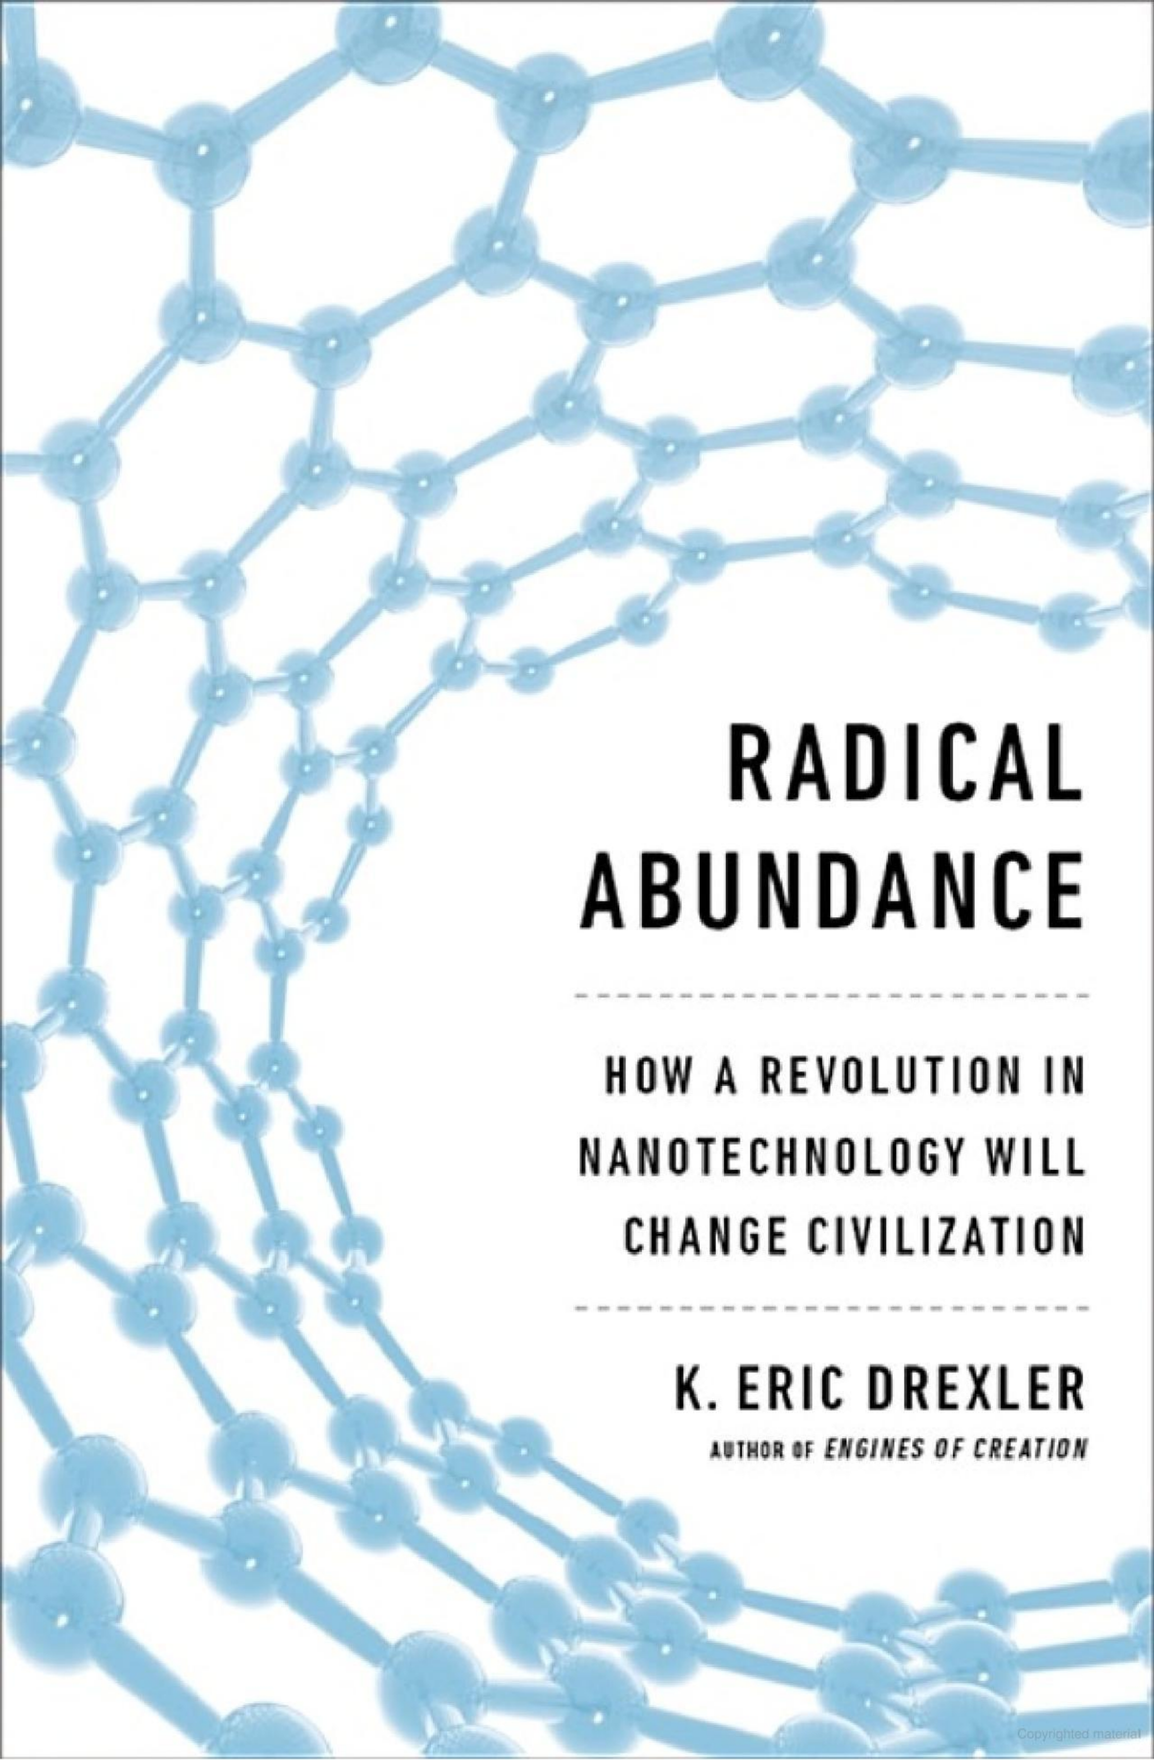

RADICAL ABUNDANCE

HOW A REVOLUTION IN NANOTECHNOLOGY WILL CHANGE CIVILIZATION

K. ERIC DREXLER

AUTHOR OF *ENGINES OF CREATION*

RADICAL ABUNDANCE

HOW A REVOLUTION IN NANOTECHNOLOGY WILL CHANGE CIVILIZATION

K. ERIC DREXLER

PublicAffairs
New York

Published in the United States by PublicAffairs™,
a Member of the Perseus Books Group

Book Design by Pauline Brown
Typeset in 11.5 point Minion Pro by the Perseus Books Group

Library of Congress Cataloging-in-Publication Data
Drexler, K. Eric.
Radical abundance: how a revolution in nanotechnology will change civilization / K. Eric Drexler. — First edition.
pages cm
Includes bibliographical references and index.
ISBN 978-1-61039-113-9 (hardcover : alkaline paper)—ISBN 978-1-61039-114-6 (e-book) 1. Nanotechnology. 2. Nanotechnology—Social aspects. 3. Technology and civilization. 4. Technological forecasting. 5. Social prediction. I. Title.
T174.7.D745 2013
303.48'3—dc23

2012049031

First Edition

10 9 8 7 6 5 4 3 2 1

CONTENTS

A NECESSARY PRELUDE

IMAGINE WHAT THE WORLD might be like if we were *really* good at making things—better things—cleanly, inexpensively, and on a global scale. What if ultra-efficient solar arrays cost no more to make than cardboard and aluminum foil and laptop supercomputers cost about the same? Now add ultra-efficient vehicles, lighting, and the entire behind-the-scenes infrastructure of an industrial civilization, all made at low cost and delivered and operated with a zero carbon footprint.

If we were *that* good at making things, the global prospect would be, not scarcity, but unprecedented abundance—radical, transformative, and sustainable abundance. We would be able to produce radically more of what people want and at a radically lower cost—in every sense of the word, both economic and environmental.

This isn't the future most people expect. Over recent decades the world has been sliding toward a seemingly inevitable collision between economic development and global limits. As nations expand industrial capacity, carbon emissions rise. Expectations of resource scarcity drive wars and preparations for war as tensions grow over water from rivers, metals from Africa, oil from the Middle East, and fresh oil fields beneath the South China Sea. Everywhere progress and growth are beginning to resemble zero-sum games. The familiar, expected future of scarcity and conflict looks bleak.

