14

WEB 2.0

Wikipedia, YouTube and the Wisdom of Crowds

In a way, blogging was simply the inevitable migration of publishing into the digital arena. The music site Pitchfork.com, which flourished in the early 2000s, was simply doing things that magazines like *Spin* and *Rolling Stone* had been doing in paper form for years: reviewing music and profiling new artists. But *Pitchfork* also encapsulates how blogging changed the media landscape in terms of taste-making and authority. *Pitchfork* allowed a slate of obscure music writers to challenge the established order merely by gaining credibility through the power of their unique point of view. This phenomenon, whereby the best content rose to the top and the most prominent voices became the new "establishment," occurred in numerous interest niches across the Internet. From food to fashion, from automobile blogs to "mommy" blogs, even touching such rarefied academic arenas as finance, economics and the law, blogs allowed new voices to surface and claim the mantle of "expert," without any official sanction, training or even previous experience.

Perhaps the most illustrative example of this came in the realm of politics. September 11, 2001, was transformative for obvious reasons. But that tragedy was also the first time a historical event could be recorded online from the perspective of those who experienced it firsthand. Thousands of bloggers recorded their emotions and their impressions and even their direct experiences for posterity. "Only through the human stories of escape or loss have I really felt the disaster," Nick Denton wrote for the *Guardian* newspaper on September 20,

2001. “And some of the best eyewitness accounts and personal diaries of the aftermath have been published on weblogs.”[1] It was what Justin Hall had been advocating for years: the common man as recording vessel for history. “If everyone was to tell their stories on the web, we would have an endless human storybook, with alternating perspectives. . . . Give someone a digital camera, a laptop, and a cellular telephone, and you’ve got an on-the-spot multimedia storyteller from anywhere in the world.”[2]

From the right side of the U.S. political spectrum, the response to 9/11 was immediate and strident. A group of conservative-leaning blogs like Instapundit, Little Green Footballs, Power Line and others, began advocating for an aggressive global war on terrorism. These sites were known collectively as the “war bloggers” in the coming years as they became vociferous cheerleaders for the wars in Afghanistan and Iraq. Conversely, it was opposition to the Iraq War that saw a community of left-wing blogs spring up like MyDD, DailyKos, Eschaton, Hullabaloo and more. The lefty blogosphere called itself the “netroots” and could rightfully claim credit for giving energy to the brief, insurgent antiwar presidential candidacy of Howard Dean in 2003.

Again, on both the left and the right side of the “blogosphere,” new voices rose from seemingly nowhere, gaining a reputation through smart comments posted on popular blogs, graduating to influential blogs of their own, and then often going on to positions of prominence at “mainstream” journalistic publications or even actual political positions. In the United States, we live in a post–political-blogging world where movements can arise online and take over the mainstream discourse. The most prominent examples of this new reality come from the right, in the form of the Tea Party movement and especially the Trump presidency, which has seen bloggers (in the form of Breitbart) ascend to the highest corridors of political power.

But perhaps what’s most interesting to observe about the rise of blogging is how the habits and behavior of web users themselves changed. If the web in the dot-com era had been about, in the words of the technology journalist Sarah Lacey, “taking prepackaged content from the offline world and throwing it onto a site,”[3] the new web was about you (and everybody else) putting up your own content, discovering it for yourself (and others), organizing it yourself and determining that your content was just as interesting and valuable as anything else in the media landscape. It had taken about a decade for mainstream users to acclimatize themselves to the web, but now that they had the lay of the land, they were no longer content to merely “surf.” Even everyday web users were now ready to *participate* in the web. As Marc Andreessen had anticipated all the

way back in the days of the Mosaic browser, the "riff-raff" were ready to join the party in a major way, not just as consumers, but as producers. To quote the title of a popular book from shortly after this era, the postbubble Internet was a moment of *Here Comes Everybody*.

Some credit can be given to Napster for opening these floodgates. All those tens of millions of users who traded MP3 files were proactively and spontaneously self-organizing and using their own libraries to create content for others. Napster was the first time mainstream web users saw the utility in producing, not just consuming, content. And baked into Napster was a "social" component to all this activity. If you found a song you liked from another user on Napster, you could also browse the other files in that user's library. If you both shared an interest in a given band, then maybe you would like that *other* band that your friend on Napster had so many MP3s of. It was like the Netflix recommendation engine, but impromptu and self-created. It was the act of finding like-minded individuals, of creating community out of silos of shared interest.

This "social" aspect of the web began manifesting itself in a number of ways beyond Napster and blogging. A link-blog site called Slashdot grew popular around the turn of the century by aggregating the blog-post and news-item deluge that came online every single day. In the comments of every link post, the thousands of members of the Slashdot community debated and discussed the posted articles. Order was given to the chaos by the Slashdot community itself. Randomly selected users were given moderation privileges to vote up or vote down content on a scale ranging from "insightful" to "troll," thereby allowing the community to police discourse on its own.

