

HERE COMES EVERYBODY

THE POWER OF ORGANIZING

WITHOUT ORGANIZATIONS

CLAY SHIRKY

THE PENGUIN PRESS | NEW YORK | 2008

CHAPTER 10 | Failure for Free 233

CHAPTER 11 | Promise, Tool, Bargain 260

EPILOGUE 293

ACKNOWLEDGMENTS 305

BIBLIOGRAPHY 308

INDEX 320

ABOUT THE AUTHOR 326

CHAPTER 1

IT TAKES A VILLAGE TO FIND A PHONE

On an afternoon in late May 2006 a woman named Ivanna left her phone in the backseat of a New York City cab. No surprise there; hundreds of phones a year show up in the New York Taxi and Limousine Commission's offices, and more than that are actually lost, since some unknown number are simply taken by the next passenger. That was the fate of Ivanna's phone, a fairly expensive multifunction version called a Sidekick, which came with a screen, keyboard, and built-in camera. Sadly for her, the Sidekick was the sole repository of much of the information for her upcoming wedding, from contact information for the catering company to the guest list.

When she realized what she'd done, Ivanna asked Evan Guttman, a friend who worked as a programmer in the financial industry, to offer a reward for its return, via an e-mail message that would show up on the phone. Getting no response after a couple of days, she shelled out more than \$300 to buy a new one. Ivanna's phone company had stored copies of her information on its servers and transferred it to her new phone. Once she had the information on her new phone, she discovered that

her original one had ended up in the hands of a girl in Queens. Ivanna knew this because the girl was using it to take pictures of herself and her friends and e-mail them around; the photos taken on her old phone had been transferred to her new one. Ivanna and Evan couldn't be sure who had taken the phone from the cab, but they knew who had it now, or rather they had her picture and her e-mail address, Sashacristal8905@aol.com (since disabled, for reasons that will become apparent).

Evan immediately e-mailed Sasha, explaining the situation and asking for the phone back. Sasha replied that she wasn't stupid enough to return it, a view punctuated with racial invective, saying that Evan's "white ass" didn't deserve it back. (She inferred Evan and Ivanna's race from pictures on the phone; Sasha is Hispanic.) The back-and-forth went on for some time. During the conversation Sasha said her brother had found it in a cab and given it to her; Evan continued to ask for it back, on the grounds that Sasha knew who its rightful owner was. Sasha finally wrote that she and her boyfriend would meet Evan, saying, in the spelling-challenged manner of casual e-mails, "i got ball this is my adress 108 20 37 av corona come n do it iam give u the sidekick so I can hit you wit it."

Evan declined to go to the listed address, both because he assumed it was fake (it was) and because of the threatened violence. Instead, he decided to take the story public. He created a simple webpage with Sasha's photos and a brief description of the events so far, with the stated rationale of delivering a lesson on "the etiquette of returning people's lost belongings," as he put it. He titled the page *StolenSidekick*, added it to his personal website at *EvanWasHere.com*, and began telling his friends about what had happened.

The original page went up on June 6, and in the first few hours it was up, Evan's friends and their friends forwarded it around the internet, attracting a growing amount of attention. Evan first updated the page later that day, noting that his friends had done some online detective work and had found a page on MySpace, the social networking website, that had photos of Sasha and a man they surmised was her boyfriend. Evan's second update provided more background on how the phone was lost and on who had it now. His third update, later that afternoon, reported that an officer from the NYPD had seen the story and had written explaining how to file a claim with the police.

That evening, two things happened. First, a man named Luis sent Evan mail, saying he was Sasha's brother and a member of the Military Police. He said that Sasha had bought the phone from a cabbie. (This story, as Evan pointed out on the webpage, directly contradicted Sasha's earlier account of her brother finding the phone.) Luis also told Evan to stop harassing Sasha, hinting violence if Evan didn't lay off. The other event that evening was that Evan's story appeared on Digg. Digg is a collaborative news website; users suggest stories, and other users rate them thumbs up or thumbs down. The Digg front page, like all newspaper front pages, is made up of stories that are both timely and important, except on Digg timeliness is measured by how recently a story was added, and importance is measured by user votes rather than by the judgment of editors. The front page of Digg gets millions of readers a day, and a lot of those readers took a look at the *StolenSidekick* page.

The story clearly struck a nerve. Evan was getting ten e-mails *a minute* from people asking about the phone, offering encouragement, or volunteering to help. Everyone who has ever

lost something feels a diffuse sense of anger at whoever found and kept it, but this time it was personal, since Evan, and everyone reading StolenSidekick, now knew who had the phone and had seen her insulting refusals to return it. When the barrier to returning something is high, we make peace with “Finders, keepers. Losers, weepers,” but when returning something becomes easier, our sympathies ebb. Finding a loose bill on the street is different from finding a wallet with ID in it, and the case of the missing Sidekick was even worse than a lost wallet. Using someone’s own phone to refuse to return it to them crossed some barrier of acceptability in the eyes of many following the saga, and the taunts and threats from Sasha and her friends and family only added insult to injury.

Evan, clearly energized by the response from his growing readership, continued posting a running commentary on his webpage. He wrote forty updates in ten days, accompanied by a growing frenzy of both local and national media attention. There was a lot to update: he and the people tuning in posted more MySpace profiles of Sasha, her boyfriend Gordo, and her brother. Someone reading the StolenSidekick page figured out Sasha’s full name, then her address, and drove by her house, later posting the video on the Web for all to see. Members of Luis’s Military Police unit wrote to inquire about allegations that an MP was threatening a civilian and promised to look into the matter.

Evan also created a bulletin board for his readers, a place online where they could communicate with one another about the attempts to recover Ivanna’s phone. Or rather, he tried to create a bulletin board, but the first such service he selected simply couldn’t cope with the crush of excited users all trying

to log in at the same time. Seeing this, he selected a second bulletin board service, but that too crashed under the sudden shock of demand, as did the third. (These kinds of failures, sometimes called “success crises,” bring to mind Yogi Berra’s famous observation about a New York restaurant: “Nobody goes there anymore. It’s too crowded.”) He finally found a service that could accommodate the thousands of people following the Sidekick saga, and those readers settled in, discussing every aspect of the events, from general speculation about Sasha’s moral compass to a forum inviting members of the military to talk about Luis, the MP, and his involvement in the events. (As is usual with these kinds of communities, much of the conversation was off-topic; the military section of the bulletin board included a conversation about whether Luis was taking sufficient care of the uniform he was wearing in the pictures Sasha had taken.)

During this period Sasha’s family and friends kept communicating with Evan about the phone, offering several inconsistent stories: her mom had bought the phone from someone, Sasha didn’t have the phone, she had sold the phone, she would sell him the phone back for \$100. Luis announced they were going to sue for harassment; her friends wrote in with more threatening e-mail. Evan and Ivanna filed a report with the police, who classified the phone as lost rather than stolen property, meaning they would take no action. Several people in the New York City government wrote in offering to help get the complaint amended, including a police officer who shared internal NYPD paperwork and explained how the complaint should have been handled. (Possession of this paperwork almost got Evan arrested when he later tried to get

the complaint reclassified.) By this point millions of readers were watching, and dozens of mainstream news outlets had covered the story. The public airing of the NYPD's refusal to treat this case as theft generated so many public complaints that the police later reversed their stand and, after dispatching two detectives to talk with Ivanna, agreed to treat the phone as stolen rather than lost.

Then on June 15 members of the NYPD arrested Sasha, a sixteen-year-old from Corona, New York, and recovered the stolen Sidekick, which they returned to its original owner, Ivanna. As Sasha's mother memorably told a reporter the day her daughter was arrested, "I never in my life thought a phone was gonna cause me so many problems." It wasn't the phone that caused the problems, though. It was the people at the other end of the phone, people who had come together around Evan's page, who found the MySpace profiles and the family's address and helped pressure the police department, all in a busy ten days, and all of it leading to Sasha's arrest. Having achieved their stated goals of publicly calling out Sasha and retrieving the phone, Evan and Ivanna declined to press charges, and Sasha was released. Ivanna's wedding went off without a hitch, and Evan, in light of his ability to gather a crowd, began getting freelance work doing PR.

"Give me a place to stand and a lever long enough, and I will move the world."

The loss and return of the Sidekick is a story about many things—Evan's obsessive tendencies, Ivanna's good fortune in

having him for a friend, how expensive phones have gotten—but one of the themes running through the story is the power of group action, given the right tools. Despite Evan's heroic efforts, he could not have gotten the phone returned if he had been working alone. He used his existing social network to get the word out, which in turn helped him find an enormous audience for Ivanna's plight, an audience willing to do more than just read from the sidelines. This audience gave Evan remarkable leverage in dealing with Sasha, and with the NYPD, leverage he wouldn't have had without such an engaged group following along. Indeed, the nature of that engagement puts many of the visitors to Evan's webpage in a category that Dan Gillmor, a journalist and the author of *We the Media*, calls "the former audience," those people who react to, participate in, and even alter a story as it is unfolding.

Consider the story from Sasha's point of view. She's a teenager in a media-saturated culture, she's given a very expensive, very cool phone that someone found in the back of a cab, and she decides to keep it rather than try to track down the owner. This isn't the most ethical behavior in the world, but neither is it premeditated theft, and in any case, what could go wrong? She's got her friends and family backing her up, and she surmises, correctly, that Evan isn't in any hurry to come out to Corona. Given all this, the combination of stories and threats from Sasha and her friends and family should have worked. After all, the phone was expensive, but it wasn't that expensive, and it's not like \$300 would buy Evan a lot of help. If what Evan wanted was to save Ivanna the price of the phone, spending more than \$300 retrieving it wouldn't make any sense.

Evan wasn't in it for the money, though. He was in it to satisfy his sense of justice. Because his commitment to the task at hand was emotional rather than financial, and because he was well-off enough, he was able to invest considerably more in the recovery effort than the phone was actually worth. His decision to present those motivations in public also helped draw people in. "This is not a religious endeavor or a moral endeavor . . . [sic] this is a HUMANITY endeavor," Evan wrote at one point. The story of righting a wrong is a powerful one and helped him generate the involvement of others that finally led to the recovery of the phone.

Sasha and her friends didn't just want Evan to fail—they assumed that he would fail. The threats from Luis and Gordo had a kind of "You and what army?" quality about them, because they were certain that the police weren't going to get involved. (Luis made this very point in his first message to Evan: "dont give me that bullshit about you going to the cops over a lost phone the nypd has better things to do then to worry about your friend losing her phone." [sic]) The turning point in Evan's quest was the moment when the police agreed to amend the complaint from "lost property" (about which they would have done nothing) to "stolen property" (which led to Sasha's arrest). The NYPD is not an easy organization to browbeat, yet days after they'd tried to close the case, there they were, sending two detectives to spend half an hour with Ivanna on the matter, then sending more officers out to Corona to collar Sasha and retrieve the Sidekick. Imagine how disorienting it must have been for Sasha to learn that the owner of the phone actually did have an army of sorts, including lawyers and cops, along with an international audience of millions.

