


WHERE GOOD IDEAS COME FROM



*the natural history
of innovation*

STEVEN JOHNSON

BESTSELLING AUTHOR OF
Everything bad is good for you

WHERE GOOD IDEAS COME FROM

A Natural History of Innovation



STEVEN JOHNSON



ALLEN LANE
an imprint of
PENGUIN BOOKS

ALLEN LANE

Published by the Penguin Group

Penguin Books Ltd, 80 Strand, London WC2R 0RL, England
Penguin Group (USA) Inc., 375 Hudson Street, New York, New York 10014,
USA

Penguin Group (Canada), 90 Eglinton Avenue East, Suite 700, Toronto,
Ontario, Canada M4P 2Y3 (a division of Pearson Penguin Canada Inc.)

Penguin Ireland, 25 St Stephen's Green, Dublin 2, Ireland (a division of
Penguin Books Ltd)

Penguin Group (Australia), 250 Camberwell Road, Camberwell, Victoria
3124, Australia (a division of Pearson Australia Group Pty Ltd)

Penguin Books India Pvt Ltd, 11 Community Centre, Panchsheel Park, New
Delhi – 110 017, India

Penguin Group (NZ), 67 Apollo Drive, Rosedale, North Shore 0632, New
Zealand (a division of Pearson New Zealand Ltd)

Penguin Books (South Africa) (Pty) Ltd, 24 Sturdee Avenue, Rosebank,
Johannesburg 2196, South Africa

Penguin Books Ltd, Registered Offices: 80 Strand, London WC2R 0RL,
England

www.penguin.com

First published in the United States of America by The Penguin Group, (USA)
Inc., 2010

First published in Great Britain by Allen Lane 2010

Copyright © Steven Johnson, 2010

The moral right of the author has been asserted

Grateful acknowledgement is made for permission to reprint from Franco
Moretti, *Graphs, Maps, Trees: Abstract Models for a Literary History*.
Verso, 2007. Copyright © Franco Moretti 2007

All rights reserved

Without limiting the rights under copyright reserved above, no part of this
publication may be reproduced, stored in or introduced into a retrieval
system, or transmitted, in any form or by any means (electronic,
mechanical, photocopying, recording or otherwise), without the prior
written permission of both the copyright owner and the above publisher of
this book

A CIP catalogue record for this book is available from the British Library

ISBN: 978-0-141-96459-1

CONTENTS

Introduction
REEF, CITY, WEB

I.
THE ADJACENT POSSIBLE

II.
LIQUID NETWORKS

III.
THE SLOW HUNCH

IV.
SERENDIPITY

V.
ERROR

VI.
EXAPTATION

VII.
PLATFORMS

Conclusion

THE FOURTH QUADRANT

Acknowledgments

Appendix: Chronology of Key Innovations, 1400–2000

Notes and Further Reading

Bibliography

Index

Introduction

REEF, CITY, WEB

*... as imagination bodies forth
The forms of things unknown, the poet's pen
Turns them to shapes and gives to airy nothing
A local habitation and a name.*

—SHAKESPEARE, *A Midsummer Night's Dream*, V.i.14–17

Darwin's Paradox

April 4, 1836. Over the eastern expanse of the Indian Ocean, the reliable northeast winds of monsoon season have begun to give way to the serene days of summer. On the Keeling Islands, two small atolls composed of twenty-seven coral islands six hundred miles west of Sumatra, the emerald waters are invitingly placid and warm, their hue enhanced by the brilliant white sand of disintegrated coral. On one stretch of shore usually guarded by stronger surf, the water is so calm that Charles Darwin wades out, under the vast blue sky of the tropics, to the edge of the live coral reef that rings the island.

For hours he stands and paddles among the crowded pageantry of the reef. Twenty-seven years old, seven thousand miles from London, Darwin is on the precipice, standing on an underwater peak ascending over an unfathomable sea. He is on the edge of an idea about the forces that built that peak, an idea that will prove to be the first great scientific insight of his career. And he has just begun exploring another hunch, still hazy and unformed, that will eventually lead to the intellectual summit of the nineteenth century.

Around him, the crowds of the coral ecosystem dart and shimmer. The sheer variety dazzles: butterflyfish, damselfish, parrotfish, Napoleon fish, angelfish; golden anthias feeding on plankton above the cauliflower blooms of the coral; the spikes and tentacles of sea urchins and anemones. The tableau delights Darwin's eye, but already his mind is reaching behind the surface display to a more profound mystery. In his account of the *Beagle's* voyage, published four years later, Darwin would write: "It is excusable to grow

enthusiastic over the infinite numbers of organic beings with which the sea of the tropics, so prodigal of life, teems; yet I must confess I think those naturalists who have described, in well-known words, the submarine grottoes decked with a thousand beauties, have indulged in rather exuberant language.”

What lingers in the back of Darwin’s mind, in the days and weeks to come, is not the beauty of the submarine grotto but rather the “infinite numbers” of organic beings. On land, the flora and fauna of the Keeling Islands are paltry at best. Among the plants, there is little but “cocoa-nut” trees, lichen, and weeds. “The list of land animals,” he writes, “is even poorer than that of the plants”: a handful of lizards, almost no true land birds, and those recent immigrants from European ships, rats. “The island has no domestic quadruped excepting the pig,” Darwin notes with disdain.

