

## From the Many to the Few

IT WAS AN ODD moment in the history of modern business. Usually when one company acquires another, the deal is announced in a carefully staged and scripted event. Reporters are shepherded into a theater or a hotel ballroom, where the chief executives of the two firms stand together at a podium. They speak in general terms of the fabulous prospects for the new company, touting the financial and organizational “synergies” that the combination will provide, and they stress that the deal should be viewed as “a merger of equals” between two businesses with rich and distinctive histories and cultures. There’s little spontaneity to the proceedings. The words tend to be as interchangeable as the CEOs’ ties.

But when, on October 9, 2006, Google purchased the fledgling Internet video network YouTube, tradition went out the window. Within hours of the deal’s completion, a two-minute video appeared on YouTube featuring the company’s twentysomething founders Chad Hurley and Steve Chen. Shot on a sidewalk with a handheld camcorder, the video had the cheap, do-it-yourself feel typical of the amateur productions uploaded to the site. For the first few seconds of the clip, the newly wealthy duo, who appeared to have gone without sleep for days and without sun for months, managed to keep

their euphoria under control as they searched for the right words to explain the acquisition to what they referred to, repeatedly, as “the YouTube community.”

“Hi, YouTube,” Hurley began. “This is Chad and Steve. We just want to say thank you. Today, we have some exciting news for you. We’ve been acquired by Google.”

“Yeah, thanks,” chimed in Chen, leaning nervously toward the camera. “Thanks to every one of you guys that have been contributing to YouTube, to the community. We wouldn’t be anywhere close to where we are without the help of this community.”

Struggling to keep a straight face, Hurley continued in a tone of strained seriousness: “We’re going to stay committed to developing the best service for you—you know, developing the most innovative service and tools and technologies so you can keep having fun on our site.”

But a minute into the video, all pretense of gravity gave way, and the announcement dissolved into a giddy slapstick routine, with the tall, gaunt Hurley playing Stan Laurel to the shorter, round-faced Chen’s Oliver Hardy.

“This is great,” Hurley said, breaking into a wide grin. “Two kings have gotten together, and we’re going to be able to provide you with an even better service.”

Chen broke into laughter at the mention of “two kings.” He ducked off camera, stumbling down the sidewalk.

“Two kings,” repeated Hurley, corralling his partner by grabbing him around the shoulders.

“Get your hand off me, king,” said Chen, still laughing.

In a vain attempt to get the founders to straighten up, the cameraman shouted out a question: “What does it mean for the user community?”

“Two kings have gotten together,” replied Hurley. “The king of

search and the king of video have gotten together. We’re going to have it our way. Salt and pepper.” Chen doubled over, and Hurley ended the video by drawing a finger across his throat. “We can’t do this,” he said. “Cut.”

The video proved a big hit on YouTube, shooting to the top of the site’s most-watched list. Within a month, it had been viewed 2 million times and had inspired a slew of parodies filmed by YouTube members and dutifully uploaded to the site. But though the subversive frivolity of Hurley and Chen’s acquisition announcement was remarkable in itself, behind it lay a much deeper break with the past. In the rise of YouTube we see a microcosm of the strange new world of online business. The company’s success reveals much about the changing economics of computing and how they’re affecting commerce, employment, and even the distribution of wealth.

CHAD HURLEY AND Steve Chen, together with a third friend, Yawad Kareem, had come up with the idea of launching an easy-to-use video-sharing service after a dinner party in early 2005. They chose the name YouTube during a brainstorming session on Valentine’s Day. Over the next few months, they designed and wrote the code for their site in the garage of the Silicon Valley house that Hurley had purchased with an earlier dotcom windfall. After successfully testing the service in May 2005 by broadcasting a video of Chen’s cat playing with a string, they received \$3.5 million in funding from a venture capital firm, enough to cover their modest startup costs. In December, the YouTube site officially opened for business, and it rapidly attracted an audience of people looking for a simple—and free—way to store, share, and view short, home-made videos (not to mention thousands of clips illegally copied from copyrighted films, TV shows, and music videos). It was just ten months later that Hurley and Chen sold the site to Google for

a staggering \$1.65 billion, bringing each of them a windfall worth about a third of a billion dollars.

At the time it was bought, YouTube had just sixty employees. They worked above a pizza joint in San Mateo, California, crammed into a single small office furnished with Ikea desks and chairs, a rubber chicken hanging from the bare metal rafters. (The company's servers were maintained off-site, in data centers run by utility hosting firms.) Despite its minuscule size and lack of formal organization, the staff was able to run one of the most popular and fastest-growing sites on the Internet. Every day, people all over the world watched more than 100 million YouTube video clips and uploaded some 65,000 new videos to the site. And those numbers were growing at an exponential pace, which explained why Google was willing to spend so much to buy the young company. At the \$1.65 billion sale price, each YouTube employee represented \$27.5 million in market value. Compare that to a traditional, and fabulously profitable, software company like Microsoft, which has 70,000 employees, each representing \$4 million in market value. Or compare it to a traditional media and entertainment company like Walt Disney, with 133,000 employees, each representing \$500,000 in value.

The abundance of cheap processing power, storage capacity, and communication bandwidth is what made it possible for YouTube to build a very large business very quickly with very few people. And YouTube's experience is far from unique. Many new companies are using the utility computing grid to create burgeoning enterprises with hardly any employees. A year before Google bought YouTube, another Internet giant, eBay, purchased the Internet telephone company Skype for \$2.1 billion. Founded just two years earlier by a pair of Scandinavian entrepreneurs, Skype had signed up 53 million customers—more than twice the number of phone customers served by venerable British Telecom—and was attracting 150,000

new subscribers every day. Morgan Stanley said the company's service "may be the fastest-growing product in history." Yet Skype employed just 200 people, about 90,000 fewer than British Telecom employed in the United Kingdom alone. A columnist for a business newspaper in Estonia, where Skype has its largest office, noted that, based on the company's sale price, "one Skype employee is worth more than the Viisnurk wood-processing company and 1.5 employees equal the entire value of the Kalev candy factory."

The classified ad site Craigslist is another example. It was created in 1995 by a software designer named Craig Newmark as an online bulletin board where people could post information about upcoming events in and around San Francisco. After Newmark incorporated his company in 1999, the site expanded rapidly. By the end of 2006, it had bulletin boards for more than 300 cities around the world, each offering a broad mix of for-sale, help-wanted, and personal ads, as well as notices of events and other local activities. More than 10 million visitors were looking at some 5 billion pages on the site every month, making Craigslist one of the most popular destinations on the Web. Yet the entire operation was being run by just twenty-two people.

Perhaps most remarkable of all is PlentyOfFish, an online service that helps people find dates. Launched in Canada in 2003, the site experienced explosive growth. By late 2006, some 300,000 people were logging on to the service every day, and they were looking at about 600 million pages a month. PlentyOfFish had become the largest dating site in Canada and one of the ten biggest in both the United States and the United Kingdom. So how many people does this booming business employ? Precisely one: its founder, Markus Frind. Frind designed and wrote all the code for the site, and he uses Google's automated AdSense service to place advertisements on it, reportedly earning the twenty-eight-year-old entrepreneur as

much as \$10,000 a day. Entirely computerized, the operation essentially runs itself. In a posting on his blog in June 2006, Frind wrote, "It amazes me to think that my 1 person company is able to compete at a level where all the competitors have . . . 300+ full time staff. I feel like I am ahead of my time, and when I look around at the companies out there today I have a hard time seeing them existing in a few years."

Companies like YouTube, Skype, Craigslist, and PlentyOfFish can grow so quickly with so few workers because their businesses are constructed almost entirely out of software code. Their products are virtual, residing in computer databases or flying across the Internet as strings of digits. The price of creating a fresh copy of their product and distributing it to a new customer anywhere in the world is essentially zero, so the companies can expand without hiring many additional employees. By relying on the public Internet as their distribution channel, moreover, they can avoid many of the capital investments traditional companies have to make. YouTube doesn't have to build production studios or transmission towers. Skype doesn't have to string miles of cable between telephone poles. Craigslist doesn't have to buy printing presses, ink, and paper. PlentyOfFish doesn't have to open offices. Although they have to pay modest fees for the bandwidth they use, they get something of a free ride on the fiber-optic cables that others paid for during the dotcom boom.

All of these businesses demonstrate an unusual sort of economic behavior that economists call "increasing returns to scale." What it means, simply, is that the more products they sell, the more profitable they become. That's a very different dynamic from the one that prevails in the industrial world, where businesses are subject to diminishing returns to scale. As a producer of physical goods increases its output, it sooner or later has to begin paying more for

its inputs—for the raw materials, components, supplies, real estate, and workers that it needs to make and sell its products. It can offset these higher input costs by achieving economies of scale—by using fewer inputs to make each additional product—but eventually the higher costs overwhelm the scale economies, and the company's profits, or returns, begin to shrink. The law of diminishing returns in effect sets limits on the size of companies, or at least on the size of their profits.

Until recently, most information goods were also subject to diminishing returns because they had to be distributed in physical form. Words had to be printed on paper, moving pictures had to be captured on film, software code had to be etched onto disks. But because the Internet frees information goods from their physical form, turning them into entirely intangible strings of ones and zeroes, it also frees them from the law of diminishing returns. A digital good can be replicated endlessly for essentially no cost—its producer does not have to increase its purchases of inputs as its business expands. Moreover, through a phenomenon called the network effect, digital goods often become more valuable as more people use them. Every new member that signs up for Skype, puts an ad on Craigslist, or posts a profile on PlentyOfFish increases the value of the service to every other member. Returns keep growing as sales or use expands—without limit.

The unique economics of running a business over the computing grid, combined with the global reach of Web sites, allow Internet-based companies to pursue a business strategy that would have been unthinkable just a few years ago: they can give their core products away for free. YouTube charges its users nothing to store or watch a video; it makes its money through advertising and sponsorships. PlentyOfFish also depends on ads for its income, letting people post and view profiles without charge. Skype allows subscribers to make

unlimited phone calls to other subscribers over the Internet—gratis—and charges just a few cents a minute for calls over traditional telephone lines. Craigslist's Newmark doesn't even seem to be interested in having his company make a lot of money. He imposes fees on only a few types of ads—commercial real-estate and job listings, mainly—and gives everything else away as a community service.

THE NEW ECONOMICS of doing business online are a boon to consumers. What used to be expensive—everything from international phone calls to classified ads to video transmissions—can now be had for free. But there's another side to the robotic efficiency, global scope, and increasing returns enjoyed by the new Internet companies. Each of these companies competes, after all, with old-line firms that have long employed and paid decent wages to many people. YouTube fights for viewers with television stations and movie theaters. Skype battles landline and mobile telephone companies for callers. Many of the classified ads that run on Craigslist would have otherwise appeared in local newspapers. Dozens of national and local dating services vie for customers with sites like PlentyOfFish. Given the economic advantages of online firms—advantages that will grow as the maturation of utility computing drives the costs of data processing and communication even lower—traditional firms may have no choice but to refashion their own businesses along similar lines, firing many millions of employees in the process.

We already see signs of the thinning out of the professional workforce in some information industries. As newspapers have lost readers and advertising revenues to Web sites, for instance, they've been forced to lay off reporters and other professionals. A study by the American Society of Newspaper Editors found that between 2001 and 2005 the newsroom staff of US papers declined by 4 percent, with a net loss of 1,000 reporters, 1,000 editors, and 300 pho-

tographers and artists. "Web 2.0 and the Net in general have been disasters for my profession," says Philip Dawdy, an award-winning journalist for *Seattle Weekly*. "Newspapers are dying. Talented people are being forced into public relations work."

