CHAPTER 5

INNOVATION: DECLINING OR RECOMBINING?

"If you want to have good ideas you must have many ideas."

—Linus Pauling

EVERYONE AGREES THAT IT would be troubling news if America's rate of innovation were to decrease. But we can't seem to agree at all about whether this is actually happening.

We care about innovation so much not simply because we like new stuff, although we certainly do. As the novelist William Makepeace Thackeray observed, "Novelty has charms that our mind can hardly withstand."[1] Some of us can hardly withstand the allure of new gadgets; others are charmed by the latest fashion styles or places to see and be seen. From an economist's perspective, satisfying these desires is great—taking care of consumer demand is usually seen as a good thing. But innovation is also the most important force that makes our society wealthier.

Why Innovation is (Almost) Everything

Paul Krugman speaks for many, if not most, economists when he says, "Productivity isn't everything, but in the long run it is almost everything." Why? Because, he explains, "A country's ability to improve its standard of living over time depends almost entirely on its ability to raise its output per worker"—in other words, the number of hours of labor it takes to produce everything, from automobiles to zippers, that we produce.[2] Most countries don't have extensive mineral wealth or oil reserves, and thus can't get rich by exporting them.* So the only viable way for societies to become wealthier—to improve the standard of living available to its people—is for their companies and workers to keep getting more output from the same number of inputs, in other words more goods and services from the same number of people.

Innovation is how this productivity growth happens. Economists love to argue with one another, but there's great consensus among them about the fundamental importance of innovation for growth and prosperity. Most in the profession would agree with Joseph Schumpeter, the topic's great scholar, who wrote that, "Innovation is the outstanding fact in the economic history of capitalist society . . . and also it is largely responsible for most of what we would at first sight attribute to other factors."[3] It is here that the consensus ends. How much of this "outstanding fact" is taking place right now, and whether it's on an upward or downward trend, is a matter of great dispute.

Why We Should Be Worried: Innovations Get Used Up

Economist Bob Gordon, one of the most thoughtful, thorough, and widely respected researchers of productivity and economic growth, recently completed a major study of how the American standard of living has changed over the past 150 years. His work left him convinced that innovation is slowing down.

Gordon emphasizes—as do we—the role of new technologies in driving economic growth. And like us, he's impressed by the productive power unleashed by the steam engine and the other technologies of the Industrial Revolution. According to Gordon, it was the first truly significant event in the economic history of the world. As he writes, "there was almost no economic growth for four centuries and probably for the previous millennium" prior to 1750, or roughly when the Industrial Revolution started.[4] As we saw in the first chapter, human population growth and social development were very nearly flat until the steam engine came along. Unsurprisingly, it turns out that economic growth was, too.

As Gordon shows, however, once this growth got started it stayed on a sharp upward trajectory for two hundred years. This was due not only to the original Industrial Revolution, but also to a second one, it too reliant on technological innovation. Three novelties were central here: electricity, the internal combustion engine, and indoor plumbing with running water, all of which came onto the scene between 1870 and 1900.

The 'great inventions' of this second industrial revolution, in Gordon's estimation, "were so important and far-reaching that they took a full 100 years to have their main effect." But once that effect had been realized, a new problem emerged. Growth stalled out, and even began to decline. At the risk of being flippant, when the steam engine ran out of steam, the internal combustion engine was there to replace it. But once the internal combustion engine ran out of fuel, we weren't left with much. To use Gordon's words,

> The growth of productivity (output per hour) slowed markedly after 1970. While puzzling at the time, it seems increasingly clear that the one-time-only benefits of the Great

> Inventions and their spin-offs had occurred and could not happen again. . . . All that remained after 1970 were second-round improvements, such as developing short-haul regional jets, extending the original interstate highway network with suburban ring roads, and converting residential America from window unit air conditioners to central air conditioning.5

Gordon is far from alone in this view. In his 2011 book *The Great Stagnation*, economist Tyler Cowen is definitive about the source of America's economic woes:

> We are failing to understand why we are failing. All of these problems have a single, little noticed root cause: We have been living off low-hanging fruit for at least three hundred years. . . . Yet during the last forty years, that low-hanging fruit started disappearing, and we started pretending it was still there. We have failed to recognize that we are at a technological plateau and the trees are more bare than we would like to think.6

