11

I'M FEELING LUCKY

Google, Napster and the Rebirth

Whhen Larry and Sergey first met, they didn't like each other much.

In the summer of 1995, Larry Page was considering a transfer to Stanford University's graduate program in computer science. Sergey Brin was already two years into the program, and he had signed up to be a tour guide of sorts to potential students. One summer day, he showed Page and a group of other potential Stanford students around the Bay Area.

"I thought he was pretty obnoxious," Page said later of his guide. "He had really strong opinions about things and I guess I did, too."

"We both found each other obnoxious," Brin agrees. They might have stepped on each other's toes a bit, but at the same time there was a degree of frisson to the encounter. "We spent a lot of time talking to each other," Brin would recall, "so there was something there. We had a kind of bantering thing going."[1]

On the surface, it might not have seemed like Page and Brin had anything in common. Page was a midwesterner, born in East Lansing, Michigan, on March 26, 1973. Brin was born in Moscow, in the Iron Curtain–era USSR, on August 21, 1973, and was brought to the United States when he was six years old. Page was reserved, quiet, contemplative. Brin was outgoing, gregarious, loud. Page was a deep thinker, a visionary. Brin, a problem solver, an engineer's engineer.

But the two had more in common than anyone knew that first day. They both

came from academic families. Page's father was a pioneering computer science professor at Michigan State University, where his mother was also a computer programming instructor. Brin's father was a mathematics professor at the University of Maryland and his mother a researcher at NASA's Goddard Space Flight Center. Larry and Sergey both grew up to respect research, academic study, mathematics and especially computers. And they both had inquisitive minds, believing in the power of knowledge to overcome any obstacle, intellectual or practical. Each had been inculcated into this spirit of intellectual fearlessness at a young age.

"You can't understand Google," early Google employee Marissa Mayer (and later, Yahoo CEO) has insisted, "unless you know that both Larry and Sergey were Montessori kids. It's really ingrained in their personalities. To ask their own questions, do their own things. Do something because it makes sense, not because some authority figure told you. In a Montessori school, you go paint because you have something to express or you just want to do it that afternoon, not because the teacher said so. This is baked into how Larry and Sergey approach problems. They're always asking, why should it be like that? It's the way their brains were programmed early on."[2]

For Larry and Sergey, their intellectual fearlessness overlapped in such a way that their conflicting personalities actually ended up complementing each other. When Page came to Stanford for the 1995–96 academic year, he and Brin became close. Friends took to calling the duo LarryandSergey, and the pair would end up debating endlessly on topics ranging from philosophy to computing to films, two equally matched polymaths thrilling to the intellectual joust. Brin's hobby project was creating a software program that could provide movie recommendations based on the tastes and viewing habits of other people who had seen similar films (not unlike what Netflix later perfected). Page's dream obsession was creating a system of networked, autonomous cars to ferry people around.

Even though they were the same age, Brin was academically two years ahead of Page because he had completed his undergraduate computer science degree at age nineteen and aced all of Stanford's required doctoral program exams on the first try.[3] But despite this head start, and despite being the recipient of a National Science Foundation fellowship that allowed him to do basically anything he wanted, Brin had stalled out in his quest to nail down a dissertation topic. Of course, the newly arrived Page also needed to decide on *his* dissertation, and so fate pushed the pair even closer together. In January 1996, LarryandSergey ended up working in the same office, number 360, in the just-completed William

Gates Computer Science Building on Stanford's campus. The building was of course named after the founder of Microsoft, who had donated $6 million to the construction. All his career, Bill Gates repeatedly predicted that one day, some student somewhere would found a company that would challenge Microsoft for dominance of the tech industry. His prediction turned out to be right, and from a building with his name on it.

■

PAGE WAS STRUCK by a fundamental truth about the web that is glaringly obvious when you state it out loud: it is built on links. One page linking to another; one idea linking to another. As of yet, no one had bothered to analyze the structure of the link ecosystem in a comprehensive way. For example, it was possible to know that webpage A linked to webpage B because you could see it—you could follow the link. But what about the reverse? What pages had linked webpage A? There was no way to know. You couldn't follow a link stream backward, only forward. Page wondered: if you analyzed all of the back links, if you mapped out the link structure of the entire web, what sort of insight might that data give you?

Page's intuition was that this might be more than just an interesting theoretical question. As he mulled over the idea with Brin, their shared upbringing as the children of academics kicked in. LarryandSergey knew the power of the academic citation. Their parents had published academic papers. They, themselves, intended to publish academic papers in order to earn their Ph.D.'s. And they knew that any academic paper worth its salt built its argument by citing other academic papers and studies. In the world of academia, those citations, the accumulated number of "votes" from paper to paper, served, over the years, to accrue value to given ideas—to essentially rank them based on the number of citations. The most cited papers were understood to be the most authoritative. "It turns out, people who win the Nobel Prize have citations from 10,000 different papers," Page would say later.[4]

Well, what was a web link but a digital citation? If you analyzed the links, analyzed the citations, you might be able to make inferences about the relative value of a given web page, and possibly even determine which webpage was more authoritative by analyzing the back links in the same way that counting the citations told you which academic paper was the definitive one. Larry Page wanted to map out the value of the web's connections by going backward through the link chain. Page went to his academic advisor, Terry Winograd, and asked for the money and machines that would allow him to map the web's links. He dubbed the project BackRub. When asked how much of the web he intended

to map, he replied: "the whole web."[5]

So, in March of 1996, Larry Page launched BackRub by sending search bots, known as "spiders," out into the web to find *all the links*. He started with a single page—the Stanford computer science department homepage—and then fanned out, following link after link, cataloging them all, and then ranking web pages based on these link citations. It was the mathematical complexity of this ranking—the complicated problem of determining which page was more valuable based on a combination of accumulated links as well as the authority passed through from pages that linked to other pages—that drew Sergey Brin to join the project. Larry and Sergey called their combined citation-ranking system PageRank, either as an ode to Page himself or as an obvious descriptor of what the system was intended to do.