Thanks to the Web, the cost of publishing globally has collapsed. That raw publishing capability, Evan's existing social contacts, the unusual nature of his story, and the fact that the audience could find Sasha's MySpace page all combined to create a kind of positive reinforcement of attention. People became interested in the story, and they forwarded it to friends and colleagues, who became interested in turn and forwarded it still further. This pattern of growth was both cause and effect for mainstream media getting involved—it's unlikely that *The New York Times* or CNN would have covered the story of a lost phone, but when it was wrapped in the larger story of national and even global attention, they picked it up, which led to still more visitors to Evan's site and still more media outlets tuning in. The story ended up in more than sixty newspapers and radio and TV stations and more than two hundred weblogs. From the humble beginnings of Ivanna's story and a handful of snapshots of Sasha and her friends, the StolenSidekick page went on to get over a million viewers.

Having the attention of this audience changed the conditions for Evan's relations with the police, and he knew it. He even said in one of his updates that the function of the StolenSidekick page was to put pressure on the NYPD. It also emboldened him. When he went down to the Ninth Precinct to get the complaint upgraded from lost to stolen property, Evan was stymied by the desk officer, who told him in no uncertain terms that it was up to the NYPD to determine what was a crime and what wasn't. Evan's update later that day read, in part, "All I want to do is report a crime. This is ridiculous. Have no fear though. I have many surprises for the NYPD tomorrow. They WILL listen to me and the thousands of you

who have written me and the millions of you who are reading this page." The surprise that he knew was coming was the appearance of the story in *The New York Times* the following morning. Later, when the police indicated a willingness to pursue the case, Evan posted an explicit request to the site: "I ask that EVERYONE come back to visit this page for updates to make sure that the NYPD stay true to what they said." Faced with the opacity of the NYPD bureaucracy, Evan had the information-age equivalent of being able to see through walls: he got insider advice, and he was able to walk into a confrontation with a New York City cop knowing that the story would be front-page news the following morning.

You can see Evan coming to accept his part of the bargain with his users—they would provide the attention that kept him going and made the story attractive to traditional media, and he would channel that attention, reporting on his every move. Many of the viewers of the StolenSidekick page were not just readers but operated as one-person media outlets, members of the former audience, and they discussed the situation on weblogs, on mailing lists, and on various electronic discussion groups Evan set up. He had lawyers, policemen, online detectives, journalists, and even his own ad hoc pressure group working on his behalf, without belonging to any organization responsible for providing those functions.

Evan's updates included mention of constant encouragement and offers of help from more people in the city government who thought he was getting a raw deal from the NYPD. Hours after he posted the first version of the page, an NYPD officer contacted him to explain how to file a complaint. Four days later another officer from the NYPD wrote Evan wanting

to meet; when they did, the officer gave Evan copies of internal NYPD paperwork to show him the kind of form he needed to file to get it treated as a theft. Finally, when Sasha's family began threatening legal action, someone from Legalmatch.org, a legal advice site, offered to help Evan get free advice.

Obviously, much about this story is unrepeatable. It isn't a worldwide media event every time someone loses a phone. The unusualness of the story, though, throws into high relief the difference between past and present. It's unlikely that Evan could have achieved what he did even five years ago, and inconceivable that he could have achieved it ten years ago, because neither the tools he used nor the social structures he relied on were in place ten years ago. Equally obviously, much about this story depends on the angle you are viewing it from. For Ivanna, the story is mostly good. She benefited from Evan's obsessive behavior and the way it was fed by the attention he received, and she had to expend little effort to get her phone back. For Evan himself, the exhilaration of fighting for what he thought was right was balanced against the investment of time and expense. And for Sasha, of course, the story was mostly bad. Of all the telephones in all the towns in all the world, the one she got had a million people at the other end of the line.

And what about us? What about the society in which this tug-of-war was happening? For us the picture isn't so clear. The whole episode demonstrates how dramatically connected we've become to one another. It demonstrates the ways in which the information we give off about our selves, in photos and e-mails and MySpace pages and all the rest of it, has dramatically increased our social visibility and made it easier for us to find

each other but also to be scrutinized in public. It demonstrates that the old limitations of media have been radically reduced, with much of the power accruing to the former audience. It demonstrates how a story can go from local to global in a heartbeat. And it demonstrates the ease and speed with which a group can be mobilized for the right kind of cause.

But who defines what kind of cause is right? Evan's ability to get help can be ascribed either to a strong sense of injustice or to a petty unwillingness to lose a fight, no matter how trivial and no matter the cost to his opponent. And for all the offensiveness of Sasha's taunting, race and class do matter. Evan is a grown-up doing work that lets him take countless hours off to work on the retrieval of a phone. Sasha is an unwed teenage mother. The recovery of the phone wasn't the only loss she suffered—Evan's bulletin board quickly became host to public messages disparaging Sasha, her boyfriend and friends, single mothers, and Puerto Ricans as a group. One conversation, headed with the subject line "[D]o something already!," noted that other people following the story had already uncovered her address, and advocated physical confrontation (though the author didn't offer to participate). Another thread, with the charming title "[W]ould you tap that?," involved discussion by the male participants as to whether Sasha was attractive enough to sleep with.

One could blame Evan for letting these kinds of racist and sexist conversations take place, but the number of people interested in talking about the stolen phone (as evidenced by the inadequacy of most software to handle the volume of users), and the standard anonymity of internet users, made the conversations effectively impossible to police. Furthermore,

though Evan was clearly benefiting from having generated the attention, he was not entirely in control of it—the bargain he had crafted with his users had him performing the story they wanted to see. Had he shut down the bulletin boards or even edited the conversations, he would have been violating his half of what had quickly become a mutual expectation. (Whether he *should* have taken this step is a judgment call; the point is that once a group has come together, those kinds of issues of community control aren't simple. Any action Evan took, either letting the conversation go or stifling it, would have created complicated side effects.)

A larger question transcends the individual events. Do we want a world in which a well-off grown-up can use this kind of leverage to get a teenager arrested, as well as named and shamed on a global platform, for what was a fairly trivial infraction? The answer is yes and no. Millions of people obviously wanted to follow the story, in part because of its mix of moral and visceral struggle. Furthermore, what Sasha did was wrong, and we want misdeeds to be punished. At the same time, though, we want the punishment to fit the crime. It's easy enough to say that Sasha shouldn't have gotten off just because other people take lost property without returning it, but that logic starts to look different if we imagine that the roles were reversed. Poor people lose phones too, and the loss hits them far harder; why should Evan have been able to browbeat the NYPD into paying attention to this of all lost property?

A few years ago Evan wouldn't have been able to get the story heard either. Before the Web became ubiquitous, he wouldn't have been able to attract an audience, much less one

in the millions, and without that audience he would not have been able to get the police to change the complaint. Given how much of our lives is spent in thrall to unresponsive bureaucracy, Evan's eventual victory seems like a shining success, but it came at a cost. Policing time is finite, yet the willingness of humans to feel wronged is infinite. Do we also want a world where, whenever someone with this kind of leverage gets riled up, they can unilaterally reset the priorities of the local police department?

Those kinds of questions are rhetorical, since that's the world we've already got. The real question is, What happens next? The story of the lost Sidekick is an illustration of the kinds of changes—some good, some bad, most too complex to label—that are affecting the ways groups assemble and cooperate. These changes are profound because they are amplifying or extending our essential social skills, and our characteristic social failings as well.

New Leverage for Old Behaviors

Human beings are social creatures—not occasionally or by accident but always. Sociability is one of our core capabilities, and it shows up in almost every aspect of our lives as both cause and effect. Society is not just the product of its individual members; it is also the product of its constituent groups. The aggregate relations among individuals and groups, among individuals within groups, and among groups forms a network of astonishing complexity. We have always relied on group effort for survival; even before the invention of agriculture, hunt-

ing and gathering required coordinated work and division of labor. You can see an echo of our talent for sociability in the language we have for groups; like a real-world version of the mythical seventeen Eskimo words for snow, we use incredibly rich language in describing human association. We can make refined distinctions between a corporation and a congregation, a clique and a club, a crowd and a cabal. We readily understand the difference between transitive labels like “my wife’s friend’s son” and “my son’s friend’s wife,” and this relational subtlety permeates our lives. Our social nature even shows up in negation. One of the most severe punishments that can be meted out to a prisoner is solitary confinement; even in a social environment as harsh and attenuated as prison, complete removal from human contact is harsher still.

Our social life is literally primal, in the sense that chimpanzees and gorillas, our closest relatives among the primates, are also social. (Indeed, among people who design software for group use, human social instincts are sometimes jokingly referred to as the monkey mind.) But humans go further than any of our primate cousins: our groups are larger, more complex, more ordered, and longer lived, and critically, they extend beyond family ties to include categories like friends, neighbors, colleagues, and sometimes even strangers. Our social abilities are also accompanied by high individual intelligence. Even cults, the high-water mark of surrender of individuality to a group, can't hold a candle to a beehive in terms of absolute social integration; this makes us different from creatures whose sociability is more enveloping than ours.

This combination of personal smarts and social intuition makes us the undisputed champions of the animal kingdom

in flexibility of collective membership. We act in concert everywhere, from simple tasks like organizing a birthday party (itself a surprisingly complicated task) to running an organization with thousands or even millions of members. This skill allows groups to tackle tasks that are bigger, more complex, more dispersed, and of longer duration than any person could tackle alone. Building an airplane or a cathedral, performing a symphony or heart surgery, raising a barn or razing a fortress, all require the distribution, specialization, and coordination of many tasks among many individuals, sometimes unfolding over years or decades and sometimes spanning continents.

We are so natively good at group effort that we often factor groups out of our thinking about the world. Many jobs that we regard as the province of a single mind actually require a crowd. Michelangelo had assistants paint part of the Sistine Chapel ceiling. Thomas Edison, who had over a thousand patents in his name, managed a staff of two dozen. Even writing a book, a famously solitary pursuit, involves the work of editors, publishers, and designers; getting this particular book into your hands involved additional coordination among printers, warehouse managers, truck drivers, and a host of others in the network between me and you. Even if we exclude groups that are just labels for shared characteristics (tall people, redheads), almost everyone belongs to multiple groups based on family, friends, work, religious affiliation, on and on. The centrality of group effort to human life means that anything that changes the way groups function will have profound ramifications for everything from commerce and government to media and religion.

One obvious lesson is that new technology enables new kinds of group-forming. The tools Evan Guttman availed himself of were quite simple—the phone itself, e-mail, a webpage, a discussion forum—but without them the phone would have stayed lost. Every step of the way he was able to escape the usual limitations of private life and to avail himself of capabilities previously reserved for professionals: he used his site to tell the story without being a journalist, he found Sasha's information without being a detective, and so on. The transfer of these capabilities from various professional classes to the general public is epochal, built on what the publisher Tim O'Reilly calls "an architecture of participation."

When we change the way we communicate, we change society. The tools that a society uses to create and maintain itself are as central to human life as a hive is to bee life. Though the hive is not part of any individual bee, it is part of the colony, both shaped by and shaping the lives of its inhabitants. The hive is a social device, a piece of bee information technology that provides a platform, literally, for the communication and coordination that keeps the colony viable. Individual bees can't be understood separately from the colony or from their shared, co-created environment. So it is with human networks; bees make hives, we make mobile phones.