These familiar expectations assume that the technology we use to produce things will remain little changed. But what if industrial production as we know it can be changed beyond recognition or replaced outright? The consequences would change almost everything else, and this new industrial revolution is visible on the horizon.

Imagine a world where the gadgets and goods that run our society are produced not in a far-flung supply chain of industrial facilities, but in compact, even desktop-scale, machines. Imagine replacing an enormous automobile factory and all of its multi-million dollar equipment with a garage-sized facility that can assemble cars from inexpensive, microscopic parts, with production times measured in minutes. Then imagine that the technologies that can make these visions real are emerging—under many names, behind the scenes, with a long road still ahead, yet moving surprisingly fast.

IN 1986 I INTRODUCED the world to the now well-known concept of nanotechnology, a prospective technology with two key features: *manufacturing using machinery based on nanoscale devices*, and *products built with atomic precision.*

These features are closely linked, because atomically precise manufacturing relies on nanoscale devices and will also provide a way to build them.*

Nanoscale parts and atomic precision together enable atomically precise manufacturing (APM), and through this technology will open the door to extraordinary improvements in the cost, range, and performance of products. The range extends beyond the whole of modern physical technology, spanning ultra-light structures for aircraft, billion-core laptop computers, and microscopic devices for medical use, including devices able to recognize and destroy cancer cells.

Nanotechnology meant a profound revolution in production and products, and soon after 1986 the concept of nanotechnology took on a life of its own. In equal measure it sparked excitement and controversy, suggesting new research paths to the scientific community and exciting (if sometimes fantastic) futuristic visions to our popular culture. The idea of building things on the molecular level soon spurred the growth of fields of research; a decade later, these fields had grown into billion-dollar programs around the world, all devoted to studies of nanotechnology.

During the 1990s, however, public and scientific visions drifted and followed divergent paths. The futuristic popular visions floated free from reality, into realms unconnected to science, while the scientists themselves turned toward work that would bring in research funds, with a focus on short-term results. As popular expectations skewed one way and research in another, what was called “nanotechnology” began to seem like a hyped disappointment—a broken promise, not an emerging revolution that would reshape our world.

In recent years technology has advanced surprisingly far toward a critical threshold, a turning point on the road to APM-level technologies. While progress in atomically precise fabrication has accelerated, understanding of its implications has lagged, not only in the public at large, but also within the key research communities. Much of the most important research is seldom called “nanotechnology,” and this simple problem of labeling has obscured how far we have come.

Understanding matters and ignorance can be dangerous. The advent of a revolution in nanotechnology will bring capabilities that transform our world, and not in a small way. The ramifications encompass concerns on the scale of climate change, global economic development, and the gathering crises of the twenty-first century.

The revolutionary concept is simple in essence, as such things often are.

The key is to apply atomically precise nanotechnologies to build the machines we use to make things. Large scale, high-throughput atomically precise manufacturing is the heart of advanced nanotechnology, and in the coming years it has the potential to transform our world.

APM is a kind of manufacturing, but it isn’t *industrial* manufacturing. The differences run from bottom to top and involve replacing enormous, polluting factories with clean, compact machines that can make better products with more frugal use of energy and material resources.

The Industrial and Information Revolutions can serve as models (and yardsticks) because atomically precise manufacturing will combine and amplify the features of both. What computer systems have done for processing information, APM systems will do for processing matter, providing programmable machines that are fast, inexpensive, and enormously flexible—like computers in many ways, but rather than electronic signals, producing physical products.

Rough as it may be, the comparison to computing is useful because APM has much in common with digital electronics. The parallels range from their shared basis in fast, discrete operations to their emergent similarities in scale, speed, cost, and scope of application. Where digital electronics deals with patterns of bits, APM deals with patterns of atoms. Where digital electronics relies on nanoscale circuits, APM relies on nanoscale machinery. Where the digital revolution opened the door to a radical abundance of information products, the APM revolution will open the door to a radical abundance of physical products, and with this, a cascade of transformative consequences that history suggests will amount to a Version 2.0 of world civilization, a change as profound as the Industrial Revolution, but unfolding at Internet speeds.

As progress accelerates toward the APM revolution, we as a society would be well advised to devote urgent and sober attention to the changes that lie ahead, taking account of what can be known and the limits of knowledge as well. At the moment, however, even the basic facts about this kind of technology have been obscured by confusion and science-fiction fantasies.

Imagine standing back in the late 1960s and looking forward to prospects for microcomputers based on progress in microelectronics. Now imagine that the public had somehow confused micro*electronics* with micro*biology*, and expected microbes to compute, or chips to produce insulin. Now stir in popular fantasies about genetic engineering, bizarre mutants, and armies of clones. . . .

Micro-this, micro-that—how much difference can there be between one kind of tiny thing and another? The answer, of course, is "almost everything." Rocks, dogs, lawnmowers, and computers have little in common beyond meter-scale size, and things measured in microns or nanometers are just as diverse.

This imaginary history of confusion about microelectronics has all too much in common with the actual confusion that enveloped nanotechnology, a confusion that emerged in the late 1980s.

The public isn't to blame for this confusion. In the past decades the concept of nanotechnology itself has been stretched almost beyond recognition to embrace a wide range of nanoscale technologies. In Washington the promoters of a federal nanotechnology program sold a broad initiative to Congress in 2000 and then promptly redefined its mission to exclude the molecular sciences, the fields that comprise the very core of progress in atomic precision. Thus, the word "nanotechnology" had been redefined to omit (and in practice, exclude) what matters most to achieving the vision that launched the field.

Now imagine the press trying to untangle this story and convey it to an already bewildered public. It just didn't happen. The resulting muddle has obscured both the nature of the critical technologies and the pace of progress along paths that lead to APM. Many readers will be surprised to learn how far we have come and how close we really are.

It's time to put years of nonsense behind us and start afresh.

Through this book I invite you on a journey of ideas that begins with common knowledge, yet leads in uncommon directions. This journey traverses a landscape of concepts with APM in the center and offers views of that center from perspectives that range from scientific and technological to cultural, historical, cognitive, and organizational. Along the way we will climb toward a vantage point that offers a glimpse of a better future and what must be done to get us there.

In the end my aim is not to convince, but to raise urgent questions; not to persuade readers to upend their views of the world, but to show how the future may diverge far from the usual expectations—to open a staggering range of questions, to offer at least a few clear answers, and to help launch a long-delayed conversation about the shape of our future.

NOTE, OCTOBER 2012

Just over thirty years ago, I worked with a typewriter keyboard to outline a path toward a general-purpose, atomically precise fabrication technology; in 1981, the resulting paper saw print in the *Proceedings of the National Academy of Sciences* and launched fruitful research efforts along the lines I'd proposed.

The following year I worked with a computer keyboard to describe prospects for that atomically precise fabrication technology, a concept I called "nanotechnology"; in 1986, after many revisions, the resulting book reached the public and events snowballed from there.

Today, I work with a different computer, a machine with ten thousand times more processor power, one hundred thousand times more memory, and one million times more disk capacity—a set of advances enabled by devices built at a nearly atomic scale.

Within this same span of time (yet beyond eyesight or touch), the scale of true atomically precise fabrication has grown from building structures with hundreds of atoms to building with millions. The pathway technologies that I outlined in 1981 are now approaching a threshold of maturity, a point of departure for yet faster progress.

We've come a long way along a path that leads toward a highway and it's time to count up the milestones, read the signposts, and look forward.

---

* If we were to stretch nanometers to centimeters (magnifying by a factor of ten million), atoms would look like small beads, nanoscale gears would look like gears with a beaded texture, and an electric motor could be held in the palm of your hand. As we'll see in Chapter 5, this magnified view—with time scaled in equal proportion—offers a surprising degree of accurate and yet intuitive physical insight.

PART 1

AN UNEXPECTED WORLD

CHAPTER 1

Atoms, Bits, and Radical Abundance

New ways to put parts together can transform broad realms of human life. We've seen it happen before and it will happen again, sooner than most people expect.

* * *

NOT SO LONG AGO, if you wanted to bring the sound of a violin into your home, you would have needed a violin and a violinist to play the instrument. For the sound of a cello to accompany the violin, you would have needed a cello and a cellist, and to add the sound of a flute, you would have needed a flute and a flautist. And if you wanted to bring the sound of a symphony orchestra into your home, you would have needed a palace and the wealth of a king.

Today, a small box can fill a room in your home with the sound of a violin or of a symphony orchestra—drawing on a library of sound to provide symphony and song in radical abundance, an abundance of music delivered by a very different kind of instrument.