"The idea behind PageRank was that you can estimate the importance of a web page by the web pages that link to it," Brin says. "We actually developed a lot of math to solve that problem. Important pages tended to link to important pages. We convert the entire web into a big equation with several hundred million variables which are the PageRanks of all the web pages, and billions of terms, which are all the links."

"It's all recursive," Page said. "In a way, how good you are is determined by who links to you and who you link to determines how good you are. It's all a big circle."[6]

LarryandSergey suddenly had a project that would make for a pretty interesting dissertation. And as soon as the pair looked at their results, they realized their intuition was dead on: the citation analogy worked. If you wanted to find the most authoritative webpage about a topic such as, say, windsurfing, BackRub/PageRank could tell you. It would know based on the accumulated links, of course, but also from the authority passed on from other authoritative sites. Thanks to Brin's math (largely linear algebra and something about the eigenvector of a weighted link matrix, for those who know what that means), citations from obviously important websites were more valuable than others. A link from some unknown person's personal webpage might be valuable, but a link from a professional windsurfer would be judged to be even *more* valuable—and a link from, say, Yahoo's homepage would be even more valuable still.

It was at this point that the really interesting application for this little math project became obvious. "It was pretty clear to me and the rest of the group," Page said later, "that if you have a way of ranking things based not just on the page itself but based on what the world thought of that page, that would be a really valuable thing for search."

■

IT TURNED OUT THAT the reason search engines had never worked very well prior to PageRank was not that they were broken, but because they were missing the key innovation that Brin and Page had stumbled upon: relevancy. If, in 1997, you did a search for "automobile company" on even the best search engine at the time (AltaVista) you'd find yourself disappointed because the websites of Ford, General Motors or Toyota would probably not show up. It's not that AltaVista couldn't find those sites. It most certainly had! Ford.com or GM.com or Toyota.com were most likely in the list of tens of thousands of results that AltaVista had found. It was just that AltaVista had no way of surfacing those most relevant results to the top. So they were on page 3 of the search results. Or page 300.

PageRank solved this problem. PageRank knew which sites were the most authoritative automotive sites already, and so when you combined its algorithmic prowess with the traditional tricks of information retrieval that all the search engines were already using, suddenly it all just worked. Indeed, as Page and Brin combined BackRub and PageRank with traditional search methods like analyzing on-page text, webpage titles or metatags and, especially, parsing the so-called anchor text of a link (someone who makes a link out of the words "flower shop" and then points it to a given website is really trying to tell you something), they realized PageRank was incredibly powerful. Page and Brin discovered that their algorithm was indeed recursive, meaning that the more data they fed it, the more webpages it analyzed, the better it got. By tweaking the math even more, LarryandSergey's search tool could reliably find people, locate the most obscure fact or data, and even *answer questions*. PageRank wasn't finding new things. It was merely finding things in a better way. The earlier search engines were already answering every query correctly. But it was finding the needle in the haystack and putting it at the top of the list that PageRank did better.

"It wasn't that they [Page and Brin] sat down and said, 'Let's build the next great search engine,' " said Rajeev Motwani, who was Brin's academic advisor. "They were trying to solve interesting problems and stumbled upon some neat ideas."[7]

■

IT WAS A GOOD THING Page and Brin had not set out to build the next great search engine, because, at the time, no one was really clamoring for one. In the late nineties, when Page and Brin began refashioning BackRub/PageRank into a

search engine, there was a universe of major search players: Yahoo, Excite, Lycos, AltaVista, AskJeeves, MSN, and on and on. In a time when Yahoo had a $100 billion market cap, who needed another entrant into an already-crowded space, no matter how superior it was? Fortunately, Page and Brin were not business-focused at that time. They were academics, more interested in defending a dissertation and publishing a paper on their research than starting a company around their idea.

So, they produced that paper: "The Anatomy of a Large-Scale Hypertextual Web Search Engine," which was presented at a conference in Australia in May of 1998. But if Page and Brin initially stayed true to their chosen academic paths, that did not mean they were blind to the financial possibilities inherent in their work. How could they have been? Students studying computer science in the heart of Silicon Valley couldn't help but notice what was going on all around them. "It was a hard time to stay in grad school," remembered Tamara Munzner, one of the students sharing room 360 of the Gates Building with Page and Brin. "Every time you went to a party, you had multiple job offers and they were all real. I had to redecide every term not to leave."[8]

The obvious move was to license PageRank to one of the existing players, and indeed, this is what Page and Brin attempted to do. They met with everyone from the Yahoo founders Jerry Yang and David Filo, to another search pioneer, Infoseek's Steve Kirsch. No one was interested. The closest they came to making a deal was when Page wrote up an extensive proposal to Excite's leadership, suggesting they replace Excite's existing algorithms with his. Doing so, he calculated, would generate an additional $47 million in revenue for the search engine. "With my help," Page wrote in his proposal, "this technology will give Excite a substantial advantage and will propel it to a market leadership position."[9] All he asked for in exchange was a seemingly reasonable $1.6 million in cash and Excite stock—a nice little payday—and then he and Brin would return to finishing their doctorate work. Excite countered with $750,000, which Page and Brin rejected.