But mere tools aren't enough. The tools are simply a way of channeling existing motivation. Evan was driven, resourceful, and unfortunately for Sasha, very angry. Had he presented his mission in completely self-interested terms ("Help my friend save \$300!") or in unattainably general ones ("Let's fight theft everywhere!"), the tools he chose wouldn't have mattered. What he did was to work out a message framed in

big enough terms to inspire interest, yet achievable enough to inspire confidence. (This sweet spot is what Eric Raymond, the theorist of open source software, calls “a plausible promise.”) Without a plausible promise, all the technology in the world would be nothing more than all the technology in the world.

As we saw in the saga of the lost Sidekick, getting the free and ready participation of a large, distributed group with a variety of skills—detective work, legal advice, insider information from the police to the army—has gone from impossible to simple. There are many small reasons for this, both technological and social, but they all add up to one big change: forming groups has gotten a lot easier. To put it in economic terms, the costs incurred by creating a new group or joining an existing one have fallen in recent years, and not just by a little bit. They have collapsed. (“Cost” here is used in the economist’s sense of anything expended—money, but also time, effort, or attention.) One of the few uncontentious tenets of economics is that people respond to incentives. If you give them more of a reason to do something, they will do more of it, and if you make it easier to do more of something they are already inclined to do, they will also do more of it.

Why do the economics matter, though? In theory, since humans have a gift for mutually beneficial cooperation, we should be able to assemble as needed to take on tasks too big for one person. If this were true, anything that required shared effort—whether policing, road construction, or garbage collection—would simply arise out of the motivations of the individual members. In practice, the difficulties of

coordination prevent that from happening. (Why this is so is the subject of the next chapter.)

But there are large groups. Microsoft, the U.S. Army, and the Catholic Church are all huge, functioning institutions. The difference between an ad hoc group and a company like Microsoft is management. Rather than waiting for a group to self-assemble to create software, Microsoft manages the labor of its employees. The employees trade freedom for a paycheck, and Microsoft takes on the costs of directing and monitoring their output. In addition to the payroll, it pays for everything from communicating between senior management and the workers (one of the *raison d'être* for middle management) to staffing the human resources department to buying desks and chairs. Why does Microsoft, or indeed any institution, tolerate these costs?

They tolerate them because they have to; the alternative is institutional collapse. If you want to organize the work of even dozens of individuals, you have to manage them. As organizations grow into the hundreds or thousands, you also have to manage the managers, and eventually to manage the managers’ managers. Simply to exist at that size, an organization has to take on the costs of all that management. Organizations have many ways to offset those costs—Microsoft uses revenues, the army uses taxes, the church uses donations—but they cannot avoid them. In a way, every institution lives in a kind of contradiction: it exists to take advantage of group effort, but some of its resources are drained away by directing that effort. Call this the institutional dilemma—because an institution expends resources to manage resources, there is a gap between what those institutions are capable of in theory

and in practice, and the larger the institution, the greater those costs.

Here's where our native talent for group action meets our new tools. Tools that provide simple ways of creating groups lead to new groups, lots of new groups, and not just more groups but more kinds of groups. We've already seen this effect in the tools that Evan used—a webpage for communicating with the world, instant messages and e-mails by the thousands among his readers, and the phone itself, increasingly capable of sending messages and pictures to groups of people, not just to a single recipient (the historical pattern of phone use).

If we're so good at social life and shared effort, what advantages are these tools creating? A revolution in human affairs is a pretty grandiose thing to attribute to a ragtag bunch of tools like e-mail and mobile phones. E-mail is nice, but how big a deal can it be in the grand scheme of things? The answer is, "Not such a big deal, considered by itself." The trick is not to consider it by itself. All the technologies we see in the story of Ivanna's phone, the phones and computers, the e-mail and instant messages, and the webpages, are manifestations of a more fundamental shift. We now have communications tools that are flexible enough to match our social capabilities, and we are witnessing the rise of new ways of coordinating action that take advantage of that change. These communications tools have been given many names, all variations on a theme: "social software," "social media," "social computing," and so on. Though there are some distinctions between these labels, the core idea is the same: we are living in the middle of a remarkable increase

in our ability to share, to cooperate with one another, and to take collective action, all outside the framework of traditional institutions and organizations. Though many of these social tools were first adopted by computer scientists and workers in high-tech industries, they have spread beyond academic and corporate settings. The effects are going to be far more widespread and momentous than just recovering lost phones.

By making it easier for groups to self-assemble and for individuals to contribute to group effort without requiring formal management (and its attendant overhead), these tools have radically altered the old limits on the size, sophistication, and scope of unsupervised effort (the limits that created the institutional dilemma in the first place). They haven't removed them entirely—issues of complexity still loom large, as we will see—but the new tools enable alternate strategies for keeping that complexity under control. And as we would expect, when desire is high and costs have collapsed, the number of such groups is skyrocketing, and the kinds of effects they are having on the world are spreading.

The Tectonic Shift

For most of modern life, our strong talents and desires for group effort have been filtered through relatively rigid institutional structures because of the complexity of managing groups. We haven't had all the groups we've wanted, we've simply had all the groups we could afford. The old limits of

what unmanaged and unpaid groups can do are no longer in operation; the difficulties that kept self-assembled groups from working together are shrinking, meaning that the number and kinds of things groups can get done without financial motivation or managerial oversight are growing. The current change, in one sentence, is this: most of the barriers to group action have collapsed, and without those barriers, we are free to explore new ways of gathering together and getting things done.

George W.S. Trow, writing about the social effects of television in *Within the Context of No Context*, described a world of simultaneous continuity and discontinuity:

Everyone knows, or ought to know, that there has happened under us a Tectonic Plate Shift [...] the political parties still have the same names; we still have a CBS, an NBC, and a *New York Times*; but we are not the same nation that had those things before.

Something similar is happening today, with newer tools. Most of the institutions we had last year we will have next year. In the past the hold of those institutions on public life was irreplaceable, in part because there was no alternative to managing large-scale effort. Now that there is competition to traditional institutional forms for getting things done, those institutions will continue to exist, but their purchase on modern life will weaken as novel alternatives for group action arise.

This is not to say that corporations and governments are going to wither away. Though some of the early utopianism around new communications tools suggested that we were heading into some sort of posthierarchical paradise, that's not what's happening now, and it's not what's going to happen. None of the absolute advantages of institutions like businesses or schools or governments have disappeared. Instead, what has happened is that most of the *relative* advantages of those institutions have disappeared—relative, that is, to the direct effort of the people they represent. We can see signs of this in many places: the music industry, for one, is still reeling from the discovery that the reproduction and distribution of music, previously a valuable service, is now something their customers can do for themselves. The Belarusian government is trying to figure out how to keep its young people from generating spontaneous political protests. The Catholic Church is facing its first prolonged challenge from self-organized lay groups in its history. But these stories and countless others aren't just about something happening to particular businesses or governments or religions. They are also about something happening to the world.

Group action gives human society its particular character, and anything that changes the way groups get things done will affect society as a whole. This change will not be limited to any particular set of institutions or functions. For any given organization, the important questions are "When will the change happen?" and "What will change?" The only two answers we can rule out are never, and nothing. The ways in which any given institution will find its situation transformed will vary,

but the various local changes are manifestations of a single deep source: newly capable groups are assembling, and they are working without the managerial imperative and outside the previous strictures that bounded their effectiveness. These changes will transform the world everywhere groups of people come together to accomplish something, which is to say everywhere.

CHAPTER 2

SHARING ANCHORS COMMUNITY

Groups of people are complex, in ways that make those groups hard to form and hard to sustain; much of the shape of traditional institutions is a response to those difficulties. New social tools relieve some of those burdens, allowing for new kinds of group-forming, like using simple sharing to anchor the creation of new groups.

Imagine you are standing in line with thirty-five other people, and to pass the time, the guy in front of you proposes a wager. He's willing to bet fifty dollars that no two people in line share a birthday. Would you take that bet?

If you're like most people, you wouldn't. With thirty-six people and 365 possible birthdays, it seems like there would only be about a one-in-ten chance of a match, leaving you a 90 percent chance of losing fifty dollars. In fact, you should take the bet, since you would have better than an 80 percent chance of *winning* fifty dollars. This is called the Birthday Paradox (though it's not really a paradox, just a

surprise), and it illustrates some of the complexities involved in groups.

Most people get the odds of a birthday match wrong for two reasons. First, in situations involving many people, they think about themselves rather than the group. If the guy in line had asked, "What are the odds that someone in this line shares *your* birthday?" that would indeed have been about a one in ten chance, a distinctly bad bet. But in a group, other people's relationship to you isn't all that matters; instead of counting people, you need to count links between people. If you're comparing your birthday with one other person's, then there's only one comparison, which is to say only one chance in 365 of a match. If you're comparing birthdays in a group with two other people—you, Alice, and Bob, say—you might think you'd have two chances in 365, but you'd be wrong. There are three comparisons: your birthday with Alice's, yours with Bob's, and Alice's and Bob's with each other. With four people, there are six such comparisons, half of which don't involve you at all; with five, there are ten, and so on. By the time you are at thirty-six people, there are more than six hundred pairs of birthdays. Everyone understands that the chance of any two people in a group sharing a birthday is low; what they miss is that a count of "any two people" rises much faster than the number of people themselves. This is the engine of the Birthday Paradox.

This rapidly rising number of pairs is true of any collection of things: if you have a bunch of marbles, the number of possible pairs will be set by the same math. The growing complexity gets much more wretched in social settings, however; marbles don't have opinions, but people do. As a group grows

to even modest size, getting universal agreement becomes first difficult, then impossible. This quandary can be illustrated with a simple scenario. You and a friend want to go out to a movie. Before you buy the tickets, you'll have to factor in your various preferences: comedy or romance, early show or late, near work or near home. All of these will have some effect on your mutual decision, but with just two of you, getting to some acceptable outcome is fairly easy.

Now imagine that you and three friends decide to go out to a movie. This is harder, because the group's preferences are less likely to overlap neatly. Two of you love action films, two hate them; one wants the early show, three the late one, and so on. With two people, you have only one agreement to make. With four, as Birthday Paradox math tells us, you need six such agreements. Other things being equal, coordinating anything with a group of four is six times as hard as with two people,




Figure 2-1: Three clusters, with all connections drawn. The small cluster has 5 members and 10 connections; the middle one has 10 members and 45 connections; and the large one has 15 and 105. A group's complexity grows faster than its size.

and the effect gets considerably worse as the group grows even moderately large. By the time you want to go to a movie in a group of ten, waiting for forty-five separate agreements is pretty much a lost cause. You could sit around discussing the possible choices all day, with no guarantee you'll get to an agreement at all, much less in time for the movie. Instead you'll vote or draw straws, or someone will just decide to go to a particular movie and invite everyone else along, without trying to take all possible preferences into account. These difficulties have nothing to do with friendship or movie-going specifically; they are responses to the grim logic of group complexity.

This complexity means, in the words of the physicist Philip Anderson, that "more is different." Writing in *Science* magazine in 1972, Anderson noted that aggregations of anything from atoms to people exhibit complex behavior that cannot be predicted by observing the component parts. Chemistry isn't just applied physics—you cannot understand all the properties of water from studying its constituent atoms in isolation. This pattern of aggregates exhibiting novel properties is true of people as well. Sociology is not just psychology applied to groups; individuals in group settings exhibit behaviors that no one could predict by studying single minds. No one has ever been bashful or extroverted while sitting alone in their room, no one can be a social climber or a man of the people without reference to society, and these characteristics exist because groups are not just simple aggregations of individuals.

