


THE LONG TAIL

Why the Future of Business
Is Selling Less of More



CHRIS ANDERSON

 HYPERION

NEW YORK

2006

INTRODUCTION

The tracking of top-seller lists is a national obsession. Our culture is a massive popularity contest. We are consumed by hits—making them, choosing them, talking about them, and following their rise and fall. Every weekend is a box-office horse race, and every Thursday night is a Darwinian struggle to find the fittest TV show and let it live to see another week. A few hit songs play in heavy rotation on the radio dials, while entertainment executives in all these industries sweat as they search for the next big thing.

This is the world the blockbuster built. The massive media and entertainment industries grew up over the past half century on the back of box-office rockets, gold records, and double-digit TV ratings. No surprise that hits have become the lens through which we observe our own culture. We define our age by our celebrities and mass-market products—they are the connective tissue of our common experience. The star-making system that Hollywood began eight decades ago has now spun out into every corner of commerce, from shoes to chefs. Our media is obsessed with what's hot and what's not. Hits, in short, *rule*.

Yet look a little closer and you'll see that this picture, which first emerged with the postwar broadcast era of radio and television, is now

starting to tatter at the edges. Hits are starting to, gasp, *rule less*. Number one is still number one, but the sales that go with that are not what they once were.

Most of the top fifty best-selling albums of all time were recorded in the seventies and eighties (the Eagles, Michael Jackson), and none of them were made in the past five years. Hollywood box-office revenue was down by more than 6 percent in 2005, reflecting the reality that the theatergoing audience is falling even as the population grows.

Every year network TV loses more of its audience to hundreds of niche cable channels. Males age eighteen to thirty-four, the most desirable audience for advertisers, are starting to turn off the TV altogether, shifting more and more of their screen time to the Internet and video games. The ratings of top TV shows have been falling for decades, and the number one show today wouldn't have made the top ten in 1970.

In short, although we still obsess over hits, they are not quite the economic force they once were. Where are those fickle consumers going instead? No single place. They are scattered to the winds as markets fragment into countless niches. The one big growth area is the Web, but it is an uncategorizable sea of a million destinations, each defying in its own way the conventional logic of media and marketing.

ITUNES KILLED THE RADIO STAR

I came of age in the peak of the mass-culture era—the seventies and eighties. The average teenager then had access to a half dozen TV channels, and virtually everyone watched a few or more of the same handful of TV shows. There were three or four rock radio stations in any town that largely dictated what music people listened to; only a few lucky kids with money built record collections that ventured farther afield.

We all saw the same summer blockbusters in the theater and got our news from the same papers and broadcasts. About the only places you could explore outside the mainstream were the library and the comic book shop. As best I can recall, the only culture I was exposed

to other than mass culture was books and whatever my friends and I made up, and that traveled no farther than our own backyards.

Contrast my adolescence with that of Ben, a sixteen-year-old who grew up with the Internet. He's the single child of affluent parents in the tony North Berkeley Hills, so he's got a Mac in his bedroom, a fully stocked iPod (and a weekly iTunes allowance), and a posse of friends with the same. Like the rest of his teenage friends, Ben has never known a world without broadband, cell phones, MP3s, TiVo, and on-line shopping.

The main effect of all this connectivity is unlimited and unfiltered access to culture and content of all sorts, from the mainstream to the farthest fringe of the underground. Ben is growing up in a different world from the one I grew up in, a world far less dominated by any of the traditional media and entertainment industries. If you don't recognize yourself in the pages to come in this book, imagine Ben instead. His reality is the leading edge of all of our futures.

From Ben's perspective, the cultural landscape is a seamless continuum from high to low, with commercial and amateur content competing equally for his attention. He simply doesn't distinguish between mainstream hits and underground niches—he picks what he likes from an infinite menu where Hollywood movies and player-created video-game stunt videos are listed side by side.

Ben watches just two hours or so a week of regular TV, mostly *West Wing* (time shifted, of course) and *Firefly*, a canceled space serial he has stored on his TiVo. He also counts as TV the anime he downloads with BitTorrent, a peer-to-peer file-sharing technology, because it was originally broadcast on Japanese television (the English subtitles are often edited in by fans).

When it comes to movies, he's a sci-fi fan, so he's pretty mainstream. *Star Wars* is a passion, as was the *Matrix* series. But he also watches movies he downloads, such as amateur machinima (movies made by controlling characters in video games) and independent productions such as *Star Wars Revelations*, a fan-created tribute film with special effects that rival the Lucas originals.

Some of the music on his iPod is downloaded from iTunes, but most comes from his friends. When one of the group buys a CD, he or

she typically makes copies for everyone else. Ben's taste is mostly classic rock—Led Zeppelin and Pink Floyd—with a smattering of video-game soundtracks. The only radio he listens to is when his parents turn on NPR in the car.

Ben's reading ranges from *Star Wars* novels to Japanese manga, with a large helping of Web comics. He, like a few of his friends, is so into Japanese subculture that he's studying Japanese in school. When I was in school, kids studied Japanese because Japan was a dominant economic power and language skills were thought to open up career opportunities. But now kids study Japanese so they can create their own anime subtitles and dig deeper into manga than the relatively mainstream translated stuff.

Most of Ben's free time is spent online, both randomly surfing and participating in user forums such as Halo and *Star Wars* discussion sites. He's not interested in news—he reads no newspapers and watches no TV news—but follows the latest tech and subculture chatter on sites such as Slashdot (geek news) and Fark (weird news). He instant messages constantly all day with his ten closest friends. He doesn't text much on his cell phone, but he has friends that do. (Texting is preferred by those who are out and about a lot; IM is the chat channel of choice for those who tend to spend more time in their own rooms.) He plays video games with friends, mostly online. He thinks Halo 2 rocks, especially the user-modified levels.

I suspect that had I been born twenty-five years later, my teenage years would have been quite similar. The main difference between Ben's adolescence and my own is simply choice. I was limited to what was broadcast over the airwaves. He's got the Internet. I didn't have TiVo (or even cable); he has all that and BitTorrent, too. I had no idea there was even such a thing as manga, much less how to get it. Ben has access to it all. Would I have watched *Gilligan's Island* reruns if I'd been able to build a clan with friends in World of Warcraft online instead? I doubt it.

TV shows were more popular in the seventies than they are now not because they were better, but because we had fewer alternatives to compete for our screen attention. What we thought was the rising tide of common culture actually turned out to be less about the triumph of

Hollywood talent and more to do with the shepherding effect of broadcast distribution.

The great thing about broadcast is that it can bring one show to millions of people with unmatched efficiency. But it can't do the opposite—bring a million shows to one person each. Yet that is exactly what the Internet does so well. The economics of the broadcast era required hit shows—big buckets—to catch huge audiences. The economics of the broadband era are reversed. Serving the same stream to millions of people at the same time is hugely expensive and wasteful for a distribution network optimized for point-to-point communications.

There's still demand for big cultural buckets, but they're no longer the only market. The hits now compete with an infinite number of niche markets, of any size. And consumers are increasingly favoring the one with the most choice. The era of one-size-fits-all is ending, and in its place is something new, a market of multitudes.

This book is about that market.

This shattering of the mainstream into a zillion different cultural shards is something that upsets traditional media and entertainment no end. After decades of executives refining their skill in creating, picking, and promoting hits, those hits are suddenly not enough. The audience is shifting to something else, a muddy and indistinct proliferation of . . . Well, we don't have a good term for such non-hits. They're certainly not "misses," because most weren't aimed at world domination in the first place. They're "everything else."

It's odd that this should be an overlooked category. We are, after all, talking about the vast majority of everything. Most movies aren't hits, most music recordings don't make the top 100, most books aren't best-sellers, and most video programs don't even get measured by Nielsen, much less clean up in prime time. Many of them nevertheless record audiences in the millions worldwide. They just don't count as hits, and are therefore not counted.

But they're where the formerly compliant mass market is scattering to. The simple picture of the few hits that mattered and the everything else that didn't is now becoming a confusing mosaic of a million mini-markets and micro-stars. Increasingly, the mass market is turning into a mass of niches.

That mass of niches has always existed, but as the cost of reaching it falls—consumers finding niche products, and niche products finding consumers—it’s suddenly becoming a cultural and economic force to be reckoned with.

The new niche market is not replacing the traditional market of hits, just sharing the stage with it for the first time. For a century we have winnowed out all but the best-sellers to make the most efficient use of costly shelf space, screens, channels, and attention. Now, in a new era of networked consumers and digital everything, the economics of such distribution are changing radically as the Internet absorbs each industry it touches, becoming store, theater, and broadcaster at a fraction of the traditional cost.