The incumbent search players' failure to scoop up the PageRank technology has become infamous in business lore as one of the great missed opportunities of all time. Larry Page has, on a few occasions, suggested that the search companies were simply myopic. "They were becoming portals. We probably would have licensed it if someone gave us the money. . . . [But] they were not interested in search," Page has said. "They did have horoscopes, though."[10] But Excite CEO George Bell has a slightly different recollection: "The thing that Larry insisted on, that we all do recall, is that Larry said, 'If we come to work for

Excite, you need to rip out all the Excite technology and replace it with [our] search.' And, ultimately, that's—in my recollection—where the deal fell apart."[11] This was Page and Brin's intellectual fearlessness demonstrating itself for the first time in a competitive setting. The pair believed—*knew*—that they had a superior way of doing things, and so they thought nothing of going to an established search company and telling them their existing product sucked. This brashness had the effect of insulting Excite. Excite was a company founded by brilliant Stanford computer scientists, after all. "We had hundreds of engineers at that point," Bell points out. Why should the company furlough their engineers just because two other engineers had come along with claims to be *more* brilliant? Bell claims that there was no way he could justify upsetting his existing talent, especially when some of them were founders of the company. "Ultimately I couldn't stomach the cultural risk that Larry insisted on," Bell says.

But if Page and Brin were confident almost to the point of being arrogant, they certainly had plenty of data to back them up. In order to fine-tune their algorithm, the pair had needed plenty of real-world feedback. Starting in 1997, they had made the search engine available, first on Stanford's internal network, and then to the general public. Through nothing but word of mouth, the service grew increasingly popular, serving more than 10,000 queries a day by late 1998.[12] Page and Brin monitored the server logs and made tweaks to their system based on the data this provided. They named the service Google, a play on the word "googol," which is a 1 followed by 100 zeros. The idea was to suggest they were capturing the whole web, everything in existence. "The name reflected the scale of what we were doing," Brin said later.[13] Googol.com was not available, so Google.com became the URL of the public service.

The popularity of the service, combined with the vast computing resources eaten up by the spidering and indexing, meant that the Google project was rapidly outgrowing the scope of a simple research project. Even when it was housed on a single machine in a Stanford dorm room, Google was hogging large amounts of the university's bandwidth. Stanford was, as ever, incredibly accommodating to an idea born within its walls, but the institution's generosity had a practical and obvious ceiling.

It was clear that if they wanted the Google experiment to continue, Page and Brin would need more resources. More computers, more bandwidth, more people to work on the algorithm—this all meant more money than a research budget, even a generous one, could provide. So the pair turned to another Stanford faculty advisor, David Cheriton. Cheriton introduced the pair to Andy

Bechtolsheim, a successful entrepreneur who had founded Sun Microsystems while also a Ph.D. student at Stanford. One morning in late 1998, Page and Brin met Bechtolsheim at Cheriton's home. Bechtolsheim made out a check on the spot for $100,000 in the name of Google Inc. The check sat in Page's dorm room desk for a number of weeks before Google Inc. was formally incorporated on September 7, 1998. Page and Brin would raise an additional $1 million when David Cheriton kicked in some money, as well as a few others, including former Netscape executive Ram Shriram and Jeff Bezos of Amazon.

Page and Brin were now entrepreneurs, if perhaps still a little reluctantly. But they were not entrepreneurs in the mold of so many others in the dot-com era. Rather than blowing Google's funds on lavish launch parties or marketing campaigns, they stayed grad students at heart, and instead invested all the money they had raised in continuing their project efficiently. Instead of building out their system by buying software from Microsoft, they used the free Linux operating system. Instead of splurging $800,000 on setups from IBM or Oracle, they spent a mere $250,000 to cobble together a rack of eighty-eight computers to meet their number-crunching needs. At Stanford they had begged, borrowed, and almost quite literally stolen the computers they needed to keep Google running. Now, they simply switched to buying computers off the shelf from Fry's, the famous Silicon Valley electronics store, and fashioned them into a strung-together system of their own design. Part of this was simple frugality, a habit that would serve them well when the dot-com bubble burst a few short years later. But a lot of it was Page and Brin's ingrained Montessori philosophy: they never met a problem they couldn't solve through smart engineering.

Google didn't take pages from the established Silicon Valley playbook because, in a way, they had never bought into it. They didn't try to Get Big Fast. Instead, Page and Brin were almost manically focused on endlessly iterating and improving upon their Big Idea, making sure it was the most comprehensive, reliable and—most important—speedy search engine in the world. Nothing Google did in its first years distracted the company from improving on its core product. This confidence that they could do everything better proved, in the coming years, to be something of Google's secret sauce. Not only did Google's search engine continue to be superior to any rival in existence, it slowly but surely widened the gap between its version of search and the competition. And their frugality paid off in efficiency. Some observers estimated that "for every dollar spent, Google had three times more computing power than its competitors."[14]

Frugality and efficiency were not just virtues, they were also philosophical

and aesthetic differentiators. Google's home page was simply the Google logo, a text field to enter a search query, a search button to execute that query and a button that said I'M FEELING LUCKY, which automatically took you to the first result returned. If you went to the search results page, you only got a list of links. And that was it. No ads, no banners, no weather, no stock quotes, no horoscopes. All the rest of the page was just copious white space. In an age of portals where every other search site was a sea of distractions meant to keep you from, you know, getting to the page you were looking for, Google stood out from the crowd with its single-minded purpose and simplicity. By keeping the pages to almost exclusively text, Page and Brin could ensure they loaded quicker than the search pages of their competitors, and expensive processing power wasn't wasted loading graphics.