As groups grow, it becomes impossible for everyone to interact directly with everyone else. If maintaining a connection between two people takes any effort at all, at some size that effort becomes unsustainable. You can see this phenom-

enon even in simple situations, such as when people clink glasses during a toast. In a small group, everyone can clink with everyone else; in a larger group, people touch glasses only with those near them. Similarly, as Fred Brooks noted in his book *The Mythical Man-Month*, adding more employees to a late project tends to make it later, because the new workers increase the costs of coordinating the group. Because this constraint is so basic, and because the problem can never be solved, only managed, every large group has to grapple with it somehow. For all of modern life, the basic solution has been to gather people together into organizations.

We use the word "organization" to mean both the state of being organized and the groups that do the organizing—"Our organization organizes the annual conference." We use one word for both because, at a certain scale, we haven't been able to get organization without organizations; the former seems to imply the latter. The typical organization is hierarchical, with workers answering to a manager, and that manager answering to a still-higher manager, and so on. The value of such hierarchies is obvious—it vastly simplifies communication among the employees. New employees need only one connection, to their boss, to get started. That's much simpler than trying to have everyone talk to everyone.

Running an organization is difficult in and of itself, no matter what its goals. Every transaction it undertakes—every contract, every agreement, every meeting—requires it to expend some limited resource: time, attention, or money. Because of these transaction costs, some sources of value are too costly to take advantage of. As a result, no institution can put all its energies into pursuing its mission; it must expend

considerable effort on maintaining discipline and structure, simply to keep itself viable. Self-preservation of the institution becomes job number one, while its stated goal is relegated to number two or lower, no matter what the mission statement says. The problems inherent in managing these transaction costs are one of the basic constraints shaping institutions of all kinds.

This ability of the traditional management structure to simplify coordination helps answer one of the most famous questions in all of economics: If markets are such a good idea, why do we have organizations at all? Why can't all exchanges of value happen in the market? This question originally was posed by Ronald Coase in 1937 in his famous paper "The Nature of the Firm," wherein he also offered the first coherent explanation of the value of hierarchical organization. Coase realized that workers could simply contract with one another, selling their labor, and buying the labor of others in turn, in a market, without needing any managerial oversight. However, a completely open market for labor, reasoned Coase, would underperform labor in firms because of the transaction costs, and in particular the costs of discovering the options and making and enforcing agreements among the participating parties. The more people are involved in a given task, the more potential agreements need to be negotiated to do anything, and the greater the transaction costs, as in the movie example above.

A firm is successful when the costs of directing employee effort are lower than the potential gain from directing. It's tempting to assume that central control is better than markets for arranging all sorts of group effort. (Indeed, during the twentieth century much of the world lived under governments

that made that assumption.) But there is a strong limiting factor to this directed management, and that is the cost of management itself. Richard Hackman, a Harvard professor of psychology, has studied the size and effectiveness of work groups in *Leading Teams*. Hackman tells a story about a man who ran a nonprofit whose board of directors numbered forty. When asked what he thought such a large board could accomplish, he replied, "Nothing," in a way that implied he liked it that way. Because of managerial overhead, large groups can get bogged down, and whenever transaction costs become too expensive to manage within a single organization, markets outperform firms (and central management generally).

Activities whose costs are higher than the potential value for both firms and markets simply don't happen. Here is the institutional dilemma again: because the minimum costs of being an organization in the first place are relatively high, certain activities may have some value but not enough to make them worth pursuing in any organized way. New social tools are altering this equation by lowering the costs of coordinating group action. The easiest place to see this change is in activities that are too difficult to be pursued with traditional management but that have become possible with new forms of coordination.

How Did All Those Pictures Get There?

On the last Saturday in June, Coney Island kicks off the summer with the Mermaid Parade, a sort of hometown procession for New York City hipsters. Hundreds of people show up to

march around Brooklyn's famously run-down amusement park in costumes that are equal parts extravagant and weird—a giant red octopus puppet, a flotilla of hula-hooping mermaids, a marcher sporting a bikini top made of two skulls. Thousands turn out to watch and photograph the festivities, taking pictures ranging from a couple of snapshots to dozens of high-quality photos.

A handful of these pictures end up in local newspapers, but for most of the history of the Mermaid Parade, most pictures were seen only by the people who took them and a few of their friends. The sponsor of the parade didn't provide any way for the photographers to aggregate or share their photos, and the photographers themselves didn't spontaneously organize to do so. That is the normal state of affairs. Given the complexities of group effort, hundreds of people don't spontaneously do much of any consequence, and it wouldn't have made much sense for anyone to expend the effort to identify and coordinate the photographers from the outside. A couple of years ago, however, the normal state of affairs stopped operating.

In 2005, for the first time, a hundred or so of the attendees pooled thousands of their Mermaid Parade photos and made them publicly available online. The photos came from all sorts of photographers, from amateurs with camera-phones to pros with telephoto lenses. The group was mainly populated by casual contributors—most people uploaded fewer than a dozen photos—but a handful of dedicated contributors shared more than a hundred pictures each, and one user, going by the online name *czarina*, shared more than two hundred photos on her own. The group pooled these photos by uploading them to a service called Flickr, giving each of the photos a free-

form label called a tag. As a result, anyone can go to Flickr today, search for the tag "mermaidparade," and see the photos. This is a simple chain of events: people take pictures, people share pictures, you see pictures. It's so simple, in fact, that it's easy to overlook the substantial effort involved behind the scenes.

Flickr is the source of the sharing, but here's what Flickr did *not* do to get the sharing to happen: it didn't identify the Mermaid Parade as an interesting event, nor did it coordinate parade photographers or identify parade photographs. What it did instead was to let the users label (or tag) their photos as a way of arranging them. When two or more users adopted the same tag, those photos were automatically linked. The users were linked as well; the shared tag became a potential stepping-stone from one user to another, adding a social dimension to the simple act of viewing. The distinction between Flickr coordinating users versus helping them coordinate themselves seems minor, but it is in fact vital, as it is the only way Flickr can bear the costs involved. Consider what it would have taken for Flickr to organize hundreds of amateur mermaid photographers. Someone at Flickr HQ would have had to know about an obscure parade on the other side of the country. (Flickr is based in California.) They would have had to propose a tag for the group to use in order to assemble the uploaded photos. Finally, they would have had to communicate the chosen tag to everyone going to the parade.

This last step is especially hard. When you are trying to address a diffuse group, you are locked into the dilemma that all advertisers face: how do you reach the people you want, without having to broadcast your message to everybody?

People in the category “Potential photographer of the Mermaid Parade” aren’t easy to find. Flickr couldn’t have known in advance who would go to the parade. Instead, they would have to send messages out to many more people than would actually attend, in hopes of reaching the right audience, advertising to photographers, hipsters, New Yorkers, and so on, in hope of getting the tiny fraction of those groups who would actually go. Most such ads would be seen by people who weren’t going to the parade, while most of the people who were going wouldn’t see (or pay attention to) the ads. Given those obstacles, no business in the world would take on the job. The profit motive is little help; no one could sell enough pictures, even the skull-bikini ones, to be able to pay the photographers, much less leave any profit afterward. Likewise, no nonprofit or government agency would touch the problem; even the porkiest of pork-barrel projects isn’t going to cover publicity for hula-hooping mermaids. The gap between effort and payoff is too large for any institution to span.

Yet there the photos are. Without spending any serious effort on any individual set of photos, and without doing anything to coordinate or even identify groups of photographers, Flickr has provided a platform for the users to aggregate the photos themselves.

The difference between the value of the photos and the cost of aggregation is a general one. Flickr isn’t just for photos of dancing mermaids, family reunions, and the effects of that third margarita; it also hosts photos of broad public interest. Flickr provided some of the first photos of the London Transport bombings in 2005, including some taken with camera-phones by evacuees in the Underground’s tunnels.

Flickr beat many traditional news outlets by providing these photos, because there were few photojournalists in the affected parts of the transport network (three separate trains on the Underground, and a bus), but many people near those parts of the transport system had camera-phones that could e-mail the pictures in. Having cameras in the hands of amateurs on the scene was better than having cameras in the hands of professionals who had to travel.

The photos that showed up after the bombings weren’t just amateur replacements for traditional photojournalism; people did more than just provide evidence of the destruction and its aftermath. They photographed official notices (“All Underground services are suspended”), notes posted in schools (“Please do not inform children of the explosions”), messages of support from the rest of the world (“We love you London”), and within a day of the bombings, expressions of defiance addressed to the terrorists (“We are not afraid” and “You will fail”). Not only did Flickr host all of these images, they made them available for reuse, and bloggers writing about the bombings were able to use the Flickr photos almost immediately, creating a kind of symbiotic relationship among various social tools. The images also garnered comments on the Flickr site. A user going by Happy Dave posted an image reading “I’m OK,” meant to alert his friends who had subscribed to his images on Flickr; he received dozens of comments from well-wishers in the comments. The “Do not inform the children” image generated a conversation about how to talk to kids about terrorism. The basic capabilities of tools like Flickr reverse the old order of group activity, transforming “gather, then share” into “share, then gather.” People

were able to connect after discovering one another through their photos.

A similar change in the broadcasting of evidence happened after the awful destruction caused by the Indian Ocean tsunami at the end of 2004. Within hours of the tsunami dozens of photos were available on the Web showing various affected places, and within days there were hundreds. As with the London bombings, there was no way to get photojournalists on the scene instantly, but here the problem was not just the speed of response but the spread of the damage, which affected thirteen countries. And as with the London bombings, the photos weren't used just for evidence; people began uploading photos of missing loved ones, and various weblogs began to syndicate these photos to aid in relocation. The most visited photo tagged "tsunami" is a picture of a little boy, age two at the time he went missing. The picture originally went up with contact information to aid in the search, but as time went on, it turned into an ongoing memorial; viewers posted hundreds of comments of support and prayers under the photo, and many commenters came back months later to check in and conversed with one another in the comments. When the boy's body was finally recovered and identified, months later, several people posted the sad news on Flickr, and the community that had formed around the photo posted expressions of grief and condolences for the family, then dissolved.

Flickr also helped provide the world with photographic documentation of the 2006 military coup in Thailand. Immediately after the coup the military placed restrictions on reporting by the media, but it didn't (and probably couldn't) place similar restrictions on the whole populace. As a result,

many of the earliest photos of tanks in front of Government House, the parliament building, came from individuals posting images from ordinary digital cameras, and they were discoverable by their tags (Bangkok, Thailand, Military, Coup). One of those users was Alisara Chirapongse, a fashion-obsessed college student going by the name gnarlykitty, who posted the coup photos to her weblog, along with running commentary on the cause and immediate aftermath of the army overthrowing Thaksin Shinawatra, then prime minister. As the army announced that it wanted to take control of communications and ban public political speech, her posts took on a new urgency:

One new little change that this law brought us is the whole new level of censorship. No political gathering, no discussing politics, and of course no voicing your opinions whatsoever about the whole mumbo jumbo coup. (Oops did I just do that?)