In early 2007, the US Department of Labor released a revealing set of statistics on the publishing and broadcasting business as a whole. Employment in the industry had fallen by 13 percent in the six years since 2001, with nearly 150,000 jobs lost. These were years when many media companies, as well as their customers and advertisers, had been shifting from physical media to the Internet. Yet the report revealed that there had been no growth in Internet publishing and broadcasting jobs. In fact, online employment had actually dropped 29 percent, from 51,500 to 36,600, during the period. "The Internet is the wave of the future," commented *New York Times* economics writer Floyd Norris. "Just don't try to get a job in it."

Markus Frind's prediction that many traditional businesses will find it impossible to compete against extraordinarily lean Web operations may be coming to pass. As a result, we could well see a hollowing-out of many sectors of the economy, as computers and software displace workers on a large scale. Anyone employed by a business whose product or service can be distributed in digital form may be at risk, and the number of companies that fit that category is growing every year.

Of course, the displacement of workers by computers is nothing new, and it would normally be welcomed as a sign of a healthy economy. Improving labor productivity is what attracted companies to computers in the first place, after all. Computerization, like electrification before it, simply continues the centuries-long trend of substituting machines for workers. As three scholars, David Autor, Frank Levy, and Richard Murnane, explain in an influential article

in the *Quarterly Journal of Economics*, “Substitution of machinery for repetitive human labor has been a thrust of technological change throughout the Industrial Revolution. By increasing the feasibility of machine substitution for repetitive human tasks, computerization furthers—and perhaps accelerates—this long-prevailing trend.” But computerization, they also point out, brings “a qualitative enlargement in the set of tasks that machines can perform. Because computers can perform symbolic processing—storing, retrieving, and acting upon information—they augment or supplant human cognition in a large set of information-processing tasks that historically were not amenable to mechanization.” Computerization extends the replacement of workers by machines from the blue-collar to the white-collar world.

Whereas industrialization in general and electrification in particular created many new office jobs even as they made factories more efficient, computerization is not creating a broad new class of jobs to take the place of those it destroys. As Autor, Levy, and Mur-nane write, computerization “marks an important reversal. Previous generations of high technology capital sharply increased demand for human input of routine information-processing tasks, as seen in the rapid rise of the clerking occupation in the nineteenth century. Like these technologies, computerization augments demand for clerical and information-processing tasks. But in contrast to [its] predecessors, it permits these tasks to be automated.” Computerization creates new work, but it’s work that can be done by machines. People aren’t necessary.

That doesn’t mean that computers can take over all the jobs traditionally done by white-collar workers. As the scholars note, “Tasks demanding flexibility, creativity, generalized problem-solving and complex communications—what we call nonroutine cognitive tasks—do not (yet) lend themselves to computerization.” That par-

enthetical “yet,” though, should give us pause. As the power and usefulness of networked computers have advanced during the few years since they wrote their paper, we’ve seen not only the expansion of software’s capabilities but the flowering of a new phenomenon that is further reducing companies’ need for workers. Commonly termed “social production,” the phenomenon is reshaping the economics of the media, entertainment, and software industries, among others. In essence, it allows many of those “nonroutine cognitive tasks” that require “flexibility, creativity, generalized problem-solving and complex communications” to be carried out for free—not by computers on the network but by people on the network.

Look more closely at YouTube. It doesn’t pay a cent for the hundreds of thousands of videos it broadcasts. All the production costs are shouldered by the users of the service. They’re the directors, producers, writers, and actors, and by uploading their work to the YouTube site they’re in effect donating their labor to the company. Such contributions of “user-generated content,” as it’s called, have become commonplace on the Internet, and they’re providing the raw material for many Web businesses. Millions of people freely share their words and ideas through blogs and blog comments, which are often collected and syndicated by corporations. The contributors to open-source software projects, too, donate their labor, even though the products of their work are often commercialized by for-profit companies like IBM, Red Hat, and Oracle. The popular online encyclopedia Wikipedia is written and edited by volunteers. Yelp, a group of city sites, relies on reviews of restaurants, shops, and other local attractions contributed by members. The news agency Reuters syndicates photos and videos submitted by amateurs, some of whom are paid a small fee but most of whom get nothing. Social networking sites like MySpace and Facebook, and dating sites like PlentyOfFish, are essentially agglomerations of the creative, unpaid

contributions of their members. In a twist on the old agricultural practice of sharecropping, the site owners provide the digital real estate and tools, let the members do all the work, and then harvest the economic rewards.

The free labor is not limited to the production of creative works. Popular news-filtering sites like Digg and Reddit rank stories based on the votes of millions of members, obviating the need for editors. The photographs uploaded to Flickr are all sorted on the basis of labels, or “tags,” placed on them by the site’s users. Del.icio.us offers a similar service for tagging Web pages, and Freebase offers one for tagging information of all sorts. Much of the processing power required to run the Skype network is “borrowed” from the micro-processors inside users’ PCs, dramatically reducing the number of computers Skype has to buy. Even the search services provided by companies like Google and Yahoo are essentially constructed from the contributions of the makers and viewers of Web sites. The search firms don’t need to hire the analysts and librarians who have traditionally categorized and filtered information for the benefit of others. More and more companies are figuring out ways to harness the power of free labor. Even police forces are getting into the act. In 2006, Texas marshals set up webcams along the border with Mexico and began streaming the video feeds over the Internet. People all around the world can now watch for illegal immigrants, clicking a button to alert the police to any suspicious activity. It’s law enforcement on the cheap.

Why do people contribute their labor in this way? There are several reasons, none of them particularly surprising. In some cases, such as the building of search engines, they contribute without even knowing it. Companies like Google simply track people’s everyday behavior online and distill valuable intelligence from the patterns the behavior reveals. No one minds because the resulting products,

like search results, are useful. In other cases, people contribute out of their own self-interest. Creating a MySpace or a Facebook page provides a social benefit to many young people, helping them stay in touch with old friends and meet new ones. Tagging photos at Flickr or Web pages at Del.icio.us helps people keep track of words and images that interest them—it serves as a kind of personal filing system for online content. Some sites share a portion of their advertising revenues with contributors (though the sums are usually trivial). In still other cases, there’s a competitive or status-seeking element to the donations. Sites like Digg, Yelp, and even Wikipedia have hierarchies of contributors, and the more you contribute, the higher you rise in the hierarchy.

But the biggest reason people contribute to such sites is no different from the reason they pursue hobbies or donate their time to charitable causes or community groups: because they enjoy it. It gives them satisfaction. People naturally like to create things, to show off their creations to others, to talk about themselves and their families, and to be part of communal projects. It’s no different on the Internet. Even very early online communities and Web sites made extensive use of free labor. In the 1990s, so many thousands of America Online members were performing unpaid jobs for the company, such as moderating chat rooms, that *Wired* magazine called AOL a “cyber-sweatshop.” Much of Amazon.com’s early appeal came from the book reviews donated by customers—and the ratings of those reviews submitted by other customers. The uploading of videos, the writing of blogs, the debugging of open-source code, the editing of Wikipedia entries—all are simply new forms of the pastimes or charitable work that people have always engaged in outside their paid jobs.

What has changed, though, is the scope, scale, and sophistication of the contributions—and, equally important, the ability of compa-

nies to harness the free labor and turn it into valuable products and services. Ubiquitous, inexpensive computing and data communication, together with ever more advanced software programs, allow individuals to make and share creative works and other information goods in ways that were never possible before, and they also enable thousands or even millions of discrete contributions to be assembled into commercial goods with extraordinary efficiency. In his book *The Wealth of Networks*, Yale law professor Yochai Benkler traces the recent explosion in social production to three technological advances. "First, the physical machinery necessary to participate in information and cultural production is almost universally distributed in the population of the advanced economies," he writes. "Second, the primary raw materials in the information economy, unlike the physical economy, are [freely available] public goods—existing information, knowledge, and culture." Finally, the Internet provides a platform for distributed, modular production that "allows many diversely motivated people to act for a wide range of reasons that, in combination, cohere into new useful information, knowledge, and cultural goods."

All three factors will become even more salient in the years ahead. The World Wide Computer will continue to give individuals new production capabilities, to expand their access to information, and to make it easier to meld their contributions into useful and attractive products and services. Benkler sees a kind of popular revolution in the making, where the means of producing and distributing information goods, formerly controlled by big companies, are put into the hands of the masses. He believes the "networked information economy" marks "a significant inflection point for modern societies," which promises to bring "a quite basic transformation in how we perceive the world around us." By changing "the way we create and exchange information, knowledge, and culture," he

writes, "we can make the twenty-first century one that offers individuals greater autonomy, political communities greater democracy, and societies greater opportunities for cultural self-reflection and human connection."

Other writers and scholars have made arguments similar to Benkler's. They see a new and liberating economy emerging—a "gift economy" that, based on sharing rather than selling, exists outside of and even in opposition to the market economy. Although the term "gift economy" dates back at least a quarter century to Lewis Hyde's 1983 book *The Gift: Imagination and the Erotic Life of Property*, its new vogue highlights how extensive free labor and its products have become on the Internet. The gift economy, it's often said, is engendering a richer, more egalitarian culture, while drawing wealth and power away from the corporations and governments that have, allegedly, monopolized the distribution of creative works and other information goods. Richard Barbrook, of the University of Westminster in London, expressed this view well in his 1998 essay "The Hi-Tech Gift Economy." He wrote of Internet users: "Unrestricted by physical distance, they collaborate with each other without the direct mediation of money or politics. Unconcerned about copyright, they give and receive information without thought of payment. In the absence of states or markets to mediate social bonds, network communities are instead formed through the mutual obligations created by gifts of time and ideas."

There's truth in such claims, as anyone looking at the Web today can see. Computers and the Internet have given people powerful new tools for expressing themselves, for distributing their work to broad audiences, and for collaborating to produce various goods. But there's a naïveté, or at least a short-sightedness, to these arguments as well. The utopian rhetoric ignores the fact that the market economy is rapidly subsuming the gift economy. The "gifts of time and

ideas” are becoming inputs to the creation of commodities. Whether it’s a big company like Rupert Murdoch’s News Corporation, which owns MySpace, or a one-man operation like Markus Frind’s Plenty OfFish, businesses are using the masses of Internet gift-givers as a global pool of cut-rate labor.

When, in 2005, the Internet giant Yahoo acquired Flickr for a reported \$35 million, the larger company freely admitted that it was motivated by the prospect of harvesting all the free labor supplied by Flickr’s users. As Yahoo executive Bradley Horowitz told *Newsweek*, “With less than 10 people on the payroll, [Flickr] had millions of users generating content, millions of users organizing that content for them, tens of thousands of users distributing that across the Internet, and thousands of people not on the payroll actually building the thing. That’s a neat trick. If we could do that same thing with Yahoo, and take our half-billion user base and achieve the same kind of effect, we knew we were on to something.”

As user-generated content continues to be commercialized, it seems likely that the largest threat posed by social production won’t be to big corporations but to individual professionals—to the journalists, editors, photographers, researchers, analysts, librarians, and other information workers who can be replaced by, as Horowitz put it, “people not on the payroll.” Sion Touhig, a distinguished British photojournalist, points to the “glut of images freely or cheaply available on the Web” in arguing that “the Internet ‘economy’ has devastated my sector.” Why pay a professional to do something that an amateur is happy to do for free?