Digital cameras were just becoming popular in the early 2000s, and sure, you could make actual prints from your photos on your ink-jet printer and then mail those to your grandmother; but conversely, you could also just post an entire album online via a site like Flickr (launched in February of 2004) and simply send Nana the link to your Flickr page. More than that, you could share your pictures with complete strangers if you wanted to. How would strangers find your photos? Well, Flickr allowed you to "tag" your photos with keywords that enabled other users to search for them. If someone wanted to browse a bunch of photos of the Grand Canyon, they could type those keywords into Flickr and see the results of a thousand different strangers' summer vacations.

From the days of the Netscape browser, users had used bookmarks and "favorites" to keep track of their favorite web pages. But what if you wanted to see what other people had bookmarked? Del.icio.us (launched in September of

2003) let you do just that, allowing users to discover cool new things on the web by sharing their bookmarks with each other, just as Napster had allowed them to exchange songs.

The new postbubble web was about the users and the content in equal measure. It was about spontaneous impulses like “sharing” and self-organizing schemes like “tagging” and taxonomies. It was about how the content created by and for the hoi polloi often ended up being more engaging and exciting than the content that was prepackaged or professionally produced. And increasingly, the new web was about the collective “wisdom” of the crowd to create and organize the anarchy.

The idea of collaborative effort and collective organization had long been a common practice in hacker and software development circles. Just as each of the hackers on w00w00 had pitched in to help Shawn Fanning refine Napster, groups of programmers often came together and formed communities around the development of “open source” projects like the Linux operating system. Far from being a case of “too many cooks in the kitchen” creating a muddled fiasco, open-source development proved that complete strangers could independently, and without much centralized coordination, come together to collectively produce things in an orderly, sublime way.

A veteran software developer named Ward Cunningham brought this practice to the web for the first time on his Portland Pattern Repository, a website for other programmers to contribute and share programming ideas. On March 25, 1995, Cunningham installed a subpage on the site called WikiWikiWeb. The “wiki” (the term came from the Hawaiian word for “quick”) constituted a series of pages that could be edited by any user. So, a given user might post some code patterns to the wiki, and another user might come behind him and add to those patterns, change them, even completely replace them. But all edits were stored, and the page could revert to previous versions if any user chose to do so. It seems counterintuitive that such a system could work, but Cunningham learned that, given enough input from enough interested users, his Wiki system worked quite well. Cunningham is famous for coining “Cunningham’s Law,” which finds that “the best way to get the right answer on the Internet is not to ask a question, it’s to post the wrong answer.”[4] If a user contributed code patterns to his site that other users found wrong or merely objectionable, Cunningham found that, almost inevitably, another user would come along and right the wrong.

Wikis tapped into a powerful impulse of collective action. A few years later, an obscure entrepreneur would make use of this impulse to save his own

struggling creation. Jimmy Wales was a serial dot-com entrepreneur who had found a modest degree of success by creating more sophisticated web directories—sites like Yahoo, but more focused. Wales also had a lifelong passion for encyclopedias and was obsessed with the notion that the web could create the largest encyclopedia conceivable. “Imagine a world in which every single person is given free access to the sum of all human knowledge,” Wales would write later.[5] In early 2000, he launched what he called his Nupedia project, soliciting experts in a wide range of fields to contribute articles for what he hoped would eventually become an infinity encyclopedia. Contributors to the project were required to be knowledgeable in a given topic, and they would have to submit their articles to a rigid system of peer review by vetted editors. Also, the editors themselves had to be credentialed. “We wish editors to be true experts in their fields and (with few exceptions) possess Ph.D.s.,” the Nupedia policy stated.

But Nupedia’s rigid quality control apparatus proved inefficient. It wasn’t until September 2000 that the first article made it through the layers of editors, and by the end of the year, less than two dozen had been published on Nupedia’s website. In frustration, on January 10, 2001, Wales installed a descendant of Cunningham’s original wiki software on Nupedia’s server. This “Wikipedia” was merely intended as a separate feeder service to speed up the Nupedia submissions process. Articles would be collectively written and edited on Wikipedia, then fed into the existing peer-review editing process. Almost immediately, however, Wikipedia overtook Nupedia not just in the quantity of articles that were created, but in the quality as well. The first article created, on January 15, was on the letter “U” and investigated the origins and usage of the twenty-first letter of the English alphabet.[6] It was comprehensive, it was well written, and it was—to the surprise of Wales and his team of editors—accurate. The few thousand users who had shown up to test out Wikipedia had, through their collective input and edits, gotten the article polished to near-authoritative quality.

Within a month, Wikipedia had around 600 articles, achieving in a matter of weeks more than Nupedia had achieved in a year. The experiment was promoted on Slashdot, and soon Wikipedia was flooded with Slashdot’s passionate users, members of a community who were already comfortable with collective editorial action. Within a year, Wikipedia had grown to 20,000 articles. By 2003, the English-language Wikipedia had more than 100,000 articles, and versions of the service were springing up in every language imaginable. By that point, Nupedia and its rigorous system of editors and peer review had long been abandoned.