Yet just a few feet away from this desolate habitat, in the coral reef waters, an epic diversity, rivaled only by that of the rain forests, thrives. This is a true mystery. Why should the waters at the edge of an atoll support so many different livelihoods? Extract ten thousand cubic feet of water from just about anywhere in the Indian Ocean and do a full inventory on the life you find there: the list would be about as “poor” as Darwin’s account of the land animals of the Keelings. You might find a dozen fish if you were lucky. On the reef, you would be guaranteed a thousand. In Darwin’s own words, stumbling across the ecosystem of a coral reef in the middle of an ocean was like encountering a swarming oasis in the middle of a desert. We now call this phenomenon Darwin’s Paradox: so many different life forms, occupying such a vast array of ecological niches, inhabiting waters that are otherwise remarkably nutrient-poor. Coral reefs make up about one-tenth of one percent of the earth’s surface, and yet roughly a quarter of the known species of marine life make their homes there. Darwin doesn’t have those statistics available to him, standing in the lagoon in 1836, but he has seen enough of the world over the preceding four years on the *Beagle* to know there is something peculiar in the crowded

waters of the reef.

The next day, Darwin ventures to the windward side of the atoll with the *Beagle*'s captain, Vice Admiral James FitzRoy, and there they watch massive waves crash against the coral's white barrier. An ordinary European spectator, accustomed to the calmer waters of the English Channel or the Mediterranean, would be naturally drawn to the impressive crest of the surf. (The breakers, Darwin observes, are almost "equal in force [to] those during a gale of wind in the temperate regions, and never cease to rage.") But Darwin has his eye on something else—not the violent surge of water but the force that resists it: the tiny organisms that have built the reef itself.

The ocean throwing its waters over the broad reef appears an invincible, all-powerful enemy; yet we see it resisted, and even conquered, by means which at first seem most weak and inefficient. It is not that the ocean spares the rock of coral; the great fragments scattered over the reef, and heaped on the beach, whence the tall cocoa-nut springs, plainly bespeak the unrelenting power of the waves ... Yet these low, insignificant coral-islets stand and are victorious: for here another power, as an antagonist, takes part in the contest. The organic forces separate the atoms of carbonate of lime, one by one, from the foaming breakers, and unite them into a symmetrical structure. Let the hurricane tear up its thousand huge fragments; yet what will that tell against the accumulated labour of myriads of architects at work night and day, month after month?

Darwin is drawn to those minuscule architects because he believes they are the key to solving the mystery that has brought the *Beagle* to the Keeling Islands. In the Admiralty's memorandum authorizing the ship's five-year journey, one of the principal scientific directives is the investigation of atoll formation. Darwin's mentor, the brilliant geologist Charles Lyell, had recently proposed that atolls are created by undersea volcanoes that have been driven upward by powerful movements in the earth's crust. In Lyell's theory, the distinctive circular shape of an atoll emerges as coral

will, on average, live 5.5 times longer, and have a heart rate that is 5.5 times slower than the woodchuck's. As the science writer George Johnson once observed, one lovely consequence of Kleiber's law is that the number of heartbeats per lifetime tends to be stable from species to species. Bigger animals just take longer to use up their quota.

Over the ensuing decades, Kleiber's law was extended down to the microscopic scale of bacteria and cell metabolism; even plants were found to obey negative quarter-power scaling in their patterns of growth. Wherever life appeared, whenever an organism had to figure out a way to consume and distribute energy through a body, negative quarter-power scaling governed the patterns of its development.

Several years ago, the theoretical physicist Geoffrey West decided to investigate whether Kleiber's law applied to one of life's largest creations: the superorganisms of human-built cities. Did the "metabolism" of urban life slow down as cities grew in size? Was there an underlying pattern to the growth and pace of life of metropolitan systems? Working out of the legendary Santa Fe Institute, where he served as president until 2009, West assembled an international team of researchers and advisers to collect data on dozens of cities around the world, measuring everything from crime to household electrical consumption, from new patents to gasoline sales.

When they finally crunched the numbers, West and his team were delighted to discover that Kleiber's negative quarter-power scaling governed the energy and transportation growth of city living. The number of gasoline stations, gasoline sales, road surface area, the length of electrical cables: all these factors follow the exact same power law that governs the speed with which energy is expended in biological organisms. If an elephant was just a scaled-up mouse, then, from an energy perspective, a city was just a scaled-up elephant.

But the most fascinating discovery in West's research came from the data that *didn't* turn out to obey Kleiber's law.

West and his team discovered another power law lurking in their immense database of urban statistics. Every datapoint that involved creativity and innovation—patents, R&D budgets, “supercreative” professions, inventors—also followed a quarter-power law, in a way that was every bit as predictable as Kleiber’s law. But there was one fundamental difference: the quarter-power law governing innovation was *positive*, not negative. A city that was ten times larger than its neighbor wasn’t ten times more innovative; it was *seventeen* times more innovative. A metropolis fifty times bigger than a town was 130 times more innovative.

Kleiber’s law proved that as life gets bigger, it slows down. But West’s model demonstrated one crucial way in which human-built cities broke from the patterns of biological life: as cities get bigger, they generate ideas at a faster clip. This is what we call “superlinear scaling”: if creativity scaled with size in a straight, linear fashion, you would of course find more patents and inventions in a larger city, but the number of patents and inventions per capita would be stable. West’s power laws suggested something far more provocative: that despite all the noise and crowding and distraction, the average resident of a metropolis with a population of five million people was almost *three times* more creative than the average resident of a town of a hundred thousand. “Great cities are not like towns only larger,” Jane Jacobs wrote nearly fifty years ago. West’s positive quarter-power law gave that insight a mathematical foundation. Something about the environment of a big city was making its residents significantly more innovative than residents of smaller towns. But what was it?