There have always been volunteers, of course, but unpaid workers are now able to replace paid workers on a scale far beyond what’s been possible before. Businesses have even come up with a buzzword for the phenomenon: “crowdsourcing.” By putting the means of production into the hands of the masses but withholding from those

masses any ownership over the products of their communal work, the World Wide Computer provides an incredibly efficient mechanism for harvesting the economic value of the labor provided by the very many and concentrating it in the hands of the very few. Chad Hurley and Steve Chen had good reason to thank the “YouTube community” so profusely when announcing the Google buyout. It was the members of that community who had, by donating their time and creativity to the site, made the two founders extremely rich young men.

THE CREATION OF the electric grid accelerated the concentration of wealth in large businesses, a trend that had been progressing, if at a slower pace, since the start of the Industrial Revolution. But as the big companies expanded and introduced new categories of consumer goods, they had to hire huge numbers of both skilled and unskilled workers and, following Henry Ford’s precedent, pay them good wages. In this way, electrification forced the companies to spread their increasing wealth widely among their employees. As we’ve seen, that played a decisive role in creating a prosperous middle class—and a more egalitarian American society.

The arrival of the universal computing grid portends a very different kind of economic realignment. Rather than concentrating wealth in the hands of a small number of companies, it may concentrate wealth in the hands of a small number of individuals, eroding the middle class and widening the divide between haves and have-nots. Once again, this would mark not the beginning of a new trend but rather the acceleration of an existing one.

Since the early 1980s, when businesses’ investments in computers began to skyrocket, American incomes have become increasingly skewed. As the incomes of the rich have leapt upward, most people’s wages have stagnated. In an extensive study of Internal Rev-

enue Service data, the economists Thomas Piketty and Emmanuel Saez found that the share of overall income that went to the richest 10 percent of households remained steady at around 32 percent between the end of the Second World War and 1980. At that point, it began to creep upward, reaching 34 percent in 1985, 39 percent in 1990, 41 percent in 1995, and 43 percent in 2000. After a brief drop following the stock market's retreat in the wake of the dotcom bust, it returned to 43 percent in 2004.

Even more dramatic, though, has been the increasing concentration of income among the wealthiest of the wealthy. The share of total income held by the richest 1 percent of Americans declined from about 11 percent in the late 1940s to approximately 8 percent in the late 1970s. But the trend reversed itself in the 1980s. By 2004, the top 1 percent were making 16 percent of the money. At the very top of the ladder, the concentration of income has been yet more striking. The share of income held by the top 0.1 percent of American families more than tripled between 1980 and 2004, jumping from 2 percent to 7 percent. When the researchers updated their statistics in late 2006, they saw a continuation of the trend, writing that preliminary tax data "suggests that top incomes have continued to surge in 2005."

In another recent study, researchers from Harvard University and the Federal Reserve examined the pay rates of the top three executives at the largest American corporations and compared them to the average compensation earned by workers in general. They found that executive pay remained fairly stable from the end of the Second World War through the 1970s but has shot up in more recent years. In 1970, according to the study, the median senior business executive earned 25 times more than the average US worker. By 2004, the median executive was receiving 104 times the average worker's pay—and the top 10 percent of executives were earning at

least 350 times the average. Ajay Kapur, an economist who serves as Citigroup's chief global equity strategist, argues that the United States, along with Canada and the United Kingdom, has become a "plutonomy" where "economic growth is powered by, and largely consumed by, the wealthy few."

Economists continue to debate the causes of the growing inequality in American incomes. Many factors are involved, including the strengthening of international trade, rising corporate profits, tax-rate changes, the declining power of trade unions, and changing social norms. But it seems increasingly clear that computerization has played a central role in the shift, particularly in holding down the incomes of the bulk of Americans. The distinguished Columbia University economist Jagdish Bhagwati argues that computerization is the main cause behind the two-decades-long stagnation of middle-class wages. "There are assembly lines today, but they are without workers," he writes; "they are managed by computers in a glass cage above, with highly skilled engineers in charge." Normally, the introduction of a new labor-saving technology would erode wages only briefly before the resulting boost in productivity pushed them up again. But information technology is different, Bhagwati contends. Unlike earlier technologies that caused "discrete changes," such as the steam engine, the ongoing advances in computer technology offer workers no respite. The displacement of workers "is continuous now," he says. "The pressure on wages becomes relentless."

In a February 2007 speech, Federal Reserve chairman Ben Bernanke backed up Bhagwati's conclusion that information technology is the primary force behind the skewing of incomes. He argued that "the influence of globalization on inequality has been moderate and almost surely less important than the effects of skill-biased technical change." With software automating ever more highly

skilled tasks, the number of people who find their jobs at risk is expanding—as newspaper reporters and editors will testify. The effect is amplified by companies' ability to shift knowledge work, the realm of well-paid white-collar workers, across national borders. Since everyone on the Internet has equal access to the World Wide Computer, a worker's location matters much less than it used to. Any job that can be done on a computer, no matter how specialized, has become portable. Even the reading of the X-rays of US patients is today being "offshored" to doctors in India and other countries. The international competition for jobs pushes salaries down in higher-wage countries like the United States, as the global labor market seeks equilibrium. Computerization hence puts many American wage-earners in a double bind: it reduces the demand for their jobs even as it expands the supply of workers ready and able to perform them.

There's a natural tendency, and a natural desire, to see the Internet as a leveling force, one that creates a fairer, more democratic society, where economic opportunities and rewards are spread widely among the many rather than held narrowly by the few. The fact that the World Wide Computer gives people new freedom to distribute their works and ideas to others, with none of the constraints of the physical world, appears to buttress that assumption. But the reality may be very different. In his 2006 book *The Long Tail*, Chris Anderson writes that "millions of ordinary people [now] have the tools and the role models to become amateur producers. Some of them will also have talent and vision. Because the means of production have spread so widely and to so many people, the talented and visionary ones, even if they're a small fraction of the total, are becoming a force to be reckoned with." This is not, as it might first appear, a vision of a world of economic egalitarianism. It's a vision of a world in which more and more of the wealth produced by mar-

kets is likely to be funneled to "a small fraction" of particularly talented individuals.

As we saw with electrification, the interplay of technological and economic forces rarely produces the results we at first expect. There are some who remain convinced that computerization, as it continues to gain momentum, will begin to close the wealth gap that up to now it has helped widen. That's the pattern that has occurred in past technological revolutions. But when we take into account the economic forces that the World Wide Computer is unleashing—the spread of the increasing-returns dynamic to more sectors of the economy, the replacement of skilled as well as unskilled workers with software, the global trade in knowledge work, and the ability of companies to aggregate volunteer labor and harvest its economic value—we're left with a prospect that is far from utopian. The erosion of the middle class may well accelerate, as the divide widens between a relatively small group of extraordinarily wealthy people—the digital elite—and a very large set of people who face eroding fortunes. In the YouTube economy, everyone is free to play, but only a few reap the rewards.

## The Great Unbundling

ELECTRIFICATION HASTENED the expansion of America's mass culture, giving people a shared set of experiences through popular television shows, radio programs, songs, movies, books and magazines, newspaper stories, and even advertisements. It opened up new nationwide channels for broadcast media and gave rise to large news and entertainment conglomerates that could afford the investments needed to make and distribute creative works. The advantages of physical scale and geographic reach enjoyed by manufacturers like Ford and General Electric also accrued to media companies like CBS, RCA, Metro-Goldwyn-Mayer, and Time Inc. And because the costs of producing mass media were so onerous, the companies had a strong incentive to pitch a small set of products to as large an audience as possible. In many cases, they had little choice but to restrict production. The airwaves could handle only so many TV and radio programs, shops could stock only so many books and records, and movie theaters could screen only so many films.

The nation's mass culture, and the sense of unity that it instilled in a motley population scattered across a vast land, was not, in other words, the expression of an essential quality of the American char-

acter. It was a by-product of the economic and technological forces that swept the country at the start of the twentieth century. The Internet, which is becoming not just a universal computer but also a universal medium, unleashes a very different set of forces, and they promise to reshape America's culture once again.

The major constraints on the supply of creative works—high costs and narrow distribution channels—are disappearing. Because most common cultural goods consist of words, images, or sounds, which all can be expressed in digital form, they are becoming as cheap to reproduce and distribute as any other information product. Many of them are also becoming easier to create, thanks to the software and storage services of the World Wide Computer and inexpensive production tools like camcorders, microphones, digital cameras, and scanners. Tasks that once required a lot of money and training, from film developing to video editing to graphic design to sound mixing, can now be done by amateurs in their dens, offices, and schoolrooms. The proliferation of blogs, podcasts, video clips, and MP3s testifies to the new economics of culture creation. And all the new digital products, whether fashioned by professionals or amateurs, can find a place in the online marketplace. The virtual shelves of the Internet can expand to accommodate everything.

The shift from scarcity to abundance in media means that, when it comes to deciding what to read, watch, and listen to, we have far more choices than our parents or grandparents did. We're able to indulge our personal tastes as never before, to design and wrap ourselves in our own private cultures. "Once the most popular fare defined our culture," explains Chris Anderson. "Now a million niches define our culture." The vast array of choices is exciting, and by providing an alternative to the often bland products of the mass media it seems liberating as well. It promises, as Anderson says,

to free us from "the tyranny of lowest-common-denominator fare" and establish in its place "a world of infinite variety."

But while it's true that the reduction in production and distribution costs is bringing us many more options, it would be a mistake to leap to the conclusion that nothing will be sacrificed in the process. More choices don't necessarily mean better choices. Many cultural goods remain expensive to create or require the painstaking work of talented professionals, and it's worth considering how the changing economics of media will affect them. Will these goods be able to find a large enough paying audience to underwrite their existence, or will they end up being crowded out of the marketplace by the proliferation of free, easily accessible products? Even though the Internet can in theory accommodate a nearly infinite variety of information goods, that doesn't mean that the market will be able to support all of them. Some of the most cherished creative works may not survive the transition to the Web's teeming bazaar.

THE TENSIONS CREATED by the new economics of production and consumption are visible today in many media, from music to movies. Nowhere, though, have they been so clearly on display, and so unsettling, as in the newspaper business. Long a mainstay of culture, print journalism is going through a wrenching transformation, and its future is in doubt. Over the past two decades, newspaper readership in the United States has plummeted. After peaking in 1984, at 63 million copies, the daily circulation of American papers fell steadily at a rate of about 1 percent a year until 2004, when it hit 55 million. Since then, the pace of the decline has accelerated. Circulation fell by more than 2 percent in 2005 and by about 3 percent in 2006. Many of the country's largest papers have been particularly hard hit. In just the six months between April and September of 2006, the daily circulation of the *Miami Herald* fell 8.8 percent; the

*Los Angeles Times*, 8.0 percent; the *Boston Globe*, 6.7 percent; the *New York Times*, 3.5 percent; and the *Washington Post*, 3.3 percent. In 1964, 81 percent of American adults read a daily newspaper. In 2006, only 50 percent did. The decline has been sharpest among young adults. Just 36 percent of 18- to 24-year-olds reported reading a daily newspaper in 2006, down from 73 percent in 1970.

There are many reasons for the long-term decline in newspaper readership. But one of the most important factors behind the recent acceleration of the trend is the easy availability of news reports and headlines on the Internet. As broadband connections have become more common, the number of American adults who get news online every day has jumped, from 19 million in March 2000 to 44 million in December 2005, according to the Pew Internet & American Life Project. The shift to online news sources is particularly strong among younger Americans. At the end of 2005, the Web had become a daily source of news for 46 percent of adults under thirty-six years of age who had broadband connections, while only 28 percent of that group reported reading a local newspaper.