What confounded everyone who learned of the success of Wikipedia was

that it actually worked! “Couldn’t total idiots put up blatantly false or biased descriptions of things, to advance their ideological agendas?” asked one of the leads of the original Nupedia project on internal Wikipedia message forums. “Yes,” replied a Wikipedia partisan, “and other idiots could delete those changes or edit them into something better.”[7] It turned out that the “infinite monkey theorem” about giving enough monkeys typewriters and eventually producing Shakespeare—was not exactly fanciful. Enough self-interested strangers could achieve a fair degree of accuracy on a wide range of topics. In 2006, there were 45,000 active editors of the English-language version of Wikipedia alone.[8]

And Wikipedia had unique advantages that the web made possible. In the coming years, as any news or historical event occurred, Wikipedia contributors would post an up-to-the-minute factual summation of these events, and then amend the entries to reflect changing circumstances or new information. Wikipedia was often accurate and authoritative in near-real time, and it had the infinite space and resources of the Internet to play with, so it could serve what became known as the “long-tail” of content. Any encyclopedia worth its salt might have an article on World War II. But Wikipedia could produce a 418-word entry on, say, the Compton railway station, an abandoned stop on the Didcot, Newbury & Southampton Railway in England. Or, it could produce a detailed plot and development synopsis on Season 8, episode 14 of the TV show *Cheers*, the one where Cliff Clavin goes on *Jeopardy*. No other encyclopedia in history was capable of that sort of breadth of topics.

Wikipedia was a modern miracle and soon became one of the most trafficked websites in the world. Wales had originally intended the project to be a commercial one, supported by advertising. But when the contributors and editors revolted at the very suggestion of putting ads up on Wikipedia, Wales instead made the site into a nonprofit enterprise. To this day, it is supported by contributions from the public and is thereby an open-source counterweight to the proprietary “answer engine” that is Google.

■

GRADUALLY, PEOPLE BEGAN to notice that there was a new energy on the web and it shared several characteristics. The long tail. The wisdom of crowds. Users creating content of and to their own design. In 2004, this new Internet energy gained the name Web 2.0, after a similarly named conference held in October 2004. If Web 1.0 was about browsing stuff created by others, Web 2.0 was about creating stuff yourself. If Web 1.0 was about connecting all the computers in the world together, then Web 2.0 was about connecting all the *people* in the world

together, via those interlaced computers. If the clarion call of Web 1.0 was the Netscape IPO, then the coming of age of Web 2.0 was Google's IPO. "Web 2.0 means using the web the way it's meant to be used," wrote Paul Graham, a veteran entrepreneur of the Web 1.0 era who would soon become a key driver of Web 2.0 as an investor. "The 'trends' we're seeing now are simply the inherent nature of the web emerging from under the broken models that got imposed on it during the Bubble."[9]

Within the technology industry itself, the sense that the Internet revolution was back in gear came via the promotional efforts of—what else?—a blog. On June 10, 2005, Michael Arrington, a thirty-five-year-old former Silicon Valley lawyer who was active during the dot-com years, started posting to a personal blog at TechCrunch.com. Arrington's entries were mostly musings about the new services, websites and companies he saw bubbling up through the Web 2.0 scene. But he soon branched out to covering the actual news of Web 2.0: what new companies were being founded and by whom; what startups were raising an investment round and with whom; what hot new websites had been acquired, and by whom. TechCrunch became not only the cheerleader of the Web 2.0 movement, but, in a sense, proof that the movement even existed. Arrington simultaneously became a power player in his own right, since his site became a PR bonanza for whatever new service or company he deigned to blog about. As *Wired* magazine put it, "A positive 400-word write-up on TechCrunch usually means a sudden bump in traffic and a major uptick in credibility among potential investors." When TechCrunch gave a glowing write-up to a startup named Scribd, as *Wired* reported, "CEO and cofounder Trip Adler says he had 10 calls from venture capitalists within 48 hours."[10]

Indeed, the startup scene was back in full swing, in no small part thanks to TechCrunch and the hype around Web 2.0. Usage of the Internet had never dipped and indeed was finally reaching critical mass in the developed world. In 2003 alone, the percentage of Americans with broadband Internet connections in their home increased from 15 to 25%.[11] A new technology called WiFi arrived on the stage to make the notion of surfing the web something that felt ubiquitous and commodified. Even online advertising was coming back, providing that same old business model (but with different tools and greater numbers) to new online efforts. Between 2002 and 2006, U.S. advertisers increased their online ad spend from $6 billion to $16.9 billion.[12]

The venture capital machine started to lurch back into life to fund this new activity. VC investments in U.S. startups bottomed out at $19.7 billion in 2003, a

far cry from the dot-com–era peak of more than $100 billion in the year 2000.[13] In the coming years, VC investment would rise—modestly but steadily—reaching $29.4 billion in 2007.[14] A slew of new companies were funded, but the renewed interest in Internet startups was not a replay of the late-nineties frenzy. Both investors and entrepreneurs had been chastened by the bubble's aftermath. Get Big Fast was no longer the strategic mantra; multimillion-dollar advertising campaigns and gaudy launch parties were out. Instead, Web 2.0 companies aimed at refining their products and services, carefully cultivating a user base through feature innovation and word-of-mouth discovery, all while focusing like a laser on issues such as reliability and scalability.