The 10/10 Rule

The first national broadcast of a color television program took place on January 1, 1954, when NBC aired an hour-long telecast of the Tournament of Roses parade, and distributed it to twenty-two cities across the country. For those lucky

enough to see the program, the effect of a moving color image on a small screen seems to have been mesmerizing. The *New York Times*, in typical language, called it a “veritable bevy of hues and depth.” “To concentrate so much color information within the frame of a small screen,” the *Times* wrote, “would be difficult for even the most gifted artist doing a ‘still’ painting. To do it with constantly moving pictures seemed pure wizardry.” Alas, the Rose Parade “broadcast” turned out to be not all that broad, given that it was visible only on prototype televisions in RCA showrooms. Color programming would not become standard on prime-time shows until the late 1960s. After the advent of color, the basic conventions that defined the television image would go unchanged for decades. The delivery mechanisms began to diversify with the introduction of VCRs and cable in the late 1970s. But the image remained the same.

In the mid-1980s, a number of influential media and technology executives, along with a few visionary politicians, had the eminently good idea that it was time to upgrade the video quality of broadcast television. Speeches were delivered, committees formed, experimental prototypes built, but it wasn’t until July 23, 1996, that a Raleigh, North Carolina, CBS affiliate initiated the first public transmission of an HDTV signal. Like the Tournament of Roses footage, though, there were no ordinary consumers with sets capable of displaying its “wizardry.”¹ A handful of broadcasters began transmitting HDTV signals in 1999, but HD television didn’t become a mainstream consumer phenomenon for another five years. Even after the FCC mandated that all television stations cease broadcasting the old analog standard on June 12, 2009, more than 10 percent of U.S. households had televisions that went dark that day.

It is one of the great truisms of our time that we live in an age of technological *acceleration*; the new paradigms keep rolling in, and the intervals between them keep shortening. This acceleration reflects not only the flood of new products, but also our growing willingness to embrace these strange new devices, and put them to use. The waves roll in at ever-

increasing frequencies, and more and more of us are becoming trained surfers, paddling out to meet them the second they start to crest. But the HDTV story suggests that this acceleration is hardly a universal law. If you measure how quickly a new technology progresses from an original idea to mass adoption, then it turns out that HDTV was traveling at the exact same speed that color television had traveled four decades earlier. It took ten years for color TV to go from the fringes to the mainstream; two generations later, it took HDTV just as long to achieve mass success.

In fact, if you look at the entirety of the twentieth century, the most important developments in mass, one-to-many communications clock in at the same social innovation rate with an eerie regularity. Call it the 10/10 rule: a decade to build the new platform, and a decade for it to find a mass audience. The technology standard of amplitude-modulated radio—what we now call AM radio—evolved in the first decade of the twentieth century. The first commercial AM station began broadcasting in 1920, but it wasn't until the late 1920s that radios became a fixture in American households. Sony inaugurated research into the first consumer videocassette recorder in 1969, but didn't ship its first Betamax for another seven years, and VCRs didn't become a household necessity until the mid-eighties. The DVD player didn't statistically replace the VCR in American households until 2006, nine years after the first players went on the market. Cell phones, personal computers, GPS navigation devices—all took a similar time frame to go from innovation to mass adoption.

Consider, as an alternate scenario, the story of Chad Hurley, Steve Chen, and Jawed Karim, three former employees of the online payment site PayPal, who decided in early 2005 that the Web was ripe for an upgrade in the way it handled video and sound. Video, of course, was not native to the Web, which had begun its life fifteen years before as a platform for academics to share hypertext documents. But over the years, video clips had begun to trickle their way online, thanks to new video standards that emerged, such as

Quick-Time, Flash, or Windows Media Player. But the mechanisms that allowed people to upload and share their own videos were too challenging for most ordinary users. So Hurley, Chen, and Karim cobbled together a rough beta for a service that would correct these deficiencies, raised less than \$10 million in venture capital, hired about two dozen people, and launched YouTube, a website that utterly transformed the way video information is shared online. Within sixteen months of the company's founding, the service was streaming more than 30 million videos a day. Within two years, YouTube was one of the top-ten most visited sites on the Web. Before Hurley, Chen, and Karim hit upon their idea for a start-up, video on the Web was as common as subtitles on television. The Web was about doing things with text, and uploading the occasional photo. YouTube brought Web video into the mainstream.

Now compare the way these two ideas—HDTV and YouTube—changed the basic rules of engagement for their respective platforms. Going from analog television to HDTV is a change in degree, not in kind: there are more pixels; the sound is more immersive; the colors are sharper. But consumers watch HDTV the exact same way they watched old-fashioned analog TV. They choose a channel, and sit back and watch. YouTube, on the other hand, radically altered the basic rules of the medium. For starters, it made watching video on the Web a mass phenomenon. But with YouTube you weren't limited to sitting and watching a show, television-style; you could also upload your own clips, recommend or rate other clips, get into a conversation about them. With just a few easy keystrokes, you could take a clip running on someone else's site, and drop a copy of it onto your own site. The technology allowed ordinary enthusiasts to effectively program their own private television networks, stitching together video clips from all across the planet.

Some will say that this is merely a matter of software, which is intrinsically more adaptable than hardware like televisions or cellular phones. But before the Web became mainstream in the mid-1990s, the pace of software innovation


and creativity are fractal: they reappear in recognizable form as you zoom in and out, from molecule to neuron to pixel to sidewalk. Whether you're looking at the original innovations of carbon-based life, or the explosion of new software tools on the Web, the same shapes keep turning up. When life gets creative, it has a tendency to gravitate toward certain recurring patterns, whether those patterns are emergent and self-organizing, or whether they are deliberately crafted by human agents.

