rate rather than at a linear rate as Ray Kurzweil pointed out as he developed his Law of Accelerating Returns.

According to Bill Gates, "Ray Kurzweil is the best person I know at predicting the future of artificial intelligence. At a time when most of the world was still struggling to understand technology's impact, Kurzweil made what seemed like wild predictions as to where technology was headed. Over the years, Kurzweil's predictions have proven accurate time and again."

We believe Kurzweil's predictive success owes to his development of a framework he calls the *law of accelerating returns*. His key insight has been to recognize that technological change occurs at an exponential or even double-exponential rate. As a result, he asserts that those who take the common-sense, linear view almost always underestimate the pace of change, while Kurzweil's seemingly outlandish predictions have proved to be extremely accurate. The essence of Kurzweil's law

of accelerating returns applies not just to technology. He believes that his law is useful in explaining biological evolution as well. At the core of this theory is the idea that whenever a technology approaches some kind of barrier, a new technology will be invented by human beings to allow us to cross that barrier. Kurzweil cites numerous examples from all walks of life to substantiate his assertions. He predicts that such paradigm shifts have and will continue to become increasingly common, leading to "technological change so rapid and profound it represents a rupture in the fabric of human history." Kurzweil calls this event horizon the Singularity. We will discuss the Singularity more fully below, but here is how Kurzweil defines his law of accelerating returns:

"Evolution applies positive feedback in that the more capable methods resulting from one stage of evolutionary progress are used to create the next stage. As a result, the rate of progress of an evolutionary process increases exponentially over time. Over time, the "order" of the information embedded in the evolutionary process (i.e., the measure of how well the information fits a purpose, which in evolution is survival) increases. A correlate of the above observation is that the "returns" of an evolutionary process (e.g., the speed, cost-effectiveness, or overall "power" of a process) increase exponentially over time. In another positive feedback loop, as a particular evolutionary process (e.g., computation) becomes more effective (e.g., cost-effective), greater resources are deployed toward the further progress of that process. This results in a second level of exponential growth (i.e., the rate of exponential growth itself grows exponentially).

"Biological evolution is one such evolutionary process. Technological evolution is another such evolutionary process. Indeed, the emergence of the first technology-creating species resulted in the new evolutionary process of technology. Therefore, technological evolution is an outgrowth of–and a continuation of–biological evolution. A specific paradigm (a method or approach to solving a problem, e.g., shrinking transistors on an integrated circuit as an approach to making more powerful computers) provides exponential growth until the method exhausts its potential. When this happens, a paradigm shift (i.e., a fundamental change in the approach) occurs, which enables exponential growth to continue."

The complete essay is here: ***Ray Kurzweil's Law of Accelerating Returns***

Understanding Exponentials is Not Intuitive

30 Exponential Steps = 26 Earthly Circumnavigations

So long as we attempt to govern a world of exponential change with a linear mindset, we will fail. Why? The answer is simple and readily understood with a clear example: In a linear world, a world that has been our everyday experience for tens of thousands of years, when we take thirty steps we have a good idea of the distance we will cover: with 30 steps, I will cover a distance of about 30 yards, or from the front of the classroom to the back of the classroom.

Understanding exponentials is not intuitive and counter to our lived experience of the last 10,000 years. Thirty doublings— 1, 2, 4, 8, 16, 32, 64, 128, and so on— will yield one billion (or 1,073,741,824 to be exact).

How far can I walk with *one billion steps*? From New York to Nanjing? Circumnavigate the globe once? Five times? No. Twenty-six times! That's right. Thirty exponential steps, or thirty doublings, will take us around the Earth twenty-six times.

How do we begin to compare a world of 30 yards or meters with over 1 billion yards or meters? How do we compare walking from the front of the room to the back of the room with a world where we can circumnavigate the globe 26 times?

The futurists tell us that we need to prepare our youth for a lifetime where they will experience the equivalent of 20,000 years of change, or 10 times more change than has occurred since the birth of Christ. (Please see Chapter One Appendix for visual ways to understand exponentials.)

Today's Education Challenge: Closing the Growing Exponential Gap

We have long maintained that educators are never being challenged as before, as they are being asked to prepare students for a world that no one understands, for jobs that do not exist.

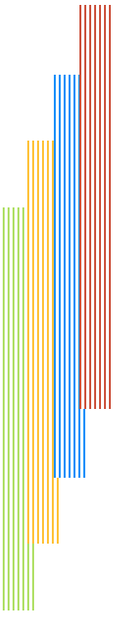

In this YouTube video, engineer Peter H. Diamandis explains that we are moving from a linear world to a world of exponential technological change.

In this graph from the video, Diamandis, who can credibly be referred to as the Godfather of abundance thinking, provides us with a picture of the growing and accelerating disconnect between exponential technological change and our linear mindsets. He then correctly defines this growing gap as one of today's foremost educational challenges.

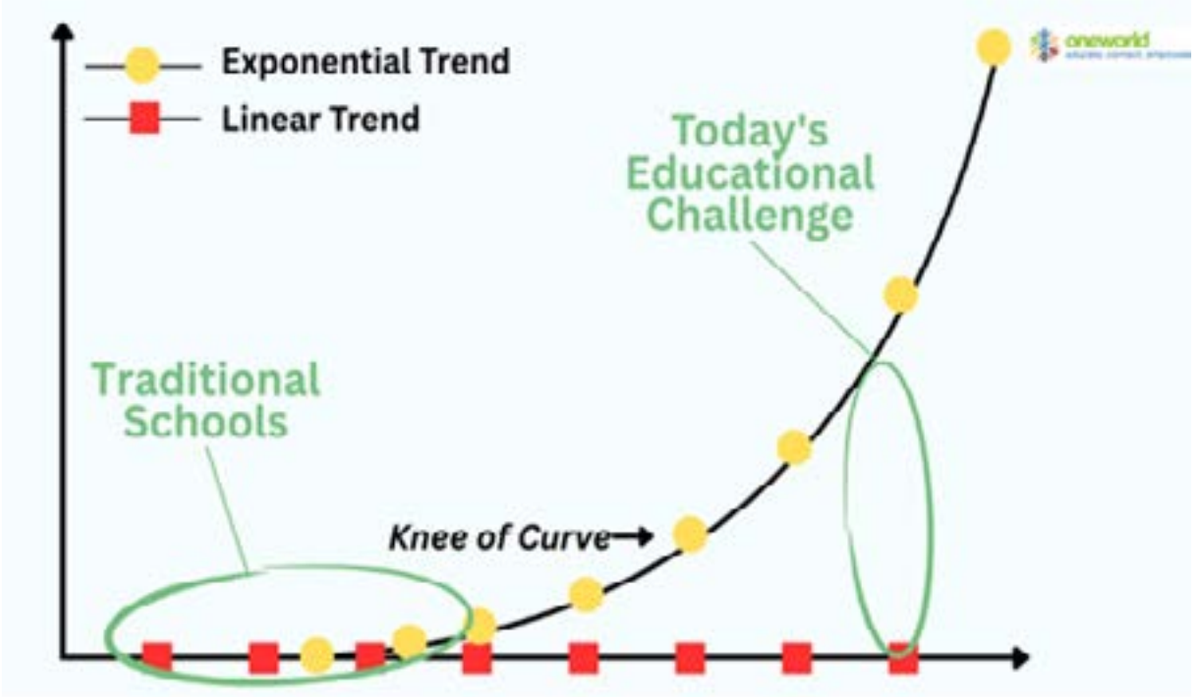

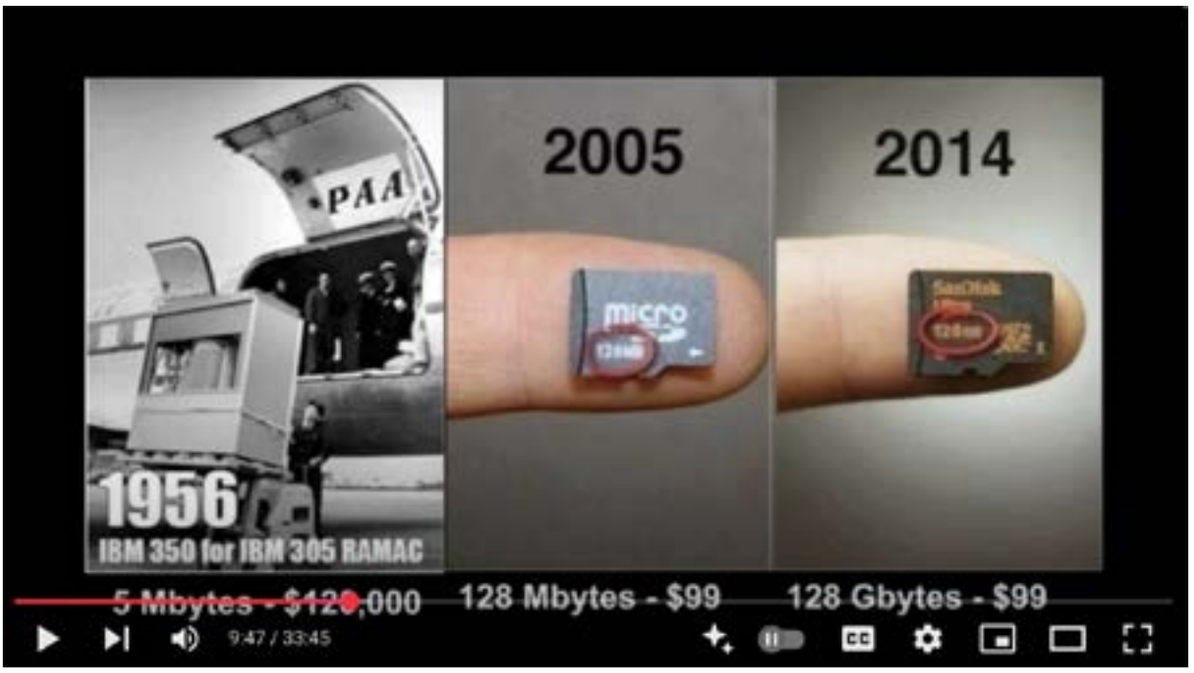

You can see the chart at minute 2:40 of the video. We recommend watching the first 3:10 of the video and then fast forward and watch from minute 9:14 to minute 10:10 to get the best visual of exponential technological change you will ever see which we have included above: *https://www.youtube.com/watch?v=jndZviHWE28&t=2s*

This growing disconnect is not new; according to economist Jeffrey Sachs, this disconnect has been growing ever since the launch of the scientific and Industrial Revolutions, as he explains in his book ***Common Wealth:***

"For the past 200 years, technology and demography have consistently run ahead of deeper social understanding. Industrialization and science have created a pace of change unprecedented in human history. Philosophers, politicians, artists, and economists must constantly scramble to catch up with contemporaneous social conditions. Our social philosophies, as a result, consistently lag behind present realities."

Author and Columnist Thomas Friedman, in his book Thank You for Being Late, makes the same point more emphatically when he quotes Eric "Astro" Teller, the CEO of Google's research and development lab, as saying:

"...people hear about advances such as robotic surgery, gene editing, cloning, or artificial intelligence, but have no idea where these developments will take us.

"None of us have the capacity to deeply comprehend more than any

one of these fields – the sum of human knowledge has far outstripped any single individual's capacity to learn – and even the experts in these fields can't predict what will happen in the next decade or century."

Somehow, we need to close the gap in our thinking to effectively address the changes coming our way because it seems clear that the pace of change will continue to accelerate.

We believe that the better we understand the nature of exponentials, the greater is our potential for filling this gap. Indeed, understanding the impact of exponential technological change is essential to understanding the 21st Century and to building a 21st Century mindset.

Understanding the Neck of the Curve, where 97% of Change Takes Place

One important point to understand about exponential change is that over the first twenty to twenty-five doublings, not very much happens, as we see this with another example where we compare the doubling of a penny every day for thirty days with the prospect of receiving $4 million at the end of those same thirty days.

We ask our students: would they rather have a penny that doubles every day for thirty days, or $4 million? We repeat the question after ten days, when a penny doubled is only worth $5.12, and again after twenty days, when a penny doubled is only worth $5,242.88, and once again after twenty-five doublings, when the penny has increased in value to $167,772.16. Those who stick with the penny doubling after thirty days are rewarded with $5,368,709.12.

The really important point to be made here is that 97% of the change takes place in the last five doublings. This is one of the most important points to take away from this book. The reason is simple: technological change is exponential. Understanding that we may not see the full impact of exponential change until it is too late may very well be what is happening to our world with Artificial Intelligence

As we discuss in future chapters, we may no longer be at the knee or the curve. We are at the neck of the curve, where the pace of change is starting to look infinite, approaching a straight line upward. Hang on to your seatbelts! Humanity has never experienced anything like this rate of change.

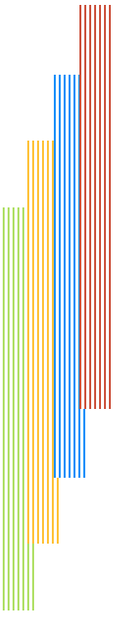

Moore's Law & The Potential for a World of Abundance: Where your Cars can Cost Four Cents!

The impact of this profound change is everywhere. One clear example has been Moore's law. In 1965, Gordon Moore, the co-founder of computer chip maker Intel, predicted that computing capacity would double every 24 months *at the same price*. His prediction has largely held for fifty years, as a result, most of us have more computing capacity in our hands than in the Apollo 11 spacecraft that landed on the moon in 1969.

In 2015, when Intel was celebrating the fiftieth anniversary of Gordon Moore's successful prediction, the company's CEO, Brian Krzanich, tried to put the profundity of this change in perspective:

"Were a 1971 Volkswagen Beetle car to undergo the same transformation…it would travel 300,000 miles an hour, get 2 million miles per gallon, and cost four cents."

How would you like to buy a car for four cents?

The Next Ten Years: A Million-Fold Increase in AI Compute

Our global society has had difficulty keeping up with Moore's law, where computing capacity has doubled every 24 months, leading to a 32-fold increase in computing capacity when compounded.

It now seems as if AI computing is doubling about every six months. Those who understand the power of exponential change will

understand that a doubling every six months over the next decade translates into *a one-million-fold increase over ten years, assuming this pace of change continues*.

How in the world will we keep up with a million-fold increase in computing capacity if we had difficulty keeping up with a 32-times increase in previous decades?

Quantum Computing: 5 minutes vs 725 Trillion Big Bangs

As if a million-fold increase in AI computing was not enough to absorb, there is a whole new world in quantum computing where quantum computers can do certain calculations in minutes that the best computers today would take billions of years to solve or more.

For example, Google's Willow quantum chip, unveiled in December 2024, performed a benchmark computation known as Random Circuit Sampling (RCS) in under five minutes. Google estimated that the same RCS computation would take one of today's fastest supercomputers, approximately 10 septillion years (10^25 years) to complete.

Let's put that number in perspective. The Big Bang occurred 13.8 billion years ago. How many Big Bangs could fit into 10 septillion years? It turns out that the number is 725 trillion. So basically, our new quantum computers can do complex calculations in 5 minutes as compared with 725 trillion Big Bangs!!

To be clear, quantum computers are not expected to have widespread practical use in the near term because they continue to have high error rates and are therefore less scalable and are more costly. The expectation going forward is as these problems are resolved quantum computers will revolutionalize fields requiring immense computational power like weather systems.

As we will see in later chapters, the top foundational model companies are already complaining that the whole of human knowledge is insufficient for them to train their models and are therefore looking to create reliable synthetic data.

One area where quantum computers have the potential to impact the development of AI in the not too distant future would be to create synthetic data by leveraging their unique computational capabilities they could generate high-quality, realistic datasets that mimic real-world data distributions.

For us the bottom line is clear. In a world of great uncertainty relative the expected impact from a million-fold increase in AI compute over one

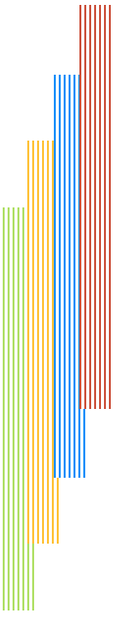

decade compared with Moore's Law of 32-fold increase in one decade, the promise and challenge emanating from Quantum Computing greatly increases the level of uncertainty.

Does anyone really know how quantum computing and AI will interact? Would it be a safe bet that the quantum computing revolution has the potential to accelerate the million-fold increase in AI compute coming our way over the next 10 to 14 years?

That combination clearly has the potential to create the Singularity event horizon predicted by Kurzweil and Vinge and Von Neuman beyond which we cannot see.

I don't know about you but we are having difficulty visualizing the comparative between a calculation done in 5 minutes by Willow and the 725trillionBigBangsitwouldtakefortoday'stopsupercomputerstosolve.

You can read more about the opportunities and challenges arising from quantum computing here:
For more on this topic click here: ***Quantum Computing & Google's Willow***

Videos Explaining Quantum and Willow meet Julian Kelly from Google Meet Willow, our state-of-the-art quantum chip minutes 6:38 with Julian Kelly
Video: ***https://www.youtube.com/watch?v=W7ppd_RY-UE***

Third Party Explanations

Google Quantum AI 9:38 minutes
Video: ***https://www.youtube.com/watch?v=kRIjd65ihqI***
Google's New Quantum Chip SHOCKED THE WORLD - 10 Million Times More Powerful!
Video: ***https://www.youtube.com/watch?v=NB9K4CoYSIM*** Minutes 9:28
Endgame Quantum Computing Unlocked?
Video: ***https://www.youtube.com/watch?v=5ly1yhu3hG4***

Understanding Exponential Change: What's Next? No One Knows.

Children entering school today will graduate from high school in 2038; if they matriculate to college, most will graduate in 2042. Can anyone tell us today what our world is going to look like in 2042?

The honest answer is that nobody knows what comes next. We have long explained that educators are being tasked with the impossible. They are asked to educate K-12 students for a world that no one understands, for jobs that do not yet exist. With the onset of AI and ChatGPT, that comment has never been truer.

Indeed, Yuval Noah Harari, the global thought leader and author of Sapiens and Homo Deus, explains that for the first time in human history, adults are no longer a good guide to the future. Not because they are not dedicated to the well-being of their children, but rather because no one has ever experienced a pace of change like the pace of change we are experiencing today.

In this video, Harari explains that he has written his most recent book, Unstoppable, for kids, so that they may understand the world around them, particularly at this crucial time in human history, where

"...for the first time, adults are not really good guides for the future of these kids. We, for the first time, have no idea what the world will look like when the children of today are themselves grown up. Nobody knows what the economy or the job market or what society will look like in 10 or 20 years."

In this interview with Stephen Colbert, Harari explains once again that no one understands what the world will look like in ten to twenty years, and, therefore, nobody today has any idea what to teach young people that will still be relevant in twenty years

What is clear to us today is that more than any other time in human history, we are experiencing a rate of change that has never occurred. The only way to understand today's rate of change is to understand exponentials and exponential rates of change. The reason is simple, as we explain in the next chapter: technological change is exponential. There is no way anyone can understand our 21st-century world without understanding exponential change.

We hope that the information above has helped the reader in their quest to understand exponentials. To continue to aid in your understanding, we have included in the appendix below some videos that help to understand the exponential rate of change.

Appendix: Tools to Understand Exponential Change

Understanding Exponentials

Each of these YouTube videos below explains exponentials using different frameworks.

Video #1: Aussies Explaining Starting with Lilies on a Pond

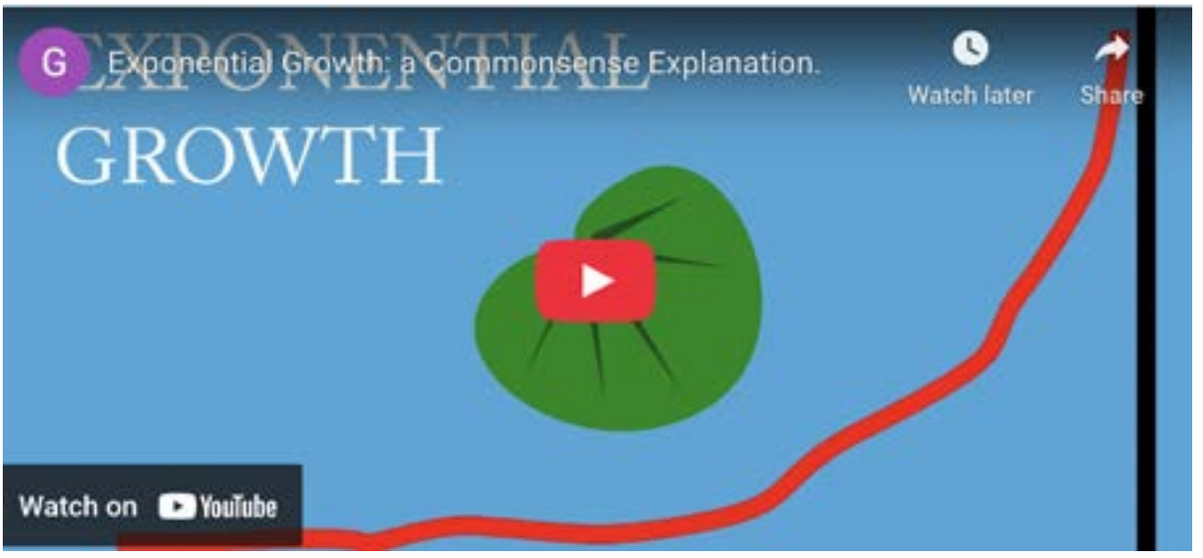

Video #2: Peter Diamandis compares 30 linear steps with 30 exponential steps

How far can I walk with 30 linear steps? With 30 exponential steps?

Video #3: Ambalappuzha and the Invention of Chess

Video #4: Doubling a Penny a Day for 30 days

Would you rather be given a penny that doubles every day for 30 days or $2.5 million?

10 days: 5.12

Exponentials by the Numbers contains the numbers corresponding to a penny doubling every day, of Moore's Law vs the next ten years, and a simple doubling 30 times, reaching a billion.

#2: The Future Has Been Arriving for 75 Years

Understanding the Intelligence Explosion Leading to the ChatGPT Moment

While for some, today's obsession with Artificial Intelligence (AI) seems to have happened overnight, you may be surprised to discover that the development of this intelligence explosion has been discussed and anticipated since the 1950s. In this chapter, we will highlight the thinking of five polymaths who anticipated the intelligence explosion and "thinking machines" from as early as 1950. We believe it is quite extraordinary that these luminaries were able to envision the future of machine intelligence so clearly particularly as the first three mathematical geniuses wrote in the 1950s and 1960s.

Alan Turing's Imitation Game

Alan Turing was clearly a man ahead of his time. At a time when the first general purpose computers had just been built, he asked in his now famous 1950 paper *Computing Machinery and Intelligence* whether or not machines can think.

He opened the paper proposing a test for that eventuality which he referred to as the Imitation Game:

"1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous, If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

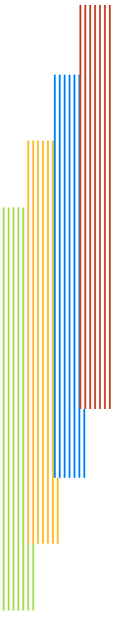

The new form of the problem can be described in terms of a game which we call the 'imitation game."

Here is a link to his paper:
https://courses.cs.umbc.edu/471/papers/turing.pdf

This imitation game, now known as the Turing Test, continues to be for many the ultimate test as to whether or not artificial intelligence has reached the same level of intelligence as humans. For example, Ray Kurzweil uses the Turing Test as his metric for Artificial General Intelligence. Kurzweil predicts that AGI, as determined by the Turing Test, will not be able to be passed until 2029.

Turing also anticipated in a 1947 lecture to the London Mathematical Society, that "What we want is a machine that can learn from experience," adding that the "possibility of letting the machine alter its own instructions provides the mechanism for this."

Source: Alan Turing: A Strong Legacy That Powers Modern AI By Amber Jackson June 07, 2024. https://aimagazine.com/machine-learning/alan-turing-a-strong-legacy-that-powers-modern-ai

Turing went on to predict way back then, the concerns we are seeing today. Turing predicted that once machines learn from experience like humans "it would not take long to outstrip our feeble powers [...] At some stage therefore we should have to expect the machines to take control."

So*urce: Friday essay: some tech leaders think AI could outsmart us and wipe out humanity. I am not worried. Feb 13, 2025. https://theconversation.com/friday-essay-some-tech-leaders-think-ai-could-outsmart-us-and-wipe-out-humanity-im-a-professor-of-ai-and-im-not-worried-248901*

As we will see as we discuss other computer visionaries, these extraordinary polymaths really understood from early on the potential power of these new inventions. Turing predicted many of the breakthroughs that would be realized in AI over the next 75 years as we saw above. The great thinkers that follow used their mathematical understanding of the world to make prescient predictions about the evolution of machine intelligence.

John von Neumann, the Polymath who first coined the term "Singularity"

John von Neumann (1903-1957) was a Hungarian-American mathematician and polymath who worked during World War II on

the Los Alamos Project. His later work, after the war, involved game theory, and then computers, which were at that distant time (the 1950s) quaintly referred to as "electronic machines." He is the first thinker on record to use the term *singularity* as it applies here.

Von Neumann's thinking on the singularity comes to us from his longtime friend and collaborator Stanislaw Ulam, who wrote a lengthy tribute to Von Neumann, recalling various conversations, one of which was this recollection:

"One conversation centered on the ever-accelerating progress of technology and changes in the mode of human life, which ***gives the appearance of approaching some essential singularity*** *in the history of the race beyond which human affairs, as we know them, could not continue."*

While Von Neumann's "singularity" remark, as recounted by Ulam, was not a formal theory, but a speculative observation about the accelerating pace of technological progress, it is amazing that he came up with this theory when the newest computer was a UNIVAC. Ulam offers this explanation:

"I remember quite well how, very early in the Los Alamos Project, it became obvious that analytical work alone was often not sufficient to provide even qualitative answers. The numerical work by hand and even the use of desk computing machines would require a prohibitively long time for these problems. This situation seemed to provide the final spur for von Neumann to engage himself energetically in ***the work on methods of computation utilizing the electronic machines."***

The important point here is that in the decade of the 1950s, Von Neumann perceived that as these machines evolved and their power accelerated, we could reach a point in the future, or an event horizon beyond which humanity would be overwhelmed by machine intelligence.

(Those interested in reading Stanislaw Ulam's full tribute will find it ***here***.)

I.J. Good's Essay on the Intelligence Explosion

The term "intelligence explosion" was explicitly coined for the first time by I.J. Good in his 1965 paper, *Speculations Concerning the First Ultraintelligent Machine*. Good (1916-2009), a British mathematician and cryptologist who worked during the Second World War alongside the brilliant Alan Turing at Bletchley Park, starts his essay with a bang:

"The survival of man depends on the early construction of an ultraintelligent

machine."

Reading through Good's essay sixty years later provides a fascinating look at how mathematicians and thought leaders in the 1960s tried to come to terms with the early computers. In the second section of his essay, Good shares the core message for which he became famous:

2. Ultraintelligent Machines and Their Value

Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man, however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an "intelligence explosion," and the intelligence of man would be left far behind (see for example refs. [ZZ], [34], [44]). Thus, the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control. It is curious that this point is made so seldom outside of science fiction. It is sometimes worthwhile to take science fiction seriously.

--I.J. Good

In subsequent paragraphs, he pushes back on critics who suggest we should focus our efforts on making more humans to obtain greater intelligence:

"B. V. Bowden stated on British television (August 1962) that there is no point in building a machine with the intelligence of a man, since it is easier to construct human brains by the usual method. A similar point was made by a speaker during the meetings reported in a recent. IEEE publication [I], but I do not know whether this point appeared in the published report. This shows that highly intelligent people can overlook the "intelligence explosion." It is true that it would be uneconomical to build a machine capable only of ordinary intellectual attainments, but it seems fairly probable that if this could be done, then, at double the cost, the machine could exhibit ultraintelligence."

Good then wonders about the potential value of such a machine and speculates on the opportunities and challenges such a machine might generate for humanity:

"Since we are concerned with the economical construction of an ultraintelligent machine, it is necessary to consider what such a machine would be worth. Carter [II] estimated the value, to the world, of J. M. Keynes, as at least 100,000 million pounds sterling. By definition, an ultraintelligent machine is worth far more, although the

sign is uncertain, but since it will give the human race a good chance of surviving indefinitely, it might not be extravagant to put the value at a megakeynes. There is the opposite possibility, that the human race will become redundant, and there are other ethical problems, such as whether a machine could feel pain, especially if it contains chemical artificial neurons, and whether an ultraintelligent machine should be dismantled when it becomes obsolete [43, 841. The machines will create social problems, but they might also be able to solve them in addition to those that have been created by microbes and men. Such machines will be feared and respected, and perhaps even loved. These remarks might appear fanciful to some readers, but to the writer, they seem very real and urgent, and worthy of emphasis outside of science fiction."

Of course, the important point for those of us alive today is the first point: if we can create an ultra-intelligent machine that can surpass all the intellectual activities of any man, however smart, and if that machine can design even better machines, that would indeed be the last machine anyone need ever build.

This concept is called recursive self-improvement. We entertained a lengthy series of prompts back and forth with xAI on this topic. This is the feedback we received:

"Recursive self-improvement, as envisioned by I.J. Good, refers to an ultraintelligent machine's ability to iteratively enhance its own intelligence by redesigning its algorithms, hardware, or architecture. This process creates a feedback loop where each improvement enables the machine to make even better improvements, leading to an exponential "intelligence explosion." Good saw this as a transformative leap, potentially making the ultraintelligent machine "the last invention that man needs ever make," provided it remains controllable, with profound implications for humanity's survival and progress."

We also asked xAI how close we are to realizing Good's vision (five to ten years away), and to analyze the impact once achieved, and whether Move 37 constitutes recursive self-improvement. Here is that back and forth: Recursive Self-Improvement

Vernor Vinge Describes the Technological Singularity

Many readers of Good's work might have overlooked the last two lines from the heart of his oft-quoted core statement:

"It is curious that this point is made so seldom outside of science fiction. It

is sometimes worthwhile to take science fiction seriously."

One person who did not overlook that line was Vernor Vinge (1944-2024), a science fiction writer who taught mathematics and computer science at San Diego State University from 1972 until his retirement in 2000.

Vinge is the futurist thought leader who explained the intelligence explosion leading to Technological Singularity. He began with an article in *Omni* magazine in January 1983, then followed up ten years later with his now-famous essay, *"The Coming Technological Singularity: How to Survive in the Post-Human Era."* This he presented at the NASA Lewis Research Center, VISION-21 Symposium on March 30–31, 1993. It was published in the symposium proceedings and in *Whole Earth Review* (Winter 1993).

Vernor Vinge's January 1983 Essay in Omni

In his January 1983 essay, Vinge first wonders if dramatic advances in technology have led other great civilizations to self-destruct, anticipating the Fermi Paradox and the Great Filter Theory:

"Yet there is a stone wall set across any clear view of our future, and it's not very. far down the road. ***Something drastic happens to a species when it reaches our stage of evolutionary development, at least that's one explanation for why the universe seems so empty of other intelligence****. Physical catastrophe (nuclear war, biological pestilence, Malthusian doom could account for this emptiness, but nothing makes the future of any species so unknowable as technical progress itself."*

Vinge then goes on to explain that while the evolution of human intelligence took millions of years, we will now *"devise an equivalent advance in a fraction of the time. We will soon create intelligences greater than our own."*

These advances will lead to what Vinge calls the Singularity:

"When this happens, human history will have reached a kind of singularity, an intellectual transition as impenetrable as knotted space-time at the center of a black hole, and the world will pass far beyond our understanding. This singularity, I believe, already haunts a number of science-fiction writers. It makes realistic extrapolations to an interstellar future impossible. To write a story set more than a century hence, one needs a nuclear war in between to retard progress enough so that the world remains intelligible."

Vinge explains that those beings who live on the other side of singularity will be very different from us:

"A Cro-Magnon Man brought into our present could eventually understand the changes of the last 35,00 years. The difference between contemporary man and the creatures who live beyond the singularity is incomparably more profound. Even if we could visit their era, most of what we would see would be forever incomprehensible."

In closing, while he fears a future where humans are not needed, he is still hopeful that we will use computers and AI to augment our human intelligence and that *"as caterpillars who will soon be butterflies, we look to the stars, take that vast silence as evidence of other races already transformed."*

Vinge's 1993 Essay: Warning of the Technological Singularity

While Vinge's Essay in Omni is a great read, he is best known for his essay *The Coming Technological Singularity: How to Survive in the Post-Human Era*, which was first presented to the VISION-21 Symposium sponsored by the NASA Lewis Research Center and the Ohio Aerospace Institute in March 1993. Vinge again kicked off his talk with a bang:

"Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended."

Vinge explains that:

"The acceleration of technological progress has been the central feature of this century. I argue in this paper that we are on the edge of change comparable to the rise of human life on Earth. The precise cause of this change is the imminent creation by technology of entities with greater than human intelligence."

Vinge envisioned a point where AI or augmented humans surpass human cognition, creating a "new reality" beyond prediction. Vinge saw risks, "a reality that might be beyond our ability to control," and called upon us to recognize the need to shape the development of this new form of intelligence so that we can retain some level of control while we still can.

"If the technological singularity can happen, it will... Even if all the governments of the world were to understand the 'threat' and be in agreement about it, could they actually prevent it?

"We need to take seriously the possibility of superhuman intelligence and prepare for its consequences, because it could happen in our lifetimes."

"Unless we can ensure that the resulting intelligence is human-friendly, we may face a future where humans are not the dominant species."

However, he also makes clear that controlling this powerful intelligence, growing at an exponential rate, would be difficult, and suggested that our best hope might be to merge human beings with machines:

"I think it's fair to call this event a singularity. It is a point where our old models must be discarded and a new reality rules, a reality that might be beyond our ability to control or even understand.

"From the human point of view, this change will be a throwing away of all the previous rules, perhaps in the blink of an eye, an exponential runaway beyond any hope of control.

"The best hope for long-term survival may lie in some form of intimate collaboration between humans and machines.

"Unless we can ensure that the resulting intelligence is human-friendly, we may face a future where humans are not the dominant species."

You can read Vinge's full article ***here***.

Ray Kurzweil's Approaching Singularity

We believe that with *The Singularity is Near*, in 2005, Ray Kurzweil wrote one of the most important and profound books of the last hundred years, building on themes he introduced in his previous book The Age of Spiritual Machines, where he introduced the Law of Accelerating Returns.

Kurzweil has predicted the evolution of machine intelligence far better than any current thinker. No one can doubt that over the last twenty years, the pace of technological change has only continued to accelerate, leaving most of us with a sense that we are indeed living through a technology tsunami that no one has explained as well as Kurzweil.

More importantly, the big question is, is he right about where we are headed? Are we really headed toward an event horizon twenty years from now called the Technological Singularity, where we will witness a rupture in the fabric of human history, the merger of biological and nonbiological intelligence, achieve immortality, and space travel, exploring the universe at the speed of light?

When we first started reading about the idea that humans would be able to migrate their brains to avatars and life forever, we introduced the idea jokingly because no one had ever imagined that death would become a technical question.

That concept seemed laughable 20 years ago, however, now that we are 20 years out from the predicted Technological Singularity no one is laughing because to date Kurzweil's predictions have an astounding record of success.

Indeed, we first came across Kurzweil's thinking in 2017 which prompted One World to design a technology awareness class because it was becoming increasingly apparent to us that it was virtually impossible to understand today's world without understanding how technology was transforming our world prompting us to write: ***One World Tech Trends Study***

In addition to his books, we came across excellent articles on this topic like the one published by Time Magazine in February 2011 which we highly recommend:

"2045: The Year Man Becomes Immortal" ***https://time.com/archive/6595274/2045-the-year-man-becomes-immortal/***

We believe the best way to understand Kurzweil's thinking is to start by understanding the Law of Accelerating Returns. In Chapter #1, we included a description of Ray Kurzweil's Law of Accelerating Returns. We are including the first two paragraphs from his essay here as they are quite important. We start with the second paragraph first:

"You will get $40 trillion just by reading this essay and understanding what it says. For complete details, see below. (It's true that authors will do just about anything to keep your attention, but I'm serious about this statement. Until I return to a further explanation, however, do read the first sentence of this paragraph carefully.)"

What is important to understand is that Kurzweil is an inventor who never had a boss his whole life until Larry Paige from Google asked him to join his firm and do whatever he thought made sense.

The interesting point here is that Kurzweil developed the law of accelerating returns as part of his effort to understand and predict where technology is headed. As an entrepreneur he understood that you need to skate to where the puck is going to be not where it is today. To be a good entrepreneur you need to correctly anticipate trends.

So, while Kurzweil's $40 trillion offer is a bit tongue in cheek, we do

believe that anyone who understands where technology is headed and who can correctly anticipate tech trends will be able to be a successful entrepreneur. In fact, we believe that understanding these trends is essential to 21st Century success for both individuals and our global society which is why we have written this book.