The loss of readers means a loss of advertising revenue. As people continue to spend more time online, advertisers have been moving more of their spending to the Web, a trend expected to accelerate in coming years. From 2004 through 2007, newspapers lost an estimated \$890 million in ad revenues to the Internet, according to Citibank research. Classified advertising, long a lucrative niche for newspapers, has been particularly hard hit, as companies and homeowners shift to using sites like Craigslist, eBay, and Autotrader to sell cars and other used goods and to list their apartments and houses. In 2006, sales of classified ads by Web sites surpassed those of newspapers for the first time.

Newspaper companies are, naturally, following their readers and advertisers online. They're expanding their Web sites and shifting

ever more of their content onto them. After having kept their print and Web units separate for many years, dedicating most of their money and talent to print editions, papers have begun merging the operations, assigning more of their top editors' time to online editions. During 2006 and 2007, the *New York Times*, *Washington Post*, and *Wall Street Journal* all announced plans to give more emphasis to their Web sites. "For virtually every newspaper," says one industry analyst, "their only growth area is online." Statistics underscore the point. Visits to newspaper Web sites shot up 22 percent in 2006 alone.

But the nature of a newspaper, both as a medium for information and as a business, changes when it loses its physical form and shifts to the Internet. It gets read in a different way, and it makes money in a different way. A print newspaper provides an array of content—local stories, national and international reports, news analyses, editorials and opinion columns, photographs, sports scores, stock tables, TV listings, cartoons, and a variety of classified and display advertising—all bundled together into a single product. People subscribe to the bundle, or buy it at a newsstand, and advertisers pay to catch readers' eyes as they thumb through the pages. The publisher's goal is to make the entire package as attractive as possible to a broad set of readers and advertisers. The newspaper as a whole is what matters, and as a product it's worth more than the sum of its parts.

When a newspaper moves online, the bundle falls apart. Readers don't flip through a mix of stories, advertisements, and other bits of content. They go directly to a particular story that interests them, often ignoring everything else. In many cases, they bypass the newspaper's "front page" altogether, using search engines, feed readers, or headline aggregators like Google News, Digg, and Daylife to leap directly to an individual story. They may not even be aware of which newspaper's site they've arrived at. For the pub-

lisher, the newspaper as a whole becomes far less important. What matters are the parts. Each story becomes a separate product standing naked in the marketplace. It lives or dies on its own economic merits.

Because few newspapers, other than specialized ones like the *Wall Street Journal*, are able to charge anything for their online editions, the success of a story as a product is judged by the advertising revenue it generates. Advertisers no longer have to pay to appear in a bundle. Using sophisticated ad placement services like Google AdWords or Yahoo Search Marketing, they can target their ads to the subject matter of an individual story or even to the particular readers it attracts, and they only pay the publisher a fee when a reader views an ad or, as is increasingly the case, clicks on it. Each ad, moreover, carries a different price, depending on how valuable a viewing or a clickthrough is to the advertiser. A pharmaceutical company will pay a lot for every click on an ad for a new drug, for instance, because every new customer it attracts will generate a lot of sales. Since all page views and ad clickthroughs are meticulously tracked, the publisher knows precisely how many times each ad is seen, how many times it is clicked, and the revenue that each view or clickthrough produces.

The most successful articles, in economic terms, are the ones that not only draw a lot of readers but deal with subjects that attract high-priced ads. And the most successful of all are those that attract a lot of readers who are inclined to click on the high-priced ads. An article about new treatments for depression would, for instance, tend to be especially lucrative, since it would attract expensive drug ads and draw a large number of readers who are interested in new depression treatments and hence likely to click on ads for psychiatric drugs. Articles about saving for retirement or buying a new car or putting an addition onto a home would also tend to throw

off a large profit, for similar reasons. On the other hand, a long investigative article on government corruption or the resurgence of malaria in Africa would be much less likely to produce substantial ad revenues. Even if it attracts a lot of readers, a long shot in itself, it doesn't cover a subject that advertisers want to be associated with or that would produce a lot of valuable clickthroughs. In general, articles on serious and complex subjects, from politics to wars to international affairs, will fail to generate attractive ad revenues.

Such hard journalism also tends to be expensive to produce. A publisher has to assign talented journalists to a long-term reporting effort, which may or may not end in a story, and has to pay their salaries and benefits during that time. The publisher may also have to shell out for a lot of expensive flights and hotel stays, or even set up an overseas bureau. When bundled into a print edition, hard journalism can add considerably to the overall value of a newspaper. Not least, it can raise the prestige of the paper, making it more attractive to subscribers and advertisers. Online, however, most hard journalism becomes difficult to justify economically. Getting a freelance writer to dash off a review of high-definition television sets—or, better yet, getting readers to contribute their own reviews for free—would produce much more attractive returns.

In a 2005 interview, a reporter for the *Rocky Mountain News* asked Craig Newmark what he'd do if he ran a newspaper that was losing its classified ads to sites like Craigslist. "I'd be moving to the Web faster," he replied, and "hiring more investigative journalists." It's a happy thought, but it ignores the economics of online publishing. As soon as a newspaper is unbundled, an intricate and, until now, largely invisible system of subsidization quickly unravels. Classified ads, for instance, can no longer help to underwrite the salaries of investigative journalists or overseas correspondents. Each piece of content has to compete separately, consuming costs

and generating revenues in isolation. So if you're a beleaguered publisher, losing readers and money and facing Wall Street's wrath, what are you going to do as you shift your content online? Hire more investigative journalists? Or publish more articles about consumer electronics? It seems clear that as newspapers adapt to the economics of the Web, they are far more likely to continue to fire reporters than hire new ones.

Speaking before the Online Publishing Association in 2006, the head of the *New York Times's* Web operation, Martin Nisenholtz, summed up the dilemma facing newspapers today. He asked the audience a simple question: "How do we create high-quality content in a world where advertisers want to pay by the click, and consumers don't want to pay at all?" The answer may turn out to be equally simple: we don't. At least one major newspaper, *The Times* of London, admits that it has already begun training its reporters to craft their stories in ways that lead to higher placements in search engines. Jim Warren, the *Chicago Tribune's* managing editor, says that "you can't really avoid the fact that page views are increasingly the coin of the realm." As long as algorithms determine the distribution of profits, they will also determine what gets published.

The unbundling of content is not unique to newspapers or other print publications. It's a common feature of most online media. Apple's iTunes store has unbundled music, making it easy to buy by the song rather than the album. Digital video recorders like TiVo and pay-per-view cable services are unbundling television, separating the program from the network and its schedule. Video sites like YouTube go even further, letting viewers watch brief clips rather than sitting through entire shows. Amazon.com has announced plans to unbundle books, selling them by the page. Google provides "snippets" of text from published works through its controversial

Book Search service. Podcasting is unbundling radio programs. Wikipedia is unbundling the encyclopedia. The "bundling of the world's computers into a single network," writes Daniel Akst, "is ushering in what may be called the unbundled age."

Economists are quick to applaud the breaking up of media products into their component pieces. In their view, it's how markets should work. Consumers should be able to buy precisely what they want without having to "waste" money on what they don't. The *Wall Street Journal* celebrates the development, saying it heralds a new era in which we'll no longer have "to pay for detritus to get the good stuff." That's true in many cases, but not in all. Creative works are not like other consumer goods, and the economic efficiency that would be welcomed in most markets may have less salutary effects when applied to the building blocks of culture. It's worth remembering, as well, that the Internet is a very unusual marketplace, where information of all sorts tends to be given away and money is made through indirect means like advertising. Once you fragment both the audience and the advertising in such a market, large investments in the production of certain creative works become much harder for businesses to justify.

If the news business is any indication, the "detritus" that ends up being culled from our culture may include products that many of us would define as "the good stuff." What's sacrificed may not be blandness but quality. We may find that the culture of abundance being produced by the World Wide Computer is really just a culture of mediocrity—many miles wide but only a fraction of an inch deep.

IN 1971, THE economist Thomas Schelling performed a simple experiment that had a very surprising result. He was curious about the persistence of extreme racial segregation in the country. He knew that most Americans are not racists or bigots, that we're generally

happy to be around people who don't look or think the same way we do. At the same time, he knew that we're not entirely unbiased in the choices we make about where we live and whom we associate with. Most of us have a preference, if only a slight one, to be near at least some people who are similar to ourselves. We don't want to be the only black person or white person, or the only liberal or conservative, on the block. Schelling wondered whether such small biases might, over the long run, influence the makeup of neighborhoods.

He began his experiment by drawing a grid of squares on a piece of paper, creating a pattern resembling an oversized checkerboard. Each square represented a house lot. He then randomly placed a black or a white marker in some of the squares. Each marker represented either a black or a white family. Schelling assumed that each family desired to live in a racially mixed neighborhood, and that's exactly what his grid showed at the start: the white families and the black families were spread across the grid in an entirely arbitrary fashion. It was a fully integrated community. He then made a further assumption: that each family would prefer to have some nearby neighbors of the same color as themselves. If the percentage of neighbors of the same color fell beneath 50 percent, a family would have a tendency to move to a new house.

On the basis of that one simple rule, Schelling began shifting the markers around the grid. If a black marker's neighbors were more than 50 percent white or if a white marker's neighbors were more than 50 percent black, he'd move the marker to the closest unoccupied square. He continued moving the pieces until no marker had neighbors that were more than 50 percent of the other color. At that point, to Schelling's astonishment, the grid had become completely segregated. All the white markers had congregated in one area, and all the black markers had congregated in another. A modest, natural preference to live near at least a few people sharing a simi-

lar characteristic had the effect, as it influenced many individual decisions, of producing a sharp divide in the population. "In some cases," Schelling explained, "small incentives, almost imperceptible differentials, can lead to strikingly polarized results."

It was a profound insight, one that, years later, would be cited by the Royal Swedish Society of Sciences when it presented Schelling with the 2005 Nobel Prize in Economics. Mark Buchanan, in his book *Nexus*, summarized the broader lesson of the experiment well: "Social realities are fashioned not only by the desires of people but also by the action of blind and more or less mechanical forces—in this case forces that can amplify slight and seemingly harmless personal preferences into dramatic and troubling consequences."

Just as it's assumed that the Internet will promote a rich and diverse culture, it's also assumed that it will bring people into greater harmony, that it will breed greater understanding and help ameliorate political and social tensions. On the face of it, that expectation seems entirely reasonable. After all, the Internet erases the physical boundaries that separate us, allows the free exchange of information about the thoughts and lives of others, and provides an egalitarian forum in which all views can get an airing. The optimistic view was perhaps best expressed by Nicholas Negroponte, the head of MIT's Media Lab, in his 1995 bestseller *Being Digital*. "While the politicians struggle with the baggage of history, a new generation is emerging from the digital landscape free of many of the old prejudices," he wrote. "Digital technology can be a natural force drawing people into greater world harmony."

But Schelling's simple experiment calls this view into question. Not only will the process of polarization tend to play out in virtual communities in the same way it does in neighborhoods, but it seems likely to proceed much more quickly online. In the real world, with its mortgages and schools and jobs, the mechanical forces of seg-

regation move slowly. There are brakes on the speed with which we pull up stakes and move to a new house. Internet communities have no such constraints. Making a community-defining decision is as simple as clicking a link. Every time we subscribe to a blog, add a friend to our social network, categorize an email message as spam, or even choose a site from a list of search results, we are making a decision that defines, in some small way, whom we associate with and what information we pay attention to. Given the presence of even a slight bias to be connected with people similar to ourselves—ones who share, say, our political views or our cultural preferences—we would, like Schelling's hypothetical homeowners, end up in ever more polarized and homogeneous communities. We would click our way to a fractured society.