This all paid dividends many times over in Google's steady growth. By 1999, usage of the search engine was increasing by as much as 50% a month.[15] From 100,000 searches a day at the beginning of that year, Google searches grew to an average of 7 million per day by the end of it.[16] Overall traffic to the Google homepage was peanuts compared to the numbers a site like Yahoo was pulling down, but in the case of Google, its users came via word of mouth alone. Not a dime was spent on marketing or promotion. Rave reviews from the media continued to turn people on to the service. The *New Yorker* said Google was "the default search engine of the digital in-crowd."[17] *Time* Digital said: "Google is to its competitors as a laser is to a blunt stick."[18] Ordinary users simply told one another about how great and useful Google was. More often than not, users would become Google converts for life.

An early article on Google in *Fortune* from November 1999 summed up a new user's experience. Describing the site as "inscrutable magic," journalist David Kirkpatrick offered this anecdote. On the occasion of the 1999 American League playoffs, Kirkpatrick typed "New York Yankees 1999 playoffs" into both Google and Alta Vista. "The first listing at Google took me directly to data about that night's game," Kirpatrick wrote. "The first two at Alta Vista linked to info about the *1998* World Series." Only by clicking the third Alta Vista link, and then visiting an additional link, did he find the information he was originally searching for. Kirkpatrick's conclusion: "Google really works."[19]

In that same article, Sergey Brin was quoted as boasting, "We're building a way to search human knowledge." If Google was meant to organize all the information in the world, it would need resources on an industrial scale. That same brashness continued to manifest itself when Google needed to raise yet more money.

Despite the glut of search companies already on the market, Google had gotten the attention of venture capitalists, and they were ready to invest in these refugees from academia. But, confident as ever, Page and Brin gave off the impression that they didn't need anyone's help or money. In meetings with potential backers, the pair refused to divulge even basic details about how their service was operating. Their stonewalling even led one prominent VC to storm out of their office in anger. "Larry and Sergey didn't have the language to say things nicely," recalled Salar Kamangar, an early employee who bore witness to Google's general evasiveness during the fundraising process. "They'd be kind of blunt and say, 'We can't tell you.' And the VCs would get very frustrated."[20] The truth was, Page and Brin did not want to take money from just any old VC. They only wanted the best: Kleiner Perkins and Sequoia Capital. The pair proposed that each firm, the blue chips of Silicon Valley venture, take a coequal stake in Google. There was usually one "lead" investor in a round of startup financing, and both KP and Sequoia had enough clout on their own that they had never before deigned to share the spotlight with another firm.

Page and Brin wanted the firms to split the round because that would allow them, as the founders, to maintain a majority share in the company, and thereby retain control of their own destiny. They even had the temerity to issue an ultimatum: each firm would invest $12.5 million in Google, for a total of $25 million, take it or leave it. On June 7, 1999, the VCs took the deal, and Kleiner's John Doerr and Sequoia's Mike Moritz joined Google's board of directors. The only concession the money men had been able to wring out of Page and Brin was a promise to hire someone experienced to take over as CEO of the company at some point in the near future.

This huge round of financing not only put Google firmly on the technology world's map, it went a long way toward ensuring the company's long-term survival. This war chest of money, coming just before the dot-com bubble burst, combined with Larry and Sergey's frugal ways, meant that Google would survive the coming nuclear winter. Had Google waited a further year to raise money, it might not have been able to. And by virtue of being flush with cash when the rest of Silicon Valley was seemingly going belly-up, Google was able to have its pick of talent when the dot-com layoffs began.

Just as it had been frugal when others were profligate, Google also bucked prevailing dot-com habits when it came to hiring. The company put off drafting an army of sales and marketing people until much later. Instead, in 1999 and 2000, Google staffed up with—what else?—brainiacs. Larry and Sergey hired software engineers, hardware engineers, network engineers, mathematicians,

even neurosurgeons. Just as with every other facet of their company, Page and Brin wanted only the very best. They wanted Ph.D.'s and scientists. Google would become notorious for the rigorous way it interviewed and screened potential hires—and for its exacting selectiveness. For many years, every new employee was personally vetted by Brin and Page themselves, who expected candidates to measure up to their own intellectual standard. "We just hired people like us," Page said.[21]

Google was able to attract talent because it was nothing short of beloved in Silicon Valley. Here was an Internet company that had solved a universally recognized problem through smart thinking alone. This created a reputational halo that was only enhanced by Larry and Sergey's increasingly bold and public enunciation of Google's mission, which was eventually formalized as an attempt "to organize the world's information and make it universally accessible and useful." While so many dot-com companies claimed to be changing the world by offering dog food online, here was a company that truly seemed revolutionary in the most expansive sense of that word. "Ultimately I view Google as a way to augment your brain with the knowledge of the world," Sergey Brin said.[22] It helped that Google positioned itself as the anti–dot-com startup. Glitz, hype and excess were out; frugality, hard work and earnestness were in. And when Google came up with its famous motto (Don't Be Evil) everyone in technology read between the lines and believed that Google was staking a claim to be the anti-Microsoft.