It may seem odd to talk about such different regions of experience as though they were interchangeable. But in fact, we are constantly making equivalent conceptual leaps from biology to culture without blinking. It is not a figure of speech to say that the pattern of “competition”—a term often associated with innovation—plays a critical role in the behavior of marketplaces, in the interaction between a swarm of sperm cells and an egg, and in the ecosystem-scale battle between organisms for finite energy sources. We are not using a metaphor of economic competition to describe the struggles of those sperm cells: the meaning of the word “competition” is wide (or perhaps deep) enough to encompass sperm cells *and* corporations. The same principle applies to the seven patterns I have assembled here.

Traveling across these different environments and scales is not merely intellectual tourism. Science long ago realized that we can understand something better by studying its behavior in different contexts. When we want to answer a question like “Why has the Web been so innovative?” we naturally invoke thoughts of its creators, and the workspaces, organizations, and information networks they used in building it. But it turns out that we can answer the question more comprehensively if we draw analogies to patterns of innovation that we see in ecosystems like Darwin's coral reef, or in the structure of the human brain. We have no shortage of theories to instruct us how to make our organizations more creative, or explain why tropical rain forests engineer so much molecular diversity. What we lack is a unified theory that describes the common attributes shared by all those

innovation systems. Why is a coral reef such an engine of biological innovation? Why do cities have such an extensive history of idea creation? Why was Darwin able to hit upon a theory that so many brilliant contemporaries of his missed? No doubt there are partial answers to these questions that are unique to each situation, and each scale: the ecological history of the reef; the sociology of urban life; the intellectual biography of a scientist. But the argument of this book is that there are other, more interesting answers that are applicable to all three situations, and that by approaching the problem in this fractal, cross-disciplinary way, new insights become visible. Watching the ideas spark on these different scales reveals patterns that single-scale observations easily miss or undervalue.

I call that vantage point the *long zoom*. It can be imagined as a kind of hourglass:



As you descend toward the center of the glass, the biological scales contract: from the global, deep time of evolution to the microscopic exchanges of neurons or DNA. At the center of the glass, the perspective shifts from nature to

culture, and the scales widen: from individual thoughts and private workspaces to immense cities and global information networks. When we look at the history of innovation from the vantage point of the long zoom, what we find is that unusually generative environments display similar patterns of creativity at multiple scales simultaneously. You can't explain the biodiversity of the coral reef by simply studying the genetics of the coral itself. The reef generates and sustains so many different forms of life because of patterns that recur on the scales of cells, organisms, and the wider ecosystem itself. The sources of innovation in the city and the Web are equally fractal. In this sense, seeing the problem of innovation from the long-zoom perspective does not just give us new metaphors. It gives us new *facts*.

The pattern of “competition” is an excellent case in point. Every economics textbook will tell you that competition between rival firms leads to innovation in their products and services. But when you look at innovation from the long-zoom perspective, competition turns out to be less central to the history of good ideas than we generally think. Analyzing innovation on the scale of individuals and organizations—as the standard textbooks do—distorts our view. It creates a picture of innovation that overstates the role of proprietary research and “survival of the fittest” competition. The long-zoom approach lets us see that openness and connectivity may, in the end, be more valuable to innovation than purely competitive mechanisms. Those patterns of innovation deserve recognition—in part because it's intrinsically important to understand why good ideas emerge historically, and in part because by embracing these patterns we can build environments that do a better job of nurturing good ideas, whether those environments are schools, governments, software platforms, poetry seminars, or social movements. We can think more creatively if we open our minds to the many connected environments that make creativity possible.


The academic literature on innovation and creativity is rich with subtle distinctions between innovations and inventions, between different modes of creativity: artistic,

scientific, technological. I have deliberately chosen the broadest possible phrasing—good ideas—to suggest the cross-disciplinary vantage point I am trying to occupy. The good ideas in this survey range from software platforms to musical genres to scientific paradigms to new models for government. My premise is that there is as much value to be found in seeking the common properties across all these varied forms of innovation and creativity as there is value to be found in documenting the differences between them. The poet and the engineer (and the coral reef) may seem a million miles apart in their particular forms of expertise, but when they bring good ideas into the world, similar patterns of development and collaboration shape that process.

If there is a single maxim that runs through this book's arguments, it is that we are often better served by *connecting* ideas than we are by protecting them. Like the free market itself, the case for restricting the flow of innovation has long been buttressed by appeals to the “natural” order of things. But the truth is, when one looks at innovation in nature and in culture, environments that build walls around good ideas tend to be less innovative in the long run than more open-ended environments. Good ideas may not want to be free, but they do want to connect, fuse, recombine. They want to reinvent themselves by crossing conceptual borders. They want to complete each other as much as they want to compete.

I.