Greatly amplifying the polarization effect are the personalization algorithms and filters that are so common on the Internet and that often work without our permission or even our knowledge. Every time we buy a book at Amazon or rent a movie from Netflix or view a news story at Reddit, the site stores information about our choice in a personal profile and uses it to recommend similar products or stories in the future. The effect, in the short run, can be to expose us to items we wouldn't otherwise have come across. But over the long run, the more we click, the more we tend to narrow the information we see.

As the dominant search engine, Google wields enormous influence over the information people find on the Web, and it has been particularly aggressive in engineering the personalization of content. In 2005, it began testing a personalized search service that "orders your search results based on your past searches, as well as the search results and news headlines you've clicked on." In 2007, it quietly made personalized search the default setting for anyone with a Gmail address or other Google account. (The company's top

three competitors, Yahoo, Microsoft, and Ask, also have personalized search tools in the works.) Google scientists have even developed an "audio-fingerprinting" system that can use the microphone in your computer to monitor the "ambient audio" in your room and use it for personalization purposes. If you have your television on, the system can identify the program you're watching by recording a sample of its audio signal and comparing it to an "audio database" stored in a Google data center. The company could then feed you stories or ads keyed to your favorite shows.

Google has said that its goal is to store "100% of a user's data" inside its utility, allowing it to achieve what it calls "transparent personalization." At that point, it would be able to automatically choose which information to show you, and which to withhold, without having to wait for you to ask. It says, for example, that people "should not have to tell us which [news] feeds they want to subscribe to. We should be able to determine this implicitly."

A company run by mathematicians and engineers, Google seems oblivious to the possible social costs of transparent personalization. Its interest, as its CEO has said, lies in "using technology to solve problems that have never been solved before"—and personalization is just one of those problems. But, of course, Google and its competitors are not imposing personalization on us against our will. They're just responding to our desires. We welcome personalization tools and algorithms because they let us get precisely what we want when we want it, with a minimum of fuss. By filtering out "the detritus" and delivering only "the good stuff," they allow us to combine fragments of unbundled information into new bundles, tailor-made for audiences of one. They impose homogeneity on the Internet's wild heterogeneity. As the tools and algorithms become more sophisticated and our online profiles more refined, the Internet will act increasingly as an incredibly sensitive feedback

loop, constantly playing back to us, in amplified form, our existing preferences.

In “Global Village or Cyber-Balkans?,” an article that appeared in the journal *Management Science* in 2005, Eric Brynjolfsson, of MIT, and Marshall Van Alstyne, of Boston University, describe the results of a mathematical model they constructed to measure how individuals’ choices influence the makeup of online communities. “Although the conventional wisdom has stressed the integrating effects of [Internet] technology,” they write, in introducing their research, “we examine critically the claim that a global village is the inexorable result of increased connectivity.”

They note that, because there are limits to how much information we can process and how many people we can communicate with (we have “bounded rationality,” to use the academic jargon), we naturally have a strong desire to use filters to screen the ideas we’re exposed to and the people we associate with. As the filters become more finely tuned, we can focus our attention—and structure our communities—with ever greater precision. Schelling’s work shows that this process naturally breeds homogeneity in the real world, and Brynjolfsson and Van Alstyne’s model confirms that the effects could be even more extreme in the virtual world. “Our analysis,” they write, “suggests that automatic search tools and filters that route communications among people based on their views, reputations, past statements or personal characteristics are not necessarily benign in their effects.” Shaped by such tools, online communities could actually end up being less diverse than communities defined by physical proximity. Diversity in the physical world “can give way to virtual homogeneity as specialized communities coalesce across geographic boundaries.”

They stress that such “balkanization” is not the only possible result of filtering. In theory, “preferences for broader knowledge,

or even randomized information, can also be indulged.” In reality, though, our slight bias in favor of similarity over dissimilarity is difficult, if not impossible, to eradicate. It’s part of human nature. Not surprisingly, then, Brynjolfsson and Van Alstyne report that their model indicates, in a direct echo of Schelling’s findings, that “other factors being equal, all that is required to reduce integration in most cases is that preferred interactions are more focused than existing interactions.” If, in other words, we have even a small inclination to prefer like-minded views and people—to be more “focused” rather than more inclusive—we will tend to create ever more polarized communities online.

We see considerable evidence of such schisms today, particularly in the so-called blogosphere. Political blogs have divided into two clearly defined and increasingly polarized camps: the liberals and the conservatives. In 2005, two researchers, Lada Adamic, of Hewlett-Packard Labs, and Natalie Glance, of Infoseek Applied Research, published the results of an extensive study of political blogs, which they titled “Divided They Blog.” They looked at the patterns of linking among the forty most popular political blogs during the two months leading up to the 2004 US presidential election, and they also examined the activity of a much broader set of political blogs—more than 1,000 in all—on one day during that period. They discovered a sharp and “unmistakable” division between the conservative and liberal camps. “In fact,” they wrote, “91% of the links originating within either the conservative or liberal communit[y] stay within that community.” In addition, the two groups “have different lists of favorite news sources, people, and topics to discuss,” with only occasional overlaps.

Another study of the political blogosphere, by Matthew Hindman, a political scientist at Arizona State University, found a similar polarization. Rather than examining the links contained in the

blogs, Hindman looked at the actual traffic flows between them. He found that the vast majority of readers tend to stay within the bounds of either the liberal or the conservative sphere. Liberals listen almost exclusively to other liberals, and conservatives listen almost exclusively to other conservatives. "Only a handful of sites," he reports, "share traffic with those on the opposite end of the political spectrum," and the small amount of interaction that does take place between the sides is dominated by what Hindman terms "name calling." His conclusion: "There's not a whole lot of great news for democratic theory here."

DURING THE SUMMER of 2005, a group of researchers assembled sixty-three Coloradans to discuss three controversial issues: same-sex marriage, affirmative action, and global warming. About half of the participants were conservatives from Colorado Springs, while the other half were liberals living in Boulder. After the participants completed, in private, questionnaires about their personal views on the three topics, they were split into ten groups—five conservative and five liberal. Each group then spent some time discussing the issues, with the goal of reaching a consensus on each one. After the discussion, the participants again filled out questionnaires.

The results of the study were striking. In every case, the deliberations among like-minded people produced what the researchers call "ideological amplification." People's views became more extreme and more entrenched:

First, the groups from Boulder became even more liberal on all three issues; the groups from Colorado Springs became even more conservative. Deliberation thus increased extremism. Second, every group showed increased consensus, and decreased diversity, in the attitudes of [its] members. . . . Third, deliberation sharply increased

the differences between the views of the largely liberal citizens of Boulder and the largely conservative citizens of Colorado Springs. Before deliberation began, there was considerable overlap between many individuals in the two different cities. After deliberation, the overlap was much smaller.

The study revealed a fact about human nature and group dynamics that psychologists have long recognized: the more that people converse or otherwise share information with other people who hold similar views, the more extreme their views become. As University of Chicago law professor Cass Sunstein, one of the organizers of the Colorado experiment, explains in his book *Infotopia*, "When like-minded people cluster, they often aggravate their biases, spreading falsehoods." They "end up in a more extreme position in line with their tendencies before deliberation began." This phenomenon, which Sunstein reports has been documented "in hundreds of studies in over a dozen countries," may in the worst cases plant "the roots of extremism and even fanaticism and terrorism."

Given how easy it is to find like-minded people and sympathetic ideas on the Internet and given our innate tendency to form homogeneous groups, we can see that "ideological amplification" is likely to be pervasive online. Here again, as Brynjolfsson and Van Alstyne note in their article, filtering and personalization technologies are likely to magnify the effect. "Individuals empowered to screen out material that does not conform to their existing preferences may form virtual cliques, insulate themselves from opposing points of view, and reinforce their biases," they write. "Indulging these preferences can have the perverse effect of intensifying and hardening pre-existing biases. . . . The effect is not merely a tendency for members to conform to the group average but a radicalization in which this average moves toward extremes."

In a further perverse twist, the very abundance of information available on the Internet may serve not to temper extremism but to amplify it further. As the Colorado study showed, whenever people find additional information that supports their existing views, they become more convinced that those views are right—and that people who hold different opinions are wrong. Each extra piece of confirming information heightens their confidence in the rectitude of their opinion and, as their confidence increases, their views tend also to become more extreme. They become single-minded.

Not only will the Internet tend to divide people with different views, in other words, it will also tend to magnify the differences. As Brynjolfsson and Van Alstyne suggest, this could in the long run pose a threat to the spirit of compromise and the practice of consensus-building that are at the heart of democratic government. “Internet users can seek out interactions with like-minded individuals who have similar values and thus become less likely to trust important decisions to people whose values differ from their own,” they conclude. Although they stress that it’s too early to know exactly how all of these forces will play out, they warn that “balkanization and the loss of shared experiences and values may be harmful to the structure of democratic societies.”

THE INTERNET TURNS everything, from news-gathering to community-building, into a series of tiny transactions—expressed mainly through clicks on links—that are simple in isolation yet extraordinarily complicated in the aggregate. Each of us may make hundreds or even thousands of clicks a day, some deliberately, some impulsively, and with each one we are constructing our identity, shaping our influences, and creating our communities. As we spend more time and do more things online, our combined clicks will shape our economy, our culture, and our society.

We’re still a long way from knowing where our clicks will lead us. But it’s clear that two of the hopes most dear to the Internet optimists—that the Web will create a more bountiful culture and that it will promote greater harmony and understanding—should be treated with skepticism. Cultural impoverishment and social fragmentation seem equally likely outcomes.

States's relatively free and open model to China's model of political control." Seemingly arcane technical standards, originally intended to create unity, will become the new terms of ideological debate and geopolitical struggle. The result, argue Goldsmith and Wu, will be "a technological version of the cold war, with each side pushing its own vision of the Internet's future."

Whether we're on the verge of Cold War 2.0 or not, the years ahead promise to be perilous ones as states and citizens struggle to come to grips with the manifold ramifications of the universal computing grid. As the Venezuelan scholar Carlota Perez has shown, governments tend to be very slow to respond to technological revolutions. Even as entrepreneurs and financiers, not to mention criminals and other bad actors, rush to exploit commercial and political disruptions, politicians, judges, and bureaucrats remain locked in the past, pursuing old policies and relying on outdated legal and regulatory schemes. The inertia magnifies the social and economic uncertainty and upheaval. In the worst cases, it lasts for decades, exacting, as Perez puts it, "a very high cost in human suffering."

## A Spider's Web

**W**HO IS 4417749? That was the question two *New York Times* reporters, Michael Barbaro and Tom Zeller Jr., set out to answer on August 7, 2006. In late July, AOL had released through its Web site a report detailing the keywords entered into its search engine by 657,000 of its subscribers over a three-month period earlier in the year. The company, a unit of media conglomerate Time Warner, thought it was performing a public service by releasing the search logs. It knew the information would be valuable to academic and corporate researchers studying the behavior of Web surfers or trying to invent new search technologies. To protect subscribers' privacy, AOL had carefully "anonymized" the data, replacing people's names with numbers and stripping out other identifying information. Said one Stanford computer science professor, "Having the AOL data available is a great boon for research."