Think of these falling distribution costs as a dropping waterline or a receding tide. As they fall, they reveal a new land that has been there all along, just underwater. These niches are a great uncharted expanse of products that were previously uneconomic to offer. Many of these kinds of products have always been there, just not visible or easy to find. They are the movies that didn’t make it to your local theater, the music not played on the local rock radio station, the sports equipment not sold at Wal-Mart. Now they’re available, via Netflix, iTunes, Amazon, or just some random place Google turned up. The invisible market has turned visible.

Other niche products are new, created by an emerging industry at the intersection between the commercial and noncommercial worlds, where it’s hard to tell when the professionals leave off and the amateurs take over. This is the world of bloggers, video-makers, and garage bands, all suddenly able to find an audience thanks to those same enviable economics of digital distribution.

THE 98 PERCENT RULE

This book began with a quiz I got wrong. One of the things I do as the editor of *Wired* is give speeches about technology trends. Because I started my career in the science world and then learned economics at *The Economist*, I look for those trends first in hard data. And, fortu-

nately enough, there has never been more data available. The secrets of twenty-first-century economics lie in the servers of the companies that are all around us, from eBay to Wal-Mart. Although it’s not always easy to get the raw numbers, the executives at those companies swim in that data every day and have a great intuitive feel for what’s meaningful and what isn’t. So the trick to trend-spotting is to ask them.

Which is what I was doing in January 2004, in the offices of Robbie Vann-Adibé, the CEO of Ecast, a “digital jukebox” company. Digital jukeboxes are just like regular jukeboxes—a big enclosure with speakers and blinking lights, often found in bars—with the difference that rather than a hundred CDs, they have a broadband connection to the Internet and patrons can choose from thousands of tracks that are downloaded and stored on a local hard drive.

During the course of our conversation, Vann-Adibé asked me to guess what percentage of the 10,000 albums available on the jukeboxes sold at least one track per quarter.

I knew, of course, that Vann-Adibé was asking me a trick question. The normal answer would be 20 percent because of the 80/20 Rule, which experience tells us applies practically everywhere. That is: 20 percent of products account for 80 percent of sales (and usually 100 percent of the profits).

But Vann-Adibé was in the digital content business, which is different. So I thought I’d go way out on a limb and venture that a whopping 50 percent of those 10,000 albums sold at least one track a quarter.

Now, on the face of it, that’s absurdly high. Half of the top 10,000 books in a typical book superstore don’t sell once a quarter. Half of the top 10,000 CDs at Wal-Mart don’t sell once a quarter; indeed, Wal-Mart doesn’t even carry half that many CDs. It’s hard to think of any market where such a high fraction of such a large inventory sells. But my sense was that digital was different, so I took a chance on a big number.

I was, needless to say, way, way off. The answer was 98 percent.

“It’s amazing, isn’t it?” Vann-Adibé said. “Everyone gets that wrong.” Even he had been stunned: As the company added more titles to its collections, far beyond the inventory of most record stores and

into the world of niches and subcultures, they continued to sell. And the more the company added, the more they sold. The demand for music beyond the hits seemed to be limitless. True, the songs didn't sell in big numbers, but nearly all of them sold something. And because these were just bits in a database that cost nearly nothing to store and deliver, all those onesies and twosies started to add up.

What Vann-Adibé had discovered was that the aggregate market for niche music was huge, and effectively unbounded. He called this the "98 Percent Rule." As he later put it to me, "In a world of almost zero packaging cost and instant access to almost all content in this format, consumers exhibit consistent behavior: They look at almost everything. I believe that this requires major changes by the content producers—I'm just not sure what changes!"

I set out to answer that question. I realized that his counterintuitive statistic contained a powerful truth about the new economics of entertainment in the digital age. With unlimited supply, our assumptions about the relative roles of hits and niches were all wrong. Scarcity requires hits—if there are only a few slots on the shelves or the airwaves, it's only sensible to fill them with the titles that will sell best. And if that's all that's available, that's all people will buy.

But what if there are infinite slots? Maybe hits are the wrong way to look at the business. There are, after all, a lot more non-hits than hits, and now both are equally available. What if the non-hits—from healthy niche product to outright misses—all together added up to a market as big as, if not bigger than, the hits themselves? The answer to that was clear: It would radically transform some of the largest markets in the world.

And so I embarked on a research project that was to take me to all the leaders in the emerging digital entertainment industry, from Amazon to iTunes. Everywhere I went the story was the same: Hits are great, but niches are emerging as the big new market. The 98 Percent Rule turned out to be nearly universal. Apple said that every one of the then 1 million tracks in iTunes had sold at least once (now its inventory is twice that). Netflix reckoned that 95 percent of its 25,000 DVDs (that's now 55,000) rented at least once a quarter. Amazon didn't give out an exact number, but independent academic research

on its book sales suggested that 98 percent of its top 100,000 books sold at least once a quarter, too. And so it went, from company to company.

Each company was impressed by the demand they were seeing in categories that had been previously dismissed as beneath the economic fringe, from the British television series DVDs that are proving surprisingly popular at Netflix to the back-catalog music that's big on iTunes. I realized that, for the first time, I was looking at the true shape of demand in our culture, unfiltered by the economics of scarcity.

That shape is, to be clear, really, really weird. To think that basically everything you put out there finds demand is just odd. The reason it's odd is that we don't typically think in terms of one unit per quarter. When we think about traditional retail, we think about what's going to sell a lot. You're not much interested in the occasional sale, because in traditional retail a CD that sells only one unit a quarter consumes exactly the same half-inch of shelf space as a CD that sells 1,000 units a quarter. There's a value to that space—rent, overhead, staffing costs, etc.—that has to be paid back by a certain number of inventory turns per month. In other words, the onesies and twosies waste space.

However, when that space doesn't cost anything, suddenly you can look at those infrequent sellers again, and they begin to have value. This was the insight that led to Amazon, Netflix, and all the other companies I was talking to. All of them realized that where the economics of traditional retail ran out of steam, the economics of online retail kept going. The onesies and twosies were still only selling in small numbers, but there were so, so *many* of them that in aggregate they added up to a big business.

Throughout the first half of 2004 I fleshed out this research in speeches, the thesis advancing with each talk. Originally the speech was called "The 98 Percent Rule." Then it was "New Rules for the New Entertainment Economy" (not one of my better naming moments).

But by then I had some hard data, thanks to Rhapsody, which is one of the online music companies. They had given me a month's worth of customer usage data, and when I graphed it out, I realized that the curve was unlike anything I'd seen before.

It started like any other demand curve, ranked by popularity. A few

hits were downloaded a huge number of times at the head of the curve, and then it fell off steeply with less popular tracks. But the interesting thing was that it never fell to zero. I'd go to the 100,000th track, zoom in, and the downloads per month were still in the thousands. And the curve just kept going: 200,000, 300,000, 400,000 tracks—no store could ever carry this much music. Yet as far as I looked, there was still demand. Way out at the end of the curve, tracks were being downloaded just four or five times a month, but the curve still wasn't at zero.

In statistics, curves like that are called "long-tailed distributions," because the tail of the curve is very long relative to the head. So all I did was focus on the tail itself, turn it into a proper noun, and "The Long Tail" was born. It started life as slide 20 of one of my "New Rules" presentations. I think it was Reed Hastings, the CEO of Netflix, who convinced me that I was burying my lead. By the summer of 2004 "The Long Tail" was not just the title of my speeches; I was nearly finished with an article of the same name for my own magazine.

When "The Long Tail" was published in *Wired* in October 2004, it quickly became the most cited article the magazine had ever run. The three main observations—(1) the tail of available variety is far longer than we realize; (2) it's now within reach economically; (3) all those niches, when aggregated, can make up a significant market—seemed indisputable, especially backed up with heretofore unseen data.

TAILS EVERYWHERE

One of the most encouraging aspects of the overwhelming response to the original article was the breadth of industries in which it resonated. The article originated as an analysis of the new economics of the entertainment and media industries, and I only expanded it a bit to mention in passing that companies such as eBay (with used goods) and Google (with small advertisers) were also Long Tail businesses. Readers, however, saw the Long Tail everywhere, from politics to public relations, and from sheet music to college sports.

What people intuitively grasped was that new efficiencies in distri-

bution, manufacturing, and marketing were changing the definition of what was commercially viable across the board. The best way to describe these forces is that they are turning unprofitable customers, products, and markets into profitable ones. Although this phenomenon is most obvious in entertainment and media, it's an easy leap to eBay to see it at work more broadly, from cars to crafts.