Google did pick up a few habits from its dot-com brethren, but in typical Larry and Sergey fashion, it did so with a twist. By the time Google moved to its first truly professional digs—an office park in Mountain View that would be dubbed the "Googleplex"—a system of perks for Google's workers were put in place, but they were instituted with an eye toward productivity. The food in the cafeteria was always free, with an in-house gourmet chef; private bus lines picked up workers from around the Valley to shuttle them to work; masseuses roamed the hallways; there were free fitness classes and gyms; and on and on. But every one of these perks was self-consciously provided as a way to keep workers motivated and productive. The free cafeteria meant that Google employees didn't have to leave the office in the middle of the day and could get back to work with ease. In the bathroom stalls were quizzes and coding tips to help people stay sharp. The shuttle buses had WiFi on them, so employees could be productive on the way to and from the Googleplex. Healthy, clear-headed workers could do better coding, or so the thinking went.

All of this combined to make Google *the* technology company to join right

as the dot-com bubble burst. If you got hired at Google, it elicited envy from your peers not only because they felt you were doing the most interesting work in technology, but because it meant you were among the best and the brightest. Anyone could get hired at a dot-com toward the end of the decade. But not everyone—even the smartest of the smart—could make the cut at Google. And when the bubble burst and it was seemingly the only company still hiring, the dream of the nineties was alive in the Googleplex.

■

GOOGLE HAD ALWAYS BEEN OBSESSED with its logs, the reams of data its users provided by their billions of searches. Google's engineers used this data to improve the algorithms, but as the company was committed to "organizing the world's information," it also had a fascination with how search behavior revealed the world's obsessions in real time. Eventually, products like Google Trends and Google Zeitgeist would allow us all to peek inside the planet's collective unconscious, surfacing perennial obsessions like "sex" or "porn" but also faddish searches like "Paris Hilton" or "Justin Bieber." In the year 2000, the hot search term was "MP3." This was because, across the country, a teenager just barely into his first year of college had dreamed up a program that would break the Internet wide open just as definitively as Google's algorithms were doing.

Shawn Fanning was a member of the first true web generation, born November 22, 1980, in the working-class Boston suburb of Brockton, Massachusetts. Earlier than most people his age, Shawn became a heavy user of online chat, especially Internet Relay Chat. It was on IRC that Shawn Fanning fell deeply in with the teenage hacker crowd.

Sometime in 1997 or 1998, Shawn was invited to join the private IRC channel called w00w00, which was the main online meeting place for a hacking collective of the same name. Members of w00w00 would go on to have a hand in the formation of dozens of technology companies ranging from WhatsApp to Arbor Networks, but at the time, they were just a bunch of kids trading hacks.[23] Under the pseudonymous login handle "napster," Fanning traded programs and coding advice, trying to impress the other hackers with exploits and programs he scratched together himself.

In the fall of 1998, Shawn enrolled at Boston's Northeastern University and saw that his new roommates and fellow students were obsessed with finding and trading music files known as MP3s. But finding these files was a complicated process of searching FTP (File Transfer Protocol) sites, Usenet newsgroups and

other online repositories. There was also no real way for users to exchange these files easily among themselves. So, late in 1998, Shawn Fanning announced to his fellow hackers on w00w00 that he was working on a program that would make finding and exchanging MP3 files a breeze.

■

FROM THE EARLIEST DAYS, people had dreamed of turning the web into a medium for music. As early as 1993, two students at UC Santa Cruz launched a website called the Internet Underground Music Archive so that artists and musicians could upload and distribute digitized recordings for others to download and listen to. This proved popular, but largely unwieldy for most users, since the size of the music files was too large for the dial-up Internet connections of the day; downloading a single song could take half a day to complete. This changed in the mid-nineties, when a new type of music file was introduced. ISO-MPEG Audio Layer-3, or MP3, was developed at the Fraunhofer Society for the Advancement of Applied Research in Germany and used audio and file compression to create music files that were much smaller in size, but without sacrificing too much in the way of sound quality.

It turns out that the human auditory system is not an instrument that scoops up all the frequencies in a given environment, like a microphone does. What we "hear" is not an accurate representation of reality, but only those sounds that the brain, over the course of millenia of evolution, has determined to be the "most important" sounds. By stripping out the unnecessary (because they were unheard) noises in a sound file, music files could be made much smaller. Most music was easily compressed and a listener was none the wiser. "That's an undergraduate project," says Karlheinz Brandenburg, the Fraunhofer researcher who is called the "father" of the MP3.[24] But the human voice was far trickier. It turned out that the key to mastering the nuances of human singing was an obscure a cappella recording of a minor hit from the 1980s, Suzanne Vega's "Tom's Diner." Brandenburg successfully tweaked the MP3's compression algorithm by listening to "I am sitting *In the morning* At the diner / On the corner . . ." over and over again, maybe 10,000 times, before he got it right. "To get it to the level that it's really perfect, or near-perfect, for everything," says Brandenburg, "*that* was work."[25]

The resulting files were small enough to be useful in a low-bandwidth era, but MP3 technology further benefited from another technological leap that was occurring at just the same time: computer storage was exploding. The web had been born in an era when the average computer hard drive was still measured in

megabytes. The first gigabyte hard drives only became commercially available in the mid-1990s,[26] and by 1999, CNN was trumpeting the arrival of 5GB, even 10GB, hard drives.[27] That amount of storage might seem woefully small for even a smartphone these days, but in the late 1990s, it was a massive amount, more than enough to store not just numerous songs, but entire albums worth of MP3s.