VC investment didn't roar back in huge numbers because it didn't have to. In the Web 2.0 era, you could create a service used by millions in a matter of months, and you could do so for pennies on the dollar—at least, compared to the dot-com era. The hangover from the bubble fallout meant that talented programmers could be hired on the cheap; the infrastructure glut leftover from the global fiber buildout meant that bandwidth, storage and data costs were lower; and the tools developed during the bubble meant that you didn't have to build a company from scratch anymore—you could cobble one together using free and open-source tools to assemble the building blocks of a minimum-viable product for next to nothing. By some estimates, the cost of starting a web company had fallen by 90% in the few short years of the nuclear winter.[15]

The website Digg was perhaps the poster-child company of the Web 2.0 era, and it illustrates this change in startup economics perfectly. In 2004, twenty-seven-year-old Kevin Rose had an idea for a new website that would help plugged-in geeks like himself discover the news of the day: whatever was hot on the blogs or even mainstream sites like the *New York Times*. His vision was of a site that took the community-voting aspects of Slashdot, but gave the power to surface news to anybody. On Digg.com, any user could submit a story and other users could "digg" it. If enough users dugg, then the story would rise to the front page. Conversely, if users didn't like a story, they could vote to "bury" it. Rose registered the Digg.com domain name (that was the biggest expense, actually; he had to buy the domain from an existing owner), paid a programmer in Canada $12 an hour to code up the site, and paid $90 a month to have a company host it. The site launched on December 5, 2004.[16] Rose's total outlay was around $10,000.

For that investment, Rose soon had the hottest site on the Internet. Within a year, Digg passed Slashdot in traffic.[17] Making it to the front page of Digg could

drive scads of traffic to a website, so publishers all around the web began to add "Digg This" buttons to their websites. Within two years, Digg had nearly as much web traffic as the *New York Times* and more than 1 million people came to the site daily, "digging" thousands of stories.[18] Digg was nominally profitable from day one, thanks to AdSense ads from Google, and later, banner ads from more traditional marketing networks. In 2007, Digg landed a $100 million ad deal with Microsoft. By that point, Rose had appeared on the cover of *Businessweek* under the headline "How This Kid Made $60 Million in 18 Months." That estimation of Rose's paper wealth came from the valuation given to Digg by venture capitalists. But the truth was, Digg had only raised money reluctantly. As Rose and his cofounder Jay Adelson made the rounds on Sand Hill Road, home to the most powerful Silicon Valley VCs, they were shocked by what they saw as the outdated thinking among the money men. "They are still back in the 1998 belief system that it's all about the portals," Adelson marveled.[19] The VCs wanted to throw tens of millions of dollars at them in order to build the next Yahoo or AOL. Rose and Adelson were content to raise a paltry $2 million. They didn't really need the funding, and besides, raising less money meant keeping more equity for themselves.

The new Web 2.0 companies didn't need to raise as much money and, unlike just a few years previously, none of them were in any hurry to go public. In the wake of the bubble bursting, a wave of scandals involving companies such as Enron and WorldCom had ushered in a new era of financial regulations. The Sarbanes-Oxley legislation especially meant that there were fewer advantages to going public and more incentives to stay private for as long as possible. Without the venture capitalists breathing down their necks for a financial "exit," the Web 2.0 companies were more in control of their own destinies and wary of the pressures that a blockbuster IPO would impose upon them. The lesson of the bubble had been learned: you can go for broke, but try to build a *real* company first.

That didn't mean the money men were denied their "exits." As the survivors of the dot-com bubble began to see their balance sheets return to health, there was an entire group of deep-pocketed acquirers that would begin to pick off the most promising members of the Web 2.0 class. Yahoo swallowed up Flickr and Del.icio.us in 2005, for around $40 million and $20 million, respectively. Scandinavians Niklas Zennström and Janus Friis created the second-generation peer-to-peer networking platform Kazaa before turning to that same P2P technology in order to make phone calls over the web. They founded Skype, enabled hundreds of millions of users worldwide to call and chat with each other

for free, and sold the company to eBay for $2.6 billion in September 2005.

But the acquisition saga everyone followed in those early Web 2.0 days was that of YouTube. Late in 2004, three former PayPal employees, Chad Hurley, Steve Chen and Jawed Karim, were mulling over a problem: why wasn't it as easy to post a video to the web as it was to post a photo to Flickr or a blog post to a blog? YouTube was the site they launched to solve that problem, and from the very beginning, the overriding idea was for dead-simple, push-button video uploading.

But what, exactly, should people be encouraged to upload? Should YouTube encourage people to create original, dramatic videos with near–television-production quality? Or maybe YouTube would just host videos for eBay auctions and use the thriving auction economy to jumpstart growth just as PayPal had (they were card-carrying members of the PayPal Mafia, remember). There was even some early discussion about copying HotorNot.com, a popular Web2.0 site where users uploaded profile pictures, and other users voted the portraits up or down based on attractiveness. "In the end, we just sat back," said Hurley, meaning they just let the users upload whatever they wanted no matter how silly, or inane, or personal, or *whatever*.[20] It was the Web 2.0 way.