General Purpose Technologies: The Ones That Really Matter

Clearly, Gordon and Cowen see the invention of powerful technologies as central to economic progress. Indeed, there's broad agreement among economic historians that some technologies are significant enough to accelerate the normal march of economic progress. To do this, they have to spread throughout many, if not most, industries; they can't remain in just one. The cotton gin, for example, was unquestionably important within the textile sector at the start of the nineteenth century, but pretty insignificant outside of it.*

The steam engine and electrical power, by contrast, quickly spread just about everywhere. The steam engine didn't just massively increase the amount of power available to factories and free them from the need to be located near a stream or river to power the water wheel; it also revolutionized land travel by enabling railroads and sea travel via the steamship. Electricity gave a further boost to manufacturing by enabling individually powered machines. It also lit factories, office buildings, and warehouses and led to further innovations like air conditioning, which made previously sweltering workplaces pleasant.

With their typical verbal flair, economists call innovations like steam power and electricity *general purpose technologies* (GPTs). Economic historian Gavin Wright offers a concise definition: “deep new ideas or techniques that have the potential for important impacts on many sectors of the economy.”[7] “Impacts” here mean significant boosts to output due to large productivity gains. GPTs are important because they are economically significant—they interrupt and accelerate the normal march of economic progress.

In addition to agreeing on their importance, scholars have also come to a consensus on how to recognize GPTs: they should be pervasive, improving over time, and able to spawn new innovations.[8] The preceding chapters have built a case that digital technologies meet all three of these requirements. They improve along a Moore’s Law trajectory, are used in every industry in the world, and lead to innovations like autonomous cars and nonhuman *Jeopardy!* champions. Are we alone in thinking that information and communication technology (ICT) belongs in the same category as steam and electricity? Are we the only ones who think, in short, that ICT is a GPT?

Absolutely not. Most economic historians concur with the assessment that ICT meets all of the criteria given above, and so should join the club of general purpose technologies. In fact, in a list of all the candidates for this classification compiled by the economist Alexander Field, only steam power got more votes than ICT, which was tied with electricity as the second most commonly accepted GPT.[9]

If we are all in agreement, then why the debate over whether ICTs are ushering in a new golden age of innovation and growth? Because, the argument goes, their economic benefits have already been captured and now most new ‘innovation’ involves entertaining ourselves inexpensively online. According to Robert Gordon:

> The first industrial robot was introduced by General Motors in 1961. Telephone operators went away in the 1960s. . . . Airline reservations systems came in the 1970s, and by 1980 bar-code scanners and cash machines were spreading through the retail and banking industries. . . . The first personal computers arrived in the early 1980s with their word processing, word wrap, and spreadsheets. . . . More recent and thus more familiar was the rapid development of the web and e-commerce after 1995, a process largely completed by 2005.10

At present, says Cowen, “The gains of the Internet are very real and I am here to praise them, not damn them. . . . Still, the overall picture is this: We are having more fun, in part because of the Internet. We are also having more cheap fun. [But] we are coming up short on the revenue side, so it is harder to pay our debts, whether individuals, businesses, or governments."[11] Twenty-first century ICT, in short, is failing the prime test of being economically significant.

Why We Shouldn't Be Worried: Innovations Don't Get Used Up

For any good scientist, of course, data are the ultimate decider of hypotheses. So what do the data say here? Do the productivity numbers back up this pessimistic view of the power of digitization? We'll get to the data in chapter 7. First, though, we want to present a very different view of how innovation works—an alternative to the notion that innovations get 'used up.'

Gordon writes that "it is useful to think of the innovative process as a series of discrete inventions followed by incremental improvements which ultimately tap the full potential of the initial invention."[12] This seems sensible enough. An invention like the steam engine or computer comes along and we reap economic benefits from it. Those benefits start small while the technology is immature and not widely used, grow to be quite big as the GPT improves and propagates, then taper off as the improvement—and especially the propagation—die down. When multiple GPTs appear at the same time, or in a steady sequence, we sustain high rates of growth over a long period. But if there's a big gap between major innovations, economic growth will eventually peter out. We'll call this the 'innovation-as-fruit' view of things, in honor of Tyler Cowen's imagery of all the low-hanging fruit being picked. In this perspective, coming up with an innovation is like growing fruit, and exploiting an innovation is like eating the fruit over time.