Alisara posted links to Wikipedia, the collaboratively produced encyclopedia, which was acting as a clearinghouse for breaking news of the coup (as is now usual). She also pointed her readers to a petition to restore freedom of speech and to a proposed demonstration, which she later attended and photographed.

Then as the initial disorientation of the coup gave way to the new normal, Alisara went back to her life as a fashion-obsessed student. As she put it,

This blog is my personal blog where I usually write things concerning my life and things I like. Since my

life is lived here in Bangkok Thailand, it should come as no surprise to anyone that I sometime blog about it. So blogging about the Coup is merely blogging about something that's currently happening in my country.

The rest of that post was about a night she spent at a club, and the post after that was about how much she likes her new camera-phone. She wasn't a full-time journalist, she was a citizen with a camera and a weblog, but she had participated in a matter of global significance at exactly the time when the traditional media were being silenced.

The content in these examples is quite varied—the gentle ridiculousness of the Mermaid Parade and the awful seriousness of the London bombings; the man-made intervention of a military coup and the natural destruction of the tsunami. The common thread is the complexity of gathering the photos. The groups of photographers were all latent groups, which is to say groups that existed only *in potentia*, and too much effort would have been required to turn those latent groups into real ones by conventional means. The mermaid photos were too unimportant to be worth any institutional effort. The London bombing photos were taken by the people on the scene. The tsunami's destruction was spread out over tens of thousands of miles of coastland, and the uses of photos included finding missing persons, something outside the purview of typical newsgathering. During the Thai coup the military rulers were able to place restrictions on organized media, giving amateur photographers an advantage in providing views of tanks in the streets. In each of those cases the cost of coordinating the

potential photographers would have defeated any institution wanting to put photos together quickly and make them available globally.

The task of aggregating and making photos available is nothing like, say, the task of putting a man on the moon. Prior to services like Flickr, what kept photo-sharing from happening wasn't the absolute difficulty but the relative difficulty. There is obviously some value to both photographers and viewers in having photos available, but in many cases that value never exceeded the threshold of cost created by the institutional dilemma. Flickr escaped those problems, not by increasing its managerial oversight over photographers but by abandoning any hope of such oversight in the first place, instead putting in place tools for the self-synchronization of otherwise latent groups.

Making the Trains Run on Time

The structure of traditional managerial oversight is often illustrated by an "org chart," a diagram of the official organizational hierarchy. This chart is the simplest possible view of an organization's reporting structure. It is usually drawn as an inverted tree of boxes and arrows. The box at the top represents the head of the organization; the lines drawn downward from that box connect her to various officers and vice presidents through the layers of management, until, at the bottom, there are the rank and file, represented by boxes with lines connecting upward but not downward. The org chart diagrams both responsibility and channels of communication—when

two boxes are connected on such a chart, the upper box is the boss; communication from the CEO flows down through the layers of management, while information from the workers flows up in the same way. Compared to the chaos of the market, the org chart draws clear and obvious lines of responsibility, and it is that very clarity that allows the firm to outperform a pure market for work.

The org chart is like institutional wallpaper—ubiquitous and not terribly dramatic. It's funny to think of it as a specific invention, but its existence and form owe quite a lot to the environment in which it was first widely used—railroad management in the 1800s. The pioneering managerial methods were meticulously documented by Alfred Chandler in his book *The Visible Hand*. The principal problem in running a railroad was arranging for eastbound and westbound trains to share the same track, because it was prohibitively expensive to lay more than one track for a particular line. By 1840 Western Railroad, a pioneer in building longer rail lines, had to deal with a dozen trains crossing in opposite directions every day. That situation created obvious safety risks, risks that were not long in moving from the theoretical to the real: on October 5, 1841, two passenger trains collided head on, with two fatalities and seventeen injured. This accident alarmed both the public and Congress and forced the railroads to rethink their management.

For the next fifteen years railroads invested in better oversight. As a result, their safety record improved, but their profitability decayed. A big firm like Western could haul more people and cargo to more places than could a smaller railroad, but the cost of managing the enterprise had risen much faster; Western was actually making less money per mile of track than its

smaller competitors. David McCallum, a railroad superintendent for the New York & Erie Railroad, proposed both an explanation and a solution for this decayed profitability. As he put it in his Superintendent's Report of 1855:

A Superintendent of a road fifty miles in length can give its business his personal attention, and may be constantly engaged in the direction of its details . . . any system however imperfect, may under such circumstances, prove comparatively successful.

In the government of a road five hundred miles in length, a very different state exists. Any system which might be applicable to the business and extent of a short road, would be found entirely inadequate to the wants of a long one.

More is different: a small railroad could function with ad hoc management, because it had so few employees and so few passing trains, but as the scale rose, the management problems rose faster. This is where the institutional dilemma meets Birthday Paradox math: not only does managing resources take resources, but management challenges grow faster than organizational size.

McCallum's proposed solution to this dilemma included making a clear delineation of the responsibility for different segments of track. Central management would oversee regional divisions and supervise the trains passing through their region. McCallum introduced several formal innovations to New York & Erie: strong hierarchical oversight, including an explicitly divisional organization of the railroad with different

superintendents responsible for different parts of the railroad. He diagrammed this form of organization with what may have been the first commercial org chart in history. This method was widely copied by other railroads, then by other kinds of firms.

In addition to revolutionizing management structure, McCallum wrote six principles for running a hierarchical organization. Most are what you'd expect (number one was ensuring a "proper division of responsibilities"), but number five is worth mentioning: his management system was designed to produce "such information, to be obtained through a system of daily reports and checks, that will not embarrass principal officers nor lessen their influence with their subordinates." If you have ever wondered why so much of what workers in large organizations know is shielded from the CEO and vice versa, wonder no longer: the idea of limiting communications, so that they flow only from one layer of the hierarchy to the next, was part of the very design of the system at the dawn of managerial culture.

Post-Managerial Organization

When an organization takes on a task, the difficulty of coordinating everyone needs to be reined in somehow, and the larger the group, the more urgent the need. The standard, almost universal solution is to create a hierarchy and to slot individuals into that organization by role. In Coasean terms, McCallum's system lowered the transaction costs of running a railroad by increasing managerial structure. This approach greatly simpli-

fies lines of responsibility and communication, making even very large organizations manageable. The individuals in such an organization have to agree to be managed, of course, which is usually achieved by paying them, and by making continued receipt of their pay contingent on their responsiveness to their manager's requests.

An organization will tend to grow only when the advantages that can be gotten from directing the work of additional employees are less than the transaction costs of managing them. Coase concentrated his analysis on businesses, but the problems of coordination costs apply to institutions of all sorts. The Catholic Church and the U.S. Army are as hierarchical as any for-profit business, and for many of the same reasons. The layers of structure between the pope and the priests, or between the president and the privates, is a product of the same forces as the layers between the general superintendent and a conductor on the New York & Erie. This hierarchical organization reduces transaction costs, but it doesn't eliminate them.

Imagine a company with fifteen hundred employees, where each manager is responsible for half a dozen people. The CEO has six vice presidents, who each direct the work of six supervisors, and so on. Such a company would have three layers of management between the boss and the workers. If you want to bring the workers closer to the boss, you will have to increase the number of workers that each manager is responsible for. This will reduce the number of layers but will also reduce average management time with each staff member (or force everyone to spend more hours per day communicating with one another). When an organization grows very

large, it reaches the limit implicit in Coase's theory; at some point an institution simply cannot grow anymore and still remain functional, because the cost of managing the business will destroy any profit margin. You can think of this as a Coasean ceiling, the point above which standard institutional forms don't work well.

Coase's theory also tells us about the effects of small changes in transaction costs. When such costs fall moderately, we can expect to see two things. First, the largest firms increase in size. (Put another way, the upper limit of organizational size is inversely related to management costs.) Second, small companies become more effective, doing more business at lower cost than the same company does in a world of high transaction costs. These two effects describe the postwar industrial world well: Giant conglomerates like ITT in the 1970s and GE in recent years used their management acumen to get into a huge variety of businesses, simply because they were good at managing transaction costs. At the same time there has been an explosion of small- and medium-sized businesses, because such businesses were better able to discover and exploit new opportunities.

But what if transactions costs don't fall moderately? What if they collapse? This scenario is harder to predict from Coase's original work, and it used to be purely academic. Now it's not, because it's happening, or rather it has already happened, and we're starting to see the results.

Anyone who has worked in an organization with more than a dozen employees recognizes institutional costs. Anytime you are faced with too many meetings, too much paperwork, or too many layers of approval (shades of

McCallum), you are dealing with those costs. Until recently, such costs have been little more than the stuff of water-cooler grumbling—everyone complains about institutional overhead, without much hope of changing things. In that world (the world we lived in until recently), if you wanted to take on a task of any significance, managerial oversight was just one of the costs of doing business.

What happens to tasks that aren't worth the cost of managerial oversight? Until recently, the answer was "Those things don't happen." Because of transaction costs a long list of possible goods and services never became actual goods and services; things like aggregating amateur documentation of the London transit bombings were simply outside the realm of possibility. That collection now exists because people have always desired to share, and the obstacles that prevented sharing on a global scale are now gone. Think of these activities as lying under a Coasean floor; they are valuable to someone but too expensive to be taken on in any institutional way, because the basic and unsheddable costs of being an institution in the first place make those activities not worth pursuing.

Our basic human desires and talents for group effort are stymied by the complexities of group action at every turn. Coordination, organization, even communication in groups is hard and gets harder as the group grows. That difficulty means that whatever methods help coordinate group action will spread, no matter how inefficient they are, so long as they are better than nothing. Small groups have several methods for coordinating action, like calling each group member in turn or setting up a phone tree, but most of these methods don't work well even for dozens of people, much less for thousands. For large-scale

activity, the methods that have worked best have been those pioneered by McCallum—hierarchical organization, managed in layers. The most common organizational structures we have today are simply the least bad fit for group action in an environment of high transaction costs.

Our new tools offer us ways of organizing group effort without resorting to McCallum's strategies. Flickr stands in a different kind of relationship to its photographers than a newspaper does. Where a newspaper is in the business of directing the work of photographers, Flickr is simply a platform; whatever coordination happens comes from the users and is projected onto the site. This is odd. We generally regard institutions as being capable of more things than uncoordinated groups are, precisely because they are able to direct their employees. Here, though, we have a situation where the loosely affiliated group can accomplish something more effectively than the institution can. Thanks to the introduction of user-generated labeling, the individual motivation of the photographers—devoid of financial reward—is now enough to bring vast collections of photos into being. These collections didn't just happen to be put together without an institution; that is the only way they could have been put together.

This is where Coasean logic gets strange. Small decreases in transaction costs make businesses more efficient, because the constraints of the institutional dilemma get less severe. Large decreases in transaction costs create activities that can't be taken on by businesses, or indeed by any institution, because no matter how cheap it becomes to perform a particular activity, there isn't enough payoff to support the cost incurred by being an institution in the first place. So long as the abso-

lute cost of organizing a group is high, unmanaged groups will be limited to undertaking small efforts—a night out at the movies, a camping trip. Even something as simple as a potluck dinner typically requires some hosting institution. Now that it is possible to achieve large-scale coordination at low cost, a third category has emerged: serious, complex work, taken on without institutional direction. Loosely coordinated groups can now achieve things that were previously out of reach for any other organizational structure, because they lay under the Coasean floor.