Seen broadly, it's clear that the story of the Long Tail is really about the economics of abundance—what happens when the bottlenecks that stand between supply and demand in our culture start to disappear and everything becomes available to everyone.

People often ask me to name some product category that does not lend itself to Long Tail economics. My usual answer is that it would be in some undifferentiated commodity, where variety is not only absent but unwanted. Like, for instance, flour, which I remembered being sold in the supermarket in a big bag labeled "Flour." Then I happened to step inside our local Whole Foods grocery and realized how wrong I was: Today the grocery carries more than twenty different types of flour, ranging from such basics as whole wheat and organic varieties to exotics such as amaranth and blue cornmeal. There is, amazingly enough, already a Long Tail in flour.

Our growing affluence has allowed us to shift from being bargain shoppers buying branded (or even unbranded) commodities to becoming mini-connoisseurs, flexing our taste with a thousand little indulgences that set us apart from others. We now engage in a host of new consumer behaviors that are described with intentionally oxymoronic terms: "massclusivity," "slivercasting," "mass customization." They all point in the same direction: more Long Tails.

A PREVIEW OF TWENTY-FIRST-CENTURY ECONOMICS

This book is partly an economic research project, with the help and involvement of students and professors from the Stanford, MIT, and Harvard business schools. It's partly the fruit of more than a hundred speeches, brainstorming sessions, and site visits with companies and industry groups that see the Long Tail changing their world. And it's

partly a collaboration with the dozens of companies and executives who shared many megabytes of internal data, giving me an unprecedented view on the emerging micro-economics of markets in the online age.

What's fascinating about this moment is that the economics of the twenty-first century are already evident in outline form in the databases of the Googles, Amazons, Netfixes, and iTunes of the world. In those many terabytes of user behavior data is a clue to how consumers will behave in markets of infinite choice, a question that hadn't been meaningful until recently but has now become essential to understand.

Surprisingly, very few economists are looking at this data, mostly because they haven't asked (most of the academics I worked with are in business schools, only a few of them are economists). There are some exceptions—University of California Berkeley economist Hal Varian works part-time at Google, and auction-theory economists unsurprisingly love eBay—but they're rare. Some of the data in this book has never before seen the light of day.

Given the uncharted waters, I solicited a lot of help from experts in all corners. As an experiment, I worked through many of the trickier conceptual and articulation issues in public, on my blog at thelongtail.com. The usual process would go like this: I'd post a half-baked effort at explaining how the 80/20 Rule is changing, for instance, and then dozens of smart readers would write comments, emails, or their own blog posts to suggest ways to improve it. Somehow this wonky public brainstorming managed to attract an average of more than 5,000 readers a day.

In software, developers release early ("beta") versions of their code to their most avid users. In exchange for the privileged early look at the program, these users test it on their own machines, in their own way, and find errors that the developer missed. Such beta-testing is essential to creating robust software applications. My hope is that the same process—stress-testing many of my ideas in public—has led to a better, or at least sounder, book.

I should note here the difference between beta-testing ideas in public and actually writing a book in public. Although many have tried

to do the latter—posting draft chapters online and sometimes even opening the text to collective editing—I chose to use the blog mostly as a public diary of my research in progress. The actual writing of the book, and most of the words in the following pages, I did offline.

Finally, one more note on parentage. Although I coined the term "The Long Tail," I can't claim any credit for creating the concept of using the efficient economics of online retail to aggregate a large inventory of relatively low sellers. That would be Amazon's Jeff Bezos, circa 1994. Most of what I've learned has come from talking to him, his counterparts at Netflix and Rhapsody, and others who have all been acting on this for years.

Those entrepreneurs are the real inventors here. What I've tried to do is synthesize the results into a framework. That is, of course, what economics does: It seeks to find neat, easily understood frameworks that describe real-world phenomena. Coming up with the framework is an advance in itself, but it pales next to the original inventions of all those who discovered and acted on the phenomena in the first place.

5

 THE NEW PRODUCERS

 NEVER UNDERESTIMATE THE POWER OF A MILLION
 AMATEURS WITH KEYS TO THE FACTORY

On the night of February 23, 1987, the underground Kamiokande II observatory in Japan detected twenty-four neutrinos in a burst lasting thirteen seconds. Although twenty-four neutrinos may not sound like a lot, the observatory usually detects only two or three an hour, and rarely in a pack. So this was something special. But what it actually meant would have to wait a few hours, for other observations to be reported.

Astrophysicists had long theorized that when a star explodes, most of its energy is released as neutrinos—low-mass, subatomic particles that fly through planets like bullets through tissue paper. Part of the theory is that in the early phase of this type of explosion, the only observable evidence is a shower of such particles; it then takes another few hours for the inferno to emerge as visible light. As a result, scientists predicted that when a star went supernova near us, we'd detect the neutrinos about three hours before we'd see the burst in the visible spectrum.

The way to test this correlation between neutrinos and visible light was to make both observations and measure the time difference be-

tween them. But the problem with the optical part of these paired observations was that you had to be looking in the right part of the sky. This wasn't much of a problem for the neutrino observatories. Because of its spherical layout, the detector hall of Kamiokande could record neutrinos penetrating the Earth regardless of which direction they came from. Yet to see the explosion in visible light, a telescope would have to be pointed at the exact right spot at the exact right time. And, needless to say, there was an awful lot of sky to watch.

There just weren't enough professional astronomers who could observe enough of the heavens to have much, if any, chance of spotting such an event. But there were thousands of amateur astronomers all too happy to do that job themselves. Armed with relatively inexpensive computer-guided telescopes with Dobsonian optics, which allow quite large apertures (twelve inches is not unusual) in telescopes less than five feet long, and sensitive CCD (charge-coupled device) sensors that can collect more light than the human eye, contemporary amateur astronomers can photograph the skies better than astronomers with house-sized telescopes could a century ago.

The first person to see Supernova 1987A was an observer somewhere between the amateurs and the pros. Ian Shelton, a Canadian grad-school dropout, was housekeeping an observatory in the Chilean Andes in exchange for time on its twenty-four-inch telescope when academic astronomers weren't using it. One of those free times was the windy night of February 23. That night Shelton decided to use the telescope to run a three-hour exposure on the Large Magellanic Cloud.

As it happened, exactly 168,000 years earlier and exactly 168,000 light-years away, a star had exploded on the edge of the Tarantula Nebula. From Earth and Shelton's view, however, the explosion looked like it was happening right then: a splash of light suddenly appearing in one corner of the Cloud where nothing of note had been before. Shelton stared at the photographic plates for twenty minutes before heading outside to see it with his own eyes. Sure enough, there it was: the first supernova to be witnessed by the naked eye since 1604.

The connection between Shelton and the Kamiokande II observatory is one of time. The neutrino observatory spotted its burst at 7:35

Universal time. Shelton observed the first bright light around 10:00 Universal time—a little less than three hours after the neutrino shower. So far, right on theory. However, could it have shown up even earlier, before Shelton was watching?

Fortunately, two other dedicated amateur astronomers were at work that night using smaller, nonprofessional telescopes. In New Zealand, Albert Jones, a veteran credited with more than half a million observations, had taken a good look at the Tarantula Nebula at 9:30 UT but had seen nothing unusual. Robert McNaught, another amateur, photographed the explosion at 10:30 UT in Australia, confirming Shelton's timing. So the light arrived somewhere between 9:30 and 10:00.

That is how one of the greatest astronomical discoveries of the twentieth century unfolded. A key theory explaining how the universe works was confirmed thanks to amateurs in New Zealand and Australia, a former amateur trying to turn professional in Chile, and professional physicists in the United States and Japan. When a scientific paper finally announced the discovery to the world, all of them shared authorship.

Demos, a British think tank, described this in a 2004 report as a key moment in the arrival of a "Pro-Am" era, a time when professionals and amateurs work side by side: "Astronomy used to be done in 'big science' research institutes. Now it is also done in Pro-Am collaboratives. Many amateurs continued to work on their own and many professionals were still ensconced in their academic institutions. But global research networks sprang up, linking professionals and amateurs with shared interests in flare stars, comets and asteroids."

As Timothy Ferris points out in *Seeing in the Dark*, his history of modern amateur astronomy: "If one were to choose a date at which astronomy shifted from the old days of solitary professionals at their telescopes to a worldwide web linking professionals and amateurs . . . a good candidate would be the night of February 23, 1987." Demos concludes: "Astronomy is fast becoming a science driven by a vast open-source Pro-Am movement working alongside a much smaller body of professional astronomers and astrophysicists."

The enabling technologies of this Pro-Am movement in astronomy are Dobsonian optics, CCDs, and the arrival of the Internet as a mechanism for sharing information. These tools have swelled the ranks of the amateur astronomers and vastly increased their impact. Over the past two decades, astronomy has become one of the most democratized fields in science, in part because it's so clear what an important role the amateurs play.