Another school of thought, though, holds that the true work of innovation is not coming up with something big and new, but instead recombining things that already exist. And the more closely we look at how major steps forward in our knowledge and ability to accomplish things have actually occurred, the more this recombinant view makes sense. For example, it's exactly how at least one Nobel Prize–winning innovation came about.

Kary Mullis won the 1993 Nobel Prize in Chemistry for the development of the polymerase chain reaction (PCR), a now ubiquitous technique for replicating DNA sequences. When the idea first came to him on a nighttime drive in California, though, he almost dismissed it out of hand. As he recounted in his Nobel Award speech, "Somehow, I thought, it had to be an illusion. . . . It was too easy. . . . There was not a single unknown in the scheme. Every step involved had been done already."[13] "All" Mullis did was recombine well-understood techniques in biochemistry to generate a new one. And yet it's obvious Mullis's recombination is an enormously valuable one.

After examining many examples of invention, innovation, and technological progress, complexity scholar Brian Arthur became convinced that stories like the invention of PCR are the rule, not the exception. As he summarizes in his book *The Nature of Technology*, "To invent something is to find it in what previously exists."[14] Economist Paul Romer has argued forcefully in favor of this view, the so-called 'new growth theory' within economics, in order to distinguish it from perspectives like Gordon's. Romer's inherently optimistic theory stresses the importance of recombinant innovation. As he writes:

> Economic growth occurs whenever people take resources and rearrange them in ways that make them more valuable. . . . Every generation has perceived the limits to growth that finite resources and undesirable side effects would pose if no new . . . ideas were discovered. And every generation has underestimated the potential for finding new . . . ideas. We consistently fail to grasp how many ideas remain to be discovered. . . . Possibilities do not merely add up; they multiply.15

Romer also makes a vital point about a particularly important category of idea, which he calls "meta-ideas":

> Perhaps the most important ideas of all are meta-ideas—ideas about how to support the production and transmission of other ideas. . . . There are . . . two safe predictions. First, the country that takes the lead in the twenty-first century will be the one that implements an innovation that more effectively supports the production of new ideas in the private sector. Second, new meta-ideas of this kind will be found.16

Digital Technologies: The Most General Purpose of All

Gordon and Cowen are world-class economists, but they're not giving digital technologies their due. The next great meta-idea, invoked by Romer, has already been found: it can be seen in the new communities of minds and machines made possible by networked digital devices running an astonishing variety of software. The GPT of ICT has given birth to radically new ways to combine and recombine ideas. Like language, printing, the library, or universal education, the global digital network fosters recombinant innovation. We can mix and remix ideas, both old and recent, in ways we never could before. Let’s look at a few examples.

Google’s Chauffeur project gives new life to an earlier GPT: the internal combustion engine. When an everyday car is equipped with a fast computer and a bunch of sensors (all of which get cheaper according to Moore’s Law) and a huge amount of map and street information (available thanks to the digitization of everything) it becomes an autopiloted vehicle straight out of science fiction. While we humans are still the ones doing the driving, innovations like Waze will help us get around more quickly and ease traffic jams. Waze is a recombination of a location sensor, data transmission device (that is, a phone), GPS system, and social network. The team at Waze invented none of these technologies; they just put them together in a new way. Moore’s Law made all involved devices cheap enough, and digitization made all necessary data available to facilitate the Waze system.

The Web itself is a pretty straightforward combination of the Internet’s much older TCP/IP data transmission network; a markup language called HTML that specified how text, pictures, and so on should be laid out; and a simple PC application called a ‘browser’ to display the results. None of these elements was particularly novel. Their combination was revolutionary.