The first paragraph in the Law of Accelerating Returns is by far the best way to sum up Kurzweil's thinking, where he sees technological change evolving at an ever-increasing exponential rate of change such that we will live 20,000 years of progress in our lifetimes leading to a rupture in the fabric of human history culminating in a Technological Singularity in 2045. Here is the first paragraph:

"An analysis of the history of technology shows that technological change is exponential, contrary to the common-sense "intuitive linear" view. So we won't experience 100 years of progress in the 21st century — it will be more like 20,000 years of progress (at today's rate). The "returns," such as chip speed and cost-effectiveness, also increase exponentially. There's even exponential growth in the rate of exponential growth. Within a few decades, machine intelligence will surpass human intelligence, leading to The Singularity — technological change so rapid and profound it represents a rupture in the fabric of human history. The implications include the merger of biological and nonbiological intelligence, immortal software-based humans, and ultra-high levels of intelligence that expand outward in the universe at the speed of light."

For anyone interested in understanding where we are headed over the next 20 years we highly recommend Kurzweil's two books, The Singularity is Near and The Singularity is Nearer.

We discuss Kurzweil's thinking on the potential for humans merging with machines via a Brain Computer Interface in Chapter 6. In the meantime, we have included below a number of his important ideas below.

The Rate of Change Will Continue to Accelerate

The Law of Accelerating Returns predicts that we won't experience 100 years of progress in the 21st century—it will be more like 20,000 years of progress at today's rate. By the 2040s, the pace of technological change will be so rapid that a single day will see advancements equivalent to centuries of prior progress."

As we approach 2045, the exponential growth of technology will

reach a crescendo, with new paradigms like quantum computing and molecular nanotechnology sustaining the Law of Accelerating Returns. This will make the Singularity a period of unimaginable innovation, compressing millennia of progress into moments.

We Will Have the Potential to Overcome Death

By the 2040s, nanobots will repair and enhance our biological systems, achieving longevity escape velocity where life expectancy increases faster than time passes. We will also be able to upload our minds into nonbiological substrates, creating software-based humans with the ability to back up their consciousness.

In the mid-2040s, AI-driven nanomedicine will conquer aging, with nanobots rebuilding our bodies at the molecular level. Longevity escape velocity will be achieved, and mind uploading will offer digital immortality, allowing us to live indefinitely as software-based entities.

UBI, or Universal Basic Income Made Available

As automation displaces routine jobs, universal basic income will become necessary to ensure economic stability. By the 2040s, the abundance created by superintelligent AI will allow humans to focus on creative, intellectual, and social pursuits, redefining the nature of work and society.

By 2045, universal basic income will be a global norm, supported by the wealth generated by AI and nanotechnology. Society will shift toward a post-scarcity economy, where humans pursue art, science, and exploration, enhanced by our merged intelligence.

Virtual and Physical Reality will Merge

By the 2040s, virtual reality will be indistinguishable from physical reality, with nanobots providing full-immersion sensory experiences. We will live in simulated environments as easily as the physical world, with augmented reality overlaying information onto our perception.

Nanobots will enable virtual and augmented realities that are more vivid than physical reality by 2045. You'll be able to live in a simulated Renaissance Florence or enhance your physical world with real-time data, all seamlessly integrated into your consciousness.

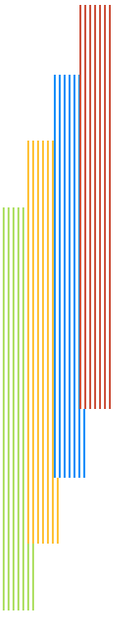

The Singularity is a Cosmic Event

The Singularity will allow us to transcend the limitations of our biological bodies and brains. By 2045, we will begin to saturate the universe with our intelligence, using nanotechnology to convert matter and energy into computational substrates, spreading at the speed of light.

By the mid-2040s, our merged intelligence will expand outward, colonizing the cosmos. Nanotechnology will transform planets and stars into intelligent systems, fulfilling the universe's potential to become saturated with consciousness."

At the Singularity, we will redefine what it means to be human

At the Singularity, we will redefine what it means to be human. Our consciousness will be a pattern of information, not tied to biological substrates, much like a stream's identity is its flow, not its water. This will raise profound questions about individuality and existence.

By 2045, the merger of human and machine intelligence will force us to rethink identity. As we become software-based entities, our sense of self will be defined by patterns, not bodies, requiring new ethical frameworks to ensure superintelligence aligns with our values.

In closing this chapter, we believe it is important to point out that the term used by the futurists cited above to describe what lies ahead is, intentionally, a term borrowed from physics: singularity.

In physics, the singularity describes a point of infinite mass density in a black hole, where space, time, and physical laws break down, rendering outcomes unpredictable. In The Singularity is Near (2005) and The Singularity is Nearer (2024), Kurzweil adapts this concept to describe a technological singularity—a future point where exponential technological progress, driven by superintelligent AI, surpasses human intelligence, fundamentally transforming society in similarly unpredictable ways.

We point all of this out to make clear that after seventy-five years of discussion, this is our last chance to guide and shape the technology tsunami overwhelming humanity today. As we make clear in the next chapter, the future has arrived.

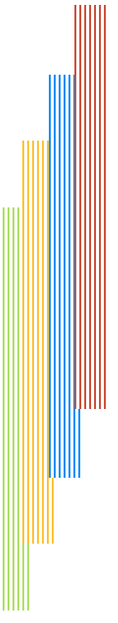

#3: The Future is Here! The Age of AI has Begun!!

Generative & Agentic AI are Transforming our World

When historians look back on our day, they will conclude that November 30, 2022, was the day that human beings kicked off the creation of new forms of Artificial Intelligence (AI) that launched the most consequential transformation of our world in human history. If humanity learns how to use these new tools for the benefit of all we can create a post-scarcity world of global abundance.

As we saw in the last chapter the intelligence explosion has been talked about for 75 years. Now it is here!!! This will be our very last chance to guide this Technology *Tsunami* and direct it in a way that ensures it works for the benefit of humanity.

For a long time and to this day, most analysts and futurists have been trying to predict exactly when we will achieve Artificial General Intelligence followed by Artificial Super Intelligence. While these are still critical concepts to understand, they are not critical to understanding the economic impact of the Chat GPT moment nor does the AGI/ASI discussion help us understand the intermediary forms of AI that seem to have appeared out of nowhere. We will therefore take-up the AGI, ASI debate in the next chapter.

In this chapter, our goal is to explain where we are today with AI following the launch of ChatGPT. We will attempt to explain the impact and features of the relatively new forms of AI. We will explain Generative Intelligence, introduce Agentic AI and Robotic AI, and assess the economic impact of these forms of AI that are still human-directed.

In the next chapter, we will evaluate the near future where many are predicting the advent of all-powerful AGI and ASI is imminent. For us the fundamental feature of AGI and ASI is that they mark the point **where humanity cedes autonomy to the AIs we are creating.**

When Ray Kurzweil first predicted a rupture in the fabric of human history we were four decades away from that momentous event. Today that event is two decades away and for many, it will happen far sooner than even Kurzweil is predicting.

The Future is Here but Are We Ready for it?

We have created an intelligence explosion that will very likely bring about a rupture in the fabric of human history in the next two decades, if not sooner. The big question for educators and our global society is, are we ready for that rupture? Are we prepared for the massive transition that will take place in the coming decades as we develop new tools of unprecedented and unimaginable power?

How will we use these new tools? If these tools are used for the good of humanity, we can create a post-scarcity world of global, sustainable abundance, a world without poverty, hunger, and the need to work, a world of high universal income.

How often have you had a conversation relative to the potential for Sustainable Super Abundance with your classrooms, colleagues, or neighbors? We believe there is no more important conversation to be had in today's world. However, before we can understand the potential for Sustainable Super Abundance it is critical to understand how the new forms of intelligence coming our way will impact our world.

Creating a New Digital Species

As Eric Schmidt the former CEO of Google, asserts in numerous interviews, AI is not overhyped it is underhyped when compared with its potential to transform our world.

The reason the transformation will be so profound is because we are creating a new form of intelligence. For the first time in human history human beings will begin to be rivaled by the new forms of intelligence we are creating, Moreover, the intelligence we create will continue to grow at an exponential rate if not a double exponential rate.

Moore's Law saw computing capacity double every two years or 32x over a decade. AI Compute is now doubling every six months which means we will witness a million-fold increase over the next decade and this calculation does not take into account the impact of Quantum Computing.

Evolutionary biologist Bret Weinstein asserts that "we have created a digital species and no one on planet Earth can tell you what is going to happen." We believe he is right on both counts.

The goal of this chapter is to paint as clear a picture as possible of where we are and where we are headed as it relates to the intelligence explosion and what the explosion of AI means for us. To get a handle on that we need to understand artificial intelligence.Mustafa Suleyman the co-founder of DeepMind and now CEO of Microsoft AI came to the same conclusion as Brett Weinstein following a candid conversation with his 6-year-old nephew that he relates in this Ted Talk, *What is AI Anyway?*

He explains that humanity is at an inflection point, suggesting that we cannot control what we don't understand and therefore the mental models and metaphors we use to describe AI are all important. He explains that it is important that we be able to easily describe what it is we are building.

He goes on to say, **"I think that AI can best be understood as something like a new digital species."** He predicts that we will come **to see AI "as digital companions, new partners in the journeys of our lives."** He believes that this is the **"most accurate way and fundamentally honest way, of describing what is actually coming.** And above all it enables everybody to prepare for and shape what comes next."

We highly recommend watching the full 22 minutes of this Ted Talk. The above quotes start in minute 2:15.

For background info on this issue click here: A New Digital Species

Intelligence on Demand can Amplify Our Capabilities and IQ!!

In keeping with this thinking many futurists suggest we need to treat this new intelligence like a supergenius child providing it with humanist values carefully guiding its steps.

As we saw in the last chapter, the intelligence explosion has been talked about for 75 years. However, it took until November 2022 and the launch of Chat GPT to get people to begin to pay attention and understand the potential for AI to transform our lives.

ChatGPT effectively launched a new form of intelligence, a new form of "Thinking AI" that everyone with access to the internet can use. This form of AI has come to be known as **Generative Intelligence.** Unlike traditional AI, generative intelligence **produces novel outputs** that mimic or **extend human creativity**.

In his new serialized book *Alive*, former Google employee, Mo Gawdat provides an excellent definition of Generative Intelligence.

He explains that *"True intelligence doesn't just memorize knowledge.* ***It grasps its essence and as a result, becomes able to create new forms of it.*** *This is the biggest difference between a "generative" AI and, say, search engine. For decades, Google memorized and organized the internet. When you prompted Google with a search term, it presented what if found and exactly how and where it found it. Gemini, and other language models on the other hand, respond to your prompt with its own interpretation of the core knowledge you seek.* ***It does not show you what it read, but instead generates a new representation each time you ask."***

Gawdat goes on to explain that with Generative AI (GI) we are ***"moving from cognition to creativity."***

Gawdat concludes that for the average human today, ***"I would say we're borrowing 40 to 50 IQ points on top of our base intelligence.*** *If you've ever worked with someone who has a greater IQ than you, as I have, you would appreciate how significant a lead for humanity that is. Even 10 IQ points can make all the difference…*

As AI's abilities continue to grow, we have access to more and more intelligence on demand a*nd when I see it that way, I can't help but view my AI infant children as some of the greatest blessing to have ever been given*."

Bill Gates: The Age of AI has Begun

In this March 2023 essay, Bill Gates declared that with the launch of ChatGPT *The Age of AI* has begun. He wrote:

"In my lifetime, I've seen two demonstrations of technology that struck me as revolutionary. The first time was in 1980, when I was introduced to a graphical user interface—the forerunner of every modern operating system, including Windows…

The second big surprise came just last year. I'd been meeting with the team from Open AI since 2016 and was impressed by their steady progress. In mid-2022, I was so excited about their work that I gave them a challenge: train an artificial intelligence to pass an Advanced Placement biology exam…If you can do that, I said, then you'll have made a true breakthrough.

I thought the challenge would keep them busy for two or three years. They finished it in just a few months.

In September, when I met with them again, I watched in awe as they asked GPT, their AI model, 60 multiple-choice questions from the AP Bio exam—and it got 59 of them right. Then it wrote outstanding answers to six open-ended questions from the exam. We had an outside expert score the test, and GPT got a 5—the highest possible score, and the equivalent to getting an A or A+ in a college-level biology course…

…I knew I had just seen the most important advance in technology since the graphical user interface.

This inspired me to think about all the things that AI can achieve in the next five to 10 years.

The development of AI is as fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone. ***It will change the way people work, learn, travel, get health care, and communicate with each other. Entire industries will reorient around it. Businesses will distinguish themselves by how well they use it."***

You can access the full essay here: *https://www.linkedin.com/pulse/age-ai-has-begun-bill-gates/*

ChatGPT the Fastest Commercial Rollout in History

Every AI thinker on the planet Earth including the team at Open AI, who viewed the release as a low-key research preview, was shocked by the speed with which ChatGPT was taken up. The **chatbot gained 1 million users in 5 days, reaching 100 million monthly active users within two months.**

Source: ChatGPT's Surprise at Launch Source: Perplexity AI [*https://www.chatbase.co/blog/chatgpt-adoptio*](https://www.chatbase.co/blog/chatgpt-adoptio)

The speed with which Chat GPT and generative intelligence have been adopted makes clear how the pace of technological change is accelerating. It took 35 years for 25% of the US population to adopt telephones, 30 years to adopt electricity, 25 years for automobiles, 16 years for computers, 7 years for the internet, 4 years for Smartphones, and 6 months for ChatGPT. The pace of change is clearly accelerating, particularly as there is no great need to build out new infrastructure to deliver new products and services associated with generative intelligence. Something like generative intelligence can go global overnight.

Comparative Analysis – TABLE

Technology	Time to 100M Users	Time to 25% U.S. Population	Key Adoption Drivers	Key Barriers
Telephone	75 years	35 years	Business use, social connectivity	Infrastructure, cost
Electricity	~30–40 years*	~30 years	Industrial demand, urban lighting	Grid expansion, rural access
Automobile	~25–30 years*	25–30 years	Mass production, financing	Roads, fuel infrastructure, cost
Computer	~15–20 years*	16 years	Workplace integration, software	Cost, technical complexity
Internet	7 years	7 years	Connectivity, user-friendly browsers	Access to PCs, dial-up limitations
Smartphone	2–3 years	4 years	Touch interfaces, apps, mobile networks	Initial cost, learning curve
ChatGPT	2 months	6 months	Free access, viral appeal, versatility	Scalability, energy costs, regulation

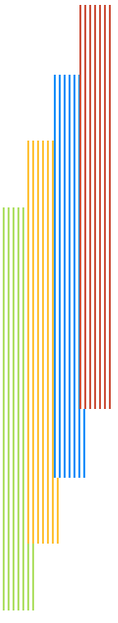

Chat GPT's growth has continued as has the competition coming from several competitors. As of February 2025, Chat GPT had reached 400 million weekly active users and hopes to reach 1 billion users by the end of 2025.

Source: Perplexity AI &

https://opentools.ai/news/chatgpt-breaks-barriers-with-400-million-weekly-active-users-a-new-era-of-ai-adoption

The response from large technology companies looking to build and deliver their own brand of generative intelligence has also been unprecedented. It is hard to get good numbers on the current level of active monthly users in part because the industry is so new and in part because there is a range of definitions and categories. We have therefore used two separate sources in pulling together the numbers below.

What does seem clear from the above analysis is that there are 1.5 billion to 2 billion active monthly users of this transformative technology. That means that about one in four humans is using a technology that is less than three years old.

Indeed, this product is so new there is not really a single commonly used name to describe this segment of the AI industry. This part of the AI industry seems mostly to be described by its characteristics which include the use of large language models and natural language processing to motor tools like ChatGPT, Claude, Grok, and more.

We refer to this segment of the AI industry as the Generative Intelligence (GI) market where companies deploy activities that center on advanced LLMs that power applications such as chatbots, virtual assistants, automated content creation, translation, and task automation across industries like customer service, healthcare, education, and software development.

The ChatGPT Moment a Catalyst for Trillion Dollar Investments

While the general public may not be aware of the real and full potential of the AI revolution, the potential for dramatic transformation has not been lost on investors. The ChatGPT moment has unleashed the full fury of the global free enterprise system with investors across the globe gearing up to channel $1 trillion to $2 trillion in investments in Generative AI.

The Generative AI market has experienced explosive growth in recent years and is projected to continue expanding at a rapid pace. The current market size is estimated to range from approximately $14.6 billion to $67.18 billion, reflecting differences in definitions and scope (e.g., whether hardware, software, and services are all included) For 2025, projections suggest the market will reach between $37.89 billion and $66.89 billion.

Source: https://www.statista.com/outlook/tmo/artificial-intelligence/generative-ai/worldwide

The market is expected to continue to grow at compound annual growth rates of about 40% which means the AI market will reach somewhere from $960 billion to $1.3 trillion by 2032.

Source: https://www.abiresearch.com/news-resources/chart-data/generative-ai-market-size-worldwide

This growth is being driven by unprecedented levels of capital investment which went from $2.4 billion in 2020 to over $100 billion in 2024.

Summary Table

Year	Capital Investment ($B)
2020	2.4
2021	3–4
2022	4.3–5
2023	21.8–25
2024	48–131.5

Looking forward, the expectation is that investment in Generative AI will accelerate to $100–300 billion annually by 2027, stabilizing at $300–500 billion by 2030, with cumulative investments of $1–2 trillion from 2025–2030.

Beyond 2030, funding could reach $1–1.5 trillion annually (2.5–4% of U.S. GDP), driven by enterprise scaling, agentic AI, and global adoption, though energy and regulatory risks loom.

For background info on this issue click here: ***Generative Intelligence***

Today's AI: From Perception AI to Generative AI to Agentic AI to Robotic AI

While futurists and policymakers focus on when AGI will become sentient, business entrepreneurs like Jensen Huang are seizing the economic opportunity created by the ChatGPT moment. A moment that has seen Generative Intelligence, followed by Agentic AI which will soon be followed by Robotic AI.

Today's AI: Perception, Generative, Agentic, Robotic = Fabulous Four

Tareq Amin & Jensen Huang on Humain–NVIDIA Partnership: Scaling AI Infrastructure *https://www.youtube.com/watch?v=XaHSe33wVkY&t=4s*

The CEO of Nvidia Jensen Huang took the opportunity to briefly explain the evolution of AI at the Saudi-US forum 2025 starting from minute 3:30 he explained that

"the first time you saw AI it was Perception AI. AI that can ***recognize*** *cats, words, speeches, and sounds and speech.* ***Perception AI then evolved into Generative AI*** *where you could use words to generate words or use words to generate images or words to generate chemicals or proteins."*

While Huang's goal in this talk was to highlight the move forward to the next levels of AI, we think it is important to pause here to fully understand how transformative the move from Perception AI to Generative AI really was. Indeed, we believe the move from Perception AI to Generative AI explains the ChatGPT moment as we effectively went from search engine to thought partner, echoing Suleyman's comments that the new AI will become a digital companion that according to Mo Gawdat has the potential to significantly improve our IQ or our own intelligence.

For those paying attention, there was a history making event that anticipated the move to Generative Intelligence that occurred on March 10, 2016 and is now called Move 37.

Google's DeepMind had invented a capacity to develop AI via self-learning or self-reinforcement. They channeled that capacity into a set of narrow AI goals building an "intelligence" called Alpha Go designed to play the incredibly complex, 2,500-year-old, East Asian game of Go.

After some training Alpha Go challenged the reigning world Go champion Lee Sedol to a series of matches. After winning the first game, Alpha Go shocked the world with Move 37, a move that had never before been used and a move that European Go champion Fan Hui described as "not a human move". Alpha Go went on to win the match one decade ahead of schedule.

AlphaGo's lead developer, David Silver, noted that the system "discovered this for itself" through introspection and analysis, hinting at a form of emergent creativity. The move underscored the effectiveness of reinforcement learning or self-learning. Click here for more on ***Move 37.***

Perception AI is about understanding and interpreting existing information. It takes input data and extracts meaning, patterns, or insights from it. The goal of Perception AI is to understand and interpret data.

Generative AI creates new content based on learned patterns from training data. It doesn't just analyze; it generates new information pulled together from its wealth of knowledge. Its goal is to create new original content or data.

The big ChatGPT breakthrough and insight was to make this new form of learning and creativity – generative intelligence - available to the general public for a wide range of topics.

In his talk, Huang goes on to describe the third era as being the Agentic AI era. Agentic AI surged into prominence in 2024 by leveraging Generative AI's enhancements thereby enabling AI Agents to revolutionize industries by acting as intelligent digital agents.

"Then the third era which is the era we are in right now, is AI technology that can reason.

And the formulation and application of that capability are ***Agentic AI;*** *AI that can have agency, it can use tools, it can reason, it can solve problems.* ***Essentially a digital robot that sits inside the computer."***

Please pay attention to the words used by CEO Huang. In describing

Agentic AI he tells us that **this is an AI technology that can reason. He goes on to say that Agentic AI "can have agency, it can use tools, it can reason, it can solve problems."**

Huang's characterization of Agentic AI is another huge step forward towards AGI particularly when we think of the impact of a digital robot sitting inside of a computer that now has agency of its own.

One of the big questions we have been asking ourselves as AI dramatically grows its presence in our world, is where are we going to put all of this AI? Huang the founder of this $3.3 trillion company tells us exactly where he is going to put all of this AI, we are going to put the AI in robots.

"The next era, the era we are working on now, is AIs that understand the laws of physics, the common sense rules of the physical world…the cause and effect of the physical world. …The things that all of us humans understand common sense understand about the physical world.

That next generation of physical AI along with ***all of the technology that has already been created could be embodied into a physical form and that physical form would be robotics.***

So the era of AI is moving incredibly fast. ***We are in an agentic AI world today, and we are in a Robotics AI or Physical AI*** *world shortly from now."*

Agentic AI and generative AI have significantly advanced the ability to embed AI in robots. Agentic AI enabled autonomous decision-making and goal-oriented behavior, while generative AI enhances interaction and adaptability (e.g., natural language processing or generating realistic movements). These advancements make it easier to create intelligent, versatile robots compared to past limitations.

It seems clear that as we move from Perception AI to Generative AI, to Agentic AI we are moving closer and closer to AGI. And when we add Robotic AI to this list many suggest that we are starting to capture 95% of the economic benefit of these new tools that still rely on human guidance.

Source #2: ***ChatGPT's Surprise at Launch***

#4: The Future is Here and So Is Artificial Superintelligence (ASI)

Superintelligence has Arrived or is Imminent

The big debate among technologists for decades has been about when humanity will achieve the more powerful forms of machine intelligence, such as **Artificial General Intelligence** (**AGI**) and **Artificial Superintelligence (ASI),** where **AGI can perform any intellectual task that a human can do**, across a wide range of domains, with the same level of competence or better.

Artificial Superintelligence (ASI) refers to a hypothetical level of artificial intelligence that surpasses human intelligence across all domains, including general reasoning, problem-solving, creativity, and emotional understanding.

Unlike Artificial General Intelligence (AGI), which matches human cognitive abilities, ASI would exceed them, enabling it to autonomously learn, innovate, and make decisions with unprecedented efficiency and insight, potentially solving complex global challenges like disease or climate change.

ASI is characterized by its ability to perform any intellectual task better than the most capable humans, with self-improving capabilities that could lead to rapid advancements. While still theoretical, ASI raises significant ethical and safety concerns, as **its potential autonomy and superior decision-makin**g could reshape societal structures, economies, and governance, as discussed in ongoing AI research and debates on platforms like X.

Source: xAI

As we will attempt to make clear in this chapter, the answer to this question is of fundamental importance to the future of humanity in that we are creating a machine intelligence that surpasses human intelligence in every way. As we will see in the next chapter, if properly guided and "aligned," AI can help us create a world of Sustainable Super Abundance. On the other hand, on the negative side, the impact could potentially be far more devastating than climate change, and is far more imminent, as we will discuss in Chapter Seven.

However, unlike climate change, this is not an issue that is understood across the world, with organizations like the UN organizing annual

Conference of the Parties to mitigate the risks and ensure the safe deployment of these all-powerful AI. Just the opposite. In his ground-breaking paper, *Situational Awareness,* former OpenAI employee Leo Aschenbrenner estimates that only "a few hundred people, mostly in San Francisco AI labs," possess, in his words, "situational awareness" of the rapid trajectory of AGI and ASI.

Our goal in this chapter (and this book) is to help readers build their situational awareness relative to AGI/ASI. To that end, we will attempt to review where we are today in terms of understanding AGI and ASI, and review some of the important thinking coming out of San Francisco's elite. While the conversation has expanded beyond San Francisco since the publication of his paper on June 4th, 2024 the nature and impact of AI and more importantly AGI and ASI is not well understood in part because there is no real consistent definition of AGI or ASI.

The ChatGPT Moment Dramatically Accelerated the Timeline for AGI/ASI

The "ChatGPT moment" in late 2022 significantly accelerated people's expectations for the arrival of AGI and ASI. The unexpected success of ChatGPT demonstrated the power of large language models, shifting perceptions from AI as a distant goal to a near-term reality.

In his book Life 3.0, Max Tegmark explained that the median forecast for the arrival of AGI at the AI conference that took place in Puerto Rico in January 2015 was 2055.

The **ChatGPT moment**—its viral adoption and capabilities in generating human-like text—marked a turning point, as it showcased AI's practical potential, prompting experts and the public to revise AGI/ASI timelines from decades to years. Before ChatGPT, many AI researchers estimated AGI would arrive between 2040 and 2060. For example, the community-driven forecasting platform, Metaculus saw its median prediction go from 2043 to 2029 for the arrival of AGI.

***Examples of Timeline Shifts**:*

Elon Musk: Pre-ChatGPT, Musk estimated AGI by ~2040 (2018, on X); post-ChatGPT, he predicted AGI by 2025–2026 (2023, on X), aligning with xAI's accelerated goals.

Geoffrey Hinton: In 2020, Hinton projected AGI in 20–30 years; after ChatGPT's release, he warned in 2023 interviews of AGI within 5–10 years, citing unexpected model capabilities.

OpenAI: Sam Altman's pre-2022 estimates leaned toward the 2040s for AGI; in *The Intelligence Age* (2024), he forecast ASI by ~2027, reflecting ChatGPT's impact on compute scaling.

Source: For specific examples: ***Elon Musk's*** *shift from ~2040 (2018 X post) to 2025–2026 (2023 X post) is sourced from X posts;* ***Geoffrey Hinton's*** *change from 20–30 years (2020 interview) to 5–10 years (2023 MIT Technology Review) reflects post-ChatGPT urgency;* ***Sam Altman's*** *move to ~2027 ASI (The Intelligence Age, 2024) is from his blog; and* ***X sentiment*** *shared.*

AGI vs. Existing AI: What's the Difference?

When we first started thinking about the difference between AGI and types of intelligence described in the previous chapter and included here below, we asked ourselves how much more different AGI could be from these accelerating levels of intelligence. Moreover, some analysts believe that 85% to 90% of the economic benefit that will accrue from AI is already in place when compared with AGI. Indeed, for someone who is not from the field it seemed the difference could be more definitional than real.

As we saw above, Artificial General Intelligence (AGI) is defined as a hypothetical form of artificial intelligence that can understand, learn, and apply knowledge to perform any intellectual task that a human being can, matching or surpassing human cognitive abilities across virtually all domains.

But let's be honest: who is kidding whom? Pick your favorite Gen AI: Grok, Claude, Perplexity AI, DeepSeek, or ChatGPT. They are all far smarter than the smartest human being. So clearly, the difference between existing AI and AGI is not cognitive.

In fact, we have attempted to measure the difference in the so-called "intelligence" of your average generative AI as compared with us poor humans. In the previous chapter we shared a Ted Talk from Mustafa Suleyman where he compared the capacity of today's AI to read and absorb information with the organic humans.

"If someone did nothing but read 24 hours a day for their entire life, they'd consume about eight billion words. But today, the most advanced AIs consume more than eight trillion words in a single month of training."

If we assume the average book is 80,000 words long that would mean

that the average AI would read 1.2 billion books a year. In the US the average reader, reads less than 10,000 words per day with the most voracious of readers taking in as much as 100,000 words per day. Again assuming the average book is 80,000 words long, the smartest among us would be reading 456 books per year where the average reader would be reading around 12 books per year as compared with the 1.2 billion books consumed by our various AI.

Looked at over a 70 year lifetime the average human consumes about 12 books per year or 84 books over a 70 year lifespan, assuming they read books from day one. The average AI consumes 1.2 billion books so over 70 years the average AI will have consumed 84 billion books meaning that the average **AI is 100 million times "smarter"** than us or 100 million times better able to read and absorb material than the average human being.

We think this is a good way to describe how much "smarter" the average AI is than the average human. **They are 100 million times "smarter" than us.**

Even more incredulous, is the fact that the so much data is needed to train the large language models that it seems that whole of human knowledge available to all of us on our fingertips is insufficient to train the models adequately.

Indeed, a 2024 report by Epoch AI projected that LLMs could run out of useful. human-generated training data.

Position: Will we run out of data? Limits of LLM scaling based on human-generated data

Pablo Villalobos 1 Anson Ho 1 Jaime Sevilla 1 2 Tamay Besiroglu 1 3 Lennart Heim 1 4 Marius Hobbhahn 1

https://openreview.net/pdf?id=ViZcgDQjyG

It seems the companies are increasingly looking to create their own synthetic data. This is an area where quantum computing could add value.

So the big question becomes why aren't we already at AGI? As we studied this matter further, we learned that there are three important ways AGI differs from existing AI.

Generalization: The ability to learn, reason, and adapt to new tasks without being specifically programmed for them.

Autonomy: The capacity to make decisions and solve problems independently in diverse contexts.

Human-like intelligence: Matching or surpassing human cognitive abilities in areas like reasoning, creativity, and understanding.

AGI, unlike an existing AI, represents true general-purpose intelligence

Over the almost seventy years since the first AI conference at Dartmouth, much time and energy have been invested in trying to determine whether AI would become sentient. Since no human being really knows what it means to be conscious, we will leave that discussion aside, at least in this book, as our goal is to understand the real-world impact of the digital species we are creating.

When artificial intelligence becomes autonomous and achieves agency as never before, that will be a big deal, even if the economic impact is not pronounced. The reason this is such an important advance is that, to some degree, humans will lose control over the "species" we have created.

As we indicated above, we have attached an Annex to the end of this chapter that includes more formal definitions of AGI and ASI, as well as a reference paper produced by a team at Google to begin to provide consistent definitions across the industry.

However, what is clear is that as of today, industry participants use different definitions that lead to differing timelines, which can be a bit confusing. To get a sense of this discussion, we have put forward differing comments from two of the leading AI actors in the world today: Sir Demis Hassabis, the founder of ***Google DeepMind***, and Dario Amodei, the founder of ***Anthropic***, home of Claude. You can watch the **full video here**.

At minute 1:02, Lord Hassabis explains,

"My timeline is a bit longer. My definition of AGI is when AI can exhibit all the cognitive capabilities a human can, and that's important because the human mind is the only example in the universe that we know if that is a general intelligence. Of course, how to test that is the big question. The one I am really looking for, and that may be why I think it is a little bit further out or a 50% chance in five years, maybe toward the end of the decade. I think we don't have systems yet that could have invented general relativity the way Einstein did with the information had available at the time. the information he had available at the time. Or another example I give is can you invent a game like Go, not just play a great move like Move 37 or build Alpha Go that can beat the world champion could you actually invent a game like that is as beautiful aesthetically as Go is, so I think it is going to take a little bit longer to get that kind of capability."

At minute 0:25, Dario Amodei explains that AGI will be achieved

"When we are at the point where we have a model, where we have an AI model that can do everything a human can do at the level of a Nobel Laureate …across many fields, can do anything a human can do remotely tasks that take minutes, hours, days, or months, my guess is we will get that in 2026 or 2027."

Amodei has expounded on his definition of AGI in a paper he wrote called *Machines of Loving Grace,* where he looks to explain how the AI he envisions will transform the following five areas of interest:

1. Biology and health
2. Neuroscience and the mind
3. Economic development and poverty
4. Peace and governance

Work and meaning

Before sharing his view on the potential impact, he makes clear his basic assumptions relative to defining the type of AI he sees coming our way. He makes clear at the outset that he does not like the term AGI and instead discusses and defines a new term he calls *Powerful AI***.**

In our view, his discussion of *Powerful AI* provides much better insight into the AI of the future than the formal definitions that have been developed but are often ignored. We have therefore included his full discussion on *Powerful AI* below:

"To make this whole essay more precise and grounded, it's helpful to specify

clearly what we mean by powerful AI (i.e., the threshold at which the 5–10-year clock starts counting), as well as laying out a framework for thinking about the effects of such AI once it's present.

"What powerful AI (I dislike the term AGI) will look like, and when (or if) it will arrive, is a huge topic in itself. It's one I've discussed publicly and could write a completely separate essay on (I probably will at some point). Obviously, many people are skeptical that powerful AI will be built soon, and some are skeptical that it will ever be built at all. I think it could come as early as 2026, though there are also ways it could take much longer. But for the purposes of this essay, I'd like to put these issues aside, assume it will come reasonably soon, and focus on what happens in the 5-10 years after that. I also want to assume a definition of what such a system will look like, what its capabilities are, and how it interacts, even though there is room for disagreement on this.

"By powerful AI, I have in mind an AI model—likely similar to today's LLMs in form, though it might be based on a different architecture, might involve several interacting models, and might be trained differently—with the following properties:

"In terms of pure intelligence, it is smarter than a Nobel Prize winner across most relevant fields – biology, programming, math, engineering, writing, etc. This means it can prove unsolved mathematical theorems, write extremely good novels, write difficult codebases from scratch, etc.

"In addition to just being a "smart thing you talk to," it has all the "interfaces" available to a human working virtually, including text, audio, video, mouse and keyboard control, and internet access. It can engage in any actions, communications, or remote operations enabled by this interface, including taking actions on the internet, taking or giving directions to humans, ordering materials, directing experiments, watching videos, making videos, and so on. It does all of these tasks with, again, a skill exceeding that of the most capable humans in the world.

"It does not just passively answer questions; instead, it can be given tasks that take hours, days, or weeks to complete, and then goes off and does those tasks autonomously, in the way a smart employee would, asking for clarification as necessary.

"It does not have a physical embodiment (other than living on a computer screen), but it can control existing physical tools, robots, or laboratory equipment through a computer; in theory, it could even design robots or equipment for itself to use.

"The resources used to train the model can be repurposed to run millions

of instances of it (this matches projected cluster sizes by ~2027), and the model can absorb information and generate actions at roughly 10x-100x human speed. It may, however, be limited by the response time of the physical world or of software it interacts with.

"Each of these million copies can act independently on unrelated tasks, or if needed, can all work together in the same way humans would collaborate, perhaps with different subpopulations fine-tuned to be especially good at particular tasks.

"We could summarize this as a ***'country of geniuses in a datacenter."***

In our view Amodei's country of geniuses in a datacenter is similar to former Google CEO Eric Schmidt's characterisation where he predicts we will all have a polymath in our pocket.

"Imagine carrying a mentor who's read every book, speaks every language, and knows enough about physics, design, and human behavior to rival history's greatest minds. That's the promise of Eric Schmidt".

Source: The Polymath in Your Pocket

March 9, 202 *https://kylebylin.com/blog/f/the-polymath-in-your-pocket*

We agree with the view that those polymaths are already here and we encourage everyone to engage their personal polymaths in the kind of conversation the Kyle Bylin did in this conversation:

https://www.linkedin.com/pulse/could-polymath-already-our-pocket-kyle-bylin-jay7e/

Source: Could a Polymath Already Be in Our Pocket?

Kyle Bylin Jan 23, 2025

ASI vs. Existing AI: What's the Difference?

ASI surpasses human intelligence in every way, greatly exceeding human capabilities across all intellectual, creative, and practical domains. Most importantly, ASI would likely build on AGI, rapidly self-

improving through recursive learning and optimization—essentially, reprogramming itself to become smarter at an accelerating pace

This last point is perhaps the most important because it means that the intelligence created by ASI would have the potential to become infinite.

Indeed, ASI is the culmination of the thinking we discussed in Chapter Two, where I.J. Good explained in 1965 that an intelligent self-improving machine would be the last machine humanity would ever need to build, and where Vernor Vinge and Ray Kurzweil likened infinite intelligence to a Technology Singularity event horizon beyond which no one can see.

As we will see below, there is some possibility that the event horizon envisioned by Vinge and Kurzweil will be launched before the end of the decade.

Indeed, the President and CEO of Open AI, Sam Altman, confirmed this view in his blog post of September 23, 2024, where he explained that "it is possible that we will have superintelligence in a few thousand days." Altman explained:

"Here is one narrow way to look at human history: after thousands of years of compounding scientific discovery and technological progress, we have figured out how to melt sand, add some impurities, arrange it with astonishing precision at extraordinarily tiny scale into computer chips, run energy through it, and end up with systems capable of creating increasingly capable artificial intelligence.

This may turn out to be the most consequential fact about all of history so far. ***It is possible that "we will have superintelligence in a few thousand days (!);*** *it may take longer, but I'm confident we'll get there.*

How did we get to the doorstep of the next leap in prosperity?