NASA often calls on amateurs to watch for specific asteroids that might be headed for Earth, an observation task coordinated via an email message group called the Minor Planet Mailing List that's run by Richard Kowalski, a forty-two-year-old baggage handler at US Airways in Florida by day and an astronomer by night. Some of the eight hundred amateurs on the list record their observations for fun; others hope to be immortalized by having an important discovery named after them. What's notable is that none of them do it for money.

Astronomy has a natural place for volunteer manpower. Again, the problem with the sky is that you need to be looking at the right place at the right time to witness most interesting new phenomena, such as asteroids or stellar evolution. It's less a matter of how big or expensive the telescope, and more a matter of how many eyeballs are transfixed on the sky at any given moment. Amateurs multiply the manpower of astronomy many times—and not just by looking at the stars from their backyards.

SETI@home ("Search for Extraterrestrial Intelligence at home") is a project that harnesses the spare computing power of more than half a million home computers. After collecting hours and hours of white noise recorded from space, the project distributes its radio telescope data to the computers of volunteers. When they're not using their computers, a special screen-saver kicks in. While it displays cosmic imagery, it scans bits of each recording in the hopes of locating a signal that may have come from alien intelligence. By divvying up its data to these volunteer computers, the project is able to examine a far greater number of signals than it would otherwise; and all anyone has to do to participate is download some software.

Another project has open-sourced the analysis of Mars imagery.

NASA put up decades-old photos snapped by the Viking orbiters and asked Web visitors to click on all the craters they could see, classifying them as fresh, degraded, or “ghost.” Usually, this is a tedious job for scientists and grad students that can take months or years, but in just three months the “Mars Clickworkers” project got volunteers to identify more than 200,000 craters. Averaged over all the clicks, this amateur collective was almost as accurate as expert planetary geologists.

In “open-source” software, where anyone can contribute to a project, the mantra is “With enough eyes, all bugs are trivial.” Likewise for astronomy: With enough eyes, we’ll see the asteroid with our name on it—and early enough to do something about it.

Of course, there are limits to what Pro-Ams can achieve. They’re largely collecting data, not creating new theories of astrophysics. Sometimes, they are unable to analyze properly the data they collect. Nevertheless, their place in the field seems assured. As John Lankford, a historian of science, put it in *Sky & Telescope* magazine, the bible of U.S. amateur astronomers: “There will always remain a division of labor between professionals and amateurs. But it may be more difficult to tell the two groups apart in the future.”

DEMOCRATIZING THE TOOLS OF PRODUCTION

What’s new about this is the way it’s done, not the concept itself. Indeed, Karl Marx was perhaps the original prophet of the Pro-Am economy. As Demos notes, “In *The German Ideology*, written between 1845 and 1847, Marx maintained that labor—forced, unspontaneous and waged work—would be superseded by self-activity.” Eventually, he hoped, there would be a time when “material production leaves every person surplus time for other activities.” Marx evoked a communist society in which “. . . nobody has one exclusive sphere of activity but each can become accomplished in any branch he wishes . . . to hunt in the morning, fish in the afternoon, rear cattle in the evening, criticize after dinner, just as I have a mind without ever becoming hunter, fisherman, shepherd or critic.”

To continue with Marx’s vocabulary, Pro-Ams are a creation of the

first force of the Long Tail, the democratization of the tools of production.

The same effect we see in astronomy is playing out in countless other fields. Just as the electric guitar and the garage democratized pop music forty years ago, desktop creation and production tools are democratizing the studio. Apple’s GarageBand, free with every Mac, greets a user with the suggestion to “Record your next big hit,” and provides the tools to do just that. Likewise, digital video cameras and desktop editing suites (free with every copy of Windows and every Mac) are putting the sort of tools into the hands of the average home moviemaker that were once reserved for professionals alone.

Then there’s the written word, always the leading edge of egalitarianism. Although it was the photocopier that first put lie to the aphorism that “the power of the press goes to those who own them,” it’s blogging that has really sparked the renaissance of the amateur publisher. Today, millions of people publish daily for an audience that is collectively larger than any single mainstream media outlet can claim. What sparked blogging was, again, democratized tools: the arrival of simple, cheap software and services that made publishing online so easy that anyone could do it.

So, too, for desktop photo editing and printing, video games that encourage players to create and share their own alternative levels, and print-on-demand book publishing. A few decades ago, there were two reasons why most of us weren’t making hit movies: (1) we didn’t have access to the necessary tools, and (2) we didn’t have the talent. Today, there’s only one excuse—and even that is not as solid as it was. Hollywood, for all its efficiencies, can’t find every potentially great filmmaker on the planet. Technology, cheap and ubiquitous, can do far better. Once upon a time, talent eventually made its way to the tools of production; now it’s the other way around.

The consequence of all this is that we’re starting to shift from being passive consumers to active producers. And we’re doing it for the love of it (the word “amateur” derives from the Latin *amator*, “lover,” from *amare*, “to love”). You can see it all around you—the extent to which amateur blogs are sharing attention with mainstream media, small-time bands are releasing music online without a record label, and fel-

low consumers dominate online reviewing. It's as if the default setting of production has shifted from "Earn the right to do it" to "What's stopping you?"

Author Doc Searls calls this a shift from consumerism to participative "producerism":

The "consumer economy" is a producer-controlled system in which consumers are nothing more than energy sources that metabolize "content" into cash. This is the absolutely corrupted result of the absolute power held by producers over consumers since producers won the Industrial Revolution.

Apple is giving consumers tools that make them producers. This practice radically transforms both the marketplace and the economy that thrives on it.

I can see it in my own young children, who are, as I write, into machinima—short computer-animated movies made with video-game software. Using the 3D rendering engines of games such as Halo 2 or the Sims for all the visuals, machinima directors need only write a script, control the characters, and voice the lines. Everything else—sets, camera, character, and vehicle models—is done by the game software. It's like having a mini-Pixar in every Xbox or PC.

The first reaction of the kids was to watch and enjoy the machinima movies as entertainment. Their second was to express curiosity as to how they're made. And their third was to ask if they could make one themselves. (The answer, of course, is *yes*.) What machinima lacks in Hollywood polish, it more than makes up for in creative inspiration. A generation is growing up watching people just like them produce impressive works of creativity. This can't help but make an impression.

It's one thing to see a movie or listen to music and to think "genius"—that some gifted person and exalted apparatus has put together this unique work of art we appreciate. However, once you know what's behind the curtain, you begin to realize that it could be *you*. It is when the tools of production are transparent that we are inspired to create. When people understand how great work is made, they're more likely to want to do it themselves.

Today, millions of ordinary people have the tools and the role models to become amateur producers. Some of them will also have talent and vision. Because the means of production have spread so widely and to so many people, the talented and visionary ones, even if they're just a small fraction of the total, are becoming a force to be reckoned with. Don't be surprised if some of the most creative and influential work in the next few decades comes from this Pro-Am class of inspired hobbyists, not from the traditional sources in the commercial world. The effect of this shift means that the Long Tail will be populated at a pace never before seen.

THE WIKIPEDIA PHENOMENON

In January 2001, a wealthy options trader named Jimmy Wales set out to build a massive online encyclopedia in an entirely new way—by tapping the collective wisdom of millions of amateur experts, semi-experts, and just regular folks who thought they knew something. This encyclopedia would be freely available to anyone; and it would be created not by paid experts and editors, but by whoever wanted to contribute. Wales started with a few dozen prewritten articles and a software application called a Wiki (named for the Hawaiian word meaning "quick" or "fast"), which allows anybody with Web access to go to a site and edit, delete, or add to what's there. The ambition: Nothing less than to construct a repository of knowledge to rival the ancient library of Alexandria.

This was, needless to say, controversial.

For one thing, this is not how encyclopedias are supposed to be made. From the beginning, compiling authoritative knowledge has been the job of scholars. It started with a few solo polymaths who dared to try the impossible. In ancient Greece, Aristotle single-handedly set out to record all the knowledge of his time. Four hundred years later, the Roman nobleman Pliny the Elder cranked out a thirty-seven-volume set of the day's knowledge. The Chinese scholar Tu Yu wrote an encyclopedia on his own in the ninth century. And in the 1700s, Diderot and a few of his pals (including Voltaire and Rousseau)

took twenty-nine years to create the *Encyclopédie, ou Dictionnaire Raisonné des Sciences, des Arts et des Métiers*.