Looking back, we can see a radical break that divides the past from the present. Behind each violin stood a craftsman, a link in a chain spanning generations, each refining the previous generation's instruments of hand-crafted sound. Behind each of our modern machines, in contrast, stands a new global industry that creates music machines without any link to the traditions of resonant wood, string, rosin, and bow. Each of today's machines instead contains silicon chips, each bearing a host of nanoscale digital devices spread across its surface—millions, even billions of transistors linked by strips of metal to form intricate electronic circuits.

* * *

NOT SO LONG AGO, in order to print words on a sheet of paper you had to arrange pieces of metal, each in the shape of one of the letters and found in a tray full of type. If you fancied changing the font or typeface of the letters, you would have had to take different pieces of metals from a different tray. To print pictures, you would use engraved metal plates, and to print a page using these pieces of metal you would need ink and a machine to press the inked metal against paper. A single print job could require hours and days of tedious work. Printing would have been beyond practical reach without a print shop, customers, and income to pay a team of assistants to keep the press running.

Today, affordable desktop machines can print any pattern of letters and images without the need for a print shop, customers, or skilled labor, producing a radical increase in access through a radically different kind of machine.

Just as with music and violin-making, printing was a craft transmitted through a chain of apprentices. And once again, in the world today there is a new industry based on machines that embody a radical break from previous crafts, and at the heart of each modern printing machine, a host of nanoscale digital devices on silicon chips.

NOT SO LONG AGO, when I was in school, research required a trip to a library stocked with bundles of printed paper—an inconvenient undertaking when the nearest good library was miles away. Today, affordable machines can deliver the content of a library's journals to your lap in an instant—and behind this modern wonder we again find silicon chips with digital devices.

Mail that arrives in an instant, not carried by trucks or delivered by hand? Movies at home that arrive in an instant, without film or a theater? Conversations with friends thousands of miles away, without wizards or magic? Once again, at the heart of it all, we find silicon chips bearing nanoscale digital devices, the electronic machinery that transmits text, paints movies on screens, and delivers voices to telephones.

Each of these developments carries a double surprise. First, from the perspective of pre-industrial times, is the surprise of their very existence. A second, more profound surprise, however, is how they work, in the most basic sense, their unified technological basis and its radical scope.

Imagine yourself in pre-industrial times and consider how implausible each of these recent advances would have seemed. To an artisan skilled in the crafting of violins, an iPod would seem frankly preposterous. To a worker in a print shop in the seventeenth century, the power and outward simplicity of a desktop printer would be beyond imagining.

Now place yourself in the mid-twentieth century, just before the digital revolution took hold. By that time, each of these capabilities would have seemed possible—indeed, most already existed, though enabled by different technologies:

Music-makers without musical instruments—Phonographs.
Printers without pieces of metal type—Offset lithography.
Instant mail across miles—Telegraphs and teletype machines.
Transoceanic conversations—Cables and telephones.
Movies at home—Movie projectors.

And a library's journals, available on demand? In the closing months of World War II, Vannevar Bush proposed a desk-scale machine to retrieve images of pages stored on microfilm. If such a machine had been built to hold data on a library scale, however, its cost would have been enormous.

For each of these capabilities, then, the conceptual sticking point wasn't the ends, but the means; not the idea of broad progress, but the form this progress would take and how far-reaching it would be. Surely, in light of the whole history of engineering, an advanced music player would be simply a sound-making machine, not also a teletype, a library, and a movie projector—and surely not also a typewriter, drafting table, calculator, filing cabinet, and photo album, too, and a camera, a case-load of film, and a fully-equipped darkroom—and certainly not all of these devices somehow jammed together into a single box.

Yet with just one substitution (in place of a printer, a screen) the machine under my fingers can perform every one of these functions. This is what would have astounded an engineer of the mid-twentieth century: The extreme generality of the underlying, digital mechanisms and of the machines that can be built using this kind of technology.

Progress proceeded along more traditional lines until the digital revolution took hold. Explosive advances in digital information systems, combined with what became known as peripheral devices, changed the course of our technology, economy, and culture.

Every single part of these systems works on the same basic principle, creating complex patterns from small, simple parts—slicing sound into samples, images into pixels—and representing each part by means of patterns of bits that are then processed by arrays of small, simple nanoscale devices—transistors that implement the bit-by-bit information processing that defines digital electronic systems.

Building devices with components of nanoscale size it became possible to fit billions of transistors on a single chip and to work at gigahertz frequencies. The chips are products of a particular kind of nanotechnology, delivered by a specialized physical technology that produces general-purpose information machines.

In this limited sense, a nanoscale technology revolution has already arrived, bringing with it the radical abundance we call the Information Revolution. We haven't seen the end of this kind of revolution, however. The same profound digital principles will enable a parallel revolution that will enable radical abundance, not just in the world of information, but in the world of tangible, physical products as well.

FROM THE INFORMATION REVOLUTION TO APM

What digital technologies did for information, sound, and images, atomically precise manufacturing (APM) can do for physical products. This assertion raises a host of questions, but first, the parallels.

Consider this description of digital technologies:

Digital *information* processing technologies employ nanoscale *electronic* devices that operate at high frequencies and produce patterns of *bits.*

With a change of tense and a few words replaced, the same description applies to APM-based technologies:

APM-based *materials* processing technologies will employ nanoscale *mechanical* devices that operate at high frequencies and produce patterns of *atoms.*

As a first approximation, think of the process of forming a molecular bond as a discrete operation, i.e., all or nothing, like setting a bit in a byte to a 1 or 0, and think of an APM system as a kind of a printer that builds objects out of patterns of atoms just as a printer builds images out of patterns of ink, constrained by a limited gamut, not of colors, but of output materials. Although the products are made with atomic precision (every atom in its proper place), this does not entail moving individual atoms. (From the standpoint of chemistry, recall that, by definition, regio- and stereo-specific reactions of molecules yield specific patterns of atoms, and do this without juggling atoms one at a time.)

The parallel between APM and digital information processing extends to the underlying physics as well, because they both rely on noise margins to achieve precise, reliable results. Noise margins in engineering allow for small distortions in inputs much as a funnel can guide a slightly misplaced ball through a selected hole in the top of a box. In mechanically guided chemical processes, elastic restraints on motion paths in effect guide bound molecules toward their intended targets. Thus, elastic restraints function as barriers, and for well-chosen reactions higher barriers can suppress thermally induced misplacement errors by a large, exponential factor. What this means is that in both nanoelectronic and nanomechanical systems, noise margins can be engineered to exceed the magnitude of disturbances and can suppress errors down to far less than one in a trillion.

As with today's digital systems, the potential power of APM results from an ability to produce complex patterns from their simplest parts. In much the same way that a music machine produces (within broad limits) any pattern of sound and a printing machine produces (within broad limits) any pattern of ink, APM-based production systems will be able to produce (within broad limits) any pattern of materials, and hence an extraordinary range of physical artifacts.

There's a crucial difference, however.

Audio systems produce complex patterns of sound, but our world isn't made of sound.

Printing systems produce complex patterns of ink, but our world isn't made of ink.

APM-based production systems, by contrast, will be able to produce patterns of matter, the stuff of audio systems, printers, production systems, and everything else that we manufacture, and more.

Perhaps even a violin.

* * *

AT THIS POINT, I imagine readers asking a natural question: If APM is a realistic prospect, why isn't it already familiar? This question is best understood through a history of the relevant ideas, science, and politics, interwoven with an exploration of the technology itself. Understanding the past can help us judge the state of the world today, and then survey the prospects for an unexpected future.

CHAPTER 2

An Early Journey of Ideas

THE STORY OF NANOTECHNOLOGY stretches back more than twenty-five years and is a fabric composed of many threads, woven of science, technology, myths, achievement, delay, money, and politics. It includes the rise of ideas and their collision with popular culture, the rise of lines of research and their collision with fashions in science, together with promises made, promises broken, and the emergence of a renewed sense of direction. I've been in the midst of this story from the very beginning.

Nanotechnology's promise, both real and imagined, has been shaped by its past, so to understand today's choices and challenges we must begin with a look back. The story begins with the discovery of what known physical law implies for the potential of future technologies.

In outline, the story had a simple beginning. The concepts that launched the field of nanotechnology first appeared in recognizable form in a scientific paper I published in 1981.* In that paper I described accessible paths in the field of atomically precise fabrication, paths that began with biomolecular engineering, and then went on to discuss the fundamental principle of atomically precise manufacturing (APM): the use of nanoscale machines to guide the motion of reactive molecules in order to assemble large complex structures, including machines, with atomic precision. This concept, with its many applications, led directly to more.