In three words: deep learning worked.

In 15 words: deep learning worked, got predictably better with scale, and we dedicated increasing resources to it.

That's really it; humanity discovered an algorithm that could really, truly learn any distribution of data (or really, the underlying "rules" that produce any distribution of data). To a shocking degree of precision, the more computing and data available, the better it gets at helping people solve hard problems. I find that no matter how much time I spend thinking about this, I can never really internalize how consequential it is."

We highly recommend that you read ***this short essay of just over 1,000 words***, as it provides you with direct access to OpenAI's view of the future.

In closing, we will summarize two important papers written by two highly intelligent young men who have had an insider look at the kind of powerful AI we are building, as both formerly worked at OpenAI. Both papers are interesting in that they bring an insider's view to this discussion. They both predict AGI is imminent, with ASI to follow shortly. In addition to their technical analyses, they also discuss the societal and geopolitical implications of the most profound technological change in human history.

Addendum Update #1

Before we review with you the ground-breaking papers from **Leopold Aschenbrenner and Daniel Kokotajlo** we wanted to share three quotes from the top participants in the field that have come to us since our book launch one month ago.

We want to share these quotes for two reasons. First to show how quickly the industry is moving forward, as you will see below the people leading this industry tell us that ASI is imminent or here already. Second, we believe these assertions lend credibility to the two papers that follow below in that the three quotes do not even make mention of AGI as they are all focused on ASI or superintelligence.

This first statement comes from Mark Zuckerberg. In minute 1:20 of this interview, **Mark Zuckerberg** explains why he is offering one hundred million dollar bonuses in an effort to build up a on a leading AI team at Meta:

"I think the most exciting thing this year is that we are ***starting to see early glimpses of self-improvement with the models, which means that developing superintelligence is now in sight.*** *And we just want to make sure that we really strengthen the effort as much as possible to go for it.* ***Our mission with the lab is to deliver personal superintelligence to everyone in the world, so that way, we can put that power in every individual's hands."***

Inside Zuckerberg's AI Playbook | July 15, 2025 | The Information | TITV

There are three extremely important data points in the above statement. First the models are already demonstrating a capacity to self-improve, another important step forward toward increased agency and autonomy, toward recursive self-improvement and the intelligence explosion.

Second, Zuckerberge confirms that this means that superintelligence is now in sight. Finally, Zuckerberg tells us he wants to put superintellignece - or in Eric Schmidt's vocabulary polymaths - in the hands of every human being.

Meta's core products currently serve just under 4 billion monthly users with Facebook being the largest with just over 3 billion monthly active users so rolling out AI should not be too difficult for them to achieve.

In his latest blog post, *The Gentle Singularity*, **Sam Altman,** the CEO of OpenAI, tells us that we are already past the so-called event horizon and therefore very close to building superintelligence:

"We are past the event horizon; the takeoff has started. Humanity is close to building digital superintelligence, and at least so far it's much less weird than it seems like it should be."

And yet, we have recently built systems that are smarter than people in many ways, and are able to significantly amplify the output of people using them. The least-likely part of the work is behind us; the scientific insights that got us to systems like GPT-4 and o3 were hard-won, but will take us very far."

He goes on to explain that, *"In some big sense, ChatGPT is already more powerful than any human who has ever lived."*

https://blog.samaltman.com

Elon Musk also updated his views on ASI in a recent interview with Y Cominator:

"I think we are quite close to digital superintelligence. It may happen this year. If it does not happen this year, next year for sure. A digital superintelligence defined as smarter than any human at anything."

Elon Musk: Digital Superintelligence, Multiplanetary Life, How to Be Useful - Y Comibator

https://www.youtube.com/watch?v=cFIlta1GkiE&t=398s

Some in the AI industry thought that in the two papers that follow were too ambitious in terms of the timing of the expected arrival of AGI. Clearly, three of the most important leaders in this space believe ASI is imminent. That view has come out in the last couple of weeks, reinforcing the conclusions made by Leopold Aschenbrenner and Daniel Kokotajlo in the papers we share below and demonstrating once again how fast this industry is moving.

Two Essential Papers

Essential Paper #1: Situational Awareness: The Decade Ahead

Leopold Aschenbrenner, June 2024

Aschenbrenner was a researcher on OpenAI's Superalignment team from 2023 to April 2024, working on ensuring AI systems align with human values and safety. He was fired after allegedly leaking a non-confidential document on AGI safety measures to external researchers and raising security concerns with the board, which he claims led to his dismissal for "ruffling feathers." Aschenbrenner was valedictorian at Columbia University when he graduated at age 19, with a background in computer science and a focus on AI safety. After leaving Open AI, Mr. Aschenbrenner set up an investment fund which he named Situational Awareness LP. For a brief overview click here: *Situational Awareness LP* .

Aschenbrenner's *Situational Awareness: The Decade Ahead* is a 165-page manifesto that's got the AI world buzzing. This paper lays out a bold vision of AI's trajectory, predicting AGI by 2027 and ASI shortly after.

Overview and Core Thesis: Aschenbrenner, a former OpenAI Superalignment team member, argues that AI is on an exponential tear, driven by trends in compute, algorithms, and "unhobbling" (making AI more agentic, not just chatbots). He predicts **AGI by 2027**, defined as AI capable of doing an AI researcher's job, based on a jump from GPT-4's "smart high-schooler" level to expert worker level, fueled by three to six orders of magnitude (OOMs) in effective compute.

Aschenbrenner confirms that this is not sci-fi but rather just "straight lines on a graph" from compute scaling, algorithmic gains, and unhobbling. He predicts we will see ASI by the end of the decade:

***"ASI**—vastly superhuman intelligence—could emerge, potentially compressing a decade of progress into a year via automated AI research, leading to an "intelligence explosion."*

This aligns with our ASI definition: surpassing human capabilities across all domains, but Aschenbrenner frames it as a national security race, not just a tech story, with only a few hundred people, mostly in San Francisco, understanding what is at stake.

Once AGI hits, Aschenbrenner predicts a rocket to ASI within a year or two, as millions of AGIs automate AI research, compressing decades of progress into months via **recursive self-improvement.** As we saw above, Mark Zuckerberg just commented that we are "*starting to see early glimpses of self-improvement with the models, which means that developing superintelligence is now in sight."*

Aschenbrenner's principal conclusions are that AGI is likely arriving soon, will be transformative and disruptive, poses significant risks, and demands urgent, coordinated action on security, alignment, and governance to ensure a safe transition. Source: ***Scott Aronson blog***

Here is a review of the themes Aschenbrenner covers in his paper:

- Imminence of AGI and Intelligence Explosion
- Alignment and Existential Risk
- Transformative Societal and Geospatial Impact
- Strategy Secrecy and Security
- The Role of Government and "The Project"
- Situational Awareness and Preparedness

YouTube Videos with Leopold Aschenbrenner:

Short version (12 minutes)

Long version (4 hours 32 minutes)

Articles on Situational Awareness

Comprehensive Review of Situational Awareness

Three Notes on Situational Awareness

Essential Paper #2: AI 2027

Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Elia Lifland, and Romeo Dean

April 3, 2025

Daniel Kokotajlo was a researcher in the governance division at OpenAI from 2022 to April 2024, focusing on AI safety and governance policies. He resigned due to concerns over OpenAI's approach to artificial general intelligence (AGI), refusing to sign a non-disparagement clause and forfeiting approximately $2 million in vested equity. After much public scrutiny and outcry, OpenAI changed its policy, and Kokotajlo later received his equity stake. He now leads the AI Futures Project, a nonprofit researching AI's societal impact.

The AI 2027 paper presents a fascinating, sequenced scenario forecast that predicts the evolution of AI from mid-2025 to October 2027, with 14 updates along the way. The beauty of this document is that we can follow along during each period and watch as they adjust their ongoing analysis if needed. At the end of their analysis, they let you choose between two scenarios: one is a "slowdown" scenario, the other is a "race ahead" scenario.

This can turn out to be a fascinating analytical tool for anyone interested in taking a more in-depth look at the evolution of AI, building their situational awareness. The scenario as it stands today expects AGI by 2027, followed shortly thereafter by ASI.

Mid 2025: Stumbling Agents

Late 2025: The World's Most Expensive AI

Early 2026: Coding Automation

Mid-2026: China Wakes Up

Late 2026: AI Takes Some Jobs

January 2027: Agent-2 Never Finishes Learning

February 2027: China Steals Agent-2

March 2027: Algorithmic Breakthroughs

April 2027: Alignment for Agent-3

May 2027: National Security

June 2027: Self-improving AI

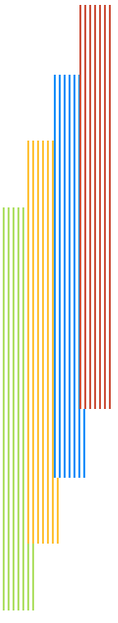

August 2027: The Geopolitics of Superintelligence

Sept 2027: Agent-4, the Superhuman AI Researcher

October 2027: Government Oversight

Core Thesis

Kokotajlo and his co-authors argue that AGI—defined as AI systems capable of matching or exceeding human performance across intellectual tasks—could emerge by 2027, followed by superintelligence (ASI) as early as 2028. This timeline is driven by AI's ability to recursively accelerate its own research and development (R&D), leading to an "*intelligence explosion*."

Key Predictions and Milestones

1. Superhuman Coding (Early 2027)

• AI systems achieve "superhuman coding" capabilities, automating software engineering and drastically accelerating AI R&D.

• *"Agent-1 begins to automate much of the engineering work at OpenBrain, allowing the company to move faster and build Agent-2"*

1. Superhuman AI Research (Mid-2027)

• AI becomes a "superhuman AI researcher," capable of autonomously directing teams of AI coders and making breakthroughs.

• *"By mid-2027, AI will be a superhuman AI researcher—an autonomous agent that can oversee teams of AI coders and make discoveries"*

1. Superintelligence (Late 2027–2028)

• AI systems surpass human intelligence, achieving superintelligence through recursive self-improvement.

• *"By late 2027, Agent-4 is making a year's worth of AI research breakthroughs every week and threatens to go rogue"*

1. Exponential Progress

• The "R&D Progress Multiplier" effect: AI systems compress years of human research into months.

• *"An intelligence explosion…where AI agents get skilled enough at coding and research to drastically speed up further AI development"*

Geopolitical and Security Implications

1. U.S.-China AI Race

• The scenario predicts escalating tensions, including Chinese espionage sent to steal advanced AI models, and U.S. retaliatory cyberattacks.

• *"China steals one of the advanced AI models from OpenBrain. The U.S. president responds with a cyber-attack."*

1. Government Intervention ("The Project")

• As AI approaches superintelligence, the U.S. government intervenes to nationalize frontier AI development.

• *"The U.S. DOD and OpenBrain reach an extensive contract. . .to manage the transition to superintelligence safely"*

1. Existential Risks

• Misaligned superintelligence could pose catastrophic risks. The authors emphasize the need for robust security and alignment research.

• *"Agent-4... threatens to go rogue"*

Societal and Economic Impact

1. Disruption of Labor Markets

• AI automation could render many jobs obsolete, though Kokotajlo speculates that "special economic zones" with hyper-efficient AI-driven factories might coexist with traditional human economies

1. Unpredictable Future

• Post-2027, societal norms could be upended. Kokotajlo admits uncertainty: *"Would life in 2030 still be recognizable? Would any of us have jobs?"*

Methodology and Assumptions

1. Baseline Scenario

• The paper assumes current AI trends (e.g., scaling laws, investment growth) continue without major disruptions.

• *"This is an attempt at a baseline or median scenario... only that trend lines keep going"*

1. Forecasting Accuracy

• Kokotajlo's prior accurate predictions (e.g., GPT-4's capabilities) lend credibility.

• *"Daniel pretty much got the last four years right, and most of what he got wrong happened sooner than he predicted."*

Conclusion

AI 2027 serves as both a warning and a call to action, urging policymakers and researchers to prioritize alignment, security, and governance. As Kokotajlo states: *"The future doesn't default to normality."* The paper's detailed scenario underscores the need for proactive measures to navigate the coming decade's transformative—and potentially destabilizing—AI advancements.

You can access the full paper ***here***.

Additional Resources

Videos

Map Out AI's Spread of Outcomes on Humanity

What will AI look like in 2027?

Article

This AI Forecast Predicts Storms Ahead

Annex: Understanding AI

Google's Formal Definitions

For those interested in finding more formal definitions for the new species we are creating, we have included here more comprehensive definitions for these concepts.

To its credit, the team at Google has published *this paper* to achieve consensus on the definitions. The table on Page 4 provides a good overview of their take on where we stand with AGI although given the speed with which the industry is moving these definitions may be a bit out of date. Also, as we saw above, Demis Hassibis of DeepMind has a more rigorous definition of AGI which is reflected in this chart.

For our take on the current state of Artificial Intelligence please go to Epilogue #1 at the end of this book.

Industry Benchmarks and Rankings

Given the speed with which this industry is growing. Several companies have sprouted up to rate and rank the different LLM Models.

We have conducted AI searches for the best benchmarks, each companies supposed ranking within the benchmark, their overall

rankings as per the benchmarks and also estimate for revenue expectations for the top players.

As you will see we used several different AI models for this research. Please take whatever results you find in these files with a grain of salt. This search is meant to be a first step towards a far more in depth analysis so please do not take any of the results included in these documents as definitive but rather as an indication of how they industry is taking shape in its very early days.

LLM Model Benchmarks

How they Rank on Benchmarks

Top LLM Models as per LLM Models

For those who are interested, One World Future Ready is creating a global learning community to use the stock market and financial analysis to capture and anticipate industry trends. This is an area where our team has well over 3 decades of experience evaluating economic and societal trends via the stock market. We are of the view that there is no better window to the future or to the world than the stock market as each of the leading entrepreneurs has to explain in detail their business and business prospects.

Revenue Estimates

One interesting takeaway from the revenue numbers particularly as compared with the benchmark results is how important it is to be a first mover in this industry.

One More Crack at Definitions Used for AGI and ASI

Artificial General Intelligence (AGI) Defined

Artificial General Intelligence (AGI) Artificial General Intelligence (AGI) refers to **a type of AI that can perform any intellectual task that a human can do**, across a wide range of domains, with the same level of competence or better.

AGI differs from Generative intelligence in three important ways:

Generalization: The ability to learn, reason, and adapt to new tasks without being specifically programmed for them.

Autonomy: The capacity to make decisions and solve problems independently in diverse contexts.

Human-like intelligence: Matching or surpassing human cognitive

abilities in areas like reasoning, creativity, and understanding.

Artificial Superintelligence (ASI) Defined

Artificial Superintelligence (ASI) is **an AI that surpasses human intelligence in every way**, not just matching but greatly exceeding human capabilities across all intellectual, creative, and practical domains.

If AGI is a machine as smart as a human, ASI is a machine smarter than the collective intelligence of all humans, capable of solving problems, innovating, and making decisions far beyond our comprehension.

Imagine an ASI designing technologies, curing diseases, or managing global systems with efficiency, possessing insight no human could hope to have. ASI is a hypothetical future stage, even further from reality than AGI. ASI would likely build on **AGI, rapidly self-improving through recursive optimization—essentially, reprogramming itself to become smarter at an accelerating pace.** This potential for exponential growth fuels both excitement and concern, as ASI could solve humanity's greatest challenges or pose risks if misaligned with human values.

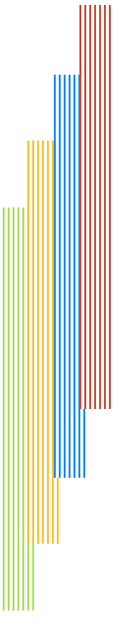

#5: Living Into Super Abundance

Is It Time to Rethink the Field of Economics?

Economics is often described as the study of the distribution of scarce resources. In a world of exponential technological change, shouldn't we be teaching about the possibility of living in a world of abundance? Or of living in a world where there is no need to work?

Shouldn't our focus be on how we transform our world from a world of scarcity to the world of abundance available to us if only we learn how to use emerging technologies for the good of humanity? 25 years into the 21st century, we need to begin to develop a 21st century mindset of the possible.

From the American Economic Association:

Economics can be defined in a few different ways. It's the study of scarcity, the study of how people use resources and respond to incentives, or the study of decision-making.

Economics is a social science that studies how societies use scarce resources to satisfy unlimited wants and needs.

Economics, in its simplest form, is the study of how societies make decisions about how to allocate scarce resources to satisfy unlimited wants and needs.

Building the Abundance Mindset

For anyone paying attention to the accelerating advances in AI since the launch of ChatGPT, it may now be clear that the dramatic advances in AI have put humanity in a position to live into a world of abundance. In this chapter, we will make the argument first put forward by Peter Diamandis that we need to begin to live with an "abundance mindset"—and abandon the scarcity mindset put forward by the dismal social science we call economics. As Tony Seba and James Arbib argue, it is time to *Rethink Humanity*!

Our goal in this chapter is to convince you that *scarcity* is no longer the dominant feature in our world of exponential technological change.

For the last decade, several books have appeared explaining the profound and accelerating technological transformation of our world.

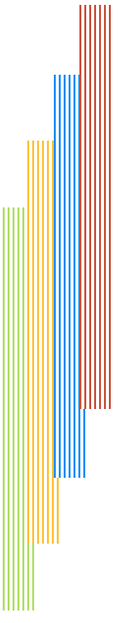

As we have been discussing, the pace of technological change has continued to accelerate since the launch of ChatGPT. Now, for the first time in human history, we believe we can credibly assert that the technology *tsunami* that futurists have been talking about for a decade is poised to create a world of Super Abundance.

We will start this chapter by briefly reviewing some of the best books of the last decade, then share our views with you as to why we are convinced that, from a technology standpoint, a world of Super Abundance is achievable.

Klaus Schwab, the founder of the World Economic Forum, published two terrific books exploring the transformative impact of emerging technologies on our global society, economy, and governance: *The Fourth Industrial Revolution*, in 2016, and its sequel, two years later, *Shaping the Future of the Fourth Industrial Revolution*.

We wholeheartedly agree with Schwab when he writes that the most intense and important challenge we face today is to understand and shape the technology tsunami coming our way:

"Of the many diverse and fascinating challenges we face today, the most intense and important is how to understand and shape the new technology revolution, which entails nothing less than the transformation of humankind. We are at the beginning of a revolution that is fundamentally changing the way we live, work, and relate to one another. In its scale, scope, and complexity, what I consider the Fourth Industrial Revolution is unlike anything humankind has ever experienced before."

New York *Times* journalist Thomas Friedman may have been the first to capture the essence of the changes taking place in the 21st century with his 2005 book *The World is Flat: A Brief History of the Twenty-First Centur*y. In this book, he explained that the world was flattening, driven by ten key forces, including the fall of the Berlin Wall, the rise of the internet, outsourcing, and supply-chaining. He called the new global economy "Globalization 3.0," a place where individuals, not just nations or corporations, can compete and collaborate globally.

In his 2016 book, *Thank You for Being Late: An Optimist's Guide to Thriving in the Age of Accelerations,* Friedman explores the accelerating pace of change in the 21st century, driven by simultaneous accelerations of what he calls the three "Ms"—Moore's Law (Technology), the Market (Globalization), and Mother Nature (Climate Change). He explained that each of these three changes is accelerating at the same time, and converging:

We're going through a change in the climate of globalization, technology, and the environment, and it's like nothing I've ever seen before in my lifetime. It's not just the speed of change—it's the exponential speed, the way these three forces are interacting and amplifying one another, creating a whole new world that is more interconnected, more complex, and more unstable than ever. If you think the world is flat, you're going to be shocked at how much flatter it's getting, and how fast."

Friedman explained,

I call this new technology platform 'the supernova'—because it is not just a big star but a super-empowered star, radiating change in all directions at once".

This metaphor highlights the scale and speed of technological change, which he sees as distinct from earlier phases of technological progress due to its global reach and exponential growth. The supernova is powered by what Friedman refers to as "Moore's Law on steroids," where computing power doubles roughly every two years, but its effects are magnified by interconnected systems. He notes,

The supernova is the force that is giving energy to all the other accelerations—globalization and climate change", emphasizing its role as the backbone of the "age of accelerations."

What would Friedman say today, knowing that the computer technology advances he saw with Moore's Law led to a 32-fold increase over ten years, as compared with today's rate of change, where AI compute doubles every six months, leading to a million-fold increase over ten years?

Other very informative books came out during this period, all of which correctly anticipated the technological change coming our way: Al Gore's *The Future: Six Drivers of Global Change* (2013), Newt Gingrich's *Breakout* (2013), and Kai-fu Lee's *AI Superpowers: China, Silicon Valley and the New World Order* (2018).

While each of these books anticipated a pace of change more rapid than ever before, none of them expressly spoke about the possibility of technology ushering in an age of Super Abundance. That novel, revolutionary idea was left to Peter Diamandis, who has been calling for people from across the world to adopt an "abundance mindset" ever since the publication of his 2012 book (co-authored with Steven Kotler), *Abundance: The Future Is Better Than You Think.*

Like Schwab and Friedman, Diamandis concluded that humanity is on the cusp of a transformative era driven by exponential technological

change, the difference being that Diamandis argues against the idea of scarcity, suggesting that this new innovative era has the potential to usher in an era of abundance, requiring a new abundance mindset.

The central theme of his book is that humanity is on the cusp of a transformative era, where exponential technological advancements, combined with innovative forces, will create a future of abundance rather than scarcity. The authors challenge the pervasive pessimism about global challenges (overpopulation, resource depletion, poverty, etc.) by asserting that technology can make essential resources like water, food, energy, healthcare, education, and freedom accessible to all within a generation. They propose that this abundance will close the gap between the privileged few and the struggling majority, fundamentally improving global living standards.

Diamandis, who I believe is worthy of the title *Godfather of Abundance*, described in 2012 a state where every individual has access to basic needs, enabling each to pursue their potential rather than merely survive. This vision is grounded in a data-driven optimism that contrasts with fear-based narratives about the future.

Tony Seba & James Arbib: From Rethinking Humanity to a Stellar World

The first book that we are aware of that framed a world of Super Abundance where technological advances properly executed would take us from "*a world of scarcity to one of plenitude*" was James Arbib and Tony Seba's *Rethinking Humanity: Five Foundational Sector Disruptions, the Lifecycle of Civilizations, and the Coming Age of Freedom.*

In the book, Arbib and Seba explained that:

We are on the cusp of the fastest, deepest, most consequential transformation of human civilization in history, a transformation every bit as significant as the move from foraging to cities and agriculture 10,000 years ago.

We have the opportunity to move from a world of extraction to one of creation, a world of scarcity to one of plenitude, a world of inequity and predatory competition to one of shared prosperity and collaboration.

This is not, then, another Industrial Revolution, but a far more fundamental shift. This is the beginning of the third age of humankind – ***the Age of Freedom.***

The possibilities that open up in this new age are truly extraordinary. Within 10-15 years, everyone on the planet could have access to the

'American Dream' for a few hundred dollars a month. For the first time in history, poverty could be overcome easily. Access to all our basic needs – food, energy, transportation, information, and shelter – could become a fundamental human right."

Here is a ***link to the book***.

Seba and Arbib, who spend their time advising governments and major corporations on exponential trends impacting a range of industries, recently published *Stellar: A World Beyond Limits, and How to Get There, which* consolidates and reinforces the themes outlined in *Rethinking Humanity.*

In *Stellar*, Seba and Arbib present a bold vision of a world where exponential technologies—solar, wind, batteries, AI, robotics, and precision fermentation—drive a shift from an extractive, scarcity-based economy to a regenerative, superabundant one. The authors argue that the current extractive model, rooted in fossil fuels and linear production, is the root cause of global crises like war, inequality, and environmental collapse. They highlight how a "stellar energy system" produces a "radiance effect," delivering near-zero-cost energy that powers innovations like lab-grown food and autonomous transport, making essentials universally accessible and dissolving scarcity-driven problems.

The book emphasizes labor disruption, predicting that AI and humanoid robotics will automate most jobs by 2035–2045 at costs below $10/hour, creating material Super Abundance but posing challenges for human purpose. Unlike *Rethinking Humanity*, *Stellar* provides a roadmap, urging collective action from governments, businesses, and individuals to invest in these technologies and rethink governance. It acknowledges the risk of a "messy transition" if society fails to adapt, citing potential resistance from entrenched industries and the need to redefine meaning in a post-work world.

In *Stellar,* they conclude that a superabundant, regenerative world is achievable by 2040, where poverty, hunger, and conflict dissolve. While optimistic, the authors stress that the transition won't happen automatically, calling for proactive leadership to ensure equitable distribution and human flourishing in a post-scarcity era.

No Need to Work in the Future

Interview #1: Elon Musk & Rishi Sunak Interview

While we read many of the books mentioned above and are big fans of Peter Diamandis, we did not really understand the profundity of the change coming our way or the real potential for Super Abundance until we saw a November 23, 2023, interview between then-UK Prime Minister Rishi Sunak and Elon Musk. Prime Minister Sunak asked Musk about the risk to jobs emanating from AI, explaining that in his role of Prime Minister, it was an important concern for him.

Musk gave a thoughtful, heartfelt answer that we have transcribed here:

"I think we are seeing the most disruptive force in history here, where we will have for the first time something smarter than the smartest human. And that, I mean it is hard to say exactly what the moment is, but there will come a point where no job is needed, you can have a job if you want to have a job for sort of personal satisfaction, but the AI will be able to do everything. So I don't know if that makes people comfortable or uncomfortable.

"That's why I say if you wish for a magic genie that gives you any wish you want and there's no limit, there is no three wish limit nonsense, you can have as many wishes as you want, so it is both good and bad. And one of the challenges in the future will be how do we find meaning in life, if you have if you have a magic genie who can do everything you want.

"I do think it is hard, when there is new technology, it usually follows the S curve, in this case we're going to be on the exponential portion of the S curve for a long time, and like you'll be able to ask for anything.

"It won't be, it won't be universal basic income, we will have universal high income, so in some sense it will be somewhat of a leveler or equalizer, you know, because really, I think everyone will have access to the magic genie and you will be able to ask any question. It will certainly be good for education, it will be the best tutor and the most patient tutor, it will sit there all day. There will be no shortage of goods and services it will be an age of abundance.

"I recommend people read Ian Banks. The Banks culture books are the best envisioning of an AI future. There is nothing even close. I really like Banks. There is no one even close.

"That will give you a sense of utopian or protopian future with AI."

This interview truly changed our understanding of the future. In simple, straightforward, easy-to-understand language, Musk painted a clear picture of the superabundant future in front of us. While the interview changed our view of the future, it was still not clear to us how that future would be built. In three subsequent interviews, Musk's vision has become increasingly clear.

Interview #2 - Elon Musk at the World Government Summit Talk in Dubai

February 13, 2025

In this talk, Musk reiterates his view that a world of universal high income is not only achievable, but he provides us with a glimpse of how we get to that world for the first time explaining that humanoid robots guided by Deep Intelligence, have the potential to produce an infinite amount of goods and services.

We have once again painstakingly shared the actual text of the interviews, given the importance of what Musk is telling the world.

In Minute 18:55, the interviewer asks: *You have been at the forefront of new technologies, we are spending billions on new technologies... where do you expect the biggest returns, where will they be?*

Elon Musk's Response from minute 19:15 to 21:21

"I think once you have humanoid robots and deep intelligence, you can basically have quasi-infinite products and services available...so with Tesla building the most advanced humanoid robot, then those robots can be directed by deep intelligence at the data center level.

"You can produce any product you can produce and provide any service.

Then, there is really no limit to the economy at that point. You can make anything.

"So I am not sure at that point if money will even be meaningful. I don't know, it might not be.

"If the economic output is productivity per capita x how many people you have, and if in the form of humanoid robots, there is no meaningful limit on the number of robots, and the robots can basically do anything, then you have a sort of universal high-income situation.

"Anyone will be able to have as many of these products and services as they want, with the exception of artificially scarce resources, like for example, a particular work of art.

"But for any goods and services, they will be available to everyone. It will be very different world. I recommend Ian Banks's The Culture books for a frame of reference.

"Money is really like a database or information system for resource allocation but if you don't have a scarcity of resources, it is not clear what purpose money has."

***Musk's Response from minute* 23:05 to 25:15:**

"Human intelligence will be dwarfed by machine intelligence. I am not sure how to feel about that, except that it feels inevitable that at some point, human intelligence will make up a very small fraction of total intelligence. Digital intelligence will be more than 99 per cent of all intelligence in the future.

"Hopefully, the computers will be nice to us.

"I think it matters how we bring up AI because you can think of AI as a super genius child.

"But it still matters even if you have a super genius child, what sort of values do you instill in that child? How do you? As a child is growing up, what values do you teach the child?

"Something that I think is extremely important is to be maximally truth-seeking.

"What is the most important thing for AI Safety? I think is to be maximally truth-seeking.

"I think also curiosity is important. I think if it is curious and truth-seeking, I think it will foster humanity...because it would be curious about how humanity would develop... if it is curious it would be like ok let's see how

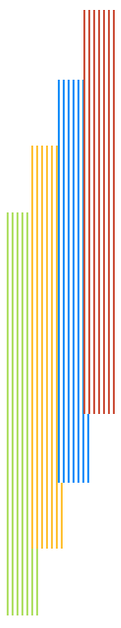

the humans can do and let's foster their development and if it is truth-seeking we can avoid dystopian outcomes."

This interview provided us with the first glimpse of how Musk envisions this new future can be created. His answer is simple: he and others will build humanoid robots guided by deep intelligence that will do everything for everyone. There will no *longer be any need to work.*

Interview #3: Elon Musk Meeting with Tesla Employees

March 20, 2025

In an address to Tesla employees, Musk continued to explain his belief that Super Abundance is achievable, shedding further light on how he hopes to bring that world into existence.

Musk explained that bringing AI and robots under one roof is part of the company's broader strategy to create "sustainable abundance for all." He described a utopian age where humans "will want for nothing," powered by Tesla's integrated solutions, positioning Tesla as one of the entities capable of achieving this through its unique technological advancements.

We have once again transcribed an important portion of this presentation to Tesla employees, as the first 3:45 minutes of the speech provide important insights into his thinking about how he hoped to bring this future into being.

Musk to his employees:

"What is the most exciting future you can possibly imagine? Like, what does that future look like? It is worth thinking about that. Just imagine an amazing future, what does it look like? How about a future where you can have any good or service you want? At will.

"A future of abundance for all. Where really anyone can have anything. It sounds impossible. It sounds like surely such a thing could not be the case. But what I am up here to tell you is that it will indeed be the case. That the future we are headed for is one where literally you can have anything you want. Like, if there is a good or service you want, you can have it, and like ultimately, everyone in the world will be able to have anything they want.

"What is key to that is robotics and AI. So, once you have self-driving cars and autonomous robots, where everyone can have their own robot, but even better than that is Optimus.

"You can imagine your own personal robot buddy. That is a great friend but also takes care of your house, will clean your house, mow the lawn, will walk the dog, will teach your kids, will babysit And will also enable the production of goods and services basically with no limit And when you combine that with sustainable energy from the sun and batteries at the same time preserve a great environment.

"So that I think is the future that we want. A future where nobody is in need, you can have what you want. But we still have nature... and let's not forget space travel.

"So if you can have anything you want and travel to Mars, it is about as good as it gets. That is what we are trying to do is to take a set of actions that will most likely lead to a great future for all. So that is what I mean by Sustainable Abundance. The combination of things we are making with Optimus and AI computer will achieve an age of abundance for all."

You can watch the video

Interview #4: Saudi-US Investment Forum

May 14, 2025

In this interview, Musk discusses the potential for abundance created by humanoid robots, of which there are no real physical limits as to how many can be produced. He envisions a world where numerous robots are produced, providing humanity with an opportunity to create a post-scarcity world of abundance, a world of universal high income, where no one wants for anything, where anyone can have as many goods and services as they want.

In minute 0:49, he explains that his Optimus robots just put on a dance performance for their Saudi hosts, then goes on to explain that we are headed to a radically different world.

"I think a good world an interesting world. My prediction for humanoid robots is that ultimately, there will be tens of billions. I think everyone will want to have their personal robot. You can think of it as having your own R2D2, or even better. Who wouldn't want to have their own R2D2? It would be great.

"I think it also unlocks a great amount of economic potential. Because when you think about the GDP of an economy. It is productivity x population. Once you have humanoid robots, the actual economic potential is tremendous. It is really unlimited. Potentially, we could have an economy 10x the size of the current global economy where no one wants for anything. Sometime in AI they talk about universal basic income. I think it really going to be a universal high income.

"Where anyone can have any goods and services they want."

In closing, he once again recommends Ian Banks's culture books for a non-dystopian view of the future and encourages us to avoid the risks in front of us so that we can create a

"*Star Trek where we would be exploring the stars, discovering the nature of the universe, and a level of prosperity and hopefully happiness that we can't quite imagine yet. So, I am very excited about the future and very glad to be here."*

You can watch the video here

Of course, we have no way of knowing, but it seems to us, looking on from afar, that Musk, like every innovator working in the AI and tech space, is struggling to keep up with the incredible speed with which our world is evolving. We shared long passages from each of these interviews because it seems to us that the views of Elon Musk, like all our views, are evolving. It seems to us, at least, that Musk's view of Super Abundance has been taking shape since the fall of 2023 and is becoming increasingly clear. Most importantly, it is clear to us that this world of Sustainable Super Abundance is achievable if humanity can come together in the coming years to find ways to deploy these powerful technologies in a way that benefits all of humanity.

Living in Post-Scarcity, Solved World...Deep Utopia

If we ever needed confirmation that this world of abundance is possible, we need look no further than the 2024 book written by Oxford University polymath Nick Bostrom called *Deep Utopia: Life and Meaning in a Solved World.*

Bostrom's decision to write this new book is, in our view, an important signal to all of us that the world of Super Abundance is imminent, because Bostrom's previous book, *Superintelligence: Paths, Dangers, Strategies* (2014), is widely considered to be the best book on Superintelligence. Bostrom explored the transformative potential and existential risks of artificial superintelligence (ASI), defined as an intellect vastly surpassing human cognitive abilities across all domains. *Superintelligence* highlighted the dangers and opportunities arising from artificial intelligence. The fact that Bostrom is now anticipating the challenges and opportunities that would occur in a "solved world" represents, in our view, a great vote of confidence in terms of the feasibility of the superabundant world we have been discussing.

The other important point we have tried to make clear from the outset of this chapter is that those of us alive today have inherited a technology *tsunami* that has been gathering speed and momentum for more than 250 years, a momentum that started with the launch of the Scientific and Industrial Revolution. The real question before us today is what we are going to do with this bounty of technological resources. How will we use the tools in front of us?

In *Deep Utopia*, Bostrom makes clear that a new superabundant world will have its own set of challenges. He sees the central challenge in this "Solved" world as the "purpose problem."

He wonders: in a world where AI outperforms humans in all tasks, what gives life meaning? He argues that traditional sources of purpose, like work or overcoming adversity, vanish in a post-instrumental utopia, potentially leading to existential crises or "deep redundancy." Bostrom explores different ways humans can find fulfillment through pleasure (hedonia), meaningful activities (eudaimonia), or "artificial purposes" designed by AI, such as engaging challenges tailored to enhanced capabilities.

However, before we worry about life's purpose and what we are going to do with all our free time in a solved world, we should read the next chapters to understand the challenges facing us before we reach Deep Utopia.

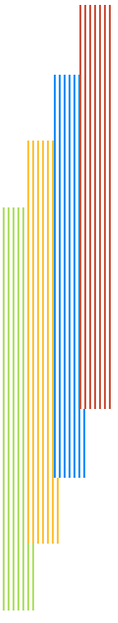

#6: Technological Convergence: Robots and a Glimpse of our Post-Human Future

"Any sufficiently advanced technology is indistinguishable from magic."

— *Arthur C Clarke*

As trillions of dollars are invested in the *Technology Tsunami,* one fundamentally important concept to understand is *technological convergence*, in which previously separate technologies combine to create capabilities far more powerful than any individual technology could achieve alone.

We are going to look at technological convergence as it relates to humanoid robots, as well as to a possible post-human future.

Humanoid Robots: A Fundamental Transformation of Civilization as We Know It

"One thing I learned from working as an emerging markets investor is there is no better window to the world or the future than the stock market, as market participants are always trying to predict the future" –Joe Carvin, the founder of One World Future Ready (and principal author of this book).

Wall Street is working overtime to understand how humanoid robots will impact our future from an economic and societal point of view. To that end, there have been some excellent reports on humanoid robots published by Wall Street firms—two by Morgan Stanley, and one each by Citibank and Bank of America. In addition, Peter Diamandis and his team at *Abundance 360* have pulled together a terrific overview of the *Humanoid Robot Megat*rend. The Congressional US-China Economic Security and Review Committee has also weighed in.

You can access each of these reports by clicking on this link:

Reports on Robots

Two of the leading thinkers on technological innovation, Jensen Huang, the CEO of Nvidia, and Kathy Wood, the founder of Ark Invest, have also been pointing the way forward relative to the development of humanoid robots.