Facebook has built on the Web infrastructure by allowing people to digitize their social network and put media online without having to learn HTML. Whether or not this was an intellectually profound combination of technological capabilities, it was a popular and economically significantly one—by July 2013, the company was valued at over $60 billion.[17] When photo sharing became one of the most popular activities on Facebook, Kevin Systrom and Mike Krieger decided to build a smartphone application that mimicked this capability, combining it with the option to modify a photo’s appearance with digital filters. This seems like a minor innovation, especially since Facebook already had enabled mobile photo sharing in 2010 when Systrom and Krieger started their project. However, the application they built, called Instagram, attracted more than 30 million users by the spring of

2012, users who had collectively uploaded more than 100 million photographs. Facebook acquired Instagram for approximately $1 billion in April of 2012.

This progression drives home the point that digital innovation is recombinant innovation in its purest form. Each development becomes a building block for future innovations. Progress doesn't run out; it accumulates. And the digital world doesn't respect any boundaries. It extends into the physical one, leading to cars and planes that drive themselves, printers that make parts, and so on. Moore's Law makes computing devices and sensors exponentially cheaper over time, enabling them to be built economically into more and more gear, from doorknobs to greeting cards. Digitization makes available massive bodies of data relevant to almost any situation, and this information can be infinitely reproduced and reused because it is non-rival. As a result of these two forces, the number of potentially valuable building blocks is exploding around the world, and the possibilities are multiplying as never before. We'll call this the 'innovation-as-building-block' view of the world; it's the one held by Arthur, Romer, and the two of us. From this perspective, unlike in the innovation-as-fruit view, building blocks don't ever get eaten or otherwise used up. In fact, they increase the opportunities for future recombinations.

Limits to Recombinant Growth

If this recombinant view of innovation is correct, then a problem looms: as the number of building blocks explodes, the main difficulty is knowing which combinations of them will be valuable. In his paper "Recombinant Growth," the economist Martin Weitzman developed a mathematical model of new growth theory in which the 'fixed factors' in an economy—machine tools, trucks, laboratories, and so on—are augmented over time by pieces of knowledge that he calls 'seed ideas,' and knowledge itself increases over time as previous seed ideas are recombined into new ones.[18] This is an innovation-as-building-block view of the world, where both the knowledge pieces and the seed ideas can be combined and recombined over time.

This model has a fascinating result: because combinatorial possibilities explode so quickly there is soon a virtually infinite number of potentially valuable recombinations of the existing knowledge pieces.* The constraint on the economy's growth then becomes its ability to go through all these potential recombinations to find the truly valuable ones.

As Weitzman writes,

> In such a world the core of economic life could appear increasingly to be centered on the more and more intensive processing of ever-greater numbers of new seed ideas into workable innovations. . . . In the early stages of development, growth is constrained by number of potential new ideas, but later on it is constrained only by the ability to process them.19

Gordon asks the provocative question, "Is growth over?" We'll respond on behalf of Weitzman, Romer, and the other new growth theorists with "Not a chance. It's just being held back by our inability to process all the new ideas fast enough."

What This Problem Needs Are More Eyeballs and Bigger Computers

If this response is at least somewhat accurate—if it captures something about how innovation and economic growth work in the real world—then the best way to accelerate progress is to increase our capacity to test out new combinations of ideas. One excellent way to do this is to involve more people in this testing process, and digital technologies are making it possible for ever more people to participate. We're interlinked by global ICT, and we have affordable access to masses of data and vast computing power. Today's digital environment, in short, is a playground for large-scale recombination. The open source software advocate Eric Raymond has an optimistic observation: "Given enough eyeballs, all bugs are shallow."[20] The innovation equivalent to this might be, "With more eyeballs, more powerful combinations will be found."

NASA experienced this effect as it was trying to improve its ability to forecast solar flares, or eruptions on the sun's surface. Accuracy and plenty of advance warning are both important here, since solar particle events (or SPEs, as flares are properly known) can bring harmful levels of radiation to unshielded gear and people in space. Despite thirty-five years of research and data on SPEs, however, NASA acknowledged that it had "no method available to predict the onset, intensity or duration of a solar particle event."[21]

The agency eventually posted its data and a description of the challenge of predicting SPEs on Innocentive, an online clearinghouse for scientific problems. Innocentive is 'non-credentialist'; people don't have to be PhDs or work in labs in order to browse the problems, download data, or upload a solution. Anyone can work on problems from any discipline; physicists, for example, are not excluded from digging in on biology problems.