The cost of all kinds of group activity—sharing, cooperation, and collective action—have fallen so far so fast that activities previously hidden beneath that floor are now coming to light. We didn't notice how many things were under that floor because, prior to the current era, the alternative to institutional action was usually no action. Social tools provide a third alternative: action by loosely structured groups, operating without managerial direction and outside the profit motive.

From Sharing to Cooperation to Collective Action

For the last hundred years the big organizational question has been whether any given task was best taken on by the state, directing the effort in a planned way, or by businesses competing in a market. This debate was based on the universal and unspoken supposition that people couldn't simply self-assemble; the choice between markets and managed effort assumed that there was no third alternative. Now there is.

Our electronic networks are enabling novel forms of collective action, enabling the creation of collaborative groups that are larger and more distributed than at any other time in history. The scope of work that can be done by noninstitutional groups is a profound challenge to the status quo.

The collapse of transaction costs makes it easier for people to get together—so much easier, in fact, that it is changing the world. The lowering of these costs is the driving force underneath the current revolution and the common element to everything in this book. We're not used to thinking of "groupness" as a specific category—the differences between a college seminar and a labor union seem more salient than their similarities. It's hard to see how Evan Guttman's quest for the return of the mobile phone is the same kind of thing as the distributed documentation of the Indian Ocean tsunami. But like a chain of volcanoes all fed by the same pool of magma, the surface manifestations of group efforts seem quite separate, but the driving force of those eruptions is the same: the new ease of assembly. This change can be looked at as one long transition, albeit one with many manifestations, unfolding at different speeds in different contexts. The transition can be described in basic outline as the answer to two questions: Why has group action largely been limited to formal organizations? What is happening now to change that?

We now have communications tools—and increasingly, social patterns that make use of those tools—that are a better fit for our native desires and talents for group effort. Because we can now reach beneath the Coasean floor, we can have groups that operate with a birthday party's informality and a multinational's scope. What we are seeing, in the amateur

coverage of the Thai coup and the tsunami documentation and the struggle over Ivanna's phone and countless other examples, is the beginning of a period of intense experimentation with these tools. The various results look quite different from one another, and as we get good at using the new tools, those results will diverge still further. New ease of assembly is causing a proliferation of effects, rather than a convergence, and these effects differ by how tightly the individuals are bound to one another in the various groups.

You can think of group undertaking as a kind of ladder of activities, activities that are enabled or improved by social tools. The rungs on the ladder, in order of difficulty, are sharing, cooperation, and collective action.

Sharing creates the fewest demands on the participants. Many sharing platforms, such as Flickr, operate in a largely take-it-or-leave-it fashion, which allows for the maximum freedom of the individual to participate while creating the fewest complications of group life. Though Flickr sets public sharing as the default, it also allows users to opt to show photos only to selected users, or to no one. Knowingly sharing your work with others is the simplest way to take advantage of the new social tools. (There are also ways of unknowingly sharing your work, as when Google reads the linking preferences of hundreds of millions of internet users. These users are helping create a communally available resource, as Flickr users are, but unlike Flickr, the people whose work Google is aggregating aren't actively choosing to make their contributions.)

Cooperation is the next rung on the ladder. Cooperating is harder than simply sharing, because it involves changing your behavior to synchronize with people who are changing their

behavior to synchronize with you. Unlike sharing, where the group is mainly an aggregate of participants, cooperating creates group identity—you know who you are cooperating with. One simple form of cooperation, almost universal with social tools, is conversation; when people are in one another's company, even virtually, they like to talk. Sometimes the conversation is with words, as with e-mail, IM, or text messaging, and sometimes it is with other media: YouTube, the video sharing site, allows users to post new videos in response to videos they've seen on the site. Conversation creates more of a sense of community than sharing does, but it also introduces new problems. It is famously difficult to keep online conversations from devolving into either name-calling or blather, much less to keep them on topic. Some groups are perfectly happy with those effects (indeed, there are communities on the internet that revel in puerile or fatuous conversation), but for any group determined to maintain a set of communal standards some mechanism of enforcement must exist.

Collaborative production is a more involved form of cooperation, as it increases the tension between individual and group goals. The litmus test for collaborative production is simple: no one person can take credit for what gets created, and the project could not come into being without the participation of many. Structurally, the biggest difference between information sharing and collaborative production is that in collaborative production at least some collective decisions have to be made. The back-and-forth talking and editing that makes Wikipedia work results in a single page on a particular subject (albeit one that changes over time). Collaboration is not an absolute good—many tools work by *reducing* the

amount of required coordination, as Flickr does in aggregating photos. Collaborative production can be valuable, but it is harder to get right than sharing, because anything that has to be negotiated about, like a Wikipedia article, takes more energy than things that can just be accreted, like a group of Flickr photos.

Collective action, the third rung, is the hardest kind of group effort, as it requires a group of people to commit themselves to undertaking a particular effort together, and to do so in a way that makes the decision of the group binding on the individual members. All group structures create dilemmas, but these dilemmas are hardest when it comes to collective action, because the cohesion of the group becomes critical to its success. Information sharing produces shared awareness among the participants, and collaborative production relies on shared creation, but collective action creates shared responsibility, by tying the user's identity to the identity of the group. In historical terms, a potluck dinner or a barn raising is collaborative production (the members work together to create something), while a union or a government engages in collective action, action that is undertaken in the name of the members meant to change something out in the world, often in opposition to other groups committed to different outcomes.

The commonest collective action problem is described as the "Tragedy of the Commons," biologist Garrett Hardin's phrase for situations wherein individuals have an incentive to damage the collective good. The Tragedy of the Commons is a simple pattern to explain, and once you understand it, you come to see it everywhere. The standard illustration of the problem uses sheep. Imagine you are one of a group of shepherds who

graze their sheep on a commonly owned pasture. It's obviously in everyone's interest to keep the pasture healthy, which would require each of you to take care that your sheep don't overgraze. As long as everyone refuses to behave greedily, everyone benefits. There is just one problem with this system: "everyone" doesn't take your sheep to market. You do. Your incentive, as an individual shepherd, is to minimize the cost of raising the fattest possible sheep. Everyone benefits from you moderating your sheep's consumption of grass, but you would benefit from free riding, which is to say letting them eat as much free grass as they possibly could.

Once you have this realization, you can still refrain from what would ultimately be a ruinous strategy, on the grounds that it would be bad for everyone else. Then another, even more awful thought strikes you: every other shepherd will have the same realization, and if even one of them decides to overgraze, all your good works will only end up subsidizing them. Seen in this light, the decision not to overgraze is provisional on everyone else making the same decision, which makes it very fragile indeed. The minute one of the other shepherds keeps his sheep out in the pasture an hour longer than necessary, the only power you have is to retaliate by doing the same. And this is the Tragedy of the Commons: while each person can agree that all would benefit from common restraint, the incentives of the individuals are arrayed against that outcome.

People who benefit from a resource while doing nothing in recompense are free riders. Societies have generally dealt with the problem of free riders in one of two ways. The first way is elimination of the commons, transferring ownership of parts of it to individuals, all of whom have an incentive to

protect their own resources. If six shepherds each own one-sixth of the former commons, the overgrazing problem is a personal one, not a social one. If you overgraze your section, you will suffer the future consequences, while your neighbor will not. The second way is governance or, as Hardin puts it, "mutual coercion, mutually agreed upon." This solution prevents the individual actors from acting in their own interests rather than in the interests of the group. The Tragedy of the Commons is why taxes are never voluntary—people would opt out of paying for road maintenance if they thought their neighbors would pay for it. It's also why restaurants often add an automatic tip for large parties—when enough people are eating, everyone feels comfortable underfunding the group's tip, even if only unconsciously.

Collective action involves challenges of governance or, put another way, rules for losing. In any group that is determined to take collective action, different members of the group will express different opinions. Whenever a decision is taken on behalf of the group, at least some members won't get their way, and the bigger the group is, or the more decisions are made, the more often this will happen. For a group to take collective action, it must have some shared vision strong enough to bind the group together, despite periodic decisions that will inevitably displease at least some members. For this reason collective action is harder to arrange than information sharing or collaborative creation. In the current spread of social tools, real examples of collective action—where a group acts on behalf of, and with shared consequences for, all of its members—are still relatively rare.

The essential advantage created by new social tools has

been labeled “ridiculously easy group-forming” by the social scientist Seb Paquet. Our recent communications networks—the internet and mobile phones—are a platform for group-forming, and many of the tools built for those networks, from mailing lists to camera-phones, take that fact for granted and extend it in various ways. Ridiculously easy group-forming matters because the desire to be part of a group that shares, cooperates, or acts in concert is a basic human instinct that has always been constrained by transaction costs. Now that group-forming has gone from hard to ridiculously easy, we are seeing an explosion of experiments with new groups and new kinds of groups.

CHAPTER 3

EVERYONE IS A MEDIA OUTLET

Our social tools remove older obstacles to public expression, and thus remove the bottlenecks that characterized mass media. The result is the mass amateurization of efforts previously reserved for media professionals.

My uncle Howard was a small-town newspaperman, publishing the local paper for Richmond, Missouri (population 5,000). The paper, founded by my grandfather, was the family business, and ink ran in Howard’s blood. I can still remember him fulminating about the rise of *USA Today*; he criticized it as “TV on paper” and held it up as further evidence of the dumbing down of American culture, but he also understood the challenge that *USA Today* presented, with its color printing and national distribution. The *Richmond Daily News* and *USA Today* were in the same business; even with the difference in scale and scope, Howard immediately got what *USA Today* was up to.

Despite my uncle’s obsession, *USA Today* turned out to be

CHAPTER 8

SOLVING SOCIAL DILEMMAS

There are real and permanent social dilemmas, which can only be optimized for, never completely solved. The human social repertoire includes many such optimizations, which social tools can amplify.

Let's say, for the sake of illustration, that you and I went out for a few drinks last Saturday night, and at around 2 a.m. one of us said, "Hey, I know! Let's steal a car!" (I think it was you who said that.) So we steal a car, one thing leads to another, mistakes are made, and half an hour later we crash right through the window of a store. We barely have time to jump out and pretend to be bystanders before the police arrive.

Now the police aren't really buying the bystander alibi, but they don't have any other witnesses, so they take us off into separate rooms for questioning. Once we are separated, they make each of us this offer: "Look, we think you're innocent, but we suspect the other person in the car was responsible. If you tell us what you know about them, we'll give you a big reward,

and file charges against them. But you gotta tell us right now, and if you don't, we're going to hold you overnight." Since each of us is getting this offer, it creates four possibilities:

1. We each stick to our stories, they've got no evidence, and they keep us both overnight.
2. I stick to the bystander story and you turn me in. You get a reward, while I get charged.
3. I turn you in while you stick to the story. I get a reward, while you get charged.
4. We turn each the other in. We both get charged.

So knowing that I face the same choice as you—sticking to my story or turning you in—what do you do?