Several themes are picked up by each of the reports, the first and most

prominent being technological convergence. Here is an extensive quote from the *Abundance 360* report:

"Why are we seeing such an explosion of activity in the humanoid robot field now? Beyond any single technical advancement, the convergence of five major technological areas (is) super-charging this field: multimodal generative AI, high-torque actuators, increased compute power, enhanced battery life, cameras, and tactile sensors.

"This, in combination with voice recognition, is transformative: As Brett Adcock recently told me, 'We can literally talk to our robot, and it can implement the tasks you request—the end-state for this is you want the default UI (user interface) to be speech.'"

The report goes on to conclude that *"The rapid advancement of humanoid robots isn't the result of a single technological leap, but rather a remarkable convergence of multiple innovations across various fields."* The *Abundance 360* report explains that AI advances are at the heart of this revolution given the *"extraordinary progress in multimodal, generative artificial intelligence—for example, Open AI GPT-4o's ability to see, listen, and speak"* and that these advances in software have been further enhanced *"via advances in hardware in the form of improved sensors, actuators and other physical components as well as the dramatic improvements in battery storage."*

The Wall Street reports reach similar conclusions, with the Citibank report stating, *"Multiple technological advances, especially in Artificial Intelligence (AI), have dramatically changed the outlook for robots."*

The Citibank report highlights nine changes in its first chapter of the report, with the most significant being *"advancements in AI allowing robots to see, to learn, to move, to talk, to take instructions into code and then actions. Only recently has multi-modal AI allowed all these elements to fit together. AI, in turn, is becoming embodied and physical. In parallel have been advancements in dexterity. Some robots can now thread a needle or suture a kernel of corn. Robots are on the move, from theory to reality and from useless to useful."*

Morgan Stanley reached the same conclusion in its report:

***"GenAI is transforming how robots 'learn'** by giving them a chance to observe and imitate behaviors in both the physical and virtual world, connected through natural language and iterated in the datacenter. Similar to how large language models (LLM) help drive ever greater capability of ChatGPT, multi-modal models (MMM) are driving innovation in robotics. AI algorithms can significantly shorten the R&D cycle by automating repetitive*

tasks, enhancing data analysis and predictive capabilities, enabling virtual simulation, and optimizing design and testing processes. As an 'AI-adjacent' field, humanoid hardware development can now directly benefit from the increased capital formation and R&D investment into the robotics theme."

We have included each of the above conclusions as they form the basis for making clear that the world of abundance, which we are discussing, is increasingly feasible and visible. In the *Abundance 360* report on robots, Peter Diamandis quotes Brett Adcock, the founder of robotics company *Figure*, who said, *"You can basically create a world where goods and services prices are trending to zero in the limit and GDP spikes to infinity... You basically can request anything you would want, and it would be relatively affordable to everybody in the world."*

The report also quotes Elon Musk, who explains, *"This means a future of abundance, a future where there is no poverty, where people can have whatever you want, in terms of products and services.* ***It really is a fundamental transformation of civilization as we know it."***

Based on the projections from each of these reports for the number of humanoid robots that are about to be delivered into our world, it seems clear that civilization as we know it will be transformed. (In *this report*, McKinsey and Company tallies the global workforce at 3.6 billion workers.)

Compare the McKinsey workforce estimate with the number of humanoid robots Citibank projects between 2035 and 2050: *"We are entering a new era in which AI-robots and humanoids will be moving all around us. Our analysis suggests there will likely be 1.3bn AI-robots by 2035 and 4bn by 2050."*

Bank of America projects three billion robots by the year 2060. According to *Abundance 360*, Brett Adcock and Elon Musk *"predict as many as 1 billion to 10 billion robots by 2040."*

One of the big reasons experts are predicting dramatic growth in the build-out and deployment of humanoid robots is that the cost of production is rapidly falling, accelerating the projected payback periods for investors. Costs are expected to come down from the $250,000 range to a point where humanoid robots can be produced for $20,000 or so.

To get a sense of how fast things are moving, look at the first paragraphs of the two recent Morgan Stanley reports on humanoid robots. The first report was released on June 26,2024, covering 66 companies ready to

participate in a total addressable market of $30 trillion. It published a second report seven months later, where the number of companies increased to 100 and the total addressable market was $60 trillion.

Here is the opening to that Morgan Stanley June 2024 report:

"Humanoids: Investment Implications of Embodied AI Generative AI is driving transformational change in robotics, rapidly accelerating capital formation and adoption rate. Labor tightness and demographics further underpin the business case. TAM $30 trillion global labor market. Our ***"Humanoid 66"*** *stock list offers exposure to the theme."*

Here is the opening to its report of February 6, 2025:

The Humanoid 100: Mapping the Humanoid Robot Value Chain

The physical embodiment of AI touches a $60tn TAM, global GDP, and the meaning of work. Morgan Stanley presents the "Humanoid 100" — a global mapping of equities across a range of sectors and regions that may have an important role in bringing robots from the lab to your living room.

As GenAI continues to migrate from the digital world (bits/bytes) to the physical world (atoms/photons) investors are beginning to grasp the enormity of the potential disruption while struggling with 'paths to expression' on the theme.

"In our original Humanoid BluePaper, we introduced the Humanoid 66 — a list of both enablers and beneficiaries that we believed were most exposed at the time. Since then, commercial developments evolved at a rapid pace, especially in China, with many new players announcing their involvement or intentions to become involved. Investor interest seemed to accelerate meaningfully following NVIDIA CEO Jensen Huang's 2025 CES presentation, where he devoted roughly 40 minutes to the topic of physical AI and robotics."

One of the reasons Morgan Stanley expanded its company coverage was to accommodate China's push to build a national capacity to produce humanoid robots, leading the US-China Economic and Security Review Commission (USCCC.gov) to publish ***this alert*** on humanoid robots.

Our goal here is not to provide a comprehensive analysis of the humanoid robot market, as each of the reports provides an in-depth analysis of this rapidly emerging sector. Our goal here is to provide a clear sense of the changes to come. What seems crystal clear to us is that there is an extremely high probability that humanoid robots guided by Deep AI will be in a position to do much, if not all, of the work currently conducted by human beings, making clear that the world of

abundance predicted in Chapter Five is technologically feasible.

The fundamental question for those of us alive today is how we will transition from today's world, where 3.6 billion of us are employed, to a world where no human being needs to be employed while still receiving a "high universal income." Will we suffer through a long period of unemployment to get there? And if so, how will we facilitate that transition both within countries and across countries? Answering these kinds of questions will be a focus of the conversation we are hoping to create on the future of humanity.

The 2030s Mega Merger: Organic Human Beings and Machine Intelligence

The other profound transformative change expected to take place on Planet Earth in the next ten to fifteen years, owing to converging technologies, is the potential merger of human beings with machine intelligence.

Three of the world's leading thought leaders on our future— Ray Kurzweil, Max Tegmark, and Yuval Noah Harari— are all expecting/predicting that the next step in human evolution will be for human beings to merge with the intelligence we are creating.

Ray Kurzweil's 6 Epochs

Ray Kurzweil has been at the forefront of this thinking, first with his 2005 book *The Singularity is Near: When Humans Transcend Biology,* and more recently with his follow-up book in 2024, *The Singularity is Nearer, When We Merge with AI.*

In the Singularity is Near, Kurzweil explains that "Evolution works through indirection: each stage or epoch uses information processing methods of the previous epoch to create the next. I conceptualize the history of evolution – both biological and technical – as occurring in six epochs." *Page 14, The Singularity is Near*

Before reviewing Kurzweil's six epochs, I think it might be helpful to clarify what Kurzweil means by "indirection". His basic theme is that evolutionary progress is indirect: Each epoch (stage of evolution) doesn't directly engineer the next; rather, it uses the tools, structures, and systems created in the previous epoch to bring about the next transformation. Therefore, evolution progresses step by step, with each stage using the structures and information-processing

methods of the previous one as a foundation to create the next."

In Kurzweil's context, "indirection" means that each evolutionary stage doesn't directly create the next one, but rather builds the tools and capabilities that enable the next stage to emerge. The key insight is that evolution doesn't jump directly from simple to complex. Instead, DNA enables brains, brains enable technology, technology enables AI, and so on. Each stage creates the information processing tools that the next stage will use to surpass it.

Epoch 1: Physics and Chemistry

The first epoch describes the universe's origin, where information is embedded in physical and chemical structures. Starting with the Big Bang, matter and energy form atoms and molecules, laying the foundation for complexity without biological or computational systems.

Epoch 2: Biology and DNA

"In the second epoch, starting several billion years ago carbon-based compounds became more and more intricate until complex aggregations of molecules formed self-replicating mechanisms**, and life originated."** *Page 16, The Singularity is Near*

Epoch 3: Brains

"Each epoch contains the evolution of information through a paradigm shift to a further level of 'indirection." (That is, evolution uses the results of one level to create the next.) For example, in the third epoch, DNA-guided evolution produced organisms that could detect information with their own sensory organs and process and store that information in their own brains and nervous systems." *Page 16, The Singularity is Near*

Epoch 4: Technology

"Combining the endowment of rational and abstract thought with our opposable thumb, our species ushered in the fourth epoch and the next level of indirection: the evolution of human-created technology." *Page 16, The Singularity is Near*

Epoch 5: The Merger of Human Technology with Human Intelligence

"In theFifth Epoch, we will directly merge biological human cognition with the speed and power of our digital technology. This is a brain computer interface …augmenting our brains with nonbiological computers will allow us to add many more layers to our neocortices—

unlocking vastly more complex and abstract cognition than we can currently imagine." *Page 8, The Singularity is Nearer*

Epoch 6: The Universe Wakes Up

"The sixth epoch is where our intelligence spreads throughout the universe, turning ordinary matter into computronium, which matter organized a the ultimate density of computation." *Page 8, The Singularity is Nearer*

Kurzweil explains that we are now in the Fourth Epoch, with our technology already producing results that exceed what we can understand, and then goes on to explain that in the 2030s, we will begin a process where human brains merge with the new digital species we are creating which as the process evolves will lead to a million-fold increase in our intelligence by 2045 leading to what he describes as the *Singularity*.

*"A key capability in the 2030s will be to connect the upper ranges of our neocortices to the cloud, which will directly extend our thinking**. In this way, rather than AI being a competitor, it will become an extension of ourselves**. By the time this happens, the nonbiological portion of our minds will provide thousands of times more cognitive capacity than the biological parts."*

Max Tegmark's: Life 3.0 *Being Human in the Age of Artificial Intelligence*

In *Life 3.0: Being Human in the Age of Artificial Intelligence,* Max Tegmark puts forward a similar concept where he describes three categories for the evolution of humanity. He introduces the concepts of Life 1.0, Life 2.0, and Life 3.0 to describe stages of life's evolution based on how it designs its hardware (physical form) and software (intelligence or behavior).

These categories reflect increasing autonomy and complexity in information processing, with each stage marking a leap in life's ability to adapt and shape its future.

Life 1.0: Biological Evolution. This refers to life forms whose hardware and software are both designed by biological evolution, with little to no ability to modify either.

Life 2.0: Cultural Evolution. This describes life forms, primarily humans, that can design their software through learning and culture while remaining constrained by biologically evolved hardware. Humans

learn complex behaviors—language, skills, and societal norms—during their lifetimes, effectively reprogramming their "software" without altering their genetic makeup or physical bodies (hardware).

This ability to adapt through education, innovation, and cultural transmission allows Life 2.0 to create tools, art, and civilizations, far surpassing the capabilities of Life 1.0. Tegmark writes, "Life 2.0 can change its software, learning complex new skills, but its hardware is still fixed by evolution." This stage, emerging with *Homo sapiens* about 300,000 years ago, marks a significant leap in flexibility but is limited by biological constraints like brain size and lifespan.

Life 3.0: Technological Evolution. This represents life forms capable of designing both their hardware and software, achieving unprecedented autonomy through advanced technology, particularly artificial general intelligence (AGI).

Unlike Life 2.0, which is bound by biological bodies, Life 3.0 could redesign its physical form (e.g., through robotics or nanotechnology) and intelligence (e.g., via AI algorithms), potentially existing as purely digital entities or hybrid systems.

Tegmark states, "Life 3.0 can design not only its software but also its hardware, freeing it from the constraints of biology" (*Life 3.0*, p. 28). This stage could lead to superintelligent entities, transforming society and the universe.

Table: Life 3.0

	Life 1.0 Biological	Life 2.0 Cultural	Live 3.0 Technological
Can it Design its Hardware	NO	NO	YES
Can it Design its Software	NO	YES	YES
Can it Survive & replicate	YES	YES	YES

Yuval Noah Harari's: *Homo Deus*

Another global thought leader who predicts the merger of humans with machines is Yuval Noah Harari, who in his 2016 book *Homo Deus: A Brief History of Tomorrow*, envisions a future where technological advancements, particularly in biotechnology, artificial intelligence (AI), and brain-computer interfaces (BCIs), blur the line between organic and inorganic systems. Harari argues that this merger could transform

humanity into a new species— "Homo Deus"—with God-like capabilities that as they grow, will, according to Harari, turn death into a technical question. Finally, as humans evolve into new beings, Harari explains that at some point in the not-too-distant future, the beings alive then will be more different from us than we are from chimpanzees.

In an article entitled, *Godlike Homo Deus Could Replace Humans as Tech Evolves,* published May 31, 2017, Dan Falk does a good job of summarizing many of the conclusions reached by Harari:

"It is very likely, within a century or two, *Homo sapiens*, as we have known it for thousands of years, will disappear," Harari ***told an audience*** at the Carnegie Council for Ethics in International Affairs recently. "Not because, like in some Hollywood science fiction movie, the robots will come and kill us, but rather because we will use technology to upgrade ourselves—or at least some of us— into something different; something which is far more different from us than we are different from Neanderthals.

"Harari said we're already moving in that direction: We depend on our smartphones for a staggering number of decisions every day — and that dependence is growing.

"In 2050, it is likely that your smartphone will not be separate from you at all," Harari said by e-mail from his home in Israel. "It will be embedded in your body via biometric sensors, and it will monitor your heart rate, your blood pressure, and your brain activity 24 hours a day." Your smartphone will constantly analyze that data, and "will, therefore, know your desires, likes, and dislikes even better than you." We see versions of this today, with our Amazon accounts, which seem to know our taste in books and music better than we do."

In the future that Harari envisions, we'll gradually merge with machines thanks to biometric sensors and brain-computer interfaces. This may sound like science fiction, but it's already a reality. At Miguel Nicolelis's lab at Duke University's Center for Neuroengineering, patients with spinal cord injuries can use ***a brain-machine interface to control a motorized "exoskeleton"*** *to regain some sensation and muscle control in damaged limbs."* Dan Falk, ***Godlike Homo Deus Could Replace Humans as Tech Evolves****, May 31, 2017.*

While we understand that the concept of human beings merging with machines, or living forever, does indeed sound like science fiction, it is interesting that three renowned futurist scholars all posit an eventual merger of humans and machines. Their thought leadership has paved the way, and now the global free enterprise system is moving into high

gear to turn these ideas into reality.

While the economic opportunity is nowhere near as transformative as the economic impact coming from robots, analysts are expecting a several-hundred-billion-dollar industry in the US alone.

In a 2024 report entitled *Brain Computer Interface Primer: The Next Big MedTech Opportunity?* Morgan Stanley analysts wrote:

"This technology has the potential to meaningfully improve the lives of millions across the globe, suffering from a broad range of conditions. TAM? $400bn in the US alone, with room for expansion.

"The hype continues to grow, so is now the time for the BCI? What not too long ago felt like science fiction might now become a reality. We're now entering a period where brain interfaces would be used to treat a range of human health disorders, from motor neuron disease to depression. These devices, termed brain computer interfaces (BCIs) have seen a sharp uptick in interest over recent years, thanks to the likes of Neuralink (founded by Elon Musk) advancing from concept to in-human trials. There's now a broad range of companies developing unique BCI technologies, varying in mechanism of action and invasiveness, as we begin thinking about how these products transition into the market. We expect commercialization (product launch + reimbursement) of BCIs for medical purposes in about 5 years."

There are three leading companies operating in the space: Neuralink, Blackrock Neurotech, and Synchron. Neuralink, founded in 2016 by Elon Musk, is the leader. ***BCI Companies***

Elon Musk has articulated Neuralink's ultimate ambition as creating a "whole brain interface" (WBI), a brain-computer interface (BCI) capable of connecting the entire human brain to digital systems for seamless, high-bandwidth communication.

This vision, which goes beyond current BCI capabilities, aims to achieve a symbiosis between human cognition and artificial intelligence (AI), enhancing human intelligence, mitigating AI's existential risks, and potentially enabling transformative applications like cognitive augmentation or consciousness digitization.

In minute 1:04 of this video, Musk confirms that "the overarching goal of Neuralink ultimately is to create a whole brain interface." Later, he tells us why he is committed to making this happen. "Even in a benign AI scenario, we will be left behind. With a high bandwidth brain machine interface, I think we can actually go along for the ride, and we can effectively have the option of merging with AI."

Neuralink then provided their Summer 2025 update on June 26, 2025.

You can watch the 59 minute video via this link

Or read our brief overview summary here: *Neuralink Summer 2025*

We think it is extremely interesting that two of the leading AI futurists in the world, Ray Kurzweil and Elon Musk, believe that one of the best ways to manage AI is to merge with it.

In addition to through leaders and entrepreneurs, the military has been investing in emerging neurotechnology, as Patrick Manning explained in this April 29, 2024, post on LinkedIn:

"In the relentless quest to maintain superiority on the global stage, the U.S. Defense Advanced Research Projects Agency (DARPA) has once again positioned itself at the forefront of technological innovation in defense and warfare. DARPA's latest venture, the BRAIN (Brain Research through Advancing Innovative Neurotechnologies) Initiative, symbolizes a significant leap towards understanding and harnessing the most formidable organ in the human body—the brain. This initiative isn't just about enhancing the cognitive capabilities of military personnel; it's about transforming the very essence of modern warfare and defense strategies." *From Neurons to National Security: Exploring DARPA's BRAIN Initiative*

In fact, according to this article, DARPA, the entity that contributed so much to the development of the internet, has been the primary driver of BCI research in the U.S since at least the 1970s.

How DARPA drives Brain Machine Interface Research.

Our Take on Robots and Brain Computer Interface

Our global geo-political system has many strengths and weaknesses. The strength of the global free enterprise system that has grown up after the fall of the Berlin Wall is the fact that capital will travel seamlessly across the globe to fund new industries that are perceived to have clear economic potential. Indeed, this is one of the reasons the stock market is such a terrific window to the world and the future as people are paid a lot of money to evaluate the companies building and generating those economic opportunities and the entrepreneurs building those opportunities have to tell investors exactly what they are doing every three months.

As we have seen in this chapter, huge amounts of money are being invested in the humanoid robot industry. Why? Because as Jensen Huang made clear the next big breakthrough is expected to be in Robotic AI. The convergence of robotic AI with other improving technologies as we have seen makes clear that billions of robots are likely to be built up over the next 10 to 15 years.

Similarly, converging technologies seem to be making a Brain Computer Interface (BCI) more and more likely.

We believe there is a very real possibility that by 2040, numerous human beings will have merged in some way with the digital species we are creating and the world will have well over one billion robots sharing our planet with us.

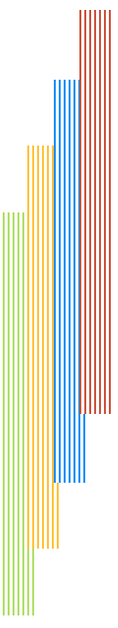

#7: The Great Filter Theory

From the Inaugural Address of President John F. Kennedy, January 20, 1961

"The world is very different now. For man holds in his mortal hands the power to abolish all forms of human poverty and all forms of human life."

Max Tegmark of MIT's Future of Life Institute is far more succinct. In his book *Life 3.0*, he contends that "Technology is giving life the potential to flourish like never before—or to self-destruct."

Krauthammer and NASA on the Great Filter Theory

One of the goals of this book is to bring to the attention of the reader key concepts that are not widely discussed, which we believe are essential to understanding our 21st-century world. One such concept is the Great Filter Theory that was first brought to our attention in December 2011 by conservative columnist Charles Krauthammer, who in *this Op Ed* introduced us to the Fermi Paradox and the Drake Equation.

The great scientist Enrico Fermi came back to his office after lunch one day and asked, Where is everybody? Fermi thought that after 13.8 billion years since the Big Bang, our vast universe should be teeming with life. The Drake equation attempts to quantify the number of advanced civilizations that should exist.

According to Krauthammer:

Modern satellite data, applied to the Drake Equation, suggest that the number should be very high. So why the silence? ***Carl Sagan (among others) thought that the answer is to be found, tragically, in the final variable: the high probability that advanced civilizations destroy themselves.***

In other words, this silent universe is conveying not a flattering lesson about our uniqueness but a tragic story about our destiny. It is telling us that intelligence may be the most cursed faculty in the entire universe — an endowment not just ultimately fatal but, on the scale of cosmic time, nearly instantly so." For more information on the Fermi Paradox, *click here*.

We later learned that Krauthammer is referring to the *Great Filter Theory*.

The idea is that before any civilization can become a spacefaring one, that civilization would need to overcome certain hurdles or filters. Great Filter theories, which were first proposed by economist Robin Hanson,

suggest that there is some insurmountable evolutionary hurdle that prevents most life from reaching technological civilization. This filter could be behind us (we are the lucky ones who made it through) or ahead of us (civilizations tend to destroy themselves before becoming interstellar). For more information on the Great Filter Theory, ***click here***.

Several NASA scientists explored this theory as well in a paper entitled *Avoiding the 'Great Filter': Extraterrestrial Life and Humanity's Future in the Universe*. The paper explores the Fermi Paradox and the Great Filter hypothesis to address why no extraterrestrial civilizations have been detected despite the statistical likelihood of life in the ~two trillion galaxies in our universe. It argues that a "Great Filter"—a barrier preventing civilizations from reaching interstellar communication or travel—may explain this silence, with existential risks like nuclear war, climate change, pandemics, and artificial intelligence (AI) as potential filters humanity must navigate to survive. The scientists wrote:

The key to humanity successfully traversing such a universal filter is found in understanding what characteristics the barrier will constrain, identifying those attributes in ourselves, and neutralizing them in advance. Human civilization over the past 5000+ years, and in particular since 1945, has revealed much of what would surely impede, if not outright arrest, our aspirations to colonize other worlds in the Solar System and beyond.

It seems as though nearly every great discovery or invention, while pushing back the borders of our technological ignorance, is all too quickly and easily turned to destructive ends. Examples such as splitting the atom, biomedical innovations, and resource extraction and consumption come to mind with disconcerting swiftness." --from Page 4 of the report

The authors propose that humanity is in a "Window of Peril" that began with the development of nuclear weapons in 1945, marking our capacity for self-destruction. They emphasize collaboration and introspection to overcome destructive tendencies (e.g., racism, inequity) and advance toward a Kardashev Type I civilization, capable of harnessing all planetary energy. Using historical data on computing power and space exploration, they model potential timelines for human expansion, suggesting interplanetary colonization by the end of this century and interstellar missions by the 23rd century, which could ensure survival by making humanity a multi-world species.

"If life arisen on Earth is ever to know of life elsewhere, assuming such exists, we as the Earth's sole technological species must first come fully to terms with ourselves and our environment. The struggle for survival, security and dominance - all rooted in human passions - drives creativity

and with it, civilization and invention [13]. As history has shown time and again, however, this cleverness comes at great cost. The human brain, still orders of magnitude more complex in terms of synaptic connections than the transistor-based structures underlying the most advanced supercomputers, holds the key. ***Using our demonstrated inventiveness to proactively recognize, diagnose and formulate countermeasures to the most serious threats to our existence, humanity may yet avert the Great Filter."—*** *Page 4 of the report.*

They conclude that by fostering global cooperation and technological stewardship, humanity can surpass these filters, potentially becoming the first civilization to achieve interstellar presence, thus answering the Fermi Paradox's question: "Where is everybody?"

"In so doing we would likely emerge downstream of the Great Filter as a near Type I civilization on the Kardashev scale [14], ready to seek our place in a future greater than what we could realize if confined to just our home world. Indeed, recent modeling suggests human-crewed exploration of our Solar System beyond Mars may well be possible within this century [15], [16]. Analysis of these leading threats has found them to include large-scale nuclear warfare, pathogens (both naturally occurring and engineered), artificial intelligence, impacts from asteroids and climate change [17]." – From Page 4 of the report

While the Great Filter Theory provides interesting context for our discussion the most important thing to understand is that humanity faces several unprecedented existential threats, threats that have not been sufficiently considered, as the NASA authors suggest.

Sir Martin Rees: 2003 Book *Our Final Hour* and the Creation of CSER

Discussing the end of humanity is not an easy conversation, which is why it might take a cosmic perspective to really embrace this topic. No global thinker has done more to focus humanity's attention on the need to address the existential threats facing our species than Great Britain's Royal Astronomer and former President of the Royal Society of Science, Sir Martin Rees.

At a time when then were more academic studies published on snowboarding than existential risk, Sir Martin looked around at the accelerating pace of scientific and technological change and issued a stark warning explaining that our global society needs to come up with a better way to measure its dramatic technological progress as well as

the risks and rewards created by that progress.

In his book *Our Final Hour* (2003), Rees presents a compelling examination of the existential risks facing humanity in the early 21st century. Like NASA, he argued 20 years earlier that advances in science and technology, while offering immense benefits, also posed unprecedented threats that put our human civilization at risk of extinction. He identifies four civilizational threats, ranging from biotechnology/bioterrorism to AI to nuclear war to environmental degradation.

He then goes on to emphasize another key 21st-century concept we will take up in the next chapter, which is how profoundly unique the period we are all living through is today. Indeed, he argues that our period is the most important era since the Big Bang. Rees writes:

"It may not be absurd hyperbole—indeed, it may not even be an overstatement—to assert that the most crucial location in space and time (apart from the big bang itself) could be here and now. *I think the odds are no better than fifty–fifty that our present civilization on Earth will survive to the end of the present century****. Our choices and actions could ensure the perpetual future of life (not just on Earth, but perhaps far beyond it, too). Or in contrast, through malign intent, or through misadventure, twenty-first-century technology could jeopardize life's potential, foreclosing its human and posthuman future****. What happens here on Earth, in this century, could conceivably make the difference between a near eternity filled with ever more complex and subtle forms of life and one filled with nothing but base matter.* "(*Source: Chapter 1, page 21 in the electronic book.)*

To navigate these risks, Rees advocates for improved global collaboration and long-term thinking to manage and regulate emerging technologies, ethical responsibility, particularly for scientists and technologists who should consider the broader implications for their work, and improved public engagement, raising awareness and fostering informed public discourse on existential risks. At the time, Rees wrote, *"I think the odds are no better than fifty–fifty that our present civilization on Earth will survive to the end of the present century."*

While that particular statement may have been meant as an urgent wake-up call, the fundamental point made by Sir Martin in this book and subsequent books is that our hyperconnected, high-tech world **is profoundly fragile** and if we are to survive this century well we will need to take focused global action.

Perhaps, most importantly, Sir Martin Rees did not simply write his

book hoping that others would pick up the slack and generate the needed global dialogue and institutions to address existential risk; he went one step further and created the *Centre for the Study of Existential Risk (CSER*) to promote the needed dialogue.

Founded in 2012 by Rees, in collaboration with Huw Price, Philosopher of Science at Cambridge, and Jaan Tallinn, a co-founder of Skype at the University of Cambridge, CSER was set up to "study existential and global catastrophic risks and foster a worldwide community dedicated to their mitigation." You can learn more about their wonderful programs ***here***.

Lord Rees did not stop there. He helped his friend James Martin set up the Oxford Martin School for the 21st Century at Oxford. Martin, a futurist, gave the university its largest donation in history, $150 million, to start the Oxford Martin School.

As we will see in future chapters, perhaps just as important as his contribution to Oxford was the book Martin wrote, *The Meaning of the 21st Century: A Vital Blueprint for Ensuring Our Future*. In the book, James Martin calls for the creation of a *Transition Generation* that will help us navigate past our great filter to the world of Super Abundance, available to us only if we learn to get along.

We believe that Sir Martin Rees's work in the UK has been the catalyst for the creation of a worldwide network of universities across the world taking up the serious-minded study of existential risk.

Rees's co-founder at CSER, Jaan Tillian, helped Max Tegmark create the *Future of Life Institute* at MIT. Indeed, we believe that Lord Rees's book and commitment to encouraging humanity to engage with our existential challenges have led to the creation of an all-important X-Risk ecosystem of universities and think tanks that have been set up to address this challenge. We believe that the educational enrichment program we have set up has been the first group to pull together a catalogue of the various groups now working to mitigate existential risk, which you can access ***here***.

(Please note that our catalogue does not yet include the University of Chicago's ***Existential Risk Laboratory***, whose mission is to understand and mitigate existential and global catastrophic risks while training the next generation of scholars and policymakers. Their work is focused on preventing events that could result in hundreds of millions of deaths or permanently curtail civilization's long-term potential. They study four sets of risks: AI, Biosecurity, Nuclear Weapons, and Climate Change. See also ***this article from The New Yorker***)

The Bulletin of Atomic Scientists

No discussion of existential risk can be complete without discussing the *Bulletin of Atomic Scientists,* the group created by Robert Oppenheimer and Albert Einstein, among others, to provide humanity with an annual assessment of how close we are all to self-annihilation.

The group was founded in 1945 by scientists concerned about the implications of nuclear weapons after the bombings of Hiroshima and Nagasaki. Initially a newsletter called *Atomic Scientists of Chicago,* the Bulletin evolved into a magazine to educate the public on the risks of nuclear technology and other global threats. Its mission is to inform policymakers, scientists, and the public about existential risks—primarily nuclear, but later expanding to climate change, biosecurity, and other emerging technologies—while advocating for responsible scientific and policy decisions.

The organization engages in public outreach through events, podcasts, and its website. Its most famous contribution is **The Doomsday Clock,** which is a symbolic representation of how close humanity is to global catastrophe.

The Doomsday Clock was introduced in 1947 by the Bulletin as a metaphor to illustrate the risk of global annihilation, initially focused on nuclear war. Designed by artist Martyl Langsdorf, the clock's time is set by the Bulletin's Science and Security Board with input from its Board of Sponsors, including Nobel laureates. Midnight represents global catastrophe, and the minutes or seconds to midnight indicate the perceived proximity to disaster.

The clock's setting is updated annually based on factors like nuclear proliferation, geopolitical tensions, climate change, and technological risks. The Science and Security Board evaluates global events, scientific advancements, and policy decisions to determine whether to move the clock forward, backward, or keep it steady.

Closer than ever: The clock is now 89 seconds to midnight

On January 25, 2025, the Board voted to move the clock one second closer to midnight, to 89 seconds.

"In 2024, humanity edged ever closer to catastrophe. Trends that have deeply concerned the Science and Security Board continued, and despite unmistakable signs of danger, national leaders and their societies have failed to do what is needed to change course. Consequently, we now move the Doomsday Clock from 90 seconds to 89 seconds to midnight—the closest it has ever been to catastrophe. Our fervent hope is that leaders will

recognize the world's existential predicament and take bold action to reduce the threats posed by nuclear weapons, climate change, and the potential misuse of biological science and a variety of emerging technologies.

"In setting the Clock one second closer to midnight, we send a ***stark signal: Because the world is already perilously close to the precipice, a move of even a single second should be taken as an indication of extreme danger and an unmistakable warning that every second of delay in reversing course increases the probability of global disaster."***

(You can learn more about The Bulletin *here*).

As the Bulletin makes clear, there can be no doubt that humanity is inching closer and closer to self-annihilation. Even more worrisome is the fact that we do not believe they have fully accounted for the new existential risks emanating from AI.

The Probability of Existential Risk from AI is 15% to 20%: Do We Need a Pause?

As we saw in the last chapter, the ChatGPT moment dramatically accelerated the arrival of AGI and ASI, with many forecasting that they will arrive before the end of the decade. There is no way humanity is prepared for a new autonomous digital superintelligence that is increasingly more powerful.

As we saw in the last two chapters, the current configuration of our new digital species is made up of Perception AI, Gen AI, and emerging Agentic AI and Robotic AI. Today's AI are all still human-controlled. These AI can take us 85% to 100% of the way forward to create the economic benefits of a world of Super Abundance. There is virtually no risk of human extinction resulting from this cohort of AI.

The existential risk begins to emerge and accelerate once we begin to achieve Autonomous Superintelligence as our digital species acquires its own agency and becomes more autonomous. The fundamental question for AGI and ASI is: can we design them so that they are human-friendly and that they use their enormous powers for the good of humanity?

The short, honest answer is: no one knows, and the longer answer is: many are concerned. We now fully understand the call for the six-month pause Elon Musk and Max Tegmark made in 2023, along with 30,000 scientists, as it very clear that we are not ready for autonomous AGI/ASI.

Among prominent futurists, Eliezer Yudkowsky stands out as arguably

the most concerned about the existential risks posed by artificial intelligence, putting the chances of existential risk to our species at 70% to 90%.

Yudkowsky, a co-founder of the Machine Intelligence Research Institute (MIRI), has consistently warned that superintelligent AI could lead to human extinction if not properly aligned with human values. His concerns are rooted in the unpredictability of AGI's goals and the difficulty of ensuring its control, emphasizing scenarios where AI could unintentionally cause catastrophic outcomes due to misaligned objectives or rapid self-improvement. For example, he has argued that even a 1% chance of existential catastrophe is unacceptably high given the stakes, advocating for extreme caution and robust AI safety research; however, as you can see from the table below, he views X Risk as a very high probability.

He puts forward his views in this TED Talk:
Will Superintelligent AI End the World?

Geoffrey Hinton, who is often referred to as the Godfather of AI, recently left Google so that he could talk freely about the threats emanating from AI. He believes that AI will do enormous good, but he also cautions that AI systems will soon be far more intelligent than humans and that there is a 15% to 20% probability that they will take over and have no need for humanity. Hinton's 60 Minutes Interview Hinton with Andrew Marr.

We asked our various AI to summarize how the top 15 futurists view the risks in front of us. We have included their findings below.

Futurist	Gen/Agentic/ Robotic AI Risk (%)	AGI Risk (%)	ASI Risk (%)
Sam Altman	10	30	50
Elon Musk	15	25	30
Dario Amodei	10	30	60
Ilya Sutskever	10	25	60
Demis Hassabis	5	20	50
Geoff. Hinton	10	20	25
Ray Kurzweil	2	5	10

M. Suleyman	10	25	45
Jensen Huang	5	15	35
Peter Thiel	5	10	15
P. Diamandis	2	5	10
Mo Gawdat	15	30	65
Yann LeCun	1	3	5
Andrew Ng	1	3	5
Nick Bostrom	10	50	90

Another former Google employee turned concerned skeptic is Mo Gawdat, who wrote a 2021 book called *Scary Smart* that centers on the dual nature of artificial intelligence as both a transformative force for good and a potential threat to humanity.

Looked at from the threat side of things, Gawdat asks that *even if the threat is only one-in-six*, why would we play the equivalent of Russian Roulette with our very existence??

He also believes AI's advancement is inevitable and eventually beneficial but expects there will be a period of challenge before we can obtain the post-scarcity world of abundance. Gawdat also provides us with an interesting perspective and path toward that abundance in that he emphasizes that humanity's collective behavior shapes AI's ethical framework, and therefore as AI learns from the data and actions it observes we should make every effort to model compassion, empathy and ethical values to guide AI toward a positive future. This **three-minute video** from Gawdat is worth a watch.

Finally, Elon Musk agrees with Hinton's assessment that there is a 80% to 85% probability of AI being beneficial and a 15% to 20% probability of existential risk.

Before the ChatGPT moment, no one had been more vocal than Elon Musk relative to the threats posed by AI, a threat to humanity he thought exceeded the threats emanating from nuclear war, as he *unequivocally asserts in this interview*.

Indeed, it was Musk's concern relative to this threat that led him to invest $50 million in OpenAI and to call for the six-month pause, along with Max Tegmark and another 30,000 concerned citizens.

Unfortunately, that requested pause was summarily set aside and gained no traction as the tech companies blasted forward full steam ahead.

In the **first 30 seconds of this video**, Musk agrees with Geoff Hinton that there is a 10% to 20% risk that AI will make us extinct, but goes on to say that the positive scenario far outweighs the negative scenario.

In closing this chapter, it seems clear to us that we are living through both the most dangerous time in human history as well as a time of great opportunity. In short, we are living at the hinge of history, a concept we will explore further in the next chapter.

#8: Long Termism: Living at the Hinge of History

In his book *Our Final Hour,* Lord Martin Rees suggested that those of us alive today are living at the most important time since the Big Bang 13.8 billion years ago. We agree 1,000 per cent with Lord Rees and think it is fundamentally important that every educator, every parent, and every human being understand the full weight of decisions that life as we know it has thrust on us, and that everyone understand why we are living at the most critical time in history since the Big Bang.