The media was there, the storage was there, and just as serendipitously, the ability to play this media arrived on the scene as well. In 1997, a nineteen-year-old college dropout named Justin Frankel released a software program called Winamp, which allowed users to easily organize and play MP3s on computers. Winamp was downloaded by more than 25 million eager MP3 devotees, and Nullsoft, Winamp's parent company (which Frankel had formed with the Internet Underground Music Archive's Rob Lord), was sold to AOL in 1999 for around $100 million.[28]

In a way, Shawn Fanning was trying to solve the final piece of this puzzle: a search engine for MP3s. But since most MP3s were sitting on individual users' computers, he needed to find a way to search other people's hard drives, not public webpages. That way, if you wanted to find a particular song, you could simply figure out who had it on their computer and get it directly from them. You would share the songs on your hard drive as well, thereby keeping the karmic cycle going. Fanning's MP3 search program would be networking in its purest form; it would be a literal peer-to-peer exchange.

"It felt like this way of sharing media between people could be used for sharing anything," Fanning would say later. "It also felt like this whole model for sharing media was superior to, like, going and buying an album. . . . Basically to have access to the entire universe of recorded music. . . . In every way it seemed like a better system."[29]

■

IN A FEW SHORT WEEKS, Fanning coded up a rough version of a program, which he named after his nom-de-hacker, Napster. As was the custom, he turned to the other hackers in w00w00 for tips and advice. Among those other w00w00 users who began chipping in to contribute to the program was a slightly older, slightly more sophisticated coder named Jordan Ritter (w00w00 handle: "nocarrier") and a less technically savvy but more ambitious w00w00 hanger-on named Sean Parker (nickname: "nob"). Ritter would eventually take over the sophisticated back end of the Napster system, developing the complicated server connections, search algorithms and networking details that would allow users to search each

other's computers and download MP3s directly among themselves. And as for Parker's contribution? Well, Sean Parker wanted to turn Napster into a business.

Despite the populist image it cultivated later, Napster was conceived of as a business from day one. The Napster phenomenon was covered in the press as some sort of grass-roots movement that bubbled up out of nowhere, largely because that was the image Napster, the company, later fed to the press. But the truth is that long before Napster was a multimillion-user phenomenon—before Napster even had users numbering in the tens of thousands—the idea was to turn Napster into a billion-dollar company. This inclination was partially a result of the time Napster was born into; 1998 into 1999, when Napster was being developed, was the height of the dot-com mania. But it was also because the brilliance of the Napster idea was immediately obvious to everyone involved: it was an entirely new way to distribute media. Imagine being able to search and instantly find any song in existence. And then imagine the instant gratification of being able to download those songs and play them right away. Oh, and by the way: all those songs were completely, 100% free of charge, because you were getting them, not from a record store, but from some other, unknown Internet user.

Napster was seeking to raise money from investors not long after it left the friendly confines of the w00w00 IRC channel. This was thanks to the precocious Parker, who took it upon himself to raise money for the project, running through a chain of connections that eventually landed Napster a $250,000 investment from a California angel investor on Labor Day 1999. By the fall of 1999, Shawn Fanning, Sean Parker, Jordan Ritter and another w00w00 regular, Ali Aydar (IRC handle: "mars") were out in California turning Napster into a real startup.

Napster was like a supernova that exploded across the tech, media and cultural landscape just as the dot-com bubble burst in the year 2000. The grand-slam idea that everyone saw in Napster's technology proved itself out spectacularly. By the spring of 2000, less than a year after launching, Napster had more than 10 million users.[30] By the end of 2000, Napster could claim more users than even mighty AOL: around 40 million. And instead of taking more than a decade and billions of dollars to do so, Napster had attracted that many users on the backs of half a dozen barely postpubescent hackers and about $400,000 worth of hardware.[31]

Napster owed its success to all those college kids with their gigabyte hard drives and broadband dorm room Internet connections. By the spring college semester of 2000, an estimated 73% of college students were using Napster regularly.[32] On some campuses, Napster was consuming nearly 85% of available

bandwidth.[33] When various institutions began enforcing Napster bans, students nearly rioted. For a long time, Napster was in *The Guinness Book of World Records* as the fastest-growing service of all time.[34] At points early on in its development, Napster's user numbers were growing 35% *a day*.[35]

But if Napster was a supernova, it was also *the* star-crossed startup of the Internet Era. Even nearly twenty years on, it's hard to imagine how Napster could have ever succeeded. And that's before taking into account all of the self-inflicted wounds the company visited upon itself.

A series of management regimes were recruited to try to build Napster into a proper company, but as Jordan Ritter has said of the quality of leadership Napster was able to bring in, "You would think the truly fastest growing Internet startup in the world would attract the best people. But it did not. It attracted the worst people."[36] Napster was not able to attract the best investors either. Unlike Google, which was raising money at almost the exact same time, Napster never landed a deal with the VC blue chips like Kleiner Perkins (though Kleiner took a hard look before passing).

It turned out that Napster's biggest problem was what it *actually did*: allow users to exchange copyrighted songs for free. It allowed people to pirate music. It was hard to argue that this was not, at least in some way, illegal, and that was what scared off the blue-chip investors and big-name management types. Napster would argue vehemently that it was merely a middleman; a technology that allowed users to connect; in some ways it was no different than an ISP like AOL or a web service like Yahoo. People could—and did—exchange copyrighted material on AOL all the time, and no one argued that AOL was illegal. To this day, Napster insiders like Jordan Ritter believe that there was a sound legal loophole for Napster.[37] In an age of computer networks, how did it make sense to blame a technology itself for how its users employed that technology? Ever since the advent of the CD, music was nothing more than ones and zeros, digital lines of computer code. When you bought a physical album, you had always been allowed to give it to your friend or make them a mix tape from it. Because you could now do the same thing digitally, because you could now store your entire music collection on your hard drive instead of on shelves —how did that suddenly make it wrong to do with your music what you wanted?