As it turned out, the person with the insight and expertise needed to improve SPE prediction was not part of any recognizable astrophysics community. He was Bruce Cragin, a retired radio frequency engineer living in a small town in New Hampshire. Cragin said that, "Though I hadn't worked in the area of solar physics as such, I had thought a lot about the theory of magnetic reconnection."[22] This was evidently the right theory for the job, because Cragin's approach enabled prediction of SPEs eight hours in advance with 85 percent accuracy, and twenty-four hours in advance with 75 percent accuracy. His recombination of theory and data earned him a thirty-thousand-dollar reward from the space agency.

In recent years, many organizations have adopted NASA's strategy of using technology to open up their innovation challenges and opportunities to more eyeballs. This phenomenon goes by several names, including 'open innovation' and 'crowdsourcing,' and it can be remarkably effective. The innovation scholars Lars Bo Jeppesen and Karim Lakhani studied 166 scientific problems posted to Innocentive, all of which had stumped their home organizations. They found that the crowd assembled around Innocentive was able to solve forty-nine of them, for a success rate of nearly 30 percent. They also found that people whose expertise was far away from the apparent domain of the problem were more likely to submit winning solutions. In other words, it seemed to actually help a solver to be 'marginal'—to have education, training, and experience that were not obviously relevant for the problem. Jeppesen and Lakhani provide vivid examples of this:

> [There were] different winning solutions to the same scientific challenge of identifying a food-grade polymer delivery system by an aerospace physicist, a small agribusiness owner, a transdermal drug delivery specialist, and an industrial scientist. . . . All four submissions successfully achieved the required challenge objectives with differing scientific mechanisms. . . .

> [Another case involved] an R&D lab that, even after consulting with internal and external specialists, did not understand the toxicological significance of a particular pathology that had been observed in an ongoing research program. . . . It was eventually solved, using methods common in her field, by a scientist with a Ph.D. in protein crystallography who would not normally be exposed to toxicology problems or solve such problems on a routine basis.[23]

Like Innocentive, the online startup Kaggle also assembles a diverse, non-credentialist group of people from around the world to work on tough problems submitted by organizations. Instead of scientific challenges, Kaggle specializes in data-intensive ones where the goal is to arrive at a better prediction than the submitting organization's starting baseline prediction. Here again, the results are striking in a couple of ways. For one thing, improvements over the baseline are usually substantial. In one case, Allstate submitted a dataset of vehicle characteristics and asked the Kaggle community to predict which of them would have later personal liability claims filed against them.[24] The contest lasted approximately three months and drew in more than one hundred contestants. The winning prediction was more than 270 percent better than the insurance company's baseline.

Another interesting fact is that the majority of Kaggle contests are won by people who are marginal to the domain of the challenge—who, for example, made the best prediction about hospital readmission rates despite having no experience in health care—and so would not have been consulted as part of any traditional search for solutions. In many cases, these demonstrably capable and successful data scientists acquired their expertise in new and decidedly digital ways.

Between February and September of 2012 Kaggle hosted two competitions about computer grading of student essays, which were sponsored by the Hewlett Foundation.* Kaggle and Hewlett worked with multiple education experts to set up the competitions, and as they were preparing to launch many of these people were worried. The first contest was to consist of two rounds. Eleven established educational testing companies would compete against one another in the first round, with members of Kaggle's community of data scientists invited to join in, individually or in teams, in the second. The experts were worried that the Kaggle crowd would simply not be competitive in the second round. After all, each of the testing companies had been working on automatic grading for some time and had devoted substantial resources to the problem. Their hundreds of person-years of accumulated experience and expertise seemed like an insurmountable advantage over a bunch of novices.

They needn’t have worried. Many of the ‘novices’ drawn to the challenge outperformed all of the testing companies in the essay competition. The surprises continued when Kaggle investigated who the top performers were. In both competitions, none of the top three finishers had any previous significant experience with either essay grading or natural language processing. And in the second competition, none of the top three finishers had any formal training in artificial intelligence beyond a free online course offered by Stanford AI faculty and open to anyone in the world who wanted to take it. People all over the world did, and evidently they learned a lot. The top three individual finishers were from, respectively, the United States, Slovenia, and Singapore.