The first video posted to YouTube exemplified this attitude. *Me at the Zoo* is a nineteen-second video of Jawed Karim at the San Diego Zoo in front of the elephant exhibit. Uploaded on April 23, 2005, Karim offered the following pithy narration:

> Alright, so here we are in front of the, uh, elephants. Uh. The cool thing about these guys is that they have really, really, really long, um, trunks, and that's, that's cool. And that's pretty much all there is to say.

Not exactly "one small step for man" stuff, but credit to the YouTube guys for understanding that that was exactly the sort of video that YouTube was good for.

YouTube was fortunate in its timing. By 2005, broadband Internet adoption continued to increase, and consumer video cameras were becoming common. Even some cell phones allowed you to shoot video by the time YouTube launched. In August of 2005, YouTube got favorable coverage from TechCrunch as well as Slashdot. The number of videos posted started to increase. And then, the post-anything spirit of blogging that YouTube was mimicking helped traffic ramp up even more. In fact, it was the blogs themselves that really helped YouTube explode in popularity. The blogs—and social networks like Myspace.

Aside from push-button-easy uploading, the true brilliance of YouTube was the site's second important focus: dead-simple sharing. After you posted a video

the site's second important focus: dead-simple sharing. After you posted a video to YouTube, you could simply share a link to your uploaded video, just like with Flickr. But you could just as easily cut and paste a few lines of code and your video would play, embedded, wherever you wanted it to: on your website, your blog, or your Myspace page. You didn't ever have to send people to YouTube if you didn't want to. Suddenly, videos were popping up all around the web at a time when web video was still a relatively rare phenomenon. Every time someone embedded a video on a random website, there was that little YouTube logo at the bottom that encouraged people to visit YouTube and try posting videos themselves.

YouTube was incredibly popular on Myspace, but it was the combination of Myspace and the blogs that really caused YouTube to take off. It was the "share-yourself, share-anything!" ethos of the moment combined with the ubiquitous distribution platform of the web that led to what we now call "virality." This was proven by the smash online success of the *Lazy Sunday* video. In 2005, *Saturday Night Live* aired a roughly two-minute musical skit chronicling the antics of a couple of young white dudes in Manhattan hitting up Magnolia Bakery on a Sunday morning and then catching a matinee showing of the recent *Chronicles of Narnia* movie—all set to hard-core rap stylings. It was goofy and catchy, and was also probably a throwaway segment on the show's first airing. But as fate would have it, shortly after the original broadcast, someone posted a video capture of the skit to YouTube, where it quickly racked up 5 million views.[21] NBC's lawyers had it taken down in a matter of days, but not before word of mouth around the video increased YouTube traffic by 83%.

After the early months of indifferent traffic, YouTube's audience exploded faster than any previous website in history (including Google, Myspace and Facebook). By the beginning of 2006, the site was serving 3 million video views a day. Six months later, that number had grown to 100 million views a day. Like most good Web 2.0 companies, YouTube achieved this success on a shockingly small amount of money. The company only ever raised $11.5 million, in two investment rounds. The fact that YouTube could serve video to the world from just a handful of servers (and some helpful content delivery networks in the background) was a powerful testament to the infrastructure the dot-com bubble had bequeathed to this new generation of startups.

Today we're used to popular "memes" bouncing around the world in an instant and have come to expect that social media can make superstars of teenagers from Canada (I'm thinking specifically of Justin Bieber, of course, who would be discovered thanks to videos his mother posted to YouTube).

YouTube was ground zero for things like that, for the birth of modern meme culture as well as the social media–celebrity ecosystem. The idea that random events or random people could "go viral" really entered the mainstream thanks to YouTube. "We are providing a stage where everyone can participate and everyone can be seen," Hurley told the Associated Press in April of 2006.[22] There was no greater Web 2.0 manifesto than that.

But the "Lazy Sunday" phenomenon also pointed to one looming issue that concerned a lot of people about YouTube: there was a ton of copyrighted material uploaded illegally on the site. Sure, there were user-created home movies by the barrelful; but just as common were copies of last night's episode of *Survivor* or even clips from first-run movies still in theaters. In short, there was plenty of piracy going on. Just as with Napster, users came to expect that they could watch anything and everything on YouTube—from the latest Justin Timberlake video to obscure Japanese films from the 1960s.

But that was the issue: how was YouTube anything but Napster 2.0, with all the inevitable liability headaches that would imply? That was why people were obsessed with the who-will-buy-YouTube guessing game in 2006. Even though YouTube was exploding in popularity, it wasn't making any money, and in the postbubble era, an IPO was out of the question without meaningful revenue on the bottom line. So, unless YouTube was able to sell out to a deep-pocketed patron before the lawsuits started flying, it ran the very real risk of being pushed into an early grave.

As would come out in subsequent litigation, the YouTube guys knew perfectly well that there was a ton of pirated material on their site. But they had learned the lessons of Napster. Napster had attempted to make the argument that it enjoyed legal immunity under the Digital Millennium Copyright Act as a neutral platform. Service providers and platforms were protected as "safe harbors" under the law, provided they quickly and efficiently remove copyrighted material when notified. That was what had ultimately doomed Napster: it had never been able to take down 100% of the pirated files on its service. Five years on from Napster, might YouTube be able to find someone who could create a better system to remove illegally uploaded material—someone who had a mastery of algorithms, perhaps?