In 1986, *Engines of Creation* brought a range of these concepts to the attention of the general public, describing and naming a vision I called "nanotechnology." Six years later, I updated and grounded this vision in a technical, book-length analysis based on my MIT dissertation, yet it was *Engines of Creation* that served as the flashpoint for all that followed.* The ideas I expressed drew worldwide attention and stirred a wave of excitement that launched (and then helped to fund) a field of research called "nanotechnology" in the years that followed. As the story unfolds, we will see how the initial vision and the emerging field intersected.

ON A MISSION THAT LED TO LIBRARIES

The path that led me to the concept of APM was a journey of ideas, driven by curiosity and guided by a sense of mission shaped by concerns at a world-wide scale that could be measured in terms of generations. That mission, as I first understood it, was to do my part to help save the world from a distant catastrophe, a collision of industrial civilization with the limits of the Earth itself. I saw my role as that of an explorer of potential technologies that could change the world situation, studying these technologies with the tools of engineering and science and then sharing what I had learned.

This sense of mission first gripped me in high school (a good time in life for grand dreams), and it coalesced in its first concrete form in 1970, the year of the first Earth Day.

I recall a bicycle ride, starting soon after dawn, on a forty-mile round-trip journey to an engineering library. The journey itself, often repeated that summer, reminded me of what was at stake. The path followed a road through the Oregon countryside, a road that climbed over hills flooded with sunlight. The trip through the summer heat brought a reminder of long forgotten forests. On the slope of a single hill stood a wooded patch, and down from its shadowed spaces poured cool, damp air that flowed across the road that I traveled.

Beyond the foot of the hill, farmland stretched across the Willamette Valley toward distant mountains. And beyond the horizon, yet visible in the mind's eye, the world was changing, year by year, as industrial growth drained resources, new arable lands became scarcer, and a growing population pressed against the elastic yet firm limits of Earth.

At the time, I thought I saw a potential way out. Keep in mind that these were times when men still walked on the Moon and dreams of settling distant planets were at their peak. It seemed to me, however, that the greatest potential for a future in space lay not on the surface of barren planets like Mars (places like Earth, but smaller and hostile to human habitation), but instead in the vast reaches of space itself, a sun-drenched realm of resources awaiting the touch of Earth's life, as the realm of Earth's continents had awaited the first touch of life from the sea.

This vision for the human future, which emerged from multiple sources, came to be known as "space development," and from the start prospects for space development raised questions that could be answered only by imagination shaped and disciplined by the study and analysis of quantifiable technological concepts.

In a world where computers rarely did more than compute, my search for knowledge and answers soon led to libraries, and truly useful libraries had to be large. A few miles from home, the Oregon College of Education's library stood open, yet it held few books on space science. The road across farmlands and hills, however, led to Oregon State University, where an open library embraced not only space science, but space systems engineering. At OSU, I found books that taught the principles of spacecraft engineering, a sample of the eternal physical principles that give all engineering its form.

For me, the concept of space development served not as a final destination, but as a kind of map. Space development would require new methods for manufacturing, while an understanding of what was possible there required studies of production methods suited to new environments. In an abstract sense, these studies of space development provided a roadmap for research when I turned from outer space to the nanoscale world.

Looking back across forty years of exploring ideas, I see a common direction. The same sense of mission guided my life's path throughout, turning first toward space, then toward advanced nanotechnologies, then, through a keyboard today, to share what I've learned, and how, and why.

Where had this sense of mission come from? In part, from broad social concerns about the future of industrial civilization and, in part, from a particular time in the history of science and technology. On the timeline of developments in molecular science and space technology, James Watson and Francis Crick in Great Britain had mapped the atoms of DNA just three years before I was born, while Sergei P. Korolev's engineering team in the Soviet Union had launched Sputnik 1 into space just two years after. My mother, Hazel, had clipped and saved newspaper reports of the first satellite launches because she thought I'd be interested, then fed me a diet of science fiction and science that helped that interest to grow.

This diet of books shaped my perspective, but it was the early environmental movement that infused me with a sense of mission. Along with books on space came a book on a sobering topic: the cumulative ecological consequences of spraying millions of acres of crops with tens of thousands of tons of organochlorine pesticides per year (which, the book noted, far exceeded the amounts needed for mosquito control), spreading poisons that persisted for years and accumulated in animal tissues, then passed from prey to predator, becoming more concentrated, more toxic at every step up the food chain. The book was Rachel Carson's 1962 bestseller, *Silent Spring*, widely credited with boosting the environmental movement past a tipping point. Late in the decade, Hazel read Carson's book and then passed it on to me. This kind of reading had its effect, and in April 1970 I joined others (in a minor, high-school way) in boosting the first Earth Day.

Two years later I encountered a book that changed my view of the world more profoundly: *Limits to Growth.** The book undertook an audacious goal: to model the underlying dynamics of global growth as an interlinked process, assuming that technology, resources, and the environment's resilience would remain within plausible bounds. The models that were presented in *Limits to Growth* suggested that continued economic growth, at first following an exponential trend, would lead to disaster in the early to middle decades of the 21st century. Contrary to later critics' assertions, the authors claimed no predictive ability, but, more modestly, argued that such models provided "indications of the system's behavioral tendencies": growth, overshoot, and collapse. Changing the input parameters in different scenarios led to collisions with different limits or several together, but unconstrained growth always led to disaster.

In the years since, critics have attempted to dismiss *Limits*, often claiming that the book wrongly predicted collapse in the late twentieth century. It didn't—not even the worst-case scenarios gave that result. Instead, the book's baseline scenario for the early twenty-first century strongly resembles the world we see today.

At the time, the Malthusian message of *Limits to Growth* seemed more than plausible, and if taken seriously, seemed to nail a lid on the human future. To my eyes, however, every model in *Limits* shared a crucial defect: When the authors framed their models of world dynamics, they included only the Earth. That is to say, the authors had set aside as irrelevant almost the whole of the universe—and at a time when men still walked on the Moon and looked far beyond. At the time, NASA promised low-cost access to space. At the time, bold dreams flourished and the world beyond Earth seemed within practical reach.

The restricted vision embodied in *Limits to Growth* raised questions that led me to explore what might be found outside the world it had framed—to look outward, at first, toward deep space, but later inward, to explore the potential of technologies in the nanoscale world.

With the end of high school less than a year away, I applied to MIT. My grades weren't outstanding, but I tested well, and that proved to be enough.

At MIT I soon felt that I had come home; people understood what I said and filled gaps in my knowledge, and the libraries seemed endless.

At first, I found few who shared my view that free-space development had potential importance, while planetary surfaces were a distraction. The seeming lack of discussion of the subject gave me reason to doubt my previous confidence. Had I been mistaken about the promise of the space frontier? Or could it be that my better-informed elders had overlooked something important, that they had asked the wrong questions?

Indeed, I found that few had asked the right questions, and therefore few had considered and weighed the potential answers. Engineers and space planners, at NASA and elsewhere, had asked "How can we explore and survive on other planets—the places in space most like the Earth?" The question I asked was, "Where can we find an environment that can sustain a vibrant industrial civilization?" This different question had a different answer, and free-space development had no connection with distant planets.

In search of someone who shared this vision of the latent potential of the world beyond Earth, I asked my freshman adviser to direct me to someone who might know someone who knew something about this sort of idea. He did, and the professors he suggested both directed me to a revered MIT physicist, Philip Morrison. After the second recommendation, I gained the courage to knock on his door. He did indeed know something, and someone.

Professor Morrison directed me to a professor of physics at Princeton, Gerard K. O'Neill, who (as it happened) was planning a conference centered on his vision of space development. This vision had something in common with my own line of thought. It set planets aside in favor of space itself as a place for Earth's life, and it proposed ways to avoid a cataclysmic collision between human civilization and Earth's limits to growth. From there, however, his vision gave less weight to concerns with materials and manufacturing, and highlighted instead a concept that strongly engaged the public's imagination—a grand and very *visual* vision of new lands built in free space itself.* O'Neill had published calculations of geometry, light, atmospheric pressure, centripetal acceleration, and structural mass for vast cylindrical structures—kilometers in scale—all based on the known properties of sunlight, glass, mirrors, and suspension-bridge steel. These space habitats were to be open spaces large enough to hold cities and farms, places with sunlit lands, the feel of gravity underfoot, and, with the passage of years, forests. Perhaps most important of all, this concept inspired artists to portray visions of places in space that looked much like home, images that gripped the public imagination

As a freshman, I found myself playing a minor role in organizing the first Princeton conference on space colonization—a term NASA later amended to "space settlement" at the request of the State Department. As a result of this meeting, a community began to coalesce, an eclectic mix that ranged from undergraduates and scientists to aerospace systems engineers and environmental activists. Study groups and summer studies followed, together with reports, conference papers, press coverage, critics, and even a popular movement of sorts.