Individual work gradually evolved into larger team efforts, especially after the arrival of the Industrial Revolution. In the late eighteenth century, several members of the Scottish Enlightenment started to apply the industrial principles of scientific management and the lessons of assembly lines to the creation of an encyclopedia such as the world had never before seen. The third edition of the *Encyclopædia Britannica*, published between 1788 and 1797, amounted to eighteen volumes plus a two-volume supplement, totaling over 16,000 pages. Groups of experts were recruited to write scholarly articles under the direction of a manager, organized by a detailed work chart.

Now Wales has introduced a third model: the open collective. Instead of one really smart guy or a number of handpicked smart guys, Wikipedia draws on tens of thousands of people of all sorts—ranging from real experts to interested bystanders—with a lot of volunteer curators adopting entries and keeping an eye on their progression. In Wales's encyclopedia calculus, 50,000 self-selected Wikipedians equal one Pliny the Elder.

As writer Daniel Pink puts it, "Instead of clearly delineated lines of authority, Wikipedia depends on radical decentralization and self-organization; open source in its purest form. Most encyclopedias start to fossilize the moment they're printed on a page. However, add Wiki software and some helping hands and you get something self-repairing and almost alive. A different production model creates a product that's fluid, fast, fixable, and free."

In 2001, that idea seemed preposterous. But by 2005, this non-profit venture had become the largest encyclopedia on the planet. Wikipedia offers more than 1 million articles in English—compared with *Britannica's* 80,000 and *Encarta's* 4,500—fashioned by more than 20,000 contributors. Tack on the editions in seventy-five other languages, including Esperanto and Kurdish, and the total Wikipedia article count tops 3.5 million.

All you need to contribute to Wikipedia is Internet access: Every entry has an "Edit This Page" button on it, available to all. Each of us

is an expert in something, and the beauty of Wikipedia is that there is practically no subject so narrow that it can't have an entry. This is in stark contrast to *Britannica*. If you open that great encyclopedia and find either no entry for what you're looking for or an entry that seems deficient, there's little you can do but shake your fist or write a letter to the editor (expecting no response). With Wikipedia, however, you fix it or create it yourself. This kind of shift from passive resentment to active participation makes the big difference. To remix the old joke about the weather, everybody complains about the encyclopedia, but now you *can* do something about it.

THE PROBABILISTIC AGE

Much is made of the fact that Wikipedia's entries are "non-authoritative," which is a way of saying they're not invariably accurate. This is, of course, inevitable when anyone can write them. Unlike *Britannica*, where each entry is scrubbed, checked, and labored over by responsible professionals, each Wikipedia entry simply arrives, conjured from the vacuum by the miracle of the "Edit This Page" button.

In late 2005, John Seigenthaler Sr. wrote an op-ed in *USA Today* about his own Wikipedia entry; the entry started this way:

John Seigenthaler Sr. was the assistant to Attorney General Robert Kennedy in the early 1960's. For a brief time, he was thought to have been directly involved in the Kennedy assassinations of both John, and his brother, Bobby. Nothing was ever proven.

Aside from the part about him being Robert Kennedy's assistant in the 1960s, virtually everything else about the entry was false and slanderous. Seigenthaler called Wales and got him to delete the entry (although he could have easily done that himself), but after he wrote about the experience it led to a national debate over whether Wikipedia could be trusted, a question that continues today.

The answer is not a simple yes or no, because it is the nature of user-created content to be as messy and uncertain at the microscale,

which is the level at which we usually experience it, as it is amazingly successful at the big-picture macroscale. It just has to be understood for what it is.

Wikipedia, like Google and the collective wisdom of millions of blogs, operates on the alien logic of probabilistic statistics—a matter of likelihood rather than certainty. But our brains aren't wired to think in terms of statistics and probability. We want to know whether an encyclopedia entry is right or wrong. We want to know that there's a wise hand (ideally human) guiding Google's results. We want to trust what we read.

When professionals—editors, academics, journalists—are running the show, we at least know that it's someone's job to look out for such things as accuracy. But now we're depending more and more on systems where nobody's in charge; the intelligence is simply “emergent,” which is to say that it appears to arise spontaneously from the number-crunching. These probabilistic systems aren't perfect, but they are statistically optimized to excel over time and large numbers. They're designed to “scale,” or improve with size. And a little slop at the microscale is the price of such efficiency at the macroscale.

But how can that be right when it feels so wrong?

There's the rub. This tradeoff is just hard for people to wrap their heads around. There's a reason why we're still debating Darwin. And why *The Wisdom of Crowds*, James Surowiecki's book on Adam Smith's invisible hand and how the many can be smarter than the few, is still surprising (and still needs to be read) more than two hundred years after the great Scotsman's death. Both market economics and evolution are probabilistic systems, which are simply counterintuitive to our mammalian brains. The fact that a few smart humans figured this out and used that insight to build the foundations of our modern economy, from the stock market to Google, is just evidence that our mental software (our collective knowledge) has evolved faster than our hardware (our neural wiring).

Probability-based systems are, to use writer Kevin Kelly's term, “out of control.” His seminal book by that name looks at example after example, from democracy to bird-flocking, where order arises from what appears to be chaos, seemingly reversing entropy's arrow. The

book is more than a dozen years old, and decades from now we'll still find the insight surprising. But it's right.

Is Wikipedia “authoritative”? Well, no. But what really is? *Britannica* is reviewed by a smaller group of reviewers with higher academic degrees on average. There are, to be sure, fewer (if any) total clunkers or fabrications than in Wikipedia. But it's not infallible either; indeed a 2005 study by *Nature*, the scientific journal, reported that in forty-two entries on science topics there were an average of four errors per entry in Wikipedia and three in *Britannica*. And shortly after the report came out, the Wikipedia entries were corrected, while *Britannica* will have to wait for its next reprinting.

Britannica's biggest errors are of omission, not commission. It is shallow in some categories and out of date in many others. And then there are the millions of entries that it simply doesn't—and can't, given its editorial process—have. But Wikipedia *can* scale itself to include those and many more. And it is updated constantly.

The advantage of probabilistic systems is that they benefit from the wisdom of the crowd and as a result can scale nicely both in breadth and depth. But because they do this by sacrificing absolute certainty on the microscale, you need to take any single result with a grain of salt. Wikipedia should be the first source of information, not the last. It should be a site for information exploration, not the definitive source of facts.

The same is true for blogs, no single one of which is authoritative. Blogs are a Long Tail, and it is always a mistake to generalize about the quality or nature of content in the Long Tail—it is, by definition, variable and diverse. But collectively blogs are proving more than an equal to mainstream media. You just need to read more than one of them before making up your own mind.

Likewise for Google, which seems both omniscient and inscrutable. It makes connections that you or I might not, because they emerge naturally from math on a scale we can't comprehend. Google is arguably the first company to be born with the alien intelligence of the Web's “massive-scale” statistics hardwired into its DNA. That's why it's so successful, and so seemingly unstoppable.

Author Paul Graham puts it like this:

The Web naturally has a certain grain, and Google is aligned with it. That's why their success seems so effortless. They're sailing with the wind, instead of sitting becalmed praying for a business model, like the print media, or trying to tack upwind by suing their customers, like Microsoft and the record labels. Google doesn't try to force things to happen their way. They try to figure out what's going to happen, and arrange to be standing there when it does.

The Web is the ultimate marketplace of ideas, governed by the laws of big numbers. That grain Graham sees is the weave of statistical mechanics, the only logic that such really large systems understand. Perhaps someday we will, too.

THE POWER OF PEER PRODUCTION

As a whole, Wikipedia is arguably the best encyclopedia in the world: bigger, more up-to-date, and in many cases deeper than even *Britannica*. But at the individual entry level, the quality varies. Along with articles of breathtaking scholarship and erudition, there are plenty of "stubs" (placeholder entries) and even autogenerated spam.

In the popular entries with many eyes watching, Wikipedia shows a remarkable resistance to vandalism and ideological battles. One study by IBM found that the mean repair time for damage in high-profile Wikipedia entries such as "Islam" is less than four minutes. This is not the work of the professional encyclopedia police. It is simply the emergent behavior of a Pro-Am swarm of self-appointed curators. Against all expectations, the system works brilliantly well. And as Wikipedia grows, this rapid self-repairing property will spread to more entries.

The point is not that every Wikipedia entry is probabilistic, but that the *entire encyclopedia* behaves probabilistically. Your odds of getting a substantive, up-to-date, and accurate entry for any given subject are excellent on Wikipedia, even if every individual entry isn't excellent.

To put it another way, the quality range in *Britannica* goes from, say, 5 to 9, with an average of 7. Wikipedia goes from 0 to 10, with an

average of, say, 5. But given that Wikipedia has ten times as many entries as *Britannica*, your chances of finding a reasonable entry on the topic you're looking for are actually higher on Wikipedia.