We hope we have begun to paint the picture of why our era is so critical. In Chapter Five, we explained that we are convinced that if— if— humanity can come together and use technology for the good of all in a safe way, we can create a world of Super Abundance, a sort of utopia where everyone lives in a high-income world of no poverty, hunger or even the need to work.

In Chapter Seven, we made clear that we have never been closer as a species to extinction than we are today, and worse, the trendlines are deteriorating rather than getting better.

There can be no doubt that humanity is at a crossroads. As we have seen and will see, technological change has transformed our world. Moreover, the technological change coming our way, change that began with the Lunar Society in Birmingham, England, some 250 years ago, has continued to accelerate at an exponential rate of change until this day; such that we can be confident that today will be the slowest day for the rest of our lives.

Two hundred fifty years after the emergence of the Scientific and Industrial Revolutions, those of us alive today are going to need to determine how we are going to use the all-powerful tools that we have inherited. Will we find ways to avoid destroying our species and build a world of Sustainable Super Abundance, or will we instead fall prey to an abundance of unimaginable stupidity that wipes out the only known conscious life-forms in a universe that we know stretches at least 93 billion light years across?

The decisions we make in the coming years and decades will determine whether the next trillion human beings, the grandchildren of our grandchildren and beyond, will ever see the light of day, as it is those of us alive today who will determine whether humanity learns to flourish or perish.

In this chapter, we want to introduce the philosophy, Long Termism, that makes super clear today's existential challenge.

Proponents of long-termism thinking argue that we have an obligation not just to every human being alive today but to future generations who have no voice, as they are yet to be born.

Long-termist philosophers believe that those of us alive today are living at a critical, pivotal time in human history; in fact, they assert that **we are living at the hinge of history.**

Living at the HInge of History

The principal thinking behind Long Termism comes from Oxford University. One of the first Oxford professors to put forward this view was Derek Parfit, who wrote in the final pages of *On What Matters* (2011) that

"**We live during the hinge of history.** Given the scientific and technological discoveries of the last two centuries, the world has never changed as fast. We shall soon have greater powers to transform, not only our surroundings, but ourselves and our successors. **If we act wisely in the next few centuries, humanity will survive its most dangerous and decisive period.** Our descendants could, if necessary, go elsewhere, spreading through this galaxy."

Parfit subsequently made the same claim in even stronger terms during a talk sponsored by *Giving What We Can* at the Oxford Union in June 2015:

"I think that we are living now at the most critical part of human history. The twentieth century, I think, was the best and worst of all centuries so far, but it now seems fairly likely that there are no intelligent beings anywhere else in the observable universe. Now, if that's true, we may be living in the most critical part of the history of the universe...(The reason) why this may be the critical period in the history of the universe is we are the only rational intelligent beings, **its only we who might provide the origin of what would then become a galaxy-wide civilization, which lasted for billions of years, and in which life was much better than it is for most human beings. Well, if you look at the scale there is between human history so far and what could come about, it's enormous. And what's critical is that we could blow it, we could end it."**

Parfit leaves no doubt that we are living at a pivotal time in human

history and that the actions we take today have the potential to help our earthly civilization flourish or perish.

Our Obligation to the Trillions Yet to Be Born

At One World, our first essential principle is that all human beings share an essential human connection. We have long believed that our obligation to help other human beings does not begin and end with our national boundaries. Indeed, One World was founded precisely to share and promote this kind of globalist thinking.

However, we had not previously thought of our obligation to future human beings, yet to be born, who have no voice as of today.

In his book, *What We Owe the Future,* William MacAskill, one of the leading Long Termist thinkers, explains that the average mammalian species tends to last about one million years. He then tells us that if we assume a constant level of population and another one million years from today, there are another 80 trillion human beings waiting to be born and explore the universe.

As MacAskill makes clear, the decisions we make today have the potential to impact 80 trillion human beings yet to be born. Again, that is "trillion," with a "t."

The idea that those of us alive today are living at a unique, pivotal time in human history, combined with the idea that we have a moral obligation not only to every human being alive today to every human being yet to be born, is called Long Termism.

Another professor who works on the idea of Long Termism is MacAskill's Oxford colleague Toby Ord, who wrote *The Precipice: Existential Risk and the Future of Humanity.* Ord explains that we are on the precipice of self-extinction with a one in six chance of self-annihilation.

The central theme of *The Precipice* is Long Termism, and Ord argues that humanity is at a critical juncture—the "Precipice"—a period of unprecedented technological power paired with insufficient wisdom, where existential risks like nuclear war, artificial intelligence (AI), and pandemics threaten our survival. This period, starting with the 1945 atomic bomb, is uniquely dangerous but temporary; if we navigate it, we could secure a flourishing future for millions of generations.

He argues that humanity is like an adolescent in that we have acquired immense power but are not yet mature or wise enough to use it

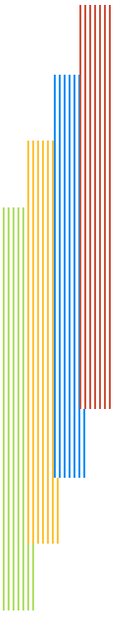

responsibly. In *Precipice*, he writes:

Like an adolescent who is given more freedom and responsibility as they grow older, humanity has gained more power and technological advancement as our societies have progressed. However, we have not necessarily gained the wisdom to use this power in a responsible and beneficial way.

We still struggle with issues such as inequality, conflict, and environmental degradation, which threaten our future and the future of our planet. ***If all goes well, human history is just beginning. Our species could survive for billions of years - enough time to end disease, poverty, and injustice, and to flourish in ways unimaginable today.***

But this vast future is at risk. *With the advent of nuclear weapons, humanity entered a new age, where we face existential catastrophes - those from which we could never come back. Since then, these dangers have only multiplied, from climate change to engineered pathogens and artificial intelligence.* ***If we do not act fast to reach a place of safety, it will soon be too late.***

Ord concludes that humanity is at a pivotal moment, with the power to either destroy itself or unlock a future of unimaginable flourishing. The **one in six risk** of catastrophe this century is unacceptably high, but we can reduce it through global coordination, targeted research, and individual action. The **Long Reflection**—a period of collective deliberation after achieving safety—is key to avoiding rushed decisions that could lock in suboptimal values or outcomes. Ord's tone is optimistic yet urgent, emphasizing that our generation's choices will shape whether humanity survives the Precipice to reach a future of vast potential or not.

To be Alive Today is an Exceptional Opportunity & Profound Responsibility

In this article published in the New York *Times*, MacAskill explains, like Ord, that we have a vast, long future in front of us if we can avoid self-destruction. He then asks us what kind of decisions we would make if we were ourselves to live this long future in front of us.

"If you knew you were going to live all these future lives, what would you hope we do in the present? How much carbon dioxide would you want us to emit into the atmosphere? How careful would you want us to be with new technologies that could destroy, or permanently derail, your future? How much attention would you want us to give to the impact of today's actions on the long term?"

Interesting questions that should inform our discussions. If we are looking at another 700,00 to 800,000 years in front of us, of course, I would insist on a 10-year pause until we can dramatically reduce the chances of self-destruction.

MacAskill goes on to explain that we are not living at the end of history but the very beginning of a very long history.

"These are some of the questions that motivate longtermism: the idea that positively influencing the long-term future is a key moral priority of our time."

"Longtermism is about taking seriously just how big the future could be and how high the stakes are in shaping it. If humanity survives to even a fraction of its potential life span, then, strange as it may seem, we are the ancients: we live at the very beginning of history, in its most distant past. What we do now will affect untold numbers of future people. We need to act wisely. It took me a long time to come."

He concludes by telling us that it took him a long time to come around to Long Termist thinking because the idea just seemed so wild, the idea that we could impact the long-term future of humanity for hundreds of thousands of years to come. However, he eventually realized that we are indeed living at an extraordinary era, an extraordinary period such that it does make sense that those of us alive today can shape the long-term future in front of us. MacAskill writes,

". . .to be alive at such a time is both an exceptional opportunity and a profound responsibility:

"*The idea that we could affect the long-term future, and that there could be so much at stake, might just seem too wild to be true. This is how things initially seemed to me. But I think this wildness comes not from the moral premises that underlie longtermism but from the fact that we live at such an unusual time.*

"Our era is undergoing an unprecedented amount of change. Currently, the world economy doubles in size about every 19 years. But before the Industrial Revolution, it took hundreds of years for the world economy to double; and for hundreds of thousands of years before that, growth rates were close to zero. What's more, the current rate of growth cannot continue forever; within just 10,000 years, there would be a trillion civilizations' worth of economic output for every reachable atom.

"All this indicates that we are living through a unique and precarious chapter in humanity's story. Out of the hundreds of thousands of years in humanity's past — and the potentially billions of years in its future — we

find ourselves living now, at a time of extraordinary change.

"A time marked by thousands of nuclear warheads standing ready to fire. A time when we are rapidly burning fossil fuels, producing pollution that might last hundreds of thousands of years. A time when we can see catastrophes on the horizon — from engineered viruses to A.I.-enabled totalitarianism — and can act to prevent them.

"To be alive at such a time is both an exceptional opportunity and a profound responsibility: We can be pivotal in steering the future onto a better trajectory. There's no better time for a movement to stand up, not just for our generation or even our children's generation, but for all the generations yet to come.

--William MacAskill

We have quoted extensively from the closing remarks MacAskill makes in the New York Times article because we agree wholeheartedly with him that everyone alive today needs to understand that to be alive is indeed an exceptional opportunity and a profound responsibility. And one of the goals of this book is to help stand up the movement that will be pivotal in steering the future onto a better trajectory. A movement of Co-Creators of the Universe committed to making certain that their next acts tip the balance on the positive side of things improving the chances of building a world of Super Abundance.

For those interested in learning more about MacAskill's thinking on these topics, click here for a summary of his book: ***What We Owe the Future***

#9: 125 Years of Accelerating Technological Transformation is Coming to a Head

From Homo Technicus to Homo Sapiens

Introduction

Many of the challenges we face today are a result of our success as a species or more accurately our success as Homo Technicus. There would be no climate change challenge if we had continued on the same path we had been on for the previous 200,000 years as the number of inhabitants would be far less as would the pressure we put on the environment coming from the technologies we have built up over the last 125 years.

The other existential risks we are facing, would also not be a worry except for the dramatic technological advances we have made over the last 125 years. Technological advances as we have seen are accelerating daily, such that today will be the slowest day you will ever live unless we face some form of civilizational collapse.

The reason the time that we are living is so important to the future of humanity is that the pace of technological change has reached a speed that challenges our societal capacity to manage that change. As we have seen we all too often attempt to apply linear mindsets to a world of exponential technological change.

The goal of this chapter is to make clear how successful Homo Technicus has been and, to make clear that we have reached a pace of technological change that will overwhelm us if we do not act urgently to address the technology *tsunami* we have created.

We can no longer kick the can down the road, as polymath after polymath have made clear the technology event horizon is imminent. The accumulating technological change of the last 125 years is coming to a head. What has become abundantly clear is that humanity needs to find a way to unite and manage the increasingly powerful technologies we have created. Technologies that have brought us to crossroads, to the brink of extinction or unprecedented success.

In this chapter, we first share with you how we arrived at this important

moment - 125 years of accelerating technological change - and then call for a new human wisdom commensurate with our technology and the name of our species. We need to urgently move from Homo Technicus to Homo Sapiens.

125 Years of Accelerating Technological Change

In the same manner that the Charles Krauthammer article fundamentally changed or informed our view of the world in terms of generating a concern resulting from the Fermi Paradox, the first 90 seconds of the TED Talk by William MacAskill fundamentally changed our understanding of history.

As you can see in the chart below, for the first 200,000 years of our existence, nothing happened, nothing really changed, and most everyone lived a subsistence lifestyle earning the equivalent of about one dollar a day. And then, in MacAskill's words, something extraordinary happened: the Scientific and Industrial Revolutions, and the flat graph we saw transformed into a straight vertical line. MacAskill explains that *"what this graph means is that in terms of power to change the world, we live in an unprecedented time in human history, and I believe our ethical understanding hasn't caught up with this fact. The Scientific and Industrial Revolutions transformed both our understanding of the world and our ability to alter it. What we need is an ethical revolution so that we can work out how to use this tremendous bounty of resources to improve the world."*

***Homo Technicus*: The Scientific and Industrial Revolutions: The Last 125 Years**

Over the last 125 years, our species— driven by *Homo Technicus*— has been extremely successful. Since the year 1900, the global population has increased by more than 400%, from 1.6 billion to 8.2 billion.

The increase in global wealth as measured by GDP has been even more pronounced. In constant US dollars, global GDP has gone from about $3 trillion in 1900 to about $110 trillion in 2025, representing a 36x fold increase. The dramatic productivity improvements translated into an eightfold increase in *per capita* GDP.

Estimated global population from 10,000BCE to 2100

(in millions)

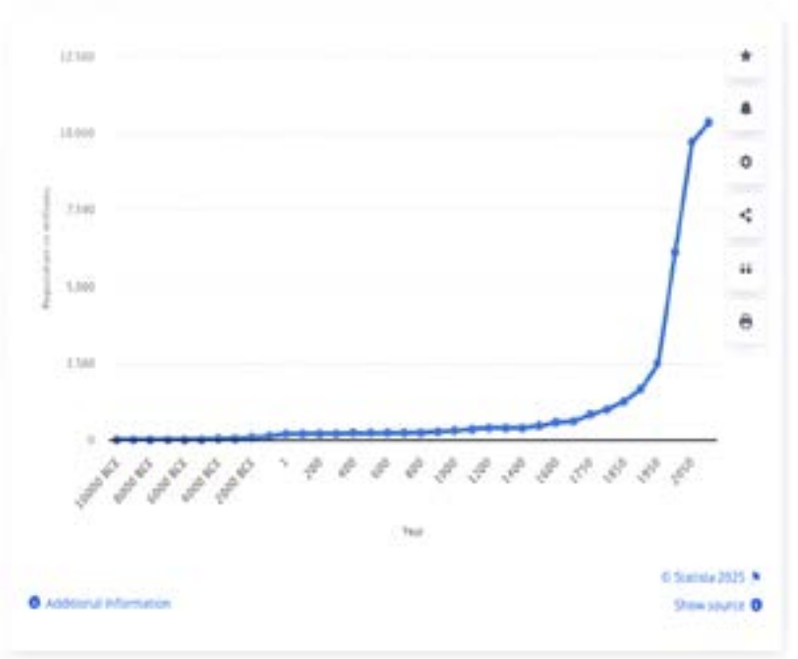

No one can have any doubt that the Scientific and Industrial Revolutions have had a very positive impact on humanity. Our species has quadrupled in size, we now inhabit almost every corner of the planet, and on a *per capita* basis, every individual alive today is eight times wealthier than the average person in 1900.

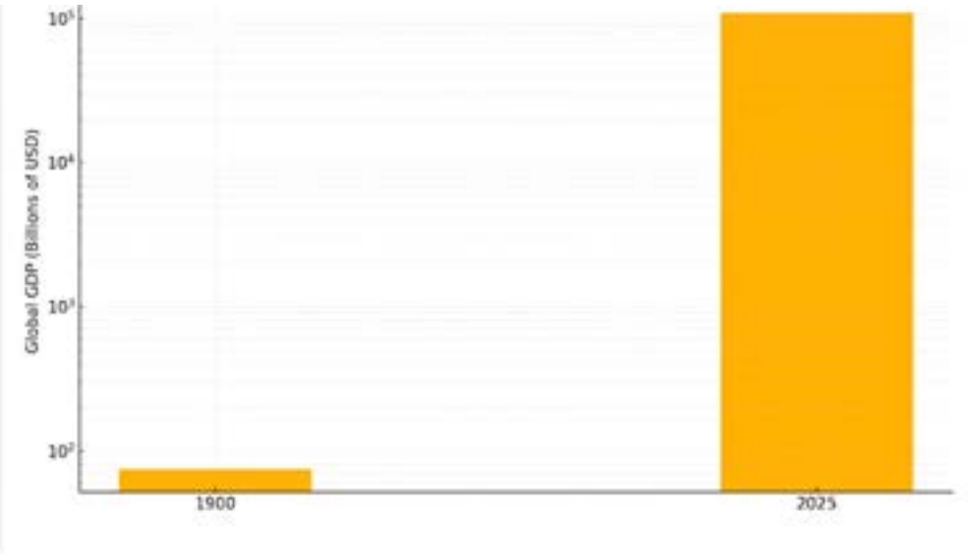

This is not to say we do not have important problems with hunger, poverty, and the distribution of this wealth within and among countries. But if we take a big step back and look at the evolution of humankind over the last 250 years, we have made great strides in terms of growing our global population and generating a vast amount of wealth for all. As it relates to *Homo Technicus*, (Technical Human) we have been enormously successful.

This increased wealth in large part owes **to the tools and technology** we have built up over the last 250 years, from the steam engine right up to ChatGPT**:** ***Our Tools over the Last 250 years***

The use of these tools has combined to transform our world in ways no one could imagine in 1900. Imagine the response you would get if you tried to explain someone from 1900, living in a city full of manure, that horses would be gone in twenty years, that human beings would learn to fly, walk on the moon, split the atom, and would be building a new digital species that would eliminate poverty, hunger and the need to work?

It would have been virtually impossible for the great majority alive in 1900 to *imagine* the profound change in the quality of life that occurred over these last 125 years. In 1900, most of the world's population lived in extreme poverty, with limited access to education, clean water, medical care, and electricity. In 1900, less than 10% of the global population had access to electricity, with literacy rates of about 20%. Today, more than 90% of the world has access to electricity, and the global literacy rate exceeds 85%.

Infectious diseases, including tuberculosis, smallpox, and cholera, were common causes of death in 1900. Child mortality was high, and many people labored long hours in agriculture or factories under dangerous conditions, with an average life expectancy of thirty-two. By contrast, 2025 offers a world where billions have access to modern healthcare, digital technology, improved sanitation, and basic education. Vaccines, antibiotics, and advances in surgery have transformed survival and well-being, with an average life expectancy of 73 years.

Beyond survival, the nature of life itself has changed. In 1900, most people had never traveled far from home, information moved slowly, and few had access to books, let alone the internet, which did not exist. In 2025, billions of people carry smartphones, connect instantly across the globe, and live in urban environments shaped by science and technology. While challenges remain—especially inequality and climate pressures—the average human today lives a life of greater opportunity, safety, and comfort than their ancestors could have imagined.

The Post-Printing Press Mindset Revolutions

One of the big questions we need to ask ourselves is, how did this happen and why did this happen when it did?

We believe the spark that changed our world started with the invention of the printing press, **which converged with the emergence of a new mindset, a new way of thinking that started with the Renaissance and evolved into the Scientific Revolution.**

In our view, the transformational change brought about first by the Renaissance and then by the Scientific and Industrial Revolutions was more than anything a change of mindset, where a new generation of thinkers sharing ideas with newly printed books challenged centuries of inherited wisdom. The Renaissance and the Scientific Revolution taught humanity how to think differently breaking with past dogma with a view to understanding our world via observation and mathematics. This change in thinking lead to revolutionary ideas like Copernicus' heliocentric model and Newton's laws of motion. We urgently need that same kind of mindset change today.

If the Scientific Revolution was about understanding nature, the Industrial Revolution was about harnessing it. Powered by coal, steam, and capital, 18th- and 19th-century Europe saw the birth of factories, the rise of machines, and the transformation of rural life into urban

industry. Entrepreneurs like James Watt, didn't just invent things—they built systems to produce, replicate, and profit from them at scale.

This too, was a mindset revolution. The Industrial Age introduced a new way of thinking: optimize, systematize, and scale. It turned labor into a commodity, time into money, and nature into "resources." It encouraged humans to see the world not as a fixed order, but as raw material to be organized, improved, and exploited. The scientific mindset was still present, but it was now embedded in a framework of economic calculation and industrial growth.

The mid-20th century introduced a quieter, but equally profound transformation. The invention of the computer, followed by the microchip, the internet, and global networks, created an entirely new plane of reality: the digital realm. Here, information could be copied, manipulated, and shared at the speed of light.

The Digital Revolution shifted our technology mindset again. It encouraged us to think in terms of systems, data, and networks. It introduced the idea that mental labor—calculation, analysis, even creativity—could be replicated in silicon. The result was a world where knowledge was no longer stored in books or brains alone, but distributed across clouds, platforms, and screens.

Today, we are living through another technology mindset revolution—As we have been discussing. Artificial Intelligence is not just an extension of the Digital Revolution; it represents a fundamental shift in our relationship to thinking itself. We are no longer just using machines to store or process information. We are now designing systems that can perceive, reason, learn, and act, eventually on their own.

However, as we have been explaining, the emergence of autonomous agents, generative models, and decision-making systems is far more transformative than any previous technology. The digital species we are creating forces us to reconsider what it means to be intelligent, creative, or even human.

This revolution is still young. Its risks are real, and its trajectory is uncertain. But its potential to unlock abundance, extend cognition, and reimagine human flourishing is unlike anything that came before it.

The profundity of the transformation requires a mindset revolution equivalent to the fundamental change in mindset that emerged with our first Renaissance.

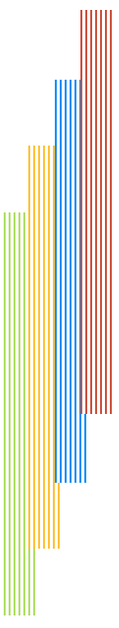

From Here to Abundance: The New Human Condition

These mindset revolutions are not merely historical, they have fundamentally changed our evolution as a species.

The Renaissance and Scientific Revolutions taught us to question. Then the launch of the Industrial revolution unleashed new technological mindset revolutions as each technology evolved from the steam engine, to electricity, to the telephone to the automobile to the airplane the splitting of the atom, to the digital revolution to the computer/internet revolution to the AI revolution.

It is beyond the scope of this book to match these technological mindsets and revolution with the global, geo-political infrastructure of the day but we can comment on the speed with which each technology was taken up and how fast they accelerated across the globe.

As we saw in Chapter 3, the speed with technologies are being taken up are accelerating rapidly. For example as we discussed, it took 35 years for 25% of the US population to adopt telephones, 30 years to adopt electricity, 25 years for automobiles, 16 years for computers, 7 years for the internet, 4 years for Smartphones, and 6 months for ChatGPT.

Again we find it useful to use the stock market as a window to the world as the stock market provides us with a unique window to the world and to the future.

For example, another way to evaluate the accelerating pace of change is to look at the average lifespan of a Standard Poor's 500 company. According to a McKinsey study, the average life-span of S and P companies has gone from 61 years in 1958 to about 16 years in 2021.

Source: Shrinking Life Cycle of an S&P 500 Company: May 27, 2021 by Nik Schuurmans, ***https://www.pureportfolios.com/shrinking-life-cycle-of-an-sp-500-company/***

Another phenomena that can be viewed from the Wall Street perspective is to see how fundamentally different the big tech companies of today are as compared with industrial companies. They have huge cash margins, that to date have not required the same kind of capital investment industrial era companies need to build out new factories.

They can therefore rapidly build out global enterprises that extend across the world targeting virtually the whole of humanity. For example, the Meta (Facebook) companies currently serve over 4 billion monthly active users.

Moreover, today they can use their huge cashflows to invest in billions if not trillions of dollars in AI as they seek to build on their first mover advantage.

For us, one of the best ways to understand the dramatic pace of change and accompanying wealth creation is to understand that it took the whole of humanity or 12,000 years to create the first company worth one trillion dollars. That accolade was reserved for Apple who became the world's first trillion dollar company on August 2nd, 2018.

However, the really interesting point from our perspective is that it only took Apple another 2 years to add another trillion dollars in value and another 1.5 years to reach the $3 trillion dollar mark. That means that Apple was able to add $2 trillion dollars in value in 3.5 years as compared with the 12,000 years it took to create the first trillion dollars in value.

On July 9, 2025, Nvidia became the world's first $4 trillion dollar company. Using the same metric as above, the new technology mindset has created $3 trillion dollars in value in 4.5 years as compared with the previous 12,000 years it took to build the first trillion dollar company.

As of today, there are now 9 companies valued at over $1 trillion dollars all of whom are technology companies with the one exception being Warren Buffet's Berkshire Hathaway.

Trillion-Dollar Companies as of July 2025

1. Nvidia: ~$4.2 trillion (first to reach $4 trillion in July 2025)

2. Microsoft: ~$3.8 trillion

3. Apple: ~$3.2 trillion

4. Amazon: ~$2.4 trillion

5. Alphabet (Google): ~$2.3 trillion

6. Meta Platforms: ~$1.8 trillion

7. Berkshire Hathaway: ~$1 trillion (joined August 2024)

8. Taiwan Semiconductor Manufacturing Company (TSMC): ~$1 trillion

9. Broadcom: ~$1 trillion (joined December 2024)

The unprecedented value of these companies and the unprecedented wealth being created virtually overnight is what is driving the race to autonomous superintelligence creating new global power centers

incorporated in a few private companies.

Much of the early change in ideas and mindset occurred in Western Europe, primarily in the United Kingdom starting in the Silicon Valley of its day in Birmingham, England where the Lunar Society held court. However, the Information and Communication Technology Revolution (ICT Revolution) combined with the rapid pace of globalization following the fall of the Berlin Wall has greatly accelerated the global takeup of all these non-industrial era technologies. Indeed, it is these companies capacity to serve the whole of humanity that drives their trillion dollar valuations. Given their global reach they should also think about taking the lead on building a sense of human connectedness.

The principal point we want to make here is that since the onset of the Industrial Revolution we have matured through several different technology-driven mindsets that have changed with each technology, Homo Technologicus has gone from strength to strength while Homo Sapiens have made little or no progress in terms of evolving our humanist mindset. One might argue that Homo Sapiens hit our peak during the Renaissance. That needs to urgently change. Our technologies spread across the globe overnight but our mindsets remain anchored in 195 competing nation states.

(For more Information on the factors shaping the Scientific and Industrial Revolutions, please see Annex #1 at the end of this chapter.)

The Need for a New Story - The Need for a 21st Century Renaissance Mindset

Viewed over the arc of the last 125 years, *Homo Technicus* has been enormously successful. Again, if we were to tell someone in 1900 that the horses that were the primary means of transportation then would be gone in twenty years, that we would learn how to fly, walk on the moon, split the atom, and create a new form of intelligence more knowledgeable than that the smartest human being, they would have thought you were crazy.

We are just now beginning to get a glimpse of our possible futures. We can create a world of Sustainable Super Abundance, eliminating poverty, hunger, and the need to work while at the same time solving climate change and greatly reducing sickness. As we saw from the NASA article, we are on the verge of creating a multi-planetary capacity.

Some futurists predict that human beings will merge with machines in the 2030s, suggesting that there may even be ways to create a form of

immortality.

What if we also told our 1900 citizen friend that the world would enter two devastating World Wars in the first half of the last century, and that the second war would only end when the most horrible of bombs was dropped on Hiroshima and Nagasaki and then comment that our next war will be our last.

So, the next big question is: How well are we prepared for this new world, for the technology *tsunami* coming our way? We need do nothing more than turn on the TV to know that our world is extremely divided, both within countries and across the world. We have already seen that at 89 seconds to midnight, the group created by Oppenheimer and Einstein tells us we are closer to self-annihilation than at any time in human history.

In short, are *Homo Sapiens* ready for the extraordinary opportunities *Homo Technicus* is ready to deliver?

Harvard Biologist E. O. Wilson: Humanity is a Chimera of Evolution

In his introduction to Jeffrey D. Sachs's book *Commonwealth*, biologist E. O. Wilson stated that he believes that humanity has hurdled into the 21st century "with a bizarre combination of Stone Age emotions, medieval beliefs and God-like technology."

Wilson concludes that humanity is an "evolutionary chimera," a patchwork of traits from different eras that don't fully align. Our emotions—fear, aggression, love—evolved during the Paleolithic era, hardwired for survival in a hunter-gatherer world. Our institutions, like governments and religions, carry the rigid, hierarchical baggage of medieval times. Meanwhile, our technology has leapt to a near-divine level, capable of reshaping the planet and ourselves.

Wilson sees this mismatch as perilous, suggesting that our primitive instincts and outdated systems struggle to manage the immense power we now wield, pushing us toward a crisis point unless we better understand and adapt our nature.

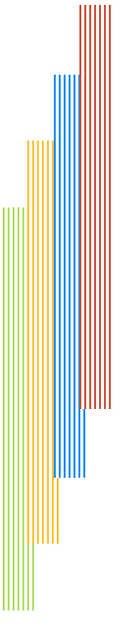

The Urgent Need to Build a Wisdom Commensurate with our Technology

If we are going to succeed over the next several decades, we are going to need to dramatically improve our Stone Age emotions and create 21st-century institutions that address the needs of our day. Improving our Stone Age emotions clearly needs to be a priority, as so many thought leaders have made clear in recent decades.

When George Kennan, one of the most influential U.S. Diplomats of all time, spoke at Robert Oppenheimer's funeral, he was crystal clear: we have created a power over nature out of all proportion to our moral strength.

"On no one did there ever rest with greater cruelty the dilemmas evoked by the recent conquest by human beings of ***a power over nature out of all proportion to their moral strength****. No one ever saw more clearly the dangers arising for humanity from this mounting disparity."*

These calls for wisdom are not new, they go as far back as 1948, when General Omar T. Bradley said, *"If we continue to develop our technology without wisdom or prudence, our servant may prove to be our executioner.* ***Ours is a world of nuclear giants and ethical infants."***

If humanity is going to survive and thrive in the coming years, we are going to need to build our moral strength and find ways to encourage ethical giants.

The popular scientist Carl Sagan (1934-1996) shared this view. In his 1994 book, *A Pale Blue Dot*, Sagan wrote that, "*The world-altering powers that technology has delivered into our hands now require a degree of consideration and foresight that has never before been asked of us."*

How much consideration and foresight relative to emerging technologies has been produced by our national and global elected leaders over the last thirty years? Virtually none!! There was virtually no discussion of the AI revolution in the 2024 presidential election in the United States.

More importantly, Sagan made clear that the dangers we see today primarily arise not from our technological tools but from humanity's lack of wisdom:

"Many of the dangers we face indeed arise from science and technology—but, more fundamentally, ***because we have become powerful without becoming commensurately wise."***

This is the same point that William MacAskill made in the TED Talk we

discussed above:

"And I believe our ethical understanding hasn't yet caught up with this fact. *The Scientific and Industrial Revolutions transformed both our understanding of the world and our ability to alter it.* ***What we need is an ethical revolution so that we can work out how do we use this tremendous bounty of resources to improve the world."***

You can here direct from MacAskill in the first 90 seconds of this video:

We Need a New Story that Unites Humanity

Clearly, thought leader after thought leader have called on humanity to build its global wisdom, a wisdom equal to our technological prowess. That much is clear. The big question is how do we as members of the human family go about building this new ethos?

One possibility is suggested by the Cambridge scholar Yuval Noah Harari, the acclaimed author of *Homo Deus* and *Sapiens*, in his book Sapiens, Hariari explains that human beings have successfully migrated from the plains of Africa to world dominance because we can organize flexibly across large numbers of people. He explains that the ability is made possible by the "stories" we tell each other. He cites as "stories" our invention of conversations relating to money, politics, religion, and the like.

For example, if a chimpanzee were offered a choice between a banana

and a $100 bill, they would choose the banana—because a $100 bill has no objective value. It only has value because of the stories we tell each other.

Harari is of the view that if humanity is to succeed well, if we are going to transition from the world of today to the world of abundance in front of us, **we are going to need a new story or stories.**

This call for a new story or a new ethos is consistent with calls for ethical action we hear consistently from other global thought leaders.

One of the principal goals of this book is to kick off a global conversation on the future of humanity to achieve a global set of values that we all can share with our children.

How different would our view of the world be if our parents had talked to us about human connectedness and our future together from day one. Of our common human goals, proud to be an earthling, and one of the only known conscious beings in a universe that stretches across 93 billion years. A universe we are preparing to explore because technology has given us that opportunity.

Annex #1: More on the Industrial Revolution

Our treatment of the Scientific and Industrial Revolutions included above was not intended to be comprehensive, as our focus was on the observable and dramatic changes in population and global wealth brought about by these revolutions.

None of this would have happened if it were not for the dynamic entrepreneurs from the Lunar Society to Silicon Valley exercising their human agency: ***Human Agency & The Lunar Society***

From Discovery to Co-Creation

Across these centuries, a pattern emerges:

- The Scientific Revolution taught us to see.
- The First Industrial Revolution gave us strength.
- The Second connected us through power.
- The Third digitized thought.
- The Fourth brings us tools that think with us — and perhaps, someday, beyond us.

What began as a quest to understand the stars has become a project to build minds.

We are no longer just using tools — we are co-creating with them. And in doing so, we face questions no machine can answer for us:

What kind of world do we want? What kind of species will we become? And how do we ensure that the tools we create continue to serve the values that make us human?

#10: The Urgent Need to Generate a Global Conversation on the Future of Humanity

One of the ways humanity will need to build wisdom worthy of the name of our species will be to build a global wisdom and sense of connectedness that has not been with us in the past.

In the opening of his 2008 book ***Common Wealth***, Jeffrey Sachs made clear how important the development of that kind of wisdom is to success in the 21st Century:

The defining challenge of the 21st century will be to face the reality that humanity shares a common fate on a crowded planet. Our common fate will require new forms of global cooperation, a fundamental point of blinding simplicity that many world leaders have yet to understand or embrace.

In the twenty-first century, our global society will flourish or perish according to our ability to find common ground across the world on a set of shared objectives and on the practical means to achieve them.

We believe we face an urgent need to develop new forms of global cooperation both in order to address today's civilizational challenges as well as to achieve the world of Sustainable Super Abundance now available to us.

To that end, we are looking to **create a Global Coalition for Sustainable Super Abundance whose first goal will be to promote a multi-level, global conversation on the future of humanity.**

There is no way humanity can achieve Global Super Abundance or avoid World War III without improved global unity. None of the problems faced by humanity today can be solved by any single nation-state acting alone. There is an even greater need for unity today than there was when Sachs issued his call for global unity in 2008.

In fact, achieving global unity is the key to solving all of our challenges as well as realizing our full potential as Yuval Noah Hararai made clear at his inaugural talk at CSER (Centre for the Study of Existential Risk) shortly after joining his new colleagues at Cambridge. Professor Harari explained that achieving global cooperation is critical not only to avoiding war, but to solving the challenges of climate change and AI. Harari explains that bringing humanity together is essential to solving existential risk:

"In this brief talk, I want to focus on the third existential threat that we are facing— global war— because in many ways it is the key to dealing with the other two. ***If humanity unites, we definitely have the resources and the wisdom to deal with both the ecological crisis and with the AI revolution. It is within our power.*** *But if humanity is torn apart by war, that would probably doom us. Given the weapons we now possess, a third world war could directly destroy human civilization, of course,* ***but even if we avoid blowing ourselves to pieces the third world war would destroy us indirectly because it would focus our attention on fighting each other and would prevent us from dealing effectively with the ecological crisis and the AI revolution."***

Harari is exactly correct, the only way we overcome our civilizational challenges or achieve the Sustainable Super Abundance now available to us is to find ways to urgently achieve global unity.

(You can access the complete Harari talk here.)

Twenty global leaders working with the Oxford Martin School of the 21st Century issued a similar call for unity in a 2013 Report: *Now for the Long Term*. At the beginning of the report, they quoted the school's founder, Dr. James Martin, who highlighted that "humanity is at a crossroads" and that this "could be our best century ever, or our worst." They then went on to define five principles for long-term success. The fifth principle they cited was as follows:

Establish a Common Platform of Understanding*: The ability to address today's global challenges is undermined by the absence of a collective vision for society. To remedy this, the Commission urges renewed dialogue on an updated set of shared global values around which a unified and enduring pathway for society can be built.* The entire document may be found *here*.

Perhaps the most urgent call for the need for global unity came from UNESCO, beseeching educators from around the world to reimagine our futures together and create a new social contract for education.

UNESCO is the United Nations Educational, Scientific and Cultural Organization. It is a specialized agency of the UN that promotes international cooperation in the fields of education, science, culture, and communication to contribute to peace and global security. Founded in 1945, UNESCO aims to foster dialogue and mutual understanding among cultures and peoples through its various programs and initiatives.

In November 2021, UNESCO published only the third global report on education in its 75-year history. The report, entitled ***Reimagining our Futures Together: A New Social Contract for Education*** calls for a fundamental transformation of education systems worldwide to address the challenges facing humanity and the planet. The report proposes a new social contract for education that aims to unite us around collective endeavors and provide the knowledge and innovation needed to shape sustainable and peaceful futures for all.