THE ADJACENT POSSIBLE



Sometime in the late 1870s, a Parisian obstetrician named Stephane Tarnier took a day off from his work at Maternité de Paris, the lying-in hospital for the city's poor women, and paid a visit to the nearby Paris Zoo. Wandering past the elephants and reptiles and classical gardens of the zoo's home inside the Jardin des Plantes, Tarnier stumbled across an exhibit of chicken incubators. Seeing the hatchlings totter about in the incubator's warm enclosure triggered an association in his head, and before long he had hired Odile Martin, the zoo's poultry raiser, to construct a device that would perform a similar function for human newborns. By modern standards, infant mortality was staggeringly high in the late nineteenth century, even in a city as sophisticated as Paris. One in five babies died before learning to crawl, and the odds were far worse for premature babies born with low birth weights. Tarnier knew that temperature regulation was critical for keeping these infants alive, and he knew that the French medical establishment had a deep-seated obsession with statistics. And so as soon as his newborn incubator had been installed at Maternité, the fragile infants warmed by hot water bottles below the wooden boxes, Tarnier embarked on a quick study of five hundred babies. The results shocked the Parisian medical establishment: while 66 percent of low-weight babies died within weeks of birth, only 38 percent died if they were housed in Tarnier's incubating box. You could effectively halve the mortality rate for premature babies simply by treating them like hatchlings in a zoo.

tapped both the local supply of parts themselves and the local knowledge of automobile repair. These were both abundant resources in the developing world context, as Rosen liked to say. You didn't have to be a trained medical technician to fix the NeoNurture; you didn't even have to read the manual. You just needed to know how to replace a broken headlight.

Good ideas are like the NeoNurture device. They are, inevitably, constrained by the parts and skills that surround them. We have a natural tendency to romanticize breakthrough innovations, imagining momentous ideas transcending their surroundings, a gifted mind somehow seeing over the detritus of old ideas and ossified tradition. But ideas are works of bricolage; they're built out of that detritus. We take the ideas we've inherited or that we've stumbled across, and we jigger them together into some new shape. We like to think of our ideas as \$40,000 incubators, shipped direct from the factory, but in reality they've been cobbled together with spare parts that happened to be sitting in the garage.

Before his untimely death in 2002, the evolutionary biologist Stephen Jay Gould maintained an odd collection of footwear that he had purchased during his travels through the developing world, in open-air markets in Quito, Nairobi, and Delhi. They were sandals made from recycled automobile tires. As a fashion statement, they may not have amounted to much, but Gould treasured his tire sandals as a testimony to "human ingenuity." But he also saw them as a metaphor for the patterns of innovation in the biological world. Nature's innovations, too, rely on spare parts. Evolution advances by taking available resources and cobbling them together to create new uses. The evolutionary theorist François Jacob captured this in his concept of evolution as a "tinkerer," not an engineer; our bodies are also works of bricolage, old parts strung together to form something radically new. "The tires-to-sandals principle works at all scales and times," Gould wrote, "permitting odd and unpredictable initiatives at any

moment—to make nature as inventive as the cleverest person who ever pondered the potential of a junkyard in Nairobi.”

You can see this process at work in the primordial innovation of life itself. We do not yet have scientific consensus on the specifics of life’s origins. Some believe life originated in the boiling, metallic vents of undersea volcanoes; others suspect the open oceans; others point to the tidal ponds where Darwin believed life first took hold. Many respected scientists think that life may have arrived from outer space, embedded in a meteor. But we have a much clearer picture of the composition of earth’s atmosphere before life emerged, thanks to a field known as prebiotic chemistry. The lifeless earth was dominated by a handful of basic molecules: ammonia, methane, water, carbon dioxide, a smattering of amino acids, and other simple organic compounds. Each of these molecules was capable of a finite series of transformations and exchanges with other molecules in the primordial soup: methane and oxygen recombining to form formaldehyde and water, for instance.

Think of all those initial molecules, and then imagine all the potential new combinations that they could form spontaneously, simply by colliding with each other (or perhaps prodded along by the extra energy of a propitious lightning strike). If you could play God and trigger all those combinations, you would end up with most of the building blocks of life: the proteins that form the boundaries of cells; sugar molecules crucial to the nucleic acids of our DNA. But you would not be able to trigger chemical reactions that would build a mosquito, or a sunflower, or a human brain. Formaldehyde is a first-order combination: you can create it directly from the molecules in the primordial soup. The atomic elements that make up a sunflower are the very same ones available on earth before the emergence of life, but you can’t spontaneously create a sunflower in that environment, because it relies on a whole series of subsequent innovations that wouldn’t evolve on earth for billions of years: chloroplasts to capture the sun’s energy, vascular tissues to circulate resources through the plant, DNA molecules to pass

on sunflower-building instructions to the next generation.

The scientist Stuart Kauffman has a suggestive name for the set of all those first-order combinations: “the adjacent possible.” The phrase captures both the limits and the creative potential of change and innovation. In the case of prebiotic chemistry, the adjacent possible defines all those molecular reactions that were directly achievable in the primordial soup. Sunflowers and mosquitoes and brains exist outside that circle of possibility. The adjacent possible is a kind of shadow future, hovering on the edges of the present state of things, a map of all the ways in which the present can reinvent itself. Yet is it not an infinite space, or a totally open playing field. The number of potential first-order reactions is vast, but it is a finite number, and it excludes most of the forms that now populate the biosphere. What the adjacent possible tells us is that at any moment the world is capable of extraordinary change, but only *certain* changes can happen.

The strange and beautiful truth about the adjacent possible is that its boundaries grow as you explore those boundaries. Each new combination ushers new combinations into the adjacent possible. Think of it as a house that magically expands with each door you open. You begin in a room with four doors, each leading to a new room that you haven’t visited yet. Those four rooms are the adjacent possible. But once you open one of those doors and stroll into that room, three new doors appear, each leading to a brand-new room that you couldn’t have reached from your original starting point. Keep opening new doors and eventually you’ll have built a palace.