The vision of space settlement had deep resonance at a time when society had begun to question the material foundations of its own existence. A common concern about terrestrial limits to growth animated the space movement. Space is large, holding room and resources enough to open up realms for life on a scale of a thousand Earths. This physical potential suggested a path for civilization that could avert overshoot and catastrophe for centuries to come.** What's more, free-space development could lift the burden of industry from the biosphere and make room in the world for restoring the Earth.

The mid-1970s was the time of "the energy crisis," when OPEC-created oil shortages had highlighted the idea of terrestrial limits and thereby spurred a search for ways to transcend them. Engineers proposed that solar power beamed from space could compete with terrestrial sources of energy, and so NASA and the Department of Energy provided research funds to aerospace firms to support exploratory design and analysis of potential solar power satellite systems. The idea of building these massive satellites using resources already in space had appeal and soon became part of the space settlement concept.

This kind of large-scale construction would require space-based manufacturing, and a comprehensive infrastructure for space industrialization.

SCIENCE AND SPACE FOR MANUFACTURING

Manufacturing makes modern society possible. Food, clothing, shelter, travel, communications, and the conveniences of daily life—in the developed world all these rely on industrial products made by what are now increasingly automated processes. Societies in space would depend on industry to an even greater extent, in fact, for producing every bit of material, even soil and air. This is why the practical questions of space settlement quickly turned toward questions of mining, refining, and manufacturing.

At MIT I majored in a program called “Interdisciplinary Science,” yet much of my study revolved around industry and agriculture—how they work on Earth and how their technologies could be recast to fit a radically different environment. To explore this area required understanding facts and principles from a wide range of fields. Some of these principles described macroscale phenomena, like the physics of heat and mass transfer; others led to the molecular world and the foundations of materials science. Yet other topics included meteoritics and planetary science, plant physiology and ecosystem engineering, photovoltaic cells and solar energy, vacuum metallurgy and the distillation of steel, the properties of materials that can withstand the incandescent temperatures of a solar furnace, and ways of stitching together terrestrial industrial processes to make glass and metals from lunar rock.

One line of study led me into the nanoscale world: designs of lightsails—rotating structures, kilometers wide, tiled with sheets of thin metal film, capable of harnessing the pressure of sunlight to drive vehicles through space, year after year, with a small but steady acceleration and no need for fuel.

Lightsails held my attention for several years (and a thesis) at MIT. Physical data showed that lightsails could catch and reflect sunlight using sheets of aluminum no more than 100 nanometers thick, or about 300 atomic diameters. Calculations and library research persuaded me that films of this thickness would serve their purpose if they could be made and installed in the vacuum of space, yet no calculation could persuade me that such delicate films could survive a manufacturing process. And so I learned to use vacuum equipment to deposit vaporized aluminum onto a surface, forming thin films, atom by atom. The films were indeed extremely delicate. In trying to free them, I tore apart one film after another. If freed and then touched, the thin metal film would mirror-coat a fingertip, wrapping each ridge in the skin and yet feeling like nothing at all. Set free in the air, a torn fragment of film would drift like a dust mote, yet reflect light like a scrap of aluminum foil. In the end, I learned to lift and mount these thin films in frames (and even took samples to Pasadena for a presentation at NASA’s Jet Propulsion Laboratory), and through this hands-on experience I learned enough to conclude that automated machines in the space environment could indeed produce lightsail film in enormous quantities.

The method I learned came from the shelves of the MIT Science Library, while the science I learned showed me how things could be built from the bottom up, atom by atom.

INTERLUDE: ARTHUR KANTROWITZ

Early in those years the MIT Space Habitat Study Group grew out of a talk I gave on space settlement. Most members were students, but at a meeting one evening, a gray-haired man walked in and stayed in my life.

Arthur Kantrowitz was a physicist and engineer, an MIT Institute Professor (Visiting), the founder and head of the Avco Everett Research Laboratory, and, I think, a wise man. Born in 1913, he was older then than I am today. He became my mentor and friend.

Over the years, Arthur shaped my view of the world, how it works, and what matters. He helped me understand the underlying nature of scientific knowledge and scientific norms, as well as the turbulent process that leads to new technologies. He shared his knowledge of the dirty side of that process, the secrecy and corruption that can flourish at the junction of policy, money, and technology. And beyond this, he shared his understanding of the underlying incentives and cultural problems, and his experience with attempts at institutional reforms.

As I look back, I can see how much of my sense of the world reflects his values.

Arthur had achieved bold and wide-ranging accomplishments in technology. In the 1950s, his inventions helped solve the problem of hypersonic atmospheric re-entry (the *New York Times* described him as "one of the first technological heroes of the space program"), yet in his youth practical aeronautics still centered on biplanes built of wood and cloth.

While leading research and development teams, Arthur pioneered a range of technologies that included high-power lasers, supersonic molecular beams, magnetohydrodynamic generators, and (with his brother, Adrian) the intra-aortic balloon pump, a heart-assist device now used in every major cardiac surgery center.

His bold visions started early. In 1939, with a colleague at what is now NASA's Langley Research Center, Arthur built the first laboratory machine to explore the potential of magnetically confined nuclear fusion power; in 1963, he reexamined the field and concluded that the entire approach faced a brick wall—nonlinear plasma instabilities—that to this day, a half century later, has not been surmounted. Arthur was bold and persistent, and he knew when to quit.

Arthur had experience with the space program from the inside and at high levels. At the inception of the Apollo program, for example, he served on a presidential commission that assessed the prospective costs, times, and development risks of competing approaches for reaching the Moon. In the 1970s, Arthur took a keen interest in the drive toward space development, giving talks, supporting organizations, leading research in high-capacity, small-payload space launch systems, and advising an MIT undergraduate who absorbed at least some of what he could teach.

It was Arthur who introduced me to the works of Karl Popper, the philosopher of science who established the principle that science can test ideas and sometimes approach the truth, yet can never prove a universally quantified theory. Popper called for an intellectual life of bold conjectures, tentatively held and subject to critical discussion and stringent testing. Grappling with Popper's view of epistemology (and with books by his critics) led me to a lifelong concern with the basis of knowledge in both science and engineering, and through this concern, to methodologies that have guided my life's work in exploring the potential of physical technologies.

Arthur was a man of both the future and the past. In a time of growing specialization, he was a generalist. In a time of growing timidity, he was bold. In a time of science increasingly driven by funding and politics, Arthur was a voice for the deeper values that make science work.

Because of Arthur, however, I misjudged the world. In a tacit, unconscious way, I assumed that science held many more people like him.

At the age of ninety-five, Arthur Kantrowitz died of a heart attack while visiting his family in New York. His last hours were good, I'm told, hours spent with his family while his life was sustained by a device he knew well, the intra-aortic balloon pump. I miss him more deeply than I would ever have guessed.

A CULTURE OF QUANTITATIVE DREAMS

My years of engagement with Arthur and others in the space systems community taught me a way of thinking that harnessed creative vision to physical, quantitative reasoning in order to explore what could be achieved in new domains of engineering.

The space systems engineering community has evolved together with the space systems themselves. Satellite launchers and moonships grew out of quantifiable engineering visions, system-level concepts that could be sketched, assessed, and discarded at a rapid pace, evolving through a kind of Darwinian competition. The best concepts would win the resources of time and attention needed to fill in more details, to optimize designs, to apply closer analysis, and after this refinement and testing, to compete again. The prize at the end would be a design refined into fully detailed specifications, then metal cut on a factory floor, then a pillar of fire rising into the sky bearing a vision made real.

For example, before President Kennedy committed the United States to the Apollo program, engineers had examined hundreds of ways of assembling rocket engines and fuel tanks to build systems that could reach the Moon. Much the same can be said of how ideas have evolved before every major space mission.

To play this game well requires creativity harnessed to skeptical evaluation, with attention to both knowledge and uncertainty. In a system-level engineering design—whether a sketch or a more detailed specification—every assumption, calculation, and uncertainty range can be critiqued. Uncertainties can be fatal or minor; some can be hedged, while others spur new research programs. In opening the space frontier, for example, one research program established how a spacecraft could survive a return to the Earth at hypersonic speeds through air heated to temperatures found on the Sun—the re-entry problem that Arthur addressed—thereby answering skeptics and squeezing engineering uncertainty into ranges narrow enough to enable more detailed and confident system design.

The space systems engineering culture shaped how I thought about problems at the junction of complexity, uncertainty, and exploratory design. It was from this milieu that I turned my attention to molecular technologies.

In those years, the space development community, supported by federal funding, explored decades-long plans for developing solar power satellites and space habitats, to be enabled by lower-cost successors to the yet-unbuilt Shuttle. Lightsails could play a role in that world, yet I found my attention drawn away in a different direction, toward the exploration of the potential of smaller and more complex things—not broad, nanometers-thick films of aluminum, but nanoscale devices and machines, the potential fruits of advances in molecular technologies.[*]

Once again, in the molecular sciences, it seemed to me that the experts were focused on different questions than those of the greatest long-term importance, that they were too close to their work to see where their fields could lead, if combined and applied to new ends. And as with exploring the potential of space, the questions and answers once again involved system-level engineering principles, and exploring how one might make things in an unfamiliar world. And once again, I found implications for the human future on a scale too large to ignore.