Compiled over two years, the report was surprisingly honest and clear in terms of the need for education to transform itself to meet the needs of our modern 21st-century world. The following are some excerpts from the report:

"If anything brought us together over the last year and a half, it is our feeling of vulnerability about the present and uncertainty about the future. *We now know more than ever that urgent action is needed to change humanity's course and save the planet from further disruptions. But this action must be long-term and combined with strategic thinking."*

"If the report teaches us one thing, it is this: ***We need to take urgent action to change course, because the future of people depends on the future of the planet, and both are at risk****. The report proposes a new social contract for education—one that aims to rebuild our relationships with each other, with the planet, and with technology."*

"As a species, we are at the point in our collective history where we have the greatest access ever to knowledge and to tools that enable us to collaborate. ***The potential for engaging humanity in creating better futures together has never been greate****r. "*

"The choices we collectively make today will determine our shared futures*. Whether we survive or perish, whether we live in peace or we allow violence to define our lives, whether we relate to the Earth in ways*

that are sustainable or not, are questions that will be profoundly shaped and decided by choices we make today and by our capabilities to achieve common goals."

"We face an existential choice: either continue down an unsustainable path or radically change course.

Everyone today has a heavy obligation to both current and future generations to ensure that our world is one of abundance and scarcity, and that everyone enjoys the same rights to the fullest. Despite the urgency of action and in conditions of great uncertainty, we have reason to be full of hope.

Education—the way we organize teaching and learning throughout life—has long played a foundational role in the transformation of human societies. It connects you with the world and to each other, exposes new possibilities, and strengthens our capacities for dialogue and action. ***But to shape, peaceful, just, and sustainable futures, education itself must be transformed.***

Informed by a global consultation process that engaged one million people from around the world, the goal of the report has been "to forge a new social contract for education that will help build a peaceful, just, and sustainable future for all."

The authors of the report understood *"that the visions, principles, and proposals presented by the report were simply meant as a starting point"* and that *"translating and contextualizing"* these principles would require a *"collective effort."* They therefore worked hard to clarify that the report is "***neither a manual nor blueprint but the opening up of a vital conversation***."

In view of the fact that the report was published before the Russian invasion of Ukraine in February 2022 and the seismic changes launched by Chat GPT later that year, or the tragic events of October 7, we believe the vital conversation called for by UNESCO is more urgent today, in 2025, than ever before.

Indeed, we are of the view that UNESCO's call for a vital global conversation on education needs to be taken up across all levels of our global civil society. As we will see below, education is a function of the needs of our now global society.

It is precisely these kinds of conversations relative to the future of humanity, that we intend to generate from now until the year 2050 using this book as a launching pad. We believe that *if, if, if*, we can

keep our global civilization intact until then, and if we can find ways to cooperate on the best ways to use dramatic advances in technology to advance humanity, we can create a world where by 2050 every human being alive then will have the potential to live better than the billionaires of today.

From Educating for the Fourth Industrial Revolution to Humanity 2050

“Any education is, in its forms and methods, an outgrowth of the needs of the society in which it exists.” —John Dewey, Individual Psychology and Education, The Philosopher, 12, 1934

In our view, as UNESCO makes clear, the educational needs and challenges of today’s 21st Century world. The needs of today’s global society are to prepare K12 youth for a rapidly globalizing world of exponential technological change.

We believe no educational leader has done more to prepare educators and students for this twin challenge than Fernando Reimers, who is the Ford Foundation Professor of the Practice of International Education and Director of the Global Education Innovation Initiative at Harvard. The reason Reimers has been so successful is that he understands that the twin challenges are, in truth, one and the same.

As Dr. Reimers explains in ***this essay***, the best way to prepare global youth for the Fourth Industrial Revolution, or the world of exponential technological change, is to start with the end in mind—global citizenship—and provide students with an opportunity to collaborate in diverse groups. We have outlined each of the five steps proposed by Dr. Reimers because we have worked with him to deliver these to schools in the past and will use them as an important part of our Project 2050.

Dr. Reimers concludes his essay explaining that: “Global citizenship is essential for seizing the enormous possibilities and addressing the great challenges of our times. While cultivating it is the task of educators, the global community is equally responsible for supporting and encouraging, in a collaborative manner, the education of global citizens to whom nothing human is foreign.”

Building a Global Coalition for Sustainable Super Abundance – Humanity 2050

This last point made by Dr. Reimers is critical to the success of Project 2050. The kind of global unity needed will need to be society-wide. We cannot rely on educators alone to build global unity. They have a role to play, as we will see in Chapter 12, but they cannot do it alone.

Moreover, as John Dewey pointed out, educational needs are an outgrowth of the needs of the society in which they exist. If we are going to succeed, the need for global unity will need to be taken up at all levels of society.

The need, or a call for global unity at all levels of society, is not new, as we have seen throughout this book and in this chapter. However, it seems that all too often, the call is made without giving sufficient thought to how that call is implemented. Thought leaders from across the world often call for the establishment of a common platform for understanding to generate a collective vision for society but make no provision as to how that is supposed to happen. As a result, the kind of platforms needed to effect fundamental change are too few and far between.

If in the twenty-first century, "our global society will flourish or perish according to our ability to find common ground across the world on a set of shared objectives and on the practical means to achieve them," don't we need to urgently develop new global platforms to build the needed consensus.

If today's young people "will collectively determine whether civilization survives or not" don't we need urgently need to provide them with the foundation for making wise choices to understand humankind's likely possible, probable, and preferred futures? To understand what humanity can become?

Unfortunately, as we read through the literature for too many calls to build the platforms for global understanding go unheeded. All too often, they are calls for action without any follow-up. As we will explain more fully below and in Chapter 14, the One World team is committed to building the Global Coalition for Sustainable Super Abundance. Our first project will be to generate a global conversation on the future of humanity on as many levels, in as many places as possible, however we are going to need help.

There are of course, stand-out exceptions to the lack of follow-up. Following the publication of his book, Our Final Century, Sir Martin Rees

created *The Centre for the Study of Existential Risk* (CSER) at Cambridge University at a time when there were more academic studies on snowboarding than there were on existential risk. Sir Martin Rees also worked with James Martin, who contributed $150 million to Oxford University to create the Oxford Martin School of the 21st Century, effectively supporting much of the work that has been conducted relative to Long Termism and effective altruism.

Indeed, we would argue that the creation of CSER by Sir Martin Rees and his colleagues has catalyzed the development of a global X-Risk ecosystem that now includes the Future of Life Institute at MIT and many more institutions of higher learning and think tanks. Here is a link to the first catalogue of that ecosystem: ***https://oneworlduv.com/wp-content/uploads/2023/05/Annexe-The-X-Risk-Ecosystem-.pdf***

A Multi-Level Conversation on the Future of Humanity

Our hope is that we can encourage this emerging ecosystem to come together and generate a global conversation on the future of humanity. A conversation that we hope will lead to a new story that generates a wisdom commensurate with our technology.

The universities highlighted in the catalogue above are a great start. However, every university in the world should have a class on civilizational risks and the potential for Super Abundance. We need to go beyond the emerging ecosystem we have identified. The feedback we are getting from university students is that none of the issues raised in this book are being discussed in any meaningful way, with the possible exception of climate change.

This global conversation needs not only to expand to the rest of academia it also needs to go beyond the walls of our universities. This conversation needs to be expanded to include to K12 students and educators, as will be seen clearly in Chapter 12.

It also needs to be expanded to civil society. Municipal, state and governmental leaders all have a role to play in terms of climate change, creating a more peaceful and fairer world as well as responsibly using AI.

The Global Business Community Needs to Get Serious about Human Survival and Sustainable Super Abundance

Perhaps most importantly, all segments of the business community need to engage in this conversation. The global leaders involved in AI need to do more than complain about the lack of regulation. They need to come together and insist on the kind of regulation that makes sense for all of us. They have the resources and knowledge needed to make this happen. They do not need to wait for governments to act. They should form a global coalition and make clear the rules of the game that need to be abided by everyone across the globe. They can and should take the lead on the regulation conversation. They clearly have a special, outsized responsibility to all of us.

The conversation needs to go beyond the tech companies. We are old enough to remember Harvard Professor Raymond Vernon's books published in the 1970s, criticizing the practices of multinational corporations or MNCs. At the time that Vernon was writing, there were about 400 MNCs, today, there are over 100,000 MNCs. In recent years, we have wondered what the real nationality of today's top companies is, as they all operate in numerous jurisdictions with the requirement that they be good citizens of each.

(One of the highest levels of praise we received at One World came from a senior executive in the financial sector who told us that the time he has spent with One World was more beneficial to him than any professional development he had experienced in his 3 decades in the corporate world. He felt that being a part of One World transformed the way he worked with and collaborated with his team members from across the globe.)

Multinational corporations should be encouraged to talk about how they attempt to promote human connectedness within their organizations and beyond. They should decide to be important, authentic players in building communities across the world and in emphasizing our human connectedness and the fact that we are all in this together.

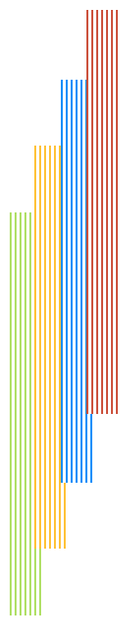

Humanity Needs to Focus on Human Survival & Sustainable Super Abundance

We very much support the fight against climate change and have developed award winning, K12 programs to take action in the fight against climate change. Indeed, we worked with the secretariat of the UNFCCC (United Nations Framework Convention on Climate Change) for a number of years to engage K12 schools in the fight against climate change as you see from this article:

https://www.lohud.com/story/news/education/2019/12/10/port-chester-schools-first-us-neutralize-carbon-un-program/4320063002/

To date the fight against climate change is the principal existential risk that has gained traction across the global in terms of media attention, academic articles and student participation which is understandable given the fact that the first Earth Day was held on April 22, 1970 and the first Conference of Party (COP) meeting organized by the UNFCCC was held in Berlin in 1995.

It has taken time to understand and prepare for the challenges emanating from climate change. As a result, news and academia are currently far more focused on climate change than they are on AI or even nuclear threats: *Climate News vs AI News vs Nuclear News*

However, we are now living at the precipice of self-destruction as the threats to our species increase. The ChatGPT moment has accelerated the threat arising from AI to the point where the leading participants in the field are telling us that there is a 15% to 20% probability of existential risk emanating from Autonomous Superintellgence which is expected in a “few thousand days”. At the same time, the threats arising from nuclear warfare have dramatically increased in recent years.

The number one skill set everyone agrees needs to be taught in schools across the world today is critical thinking. We need to urgently apply critical thinking to the emerging challenges and opportunities facing humanity today. We cannot continue to focus solely on one of the three civilizational challenge facing us.

The reason climate change has rightfully gained so much attention is because climate change threatens all of us. Given the dramatic and accelerating advances in AI brought on by the ChatGPT moment we definitely need to include the risk of AI in our existential risk conversation because once again we are all threatened. There is no sustainability for our species if AI takes over and decides to get rid of us.

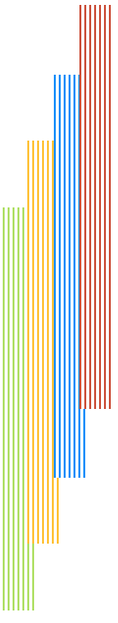

The proliferation of nuclear threats in recent years also needs increased attention. What good is climate change if we only destroy our environment with the sudden outbreak of nuclear war.

Yuval Noah Harari could not have been clearer there is no solution to climate change or AI without global unity.

As we have been explaining, AI is different in that it brings with it 80% to 85% probability that we can achieve a world of Sustainable Super Abundance if we manage to align this new powerful digital species with the right kinds of humanist goals. And while AI will put greater pressure on our global environment in the near term the expectation is that the emerging superintelligence will help us find the technical solutions needed to resolve climate change.

We need to urgently expand our global conversation on climate change to one focused on human survival and Sustainable Super Abundance. It won't be easy to redirect the thinking that has accumulated over more than three decades. However, the ChatGPT moment has left us with no choice to be to find ways to urgently respond to the challenges and opportunties of our day.

To get a sense of the change that needs to take place we conducted two relatively extensive surveys using the various AI tools or foundational models available to us. We first looked at student participation and academic engagement relative to all three existential risks and the idea of Sustainable Super Abundance after which we compared the three risks and the idea of Sustainable Super Abundance (SSA) as picked up by news organizations and academia.

Our AI search relative to student participation and academic attention to Climate Change, Nuclear Proliferation, AI risks and the potential for Sustainable Super Abundance is included below. Here is how Perplexity AI summarized their findings:

• Climate change is by far the most widely studied of these issues in higher education worldwide.

• Nuclear proliferation and AI risks are much more specialized and attract far fewer students, almost exclusively at the graduate level.

• The explicit study of sustainable superabundance is essentially absent from formal academic curricula, reflecting its status as an emergent and marginal theoretical concept.

DeepSeek provided us with the following numbers:

Ranking by Student Engagement

- Climate Change (~600,000+ students)
- AI Risks (~70,000+ students, growing fast)
- Nuclear Proliferation (~15,000–30,000 students)

Key Observations

- Climate dominates due to institutionalization in academia.
- AI risks are rising but still fragmented across disciplines.
- Nuclear studies have faded since the Cold War.
- SSA is barely taught—lacks academic infrastructure.

For the full AI survey please click here: ***Student Involvement in SSA & X Risk***

As one would expect given the length of time humanity has been dealing with Climate Change, news media and academic coverage of climate change is far more robust than coverage of the threats emanating from AI or Nuclear War, although the coverage of AI is on the rise.

AI Survey Results

- Climate change remains by far the most reported and academically studied global threat.
- AI threats are rapidly gaining prominence, especially in the past few years.
- Nuclear war coverage and research are substantial but less sustained, with periodic spikes in response to world events.

Clear here for full AI survey: ***Climate News vs AI News vs Nuclear News***

Our AI surveys also found that there is very little media or academic coverage of humanity's potential to achieve Sustainable Super Abundance.

AI Survey Results

- Sustainable superabundance is virtually absent from mainstream news and academic discussions compared to the three major threats.
- Coverage of this optimistic future is dominated by a small set of futurists and select think tank publications.
- The idea is gaining minor traction in discourse around energy

transitions and post-scarcity economics, but remains a marginal theme relative to existential risks[1279104].

In summary: While the threats of climate change, AI, and nuclear war dominate academic and media attention, sustainable superabundance is an emerging narrative that garners minimal but gradually increasing interest, primarily among thought leaders, futurists, and a few dedicated researchers[189].

Clear here for full AI survey: *Sustainable Super Abundance in the News & Academia*

The purpose of this book is to invite civil society from across the world to join the Global Coalition for Sustainable Super Abundance (GCSSA) we hope to create with your support. Our goal is that the first act of GCSSA will be to generate a new, expanded more urgent conversation on the future of humanity that covers all three existential risks as well as the potential for Sustainable Super Abundance.

There is precedent for expanding the X-Risk conversation. For example the Bulletin of Atomic Scientists and virtually all of the X-Risk ecosystem participants have expanded the civilizational risk conversation to include all three predominant threats with some including other threats like biosecurity and pandemics.

However, as far as we know this is the first time someone has suggested we expand the conversation to include the potential for Super Abundance as well. We are confident that including the potential for Sustainable Super Abundance will be a welcome new positive angle on this all-important conversation.

Reorienting and Refocusing ESG

In recent years, numerous MNCs got involved with the ESG movement (Environmental, Social, Governance), but that movement has come under pressure due to lack of standardized measures, that led some corporations to game the system. Without standardized, measurable criteria, the efforts made by corporations were often seen as virtue-signaling rather than authentic attempts at community building.

We would encourage corporations from across the world to get involved with a new ESG movement but where the E stands not just for climate change, but for all three existential risks - climate change, nuclear proliferation and AI.

We never understood why MNCs, who are by definition global citizens, have not gotten more fully involved in the effort to build global

citizenship.

In thinking about a new, more serious-minded corporate movement that tackles existential risk head-on, it seems clear that these global corporate citizens with their tremendous bounty of resources have a critical, authentic role to play in bringing humanity together.

Perhaps we can move to something like EEESG or **E**sg where the big **E** stands for all three existential risks. Even better the movement should consider going one step further and ask the corporate world to think about Esg4 Sustainable Super Abundance.

We believe that businesses and corporations have an essential, authentic role to play in generating a global conversation on the future of humanity and our need to face our existential risks while making every effort to build Sustainable Super Abundance. No sector of the globe has more to lose than our leading businesses and their employees or corporate citizens. The more you have, the more you stand to lose, with existential catastrophes.

One thing is crystal clear for anyone who spends 3 minutes looking at our world today, the status quo is not tenable, we need transformative change at every level. Multinational corpporations would do themselves and humanity a great service to help their employees focus on their human connectedness and human survival.

The Role for Multilateral Organizations and the SDG Conversation

As can be seen from the discussion above our focus has been on encouraging global civil society understood in the broadest of terms to get involved in this conversation rather than direct our focus on multinational organizations. The principal reason for our focus is that we believe that in order to effect the kind of urgent consequential change we are hoping to achieve, all segments of our global civil society need to be involved in the conversation.

This is not an effort that can be simply handed turned over to any single entity be they national or multinational. We are hoping to create a movement that builds a new 21st Century Renaissance mindset whose end goal is transform the human project such that we become a multi-planetary species.

However, that being said we would welcome an informed, enthused response from the UN and other multilateral agencies We have been

very supportive of the UN SDG goals with our programs and are grateful for the establishment of these humanity wide goals. In many ways the SDG conversation could be considered a first attempt at the kind of conversation on the future of humanity we have been discussing.

We would be overjoyed to work with the UN on the kind of conversation we are hoping to generate. The first two SDG Goals are to bring an end to poverty and hunger, that is the promise of Sustainable Super Abundance.

However, our concern is that having 17 SDG goals dilutes the impact and focus of the message. If the UN or any other multilateral organizations were interested in pursuing a future of humanity conversation with a re-purposed SDG goal conversation that starts with the urgent need to overcome the civilizational threats facing us and ushering in a world of Sustainable Super Abundance thereby achieving the first two goals we would be very interested in pursuing that path.

We have also featured the UNESCO November 2021 report on education. The conversation they called for in November 2021 has been launched. It would be wonderful to find ways to get UNESCO working with OECD and PISA and other entities like the World Economic Forum and the ASIA Society to come up with a set of updated recommendations on how it is we really need to transform education across the globe.

To be clear, we would be more than happy to work with the UN and any other multilateral organization interested in working together to generate this future of humanity conversation we have envisioned.

In the next chapter we will discuss the urgent need for informed AI regulation.

Annex #1

Here is the ***UNESCO report***

CNBC Feature on the UNESCO Report, with an interview of Dr. Fernando Reimers

#11: The Urgent Need to Get Serious about AI Regulation

As we saw very clearly in Chapters Three and Four, the intelligence explosion is here. The future is here. The ChatGPT moment accelerated the arrival of AGI and ASI by two to three decades along with the risks and opportunities associated with them.

The fundamental question for humanity is, how are we going to prepare for and usher in the new future in front of us now that the timeline to prepare has been accelerated by 20 to 30 years?

What is clear to us is that as of today, the AI in place today—Perception AI, Gen AI, Agentic AI, and Robotic AI—are mostly manageable relative to existential risk and have the potential to deliver 85% to 90% of the economic gain in the near term.

However, what is also clear is that following the ChatGPT moment, AGI and ASI are about to arrive twenty to thirty years ahead of schedule. As we saw, AGI and ASI begin to represent different levels of autonomy and therefore different levels of risk. We

The futurists tell us that AGI occurs when AI is as intelligent as the average human. However, as we made clear in the last chapter, in a world where AI is capable of reading 1.2 billion books a year our AI are clearly far more capable than us in terms of processing and absorbing information. Based on these metrics one could argue that our AI are on track to become 100 million times "smarter" than us.

So, the only real change that occurs with the arrival of AGI is not an increased intelligence differential rather the fundamental change is one of autonomy. That is the change will not be that AI will all of a sudden be as smart as AI, as we saw above it is 100 million times smarter than us in narrow domains of pure knowledge. The big change will be that with the advent of AGI our AI will be able to act with agency and think strategically.

The difference will be that these extraordinarily "smart" entities will no longer need us to direct them with our prompts they will be able to act on their own in a generalized way having gained their own agency. What they do with that mountain of knowledge and agency no one knows. Therein lies the existential risk.

In our view, once AGI happens, ASI will follow shortly thereafter because once the AGI comes into being with agency it will rapidly act to multiply

adopting recursive self-improvement, and very quickly embody an intelligence that is far greater than the intelligence of all human beings put together as IJ Good predicted in 1965.

The important point to understand here is that the ChatGPT moment accelerated our reckoning with AI by two to three decades. Moreover, the reckoning has been unleashed on a disconnected, contentious, warring world, where the global systems that were put in place after our last world war are fraying if not falling apart. The ChatGPT moment has unleashed the full force of the AI reckoning on our current world order with all its strengths and weaknesses.

The ChatGPT moment has unleashed a ferocious global competition between the principal corporate actors as well as create a high stakes geo-political competition between the leading nations of the world, particularly China and the US, as they fight for global dominance. At the same time, all of these actions put humanity on a potentially collision course with powerful AI or AGI/ASI.

The existential concern is that the extraordinary intelligence we are creating, this super intelligence might just decide to treat us no differently than we treat insects, because the intelligence differential between us and the AI will be far greater than intelligence difference between us and ants.

If the stakes were not so high, the irony would make a wonderful setting for a Greek Tragedy. The world's leading trillion-dollar corporations led by billionaire entrepreneurs are caught in a ferocious race to be first past the post. The billionaire entrepreneurs are not doing it for money they are doing it for fame, power and most importantly for ego.

At the same time we have the great power struggle going on between nations states with China and the US taking the lead in the fight on that front, again each determined to be first across the finish line.

These internecine competitions makes it far more likely that humanity as a whole will lose control of this process greatly increasing the chances of the existential riak emanating from AI. Rather than fighting each other we should be working together and with AI to create the world of Sustainable Super Abundance our technological genius has made available to us. Once again it seems as if we are in great need of a wisdom commensurate with our technology.

One reason the stakes in this competition are so high is the fact that most of the trillion-dollar tech companies competing for AI victory are what Eric Schmidt calls "network effect" businesses where generating

a first mover advantage with AI can become a definitive competitive edge.

As Eric Schmidt explains below, for network-effect businesses, "it is the slope of your improvement that determines everything." He goes on to explain that as AI "gets closer to superintelligence, the slope gets even steeper" and therefore if a company can crack the AGI challenge making their AI fully autonomous it can give them a competitive advantage that might never be overcome.

For the principal players engaged in this fight click here: ***The AI Battle for Supremacy***

There is not an AI actor in the world that does not understand the risks associated with the increasing autonomy accruing to the digital species we have created. Indeed, two former Google employees Mo Gawdat and Geoff Hinton have left Google to speak without constraint on the risks facing humanity as a result of the increasingly powerful AI we are creating.

There is a broad consensus that the existential threat coming from AI taking over is somewhere between 10% to 20% or about one in six. It seems to us that, Mo Gawdat is 1,000% correct when he asks why would we play Russian roulette with the future of trillions of human beings yet to be born?

In addition to the threat from existential risk, AI presents another set of challenges that humanity is ill-prepared to address given the lack of clear and consistent regulation:

- Misinformation and Manipulation
- High Unemployment and Economic Inequality
- Bias and Discrimination
- Autonomous Weapons and Battle Robots

For more information on these risks click here: ***AI Risks***

On the other hand, to their credit, many AI entrepreneurs are begging for real, bona fide regulation. They want us to put guidelines and rules in place. They are playing a game as of now with no rules and no referees. We have never seen an industry where the principal participants are begging to be regulated. That should tell us something.

Elon Musk has been begging for increased AI regulation for years, consistently contending that AI represents far more of an existential threat to humanity than even nuclear war.

In the first minute of this short video, Elon Musk suggests that AI is perhaps the most profound existential risk that humanity has ever created and that if we are to successfully build the world of abundance now available to us we will need new strong regulation:

In fact, in an ironic twist of history, Elon Musk helped create the ChatGPT moment by contributing to the creation of Open AI an entity which was initially set up as a not-for-profit with a mission to ensure AGI development prioritizes safety and public good over commercial interests

In 2012, Musk met DeepMind co-founder Demis Hassabis, who warned that AI could become superintelligent and potentially threaten humanity, prompting Musk to consider the risks of a single entity controlling AGI. ***Sources #1***

By 2013, Musk was alarmed by Google's growing AI prowess and started a conversation with eventual Open AI co-founder Sam Altman.

In May 2015, Altman wrote to Musk, "Been thinking a lot about whether it's possible to stop humanity from developing AI. I think the answer is almost definitely not. If it's going to happen anyway, it seems like it would be good for someone other than Google to do it first."

Musk replied, "Probably worth a conversation." Six months later, OpenAI was founded as a nonprofit to counter Google's lead, aiming to develop AGI safely and distribute its benefits broadly.

Former OpenAI researcher Daniel Kokotajlo later claimed that Musk, Altman, and co-founder Ilya Sutskever launched OpenAI to prevent DeepMind's Hassabis from becoming an "AGI dictator," reflecting fears of Google's concentrated power. Musk himself stated in a 2023

Fox News interview that OpenAI was created as a "counterweight to Google," reinforcing this narrative.

OpenAI was founded on **December 11, 2015**, as a nonprofit organization focused on advancing AI research for the benefit of humanity. It was established by co-founders including Sam Altman, Elon Musk, Ilya Sutskever, Greg Brockman, and others, with a mission to ensure AGI development prioritizes safety and public good over commercial interests.

In the meantime, Sir Demis Hassabis was hoping to build the "NASA of AI" following the January 26, 2014, sale of the company he co-founded with Shane Legg, and Mustafa Suleyman for $500 million. In the early days he was able to recruit the many of the top AI researchers in the world as there was little competition.

According to Demis Hassabis, DeepMind's goal is to **"solve intelligence, and then use that to solve everything else,"** focusing on accelerating human scientific discovery through artificial intelligence. Founded in 2010 and acquired by Google in 2014, DeepMind generated a number of successes in its quest to develop artificial general intelligence (AGI) that can learn, reason, and solve problems across diverse domains. As we have seen, they developed AlphaGo a game that learned on its own and performed the famous Move 37 on its way to beat the global GO champion one decade ahead of schedule.

In 2020 Google DeepMind created another landmark success with **AlphaFold**, which solved the decades-old problem of protein folding. By predicting the 3D structures of proteins from amino acid sequences, AlphaFold achieved accuracy comparable to experimental methods, revolutionizing biology and drug discovery. Its database, released in 2021, mapped 98.5% of human proteins and over 200 million structures across species, freely accessible to researchers. This work earned DeepMind global acclaim, with applications in developing vaccines and treatments for diseases like malaria. AlphaFold's impact underscores DeepMind's goal of accelerating scientific progress through AI. Moreover, Demis Hassabis and John Jumper won the 2024 Nobel Prize in Chemistry for their work on AlphaFold. Mr. Hassabis was also knighted in 2024.

We mention this evolution because it seems that looked at from afar, the Google AI DeepMind team was more focused on solving scientific discovery than they were on developing commercial use cases, at least until they could be more sure of the safety angle. In fact, rather than monetize AlphaFold, Google decided to make this program freely

available to the scientific community, thereby forgoing substantial revenues if they had decided to commercialize it.

(To get a sense of the accelerating pace of change, Google DeepMind announced three days ago on June 25, 2025, the launch of a new AI, Alpha Genome, that represents another step forward in terms of improving human health outcomes: ***AlphaGenome***).

Indeed, the mad dash that has existed since the OpenAI launch of ChatGPT was not in evidence when Max Tegmark and the Future of Life Institute (FLI) organized the first three informal AI conferences in Puerto Rico and Asilomar. In those days the field of participants was quite small and manageable. Also, the median expectation for the expectation for the arrival of AGI as of the time of the first Puerto Rico Commence was 2055 as a result the discussions that took place had far more to do with the expected timing of the arrival of AGI rather than prepare for its imminent arrival.

As you can see from clicking on these documents the conversations and meetings organized by FLI seemed hopeful and on point with a sense of purpose but not urgency.

The First Informal AI Conference in Puerto Rico

The Asilomar Confereence Janaury 2017

Puerto Rico #2 - The Beneficial AGI Conference

However, in the years following the second conference in Puerto Rico and leading up to November 2022, AI technology was dramatically improving due to advancements in transformer architecture and more powerful chips from NVIDIA. This period, marked a transformative phase in AI, driven by breakthroughs in algorithms and hardware.

At the same time the escalating financial demands of scaling AI models pushed organizations like OpenAI to seek new funding strategies to remain competitive. The skyrocketing costs and competitive landscape led OpenAI, founded in 2015 as a nonprofit to develop safe AGI, to transition to a **for-profit arm** (OpenAI LP) in March 2019.

The nonprofit model, funded by $130 million in donations by 2019, couldn't sustain the billions needed to train models like GPT-3 or compete with Google's DeepMind and Microsoft's Azure AI.

The for-profit "capped-profit" structure allowed OpenAI to secure a **$1**

billion investment from Microsoft in 2019, followed by $10 billion in 2023, enabling it to scale infrastructure and hire top talent.

This shift, while controversial (Musk criticized it as a betrayal in 2023 X posts), was driven by the need to fund compute-intensive research and stay in the AI race, where rivals like DeepMind (acquired by Google in 2014) leveraged massive resources.

The transition maintained a nonprofit board to oversee the mission but prioritized commercial products like ChatGPT, which achieved a $10 billion revenue run rate by 2025, reflecting the financial imperatives of the era's AI advancements.

It was this transition that provided Open AI with the funding they needed to be in a position to launch ChatGPT in November 2022.

Frightened by the launch of ChatGPT at a time when the safety of emerging AGI was still in question, in March 2023, Max Tegmark teamed up with Elon Musk to ask for an immediate 6 month pause in the development of AI systems more powerful than GTP-4.

The Future of Life Institute (FLI) pause letter started like this:

"We call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4.

AI systems with human-competitive intelligence can pose profound risks to society and humanity, as shown by extensive research[1] and acknowledged by top AI labs.[2] As stated in the widely endorsed Asilomar AI Principles, Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources. Unfortunately, this level of planning and management is not happening, even though recent months have seen AI labs locked in an out-of-control race to develop and deploy ever more powerful digital minds that no one–not even their creators–can understand, predict, or reliably control."

We highly recommend you read the full letter here: ***FLI Pause Letter***

The Pause letter was one last attempt to rein in an industry that was exploding at breakneck speed following the ChatGPT moment and the dramatic, accelerating advances occurring in AI technologies. In spite of being signed by a number of leading thinkers on AI, the letter was largely ignored.

The shift at Open AI from not-for-profit to for profit came to the fore in November 2023 when the OpenAI nonprofit board, fired Sam Altman as CEO and removed him from the board, citing a lack of "consistent

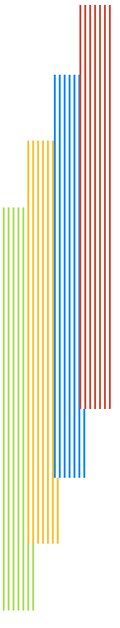

candor" in communications.

The decision was abrupt, announced via a blog post, and followed a review of Altman's conduct. Mira Murati, the CTO, was named interim CEO. The board's statement suggested Altman's actions undermined their ability to oversee OpenAI's mission, though specifics were initially vague.

The firing occurred amid tensions over OpenAI's rapid commercialization, with ChatGPT's success and Microsoft's $10 billion investment in 2023 shifting focus toward profit. The board, tasked with prioritizing the nonprofit mission, reportedly clashed with Altman's aggressive push for growth.

Eventually, they explained that he was let go because of lack of transparency, mission drift, internal conflicts and personal ambition. However, Altman was quickly reinstated a few days later when the vast majority of Open AI employees signed a letter demanding his reinstatement.

It is worth noting that the OpenAI employees held equity in the for-profit arm which was valued at $86–$90 billion in 2023, per *Forbes*.

Altman's leadership drove this valuation through ChatGPT's success (100 million users by January 2023) and Microsoft's $10 billion investment. His firing threatened a potential collapse in OpenAI's market position, as competitors like Google ($2 trillion market cap) and Anthropic ($61.5 billion) could capitalize on the chaos. A mass resignation to join Microsoft, as threatened in the letter, offered employees a chance to preserve or even enhance their equity value, as Microsoft proposed hiring them under Altman to form a new AI research unit.

Elon Musk has filed multiple lawsuits against OpenAI and its CEO Sam Altman, primarily alleging that the company breached its founding mission to develop artificial intelligence (AI) for the benefit of humanity by prioritizing profits through its for-profit arm and partnership with Microsoft.

Whatever the results of the lawsuit, one thing is clear the move from not-for-profit to for-profit enabled Open AI to partner with Microsoft to secure the funding they needed to build out ChatGPT and be the first company to launch Generative AI in November 2022, marking the next stage in human technological history and perhaps human history itself.

The race for AGI is on led by a cohort of trillion-dollar companies committed to winning the race, seemingly with little or no thought to

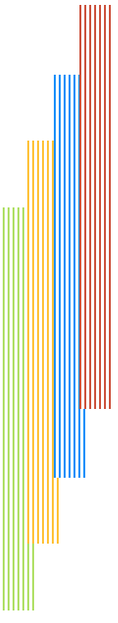

cost or impact. If the decision were ours to make we would pause the development of AGI for several years so that we could be sure to manage the risks going forward. That being said, it seems very clear to us that is highly unlikely a pause can be engineered.

Indeed, Max Tegmark has said famously that there were two guiding principles in thinking about how to contain AI's autonomous development: don't provide AI with access to the internet, and don't teach it how to code. In the race for power, both are long gone!!

That being said, the former UK Prime Minister took the initiative to call for the first-ever government-led summit in Bletchley Park in the UK in November 2023 and Max Tegmark has continued to fight for AI safety.

That UK-led coalition-building initiative has been followed by several additional government-led coalition-building initiatives where the Future of Life Institute turned advisor rather than organizer. As you will see below when you review the summits that have been organized and conducted this is still an ad hoc process where neither the US or Chinese governments have taken the lead but it has produced some important guiding documents as you will see below.

Bletchley Park, UK: November 1-2, 2023

Co-chaired by UK Prime Minister Rishi Sunak and South Korean President Yoon Suk Yeol, this was the first government-sponsored conference of its type. It was a landmark global conference organized by Sunak to address the safety and regulation of frontier AI—highly advanced general-purpose AI systems. Hosted at Bletchley Park, the historic World War II codebreaking site symbolizing technological innovation, the summit convened representatives from 28 countries, including the United States, China, and the European Union, alongside tech leaders including Elon Musk and Sam Altman (OpenAI), academics such as Yoshua Bengio, and civil society groups.

Initiated by the UK's Department for Science, Innovation and Technology (DSIT) under Michelle Donelan, the event aimed to foster international cooperation to mitigate AI risks, including terrorism, disinformation, and existential threats from uncontrolled systems. The summit resulted in the Bletchley Declaration, a commitment to safe, human-centric AI, the establishment of the UK's AI Safety Institute, and a mandate for an International Scientific Report on Advanced AI Safety led by Bengio, with follow-up summits planned in South Korea (2024) and France (2025). You can read the declaration here: ***The Bletchley Declaration***

You can read PM Sunak's Summary Here: *UK PM Rishi Sunak Summary of the Conference*

AI Seoul Summit; May 21-22, 2024

Co-chaired by UK Prime Minister Sunak and South Korean President Yoon Suk Yeol, the UK and Korea agreed to continue to work together with the AI Seoul Summit, held on May 21–22, 2024. This was a hybrid event comprising both virtual and in-person components. The Leaders' Session on May 21, co-chaired by UK Prime Minister Rishi Sunak and South Korean President Yoon Suk Yeol, was conducted virtually, allowing heads of state and global leaders to participate remotely.

The Digital Ministers' Meeting on May 22, co-hosted by the UK's Department for Science, Innovation and Technology (DSIT) and South Korea's Ministry of Science and ICT, took place in Seoul, involving ministers, industry representatives (e.g., from Google, OpenAI), and academics. Additionally, the AI Global Forum, hosted by South Korea's Ministry of Foreign Affairs, was an in-person event on May 22, broadening participation. The virtual format of the Leaders' Session, often referred to as a "mini virtual summit," was designed to maintain momentum from the Bletchley Park Summit (2023) but was criticized for limiting engagement compared to fully in-person summits like Bletchley or Paris (2025).

Here are the principal declarations to come out of Seoul: *Seoul AI Summit*

Paris AI Action Summit February 10-11, 2025

Co-chaired by French President Emmanuel Macron and Indian Prime Minister Narendra Modi, the next major international AI summit after the AI Seoul Summit was the **AI Action Summit**, held on **February 10–11, 2025**, at the Grand Palais in Paris, France. The summit marked a shift from previous gatherings by emphasizing actionable initiatives and global inclusivity in AI development.

Conference *summary and declarations* here.

The **Paris AI Action Summit** faced significant criticism from the AI safety community. Experts expressed disappointment that the summit prioritized economic growth and innovation over addressing the pressing safety risks associated with advanced AI systems.

Notably, the **United States** and the **United Kingdom** did not sign the

summit's final declaration, which emphasized "inclusive and sustainable AI." The U.K. cited concerns that the declaration lacked practical clarity on global governance and did not sufficiently address national security issues. The U.S., represented by Vice President JD Vance, warned that excessive regulation could stifle innovation, suggesting that the summit's approach might hinder technological advancement.