Basic fatty acids will naturally self-organize into spheres lined with a dual layer of molecules, very similar to the membranes that define the boundaries of modern cells. Once the fatty acids combine to form those bounded spheres, a new wing of the adjacent possible opens up, because those molecules implicitly create a fundamental division between the inside and outside of the sphere. This division is the very essence of a cell. Once you have an “inside,” you can put things there: food, organelles, genetic code. Small molecules

can pass through the membrane and then combine with other molecules to form larger entities too big to escape back through the boundaries of the proto-cell. When the first fatty acids spontaneously formed those dual-layered membranes, they opened a door into the adjacent possible that would ultimately lead to nucleotide-based genetic code, and the power plants of the chloroplasts and mitochondria—the primary “inhabitants” of all modern cells.

The same pattern appears again and again throughout the evolution of life. Indeed, one way to think about the path of evolution is as a continual exploration of the adjacent possible. When dinosaurs such as the velociraptor evolved a new bone called the semi-lunate carpal (the name comes from its half-moon shape), it enabled them to swivel their wrists with far more flexibility. In the short term, this gave them more dexterity as predators, but it also opened a door in the adjacent possible that would eventually lead, many millions of years later, to the evolution of wings and flight. When our ancestors evolved opposable thumbs, they opened up a whole new cultural branch of the adjacent possible: the creation and use of finely crafted tools and weapons.

One of the things that I find so inspiring in Kauffman’s notion of the adjacent possible is the continuum it suggests between natural and man-made systems. He introduced the concept in part to illustrate a fascinating secular trend shared by both natural and human history: this relentless pushing back against the barricades of the adjacent possible.

“Something has obviously happened in the past 4.8 billion years,” he writes. “The biosphere has expanded, indeed, more or less persistently exploded, into the ever-expanding adjacent possible.... It is more than slightly interesting that this fact is clearly true, that it is rarely remarked upon, and that we have no particular theory for this expansion.” Four billion years ago, if you were a carbon atom, there were a few hundred molecular configurations you could stumble into. Today that same carbon atom, whose atomic properties haven’t changed one single nanogram, can help build a sperm whale or a giant redwood or an H1N1 virus, along with a near-

infinite list of other carbon-based life forms that were not part of the adjacent possible of prebiotic earth. Add to that an equally formidable list of human concoctions that rely on carbon—every single object on the planet made of plastic, for instance—and you can see how far the kingdom of the adjacent possible has expanded since those fatty acids self-assembled into the first membrane.

The history of life and human culture, then, can be told as the story of a gradual but relentless probing of the adjacent possible, each new innovation opening up new paths to explore. But some systems are more adept than others at exploring those possibility spaces. The mystery of Darwin's paradox that we began with ultimately revolves around the question of why a coral reef ecosystem should be so adventurous in its exploration of the adjacent possible—so many different life forms sharing such a small space—while the surrounding waters of the ocean lack that same marvelous diversity. Similarly, the environments of big cities allow far more commercial exploration of the adjacent possible than towns or villages, allowing tradesmen and entrepreneurs to specialize in fields that would be unsustainable in smaller population centers. The Web has explored the adjacent possible of its medium far faster than any other communications technology in history. In early 1994, the Web was a text-only medium, pages of words connected by hyperlinks. But within a few years, the possibility space began to expand. It became a medium that let you do financial transactions, which turned it into a shopping mall and an auction house and a casino. Shortly afterward, it became a true two-way medium where it was as easy to publish your own writing as it was to read other people's, which engendered forms that the world had never seen before: user-authored encyclopedias, the blogosphere, social network sites. YouTube made the Web one of the most influential video delivery mechanisms on the planet. And now digital maps are unleashing their own cartographic

considered by most technology historians to be the father of modern computing, though he should probably be called the great-grandfather of modern computing, because it took several generations for the world to catch up to his idea. Babbage is actually in the pantheon for two inventions, neither of which he managed to build during his lifetime. The first was his Difference Engine, a fantastically complex fifteen-ton contraption, with over 25,000 mechanical parts, designed to calculate polynomial functions that were essential to creating the trigonometric tables crucial to navigation. Had Babbage actually completed his project, the Difference Engine would have been the world's most advanced mechanical calculator. When the London Science Museum constructed one from Babbage's plans to commemorate the centennial of his death, the machine returned accurate results to thirty-one places in a matter of seconds. Both the speed and precision of the device would have exceeded anything else possible in Babbage's time by several orders of magnitude.

For all its complexity, however, the Difference Engine was well within the adjacent possible of Victorian technology. The second half of the nineteenth century saw a steady stream of improvements to mechanical calculation, many of them building on Babbage's architecture. The Swiss inventor Per Georg Scheutz constructed a working Difference Engine that debuted at the Exposition Universelle of 1855; within two decades the piano-sized Scheutz design had been reduced to the size of a sewing machine. In 1884, an American inventor named William S. Burroughs founded the American Arithmometer Company to sell mass-produced calculators to businesses around the country. (The fortune generated by those machines would help fund his namesake grandson's writing career, not to mention his drug habit, almost a century later.) Babbage's design for the Difference Engine was a work of genius, no doubt, but it did not transcend the adjacent possible of its day.

The same cannot be said of Babbage's other brilliant idea: the Analytical Engine, the great unfulfilled project of

Babbage's career, which he toiled on for the last thirty years of his life. The machine was so complicated that it never got past the blueprint stage, save a small portion that Babbage built shortly before his death in 1871. The Analytical Engine was—on paper, at least—the world's first programmable computer. Being programmable meant that the machine was fundamentally open-ended; it wasn't designed for a specific set of tasks, the way the Difference Engine had been optimized for polynomial equations. The Analytical Engine was, like all modern computers, a shape-shifter, capable of reinventing itself based on the instructions conjured by its programmers. (The brilliant mathematician Ada Lovelace, the only daughter of Lord Byron, wrote several sets of instructions for Babbage's still-vaporware Analytical Engine, earning her the title of the world's first programmer.) Babbage's design for the engine anticipated the basic structure of all contemporary computers: "programs" were to be inputted via punch cards, which had been invented decades before to control textile looms; instructions and data were captured in a "store," the equivalent of what we now call random access memory, or RAM; and calculations were executed via a system that Babbage called "the mill," using industrial-era language to describe what we now call the central processing unit, or CPU.