Nonetheless, the legal aspects of what was happening on Napster's network were new and untested by precedent. Everyone knew that it was only a matter of time before Napster wound up in court, and sure enough, on December 6, 1999, the Recording Industry Association of America filed a lawsuit against Napster in San Francisco's U.S. District Court. Napster was not even six months old.

This is another point that's widely misunderstood about the Napster story. The lawsuits and the media publicity that came with them helped *create* the Napster sensation. It was almost a textbook example of the Streisand Effect, the phenomenon (as Wikipedia describes it) whereby an attempt to hide, remove or censor a piece of information has the unintended consequence of publicizing the information more widely. Before the lawsuit, there were maybe 50,000 users on Napster; a month after the lawsuit, that number had tripled to 150,000.[38] By the summer of 2000, there were more than 20 million.[39] The phenomenon of Napster, this seemingly organic impulse that suddenly inspired millions of everyday people to skirt copyright laws and social conventions and begin exchanging music freely with one another—it was largely inspired by the publicity surrounding Napster's legal battles.

Napster played up the publicity for all it was worth. It cast itself simultaneously as (1) the little guy getting beat up on by greedy corporations, (2) the cutting-edge technology company that the dinosaurs of old media were threatened by and (3) the champion of everyday users who just wanted to consume their music the way they wanted. Napster quietly encouraged the campus protests when the RIAA pressured colleges to block Napster from their networks. As would later come out during litigation, Napster even paid some musicians to publicly support the service, encouraging them to laud Napster as a foil to the rapacious record industry. And when the vociferously anti-Napster band Metallica showed up at Napster's offices to deliver a list of more than 300,000 Napster users it claimed were pirating the band's tracks online, Napster organized a "spontaneous" same-day counterprotest to ensure that the event made front-page headlines. "Fuck you, Lars, it's our music too!" protesters shouted at Metallica's Lars Ulrich as he delivered the list of usernames.[40]

Napster also played up the by now well-worn angle of a young company founded by a bunch of kids who just wanted to change the world. Shawn Fanning and Sean Parker were paraded regularly on MTV and other television outlets. Napster made the cover of magazines from *Rolling Stone* to *Time*. Shawn Fanning introduced Britney Spears at the 2000 MTV Video Music Awards and hobnobbed publicly with famous artists such as Billy Corgan and Courtney Love. Fanning even testified before Congress alongside Metallica's Ulrich.

But the bottom line was that Napster users *were* pirating copyrighted songs, and it was this simple fact that Napster couldn't escape. Napster hired lawyer David Boies, fresh off his victory over Microsoft, to argue that Napster didn't

have any control over what its users did, that its servers didn't touch, much less store, any of the copyrighted material, that it was no more liable for crimes committed because of its technology than the phone company was for allowing users to dial in to Napster in the first place. But none of it mattered in the end, because the courts decided that Napster *knew*; it knew what its users were up to, and that made all the difference in the world.

Napster was ultimately done in by internal documents that were uncovered during the RIAA trial. In a key email exchange between Shawn Fanning and Sean Parker (who was, ostensibly, the strategic visionary of the early Napster), Parker wrote about the need for Napster users to protect their anonymity: "Users will understand that they are improving their experience by providing information about their tastes without linking that information to a name or address or other sensitive data that might endanger them (*especially since they are exchanging pirated music*)."[41] The emphasis on that last statement is mine, but at trial, the RIAA stressed that section as well. In her initial ruling against Napster, the judge in the case, Marilyn Hall Patel, ruled that the evidence "overwhelmingly establishes that the defendant had actual or, at the very least, constructive knowledge" that users were using Napster to pirate copyrighted music.[42] Napster briefly got relief on appeal, but ultimately, rulings came down that said the company either had to put a system in place that blocked copyrighted material on its network, or else it had to shut down the entire network. Fanning and the other Napster engineers tried gamely to implement algorithms to do just that, and they succeeded in blocking 98% to 99% of the offending material. But the judge was ultimately not satisfied unless the percentage of blocked material reached 100%, and Napster was never quite able to achieve that. When all legal options were exhausted, Napster filed for bankruptcy on May 14, 2002, and fired all seventy employees, including Shawn Fanning, who had stayed with his brainchild until the bitter end (Jordan Ritter had left in October of 2000, and Sean Parker had been quietly shown the door after his damning emails had come to light).[43]

Napster was perhaps the victim of its own naïve faith in technology. Did Napster know that people were largely using its technology for pirating music? "Yeah we knew," Napster engineer Ali Aydar would say years later. "But we also knew that this thing called the Internet existed. And it was new. And as it evolved, these things were going to start to happen. And things were going to have to change. And the way in which the world worked was going to have to change."[44] The hope was that if the majority of the music-buying public could be converted to this new way of consuming music—of downloading, of storing

songs on your hard drive, of every song in the world being available at your fingertips—that Napster could then cut a deal with the record companies, something along the lines of "Hey, all your customers are now on our platform. Let us help you reach them, in a mutually beneficial, profitable way." In internal strategy documents drawn up by Parker, this was laid out explicitly: "We use the hook of our existing approach to grow our user base, and then use this user base coupled with advanced technology to leverage the record companies into a deal."[45]