What makes Wikipedia really extraordinary is that it improves over time, organically healing itself as if its huge and growing army of tenders were an immune system, ever vigilant and quick to respond to anything that threatens the organism. And like a biological system, it evolves, selecting for traits that help it stay one step ahead of the predators and pathogens in its ecosystem.

The traditional process of creating an encyclopedia—professional editors, academic writers, and peer review—aims for perfection. It seldom gets there, but the pursuit of accuracy and clarity results in a work that is consistent and reliable, but also incredibly time-consuming and expensive to produce. Likewise for most other products of the professional publishing industry: One can expect that a book will, in fact, have printing on both sides of the pages where intended and will be more or less spelled correctly. There is a quality threshold, below which the work does not fall.

With probabilistic systems, though, there is only a statistical level of quality, which is to say: Some things will be great, some things will be mediocre, and some things will be absolutely crappy. That's just the nature of the beast. The mistake of many of the critics is to expect otherwise. Wikipedia is simply a different animal from *Britannica*. It's a living community rather than a static reference work.

The true miracle of Wikipedia is that this open system of amateur user contributions and edits doesn't simply collapse into anarchy. Instead, it has somehow self-organized the most comprehensive encyclopedia in history. Reversing entropy's arrow, Jimmy Wales's catalytic moment—putting up a few initial entries and a mechanism for others to add to them—has actually created order from chaos.

The result is a very different kind of encyclopedia, one completely unbounded by space and production constraints. It offers all the expected entries of any world-class reference work and then hundreds of thousands of unexpected ones, ranging from articles that go into textbook-like depth in fields such as quantum mechanics to biograph-

ical entries on comic book characters. Or, to put it another way, it's got all the hits plus a huge number of niches.

The classic model of the encyclopedia is a curated list of received cultural literacy. There is the basic canon, which must be recognized by authorities. Then, there are other entries of diminishing length until you get to that line at which the priests of *Britannica* decide "This is not worthy." There, the classic encyclopedia ends. Wikipedia, on the other hand, just *keeps going*.

In a sense, you can think of Wikipedia as equivalent to Rhapsody, the music site. There are the popular top 1,000, which can be found in any encyclopedia: Julius Caesar, World War II, Statistics, etc. These are like the hit songs. With these, Wikipedia is competing with professionals at their best, who produce well-written, authoritative entries that deploy facts with the easy comfort that comes with great scholarship. The main advantage of the user-created Wikipedia model for these entries is its ability to be up-to-date, have unlimited length and visual aids (such as photos and charts), include copious links to support material elsewhere, and perhaps, better represent alternate views and controversies.

In the middle of the curve, from the 1,000th entry to where *Britannica* ends at 80,000, are the narrower subjects: Caesarian Section, Okinawa, Regression Analysis, etc. Here, the Wikipedia model begins to pull ahead of its professional competition. Unlimited space means that the Wikipedia entries tend to be longer and more comprehensive. While the average length of a *Britannica* entry is 678 words, more than 200,000 Wikipedia entries (more than two entire *Britannicas*) are longer than that. Meanwhile, the external links and updated information emerge as a key advantage as Wikipedia becomes a launching place for further research.

Then there is the Tail, from 80,000 to 1 million. These are the entries that Wikipedia has that no other encyclopedia even attempts to include. Its articles on these subjects—Caesar Cipher, Canned Spam, Spearman's Rank Correlation Coefficient—range from among the best in Wikipedia (those written by passionate experts) to the worst (self-promotion, score-settling, and pranks). While many critics focus on the worst entries, the really important thing about Wikipedia's Tail is that

there is nothing else like it *anywhere*. From hard-core science to up-to-the-minute politics, Wikipedia goes where no other encyclopedia—whether constrained by paper or DVD limitations—can. *Britannica* doesn't have an entry about the Long Tail phenomenon (yet), but Wikipedia's entry is not only well written and thorough, it's also 1,500 words long (and none of it was written by me!).

Wikipedia authors tend to be enthusiastically involved, liberated, and motivated by the opportunity to improve public understanding of some subject they know and love, a population that has, in five short years, grown a thousandfold with an invasion of empowered amateurs using the simple, newly democratized tools of encyclopedia production: a Web browser and an Internet connection.

This is the world of "peer production," the extraordinary Internet-enabled phenomenon of mass volunteerism and amateurism. We are at the dawn of an age where most producers in any domain are unpaid, and the main difference between them and their professional counterparts is simply the (shrinking) gap in the resources available to them to extend the ambition of their work. When the tools of production are available to everyone, everyone becomes a producer.

THE REPUTATION ECONOMY

Why do they do it? Why does anyone create something of value (from an encyclopedia entry to an astronomical observation) without a business plan or even the prospect of a paycheck? The question is a key one to understanding the Long Tail, partly because so much of what populates the curve does not start with commercial aim. More important, this question matters because it represents yet another example of where our presumptions about markets must be rethought. The motives to create are not the same in the head as they are in the tail. One economic model doesn't fit all. You can think of the Long Tail starting as a traditional monetary economy at the head and ending in a non-monetary economy in the tail. In between the two, it's a mixture of both.

Up at the head, where products benefit from the powerful, but ex-

pensive, channels of mass-market distribution, business considerations rule. It's the domain of professionals, and as much as they might love what they do, it's a job, too. The costs of production and distribution are too high to let economics take a backseat to creativity. Money drives the process.

Down in the tail, where distribution and production costs are low (thanks to the democratizing power of digital technologies), business considerations are often secondary. Instead, people create for a variety of other reasons—expression, fun, experimentation, and so on. The reason one might call it an economy at all is that there is a coin of the realm that can be every bit as motivating as money: *reputation*. Measured by the amount of attention a product attracts, reputation can be converted into other things of value: jobs, tenure, audiences, and lucrative offers of all sorts.

Tim Wu, a Columbia University law professor, calls this the “exposure culture.” Using blogs as an example, he writes,

The exposure culture reflects the philosophy of the Web, in which getting noticed is everything. Web authors link to each other, quote liberally, and sometimes annotate entire articles. E-mailing links to favorite articles and jokes has become as much a part of American work culture as the water cooler. The big sin in exposure culture is not copying, but instead, failure to properly attribute authorship. And at the center of this exposure culture is the almighty search engine. If your site is easy to find on Google, you don't sue—you celebrate.

Once you think of the curve as being populated with creators who have different incentives, it's easy to extend that to their intellectual property interests as well. Disney and Metallica may be doing all they can to embrace and extend copyright, but there are plenty of other (maybe even more) artists and producers who see free peer-to-peer (“P2P”) distribution as low-cost marketing. Musicians can turn that into an audience for their live shows, indie filmmakers treat it as a viral resume, and academics treat free downloads of their papers as a way to increase their impact and audience.

Each of these perspectives changes how the creators feel about copyright. At the top of the curve, the studios, major labels, and publishers defend their copyright fiercely. In the middle, the domain of independent labels and academic presses, it's a gray area. Farther down the tail, more firmly in the noncommercial zone, an increasing number of content creators are choosing explicitly to give up some of their copyright protections. Since 2002, a nonprofit organization called Creative Commons has been issuing licenses of the same name to allow for a flexible use of certain copyrighted works for the sake of the greater value (for the content creators) of free distribution, remixing, and other peer-to-peer propagation of their ideas, interests, and fame. (Indeed, I've done that with my own blog, for all of the reasons above.)

In short, some creators care about copyright and some don't. Yet the law doesn't distinguish between them—copyright is automatically granted and protected unless explicitly waived. As a result, the power of “free” is obscured by fears over piracy and is often viewed with suspicion, not least because it evokes unfortunate echoes of both communism and hippie sloganeering.

Regardless, it's something we're starting to reconsider as the power of the “gift economy” becomes clear—in everything from the blogosphere to open source. In one part of my professional life (the 600,000-circulation magazine I edit), I'm near the head of the curve, and in another (my 5,000-reader blog) I'm in the tail. My decisions on intellectual property are different in each. Someday soon, I hope, marketplace and regulation will more accurately reflect this reality.

SELF-PUBLISHING WITHOUT SHAME

We think of books through a commercial lens, assuming that most authors want to write a best-seller and get rich. But the reality is that the vast majority of authors not only won't become best-sellers, but also aren't even trying to write a hugely popular book. Each year, nearly 200,000 books are published in English. Fewer than 20,000 will make it into the average book superstore. Most won't sell.