Babbage had most of this system sketched out by 1837, but the first true computer to use this programmable architecture didn't appear for more than a hundred years. While the Difference Engine engendered an immediate series of refinements and practical applications, the Analytical Engine effectively disappeared from the map. Many of the pioneering insights that Babbage had hit upon in the 1830s had to be independently rediscovered by the visionaries of World War II-era computer science.

Why did the Analytical Engine prove to be such a short-term dead end, given the brilliance of Babbage's ideas? The fancy way to say it is that his ideas had escaped the bounds of the adjacent possible. But it is perhaps better put in more prosaic terms: Babbage simply didn't have the right spare

parts. Even if Babbage had built a machine to his specs, it is unclear whether it would have worked, because Babbage was effectively sketching out a machine for the electronic age during the middle of the steam-powered mechanical revolution. Unlike all modern computers, Babbage's machine was to be composed entirely of mechanical gears and switches, staggering in their number and in the intricacy of their design. Information flowed through the system as a constant ballet of metal objects shifting positions in carefully choreographed movements. It was a maintenance nightmare, but more than that, it was bound to be hopelessly slow. Babbage bragged to Ada Lovelace that he believed the machine would be able to multiply two twenty-digit numbers in three minutes. Even if he was right—Babbage wouldn't have been the first tech entrepreneur to exaggerate his product's performance—that kind of processing time would have made executing more complicated programs torturously slow. The first computers of the digital age could perform the same calculation in a matter of seconds. An iPhone completes millions of such calculations in the same amount of time. Programmable computers needed vacuum tubes, or, even better, integrated circuits, where information flows as tiny pulses of electrical activity, instead of clanking, rusting, steam-powered metal gears.

You can see a comparable pattern—on a vastly accelerated timetable—in the story of YouTube. Had Hurley, Chen, and Karim tried to execute the exact same idea for YouTube ten years earlier, in 1995, it would have been a spectacular flop, because a site for sharing video was not within the adjacent possible of the early Web. For starters, the vast majority of Web users were on painfully slow dial-up connections that could sometimes take minutes to download a small image. (The average two-minute-long YouTube clip would have taken as much as an hour to download on the then-standard 14.4 bps modems.) Another key to YouTube's early success is that its developers were able to base the video serving on Adobe's Flash platform, which meant that they could focus on the ease of sharing and discussing clips, and not spend

millions of dollars developing a whole new video standard from scratch. But Flash itself wasn't released until late 1996, and didn't even support video until 2002.

To use our microbiology analogy, having the idea for a Difference Engine in the 1830s was like a bunch of fatty acids trying to form a cell membrane. Babbage's calculating machine was a leap forward, to be sure, but as advanced as it was, the Difference Engine was still within the bounds of the adjacent possible, which is precisely why so many practical iterations of Babbage's design emerged in the subsequent decades. But trying to create an Analytical Engine in 1850—or YouTube in 1995—was the equivalent of those fatty acids trying to self-organize into a sea urchin. The idea was right, but the environment wasn't ready for it yet.

All of us live inside our own private versions of the adjacent possible. In our work lives, in our creative pursuits, in the organizations that employ us, in the communities we inhabit—in all these different environments, we are surrounded by potential new configurations, new ways of breaking out of our standard routines. We are, each of us, surrounded by the conceptual equivalent of those Toyota spare parts, all waiting to be recombined into something magical, something new. It need not be the epic advances of biological diversity, or the invention of programmable computing. Unlocking a new door can lead to a world-changing scientific breakthrough, but it can also lead to a more effective strategy for teaching second-graders, or a novel marketing idea for the vacuum cleaner your company's about to release. The trick is to figure out ways to explore the edges of possibility that surround you. This can be as simple as changing the physical environment you work in, or cultivating a specific kind of social network, or maintaining certain habits in the way you seek out and store information.

Recall the question we began with: What kind of environment creates good ideas? The simplest way to answer it is this: innovative environments are better at helping their

inhabitants explore the adjacent possible, because they expose a wide and diverse sample of spare parts—mechanical or conceptual—and they encourage novel ways of recombining those parts. Environments that block or limit those new combinations—by punishing experimentation, by obscuring certain branches of possibility, by making the current state so satisfying that no one bothers to explore the edges—will, on average, generate and circulate fewer innovations than environments that encourage exploration. The infinite variety of life that so impressed Darwin, standing in the calm waters of the Keeling Islands, exists because the coral reef is supremely gifted at recycling and reinventing the spare parts of its ecosystem.

There's a famous moment in the story of the near-catastrophic *Apollo 13* mission—wonderfully captured in the Ron Howard film—where the mission control engineers realize they need to create an improvised carbon dioxide filter, or the astronauts will poison the lunar module atmosphere with their own exhalations before they return to Earth. The astronauts have plenty of carbon “scrubbers” on board, but these filters were designed for the original, damaged spacecraft, and don't fit the air ventilation system of the lunar module they are using as a lifeboat to return home. Mission Control quickly assembles what it calls a “tiger team” of engineers to hack their way through the problem, and creates a rapid-fire inventory of all the available equipment currently on the lunar module. In the movie, Deke Slayton, head of Flight Crew Operations, tosses a jumbled pile of gear on a conference table: suit hoses, canisters, stowage bags, duct tape, and other assorted gadgets. He holds up the carbon scrubbers. “We gotta find a way to make this fit into a hole for this,” he says, and then points to the spare parts on the table, “using nothing but *that*.”