The worst outcome would clearly be getting charged with a crime, and the best outcome would be getting the reward. You know that I know that too, and if we both try to get the reward, we both get charged. The second best outcome is spending the night in jail, but you know that I know that too, and if you stick to your story in an attempt to get this outcome, I can go for the reward by turning you in. Similarly, if I stick to my story in an attempt to get the night in jail, you can turn me in to try to get the reward, but if we both try to get the reward, we both get charged—back to the worst outcome again.

This is a simplified version of the Prisoners' Dilemma, a social science thought experiment about how people make decisions. (The payoff matrix is bit more complex in the standard version, but the dilemma is the same.) Assuming that the two people can't communicate with each other and don't trust each other (about which more in a moment), the worst

outcome—number four—is the rational one, an outcome called a Nash equilibrium. The dilemma of the Prisoners' Dilemma is that, because it is a one-off transaction in which you and I can't communicate with each other, we can't coordinate any outcome better than the dismal Nash equilibrium. (This is the same math underlying the Tragedy of the Commons, where the Nash equilibrium encourages individual defection, even as it damages the group.) Things change, though, when the prisoners interact with each other repeatedly, a version called an iterated Prisoners' Dilemma.

Robert Axelrod, a sociologist at the University of Michigan who studied the iterated version extensively, staged tournaments for different software programs emulating the prisoners. Each program was given a strategy for when to cooperate and when to defect (the same two choices you and I faced in our notional interrogation rooms). These strategies were measured by adding or deducting points for the various outcomes. After running the tournament with many different participating strategies, ranging from "always defect" to "cooperate or defect at random," Axelrod found that a single strategy, called Tit-for-Tat, was most successful against every other strategy tried. Tit-for-Tat started by trying to cooperate the first time it was paired with any other program. If that program also cooperated, then Tit-for-Tat would offer to cooperate in the next round, and so on. As long as another program offered to cooperate, Tit-for-Tat would continue to do so as well. If the other program defected, though, taking advantage of Tit-for-Tat's trusting behavior, then Tit-for-Tat would defect against that program in the next round, effectively punishing the

other program as a way of communicating that its trusting nature extended only to those who reciprocate.

This strategy is a highly simplified version of real life—the more general lesson is that people who interact with one another repeatedly communicate through their actions, introducing what Axelrod calls "the shadow of the future." We all face the Prisoners' Dilemma whenever we interact with people we could take advantage of, or people who could take advantage of us, yet actually manage to trust one another often enough to accomplish things in groups. The shadow of the future makes it possible for me to act on your behalf today, even at some risk or cost to me, on the expectation that you will remember and reciprocate tomorrow.

New Tools to Create Social Capital

Over on University Place in lower Manhattan, a few blocks from my office, is the local bowling alley. Bowling often conjures up an era of picket fences and twenty-five-cent Cokes, and our local bowling emporium even has a name reminiscent of that time—Bowlmor Lanes. On any given Friday night, though, Bowlmor is very much an institution of the moment, catering to martini-sipping twentysomethings instead of factory workers unwinding with a beer. Through the decades bowling has been persistently reinvented, and it remains a durably popular activity. But between the 1950s and now there has been one significant change—a precipitous decline in league bowling, with its memberships and seasons and uniforms and all the

rest. Though plenty of groups bowl at Bowlmor Lanes, they are mainly people who already know one another; the bowling is more a consequence of group interaction than a source of it. The gradual disappearance of bowling leagues is one of many reductions in social mechanisms whereby people may be introduced to one another as a consequence of shared activity. This doesn't matter much for the fate of Bowlmor Lanes—a customer is a customer, league or no—but it may matter for the country.

When Robert Putnam, a Harvard sociologist, published *Bowling Alone* in 2000, it was an immediate sensation. His account of the weakening of community in the United States, based on a huge number of indicators from the decline of picnicking to the abandonment of league bowling, offered two provocative observations. First, much of the success of the United States as a nation has had to do with its ability to generate social capital, that mysterious but critical set of characteristics of functioning communities. When your neighbor walks your dog while you are ill, or the guy behind the counter trusts you to pay him next time, social capital is at work. It is the shadow of the future on a societal scale. Individuals in groups with more social capital (which is to say, more habits of cooperation) are better off on a large number of metrics, from health and happiness to earning potential, than those in groups with less social capital. Societies characterized by a high store of social capital overall do better than societies with low social capital on a similarly wide range of measurements, from crime rate to the costs of doing business to economic growth.

This is the shadow of the future at work: direct reciprocity assumes that if you do someone a favor today, that person will

do you a favor tomorrow. Indirect reciprocity is even more remarkable—it assumes that if you do someone in your community a favor today, someone in your community will be around to do you a favor tomorrow, even if it isn't the same person. The set of norms and behaviors that instantiates the shadow of the future is social capital, a set of norms that facilitate cooperation within or among groups.

It was Putnam's second observation, however, that generated the real reaction. Across a remarkably broad range of measures, participation in group activities, the vehicle for creating and sustaining social capital, was on the decline in the United States. Putting the two observations together, he concluded that one of the greatest assets in the growth and stability of the United States was ebbing away. One cause of the decline in social capital was a simple increase in the difficulty of people getting together—an increase in transaction costs, to use Coase's term. When an activity becomes more expensive, either in direct costs or increased hassle, people do less of it, and several effects of the last fifty years—including smaller households, delayed marriage, two-worker families, the spread of television, and suburbanization—have increased the transaction costs for coordinating group activities outside work. For most people the only possible reaction to Putnam's conclusion was nostalgia for a lost world of Rotary clubs and ice cream socials. One person, though, took it as an opportunity. In the 1990s Scott Heiferman had founded and sold a successful web business in New York City, and he was looking for his next business idea when he read *Bowling Alone*. Instead of regarding it as news of an inevitable decline, he set about trying to reinvigorate the creation of social capital through

real-world interaction. The solution he came up with was surprisingly simple.

First Heiferman assumed that people knew what they were missing and would want it back if they could get it; in an era of declining social capital, people would take steps to increase their communal participation if someone could make it easy again. Second, he recognized that treating the internet as some sort of separate space—cyberspace, as it was often called—was part of the problem. That word, coined by William Gibson in his novel *Neuromancer*, refers to a kind of alternate reality mediated by the world's communications networks. The cyberspace of *Neuromancer* is a visual representation of all the world's data; John Perry Barlow, a digital rights activist, later used the word to refer to the social spaces of the internet. Whether visual or social, though, the basic sense of cyberspace was that it was a world separate and apart from the real world. The predicted end point of this process was a progressive disassociation of social life from real space, leading to the death of cities as the population spread out to more bucolic spots.

The assumption that communications tools are (or will someday be) a good substitute for travel assumes that people mainly gather together for utilitarian reasons of sharing information. Companies have been selling us this idea since the invention of the telegraph, and AT&T's famous Picturephone, first launched at the 1964 World's Fair, was pitched as a way to reduce the need for travel. This reduction did not happen, not in 1964 or ever. If communication were a substitute for travel, then the effects would have shown up by now, but they haven't. In 1978 President Carter deregulated the airlines, causing travel prices to fall, but telecommunications stocks

didn't collapse; they rose. Similarly, in 1984 Judge Harold Greene broke up AT&T, leading to a rapid decrease in long-distance phone call costs; airline customers increased that year. Communication and travel are complements, not substitutes. Chris Meyer, a globe-trotting consultant for the Monitor Group, observes that "better communications make it easier for me to keep in touch with the office, so I spend *more* time on the road, talking to clients."

We gather together because we like to, and because it is useful. Assuming that videophones or e-mail or virtual reality will reduce the overall amount of travel is like assuming that liquor stores will kill bars, since liquor stores sell drinks much more cheaply than bars do. In fact, the reason people go to bars is not simply to get a drink, but to do so in a convivial environment. Similarly, cities don't exist just because people have had to be nearby to communicate; cities exist because people like to be near other people, and it is this fact, rather than the mere trading of information, that creates social capital. (Anyone who predicts the death of cities has already met their spouse.) This obvious human preference was overlooked during the early public spread of the internet, in large part because the average user interacted with different people online and offline.

What seemed like a deep social change in the 1990s was revealed to be a temporary accident by the year of Meetup's founding. The idea of cyberspace made sense when the population of the internet had a few million users; in that world social relations online really were separate from offline ones, because the people you would meet online were different from the people you would meet offline, and these worlds would rarely overlap. But that separation was an accident of

partial adoption. Though the internet began to function in its earliest form in 1969, it was not until 1999 that any country had a majority of its citizens online. (Holland was first, but that condition now applies to most countries in the developed world.) In the developed world, the experience of the average twenty-five-year-old is one of substantial overlap between online and offline friends and colleagues. The overlap is so great, in fact, that both the word and the concept of “cyberspace” have fallen into disuse. The internet augments real-world social life rather than providing an alternative to it. Instead of becoming a separate cyberspace, our electronic networks are becoming deeply embedded in real life.

Heiferman realized that if enough people are online, you don't have to group them solely by affinity (pug lovers, White Stripes fans, libertarians, whatever). Instead you can group them by affinity *and* proximity (pug lovers in Poughkeepsie, White Stripes fans in Walla Walla). He designed Meetup to help people find each other online and then meet in the real world, taking the burden of coordination off the hands of the potential users. Meetup users can search by interest (Are there any relevant Meetups in my town?) or they can look by area (I live in Milwaukee, what Meetups are nearby?)

By registering people's interests and location, Meetup can identify latent groups and help them come together. Heiferman bet that all over the United States (and later, the world) latent groups would be happy to get together if someone solved the coordination problem. Armed with this intuition (and the work of a talented group of programmers and designers), he launched the service. In early talks to potential users or inves-

tors he sometimes presented Meetup as a kind of time machine, reinvigorating classic American interest groups—people who shared an interest in bowling, cars, or Chihuahuas. (He talked about people who liked Chihuahuas so often, in fact, that it became a trademark bit of his spiel.)

The groups that actually ended up using Meetup didn't look anything like Heiferman expected. Here's the list of the fifteen most active Meetups the year after the site launched:

Topic	Total Meetups	Total Members
Witches	442	6,757
Slashdot	401	11,809
LiveJournal	311	10,691
Bloggers	136	4,222
Pagans	90	2,841
Fark	81	4,621
Ex-Jehovah's Witness	67	1,609
Bookcrossing	56	4,414
Xena	51	1,641
Tori Amos	47	2,261
Ultima	38	2,467
Star Trek	35	1,196
Radiohead	32	1,986
Vampires	28	1,339
Atheists	27	1,338

This list is unlike any list of American groups ever assembled. It measures something important (or rather it collates several different important things) because it demonstrates

that Meetup's convening power lies not in recreating older civic groups but in creating new ones.

The groups represented here can be divided into three broad categories. The first, including Witches, Pagans, Ex-Jehovah's Witnesses, and Atheists, are people who share some religious or philosophical outlook but have no support from the broader U.S. culture. There are many more Presbyterians than pagans in the United States, but the Presbyterians aren't on this list because they don't need Meetup to figure out when and how to assemble; they meet every Sunday morning at the Presbyterian church. Because they are both internally organized and externally supported, Presbyterians suffer less than pagans from transaction costs, who have no culturally normal place and time to meet and no ready way to broadcast their interests without censure. Jehovah's Witnesses enjoy advantages similar to those of other Christian sects, but ex-Witnesses turn to Meetup because they don't enjoy those socially supported advantages of coordination.