The same sense of mission that led me to explore the potential of space now pulled me toward the molecular world. The scientific knowledge already in place was enormous, and growing.

---

[*] “Molecular Engineering: An Approach to the Development of General Capabilities for Molecular Manipulation,” *Proceedings of the National Academy of Sciences USA*, 78, no. 9 (1981): 5275–5278.

[*] *Engines of Creation: The Coming Era of Nanotechnology* (New York: Doubleday, 1986); and *Nanosystems: Molecular Machinery, Manufacturing, and Computation* (Hoboken, NJ: Wiley/Interscience, 1992).

[*] Donella H. Meadows, *The Limits to Growth. A Report for the Club of Rome’s Project on the Predicament of Mankind* (New York: Universe, 1972).

[*] In fact, O’Neill imagined building with resources mined from a visible source that loomed large in the human imagination, the Moon, while I advocated using the more attractive resources offered by the less charismatic asteroids; the concept of asteroid mining at first gained little traction, yet missions to asteroids have become part of NASA’s plans, now slated for 2025, before any return to the Moon.

[**] Not forever, of course, because in the end Malthus was right. Buying time for dozens of generations, however, seemed reason enough, and with this, perhaps time enough for humanity to gain wisdom enough to face future limits with a measure of grace. Stranger things have happened in the long arc of history.

[*] The learning that prompted this line of thought came from reading journals like *Science* and *Nature*, and from dipping into more specialized journals, such as *Angewandte Chemie* (all of which I found, of course, on MIT’s library shelves).

CHAPTER 3

From Molecules to Nanosystems

THE IDEA OF BUILDING THINGS with atomic precision often strikes people as futuristic, yet atomically precise fabrication has a longer history than spaceflight, or even wooden biplanes. The story of atomically precise fabrication begins more than a century ago, at the start of an arc of accelerating progress.

By 1899, chemists had gained skill in building structures with atomic precision, and they knew what they were doing well enough to draw correct diagrams of molecules, atom by atom and bond by bond. Chemists knew, for example, that carbon atoms form four bonds, typically directed toward the corners of a tetrahedron, and that molecules therefore can have chirality, that is, they can have both left- and right-handed forms. They knew that carbon can form double bonds and that benzene consists of a ring of six equivalent carbon atoms, and they had also inferred how methyl groups could be attached to those rings in the patterns that define different molecular isomers. This was a remarkable degree of knowledge, considering that no one could yet actually see a molecule.

During that time, chemists were developing the first systematic methods of altering molecular structures, and they used the resulting changes to infer the structures of the molecules themselves—a special form of learning by doing.

The idea of atoms, of course, had been around since antiquity. In Greece circa 400 BCE Democritus argued that matter must ultimately consist of indivisible particles—as indeed atoms are, barring nuclear reactions. In Rome circa 50 BCE Lucretius argued the same case in considerable depth and suggested that dust motes that could be seen dancing in sunbeams were, in fact, driven by what is now called “Brownian motion,” the effect of collisions with atoms (and for some of the motions he saw, he was right). Today, the most advanced forms of atomically precise fabrication rely on this Brownian dance to move molecules.

After classical times, centuries passed before any further progress was achieved in understanding the atomic basis of the material world. Inquiry reached a landmark in England in the early 1800s when John Dalton observed that chemical reactions combined substances in fixed proportions and explained these proportions in terms of atoms. Dalton postulated that each pure chemical compound consisted of particles —“molecules”—each containing a fixed number of one or more kinds of atoms. Reasoning from observed proportions, chemists applied this principle to infer the atomic composition of molecules, eventually deriving the chemical formula CO_2 for carbon dioxide, H_2O for water, and so on. Along another line of inquiry, the laws that describe how gases expand and contract in response to changes in pressure and temperature were explained in terms of molecular motions driven by thermal energy, the same thermal motions that cause the Brownian dance.

This kind of indirect evidence, steadily augmented with observations of the results of chemical reactions (in fact, tens, hundreds, thousands of reactions), was what led chemists to formulate and test hypotheses about the atomic structure of molecules. The systematic experiments gave rise to systematic techniques for organic synthesis, a technology of atomically precise fabrication that has changed industry, medicine, and daily life. The most impressive examples of atomically precise structures, however, came from biology. And what's more, some of these structures were functional devices that came to be known as *molecular machines*.

The concept of molecular machines dates to the mid-twentieth century and emerged out of efforts to understand how enzymes worked and how biomolecules fit together. Indeed, as early as 1890, the German chemist Hermann Emil Fischer had suggested a "lock-and-key" model for how enzymes select specific molecular substrates out of the sea of different molecules in a cell; his suggestion provided the first inkling of how complementary macromolecular shapes could enable specific parts to fit together and perform useful operations.

Since the 1950s, molecular biologists have expanded and deepened our understanding of how large molecules—including nanoscale objects made of protein—bind, move, and perform useful functions, like copying a strand of DNA in a cell's nucleus, or pulling protein fibers in a muscle to move a leg. Over the years, more and more biomolecular structures have been mapped in atomic detail, first a few in the 1950s, and today, tens of thousands, earning Nobel Prizes for James Watson, Francis Crick, and Maurice Wilkins for their discovery of the structure of DNA and for John Kendrew and Max Perutz for their use of X-ray diffraction techniques to provide the first atomically precise maps of protein structures.

This emerging knowledge of biomolecular machinery intrigued Richard Feynman. At Caltech in 1959, speaking after dinner at a meeting of the American Physical Society, Feynman discussed the physics of artificial micro- and nanoscale machinery, "inspired by the biological phenomena in which chemical forces are used in a repetitious fashion to produce all kinds of weird effects (one of which is the author)." In his talk, "There's Plenty of Room at the Bottom," Feynman proposed the idea of using machine-guided motion to assemble molecular structures with atomic precision.

Thus, in 1959, Feynman had outlined the fundamental physical principle of atomically precise manufacturing. The idea of using machines to build with atomic precision, however, then lay fallow for more than a decade and a half, while the biomolecular sciences moved forward.

By the mid-1970s, biomolecular machine engineering was already on the horizon. Scientists were beginning to write instructions coded in DNA, founding a field they called "genetic engineering." Through this technology, scientists learned how to reprogram the molecular machinery of cells to produce new proteins—or, to speak more precisely, they had learned how to program cells to produce proteins already made by other cells.

Genetic engineering and molecular biology pioneered new fields of science and technology, but they could also be seen as opening the door to manufacturing in a new environment and on a scale that could be important for the future of humanity. I followed this field with particular attention, and by 1976 my thoughts were drawn to the question of where it might lead. (Libraries were the cause once again. As an information omnivore, I'd been casting a net into the flow of knowledge that crossed the new-journals shelves of the MIT Science Library.)

The following spring, after toying with ideas about computing with molecular devices, I found myself asking several crucial questions—not just "What could be built by programming nature's machines?" but a question a step beyond: "What could be built using the machines that nature's own machines could be programmed to build?" And then, another question a further step beyond: "What could be built using machines that could be built using *those* machines?" and so on, looking up toward the heights of a dizzying spiral of ever more capable fabrication technologies.

This upward spiral leads toward powerful manufacturing capabilities, atomically precise and yet thoroughly nonbiological, capabilities limited not by the properties of the biomolecular materials and devices that nature has evolved, but by the properties of materials and devices within bounds set only by the limits of physical law. In other words, this concept of an upward spiral suggested a path from today's technologies to advanced APM, a path based on using atomically precise fabrication technologies to build better tools for atomically precise fabrication—and a path that could begin by building with tools already at hand.

Much of the progress that has been achieved since the 1970s builds directly on biomolecular machines and materials, and this includes the rise of a field called "protein engineering." To understand the implications of progress in protein engineering, however, requires breaking away from a natural misconception.

It is tempting to think of protein molecules as watery, gelatinous stuff like meat, but this idea is highly misleading. Protein molecules are solid, nanoscale objects, much like bits of plastic, but with more diverse and intricate structures. In fact, they consist of folded polymer chains built from a kit of twenty distinct monomers that differ in terms of size, shape, and physical properties. In different combinations and sequences, these monomers can form materials as diverse as soft rubber, hard plastic, and fibers stronger than steel (spider silk, for example, is made of protein, as is the horn of a bull). But what matters most to an engineer, however, is what these nanoscale objects can *do*.