The space gear on the table defines the adjacent possible for the problem of building a working carbon scrubber on a lunar module. The device they eventually concoct, dubbed the “mailbox,” performs beautifully. The canisters and nozzles are like the ammonia and methane molecules of the

achievements of the human brain) if they weren't capable of making such elaborate connections with each other. The average neuron connects to a thousand other neurons scattered across the brain, which means that the adult human brain contains 100 trillion distinct neuronal connections, making it the largest and most complex network on earth. (By comparison, there are somewhere on the order of 40 billion pages on the Web. If you assume an average of ten links per page, that means you and I are walking around with a high-density network in our skulls that is orders of magnitude larger than the entirety of the World Wide Web.)

The second precondition is that the network be *plastic*, capable of adopting new configurations. A dense network incapable of forming new patterns is, by definition, incapable of change, incapable of probing at the edges of the adjacent possible. When a new idea pops into your head, the sense of novelty that makes the experience so magical has a direct correlate in the cells of your brain: a brand-new assemblage of neurons has come together to make the thought possible. Those connections are built by our genes and by personal experience: some connections help regulate our heartbeat and trigger reflex reactions; others conjure up vivid sense memories of the cookies we ate as children; others help us invent the concept of a programmable computer. The connections are the key to wisdom, which is why the whole notion of losing neurons after we hit adulthood is a red herring. What matters in your mind is not just the number of neurons, but the myriad connections that have formed between them.

Of course, everything that happens in your brain is, technically speaking, a network. Remembering to cut your toenails involves a network of neurons firing in some kind of orderly fashion. But that doesn't make it an epiphany. It turns out that good ideas have certain signature patterns in the networks that make them. The creating brain behaves differently from the brain that is performing a repetitive task. The neurons communicate in different ways. The networks take on distinct shapes.

The question is how to push your brain toward those more creative networks. The answer, as it happens, is delightfully fractal: to make your mind more innovative, you have to place it inside environments that share that same network signature: networks of ideas or people that mimic the neural networks of a mind exploring the boundaries of the adjacent possible. Certain environments enhance the brain's natural capacity to make new links of association. But these patterns of connection are much older than the human brain, older than neurons even. They take us back, once again, to the origin of life itself.

As far as we know, “carbon-based life” is a redundant expression: life would be impossible without the carbon atom. Most astrobiologists—scientists who study the possibility of life elsewhere in the universe—believe that if we are ever to discover convincing evidence of extraterrestrial life, be it on Mars or in some distant galaxy, it, too, will turn out to be carbon-based.

Why are we so confident about carbon's essential role in creating living things? The answer has to do with the core properties of the carbon atom itself. Carbon has four valence electrons residing in the outermost shell of the atom, which, for complicated reasons, makes it uniquely talented at forming connections with other atoms, particularly with hydrogen, nitrogen, oxygen, phosphorus, sulfur—and, crucially, with other carbon atoms. These six atoms make up 99 percent of the dry weight of all living organisms on earth. Those four valence bonds give carbon a strong propensity for forming elaborate chains and rings of polymers: everything from the genetic information stored in nucleic acids, to the building blocks of proteins, to the energy storage of carbohydrates and fats. (Modern technology has exploited the generative potential of the carbon atom via the artificial polymers we call plastics.) Carbon atoms measure only 0.03 percent of the overall composition of the earth's crust, and yet they make up nearly 20 percent of our body mass. That

abundance highlights the unique property of the carbon atom: its combinatorial power. Carbon is a connector.

Those connections are essential for the day-to-day functioning of life: chains of nucleic acids instructing amino acids to assemble into long strings of protein, powered by the stored energy of carbohydrates. But the connective properties of carbon were essential to the original innovations of life itself. Without carbon's innate talent for forming new complex molecules with other atoms, it is hard to imagine how the first organisms would have evolved. Those four valence electrons allowed the prebiotic earth to explore its own adjacent possible, sifting through the long list of potential molecular combinations until it hit upon a series of stable chemical reactions that blossomed into the first organisms. Without the generative links of carbon, the earth would have likely remained a lifeless soup of elements, a planet of dead chemistry.

Carbon's connective talents lie at the center of one of the most famous scientific experiments of the twentieth century. In 1953, two University of Chicago professors, Stanley L. Miller and Harold C. Urey, created a closed system of glass tubes and flasks that simulated the early conditions of the prebiotic earth. The main ingredients were methane (CH_4), ammonia (NH_3), hydrogen (H_2), and water (H_2O). Only the methane contained carbon atoms. One flask connected to the chemical soup contained a pair of electrodes, which Miller and Urey used to simulate lightning by triggering a series of quick sparks between them. They ran the experiment for seven straight days, and by the time they had completed the first cycle, they found that more than 10 percent of the carbon had spontaneously recombined into many of the organic compounds essential to life: sugars, lipids, nucleic acids. Miller claimed at the time that "just turning on the spark in a basic pre-biotic experiment" produced half of the twenty-two amino acids. Several years ago, a team reanalyzed the original flasks from the Miller-Urey experiments, and found that in one version—which simulated the environment around an undersea volcano—all twenty-two amino acids had