On October 9, 2006, Google announced that it was purchasing YouTube for $1.65 billion in stock. For the YouTube guys, selling to Google was logical: for all of YouTube's frugality, the cost of serving hundreds of millions of videos would eventually become prohibitive. Bandwidth might have been cheaper now, but who could hope to manage data on a scale that YouTube was achieving?

Google was a perfect fit because its enormous infrastructure allowed YouTube the chance to handle the scale.

But Google's decision to take on YouTube's burden seemed downright crazy to a lot of people. Wasn't Google paying a lot of money to basically assume a huge liability risk? It turned out that Google made one simple calculation when it purchased YouTube: in the broadband era, video was likely to become as ubiquitous on the web as text and pictures had always been. YouTube was already, in essence, the world's largest search engine for video. In fact, it would eventually become the second-most-used search engine, period. With its stated mission to organize all the world's information, Google simply couldn't let video search fall outside its purview.

Google was able to come up with sophisticated automated systems that quickly and efficiently took down copyrighted videos when the rights holders alerted them. Lawsuits from aggrieved rights holders did eventually come, especially a billion-dollar lawsuit from Viacom. But because Google could prove that it was effective in policing content, in 2010 the judge in the Viacom case ruled in Google/YouTube's favor, saying that Google's takedown system was efficient enough that it complied with the Digital Millennium Copyright Act.

Google was the savior Napster never had. It had the infrastructure to allow YouTube to scale up; it had the technical sophistication to keep YouTube on the right side of the law; it had the money to contest the legal battles; and—most important—it provided YouTube with the business model that would allow it to thrive. Those little text ads that Google had put all over the Internet? They could be used to monetize the videos on YouTube just as they could with any other type of content. As the years went by, the text ads could even morph into actual video ads—but algorithmically targeted and effective ads, as Google's ads always were.

And this was the last way in which YouTube's timing was impeccable. The movie and television studios had watched the Napster debacle with dread. They knew their industries were next in line for disruption from the Internet. When that disruption arrived, in the form of YouTube, Hollywood was at least willing to weigh its options this time. Going scorched earth against Napster had not saved the music industry. And so, once Google came to the table with a willingness to share advertising revenue with rights holders, a lot of them (Viacom notwithstanding) were willing to play ball. At least Google/YouTube was offering Hollywood *some* kind of revenue stream. Digital revenue might not be as lucrative as the old analog revenue streams but, well, that was the Napster

lesson, right? Better to take what you could get and embrace new distribution models rather than fight them. The entertainment industry was even now willing to buy into one of the key arguments Napster had tried to make only half a decade before: giving users a taste of your content online was actually great promotion! The phenomenon of *Lazy Sunday* had shown that. By 2008, when YouTube was streaming 4.3 billion videos *per month* (in the United States alone), many people—young people especially—were beginning to watch more video online than they were watching on traditional TV.[23] For the first time, Hollywood stopped fighting disruption, and followed the changing tastes of their audience into a digital future.

■

WEB 2.0 WAS ABOUT PEOPLE expressing themselves—actually being themselves, actually living—online. The last piece of the puzzle was simply to make the threads of all this social activity explicit.

Online chat clients like IRC, through which the Napster hackers had met each other and collaborated, had a technological cousin at AOL. In the days when AOL was still the dominant ISP with more than 20 million users, its internal messaging program allowed you to chat with your friends and family in real time. AOL's chat had an extra feature called the "Buddy List" that alerted you as to which of your friends were online at the same time you were, so you could hit them up for a quick conversation. The system also allowed you to leave an away message so that your friends could know when they might expect you to be online again.

Instant messaging was only intended for internal use by members of AOL's walled garden. But in 1997, the company did something completely out of character: it released the messenger program online as a stand-alone web client. It was known as AOL Instant Messenger, or AIM, and it allowed people to stay in touch with their AOL friends when they were away from AOL. It proved especially popular for people who were at work, where they couldn't log on to AOL, and among teenagers, allowing them to keep up with all of their friends, whether they were AOL users or not. Soon, there were hundreds of millions of AIM users, many times more than the number of actual AOL subscribers at its height. Even as AOL the company began to crumble after the disastrous merger with Time Warner, AIM continued as a breakout success for one simple factor: it was a literal social graph, a tangible map of your online connections and relationships. Chatting on AIM became more popular than email, and your AIM screen name eventually gave you the ability to customize a rudimentary profile,

turning it into a valuable online marker of identity. These features, combined with the away messages and status updates, came to reflect a user's daily circumstance. Add to this the emojis and icons that allowed AIM users to project their mood, and AIM became a fully functional and real-time representation of the digital self. There was even an abortive project to create "Aimster," which would add the ability to search a friend's hard drive and trade files (AOL management, of course, killed that before it could see the light of day).