But others wondered whether the data was really as anonymous as it seemed. Could the identities of the subscribers be inferred simply by examining what they searched for? Barbaro and Zeller, along with their editor, David Gallagher, decided to find out. They took a close look at one set of keywords, those entered by an AOL

subscriber known only as "4417749." The terms were a mishmash, ranging from "swing sets" to "single dances in Atlanta" to "dog who urinates on everything" to "school supplies for Iraq children." They formed what the reporters called "a catalog of intentions, curiosity, anxieties and quotidian questions." But there were enough clues in that catalog for Barbaro, Zeller, and Gallagher to track down the name, address, and phone number of the searcher. The search took only "a couple of hours," according to Gallagher. Number 4417749 turned out to be Thelma Arnold, a sixty-two-year-old widow living in Lilburn, Georgia. On August 9, Arnold woke up to find her name and picture on the front page of the national edition of the *Times*.

She was shocked to discover that her searches had been monitored by AOL, each keyword meticulously collected and connected to her account. "My goodness, it's my whole personal life," she told the reporters. "I had no idea somebody was looking over my shoulder." But however embarrassing she found the disclosure, Arnold had some cause to be relieved. The terms she had searched for were innocuous. Other subscribers had divulged much more intimate information about themselves. Subscriber 11574916 searched for "cocaine in urine" and "florida dui laws." Subscriber 1515830 searched for "how to tell your family you're a victim of incest" and "can you adopt after a suicide attempt." Subscriber 59920 searched for "what a neck looks like after its been strangled" and "rope to use to hog tie someone." Along with the quotidian came the peculiar and the perverse.

Like Thelma Arnold, most of us assume that we're anonymous when we go about our business online. We treat the Internet not just as a shopping mall and a library but as a personal diary and even a confessional. Through the sites we visit and the searches we make, we disclose details not only about our jobs, hobbies, families, politics, and health but also about our secrets, fantasies, obsessions,

peccadilloes, and even, in the most extreme cases, our crimes. But our sense of anonymity is largely an illusion. Detailed information about everything we do online is routinely gathered, stored in corporate or governmental databases, and connected to our real identities, either explicitly through our user names, our credit card numbers, and the IP addresses automatically assigned to our computers or implicitly through our searching and surfing histories. A famous 1993 *New Yorker* cartoon bore the caption "On the Internet, nobody knows you're a dog." In reality, not only is it known that you're a dog, but it's probably known what breed you are, your age, where you live, and the kind of treat you prefer.

Linking sensitive information to people's names doesn't require a team of *New York Times* reporters sifting through search logs and phone books. Nor does it require the inadvertent or intentional disclosure of data. As online databases proliferate and as analytical technologies advance, it becomes ever easier to use the World Wide Computer to "mine" personal information. A few months before AOL released its search logs, the writer Tom Owad provided a chilling lesson in just how easy it has become to extract private data from the Internet. Owad, who publishes a Web site for Macintosh users, wrote a simple piece of software—a "script"—to download the wish lists posted by Amazon.com customers. Millions of people maintain such lists on the online retailer's site, using them to catalog products that they plan to purchase in the future or that they'd like to receive as gifts. These lists can be searched by anyone, and they usually include the name of the list's owner and the city and state in which he lives.

Using two five-year-old PCs and a standard home Internet connection, Owad was able to download more than a quarter of a million wish lists over the course of a day. "I now had documents describing the reading preferences of 260,000 US citizens," he later wrote

on his site. Encoded into each list was the owner's unique Amazon customer identification number, allowing Owad to easily sort the lists by individual. He could then search the resulting database for book titles as well as various other keywords. He performed searches for several controversial or politically sensitive books and authors, from Kurt Vonnegut's *Slaughterhouse-Five* to the Koran, from the right-wing pundit Rush Limbaugh to his left-wing counterpart Michael Moore. Knowing the names and home cities of the list owners, he was then able to use Yahoo People Search to identify addresses and phone numbers for many of them. He took one final step and used Google Maps to plot their street addresses. He ended up with maps of the United States showing the precise locations of people interested in particular books and ideas. He posted on his site, for instance, a map of the homes of Amazon customers who had expressed an interest in George Orwell's *1984*. He could just as easily have published a map showing the residences of people interested in books about growing marijuana or giving up a child for adoption. "It used to be," Owad concluded, "you had to get a warrant to monitor a person or a group of people. Today, it is increasingly easy to monitor ideas. And then track them back to people."

Owad spent a fair amount of time organizing and conducting his information-gathering experiment. He had to write a custom code to download the data, and he had to manually perform his database searches. But what Owad did by hand can increasingly be performed automatically, with data-mining algorithms that draw information from many different sites simultaneously. One of the essential characteristics of the computing grid is the interconnection of diverse stores of information. The "openness" of databases is what gives the World Wide Computer much of its power. But it also makes it easy to discover hidden relationships among far-flung

bits of data. Analyzing those relationships can unlock a surprisingly large trove of confidential information about Web users.

At a 2006 conference of computer scientists, held in Seattle during the same week that Thelma Arnold's identity was revealed in the press, five scholars from the University of Minnesota presented a paper titled "You Are What You Say: Privacy Risks of Public Mentions." They described how software programs can be used to make connections among online databases. By uncovering overlaps in the data, the programs can often create detailed personal profiles of individuals—even when they submit information anonymously. The software is based on a simple and obvious principle: people tend to express their interests and discuss their opinions in many different places on the Internet. They may, for instance, buy an album at the iTunes Music Store, include that album on their playlist at Last.fm, rate it at the Rate Your Music site, and mention it in a comment on a music blog. Or they may edit a Wikipedia entry on their favorite actor, write a review of a new biography of him at Amazon, become his "friend" at MySpace, and tag pictures of him at Flickr. Sophisticated algorithms can identify such correspondences and use them to identify individuals with extraordinary precision, the Minnesota researchers discovered. In analyzing just two databases—one drawn from a movie ratings site, the other from a movie discussion forum—they found that an algorithm could successfully identify 60 percent of the people who mentioned eight or more films.

"In today's data-rich networked world," they explain in their paper, "people express many aspects of their lives online. It is common to segregate different aspects in different places: you might write opinionated rants about movies in your blog under a pseudonym while participating in a forum or Web site for scholarly discussion of medical ethics under your real name. However, it may

be possible to link these separate identities” using data-mining algorithms. Such an automated process of identification, they say, “creates serious privacy risks for users.” Even if people don’t divulge their real identities anywhere, their names can often be easily discovered if they disclose a very small number of identifying characteristics. The authors note, for example, that the vast majority of Americans can be identified by name and address using only their zip code, birthday, and gender—three pieces of information that people routinely divulge when they register for a user name at a Web site.

“You have zero privacy,” Scott McNealy, the former chief executive of Sun Microsystems, remarked back in 1999. “Get over it.” The idea that the loss of privacy is the price we have to pay for the handiness of the Internet is a common one, and there’s some truth to it. But few of us are aware of the extent to which we’ve disclosed details about our identities and lives or the way those details can be mined from search logs or other databases and linked back to us. And whether or not we’re comfortable with the possible compromise of our privacy, that’s far from the only or even the most disturbing threat posed by today’s Internet. As mathematicians and computer scientists continue to refine data-mining algorithms, they are uncovering new ways to predict how people will react when they’re presented with information or other stimuli online. They’re learning not just how to identify us but how to manipulate us, and their discoveries are being put to practical use by companies and governments, not to mention con artists and crooks.

It’s natural to think of the Internet as a technology of emancipation. It gives us unprecedented freedom to express ourselves, to share our ideas and passions, to find and collaborate with soul mates, and to discover information on almost any topic imaginable. For many people, going online feels like a passage into a new and

radically different kind of democratic state, one freed of the physical and social demarcations and constraints that can hobble us in the real world. The sense of the Web as personally “empowering,” to use the common buzzword, is almost universal, even among those who rue its commercialization or decry the crassness of much of its content. In early 2006, the editors of the Cato Institute’s online journal *Cato Unbound* published a special issue on the state of the Net. They reported that the “collection of visionaries” contributing to the issue appeared to be “in unanimous agreement that the Internet is, and will continue to be, a force for liberation.” In a July 2007 essay, the media scholar Clay Shirky wrote, “The internet’s output is data, but its product is freedom, lots and lots of freedom.” David Weinberger, in his book *Small Pieces Loosely Joined*, summed up the Internet’s liberation mythology in simple terms: “The Web is a world we’ve made for one another.”

It’s a stirring thought, but like most myths it’s at best a half-truth and at worst a fantasy. Computer systems in general and the Internet in particular put enormous power into the hands of individuals, but they put even greater power into the hands of companies, governments, and other institutions whose business it is to control individuals. Computer systems are not at their core technologies of emancipation. They are technologies of control. They were designed as tools for monitoring and influencing human behavior, for controlling what people do and how they do it. As we spend more time online, filling databases with the details of our lives and desires, software programs will grow ever more capable of discovering and exploiting subtle patterns in our behavior. The people or organizations using the programs will be able to discern what we want, what motivates us, and how we’re likely to react to various stimuli. They will, to use a cliché that happens in this case to be true, know more about us than we know about ourselves.

Even as the World Wide Computer grants us new opportunities and tools for self-expression and self-fulfillment, it is also giving others an unprecedented ability to influence how we think and what we do, to funnel our attention and actions toward their own ends. The technology's ultimate social and personal consequences will be determined in large measure by how the tension between the two sides of its nature—liberating and controlling—comes to be resolved.

ALL LIVING SYSTEMS, from amoebas to nation-states, sustain themselves through the processing of matter, energy, and information. They take in materials from their surroundings, and they use energy to transform those materials into various useful substances, discarding the waste. This continuous turning of inputs into outputs is controlled through the collection, interpretation, and manipulation of information. The process of control itself has two thrusts. It involves measurement—the comparison of the current state of a system to its desired state. And it involves two-way communication—the transmission of instructions and the collection of feedback on results. The processing of information for the purpose of control may result in the release of a hormone into the bloodstream, the expansion of a factory's production capacity, or the launch of a missile from a warship, but it works in essentially the same way in any living system.

When Herman Hollerith created his punch-card tabulator in the 1880s, he wasn't just pursuing his native curiosity as an engineer and an inventor. He was responding to an imbalance between, on the one hand, the technologies for processing matter and energy and, on the other, the technologies for processing information. He was trying to help resolve what James R. Beniger, in *The Control Revolution*, calls a "crisis of control," a crisis that threatened to

undermine the stability of markets and bring economic and technological progress to a standstill.

Throughout the first two centuries of the Industrial Revolution, the processing of matter and energy had advanced far more rapidly than the processing of information. The steam engine, used to power ships and trains and industrial machines, allowed factories, transportation carriers, retailers, and other businesses to expand their operations and their markets far beyond what was possible when production and distribution were restricted by the limitations of physical strength. Business owners, who had previously been able to observe their operations in their entirety and control them directly, now had to rely on information from many different sources to manage their companies. But they found that they lacked the means to collect and analyze the information fast enough to make timely decisions. Measurement and communication both began to break down, hamstringing management and impeding the further growth of businesses. As the sociologist Emile Durkheim observed in 1893, "The producer can no longer embrace the market in a glance, nor even in thought. He can no longer see limits, since it is, so to speak, limitless. Accordingly production becomes unbridled and unregulated." Government officials found themselves in a similar predicament, unable to assemble and analyze the information required to regulate commerce. The processing of materials and energy had progressed so rapidly that it had gone, quite literally, out of control.