In 2004, 950,000 books out of the 1.2 million tracked by Nielsen BookScan sold fewer than ninety-nine copies. Another 200,000 sold fewer than 1,000 copies. Only 25,000 sold more than 5,000 copies. The average book in America sells about 500 copies. In other words, about 98 percent of books are noncommercial, whether they were intended that way or not.

The quest for mass-market acceptance requires compromise—a willingness to pick topics of broad rather than narrow interest, and to write in conversational rather than academic style. Most writers can't do that and many others won't. Instead, the vast majority of authors choose to follow their passions and assume they won't make money. Many want no more than to be read by some group that matters to them—from their peers to like-minded souls.

Such profitless publishing can be lucrative all the same. The book becomes not the product of value but the *advertisement* for the product of value—the authors themselves. Many such noncommercial books are best seen as marketing vehicles meant to enhance the academic reputation of their authors, market their consultancy, earn them speaking fees, or just leave their mark on the world. Seen that way, self-publishing is not a way to make money; it's a way to distribute your message.

To get a glimpse of that world, consider Lulu.com, which is a new breed of DIY publisher. For less than two hundred dollars, Lulu can not only turn your book into a paperback or hardcover and give it an ISBN number, but also ensure that it gets listed with online retailers. Once it's listed, the book will be available to an audience of millions and potentially side by side with *Harry Potter*, if the winds of the recommendation engines blow that way. With Lulu, the copies are printed in batches as small as a few dozen and the inventory is replenished as needed via print-on-demand. It's an extraordinary improvement over the scorned "vanity" publishing model of just a few years ago. As a result, thousands of authors are now choosing this route.

Here are the top five self-published books on Lulu, as I write:

1. *Raw Foods for Busy People: Simple and Machine-Free Recipes for Every Day*

2. *The Havanese* ("The quintessential handbook for Havanese dog owners, breeders and fanciers.")
3. *Investigating Biology—A Laboratory Manual for BIO 100*, 12th Edition
4. *Maximum SAT*
5. *How to Start a Wedding Planning Business*

All of them have sold between 5,000 and 50,000 copies, which is not bad. Eighty percent of the profits from these sales go directly to the authors, compared to 15 percent for standard publishers. So much for the notion that self-publishing is just for losers.

Still, most authors don't use such self-publishing services to make money, nor do they expect to hit it big. The vast majority of Lulu's other few thousand customers choose to self-publish because they know that what they're writing isn't likely to sell enough to make the search for a commercial publisher worthwhile. That doesn't mean they don't have a potential audience; it's just that it's a small one.

A few years ago, most of these authors wouldn't have been published at all—and that would have been enough to discourage many of them from writing a book in the first place. But today, the economics of publishing have fallen so low that nearly everyone can do it. That means people can write books for whatever reason they want, and they don't need to depend on some publisher deciding if the book is worth taking to market.

The effect of this is being felt throughout the industry, right up to the giant booksellers. In 2005, Barnes & Noble sold 20 percent more unique titles than it had in 2004, something its CEO, Steve Riggio, attributes to three forces: (1) the efficiencies of print-on-demand, which keeps more books in print; (2) the increase in the number of smaller and independent publishers; and (3) self-publishing.

"Over the next few years, the traditional definition of what a 'published book' is will have less meaning," he says. "Individuals will increasingly use the Internet as a first stage to publish their work, whether they are books, short stories, works in progress, or articles on their area of expertise. The best of this work will turn into physical books. I tend to be sanguine about the book industry's prospects be-

cause a whole new and efficient means of first-step publishing is emerging and rapidly becoming more sophisticated.”

One of the big differences between the head and the tail of producers is that the farther down you are in the tail, the more likely you are to have to keep your day job. And that’s okay. The distinction between “professional” producers and “amateurs” is blurring and may, in fact, ultimately become irrelevant. We make not just what we’re paid to make, but also what we *want* to make. And both can have value.

South Korea’s “citizen journalism” phenomenon, created in 2000 by OhmyNews, is another example. At OhmyNews, about fifty professional reporters and editors screen, edit, and complement news articles written by more than 40,000 amateurs, from elementary school students to professors. These volunteers submit between 150 and 200 articles a day, which account for more than two-thirds of OhmyNews’s content. For this, they receive a pittance: If the article goes to the front page, which only a small fraction do, the author gets around \$20. Why do they do it? “They are writing articles to change the world, not to earn money,” says Oh Yeon Ho, the site’s founder.

From filmmakers to bloggers, producers of all sorts that start in the Tail with few expectations of commercial success can afford to take chances. They’re willing to take more risks, because they have less to lose. There’s no need for permission, a business plan, or even capital. The tools of creativity are now cheap, and talent is more widely distributed than we know. Seen this way, the Long Tail promises to become the crucible of creativity, a place where ideas form and grow before evolving into commercial form.

CASE STUDY: LONELY ISLAND

One size of incentive doesn’t fit all. People create things for all sorts of reasons, ranging from expression to reputation. What makes this important is that there is increasingly frictionless mobility in the Long Tail. In a seamless digital marketplace, from iTunes to the Web itself, content that starts at the bottom can easily move to the top if it

strikes a chord. Understanding the diverse incentives that can motivate the creators of such content becomes essential in finding and encouraging it.

Speaking at a conference in mid-2005, Barry Diller, the media mogul chairman of IAC/InterActiveCorp, acknowledged that peer production is interesting, but he scoffed at the idea that it is a force capable of rivaling Hollywood. “People with talent won’t be displaced by 18 million people producing stuff they think will have appeal,” he confidently predicted.

What are the odds that he’s right? Well, if you define “people with talent” only as those who have a proven ability to make mass-market blockbusters, Diller may have a point. But there’s more to creativity than Hollywood hits, and people who can strike a chord can come from anywhere, via any path.

Take Akiva Schaffer, Jorma Taccone, and Andy Samberg. Until recently, they fit nicely into the category of people Diller’s talent-identification machine had efficiently filtered out.

After college, the three high school buddies relocated to Hollywood together. They moved into a big house with low rent on Olympic Boulevard and dubbed it the Lonely Island. Then they tried to figure out how to break into the entertainment industry as a comedy troupe.

It isn’t easy for an individual comic to make it in TV—even as a writer—but it’s even harder for a preassembled team. Sure enough, the threesome quickly ran up against all the usual barriers in their hunt for work in Hollywood. However, rather than subject themselves to endless rejection, the three took their act—now named after their home—online. Borrowing some video gear, the Lonely Island crew started producing short-form comedy videos and songs. Schaffer’s kid brother Micah—a tech consultant and Internet agitpropster—threw together their Web site, thelonelyisland.com, in 2001.

The Lonely Islanders started with white-boy rap music videos, presented with signature deadpan humor. One of the first videos was about things that are “ka-blamo!” (as in, “You kissed Shannen Doherty”) and things that aren’t (“I majored in pottery”). As is sometimes the case for such amusing ephemera, the video circulated widely on

the Internet. At one point, a Dutch DJ “mashed” it up (mixed it with other video footage), further boosting its popularity.

Soon more videos and fan mashups followed, something the group encouraged by releasing their videos under a Creative Commons license that freely permitted creative reuse. In just a few years, the Lonely Island was “Internet famous,” which is to say they were big with the demographic that has traded its TV time for online time, constantly surfing the contours of online subculture.

Capitalizing on their online celebrity, the Dudes—as they’re known to fans—scored better writing and performing gigs. Still, their main show continued to be online. The first episode of their “Internet prime time” series was called “The ‘Bu.” “Young, sexy people that live in Malibu call it The ‘Bu,” reads thelonelyisland.com, “because when you say the entire word, it takes time, and then you wouldn’t be young anymore.”

As the group’s cult following grew, word of their shorts got to *Saturday Night Live* star Tina Fey and the show’s creator, Lorne Michaels. In mid-2005, the threesome flew to Manhattan for auditions with the most famous team in comedy. In short order, all of the Dudes were hired.

In December 2005, the Lonely Island crew did another one of their white-boy rap sendups on *SNL*. Riffing on the *Chronicles of Narnia* film, the sketch was, as expected, twisted, wrong, and very, very funny. Now that the crew is on network TV, the skit went out as broadcast on a Saturday night, when it was watched by the usual (dwindling) audience, most of whom no doubt laughed and forgot about it.

But some people had recorded the show to their DVRs, and a few of them recognized a flash of brilliance in the *Narnia* skit. So they uploaded the video to the Internet. After it started to take off in the usual link frenzy, NBC heard the stampede and put the video on the official *SNL* site and even iTunes. Then, once again, the viral video effect kicked in—this time bigger than ever.