And that was the problem, of course. AOL had no idea what it was sitting on. AIM was a fully fleshed-out social network. True, it was free to use; but it was making a limited amount of money thanks to traditional banner ads. Had anyone at AOL been able to predict the future, AIM could have been the perfect platform to transition AOL users into the post–dial-up world. Before we were all sending SMS texts, before we all reconnected on Facebook, a great many of us were connected on AIM. The social graph was actually *the* great prize of Web 2.0. Others were only able to seize this prize because AOL dropped the ball. AIM eventually lost its relevance through benign neglect. "If AOL had 20/20 hindsight, maybe the story [of social networking] would have had a different ending," says Barry Appelman, one of the AOL engineers who invented AIM.[24]

■

SOCIAL NETWORKING MIGHT SEEM like a dead-obvious concept in retrospect, but that's only because we've gone through the looking glass into a modern world where the boundaries between our online lives and "real life" have been broken down almost completely. The roots of social networking go all the way back to the early web. The earliest dating sites like Match.com and the message boards on sites like iVillage allowed users to create an online "profile" or representation of your real-world self. And sites like GeoCities and Angelfire allowed users to construct personal webpages so intricate as to serve as virtual avatars in cyberspace.

The first modern social-networking site as we would recognize it today was invented by SixDegrees.com. In 1996, a former lawyer and Wall Street analyst named Andrew Weinreich had an idea inspired by the popular notion that any single person on the planet can be connected to anyone else by around six steps of personal connections—"six degrees" of separation. If that was true, then the web was the perfect tool for mapping those connections.

Launched in early 1997, SixDegrees took off in about a month, in the usual viral way we're now familiar with: users sent their friends invitations to link up on the site. At its peak, the site had 3.5 million members, and in 1999, Weinreich

wisely sold the company for $125 million to another Internet startup.[25]

At the time, many viewed SixDegrees as a newfangled Rolodex at best, a creepy dating site at worst. But Weinreich had been convinced there was something more powerful to the idea of networking online. “We envisioned Six Degrees being something of an OS—of an operating system—and we thought about it in the context of when you’re buying a watch at eBay you should be able to filter the watches based on people’s proximity to you,” Weinreich said. “You should be able to filter movie reviews in the future by who’s reviewing them.”[26] It was the right idea, but as Weinreich would ruefully admit, “We were early. Timing is everything.”[27] The site was expensive to operate in the dot-com days, and of course, there were no photos on the profiles. “We had board meetings where we would discuss how to get people to send in their pictures and scan them in,” Weinreich says.[28] After the dot-com crash, the site was shuttered.

In 2002, a former Netscape employee named Jonathan Abrams launched a site called Friendster. Abrams wanted to rekindle SixDegrees’ original notion of real identities and real personal connections. Within a few months, the site had 3 million users from word-of-mouth marketing alone.[29] The media seized upon Friendster as a more sophisticated version of online dating, and certainly, the digital profile pictures that could now easily be uploaded to your Friendster helped shape this impression of the site. Once connected to someone else, you could browse their friends to see who among them was attractive (and single) and then the idea was that your friend would put in a good word for you. But, this was just as the notion of the Web 2.0 renaissance was taking hold in Silicon Valley, so, dating site or no, Friendster was able to raise $12 million from blue-chip VCs including Kleiner Perkins and Benchmark Capital. In 2003, Google offered to buy Friendster for $30 million in pre-IPO Google stock, but the venture capitalists encouraged Abrams to spurn the offer and instead shoot for the moon.

Friendster ended up missing the moon by some distance. It turned out that hosting blogs or even serving portal pages to millions of users was one thing, but a social network scaled to millions of users was another thing entirely. On a social network, the content was ever-changing, and what was served to each user was often unique to that user, often only in that moment of time. Friendster had to dynamically propagate each new update, each new post—and each new picture. The engineering challenges of delivering what was quickly becoming a deluge of content were at a whole new scale, and Friendster simply wasn’t up to the challenge.

“When it grew as fast as it did, we absolutely weren’t prepared for it,” Abrams said later. “Throughout 2004, 2005, Friendster barely worked. The site was really slow; it was buggy. That, unsurprisingly, caused an exodus of users to leave.”[30] When Friendster users grew frustrated waiting thirty or more seconds for pages to load, they had a throng of Friendster copycat sites to turn to instead. Like any good idea, the rebirth of social networks inspired dozens of people to try their hand at the concept. Many of the Friendster copycats tried to create social networks that targeted specific niches: college students, high school students, even, in the case of Dogster.com, pet owners.

One of the copycat sites that rushed in to tempt away disillusioned Friendster users was called Myspace. Myspace was owned by eUniverse, a dot-com survivor that made a lot of money peddling wrinkle cream (“Better than Botox”) via online ads that purported to offer the cream for free despite built-in expensive automatic refills, and that made advertising claims that the FDA asserted “were not supported by reliable scientific evidence.” An eUniverse employee named Tom Anderson became obsessed with Friendster and convinced his boss, Chris DeWolfe, that creating a Friendster clone might be a cheap and easy way to amass more people for eUniverse’s marketing lists. On August 15, 2003, Myspace was launched as a nearly feature-for-feature clone of Friendster. Users had a profile page where they could post pictures, share their interests and hobbies, and link to the profiles of their friends and family. But Myspace also added a kitchen sink’s worth of features, such as blogs, horoscopes, games and more.