During the second half of the nineteenth century, a series of technological advances in information processing helped administrators, in both business and government, begin to reimpose control over commerce and society, bringing order to chaos and opening the way for even larger organizations. The construction of the telegraph system, begun by Samuel F. B. Morse in 1845, allowed information to

be communicated instantaneously across great distances. The establishment of time zones in 1883 allowed for more precise scheduling of trains, speeding shipments and reducing accidents. The most important of the new control technologies, however, was bureaucracy—the organization of people into hierarchical information-processing systems. Bureaucracies had been around as long as civilization itself, but as Beniger writes, “bureaucratic administration did not begin to achieve anything approximating its modern form until the late Industrial Revolution.” Just as the division of labor in factories provided for the more efficient processing of matter, so the division of labor in government and business offices allowed for the more efficient processing of information.

But bureaucrats alone could not keep up with the flood of data that needed to be processed—the measurement and communication requirements went beyond the capacities of even large groups of human beings. Just like their counterparts on factory floors, information workers needed new tools to do their jobs. That requirement became embarrassingly obvious inside the US Census Bureau at the end of the century. During the 1870s, the federal government, struggling to administer a country and an economy growing rapidly in size and complexity, had demanded that the Bureau greatly expand the scope of its data collection, particularly in the areas of business and transport. The 1870 census had spanned just five subjects; the 1880 round was expanded to cover 215. But the new census turned into a disaster for the government. Even though many professional managers and clerks had been hired by the Bureau, the volume of data overwhelmed their ability to process it. By 1887, the agency found itself in the uncomfortable position of having to begin preparations for the next census even as it was still laboring to tabulate the results of the last one. It was in that context that Hollerith, who had worked on the 1880 tally, rushed to invent his

information-processing machine. He judged, correctly, that it would prove invaluable not only to the Census Bureau but to large companies everywhere.

The arrival of Hollerith's tabulator was a seminal event in a new revolution—a “Control Revolution,” as Beniger terms it—that followed and was made necessary and inevitable by the Industrial Revolution. Through the Control Revolution, the technologies for processing information finally caught up with the technologies for processing matter and energy, bringing the living system of society back into equilibrium. The entire history of automated data processing, from Hollerith's punch-card system through the mainframe computer and on to the modern computer network, is best understood as part of that ongoing process of reestablishing and maintaining control. “Microprocessor and computer technologies, contrary to currently fashionable opinion, are not new forces only recently unleashed upon an unprepared society,” writes Beniger, “but merely the latest installment in the continuing development of the Control Revolution.”

It should come as no surprise, then, that most of the major advances in computing and networking, from Hollerith's time to the present, have been spurred not by a desire to liberate the masses but by a need for greater control on the part of commercial and governmental bureaucrats, often ones associated with military operations and national defense. Indeed, the very structure of a bureaucracy is reflected in the functions of a computer. A computer gathers information through its input devices, records information as files in its memory, imposes formal rules and procedures on its users through its programs, and communicates information through its output devices. It is a tool for dispensing instructions, for gathering feedback on how well those instructions are carried out, and for measuring progress toward some specified goal. In using a computer, a person

becomes part of the control mechanism. He turns into a component of what the Internet pioneer J. C. R. Licklider, in his seminal 1960 paper "Man-Computer Symbiosis," described as a system integrating man and machine into a single, programmable unit.

But while computer systems played a major role in helping businesses and governments reestablish central control over workers and citizens in the wake of the Industrial Revolution, the other side of their nature—as tools for personal empowerment—has also helped shape modern society, particularly in recent years. By shifting power from institutions to individuals, information-processing machines can dilute and disturb control as well as reinforce it. Such disturbances tend to be short-lived, however. Institutions have proven adept at reestablishing control through the development of ever more powerful information technologies. As Beniger explains, "information processing and flows need themselves to be controlled, so that informational technologies continue to be applied at higher and higher levels of control."

The arrival of the personal computer in the 1980s posed a sudden and unexpected threat to centralized power. It initiated a new, if much more limited, crisis of control. Pioneered by countercultural hackers and hobbyists, the PC was infused from the start with a libertarian ideology. As memorably portrayed in Apple Computer's dramatic "1984" television advertisement, the personal computer was to be a weapon against central control, a tool for destroying the Big Brother-like hegemony of the corporate mainframe and its dominant producer, IBM. Office workers began buying PCs with their own money, bringing them to their offices, and setting them up on their desks. Bypassing corporate systems altogether, PC-empowered employees seized control of the data and programs they used. They gained freedom, but in the process they weakened the ability of bureaucracies to monitor and steer their work. Business

executives and the IT managers that served them viewed the flood of PCs into the workplace as "a Biblical plague," in the words of Paul Ceruzzi, the computer historian.

The breakdown of control proved fleeting. The client-server system, which tied the previously autonomous PCs together into a network connected to a central store of corporate information and software, was the means by which the bureaucrats reasserted their control over information and its processing. Together with an expansion in the size and power of IT departments, client-server systems enabled companies to restrict access to data and to limit the use of software to a set of prescribed programs. Ironically, once they were networked into a corporate system, PCs actually allowed companies to monitor, structure, and guide the work of employees more tightly than ever. "Local networking took the 'personal' out of personal computing," explains Ceruzzi. "PC users in the workplace accepted this Faustian bargain. The more computer-savvy among them resisted, but the majority of office workers hardly even noticed how much this represented a shift away from the forces that drove the invention of the personal computer in the first place. The ease with which this transition took place shows that those who believed in truly autonomous, personal computing were perhaps naïve."

The popularization of the Internet, through the World Wide Web and its browser, touched off a similar control crisis. Although the construction of the Internet was spearheaded by the Defense Department, a paragon of centralized power, it was designed, paradoxically, to be a highly dispersed, loosely organized network. Since the overriding goal was to build as reliable a system as possible—one that could withstand the failure of any of its parts—it was given a radically decentralized structure. Every computer, or node, operates autonomously, and communications between computers don't have to pass through any central clearinghouse. The Net's "internal

protocols," as New York University professor Alexander Galloway writes, "are the enemy of bureaucracy, of rigid hierarchy, and of centralization." If a corporate computer network was akin to a railroad, with tightly scheduled and monitored traffic, the Internet was more like the highway system, with largely free-flowing and unsupervised traffic.

At work and at home, people found they could use the Web to once again bypass established centers of control, whether corporate bureaucracies, government agencies, retailing empires, or media conglomerates. Seemingly uncontrolled and uncontrollable, the Web was routinely portrayed as a new frontier, a Rousseauian wilderness in which we, as autonomous agents, were free to redefine society on our own terms. "Governments of the Industrial World," proclaimed John Perry Barlow in his 1996 manifesto "A Declaration of the Independence of Cyberspace," "you are not welcome among us. You have no sovereignty where we gather." But, as with the arrival of the PC, it didn't take long for governments and corporations to begin reasserting and even extending their dominion.

The error that Barlow and many others have made is to assume that the Net's decentralized structure is necessarily resistant to social and political control. They've turned a technical characteristic into a metaphor for personal freedom. But, as Galloway explains, the connection of previously untethered computers into a network governed by strict protocols has actually created "a new apparatus of control." Indeed, he writes, "the founding principle of the Net is control, not freedom—control has existed from the beginning." As the disparate pages of the World Wide Web turn into the unified and programmable database of the World Wide Computer, moreover, a powerful new kind of control becomes possible. Programming, after all, is nothing if not a method of control. Even though the Internet still has no center, technically speaking, control can

now be wielded, through software code, from anywhere. What's different, in comparison to the physical world, is that acts of control become harder to detect and those wielding control more difficult to discern.

IN EARLY 2000, a Frenchman named Mark Knobel sued Yahoo for selling Nazi memorabilia through its online auction pages. Distributing such goods had long been illegal in France, and Knobel saw no reason why the law shouldn't apply to Internet merchants as it applied to local shopkeepers. "There is this naïve idea that the Internet changes everything," said one of Knobel's lawyers in filing the suit in a Paris court. "It doesn't change everything. It doesn't change the laws in France." But Yahoo's founder Jerry Yang, viewing the suit as a silly and futile attack on the supranational imperium of the Internet, believed that it was Knobel and his lawyers who were the ones afflicted by naïveté. "The French tribunal wants to impose a judgment in an area over which it has no control," he said. "Asking us to filter access to our sites is very naïve."

The suit went forward, nonetheless, and Knobel won. On November 20, 2000, a French judge ruled that Yahoo had broken the law and ordered the company to use its "best effort" to remove Nazi merchandise from any Web page that could be viewed in France. He noted that the company, contrary to its initial claims, was already using software to identify the locations of visitors to its sites in order to serve them customized advertisements. If it could control the ads it showed to people in different countries, he reasoned, it could control the other content they saw as well.

Unbowed, Yahoo announced that it would ignore the decision, claiming that the French courts had no authority in the matter. The judge was not amused. If the company did not comply with his order by February 1, 2001, he said, Yahoo's French assets would

be subject to seizure and its executives would be subject to arrest if they set foot in Europe. At that point, Yahoo retreated. A public company with interests around the world, it had little choice but to give in. On January 2, it announced that it would ban from all its sites the sale of products "associated with groups which promote or glorify hatred and violence."

A year later, in the summer of 2002, Jerry Yang had little to say, at least publicly, when Yahoo signed an agreement with the Chinese government that required the company to actively monitor and censor the contents of its sites in China. Yang and other Yahoo executives also remained quiet in 2005 when the company obeyed a demand by Chinese authorities that it reveal the identity of a Chinese citizen who had used a Yahoo email account to send a message about the anniversary of the Tiananmen Square massacre to an organization in the United States. A journalist named Shi Tao was promptly arrested and jailed. "The Yahoo story," write Jack Goldsmith and Tim Wu, "encapsulates the Internet's transformation from a technology that resists territorial law to one that facilitates its enforcement."

Not only are governments in general beginning to partition the online world along old geopolitical lines, but authoritarian regimes are coming to realize that the Internet may not pose as large a threat to their power as they initially feared. While the Net offers people a new medium for discovering information and voicing opinions, it also provides bureaucrats with a powerful new tool for monitoring speech, identifying dissidents, and disseminating propaganda. In a country like China, anyone who assumes that he can act anonymously on the Web opens himself to dangers far beyond embarrassment. In a 2007 speech, China's president, Hu Jintao, spoke glowingly of the Internet's potential for reinforcing the Communist Party's influence over the thoughts of his country's people.

"Strengthening network culture construction and management," he told a group of high-ranking party officials, "will help extend the battlefield of propaganda and ideological work. It is good for increasing the radiant power and infectiousness of socialist spiritual growth."

Democratic governments, as well, have begun scouring Web databases and monitoring Internet traffic for purposes of domestic surveillance. In 2004, federal agencies in the United States were conducting or planning 199 data-mining programs, according to a survey by congressional auditors. In late 2005 and early 2006, press reports indicated that the shadowy National Security Agency had been tapping into the commercial switches that route Internet traffic in a sweeping data-mining operation aimed at uncovering terrorists. As phone calls and other conversations are digitized and routed over the Internet and as ge positioning chips proliferate, the ability of governments of all stripes to monitor their citizens' words and movements will only increase.

Businesses have also found that the Internet, far from weakening their control over employees, actually strengthens their hand. Corporate influence over the lives and thoughts of workers used to be bounded by both space and time. Outside the walls of a company's offices and outside the temporal confines of the workday, people were largely free from the control of their bosses. But one of the consequences of the Net's boundary-breaking is that the workplace and the workday have expanded to fill all space and all time. Today, corporate software and data can be accessed from anywhere over the Internet, and email and instant-messaging traffic continues around the clock. In many companies, the *de facto* assumption is that employees are always at work, whether they're in their office, at their home, or even on vacation.