Nature shows some of the possibilities. Looking at the molecular machinery of life, we find that proteins can fit together to form motors, sensors, structural frameworks, and catalytic devices that transform molecules; protein-based devices also copy and transcribe data stored in DNA. Most important of all, machine systems built of biomolecules can serve as programmable manufacturing systems that build components for new molecular machines.

Examples like these made it clear from the start that genetic engineering offered access to tools that could make all of these things, and much more, if we mastered the arts of protein engineering.

Confronted with these facts, my thinking went something like this:

1. Nature shows that molecular machine systems can be programmed by instructions encoded in DNA to build complex, atomically precise structures, including components that fit together to form molecular machine systems.
2. Nature also shows that molecular machine systems can bind and position a wide range of reactive molecules, guiding their encounters in order to build atomically precise biomolecular structures and machine components.
3. Similar machine systems could be used to bind, position, and combine an even wider range of reactive molecules, not all found in biology, and thereby build a greater range of atomically precise structures, including machine components that are more densely bonded and hence more robust.
4. These more robust next-generation components could be used to build robust and higher-performance production machinery, which in turn could be used to build a yet wider range of components, and from these components yet more capable production machines, and so on, extending toward a horizon far beyond biology.

In the end, to look toward this horizon means asking what physical law itself allows, and from this perspective—and viewing the landscape through the lens of systems-level engineering—I got a first glimpse of the potential power and scope of atomically precise manufacturing.

The prospects were startling. Indeed, I found them hard to believe, yet over time, studies based on exploratory design concepts, calculations, and the knowledge I gleaned from textbooks and journals persuaded me that the startling prospects were entirely realistic.

Driven by a mission to explore and share insights about the potential of these world-changing technologies, I moved to publish. First came a scientific paper in 1981,* followed five years (and three drafts) later by a book for the general public, *Engines of Creation*.

Published in September, 1986, *Engines of Creation* introduced a new concept and word into public discussion: “nanotechnology.” The press formed an impression of what this might mean and ran with it.

Two months later, a leading general-audience, science-oriented magazine of the day, *OMNI*, confronted a million readers with a blazing cover headline: “NANOTECHNOLOGY: MOLECULAR MACHINES THAT MIMIC LIFE.”

The article’s author, Fred Hapgood, had been with the MIT Nanotechnology Study Group I had founded the year before. This enormous (and unsolicited) kickoff launched a wave of stories in newspapers and magazines that brought the concept to a wide audience, while the article’s biological spin (MACHINES THAT MIMIC LIFE) marked a trend that grew into a problem—analogies to biology became a simplistic distorting lens though which nanomachines were mistaken for nanobugs.

Time passed, and as ideas about nanotechnology echoed and spread throughout society, they evolved and diversified to exploit a range of memetic niches. By the early 1990s, the initial, revolutionary vision of nanotechnology had launched a wave of excitement for everything “nano,” and although that excitement took various forms, one form became central. As it made its way into science, the vision of nanotechnology spurred a fresh, more unified focus on nanoscale phenomena, both within science and among the general public, and it gradually grew into a surge of support for new research initiatives.

“Nanotechnology” broadened to embrace far more than nanomachines and atomically precise fabrication. It became a generic term defined primarily by size. This new, generic brand of nanotechnology (often better called “nanoscience”) spanned a host of fields that worked with nanoscale structures, and it brought together researchers working with materials, surfaces, small particles, and electronic devices. They shared concepts and techniques, formed collaborations, and explored new frontiers in science and technology. The long-range vision of advanced nanotechnology excited the public, while a growing understanding of the practical importance of nanoscale phenomena stimulated the growth of both research and funding.

The story of nanotechnology, and of APM, soon became entangled in the special kinds of confusion that thrive at the borders between engineering and science. The problems stemmed from contrasts between the two that are profound yet often unrecognized.

Scientists and engineers, on the whole, are different species and have different approaches to knowledge. Scientists inquire; engineers design. Scientists study physical things, then describe them; engineers describe physical things, then build them. Scientists and engineers ask different questions and seek different answers.

In saying this, I am painting with a broad brush; a more nuanced approach recognizes that inquiry and design are often integral to a single line of research and may mesh within a single mind as it follows a single train of thought. In Chapter 8 I will draw more careful contours around questions of knowledge, practice, and culture, exploring the contrasting parts of the engine of progress that drives the modern world.

Experience shows that mistaking engineering questions for scientific questions can create a conceptual muddle. In the molecular sciences, in particular, these two modes of thought are in essence distinct, yet inextricably linked in practice and easily blurred. The costs have included missed opportunities to apply scientific knowledge to open new fields.

Scientific inquiry long ago uncovered the fundamental principles of molecular physics, and these principles enable a vast range of reliable, predictive calculations.

Experimental science, however, provides knowledge—and *know-how*—beyond reach of calculation. Laboratory researchers develop what are often hard-won techniques for building atomically precise nanoscale structures, and in the course of their research the tasks of learning and making become deeply entangled. Indeed, in the beginning, it was making molecules that enabled chemists to discover atoms and bonds, long before quantum mechanics provided an explanation.

Thus, when my work led me into the land of the molecular sciences, I found a culture in which questions of inquiry and design were often confused, a culture in which most researchers scarcely recognized the concepts and methods of systems-level engineering. Nonetheless, I found that abstract engineering concepts had direct applications.

Consider, for example, the pattern of thought then prevalent in protein science. A scientific problem—given a monomer sequence, predict how a protein will fold—had been confused with an engineering problem—given a desired fold, design a monomer sequence that will produce it. At the time, fold prediction was an unsolved problem (and remains only partially solved today), and researchers had implicitly assumed that successful fold prediction must precede fold design.

My 1981 paper, however, explained why design and prediction were fundamentally different problems and why design should be less challenging. Fold design was soon dubbed "the inverse folding problem," and this deep, elementary idea launched the field of protein engineering.

Protein engineering, however, remained embedded in science. Speaking at a conference a decade later, I asked for a show of hands: "How many of you consider yourselves to be scientists?" and about one hundred hands went up. When I then asked, "How many of you consider yourselves to be engineers?" the total was no more than three. And this at a conference convened with the title topic, "Protein Engineering."

Fields differ, of course. If I had asked the same questions when speaking to an audience of experimental physicists, I suspect that many would have raised their hands twice, and likewise in talks I have given to space scientists tasked (for example) with sending instrument systems to Mars.

As a rule, however, one finds that engineers and scientists have contrasting cognitive habits, intellectual values, and cultures; the contrast is particularly sharp where science centers, not on complex systems like spacecraft or particle accelerators, but instead on laboratory experiments using equipment appropriate to the molecular sciences (think of beakers, pipettes, commercially available infrared spectrographs, and the like). In the molecular sciences, most researchers have had no reason to learn the arts of systems-level engineering design.

Thus, the modes of thought that were best suited to the needs of molecular research were ill-suited to the task of grasping or judging abstract engineering analysis, while the researchers who could easily understand the scientific basis for atomically precise manufacturing were liable to slip into confusion about the concept itself, and misjudge it. Then came a push, a pull, and a slide into a conceptual pit.

THE ENGINES OF CONFLICT

The vision presented in *Engines of Creation* unleashed forces that soon came into conflict. The first, of course, was the force of vision itself, which spurred studies that deepened our understanding of the prospects. As part of this effort, my own attention had turned toward preparing my doctoral dissertation, and from this, *Nanosystems.*

The many-sided force of popularization, however, had a head start of six years. As they echoed through the press and Internet newsgroups, the concepts presented in *Engines* devolved into vivid, simplistic stories and images and the science and engineering lost ground to fiction that shaded off into ideas that amounted to magic. Utopias and scare stories took form and gained strength. From a distance, the concepts in *Engines* were perceived through layers of distortion and science-free fiction.

Excitement and popularization presented a risk because when presented in a brief or distorted summary, APM-level technologies seemed like hype gone wild, and when wrapped in popular enthusiasm, the concepts seemed to be no more than another delusion of crowds. Ironically, what makes these technologies important—the radical scope of their implications—makes them hard to credit for reasons that are correct at least 99 percent of the time. Heuristics, however, sometimes go wrong, and this is an instance.

In the early days, however, from where I stood, it seemed that enthusiasm was primarily a positive force—that it would (as it did) channel support toward scientific progress, and that scientists, in turn, would surely help to channel enthusiasm toward reality and gradually push nonsense into the background.

Indeed, this happened to some extent as reality-based thinking continued to advance. Students read *Engines* and turned their careers toward nanotechnology; researchers at Caltech and elsewhere applied computational methods to study advanced AP machinery; I spoke at scientific conferences, corporate meetings, the White House Office of Science and Technology, the Pentagon, NSA, the Congressional Office of Technology Assessment, and at a Senate hearing convened by then Senator Al Gore.