Quirky, another Web-based startup, enlists people to participate in both phases of Weitzman’s recombinant innovation—first generating new ideas, then filtering them. It does this by harnessing the power of many eyeballs not only to come up with innovations but also to filter them and get them ready for market. Quirky seeks ideas for new consumer products from its crowd, and also relies on the crowd to vote on submissions, conduct research, suggest improvements, figure out how to name and brand the products, and drive sales. Quirky itself makes the final decisions about which products to launch and handles engineering, manufacturing, and distribution. It keeps 70 percent of all revenue made through its website and distributes the remaining 30 percent to all crowd members involved in the development effort; of this 30 percent, the person submitting the original idea gets 42 percent, those who help with pricing share 10 percent, those who contribute to naming share 5 percent, and so on. By the fall of 2012, Quirky had raised over $90 million in venture capital financing and had agreements to sell its products at several major retailers, including Target and Bed Bath & Beyond. One of its most successful products, a flexible electrical power strip called Pivot Power, sold more than 373 thousand units in less than two years and earned the crowd responsible for its development over $400,000.

Affinnova, yet another young company supporting recombinant innovation, helps its customers with the second of Weitzman’s two phases: sorting through the possible combinations of building blocks to find the most valuable ones. It does this by combining crowdsourcing with Nobel Prize–worthy algorithms. When Carlsberg breweries wanted to update the bottle and label for Belgium's Grimbergen, the world's oldest continually produced abbey beer, it knew it had to proceed carefully. The company wanted to update the brand without sacrificing its strong reputation or downplaying its nine hundred years of history. It knew that the redesign would mean generating many candidates for each of several attributes—bottle shape, embossments, label color, label placement, cap design, and so on—then settling on the right combination of all of these. The 'right' combination from among the thousands of possibilities, however, was not obvious at the outset.

The standard approach to this kind of problem is for the design team to generate a few combinations that they think are good, then use focus groups or other small-scale methods to finalize which is best. Affinnova offers a very different approach. It makes use of the mathematics of choice modeling, an advance significant enough to have earned a Nobel Prize for its intellectual godfather, economist Daniel McFadden. Choice modeling quickly identifies people's preferences—do they prefer a brown embossed bottle with a small label, or a green non-embossed one with a large label?—by repeatedly presenting them with a small set of options and asking them to select which they like best. Affinnova presents these options via the Web and can find the mathematically optimal set of options (or at least come close to it) after involving only a few hundred people in the evaluation process. For Grimbergen, the design that resulted from this explicitly recombinant process had an approval rating 3.5 times greater than that of the previous bottle.[25]

When we adopt the perspective of the new growth theorists and match it against what we see with Waze, Innocentive, Kaggle, Quirky, Affinnova, and many others, we become optimistic about the current and future of innovation. And these digital developments are not confined to the high-tech sector—they're not just making computers and networks better and faster. They're helping us drive our cars better (and may soon make it unnecessary for us to drive at all), allowing us to arrive at better predictions of solar flares, solving problems in food science and toxicology, and giving us better power strips and beer bottles. These and countless other innovations will add up over time, and they'll keep coming and keep adding up. Unlike some of our colleagues, we are confident that innovation and productivity will continue to grow at healthy rates in the future. Plenty of building blocks are in place, and they're being recombined in better and better ways all the time.

---

* In reality, many of the countries that *do* have large amounts of mineral and commodity wealth are often crippled by the twin terrors of the "resource curse": low growth rates and lots of poverty.

* Some have tied the invention of the cotton gin to increased demand for slave labor in the American South and therefore to the Civil War, but its direct economic effect outside the textile industry was minimal.

* Keep in mind that if there are only fifty-two seed ideas in such an economy, they have many more potential combinations than there are atoms in our solar system.

* Improvements in this area are important because essays are better at capturing student leaning than multiple-choice questions, but much more expensive to grade when human raters are used. Automatic grading of essays would both improve the quality of education and lower its cost.