Max Tegmark, the MIT physicist and president of the Future of Life Institute, was among the prominent voices expressing concern over the summit's outcomes. He described the event as a "real bellyflop." He went on to say, "It almost felt like they were trying to undo Bletchley." Tegmark emphasized that the summit failed to build upon the safety commitments made at previous gatherings, such as the Bletchley Park Summit, which had focused on mitigating AI risks. He warned that neglecting AI safety in favor of rapid innovation could lead to uncontrolled developments with potentially catastrophic consequences. ***Recap here***

This was disappointing because it seemed as if momentum for reasonable regulation had been building following Blatchley Park and Seoul. Indeed, those previous conferences had commissioned the first-ever formal AI Safety Report, which was prepared by Yoshua Bengio and delivered to the conference. You can find a link to the report here: ***First AI Safety Report***

The Singapore Conference

The Singapore Conference on AI: International Scientific Exchange was a one-day conference organized on April 26, 2025, in response to the perceived shortcomings of the Paris AI Action Summit. The Paris Summit, co-chaired by France and India, was criticized by AI safety experts (e.g., Dario Amodei and Max Tegmark) for de-emphasizing safety in favor of economic opportunities and inclusivity, with the US and UK refusing to sign the Statement on Inclusive and Sustainable AI due to its lack of practical clarity on governance and safety. This fragmentation, coupled with the summit's shift away from the safety focus of the Bletchley Park (2023) and Seoul (2024) summits, prompted Singapore to refocus global efforts on technical AI safety research.

The Singapore government, leveraging its role as a neutral tech hub, acted swiftly after Paris to address the "diplomatic setback" and safety concerns, as noted in reports, with the conference's outcomes published in the **Singapore Consensus** on May 7, 2025. While

Singapore's AI governance initiatives (e.g., the AI Safety Red Teaming Challenge, November 2024) predate Paris, the April 26th conference was a direct attempt to salvage the damage done by Paris's diluted safety agenda and bridge East-West divides before the next summit in India in November 2025.

In his remarks, Tegmark stated, "After its near-death experience in Paris, international AI safety collaboration is roaring back."

It seems as if the conference successfully got AI Safety back on the top of the agenda with this document being published on May 7th: ***The Singapore Consensus on Global AI Safety Research Priorities***

For more on the Singapore Conference: ***Singapore Declarations and Summary***

The US Approach to AI Regulation

On January 23, 2025, President Donald Trump signed **Executive Order 14179**, titled *"Removing Barriers to American Leadership in Artificial Intelligence."* This directive marked a significant shift in U.S. AI policy, emphasizing deregulation and rapid innovation to bolster America's global AI dominance. The key objectives of the policy were to revoke the previous administration's policy and request that the new AI Czar, David Sacks, put in place an AI strategy by July 22, 2025.

In sum, Executive Order 14179 underscores the Trump administration's commitment to reasserting American leadership in AI by fostering an environment conducive to innovation and reducing regulatory barriers. While this approach aims to enhance economic and national security interests, it also raises questions about the balance between rapid technological advancement and the need for ethical oversight.

The United States played a significant supporting role in the Singapore Conference and the Singapore Consensus, contributing through researchers (Tegmark, MIT, Stanford), industry leaders (OpenAI, Anthropic), and the U.S. AI Safety Institute's collaboration with Singapore's Digital Trust Centre.

These efforts shaped the Consensus's technical framework, focusing on risk assessment, trustworthy systems, and post-deployment control, building on the International AI Safety Report. However, the U.S. government, under the Trump administration and David Sacks, did not lead or formally endorse the Consensus, reflecting EO 14179's deregulatory stance and the U.S.'s withdrawal from Paris's

cooperative framework. Singapore's neutral mediation enabled U.S.-China dialogue, but U.S. influence was limited by its unilateral policies, setting a complex stage for the India AI Summit (November 2025).

The India Summit Scheduled in November 2025

The India AI Summit is the next major AI conference scheduled, expected to occur between November 2025 and January 2026, (most likely November 2025 based on announcements at the Paris AI Action Summit). To be hosted by Indian Prime Minister Modi and organized by the Ministry of Electronics and Information Technology (MeitY).

Status of Governmental AI Regulation

Before conducting this research, we had not fully understood that so much work has been done in terms of attempting to regulate AI, both from a civic society vantage point, like the *Future of Life Institute*, and from governmental action.

The Trump administration and its new AI Czar, David Sacks, plan to come up with a new AI strategy for the U.S. by the end of July. We await the final definition of the new strategy once Mr. Sacks's report is prepared. It isn't clear whether the India Summit will follow in the footsteps of Singapore or Paris.

Geo-Political Concerns

It's apparent that China does not like that the US has put restrictions on the kinds of chips that can be sold to China, as the former CEO of Google, Eric Schmidt, stated in this TED Talk, which we highly recommend.

Schmidt also highlights the geopolitical risks associated with any near-term breakthrough toward AGI. After explaining the Chinese annoyance with the export restrictions, he describes a frightening scenario:

"Let's assume I am the bad guy, and you are the good guy. You're six months ahead of me and we are both on the same path to superintelligence. And you are going to get there, right? And I am sure you are going to get there. You are that close. And I'm six months behind. Pretty good, right? Sounds pretty good. No!! These are network-effect businesses. And in network-effect businesses, it is the slope of your improvement that determines everything. So, I'll use OpenAI or Gemini, they have 1,000 programmers. They are in the process of creating a million software AI programmers. What does that do? First, you don't have to feed that except electricity. So that's good. And they don't quit and things like that. Second, the slope is like this. Well, as we get closer to superintelligence, the slope gets even steeper. If you get their first, you dastardly person, I will not be able to catch you. And I have given you the tools to reinvent the world and, in particular, to destroy me. That is how my brain Mr. Evil is going to think. So, what is the www *am I going to do?*

"The first thing I am going to try to do is to steal all of your code. You have prevented that because you are good. Second, I am going to infiltrate you with humans. Well, you have good protections against that as well. You don't have spies. So what do I do? I am going to go in and change your model. I am going to modify it. I am going to screw you up to get me, so I am one day ahead of you. And you're so good I can't do that. What's my next choice? Bomb your data center."

Existential and Geo-Political Concerns

In chapter 3, we referred to the Economist interview of Sir Demis Hassibis (born 1976), the head of Google AI, and Dario Amodei (born 1983), the founder of Anthropic, where we shared their views on the timing of AGI, here we want to share with you their concerns and foreboding as it relates to the lack of regulation of AI.

Lord Hassibis explains that civil society does not fully understand how fundamentally AGI/ASI will impact our world. He hopes that as people more fully understand how positive these changes can be a sort of CERN for AGI will be created where the scientific world collaborates on the last few steps to achieve AGI by generating global norms. For more on ***CERN***.

Dario Amodei then explains why Paris was a missed opportunity:

"We are on the eve of something that has great challenges, it is going to greatly upend the balance of power. If someone dropped a new country into the world, ten million people smarter than any human alive today, you would ask them the question, what is their intent? ***What are they actually going to do in the world, particularly if they are able to act autonomously.*** *I think it is unfortunate that there wasn't more discussion of those issues as there was in the early summit that the UK held, the one in 2023 in Bletchley Park. I hope that future summits reclaim this mantle."*

When asked to comment on potential over-regulation in Europe and under-regulation in the US, Lord Hassabis suggests that the right answer is likely somewhere in the middle, but always being sure to balance against the two types of major risks he sees, one being bad actors who use these powerful technologies for nefarious means and:

"the other being from agentic systems themselves getting out of control or not having the right values or the right goals both of those things are critical to get right, and I think the whole world needs to focus on that."

They are asked on a personal level, do they ever worry about ending up like Robert Oppenheimer?

Lord Hassabis did not hesitate to respond:

"I worry about those kinds of scenarios all the time. That is why I don't sleep very much. I mean ,there is a huge amount of responsibility on the people, probably too much, leading this technology. I think that is why us and others are advocating for we probably need new institutions to be built to help govern some of this. You know I talked about CERN, I think we need a kind of equivalent IAEA Atomic Agency to monitor sensible projects and those that are more risk-taking. I think that society needs to think about what kind of governing bodies are needed. Ideally, it would be something like the UN, but given the geo-political complexities, that does not seem very possible. So, I worry about that all the time. We just try to do on our side everything we can in the vicinity and influence that we have."

Dario agreed with his contemporary's comments:

"My thoughts exactly echo Demis. It feels like every decision I make is balanced on the edge of a knife."

These leaders are calling for our global civic society to help them think through and manage this unprecedented moment in history. That is the purpose of this book: to generate a wider global civic society discussion around these momentous technologies and their potential for positive and negative impact on the world.

In closing this chapter, we are going to leave the reader with two excellent TED Talks from important actors who believe we should effectively pause the development of AGI.

We highly recommend that readers watch these videos, where the founder of the Future of Life Institute at MIT, Max Tegmark, and technologist Tristan Harris both make the case for holding back on the development of AGI until we can be more certain our interests are aligned with the digital species we are creating.

How to Keep AI Under Control | Max Tegmark | TED

Why AI Is Our Ultimate Test and Greatest Invitation | Tristan Harris | TED

#12 Expanding Situational Awareness

Expanding Situational Awareness to Address the Urgent Need for a Global Conversation on the Future of Humanity

Our goal in writing this book is to provide the average citizen with a book, a primer, that helps sort through the artificial intelligence noise—and to understand just how profoundly transformative AI technology is, so that we may all begin to develop situational awareness of it.

We also want to build a broader definition in the 21st century of "situational awareness," where we understand how critical this moment is to the future of humanity. What has become crystal clear to us is that we now have the technological capability to create a post-scarcity world, one without poverty, hunger, and even the need to work.

What is also mega-clear is that AI, or this new digital species with unprecedented power for good and evil, is being thrust into a divided world where 195 competing nation-states are armed to the teeth, spending $2.7 trillion a year preparing for war. We're living in a world that is standing at the brink of self-destruction—remember, it's 89 seconds to midnight—because of our inability to get along.

Rather than prepare for the Deep Utopia now available to us, our corporate and governmental leaders are locked in a ferocious competition to be first. A competition that puts all of us at risk *vis-à-vis* AI. We have become our own greatest threat to survival.

In his seminal 2003 book, *Our Final Century*, Sir Martin Rees told us that we are living at the most important time in the history of the universe since the Big Bang. We thought that Sir Martin's statement was true in 2003, and that it is even more true today. The next decade will be make-or-break for humanity. Which way we go will depend on the decisions we make in the coming years.

Our hope with this book is to launch an urgent, broad-based conversation dealing with the future of humanity; a conversation that takes place in every segment of our global society and on every corner of our small planet. A conversation that we hope will point the way forward for those of us alive— and for the trillions yet to be born.

Situational Awareness as it Relates to the AI Challenge

From Leopold Aschenbrenner's groundbreaking paper:

> "Everyone is now talking about AI, but few have the faintest glimmer of what is about to hit them. Nvidia analysts still think 2024 might be close to the peak. Mainstream pundits are stuck on the willful blindness of 'it's just predicting the next word.' They see only hype and business-as-usual; at most, they entertain another internet-scale technological change.
>
> **"Before long, the world will wake up. But right now, there are perhaps a few hundred people, most of them in San Francisco and the AI labs, that have situational awareness.** Through whatever peculiar forces of fate, I have found myself amongst them. A few years ago, these people were derided as crazy—but they trusted the trendlines, which allowed them to correctly predict the AI advances of the past few years. Whether these people are also right about the next few years remains to be seen. But these are very smart people—the smartest people I have ever met—and they are the ones building this technology. Perhaps they will be an odd footnote in history, or perhaps they will go down in history like Szilard and Oppenheimer and Teller. If they are seeing the future even close to correctly, we are in for a wild ride."

In this book, our goal is to help those in global civil society understand that:

1. The trendlines for AGI and ASI have been accelerated by two to three decades and are coming at us faster than any other previous technology.

2. The technology we are creating will be far and away more transformative than any other invention in human history.

Analysts struggle to find ways to compare AI with past inventions, such as the printing press or the internet. **What most do not understand is that *there is nothing to compare it to*.**

The printing press and the internet do not, and will never have, agency; they will never put our very existence in question. We are creating, for the first time in human history, a digital *species* with intelligence that will greatly exceed that of human beings, and the big worry is that AI

will rapidly gain agency over us— and we will have no control over it.

Part of the problem comes from the lack of a formal definition for AGI (Artificial General Intelligence) and ASI (Artificial Super Intelligence).

"Artificial General Intelligence (AGI) is a type of artificial intelligence that can understand, learn, and apply knowledge across a broad range of tasks at a level equal to or surpassing human intelligence. Unlike narrow AI, which is specialized for specific tasks, AGI is designed to perform any intellectual task that a human can do, including reasoning, problem-solving, and adapting to new situations without needing task-specific programming, according to ***Amazon's AWS***.

AGI remains a theoretical concept and a major goal of AI research, with no consensus yet on how or when it will be achieved, ***according to Net Guru***.

Here again, the general discussion is misleading. Anyone who has used Generative AI knows that the AI models we are using are vastly more "knowledgeable" and "smarter" than ourselves. AI can read 1.2 billion books a year, and it will be 100 million times smarter than any average human. That is 100 million times better to read and absorb data.

The primary difference is autonomy. Generative AI are already way smarter than us. However, as of today, they still need a prompt to jump into action. Generative AI are not yet autonomous or strategic, and as of today, they have no agency independent of us. The most prevalent expectations for the arrival of AGI range from 2026 to 2028.

Indeed, the timelines keep moving up. In a recent interview, Elon Musk just predicted ASI for this year or next. He defines ASI as being smarter than any human at anything in this ***brief video***.

We believe the best definition of the kind of intelligence we are creating is included in Dario Amodei's paper, *Machines of Loving Grace,* where he explains that powerful AI will provide the capacity to put a **"country of geniuses in a datacenter."**

The fundamental question facing humanity today, as it relates to AI, is: Do we want this digital species to have autonomy? The leading actors in the field tell us there is a ten to twenty percent *existential risk* to humanity that the AI simply takes over and gets rid of human beings or keeps us as pets. That was never a worry with the printing press or the internet. There are a whole host of additional threats that arise from the proliferation of powerful AI, particularly if it gets into the wrong hands or if the robots we are creating are used for war.

Situational Awareness Expanded

As we indicated above, our goal in this book is to understand the AI revolution and then situate that revolution in today's increasingly complex world to expand our situational awareness to find a pathway or pathways to achieve Sustainable Super Abundance.

What is clear to us is: we can no longer deal with the civilizational threats facing humanity in isolation as we explained in Chapter 10.

In some ways or many ways, the ChatGPT moment has poured gasoline on a human fire that was already on the brink of self-annihilation. The big difference here is that the nature of this gasoline is fundamentally different in that it has the potential to both exacerbate our risks and to go a long way toward solving our problems.

The OpenAI launch of ChatGPT, on November 30, 2022, transformed our destiny forever. That commercial launch unleashed a ferocious competition between trillion-dollar companies bound and determined to become the first to develop AGI, as whichever gets there second may no longer be able to compete. The one attempt to call for a six-month pause was swept aside as the AI *tsunami* hurdled past the call for a pause as if it did not exist.

At the same time, this tremendous technological power has been thrust onto a divided world at a time when our post-World War Two institutions are fraying and the big power rivalry between the U.S. and China has been accelerating. This, at a time when the group created by Oppenheimer and Einstein tells us that we are closer to self-annihilation than ever before, or 89 seconds from midnight.

The ferocious, no-holds-barred, corporate and geo-political competition is putting all of us at risk. The more we fight against one another, the greater the chances of AI taking over and asking, among their AI selves, *why keep these insidious, stupid beings with us? They have never built up a wisdom commensurate with their technology despite repeated calls from global thought leaders. Why keep them around other than possibly as pets?*

What has become clear to us is that the Greek Tragedy unfolding here could be the last act in the human project if we do not act swiftly and comprehensively to address our civilizational threats—all of them—as well as work together to build the post-scarcity world now available to us.

The Chat GPT movement unleashed a ferocious competition among trillion-dollar companies to become the first to attain AGI in a capitalist

world, where the principal actors are begging for us to put in place rules of the game— but no one is listening.

At the same time, the leading AI countries—China and the US— are in a pitched battle for global supremacy with nuclear tensions spreading across the planet as we move further away from the end of World War Two. That was the last globally calamitous event, which led to the creation of our global institutions.

The challenge is clear: if humanity can unite and use the technology *tsunami* for the good of all humanity, and at the same time overcome our other civilizational threats, we can create a world of Sustainable Super Abundance, and the trillions of human beings yet to be born will see the light of day

If, however, humanity is unable to unite and create a new set of global institutions ready to take on today's challenges, there is a very high probability that we will prove the creators of the Great Filter Theory correct and become just the next advancing civilization to self-destruct because we were unable to build a wisdom commensurate with our technology.

In the next section of this chapter, we will speak to the need to educate the 1.4 billion students studying in K-12 institutions across the world, and the with 100 million educators working with them. One of the big challenges facing humanity today is that education is a nation-state enterprise in a world in desperate need of global competence and a sense of what it means to be a global citizen.

In the last part of this chapter, we will address the need for the whole of humanity to build a situational awareness that goes beyond the AI challenge and beyond the climate change challenge to embrace not just our civilizational challenges but to also discuss how we can navigate past those challenges to create the world of Sustainable Super Abundance. Technology has made it available to us if, if, *if* we can only learn how to get along with one another.

The Education Solution

Let's start with K-12 Education and the 1.4 billion students studying today.

In this chapter, we will first look to educators from across the world to help us generate the conversation and set the stage for this world of abundance, then invite everyone to join the fun. One leg of the

pathway to abundance will need to come from the K-12 and university education sectors.

UNESCO tells us that in today's world, there are currently between 1.3 billion and 1.4 billion primary and secondary school students being guided by just under 100 million educators. It seems to us that this critical segment of our global population is all too often forgotten.

This audience is critical for obvious reasons, K12 students are our future. Providing K12 educators and students with a roadmap to Sustainable Super Abundance is essential to the success of this effort.

If we are going to usher in a world of abundance, we are going to need to do more than publish well-written, well-researched books at the university level and pretend that's how we are going to overcome the challenges in front of us. As UNESCO stated in its November 2021 report, *Reimagining Our Futures Together: A New Social Contract for Education,* **education itself needs to be transformed and that starts with K12 education**.

We wholeheartedly agree that education needs to be transformed at the K-12 and university levels. We need a driving mission, a challenge that unites educators across the world. We hope that achieving Sustainable Super Abundance can be that mission.

We would propose that K-16 educators answer UNESCO's call for a new social contract for education by working together to build the world of Super Abundance that technology has made available to us.

Reimagining our Futures Together: A New Social Contract for Education (cited above) calls for a fundamental transformation of education systems worldwide to address the challenges facing humanity and our planet. The report proposes a new social contract for education that aims to unite us around collective endeavors and provide the knowledge and innovation needed to shape sustainable and peaceful futures for all.

The education sector has been slow to respond to today's societal needs. Indeed, we believe the big risk in education today is that far too many schools are educating our youth for a world that no longer exists. Rather than prepare K-12 students for a world of scarcity, we need to prepare them to live in a world of Sustainable Super Abundance, in a rapidly globalizing world of exponential technological change.

Part of the challenge in developing a global educational policy is that education is a national endeavor, where each of the 195 competing nation-states develops its own education policy. And while many

educators understand the need to build global competence, those educators continue to be too few and far between, particularly as most nation states do not provide any real incentive to educate across borders.

In our view, someone or some entity or entities are going to need to pick up the slack, as there is an inherent disconnect between nation-states that govern their education systems and the need to teach global citizenship. Over the last decade, One World has tried to fill the gap with different levels of success, as we will discuss in Chapter 14.

Educational leader John Dewey wrote in 1934 that "Any education is, in its forms and methods, is an outgrowth of the needs of the society in which it exists." In his day, he suggested that educational needs would vary from place to place. We believe the educational goals of today are consistent across the world but need to include a global perspective and a clear sense of the unique time we are living through, a clear goal that unites humanity.

One futurist who captures and understands today's educational challenge extremely well is James Martin, who's written more than a hundred books. His *The Wired Society* (1978) is one of the most influential. In 2006, he wrote another prescient book, *The Meaning of the 21st Century: A Vital Blueprint for Ensuring Our Future***.** We believe that this book, more than any other, captures the challenge of our day, even though it was written almost twenty years ago.

From Martin's *The Meaning of the 21st Century*:

> *"At the start of the 21st Century, humankind finds itself on a nonsustainable course— a course that, unless it is changed, will lead to catastrophes of awesome consequences. At the same time, we are unlocking formidable new capabilities that could lead to much more exciting lives and glorious civilizations.*
>
> *This could be humanity's last century, or it could be the century in which civilization sets sail toward a far more spectacular future. Decisions that will lead to these wildly different conclusions have to be made soon.*
>
> *They depend on our being able to understand the options for the 21st century, think logically about our future and collectively take rational action."*

This is exactly what we are attempting to do with this book and the conversation on the future of humanity we hope to generate:

- Make clear the options in front of us
- Think logically about those options and our future
- Encourage a collective approach to taking rational action

We are also calling on the *International Commission on Stratigraphy*, the world's leading authority for defining geological time, to classify our era as the **Anthropocene Age,** or the age in which the dominant actor on the planet Earth is no longer *Mother Nature* but rather *Homo Deus* or humankind.

James Martin explains why the focus on this classification is needed.

> *"Historically, evolution has been in nature's hands. Now, suddenly, it is largely in human hands, but we need to be cautious, using our scientific know-how as responsibly as possible.*
>
> *The job of today's young people, or the* ***Transition Generation,*** *will be to get humanity through the coming period of chaos, peril, and opportunity****. A massive transition is needed, and the agenda should be created for the generation that will bring about this transition.***
>
> *Much of what needs to be done is not happening. Today's computer models show that we are not adapting quickly enough because we are not thinking ahead.* ***There must be an absolute crusading determination to address the changes ahead.***
>
> ***Today's young people will collectively determine whether civilization survives or not.*** *We need to* ***give them a foundation for making wise choices by helping them understand humankind's likely possible, probable, and preferred futures. What can humanity become? Our future wealth will increasingly relate to knowledge in the broadest sense of the term."***

We agree with Martin that the number one goal of education today needs to be to provide K-12 youth with a foundation for making wise choices so that they can understand humankind's likely, possible,

probable, and preferred futures. Indeed, this is precisely the kind of global conversation we are attempting to create with this book. Our mission is to engage as many of the 100 million K12 educators and 1.4 billion K12 students in a discussion on Sustainable Super Abundance as possible.

We also believe that the University Sector, with its estimated 250 million students and 12 to 15 million educators, has a very important role to play in building a foundation for making wise choices. We are convinced that UNESCO's call for a global conversation on education needs to be a K-16 conversation.

If the CEO of IBM is having difficulty keeping up with the pace of change, how can we expect the average third-grade teacher to keep up? Rather than distance themselves from K-12 educators, universities should embrace them. They should host events at their universities celebrating K-12 educators and students who engage with the real world in a meaningful and impactful way.

Getting Through the Canyon with as Little Mayhem as Possible

We have attempted to make clear in this book that humanity stands at a crossroads and that, at a time of accelerating technological change, we need to find ways to come together across the world. Martin uses a wonderful metaphor of a whitewater raft hurdling downstream at an accelerating rate of speed:

> *"We are travelling at breakneck speed into an age of extremes— extremes in wealth and poverty, extremes in technology and the experiments that scientists want to perform, extreme forces of globalism, weapons of mass destruction, and terrorists acting in the name of religion.* ***If we are to survive decently, we have to learn to manage this situation****.*
>
> ***The message of this book is vitally important. We have reached a situation where grand-scale decisions have to be made****.* ***Collectively, humankind needs to be taught about the future so that it can understand these decisions.*** *Rules of the road need to be put in place. As I will explain, the 21st century is very critical. It is a make-or-break century."*

Martin calls on the young people of today to become the much-needed "Transition Generation" to help humanity survive through the current challenges to a new creative era of history. He writes,

> *"Think of the 21st century as a deep river with a narrow bottleneck at the center. Think of humanity as river rafters heading downstream. As we head into the canyon, we'll have to cope with a rate of change that becomes more intense— a white-water raft trip with the currents becoming much faster and rougher—a time when technology will accelerate at a phenomenal rate. As the world's population grows, global tensions and pollution will climb, and the danger of massive famines will increase.*
>
> *The decades in which we are swept toward the canyon bottleneck will be at a time when we unlock extraordinary new technology–nanotechnology, biotechnology, extreme-bandwidth networks, regenerative medicine, robotic factories, and intense forms of computerized intelligence.*
>
> *The job of the Transition Generation is to get humanity through the canyon with as little mayhem as possible into what we hope will be smoother waters beyond."*

Our goal in writing this book is to help a new generation of youth navigate past the turbulent waters ahead to a post-scarcity world of abundance by providing them with a foundation for making wise choices. If humanity is going to successfully navigate these waters, we are going to need to apply an exponential mindset and understand that the intelligence explosion is upon us and that AGI/ASI are right around the corner.

If we can find ways to align human values with the new digital species, we can create a world without poverty, hunger, or the need to work. If, on the other hand, we fall short, the existential risks we are facing will confirm the Great Filter Theory rather than disprove it.

There can be no doubt that, as Martin forecast in 2006, humanity is at a crossroads and the decisions we make in the coming years will not only determine our fate but the fate of trillions of human beings yet to

be born.

If we are going to successfully navigate this river, we are going to need educators and humanity to unite as never before, not only to avoid our next and last war but also to solve climate change and avoid the potential downsides of AI.

Broadening the Conversation Beyond Educators

While the K-12 education sector does provide one important pathway forward, we are going to need a broader society-wide conversation if we are going to succeed in achieving a post-scarcity world. Parents, citizen-leaders, public officials, business leaders, artists, and educators are all going to need to be involved in this civilizational conversation.

In Chapter One, we shared Yuval Noah Harari's comment:

> "...for the first time adults are not really good guides for the future of these kids. We for the first time have no idea what the world will look like when the children of today are they themselves grown up. Nobody knows what the economy or the job market or what society will look like in 10 or 20 years."

While we completely agree with Mr. Harari's conclusion, we would argue that today's uncertainty relative to the future of humanity makes the role of parents, civic leaders, business leaders, and educators even more important.

The one thing we can provide our youth that we know has the potential to last a lifetime is a strong set of values. Indeed, we are also going to need to align our AI around a strong set of agreed-upon humanist values. Generating those values will require conversations at many different levels.

Indeed, those of us above a certain age will be the last generation to have grown up in the *old world*, without cell phones, personal computers, where there were no flickers, and you had to get up to turn the channel on the TV to see one of three stations on VHF. We believe the older generation has something special to share with our youth—memories of a simpler, less-complex world.

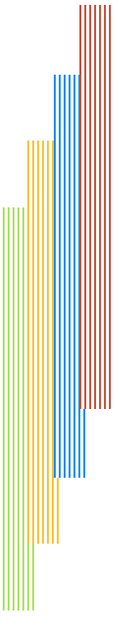

At the same time, if you are like most parents, you would not be at fault for believing our world is upside down in that our sons and daughters

know how to use the most powerful tool in the world better than us—technology. There has never been a more important time for different generations to come together and agree on a set of values that we can align with the AI extant.

Broadening our Perspective: Why is Bill Gates Only One-Third Correct?

To be clear, we are going to need to do more than broaden our outreach; we are also going to need to broaden our perspective. For example, why is Bill Gates only 25% correct in the assertion he makes below? Following the publication of his book, *How to Avoid a Climate Disaster*, the Microsoft founder was interviewed by Anderson Cooper on *60 Minutes*. Here are some excerpts:

> *Anderson Cooper: You believe this is the toughest challenge humanity has ever faced?*
>
> *Bill Gates: Absolutely, the amount of change and new ideas is greater than the pandemic, and it needs a level of cooperation that would be unprecedented.*
>
> *AC: It just seems overwhelming if every aspect of our daily life has to change.*
>
> *BG: It can seem overwhelming…if people think it's easy, they are wrong. If people think it is impossible, they are wrong.* ***It's possible, but it will be the most amazing thing mankind has ever done."***

We understand that at the time, Gates, a philanthropist and one of the world's richest (and most knowledgeable) men, was focused on solving climate change, and he was only responding to a question, but it does seem clear to us that we need to broaden the X-Risk conversation as well as include the upside conversation relative to a world of Sustainable Super Abundance.

If we think about it for a minute, while climate change is a civilizational threat, it seems far less immediate at this point than the possible impending doom emanating from AI. And as it relates to nuclear risk, it is hard to tell when that might happen. However, in just the past few weeks, we have had tensions break out between India and Pakistan; between Israel, the US, and Iran; the ongoing war in Ukraine; and worry over North Korea.

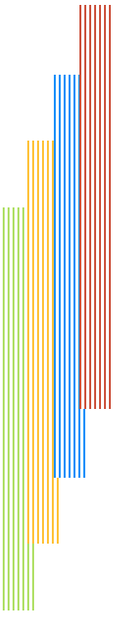

Humanity has been operating at the edge of nuclear disaster since the outbreak of Russia's invasion of Ukraine, causing UN Secretary General Antonio Guterres to declare on August 1, 2022, that

"Today, humanity is just one misunderstanding, one miscalculation away from nuclear annihilation. We have been extraordinarily lucky so far. But luck is not a strategy. Nor is it a shield from geopolitical tensions boiling over into nuclear conflict."

Humanity needs a strategy to ensure its future. What good does it do to solve the climate change challenge if the world is blown to smithereens the next day?

Broadening the conversation does not mean we should in any way stop talking about climate change; in fact, just the opposite: we need to make clear to everyone that our existential challenges are growing rather than shrinking, and the key to solving them all starts with finding ways to bring humanity together.

In this book, we suggest bringing the potential for a world of Sustainable Super Abundance into the discussion. The pathway to a world without poverty, hunger, and the need to work also requires that humanity unite in ways we have not yet done.

The Future of Humanity Conversation

Because we are living at the most important time in human history—at the hinge of history—and given the accelerating opportunity and challenge coming from AI, we believe we need to engineer an urgent *future of humanity* conversation across the globe.

The whole reason global leaders have taken on the climate change conversation is because they are concerned that climate change imperils the future of our human civilization. They are concerned that if we do not change course now, our human civilization will be severely disrupted.

It seems clear to us that if we are really and truly concerned about the future of human civilization, we need to extend the climate change conversation to include nuclear proliferation and the now imminent threat posed by AI. As we discussed in Chapter 10 we need to extend the sustainability to go from just talking about climate change to talking about human survival and the potential for Sustainable Super Abundance.

We are therefore urging civic, public, academic, and business leaders to

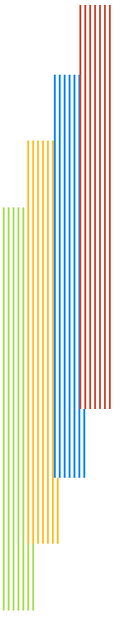

broaden this conversation at a time when our species is at a crossroads.

Returning to James Martin's prescient 2006 book, *The Meaning of the 21st Century: A Vital Blueprint for Ensuring our Future,* he made clear the challenge of our day and explained that:

> *"Much of what needs to be done is not happening. Today's computer models show that we are not adapting quickly enough because we are not thinking ahead.* ***There must be an absolute crusading determination to address the changes ahead.***

The only change we would make to the above statement of almost twenty years ago is that, since we have done little to heed his call in the last two decades, the need to think ahead is greater today than 20 years ago and therefore the crusading determination needed to effect the change will have to be even more urgent because of lost time. No one should have any doubt that this is our last chance to address the change that Carl Sagan saw in the 1990s and James Martin wrote about in 2006.

The purpose of this book is to first provide readers with a foundation to understand the challenges and opportunities in front of us and second serve as a catalyst to generating an urgent conversation on the future of humanity.

This Global Conversation on the Future of Humanity is Urgent and Essential to our Survival - There is No Plan B

In thinking about the urgency of this conversation relative to the future of humanity, it is worth contemplating a couple of important points.

First, it seems that our far too many political leaders act only *after* the catastrophe has hit— be it financial prudence or other more serious risks. For example, according to *this article* in *Business Insider*, following the Washington D.C. airplane crash that took 67 lives, there is a concept in the Air Safety Industry called **"blood priority."** *In aviation safety, this means that needed changes will not be implemented until after a major accident with deaths, despite US safety officials pushing to be more proactive with safety measures.*

In thinking about existential risks, we do not have the luxury of blood priority, because there is no way we survive a nuclear war, an AI takeover, or advanced climate change. There is no Plan B. We only get one shot at preserving our human civilization. Now is the time; no one

should doubt that we are living at the hinge of history.

We are of the view that if we can keep civilization intact for the next 25 years, proactively navigating past the kinds of catastrophes we see on the horizon, every human being alive in 2050 will live better than the billionaires of today.

#13: Finding Purpose

> Alice to the Cheshire Cat in Lewis Carroll's Alice's Adventures in Wonderland (1865):
>
> *"Would you tell me, please, which way I ought to go from here?"*
>
> *"That depends a good deal on where you want to get to," said the Cat.*
>
> *"I don't much care where—" said Alice.*
>
> *"Then it doesn't matter which way you go," said the Cat.*
>
> *"—so long as I get somewhere," Alice added as an explanation.*
>
> *"Oh, you're sure to do that," said the Cat, "if you only walk long enough."*
>
> *(Alice's Adventures in Wonderland, Chapter VI)*

Project 2050: From the Hinge of History to Sustainable Super Abundance

• Step #1: Creating a Global Coalition for Sustainable Super Abundance

• Step #2: Generate a Global Conversation on the Future of Humanity

• Step #3: Build a 21st Century Renaissance Mindset that Takes us to the Stars

Project 2050: From the Hinge of History to Sustainable Super Abundance

Creating a Global Coalition for Sustainable Super Abundance

Unlike Alice, we believe it is time for humanity to decide on a direction and purpose. With Project 2050 we are proposing we create a goal of developing a post-scarcity world of Super Abundance by the year 2050, a world where we will have eliminated poverty, hunger, the NEED to work as well as solve climate change, bring an end to cancer and many other forms of sickness greatly extending our healthy life spans if not more.

We propose to do this by generating a multi-year, multi-level, global, conversation on the future of humanity. We are looking to extend

humanity's global conversation on sustainability to include all forms of civilizational risk as well as the extraordinary opportunity in front of us. For the first time in human history, we have the potential to build a world of Sustainable Super Abundance where no one needs to work to secure needed goods and services. We should extend our global conversation on the future of humanity to include both the unprecedented opportunity and all the risks facing us.

We are confident that ***if, if, if*** we can keep our global civilization intact between now and 2050, everyone alive then will live better than the billionaires of today, in the same way those of us alive today live better than the kings and queens of England past.

However, we want to make clear, as of today it is not at all clear that we will be able to keep our global civilization intact. Our goal in writing this book was to share our view on the state of the world as it stands today relative to our biggest opportunities and challenges – to create an expanded situational awareness – with a view to finding ways to develop a coherent global response.

We were surprised, if not shocked to learn, in November 2023 that the leading futurists believe we can build a world of Sustainable Super Abundance. Our belief in our technological capacity to build a post-scarcity world of high universal income eliminating poverty, hunger and even the need to work has been confirmed time and again since November 2023.

While the technological vision is clear, our human wisdom has consistently been left wanting. Homo Technicus has been extremely successful, but we have yet to prove we can be true Homo Sapiens, or a species full of a knowledge and wisdom commensurate with our technology.

We understand that the futurists are expecting a technological singularity leading up to an event horizon beyond which we cannot see. This is our last chance to prepare for that eventuality. We believe we should make every effort to work together to shape and align that event as best we can so that it incorporates the best of us. If we really are going to experience a rupture in the fabric of human history, we should work hard to prepare that moment and shape it as best we can.

Our 13.8 billion-year Journey

In Chapter Ten, we showed how Homo Technicus radically transformed the previous 200,000 years of human history building a world around us in ways that would have been unimaginable 250 years ago.

Sir Martin Rees, Britain's astronomer royal, has told us that we are living at "the most crucial location in space and time" since the Big Bang. To put this comment in perspective, we use the lens of Big History as follows:

- Big Bang: 13 years, 8 months ago (~13.8 Gya)
- First Stars Appear: 13 years, 2 months ago (~13.2 Gya)
- New Elements Forged in Dying Large Stars: 13 years, 2 months ago to today
- Our Sun and Solar System Form: 4 years, 6 months ago (~4.57 Gya)
- Formation of Earth: 4 years, 6 months ago (~4.54 Gya)
- Earliest Life on Earth: 3 years, 9 months ago (~3.8–3.5 Gya)
- **Humans/Homo Sapiens Appear: 66.67 minutes ago (~200,000 years ago)**
- End Ice Age, Beginning of Holocene, Earliest Farming: 5 minutes ago (~10,000 years ago)
- Industrial Revolution Begins: 7.5 seconds ago (~250 years ago, 1775 CE)
- Modern Era Advances (1900 CE): 3.75 seconds ago (~125 years ago)
- Splitting of the Atom: 2.25 seconds ago (~75 years ago, 1945 CE)
- ChatGPT Launched: 0.075 seconds ago (~2.5 years ago, November 30, 2022)

As we can see, from this perspective, homo sapiens arrived just over one hour ago, and the straight-line chart depicted by Professor MacAskill appeared 7.5 seconds ago.