The BlackBerry has become the most visible symbol of the expan-

sion of corporate control over people's lives. Connected wirelessly to corporate servers, the ubiquitous gadget forms an invisible tether tying employees to their jobs. For many of today's knowledge workers, turning off the BlackBerry is the last thing they do before going to bed and turning it back on is the first thing they do upon waking. The *Wall Street Journal*, in a 2006 feature titled "BlackBerry Orphans," told the story of a typical BlackBerry-addicted executive whose young children demanded that she not check her email while at home with them in the evening. "To get around their dictates," reported the paper, "the mother hides the gadget in the bathroom, where she makes frequent trips before, during and after dinner." The woman sheepishly told the *Journal* that her children "think I have a small bladder."

The story is as amusing as it is poignant, but it underscores the striking change in the relationship between employers and employees that the Net has already produced. And it reveals another Faustian bargain that employees have struck with computer technology. Many people feel a genuine sense of empowerment when they use their BlackBerry or otherwise connect from afar to their corporate network. They welcome the technology because it "frees" them to work whenever and wherever they want, making them more productive and successful in their jobs. The price they pay, of course, is a loss of autonomy, as their employers gain greater control over their time, their activities, and even their thoughts. "Even though I'm home," another BlackBerry user told the *Journal*, "I'm not necessarily there."

With the data-collection and analytical tools of the World Wide Computer, employers will be able to extend their influence even further. Some companies have already begun to create mathematical models of their workforces, reducing each employee to a set of numbers that can be "optimized" by computer algorithms. IBM,

which in recent years has developed sophisticated software for modeling the workings of industrial supply chains, is now creating similar models for managing people. A forty-person group of IBM statisticians and data-mining experts is working "to refocus the supply-chain programs on 50,000 of the consultants in IBM's services division," reports *BusinessWeek*. "That means that instead of modeling machines, furnaces, and schedules, they're building models of their colleagues." The team is drawing information on employees from IBM's many corporate databases, and it's also looking at incorporating data from workers' email messages, online calendars, and mobile phone calls. The company hopes to use the model to assign consultants automatically to projects and to direct their work for optimum efficiency.

Google, too, has launched an experiment in using mathematical modeling for personnel management. During the summer of 2006, it asked its employees to fill out an extensive online survey about themselves, answering some 300 questions on everything from the programming languages they know to the magazines they read to the pets they keep. The company ran the answers through a computer, comparing them to various measures of the employees' skills and achievements in order to create algorithms that might predict performance. In 2007, it began using the algorithms to evaluate all job applicants, who are now also required to fill out a long questionnaire on the Internet. As such modeling techniques progress, they will come to be used much more broadly by companies. "This mathematical modeling of humanity promises to be one of the great undertakings of the 21st century," *BusinessWeek* concludes. "And it doesn't take much imagination to see where that can lead. Managers will operate tools not only to monitor employees' performance but also to follow their movements and drive up productivity."

THE MOST far-reaching corporate use of the World Wide Computer as a control technology will not be for optimizing what we do as employees. It will be for optimizing how we act as consumers. Despite the resistance of the Web's early pioneers and pundits, consumerism long ago replaced libertarianism as the prevailing ideology of the online world. Restrictions on the commercial use of the Net collapsed with the launch of the World Wide Web in 1991. The first banner ad—for a Silicon Valley law firm—appeared in 1993, followed the next year by the first spam campaign. In 1995, Netscape tweaked its Navigator browser to support the “cookies” that enable companies to identify and monitor visitors to their sites. By 1996, the dotcom gold rush had begun. More recently, the Web's role as a sales and promotion channel has expanded further. Assisted by Internet marketing consultants, companies large and small have become much more adept at collecting information on customers, analyzing their behavior, and targeting products and promotional messages to them.

The growing sophistication of Web marketing can be seen most clearly in advertising. Rather than being dominated by generic banner ads, online advertising is now tightly tied to search results or other explicit indicators of people's desires and identities. Search engines themselves have become the leading distributors of ads, as the prevailing tools for Web navigation and corporate promotion have merged into a single and extraordinarily profitable service. Google originally resisted the linking of advertisements to search results—its founders argued that “advertising-funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers”—but now it makes billions of dollars through the practice. Search-engine optimization—the science of using advanced statistical techniques to increase the likelihood that a person will visit a site or click on an ad—has become an important corporate function, which Google and other search

engines promote by sharing with companies information on how they rank sites and place ads.

In what is perhaps the most remarkable manifestation of the triumph of consumerism on the Web, popular online communities like MySpace encourage their members to become friends with corporations and their products. During 2006, for example, more than 85,000 people “friended” Toyota's Yaris car model at the site, happily entangling themselves in the company's promotional campaign for the recently introduced vehicle. “MySpace can be viewed as one huge platform for ‘personal product placement,’” writes Wade Roush in an article in *Technology Review*. He argues that “the large supply of fake ‘friends,’ together with the cornucopia of ready-made songs, videos, and other marketing materials that can be directly embedded in [users'] profiles, encourages members to define themselves and their relationships almost solely in terms of media and consumption.” In recognition of the blurring of the line between customer and marketer online, *Advertising Age* named “the consumer” its 2007 Advertising Agency of the Year.

But the Internet is not just a marketing channel. It's also a marketing laboratory, providing companies with unprecedented insights into the motivations and behavior of shoppers. Businesses have long been skilled at controlling the supply side of their operations, thanks in large part to earlier advances in information technology, but they've struggled when it comes to exerting control over the demand side—over what people buy and where and when they buy it. They haven't been able to influence customers as directly as they've been able to influence employees and suppliers. Advertising and promotion have always been frustratingly imprecise. As the department store magnate John Wanamaker famously said more than a hundred years ago, “Half the money I spend on advertising is wasted. The trouble is, I don't know which half.”

The World Wide Computer is beginning to change that. It promises to strengthen companies' control over consumption by providing marketers with the data they need to personalize their pitches precisely and gauge the effects of those pitches accurately. It optimizes both communication and measurement. In a 2006 interview with the *Economist*, Rishad Tobaccowala, a top executive with the international ad agency Publicis, summed up the change in a colorful, and telling, metaphor. He compared traditional advertising to dropping bombs on cities—a company can't be sure who it hits and who it misses. But with Internet ads, he said, companies can "make lots of spearheads and then get people to impale themselves."

From time to time, in response to public or governmental concerns, the leading search engines and other top Internet companies roll out new "privacy safeguards" with great fanfare. But the moves are rarely more than window-dressing. In the summer of 2007, Google announced it would delete the cookies it uses to gather information on a person's searches after two years—but only if the person doesn't perform a Google search or visit any other Google site during that period, an unlikely scenario given the company's dominance on the Web. Around the same time, Microsoft issued a press release saying it would "anonymize" data on searches after eighteen months. As Thelma Arnold's experience shows, anonymization provides little real protection in the face of sophisticated data-mining techniques. In fact, even as Microsoft was making its announcement, the company had a team of Chinese researchers hard at work developing analytical software for distilling personal demographic information from anonymous online data. The team had already developed an algorithm able to predict with considerable accuracy a Web surfer's age and gender based on the sites he or she visits, and it was working on prediction algorithms for other characteristics such as occupation, education, and location. In

a paper on their work, the researchers wrote that "the diversity of [a] user's online browsing activities can be exploited to determine an unknown user's demographic attributes."

The ability of businesses to gather and analyze rich data on individual customers comes at the same time that psychologists and economists are making progress in a new discipline called neuromarketing. Neuromarketers use brain scans to uncover the mental triggers that determine what we buy. In a landmark 2007 article, "Neural Predictors of Purchases," published in the journal *Neuron*, a group of scholars from MIT, Stanford, and Carnegie Mellon reported that they could use MRI machines to monitor the brain activity of shoppers as they evaluated products and prices on computer screens. By pinpointing which circuits in the brain "lit up" at different stages in the buying process, the researchers found they were able to predict whether a person would buy a product or pass it up. They concluded, after further analysis of the results, that "the ability of brain activation to predict purchasing would generalize to other purchasing scenarios." *Forbes* heralded the study as a milestone in business, saying it marked the first time researchers have been able "to examine what the brain does while making a purchasing decision." It's not hard to see that we're entering a new era of commerce in which companies will wield far greater influence over our choices than ever before—without our knowing it.

As has been the case so often in the history of information processing, many of the control tools companies are coming to deploy on the Internet have their origins in military research. In the wake of the September 11, 2001, terrorist attacks, the Department of Defense, the Department of Homeland Security, and other US government agencies began investing millions if not billions of dollars into the development of data-mining and analysis technologies that can draw valuable intelligence out of the Net's cloud of data.

Beyond uncovering terrorists, many of these technologies will have applications in the commercial realm. The government has, for instance, funded research at the University of Arizona to develop “stylometry” software that can be used to identify the authors of textual messages appearing on the Internet by analyzing their diction and syntax. Although the research is aimed at “the application of authorship identification techniques to English and Arabic extremist group forum messages,” the software could also aid marketers in profiling customers or tracing connections among anonymous product reviews.

None of this means that the computer is about to become purely a control technology. It will always have a dual nature, giving new powers to individuals as well as institutions. We will continue to see advances in information technology that weaken central control, but every disruption will almost certainly be followed by the reassertion of control, whether through legal or technological means. We see this process playing out again today with Napster’s heirs, the powerful and highly decentralized peer-to-peer networks that people use to trade movies, software programs, and other large files. Long the realm of libertarians, pirates, and anti-copyright activists, the freewheeling networks have faced a series of lawsuits and police raids in recent years. Now, they’re beginning to turn themselves into mainstream media. BitTorrent, the developer of the most popular of the peer-to-peer systems, struck deals in 2006 with several Hollywood studios, agreeing to distribute their movies and TV shows in formats that prevent unauthorized copying. It then turned its home page into a slick online store, required users to sign up for accounts and agree to a long list of “terms of use,” and adopted a new name: the BitTorrent Entertainment Network.

History tells us that the most powerful tools for managing the processing and flow of information will be placed in the hands not

of ordinary citizens but of businesses and governments. It is their interest—the interest of control—that will ultimately guide the progress and the use of the World Wide Computer.

“AS EVERY MAN goes through life he fills in a number of forms for the record, each containing a number of questions,” Alexander Solzhenitsyn wrote in his novel *Cancer Ward*. “A man’s answer to one question on one form becomes a little thread, permanently connecting him to the local center of personnel records administration. There are thus hundreds of little threads radiating from every man, millions of threads in all. If these threads were suddenly to become visible, the whole sky would look like a spider’s web. . . . Each man, permanently aware of his own invisible threads, naturally develops a respect for the people who manipulate the threads.”

As we go about our increasingly digitized lives, the threads that radiate from us are multiplying far beyond anything that even Solzhenitsyn could have imagined in the Soviet Union in the 1960s. Nearly everything we do online is recorded somewhere in the machinery of the World Wide Computer. Every time we read a page of text or click on a link or watch a video, every time we put something in a shopping cart or perform a search, every time we send an email or chat in an instant-messaging window, we are filling in a “form for the record.” Unlike Solzhenitsyn’s *Everyman*, however, we’re often unaware of the threads we’re spinning and how and by whom they’re being manipulated. And even if we were conscious of being monitored or controlled, we might not care. After all, we also benefit from the personalization that the Internet makes possible—it makes us more perfect consumers and workers. We accept greater control in return for greater convenience. The spider’s web is made to measure, and we’re not unhappy inside it.