Jeff Jarvis, a media commentator, described the impact like this: “I haven’t heard anyone buzz about, recommend, or admit to watching *SNL* in, oh, a generation. But suddenly, I hear lots of buzz about the show. And it’s not because millions happened to start watching when

the show happened to actually be funny again. No, the buzz is born because folks started distributing the *Narnia* bit, which indeed is funny, on the Internet, and people are linking to it. NBC is learning the power of the network that no one owns.” And sure enough, links to the *SNL* site increased more than 200-fold in the two weeks after the video started circulating.

The Lonely Island tale has come full circle. Misfits rejected by the entertainment industry go online and get popular. Entertainment industry wakes up to this phenomenon in the hard-to-reach demographic of influential twenty-somethings and hires the misfits. The kids do the same thing on broadcast TV, but since that influential demographic doesn’t actually watch much TV, it isn’t until the skit *goes back online* (now amplified by the net-kids-make-it-big appeal) that the skit gets really popular. Thus *SNL*, previously scorned by the online generation, suddenly gets cool again by tapping into the authentic underground spirit blossoming online. Once upon a time, the show used to handpick its talent pool from obscure regional theaters and improv troupes. Now they also find it online.

So what’s the lesson in this story? Well, on one hand, the existing entertainment industry filters did recognize the appeal of the Lonely Island and found a way to tap it. In that sense, maybe the system works. Yet if three kids with a video camera doing goofy raps and putting them on their Web site isn’t “18 million people producing stuff they think will have appeal”—to borrow Diller’s scornful phrase—I really don’t know what is.

The truth is that the next generation of talent will probably come from the 18 million people doing their own thing—and these are the people who are most likely to save Hollywood and the rest of the entertainment industry from grinding formula. Maybe Diller is right. Maybe there are only a small number of people who can write *Friends*. But just think about how many people can produce quirkier fare, like the *Narnia* sketch, content that can resonate with an audience that has grown up online—the place where niches, not networks, rule. Think about how many of those potential talents now have a chance to find a real audience, thanks to the democratized distribution of the Internet.

It may still require the full might of the Hollywood machine to make a multiseason drama with high production qualities. But over that same time hundreds of grassroots videos can collectively capture a similar size audience. That comparison would seem like apples and oranges—lasting commercial brands versus transient amateur amusements—were it not for the fact that the two compete for the time of a generation of Web-savvy viewers. If they're watching one kind of video, they're not watching the other.

What Diller neglects to consider is that today there seems to be less demand for blockbusters than there is for focused or targeted content that *isn't* for everybody. As the audience continues to move away from Top 40 music and blockbusters, the demand is spreading to vast numbers of smaller artists who speak more authentically to their audience. So what if 99 percent of blogs will never attract an audience of more than a few dozen? The fraction of a percent that *do* emerge with broader reach still number in the thousands. And collectively, that 1 percent can draw as much traffic as many mainstream media. The typical "viral video" sensation is seen by several million people, something that can only be said for the most popular TV shows.

As with authors who self-publish their books via Lulu, the products themselves aren't usually making much if any money, but that's not the point. The point is simply that the product exists and it's taking audience share. It isn't a creation of the traditional commercial industry, but it competes with it. Today, the number of people who produce content is far more than the usual talent finders of the media can process—the wave of grassroots creativity would overwhelm the script-readers and tape-listeners of any studio and label. Because the tools of production have entirely democratized, the population of producers is expanding exponentially, and now there's little stopping those with the will and skill to create from doing just that.

THE ARCHITECTURE OF PARTICIPATION

We've seen parts of this story before. In the late 1970s and early 1980s, the combination of the electric guitar, the arrival of cheap

multitrack recorders, and the fine example set by the Sex Pistols gave license to a generation of kids with no musical training, obvious talent, or permission from anyone to start bands and record music. When punk rock exploded onto the scene, it was a shocking epiphany for a generation of kids in the mosh pit. Watching someone your age play three chords badly, while jumping around on stage, one couldn't but think: "I could do *that*."

For a while, the assumption was that to be a musician, the right way of learning was to copy the masters. So you should start by playing covers, reading music, and maybe going to music school. This was the notion of paying your dues: Do the circuit, and play the standards, because that's what people want (no one wants to hear your crappy original compositions). Do it *right*.

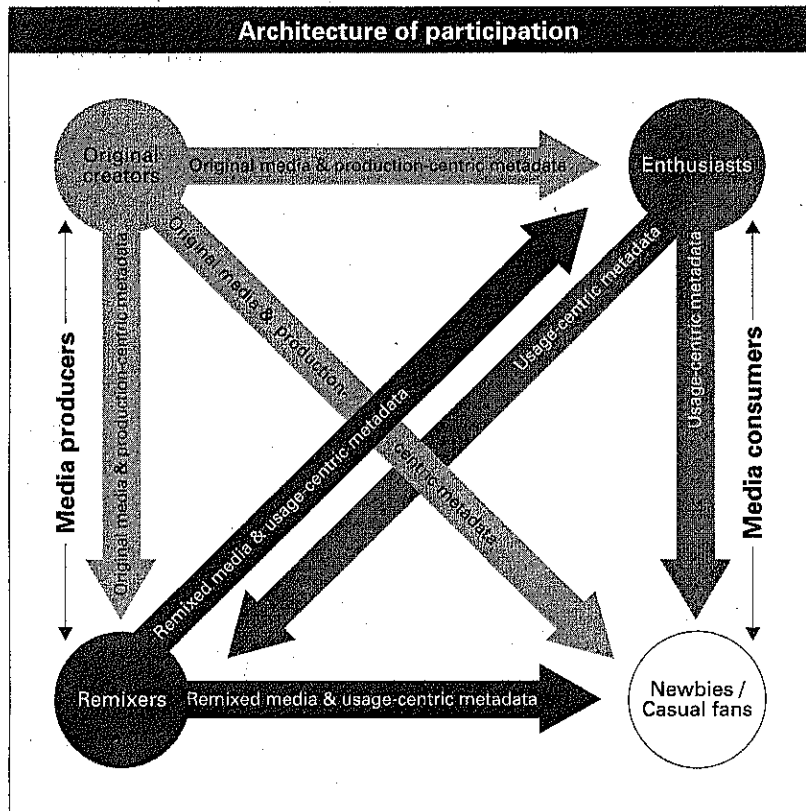
But punk rock changed the game. Punk rock said: "Okay, you have your guitar, but you *don't* have to do it right. You can do it wrong! It doesn't matter one bit if you're a skilled musician; it just matters if you have something to say."

Through punk rock, we saw a premium on fresh voices, new sounds, vigor, and an antiestablishment sentiment that could have only come from outside the system. It was inspirational to see people out there with no more talent than you, having fun, being admired, doing something novel. To put it in economic terms, punk rock lowered the barriers of entry to creation.

The traditional line between producers and consumers has blurred. Consumers are also producers. Some create from scratch; others modify the works of others, literally or figuratively remixing it. In the blog world, we talk about "the former audience"—readers who have shifted from passive consumers to active producers, commenting and blogging right back at the mainstream media. Others contribute to the process nothing more than their Internet-amplified word of mouth, doing what was once the work of radio DJs, music magazine reviewers, and marketers.

The result is starting to look like what Tim O'Reilly, a book publisher and seer of the DIY age, calls the "The New Architecture of Participation."

A team at the University of California, Berkeley, illustrated this with a new map of creation, as follows.



As this figure shows, a once-monolithic industry structure where professionals *produced* and amateurs *consumed* is now a two-way marketplace, where anyone can be in any camp at any time. This is just a hint of the sort of profound change that the democratized tools of production and distribution can foster.

6

THE NEW MARKETS

HOW TO CREATE AN AGGREGATOR THAT
CAN STRETCH FROM HEAD TO TAIL

In 1982, a bookseller named Richard Weatherford realized that the then-new personal computer could revolutionize the used-book business. There are thousands of used-book stores around the country, all with different inventories. Virtually any book you might want is out there somewhere, but good luck in finding it. Weatherford saw this as primarily an information problem, exactly the sort of thing computers are good at solving, and he wrote a business plan for a company that would build an online database for antiquarian booksellers. He called it Interloc, short for interlocutor, a fancy way of saying “go-between.”

Weatherford was a few decades ahead of his time, and he failed to get funding. But in 1991, he was hired by Faxon, a book and magazine service firm, to salvage BookQuest, which had attempted to do the same thing. It didn’t work—this was still about a decade too early—but the funding at least was starting to become available. With \$50,000 from other booksellers, Weatherford launched Interloc in 1993, before the Web. It was a closed network to enable booksellers to search other merchants’ inventory to find books for their own customers. It created a data standard (which is still in use today) and soft-