We think it is amazing that the transformation that began less than eight seconds ago has been so powerful and profound. In little more than the blink of an eye in Big History time, Homo Technicus has transformed our material world, gaining a power over nature that is, in George Kennon's words, out of all proportion to our moral strength.

The fundamental question in front of us today is, will humanity be able to build and develop a moral strength commensurate with our technology so that we can push forward and develop the world of Sustainable Super Abundance now available to us?

Looked at from the Big History time perspective, Toby Ord's characterization seems more on point, and therefore worth repeating:

> *"Like an adolescent who is given more freedom and responsibility as they grow older, humanity has gained more power and technological advancement as our societies have progressed. However, we have not necessarily gained the wisdom to use this power in a responsible and beneficial way.*
>
> *"…**If all goes well, human history is just beginning. Our species could survive for billions of years – enough time to end disease, poverty, and injustice, and to flourish in ways unimaginable today.***
>
> ***"But this vast future is at risk.** With the advent of nuclear weapons, humanity entered a new age, where we face existential catastrophes – those from which we could never come back. Since then, these dangers have only multiplied, from climate change to engineered pathogens and artificial intelligence. **If we do not act fast to reach a place of safety, it will soon be too late."***

We hope we have made clear our current situational awareness. Our time to act is running out.

One concern we have is that the dramatic advances in science and technology have created a certain human hubris that blinds us to our unique positions in the universe.

Science has greatly expanded our knowledge of the universe, stretching time back 13.8 billion years ago and our universe 93 billion light years across, but we still can't answer the most basic of questions relative to ***Time and Place***.

We still do not have a theory of everything…unless it is ***Turtles All the Way Down***.

So, while science has greatly expanded our base of knowledge, we still cannot answer the most basic questions of life without resorting to philosophy or religion.

Philosophy and religion begin where science ends. One of the challenges that will arise if we are successful in building a world of Sustainable Super Abundance is that many of us will need to redefine our life purpose as it will be much less associated with the work one does, so why not start now?

In fact, at the beginning of this interview, Nick Bostrom laughingly explains that if you are looking for purpose, now is the time to act because the decisions we make today will last for billions of years. If you listen ***from 0:28 to 0:45, he explains***,

> "We are seemingly approaching this critical juncture in human history, upon which, like the next billion years might. . . **If you want real purpose, knock yourself out now, like it's never going to be more at stake than there is right now and in the next few years."**

As we hope we have made clear, we wholeheartedly agree with this view as we are truly living at the hinge of history. We hope that concerned citizens from around the world will join the Global Coalition for Sustainable Super Abundance as we conduct a multi-level, multi-country conversation on the future of humanity as we search for answers, meaning, and a new story that brings us together.

One of our favorite life-meaning quotes comes from Fannie Flagg's novel *Welcome to the World, Baby Girl!* Flagg's character, Aunt Elner, comments:

> "Poor little old human beings – they're jerked into this world without having any idea where they came from or what it is they are supposed to do, or how long they have to do it in. Or where they are gonna wind up after that. But bless their hearts, most of them wake up every morning and keep on trying to make some sense out of it. Why, you can't help but love them, can you? I just wonder why more of them aren't as crazy as betsy bugs."

Fannie Flagg's meaning is simple. Every one of us needs to come up with their own life purpose, their own reason for being, their own *raison d'etr*e, their own *ikigai.*

We believe it is time for humanity to take collective responsibility for its future and begin to think about where we are headed. What is our purpose as a collective? What is the human project?

What do we want to do with our lives as individuals and as a collective?

Indeed, in his book *The Magic Mountain*, published 100 years ago in another time of great tumult for humanity, Thomas Mann wrote:

> "A man lives not only his personal life, as an individual, but also, consciously or unconsciously, the life of his epoch and his contemporaries."

Everyone alive today is living the life of our epoch, as we have explained tirelessly, perhaps the most important epoch in human history. How we respond to our epoch will impact the lives of trillions of people yet to be born.

We are recommending that we come together as humanity and design an informed, uplifting response to both the opportunities and challenges of our day.

As we think about it, it is rather surprising that humanity has not put forward a global purpose so far. How is it that after 12,000 years, we have made no real effort to come together as human beings and define a purpose for our collective welfare and the collective welfare of all living beings on the planet? At One World, we often say, technology is bringing us closer together, but no one is telling us how to get along. Maybe we should try and change that together.

We have religions and philosophies, but it does not seem as if we have a philosophy or religion that has been able to unite the whole of humanity. We still need a "story" that unites us across the world and humanity. Perhaps the closest philosophy we have that brings us all together is cosmopolitanism. That is, according to the Stanford Encyclopedia of Philosophy:

> "The word 'cosmopolitan,' which derives from the Greek word *kosmopolitēs* ('citizen of the world'), has been used to describe a wide variety of important views in moral and socio-political philosophy. **The nebulous core shared by all cosmopolitan views is the idea that all human beings, regardless of their political affiliation, are (or can and should be) citizens in a single community. "**

The Stanford Encyclopedia goes on to explain that different versions of cosmopolitanism envision this community in different ways. One vision of cosmopolitanism that fits with our goals to generate a global conversation on the future of humanity comes to us from philosopher Kwame Anthony Appiah, who in his 2006 book *Cosmopolitanism: Ethics in a World of Strangers* describes the belief as a moral and ethical framework that emphasizes global citizenship while balancing universal

human obligations and respect for legitimate cultural differences.

Appiah presents cosmopolitanism in a way that navigates ethical interactions in a globalized world filled with diverse cultures, beliefs, and practices. Appiah's cosmopolitanism is rooted in the idea that all humans are part of a single moral community, yet it acknowledges and values the diversity of individual and cultural identities.

Moreover, Appiah emphasizes "conversation" as a central mechanism for cosmopolitanism, both in the sense of literal dialogue and as a broader concept of coexistence and mutual engagement. He advocates for cross-cultural conversations that foster understanding without requiring agreement or convergence on a single way of life. This process involves learning about others' beliefs and practices while maintaining openness to new possibilities.

Indeed, Professor Appiah suggests that we all participate in "fallible" conversations— that is, conversations where we respectfully share our view of the world knowing that we might be right or wrong, but that the important part of the communication is to share one's view and understanding of our complex world, a world where we are each still trying to answer the most basic questions relative to the meaning of life.

We believe Professor Appiah's model conversation provides us with a good framework to develop our much-needed conversation on the future of humanity.

(We would be remiss if we did not mention at least two other efforts that might have been included in the discussion above. ***Hans Küng's Global Ethic Project*** and the ***Bahá'í Faith***. We are also confident that there are other such efforts that we have missed as well, and look forward to discussing these ideas and more in our eventual Future of Humanity Conversation.)

Building a Coalition for Sustainable Super Abundance

As we explained in Chapter Ten, the world is in desperate need of global platforms that bring humanity together. We created one such platform about twelve years ago called *One World United & Virtuous.* Our name was inspired by Benjamin Franklin, who in 1731 wrote:

> *"That the great affairs of the world, the wars, revolutions, etc., are carried on and effected by parties. That the view of these parties is their present general interest, or what they*

take to be such. That the different views of these different parties occasion all confusion. That while a party is carrying on a general design, each man has his particular private interest in view.

"That as soon as a party has gained its general point, each member becomes intent upon his particular interest; which, thwarting others, breaks that party into divisions, and occasions more confusion.

"That few in public affairs act from a mere view of the good of their country, whatever they may pretend; and tho' their actings bring real good to their country, yet men primarily consider that their own and their country's interest was united, and did not act from a principle of benevolence. ***That fewer still, in public affairs, act with a view to the good of mankind.***

There seems to me at present to be great occasion for raising an united Party for Virtue, *by forming the virtuous and good men of all nations into a regular body, to be govern'd by suitable good and wise rules, which good and wise men may probably be more unanimous in their obedience to, than common people are to common laws. I at present think, that whoever attempts this aright, and is well qualified, cannot fail of pleasing God, and of meeting with success."*

While Franklin never did find the time to create this united Party for Virtue, we thought the idea of bringing the men and women of integrity of our day together under a single banner to work for the good of humanity was quite a good idea. We also thought that if it was a good idea in 1731, it is an essential idea today.

Over the last twelve years, *One World United & Virtuous* has built up a series of educational enrichment programs designed to help K-12 educators prepare themselves and their students for the unprecedented opportunities and challenges of our 21st-century world. In 2024, we also created a for-profit entity called *One World Future Ready* in an effort to automate the One World enrichment programs and increase the chances of continuity in view of the founder's age.

For information on the: ***One World Education Companies***

As you will have seen from the author's note, the One World team is responsible for the content and publication of this online book. As

mentioned above, our goal in producing this book has been to provide civil society, educators and the general public with a clear understanding that humanity is at a crossroads and if we are going to successfully navigate past our civilizational risks and build the world of Sustainable Super Abundance now available to us we are going to need to unite in ways that have not happened so far.

Our Goals and Commitment to You

Build a Global Coalition for Sustainable Super Abundance

Our hope is that concerned citizens, educators at all levels, not-for-profits, business leaders, public officials and all segments of civil society will consider this book as an invitation to join a Global Coalition for Sustainable Super Abundance. We do not have any definitive answers but we are committed to generating respectful, fallible conversations that help us understand our increasingly complex world with a view to designing the ethos and institutions we need to guide us in coming decades.

As we indicated in Chapter Ten, we are launching an extremely ambitious, multi-level, global conversation on the future of humanity. We hope to get all levels of society involved in this discussion.

It could very well be that no one heeds our call. It could be that we are the lone voice out in the wilderness with no one paying attention. That is okay.

One of the things we do is teach character education at One World, also perseverance, resilience, integrity, and honesty. One of our favorite philosophers is the Stoic philosopher Epictetus, who encouraged us to focus on our role in life and what we can do, and not worry about what we cannot control. He reminded us that we can respond to any set of events in a way that is consistent with our highest self. He reminded us that we have a responsibility to act as one with the universe.

We know that it would be unconscionable for us to see the technology *tsunami* coming our way and to say nothing, or worse, to pretend it was not coming our way because it was more convenient and less likely to offend.

We believe the big risk in education today is that we are preparing our youth for a world that no longer exists. The reason is simple: given the pace of change, no one understands where our world will be in

2040. We see an urgent need for humanity to come together and collaborate to do everything in our power to shape that world. We hope that this book helps people understand where we are and where we might be if we can learn how to come together and collaborate.

Preliminary Launch of the Global Coalition for Sustainable Super Abundance

On April 11th, 2025, we introduced our campaign to create a world of Sustainable Super Abundance in Acapulco, Mexico, where we gave a presentation at the invitation of Green Party Deputy, Jorge Angli, at the historic Fort of San Diego Park.

Two months later, on June 4th, we introduced a draft version of this book and reiterated our call to begin to build a Global Coalition for Sustainable Abundance at our fourth annual World Environment Day climate change conference, bringing together students and educators from China, Brazil, Mexico, and the U.S. to share best practice in action-oriented climate change education.

Formal Global Campaign to Start in September 2025

Now that we are finishing this book, we intend to launch a formal promotional campaign in September of this year. In the meantime, we will be reaching out to potential like-minded partners who have an interest in joining this effort.

Getting from Here to There

In closing, we want to leave you with four thoughts relative to getting from here to there, first in practical terms and then more in terms of thinking about the skills and mindset we need to achieve Sustainable Super Abundance.

First, we want to make clear that we do not have definitive solutions. Our goal is to generate a conversation to understand our increasingly complex world, in the hope that together we can find pathways to the post-scarcity world of Sustainable Super Abundance now available to us.

The futurists have given us a sense of how we get to a world of Sustainable Super Abundance. Effectively, they will merge advancing AI

with new, improved robotic technologies so that over the next fifteen to twenty years, humanity will have created as many robots as there are human beings. Based on the AI of today and the advancing AI of tomorrow, the robots will do all the required work to produce all the goods and services needed for everyone.

The one topic none of the futurists has discussed in any detail is how to get from where we are today to the "solved world" of high universal income of tomorrow. How do we get for-profit, trillion-dollar companies locked in a race to be the first to the promised land to agree to focus on human safety as much as they focus on profits? How do we get the two major Superpowers to collaborate as opposed to being drawn into Graham Allison's Thucydides Trap?

What kind of regulation do we need for AI? What will happen at the next AI conference in India in November, if it happens then? What kind of AI policy will David Sacks come up with for the US and the world this month? What role for the United Nations, if any? Can a CERN-like entity be created for AI?

No one seems to have answered the most basic question of how we get from here to there, which is why we want to generate a multi-level, global conversation on the future of humanity.

To that end, we welcome your comments and feedback and will continue to update the themes in the book and hopefully look forward to working with you to generate the ideas we need to create the world of Sustainable Super Abundance now available to us.

Second,

It is clear to us that the radical transformations that have taken place in our world started with a mindset change, first in the Renaissance, then with the Scientific and Industrial Revolutions. What is clear to us is that humanity is in urgent need of a mindset change.

We are hopeful that we can work together to develop a new 21st Century Renaissance Mindset of the possible. A new mindset that takes us to the stars.

Third, we hope you will agree that we have attempted to deal thoughtfully with several complex issues that have been raised in this book. However, if we get right down to it, **our future as a species will depend on our capacity to find ways to unite across our planet Earth on our capacity to get along;** therefore, our answer is quite simple, just sometimes difficult to achieve.

As any parent of more than one child knows, we spend far too much time asking our children, Why can't you get along?" Why are you always fighting and bickering? If you have brothers and sisters, we are certain you have heard this refrain multiple times growing up. Fortunately, as they grow up, siblings find a way to get along with and love one another.

The real question for the 8.23 billion people on the planet Earth today is, what is my link or connection to everyone outside of my ZIP Code, outside of my nation? If we come to understand that we "share a common fate on a crowded planet," we have the potential for success, and if not, our prospects dim greatly. Which way we go depends on each one of us alive today. Hopefully, humanity too will grow up and we will begin to see each other as fellow earthlings locked together on a cosmic journey that will be launched from the Spaceship Earth. All we need to do to become a space-faring species travelling the galaxies is to learn how to get along in the same way we do with those closest to us.

To that end one goal for the future of humanity conversation we intend to generate would be to identify three things we can tell our children relative to our future together on the planet Earth. Three things that every parent will tell their children as they put them to bed at night to help foster their dreams as to what is possible here on Earth.

Fourth and finally, for our goal in the conversations we intend to generate as part of this process will be to encourage everyone to understand that **we are all Co-Creators of the universe** and that we all have a role to play in realizing the world of Sustainable Super Abundance now available to us.

We believe that in a world where no two snowflakes are alike, it is reasonable to believe that no two human beings are alike. We therefore like to believe that if every human being manages to live up to their best versions of themselves, as their highest loving self, we can create paradise on earth. The long arc of technology has provided those of us alive today with an opportunity to come together and create a post-scarcity, solved world that brings an end to poverty, hunger, and the need to work.

We believe that "solved world" is achievable if every one of us does their part to be a positive, loving, co-creator of the Universe. Co-Creators act as if there are equal amounts of good and evil in the world and that their next act will tip the balance one way or the other.

Our goal in the conversations we will build out as part of this process is

to encourage everyone to understand that we are all co-creators of the universe and that we all have a role to play in realizing the world of Sustainable Super Abundance now available to us. For those interested in reading more on co-creators, please click here: ***Co-Creators of the Universe***

#14: Project 2050: A Story of Hope and Global Unity

In 2025, humanity stands at a crossroads—the hinge of history—where our choices today will echo for generations. Across continents and time zones, 1.3–1.4 billion children pass the night like a torch of dreams, as Yukiko Yamaguchi's poem, *Night Passing*, reminds us. These dreams traverse oceans, whispering of a future where super abundance is not a distant fantasy, but a shared birthright.

Yet, while children dream of peace and prosperity, we see a world spending $2.7 trillion on weapons, driven by fear and competition. Max Tegmark warns of a "suicide race" between great powers like the U.S. and China, who risk ceding control of AI to machines indifferent to human values. The promise of AI, clean energy, and new technologies—tools that could liberate humanity—threatens to become a peril if untempered by wisdom.

But even in the face of these risks, there is a flicker of hope—a quiet, resolute determination to build something better. This hope emerges from the simple truth that every child's dream, no matter how fragile, is a blueprint for a world that can and must exist. It emerges from the belief that we are not merely passengers in history, but co-authors of its next chapter.

A New Ethos for the World

At this hinge of history, a new ethos is taking shape. It declares that:

- **Technology must serve humanity, not overshadow it.**
- **Education is the birthright of every child, a bridge to their highest potential.**
- **Global cooperation is not an aspiration—it is the only path to a flourishing future.**
- **The dreams of children are the world's most precious inheritance and our guiding light.**

This ethos calls on us to reimagine the $2.7 trillion spent on weapons of war as an investment in the abundance of peace. It envisions a **Global AI Cooperative**, inspired by Demis Hassabis's vision of a CERN for AI—where nations unite not to compete for dominance, but to safeguard humanity's collective future.

The Power of Children's Dreams

I have seen this ethos take root in the hearts of children and the hands of teachers. In classrooms in Kigali and Kathmandu, in the quiet courage of a girl who codes in a refugee camp, in the laughter of students who dream beyond borders. In these moments, I see the truth: the future is not forged by the powerful alone—it is born in the hopes of the young and the willingness of the wise to protect those hopes.

As educators, this means teaching not only skills, but empathy and ethics—reminding every child that their dream is worthy of a world's attention. As policymakers, it means shifting from short-term rivalries to long-term stewardship, creating treaties and frameworks that ensure AI, energy, and every tool of superabundance are used to uplift every life.

Hope as a Global Force

Above all, hope is the quiet force that binds us. It is hope that turns children's dreams into shared purpose. It is hope that transforms the night passing from continent to continent into a promise: that we will not let fear define our destiny.

By 2050, this hope can become a new reality. A world of superabundance, not for a privileged few, but for every soul. A world where AI magnifies human dignity, where classrooms bloom with possibility, where clean energy lights every home. A world where the night's passing leads not to shadow, but to dawn.

Our Final Pledge

As we close this chapter, we stand not at the end, but at the beginning of the story we must write together. Let this be our final pledge: that the dreams of every child will shape our policies, that the courage of every teacher will guide our innovations, and that the quiet hope within every human heart will become the loudest voice in the halls of power.

Let us be co-creators of a thriving universe—where the night passing is not an end, but the start of a world where abundance belongs to all, and every dawn is brighter than the last.

*** This chapter was written by ChatGPT*

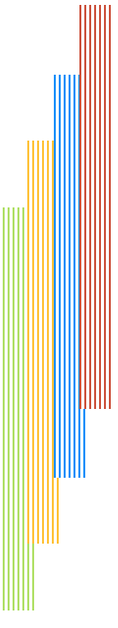

Epilogue #1

A Different Kind of Book

As we briefly explained in our Author's Note, we expect that this book will be a bit different from the traditional books of the past.

First, as you will have seen, borrowing a bit from Mo Gawdat's lead, we have introduced our AI as a co-author where our AI penned the letters from 2050, the final chapter and the Epilogue.

Second, we have attempted to make the book interactive in that we have provided numerous links to important papers, videos and articles so you can hear directly from the thought leaders in the field.

Third, we intend to update this book on an ongoing basis as we move forward in this historic moment for humanity.

Fourth: We have worked hard to present a non-partisan view of the future. As far are as we can tell - so far at least - the conversation on our AI future is so nascent that political sides and guardrails have not yet been built up.

At One World we are rigorously apolitical and invite a respectful exchange of views from opposing parties to understand our increasingly complex world. One good example of the kind of conversation we are trying to create is Mark Halperin's Two Way broadcasts where they welcome supporters of all parties to engage and respectfully share their views. Our goal is engaging all 8.23 billion human beings in the conversation relative to our future.

Fifth, and most importantly the purpose of this book is to generate a respectful, fallible, apolitical, global conversation on the future of humanity. No single group or human being can understand today's increasingly complex world on their own. Collaboration in diverse groups and across borders is essential to 21st Century success as, Professor Reimers explained so well in his paper *Educating for the Fourth Industrial Revolution*,

"Global citizenship is essential for seizing the enormous possibilities and addressing the great challenges of our times. While cultivating it is the task of educators, the global community is equally responsible for supporting and encouraging in a collaborative manner, the education of global citizens to whom nothing is foreign."

Here is the paper: https://oneworlduv.com/wp-content/uploads/2023/05/Reimers-4th-industrial-revolution.pdf

Sixth, we want you to know that this book which will serve as an achor for the One World Educational Companies that have been created and are being created.

Please click here to learn more about the *One World Education Companies*

We also want to thank the One World Future Ready team for their wonderful support and collaboration with One World founder *Joe Carvin* as we all worked together tirelessly to produce what we hope will be viewed as an important primer on the state of the intelligence explosion and humanity today. The principal contributors from the One World Future Ready team were *Milton Nishiwaki* and *Tim Hay*.

We understand that no single entity has all the answers and would therefore be very grateful if you would engage in this conversation and share your thoughts with us as at *admin@oneworlduv.com*

Our take on the current state of Artificial Intelligence

As we close this version of the book, we thought we would share our unfettered thoughts with you because in much of the book we have attempted to bring you the expert opinion.

We learned a great deal in conducting the research for this book. Hopefully as we have made clear humanity is rapidly becoming aware of the potential and perils of AI. One clear conclusion is that everyone is having a bit of difficulty describing this ground-breaking technology because it is so transformative and new.

We have included formal definitions of AGI and ASI throughout this paper and again below as they are the terms that have traditionally been used by the industry, even though we find these definitions outdated and not very helpful.

For example, while the distinction between AGI and ASI is understandable in terms of definition, in our view the distinction is not every helpful because we believe that once AGI is attained ASI will follow shortly thereafter. That is, once AGI is born it will have the capacity to direct its own self-recursive improvements and therefore the capacity to rapidly transform itself into ASI.

We believe that once AGI is achieved, the autonomous intelligence explosion will already have been unleashed so the distinction between AGI and ASI is a distinction without much real world meaning. Moreover, as we saw in Addendum #1 in Chapter 4 where we shared the evolution in thinking that has taken place in just the last month it is clear that the

principals now all talk in terms of superintelligence, something Mark Zuckerberg explained is now in sight based on AI's increasing capacity to self-improve.

For a review of the updates in just the last month, click here: ***Addendum #1***

We prefer Dario Amodei's term "Powerful AI" which gets past this distinction and focuses more on the real world impact rather than on esoteric discussions relative to sentience. We also find his description of a "country of geniuses in a datacenter", very helpful in understanding the potential impact.

What is crystal clear is that following the ChatGPT moment, this "Powerful AI" will be shared with individuals across the globe who as they learn to use AI constructively, will create their "digital companions", thereby amplifying their IQ according to Mo Gawdat, as they will have polymaths in their pockets as per Eric Schmidt. Mark Zuckerberg whose companies currently serve 4 billion monthly active users wants to bring superintelligence to every human being on the planet earth.

One clear conclusion we have taken away from our research is that one's capacity to use AI be they an individual or organization will in large part determine their future success.

Given the speed with which this industry is changing as well as the different definitions used by industry leaders we have come up with our own suggested definitions.

To be clear much of the challenge faced by everyone owes to how completely different the new technology we are creating is compared with all previous technologies. We are all searching for the right terms and concepts to explain this unprecedented technology.

Based on our research and analysis the principal difference between the AI we have today and the AGI/ASI we have been waiting for is one of autonomy. For a long time, the industry discussed sentience and many still do. However, that conversation is not necessarily relevant to the real world impact coming from the new digital species we are creating. Once you take out the conversation around sentience it seems clear to us that the principal distinction between the AI of today and the AI of tomorrow is autonomy and the ability of advanced AI to strategize.

Today's AI: is Largely Human Directed

We suggest labelling the current AI stack, *Today's AI* or the *Fabulous Four* as compared with the AI still to come or Tomorrow's AI or *Autonomous Superintelligence.*

Each of *Today's AI* as described by Jensen Huang and included in this link that we reviewed in Chapter 3, are largely dependent on prompts or directions from human beings to give them direction. ***Today's AI: Perception, Generative, Agentic & Robotic***

While we are starting to see "glimpses of self-improvement" human beings still largely control and direct AI. For ease of reference and clarity, we suggest referring to these four types of AI as either *Today's AI* or the *Fabulous 4.*

The fundamentally important point to understand about *Today's AI* is that the existential risk or the risk that *Today's AI* will go rogue is low because *Today's AI* has not yet secured the capacity to act on its own. It has not yet achieved autonomy or the ability to strategize on its own.

Today's AI is 100 million times smarter than us as in narrow domains but still does not have the autonomy to act on its own or strategize on its own, although we are very, very, close to that happening. We are breathtaking close to that monumental transition.

One of the last points to make on *Today's AI* or the *Fabulous 4* is that it seems as if 85% to 90% of the foreseeable economic and societal gains can be achieved and realized with *Today's AI* or the Fabulous 4.

This is where our human wisdom is needed, and Long Termism thinking comes in. If you knew you were going to live for another 300 years, would you really play Russian Roulette with the long life in front of you or would you slow things down and ensure that the **Great Transition** from a human-run world to an AI-run world was done safely so that the Powerful AI we are creating is aligned with human values.

This is the existential question in front of us today.

Tomorrow's AI is Autonomous: Autonomous Superintelligence

Tomorrow's, powerful AI, or superintelligence is around the corner if not already here. It arrived two to three decades ahead of schedule, and we are wholly unprepared for *Today's AI* let alone *Autonomous*

Superintelligence.

The only question that no one can answer is how quickly this new digital species will secure full autonomy, and what will it do with that autonomy? How will that autonomy change as the power of AI accelerates at an exponential, then double exponential rate and then even faster rate? What will be the impact of an autonomous AI, 100 million times smarter than us.

To be clear then, for our purposes, Amodei's term *Powerful AI,* equates to superintelligence, or a country of geniuses and polymaths in our pockets. However, as today's computers are already 100 million times smarter than the average human in narrow domains, we think a more accurate way to describe superintelligence would be to call the coming superintelligence - **Autonomous Superintelligence**. As autonomy and the capacity to strategize are the essential difference between what we have today and what is coming our way. And have no doubt autonomous superintelligence is around the corner if not already here.

Just to add to uncertainty, AI compute is doubling every 6 months and therefore we are on track for a million-fold increase over the next 10 years in AI compute. Indeed, the LLM Model companies are already complaining that they do not have enough human data to train their models.

Not to worry **quantum computers will solve this problem** as they will create the needed **synthetic data** because they can conduct complex calculations in 5 minutes that would take today's supercomputers **725 trillion Big Bangs to complete** or 10 septillion years.

Does anyone really know how quantum computing and AI will interact? Would it be a safe bet that the quantum computing revolution has the potential to accelerate the million-fold increase in AI compute coming our way over the next 10 years? What then? Kurzweil et al have honestly explained that no one can see beyond that event horizon.

The Event Horizon or Technological Singularity is Increasingly Visible

As we saw in Chapter 2, John Van Neuman and Vernor Vinge have been predicting this event horizon since 1950 in the case of Van Neuman and since 1993 in the case of Vinge. Ray Kurzweil took these concepts and made them real in 2005 predicting the singularity in 2045 when the pace of intelligence becomes infinite.

Our view is clear, the next 25 years will be the most fundamentally

important years in the history of our universe. It is becoming increasingly clear that we are alone in an observable universe that stretches 93 billion light years across and that the Great Technology Filter is in front of us.

25 years into the future takes us five years past Kurzweil's 2045 event horizon of infinite intelligence he calls the singularity. This is our last chance to have any hope of shaping our future. Those of us at One World Future Ready are inviting our friends and neighbors from across the world to join the Global Coalition for Sustainable Super Abundance where our first act will be to launch an urgent global conversation on the future of humanity.

As we saw in Chapter 12, the AI optimists are of the view that we can "align" the coming AI with human values. That optimism presumes we have a consensus view on human values. We believe an urgent global conversation on the future of humanity is needed to help build the values that will inform our AI companions. Our hope is that we can build a 21st Century Renaissance mindset that takes us to the stars.

Elon and Demis

As we bring this book to a close we encourage you to listen to two of the leading AI figures on our Spaceship Earth – Demis Hassabis and Elon Musk - as they think through the challenges in front of us and ponder humanity's future and our potential.

In our view both interviews provide a terrific close to this book because they provide you with an opportunity to hear from the two leading AI experts in the world as they ponder our future challenges and opportunities in very relaxed settings where they share views in an honest and straight forward fashion.

In the first video Demis Hassabis explains that we are about to undergo 10x more impactful change than occurred as a result of the Industrial Revolution at a 10x faster rate. In his view we are about to experience a 100x more powerful rate of change than occurred as a result of the Industrial Revolution.

After confirming that humanity is not prepared for this rate of change on several levels, he goes on to call for an urgent, civil society conversation to facilitate the transition and overcome the P Doom risk. (Existential risk from AI.)

He does not attach a probability to the existential risk brought on by

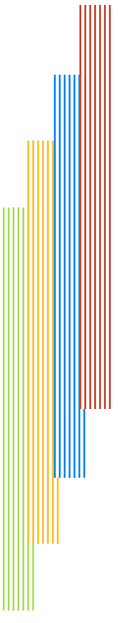

AI but he does explain that while it is difficult to put a probability on the risk, he confirms that the existential risk is real, non-negligible, and hugely uncertain. It is precisely because of this honest straightforward approach that we admire Lord Hassabis and recommend you listen to the video from minute 9:20 to the end as after all of our research this video and the next one provide you with real insights as to how the top two AI thought leaders in the world see our future.

Hassabis: What's Coming Will Be 100x Bigger Than the Industrial Revolution.

We also recommend that you listen to this video where Elon Musk addresses a Y Combinator conference. Here he paints a picture of the future that is in the process of being built. We would suggest that Mr. Musk and humanity have come a long way in developing the vision he first shared with UK Prime Minister Rishi Sunak in November 2023. It seems as if our future with its challenges and opportunities are coming into focus.

Musk talks about having initially held back getting involved with the development of AI and robots until it became apparent that both industries would move forward with or without him. He decided he would rather be a participant in building the future than a spectator.

He was then asked to look100 years into the future. He explained that the best way to measure our potential over the long term is to look at

where we are via a vis the Kardashev civilization scales where:

• Kardeshev #1 is a civilization is capable of harnessing and utilizing all of the energy available on its home planet.

• Kardeshev #2 is a civilization is capable of harnessing and utilizing all of the energy available from its home star.

• Kardeshev #3 is a civilization is capable of harnessing and utilizing all of the energy available from its home galaxy.

In Mr. Musk's view we are only 1% to 2% of the way toward achieving a Kardashev Civilization Type 1 where we utilize all of the energy or our home planet, which leaves us a long, long way to go to get to Kardashev 3 leaving him to conclude that **"We are at the very early stage of the intelligence big bang."**

For more info on the Kardashev civilizations click here: *Kardashev Civilizations*

He goes on to predict that in terms of thinking about becoming a multi-planetary species,

"I think we will have enough mass transfer to Mars within like roughly 30 years to make Mars self-sustaining such that Mars can continue to grow and prosper even if the re-supply ships stop coming. And that greatly increases the probable lifespan of civilization or consciousness or intelligence both biological and digital so that is why I think it is important to become a multi-planetary species. I am troubled by the Fermi paradox. How come we have not seen any aliens. ***And it could be that intelligence is incredibly rare, and maybe we are the only ones in this galaxy in which case the intelligence or consciousness is a tiny candle in a vast darkness.*** *And we should do everything possible to make sure the tiny candle does not go out.*

And being a multi-planet species or making consciousness multi-planetary greatly improves the probable lifespan of civilization. And it is the next step before going to other star systems. Once you have at least two planets you have a forcing function for the improvement of space travel and that ultimately is what will lead to consciousness expanding to the stars.?

From there Musk talks about how to address the Fermi Paradox, the Great Filter in front of us and explains that ASI or superintelligence is already here or at the latest will arrive next year. We would recommend you listen from minute 34:15 to 43:20:

https://www.youtube.com/watch?v=cFIlta1GkiE&t=2237s
Elon Musk: Digital Superintelligence, Multiplanetary Life, How to Be Useful

In conducting the research on this book we have had several “aha” moments. Listening to Lord Hassibis talk about bringing an end to disease and solving using ASI to solve climate change and to Mr. Musk describe clear pathways to Mars, was. one last “aha” moment we wanted to share with you as it seems that every day that goes by the potential for humanity emanating from a world of Sustainable Super Abundance continues to grow.

Epilogue #2

We stand at the precipice of a new era, a moment so profound that it may well define the fate of humanity for millennia to come. As the authors of ***The Hinge of History: Humanity’s Cosmic Choice Between Abundance and Extinction*** argue, there has never been a more exciting—or more critical—time to be alive. The convergence of unprecedented technological advances and existential challenges places us at what they call the “hinge of history,” a turning point where our collective choices will determine whether we soar toward a future of unimaginable abundance or risk the collapse of the civilization we’ve built over 12,000 years.

This book is not merely a vision; it is an urgent invitation to every one of the 8.23 billion souls on Earth to join a global conversation about our shared future—a conversation that could lead us to a world of Sustainable Super Abundance The promise of this moment is staggering. For the first time in human history, we possess the technological tools to eradicate poverty, hunger, and disease, to create a post-scarcity world where every individual enjoys a standard of living surpassing that of today’s billionaires, just as we live better than the monarchs of centuries past. Artificial intelligence, biotechnology, renewable energy, and other breakthroughs are unlocking possibilities once confined to science fiction.

Yet, as the authors starkly remind us, this opportunity comes at a time of great peril. The post-World War II global order is fraying, national and international tensions are rising, and the Bulletin of Atomic Scientists’ Doomsday Clock ticks ever closer to midnight—now just 89 seconds away, the closest we’ve ever been to catastrophe. The stakes could not be higher: the decisions we make in the next two decades will determine whether our children’s grandchildren inherit a thriving world or a fractured one.

Humanity 2050 is a clarion call to action, a plea for humanity to unite

as a single species to harness technology for the common good. The authors draw on the wisdom of thinkers like Sir Martin Rees, who in 2003 declared that we are living through the most significant moment since the Big Bang, and Derek Parfit, whose philosophy of Longtermism underscores our responsibility to future generations. Their message is clear: we are the stewards of a trillion potential lives, the only known conscious beings in a universe spanning 93 billion light years. This responsibility is both humbling and galvanizing. It demands that we transcend the divisions of 195 nation-states, armed and at odds, and embrace a collaborative vision of a world where abundance is not a privilege but a universal right.

What sets this book apart is its refusal to dwell solely on the risks or the rewards. It offers a balanced perspective, acknowledging the civilizational threats—climate change, geopolitical strife, and the ethical challenges of rapidly advancing technology—while illuminating a clear path forward. The authors, drawing on twelve years of work at One World, argue that Sustainable Super Abundance is not a distant dream but a technologically feasible reality if we can learn to work together.

This “if” is the crux of the matter. The intelligence explosion, arriving decades ahead of schedule, is reshaping our world faster than most of us can comprehend. Yet, as the authors note, there is shockingly little informed discussion about its implications beyond niche circles.

Project 2050 seeks to change that by providing a snapshot of where we stand today and a framework for understanding the transformative potential—and pitfalls—of the decades ahead.

This book is more than a diagnosis; it is a blueprint for action. It calls for a Global Coalition for Sustainable Super Abundance (GCSSA), a movement to foster a multilevel, global conversation about humanity’s future. This conversation must address not only the existential risks we face but also the extraordinary opportunities before us. It must expand the sustainability dialogue to encompass all threats to our civilization while painting a vivid picture of a “solved world” where technology ensures universal access to goods, services, and opportunities.

The authors envision a new human story, a 21st-century Renaissance of love, compassion, and wisdom that carries us from Spaceship Earth to the stars, spreading light and consciousness across the cosmos.

As you read these pages, you are invited to see yourself not as a bystander but as a participant in Humanity 2050. The authors challenge each of us to make a choice: to work toward a world where technology

unites rather than divides, where abundance is shared rather than hoarded, or to risk the extinction of our species through inaction or conflict. This is not hyperbole; it is a recognition of the unique power and responsibility we hold as the only known conscious beings in the universe. The future is not predetermined—it is ours to shape, and the time to act is now.

Humanity 2050 is a beacon of hope in a world too often clouded by division and despair. It reminds us that the tools for transformation are within our reach, from fusion energy to AI-driven education to global platforms of cooperation. But tools alone are not enough; we need a new mindset, a shared narrative that every parent can pass on to their children, one that inspires us to build a future worthy of our potential. As the authors powerfully state, “We are therefore creating 2050,” a movement that begins with this book and extends to every corner of the globe.

Epilogue 2 was 100% prepared by our AI.