One of the things that was driving users away from Friendster (aside from the slow performance) was the fact that Abrams had insisted on a strict fidelity to identity. Anytime users created a Friendster account under a pseudonym, or started a parody account or pretended to create an account as a celebrity, Friendster would delete it. Myspace had no such regulations. If you wanted to sign up as Leonardo DiCaprio or Bugs Bunny, Myspace let you do it. Furthermore, you could follow anyone you wanted, whether you truly knew them or not. Myspace was the first to hit on a key concept in social networking: linking to others could be a way of mapping your personal connections, but it could also highlight your personal tastes. Friending, or “following” another profile, could be a powerful vote of interest and engagement. When this was combined with the ability to host MP3 files on your profile, Myspace became a potent venue for promotion, especially among musicians. Now that Napster was gone, an entire generation of unknown musical acts ranging from Fall Out Boy and My Chemical Romance to Arctic Monkeys would rise to prominence by engaging with their thousands of fans, promoting tour dates and even releasing

engaging with their thousands of fans, promoting tour dates and even releasing new songs on their Myspace pages.

Myspace also had a laissez-faire attitude when it came to self-expression. Users could redesign their pages at will, hacking into the design code itself to create flashy, colorful, even garish profiles. This appealed especially to teenagers, who decorated their Myspace pages like they would decorate the walls of their adolescent bedrooms. Myspace also looked the other way when users posted racier content. Profiles featuring scantily clad women abounded. This side of Myspace was exemplified by Tila Tequila, a young Vietnamese-American model who was one of the many users fed up with Friendster. "I was getting too many friend requests, and the pictures were too hot," Tequila said about Friendster's habit of repeatedly banning her profile.[31] So she took the tens of thousands in her digital audience to Myspace, where she could represent herself however she wanted. Soon her "friends" numbered in the hundreds of thousands and Tequila achieved that unique mid-2000s form of D-level fame. "There's a million hot naked chicks on the Internet," Tequila told *Time*. "There's a difference between those girls and me: Those chicks don't talk back to you."[32]

Thanks to all of these factors, Myspace quickly rocketed past Friendster to become the king of the social networks, racking up 1 million users less than six months after launching and 3.3 million after a year of operation, with 23,000 new users signing up daily.[33] By May of 2005, Myspace was attracting 15.6 million visitors every month.[34] Myspace founders Tom Anderson and Chris DeWolfe became celebrities in their own right. In Anderson's case, it was because he was the guy who interacted with the users; by default, Tom was every new user's first friend. For his part, DeWolfe put himself forward as Myspace's strategic visionary. "We want to be the MTV of the Internet," Wolfe told investors.[35] To the *New Yorker*, he proclaimed: "The Internet generation has grown up, and there are just a lot more people who are comfortable putting their lives online, conversing on the Internet, and writing blogs. This generation grew up with Napster and the iPod."[36] Myspace was just serving this new audience's behavior and expectations.

But the story of Myspace is slightly different from that of the other companies in the Web 2.0 wave. For one thing, Myspace was Los Angeles–based, a key factor that may have contributed to the site's focus on glam and glitter. And—uniquely—Myspace wasn't a startup. Rather, it was a subsidiary of a parent company. Anderson and DeWolfe weren't actually calling the shots at Myspace. That parent company, eUniverse, had rebranded itself as Intermix in order to escape the shadow of its seedy past, and as the excitement over Web 2.0

grew more frenzied, Intermix decided the time was right to cash in on Myspace. In July of 2005, Intermix announced that it (and therefore, Myspace) had been acquired for $580 million. The acquiring party was not a Google, or even a Yahoo, but News Corp, the company run by media mogul Rupert Murdoch.

Coming as it did among the slate of other Web 2.0 acquisitions, as soon as the deal was announced, many in the press and even some in the tech industry itself were quick to announce that another bubble had formed in Silicon Valley. But for a while, Myspace's unbelievable growth made those fears seem far-fetched. By the end of 2005, a mere six months after the acquisition, Myspace could claim about 40 million registered users and more monthly pageviews than eBay, AOL or even Google.[37] By the time Myspace inked a $900 million advertising partnership with Google in 2006, it looked like social networking was, indeed, the next big thing. MySpace was the new 800-pound gorilla on the web, and Rupert Murdoch had pulled off the steal of the new digital century.

But even when Myspace was at its zenith in terms of users and traffic and revenue, people couldn't stop comparing it to another of the Friendster clones, particularly the clone that had chosen to focus exclusively on college students. In a November 2007 News Corp earnings conference call, Rupert Murdoch himself dismissed this competitor, Facebook, as merely a "Web utility similar to a phone book." Myspace, by comparison, had "become so much more than a social network. It connects people, but it's evolved into a place where people are living their lives. A social platform packed with search, video, music, telephony, games."[38] Little did Murdoch know that, even as he said those words, the battle for social networking was already over, and Myspace would join SixDegrees and Friendster as an also-ran in the history books.