Surely the record companies would see that digital distribution was more efficient. They would see that Napster could help people discover new artists and promote existing ones by creating a central hub. In retrospect, there is no shortage of people, even inside the music industry, who imagine how different the world would be if it had worked out that way—if the music companies had partnered with Napster and accepted the inevitability of technology. "Something like thirty million–plus music fans were in one spot online," says Jeff Kwatinetz, a former representative of music artists ranging from Linkin Park to Mandy Moore and Ice Cube. "At the time, the idea of all the music you would want for $15 a month was an appealing thing and studies showed most users would have paid it."[46] Napster could have been *the* portal for all of music, a Yahoo of music, a Google of music, maybe even a Facebook of music.

Of course, a less polite word for "leverage" is "extortion." Perhaps Napster's biggest misstep was trying to leverage the record companies into a deal, given that the music business has always been known as one of the most notoriously cutthroat and aggressive in the world. This was an industry with quite literal mob ties throughout much of its existence. Napster simply picked a fight with the wrong adversary. The music industry was never interested in a deal. The music industry was only ever interested in suing Napster dead.

The RIAA would follow up its victory over Napster by attempting to sue other digital technologies out of existence, and even, eventually, suing music consumers themselves—tens of thousands of them, in fact. Of course, all this did nothing to halt the advance of file-sharing technology. In Napster's wake, first came Gnutella, from Justin Frankel, who had created Winamp. Gnutella spawned a whole ecosystem of next-generation file-sharing networks like LimeWire, BearShare, Morpheus and many more. A few years later, in 2003, a twenty-five-year-old coder named Bram Cohen released the BitTorrent protocol, which took file sharing to new frontiers like movies, TV shows, and video games.

If Napster had been naïve to think it could have done a deal with the record

companies, then the record companies were certainly naïve to think destroying Napster would somehow make the threat of digital technology go away. But, as has been endlessly discussed and is widely understood, the music industry was caught in a classic innovator's dilemma, tied to a highly lucrative business model it was loath to give up, even in the face of an existential threat presented by new technology. Everyone knew the music industry had gotten filthy stinking rich on the back of the compact disc in the 1980s and 1990s. Having convinced all of us to repurchase our record collections in digital form, the music industry went from selling 800,000 CDs in 1983 to 288 million in 1990 and nearly a billion in the year 2000.[47] Unlike with most digital technologies (the price of which almost always declines over time), the price of the average CD seemed only to inch upward every year, approaching nearly $20 a disc by the turn of the century.

But even this analysis—the record companies were wedded to the cash cow of the CD—doesn't quite get at the truth behind the revolution that Napster began. Napster was the first signal that the web had changed consumer behavior in a fundamental way. Today, we live in a world where consumers not only expect, but demand, infinite selection and instant gratification. Amazon had first introduced the concept of infinite selection, and now Napster was training an entire generation to require the instant gratification. Shawn Fanning had been right from the very beginning: digital really was a better way to distribute music. Computers (at least, the gadgetry computers would evolve into) would turn out to be pretty damn good music consumption machines.

Advertising might have been the first industry the web disrupted, but Madison Avenue adapted to the change, quickly following our attention spans and our eyeballs as they drifted online. The record companies, in contrast, refused to budge as the habits and preferences of music consumers changed. It was never piracy that was the problem for the music industry (at least, not entirely). But rather, it was the stubborn refusal to adapt to a revolution in consumer expectations that has, at its root, truly bedeviled the record companies, and the television companies and the movie companies, and on and on and on over the course of the Internet Era.

Infinite selection. Instant gratification. On any device. When it comes to digital disruption of media, it is almost never about free content or piracy, not at the core. It is *always* about giving people what they want, when they want it, how they want it. Napster seemed to understand this intuitively, even if its execution on this insight was bungled. In early interviews where Shawn and Sean were trotted before the media to explain what Napster was trying to do,

Sean Parker would say things that, in retrospect, were completely dead-on. "Music will be ubiquitous and we believe you'll be able to get it on your cell phone, you'll be able to get it on your stereo, you'll be able to get it on whatever the device of the future is. And . . . I think people are willing to pay for convenience."[48] The Internet and the web and Google had already made information ubiquitous. Napster was the first company to prove that, in the future, media would be ubiquitous as well.

■

EVERYONE TENDS TO FOCUS on the Napster trial as the pivot point in the history of modern technology versus traditional media. But there was another trial, from around the same time, that would ultimately have a larger impact on how we consume media in the digital era. In September 1998, a small company called Diamond Multimedia released one of the first portable MP3 players, the Rio PMP300. The PMP300 had only 32 megabytes of storage, so it could only hold about 30 minutes of music—half an album or so, at decent sound quality; a whole album and a couple extra songs if you didn't mind compressing everything to a level of barely tolerable sound quality.[49] About a year before it sued Napster, the RIAA sued Diamond Multimedia. Before it had even heard of Napster, the record industry knew it didn't want MP3 as a technology to catch on. But while Napster was eventually defeated, the RIAA *lost* the Diamond Multimedia case. The Rio PMP300 went on to become the first commercially successful portable MP3 player.

As the author Stephen Witt has noted in his book *How Music Got Free: A Story of Obsession and Invention*, from the perspective of history, the music industry won the wrong lawsuit.[50]