By the end of the decade, researchers in a host of fields (but weighted toward materials science) had gathered under the banner of nanotechnology, both promoting and stretching the vision attached to the word. By the late 1990s, support for the resulting, greatly broadened kind of nanotechnology had reached the threshold of launching a federal program, making ownership of “nanotechnology” a billion-dollar prize.

Near-term research and longer-term objectives were entirely compatible, or should have been, yet as events unfolded, a conflict emerged, feeding on clashes between popular visions and near-term scientific realities. This conflict polarized, taking on an us-vs.-them, and science-vs.-fantasy tone, while the distinction between fantasy and genuine prospects was increasingly lost in the noise. A turning point came when the new federal program’s promoters secured funding to develop atomically precise fabrication: They turned against the vision they had sold to Congress, redefined their mission, and launched a strange and confused war of ideas that still echoes today. I will pick up this story in Chapter 13.

The conflict had a particularly perverse effect. It severed “nanotechnology,” as widely perceived, from the concept of atomic precision and its natural roots in the molecular sciences. I had outlined a path toward APM that led forward from existing capabilities for atomically precise fabrication, yet strong, continuing progress on that very path somehow slipped from attention. This has set up the world for a surprise.

Consider how far we have already come. In 1986 neither protein engineering nor structural DNA technologies had yet been demonstrated, no one had yet used a machine to move individual atoms, and the largest complex, artificial atomically precise structures were no more than a few hundred atoms in scale. Since then, atomically precise fabrication technologies have made great progress on multiple fronts:

- Researchers now routinely use scanning probe instruments to image and place individual atoms and to maneuver and bond individual molecules. This level of control has demonstrated the principle of mechanically directed atomically precise fabrication.
- Organic chemists have built steadily larger and more complex structures along with motors and other machines; their techniques now provide a rich toolkit for building molecular systems, while inorganic chemists and materials scientists have expanded a complementary toolkit of nanoscale structures.
- Protein engineering has flourished, supported by computer-aided design software, and now enables the routine design of intricate, atomically precise nanoscale objects, including structural components and functional devices.

- Structural DNA nanotechnology has emerged and now enables rapid and systematic fabrication of addressable, atomically precise frameworks on a scale of hundreds of nanometers and millions of atoms.
- Quantum methods in chemistry have advanced together with the power of computers and algorithms, providing powerful, physics-based tools for scientific modeling and molecular engineering.
- Molecular mechanics methods in chemistry can now describe the structure and dynamics of molecules on scales that reach millions of atoms, a range that can enable the design and development of complex, atomically precise systems.

The current state of the art is more than enough to support a drive for next-generation molecular systems on the road to atomically precise manufacturing. Indeed, I am persuaded that advances in recent years now provide a platform that can support rapid progress. The greatest challenge today is to put the pieces together—not only components and computational tools, but also engineering concepts and the research teams that can bring them into physical reality.

INTERLUDE: PROSPECTS AND CHALLENGES

Stepping back for a moment, let us ask a question: “Where do we stand today as we consider APM-linked technological choices that could change the shape of our future?” The issues reach far beyond laboratories, politics, and molecules.

In brief, the APM-based production revolution promises to transform the material basis of human life with far-reaching consequences that include both new solutions and new problems of global scope.

Consider the challenge of resource scarcity (minerals, petroleum, water) and the challenge of environmental problems that range from toxic metal emissions to global climate change. These are physical problems that have potential physical solutions. Through a chain of physical and economic links, the APM-based production revolution can transform global problems by slashing resource consumption and toxic emissions and by providing the infrastructure for low-cost solar energy and a carbon-neutral economy (and even more remarkably, by providing affordable means for removing carbon already released into the atmosphere).

These physical capacities could solve critical problems and enable us to live more lightly on Earth, while radically raising the material standard of living worldwide. These solutions bring problems, however. In particular, rapid deployment of this range of capabilities would lead to deep, pervasive disruptions in the global economy, beginning with mining, manufacturing, and trade, and spreading outward from there.

How would APM have these far-reaching effects? In manufacturing, APM-based technologies can produce better products at far lower cost than today, out-competing existing industries. As for mining, APM naturally consumes and produces a different mix of materials (no need for the iron and chromium in stainless steel, no need for lead and tin in solder) and as it happens, the most useful elements—including carbon, nitrogen, oxygen, and silicon—are not at all scarce. On grounds of performance and cost, even common structural materials will be subject to widespread displacement, and with them, most mines. (Chapter 11 explores questions of APM-based product performance, cost, and resource requirements in greater depth.)

Today, trade builds global supply chains that lead from mines and wells to smelters and refineries, to materials processing plants, to networks of factories that shape and assemble components to make final products. As we will see, with APM-based technologies it would be natural for these long, specialized supply chains to collapse to a few steps of local production, progressing from common materials to simple chemical feedstocks; from simple feedstocks to generic, microscale building blocks;

and then from generic components to an endless range of products, much as printers can be used to arrange generic pixels to form an endless range of images.

Long, specialized supply chains drive the physical trade that today joins the world into a global economy, and collapsing supply chains would cause that trade to decline. One can easily imagine disruptions in trade that would affect the livelihood of half the planet or more. And one can easily imagine a level of suffering and scarcity in the midst of potential abundance.

I think that this prospect points to a need for exploring policies for managing what could be a catastrophic success. In other words, it calls for a conversation that considers prospects for our world as the physical potential of APM-level technologies crosses the threshold into physical reality, a conversation that was interrupted more than a decade ago and must now be renewed.

Implications for the military sphere, in particular, demand careful consideration because easy, unconsidered policies would bring great and needless risks. Here the nature of potential products (and of the potential dynamics of their development, production, deployment, and use) will have profound implications that call for fresh thinking. The economic implications of an APM transition likewise call for a reassessment of national interests as deep and broad as the prospective changes in the material economy.

Enough can be understood today to reframe global problems and raise new concerns. In a world on the path to profound transformations, our situation calls for asking unusual questions about our prospects and how we might best respond—new questions of how to avoid needless risks, resolve difficult global problems, grasp unexpected opportunities, and manage disruptive change.

In short, we need to begin to broaden the agenda for conversations about the future—not to change widespread premises in an instant, but to begin to assess the prospects for APM-level technologies and the questions they raise regarding challenges and prudent near-term choices.

* * *

MY AIM, HOWEVER, is not to overturn anyone’s worldview. Prospects for radical abundance deliver a banquet of almost indigestible truths—or so the prospects seem to me, even now. Digesting and integrating new information will take time and the contributions of many minds.

My aim is, therefore, more modest: to outline facts about what is truly possible, to discuss where technology may lead in the coming years, and to consider some critically important questions that have not yet been asked. In light of the prospects ahead, I think that it’s time to begin a new conversation about our future, a conversation that begins to explore the prospects for radical abundance.

---

[*] The paper, in the *Proceedings of the National Academy of Sciences*, came to be widely cited in the scientific literature as a foundation for the concepts of both protein engineering and advanced nanotechnologies based on machine-guided molecular assembly.

PART 2

THE REVOLUTION IN CONTEXT

CHAPTER 4

Three Revolutions, and a Fourth

THROUGH NEW TECHNOLOGIES, human history has repeatedly changed directions, and with unimaginable consequences. The Agricultural Revolution of the Neolithic era marked the advent of a way of life that enabled dense, settled populations to feed themselves and opened the way for the rise of civilization. The Industrial Revolution launched an explosive wave of physical products that brought with it the modern world. The Information Revolution, still unfolding today, has rewoven the fabric of knowledge, commerce, and society, setting the stage for a future whose outlines are still beyond knowing.

The approaching APM Revolution will provide the driving force for a fourth revolution, and like the preceding revolutions, it will transform daily life, labor, and the structure of society on Earth.

Past revolutions offer lessons that can help us grasp the nature of the revolution ahead. Each has had pervasive impacts on the human world, bringing profound and far-reaching change, and the very nature of their technologies can help us understand how the APM revolution compares with those of the past.

The Agricultural Revolution was based on molecular machinery (though not machines made by design); the Industrial Revolution was based on machines made by design (yet far from molecular); and the Information Revolution was based on digital, nanoscale devices (but devices that process only information, not matter). In a sense, the APM Revolution draws from all three of these. It employs artificial, molecular, nanoscale machinery that operates on digital principles (but this time, in order to process matter). It should be no surprise, then, that each revolution holds lessons for understanding the APM Revolution and why its advent will mark a divide in human history.

As Winston Churchill once said, “The farther backward you can look, the farther forward you are likely to see.” Here, we can begin with the Neolithic era.

THE FIRST AGRICULTURAL REVOLUTION

The first agricultural revolution—*the* Agricultural Revolution—defines the dawn of the Neolithic era, more than ten thousand years in our past.

The Agricultural Revolution gave human beings a new way to exploit the productive capabilities of the nanoscale machines found in living organisms, the molecular metabolic machinery that makes complex structures out of nothing more