The second category of Meetup groups includes the members of websites and services who would like to assemble with other users of those services in real life. This group includes Slashdot, LiveJournal, Bloggers, Fark, Ultima, and Bookcrossing. (Interestingly, the numbers show how clustered these groups are; though Slashdot and LiveJournal had more members than Witches did, they met in fewer cities; or put another way, Witches are more evenly distributed in U.S. society than are geeks or bloggers.) This is what the end of cyberspace looks like: the popularity of these Meetup groups suggests that meeting online isn't enough and that after communicating with one another using these various services, the members become convinced that they share enough to want to get together in the

real world. Especially relevant to this thesis is the Ultima group. Ultima is an online game set in an imaginary world, Britannia, rendered in 3D, where players interact with one another. It is one of a class of games called "massively multiplayer online role-playing games," or MMOs for short. If virtual interactions were ever enough to be completely satisfying, we'd expect them to work best in these virtual worlds. But the popularity of Meetup groups for virtual contacts shows that even online communication that emulates face-to-face interaction still leaves people wanting real human contact.

The third category includes fans of cultural icons whose work is quirky enough that those fans want to be in one another's presence. LiveJournal users can at least potentially come in contact with one another on the website, but Tori Amos fans are simply guessing that they will get along. (The Vampires group falls into both the first and third categories.) To want to be in other people's company without having spoken before, on the basis of a shared cultural affinity, is a pretty good advertisement for Heiferman's initial thesis—that even in a mediated age, people crave real human contact.

These three categories have several things in common. First, they represent not just things people do but ways they think of themselves (and of other people). Many more people use Google than LiveJournal, but there is no broad interest in a Google users' Meetup group. Second, this self-conception translates into a desire to meet with other people who share the same interests. Many more people were watching *Everybody Loves Raymond* in 2002 than were watching *Xena: Warrior Princess*, but *Xena*-fandom was a better predictor of real commonality. Finally, the world provided no easy way for these

people to find one another prior to Meetup. Because the audience for *Xena* was passionate but small, the likelihood that *Xena* fans would find one another at random was similarly small, but precisely because of this minority status, the likelihood that, once they did, they would feel some sense of kinship was higher than average. This effect is general. Lada Adamic, a researcher at HP Labs, studied the users of an online student center at Stanford called Club Nexus, and they found that two students were likely to be friends if their interests overlapped, and that the likelihood rose if the shared interests were more specific. (Two people who like fencing are likelier to be friends than two people who like football.) The net effect is that it's easier to like people who are odd in the same ways you are odd, but it's harder to find them. Meetup, by solving the finding problem, created an outlets for many new groups—groups that had never been able to gather before.

Meetup didn't end up recreating the old model of community, because it provided a different set of capabilities; the groups that took first and best advantage of those capabilities were the groups with a latent desire to meet but had faced previously insuperable hurdles. These groups aren't the classic American interest groups of yore; many of the most popular groups tell us surprising things about what our society is like right now.

Stay at Home Moms and the Politics of Exclusion

One of the most popular current groups on Meetup is Stay at Home Moms (SAHM). Mothers with young children have

been gathering in groups since before the invention of the internet, in fact before the invention of agriculture. This is an old pattern, so why would SAHM Meetups be so popular? The answer, in one sentence, is that modern life has raised transaction costs so high that even ancient habits of congregation have been defeated. As a result, things that used to happen as a side effect of regular life now require some overt coordination.

Some of the hurdles to be overcome are physical. As of the 2000 census, a majority of the U.S. population lived in the suburbs, and in the suburbanized United States, physical distance raises several barriers. Houses are often separated from commerce, so much of the time spent doing errands or ferrying children from hither to yon is spent in a car. In a pedestrian setting, running into someone is a good thing; in a car, not so much. Both the distance between the grocery store and home, and the fact that travel between the two is highly enclosed, reduce the likelihood of chance social encounters (and as a result reduces the raw material for building social capital).

As the two-income family has become more normal, the center of gravity for social interaction has shifted from the neighborhood to the workplace. Not only have the suburbs reduced the likelihood of chance encounters, but the increased percentage of the population with jobs, including especially a sharp increase in the number of women, means that the workplace now has many of the characteristics that the neighborhood used to have. You are likelier to be introduced to new coworkers than to new neighbors, and interactions at work produce the kind of familiarity and trust that used to be more a part of the fabric of our communities.

Meetup makes the coordination of groups simple, offering

a way of undoing at least some of the damage inflicted on that fabric. This is one reason groups like Stay at Home Moms matter so much. Some groups we expect to be technology-obsessed; maleness, singleness, and youth all correlate with technophilia, while femaleness, age, and family life don't. So when a group of mothers adopts a piece of technology, it indicates an expression of preference far more serious than seeing a thirteen-year-old boy go wild over an Xbox. The popularity of groups like Stay at Home Moms indicates that Meetup's utility in helping people gather in the real world is valuable enough to get the attention of people who are too busy for most new tools.

The most successful Meetup parents' group didn't turn out to be the most general one. Meetup also lists a Parents and Kids Playgroup, which describes a much larger class of potential members than Stay at Home Moms does, but the Parents and Kids group is significantly less popular. This is one of the essential conundrums of social capital—inclusion implies exclusion. The very name Stay at Home Moms is a salvo in the decades-long conversation about the ideal structure of a family—this group is for mothers who are playing a relatively traditional role in child-raising. Though it is hard to imagine a man with a child being turned away from the North Charlotte Stay at Home Moms Meetup, say, it's also hard to imagine that a lot of dads show up in the first place.

Self-Help We Don't Approve Of

In 2002 I taught a graduate course at New York University called "Social Weather," about the experience of participating

in online groups. The course's title was an analogy to the way the weather affects our mood; in the class we were looking at how social groups create an emotional environment that affects all the participants. One of my students in that class, Erika Jaeggli, was also working on the magazine YM's website. YM (formerly *Young Miss*, then *Your Magazine*, then just YM) is designed to appeal to teen girls. In 2002, like almost every other magazine in the country, YM was wrestling with how to embrace the Web. In addition to putting the magazine's articles online, the staff created a set of online bulletin boards where YM readers could go online and talk to one another about whatever was on their mind. Popular topics included clothes, school, romance, and health and beauty—pretty standard fare for teen girls. Erika's job was half host, half chaperone, working to draw the girls out and make them feel comfortable talking to one another, while also keeping the conversation from devolving into name-calling or turning to inappropriate subjects. Particularly at an age when readers were exploring previously off-limits subjects like sex or the use of alcohol and other drugs, the role of an editor was a balancing act. Too little intervention, and the conversation would turn into bedlam; too much would seem like a ham-handed attempt to bring the girls into line—precisely the kind of treatment from adults they were coming to the YM website to escape.

A few months into the semester Erika stopped me in the hallway to tell me YM was shutting down its health and beauty bulletin board. When I expressed surprise that a magazine focused on teen girls would kill off those discussions, she said, "Most of the girls were fine, but we couldn't figure out how to stop this one group of girls from swapping tips on remaining

anorexic." These Pro-Ana girls (short for pro-anorexia) were posting pictures of models and actresses whose rib cages were showing as "thinspiration" and exhorting each other with "You've made a decision—you won't stop. The pain is necessary, especially the pain of hunger. It reassures you that you are strong—can withstand anything—and that you are NOT a slave to your body; you don't give into its whining."

Most dangerously, the Pro-Ana girls were trading practical advice (though the word "practical" is odd in this context):

You can train yourself to forget hunger by gently punching your stomach every time you get hungry because you'll hurt too bad to eat.

Take TUMS to help with hunger pains; they have calcium so they'll help in that area also.

Clean something you find truly disgusting. Afterwards, you won't feel like eating for another couple of hours.

The problem for YM wasn't that the bulletin board had failed to get the interest of their readers. The problem was that it had succeeded in a way for which YM was unprepared.

Whenever individuals want to find one another, the larger society in which they are embedded can provide or withdraw support for their association. Much of the way we talk about identity assumes it is a personal attribute, but society maintains control over the use of identity as an associational tool. A recovering addict would find it very risky to ask coworkers for

help finding a support group, as might someone looking for the local gay community. Whether society offers or withholds this support, however, matters less with each passing year.

Here is the dilemma the YM staff found themselves in. To host a conversation among their most active and engaged readers, they had to monitor the site, but if Erika and the other online editors had weeded out every mention of anorexia, they would come to seem like bullies, especially as some of the conversations were genuinely about avoiding anorexia. Further complicating things, the Pro-Ana girls were willing to go to great lengths to have their discussions out in the open. In the end, the possible sweet spot between too little intervention and too much came to seem illusory, and YM simply shut down the conversation, rather than engage in daily censorship or risk having the girls who congregated at YM get sick. But what exactly had the girls done that presented such a novel challenge? Anorexia has been a source of public worry since the 1960s, and groups of girls have been hanging out together for decades, talking about everything from sex and drugs to fashion and food. Did YM just act on the standard fear that new technology would bring ruin to society? Or was something different?

Something is different. It is easier for groups to form without social approval. Predictably, the Pro-Ana movement has simply moved from hosted conversation spaces like that on YM to more open tools like weblogs and social networking sites like MySpace. YM was able to withdraw its support for the group on its own site, but neither it nor any other organization could prevent the girls from forming groups and conversing with one another if they wanted to. Before we had any

real group-forming technologies, merely finding people who were interested in the same things was hard, and most of the ways we had for doing so—from putting up flyers around the neighborhood to taking out an ad in the local paper—were expensive and time-consuming. Because of these difficulties, social approval could make group-forming much easier, and social disapproval could make it much harder. Formal mechanisms like the law are one factor: it is easier to find a group of people to drink with than to shoot up with, because the law treats alcohol and heroin differently. But legal strictures account for only a small number of these cases; there are many more informal mechanisms for creating the same effect.

Remember the Mermaid Parade photographers? Or Voice of the Faithful? Or the Ex-Jehovah's Witnesses? All these groups, different as they are in membership, outlook, and goals, share two key characteristics. First, they all started out as latent groups—they had things in common, but the cost and hassle of finding one another was too high. Second, the society they lived in didn't make it easy for them to find one another. In some cases, as with the Mermaid Parade attendees, it was simply because of the old mismatch between effort and outcome. In other cases, though, it was because the institutions best positioned to do the introducing were actively opposed to the goals of the latent group. You could hardly expect the Jehovah's Witnesses or the Catholic Church to spend time or money helping coordinate people who want to criticize them or force them to change their ways of doing business.

Groups like Ex-Jehovah's Witnesses and the Pro-Ana girls

no longer need social support to gather; they all operate under the Coasean floor, where lowered transaction costs have made gathering together so simple that anyone can do it. Recording, searching, and transmitting information, including especially information about ourselves, is something our communications networks are effortlessly good at. The enormous visibility and searchability of social life means that the ability for the like-minded to locate one another, and to assemble and cooperate with one another, now exists independently of social approval or disapproval. The gathering of the Pro-Ana girls isn't a side effect of our social tools, it's an effect of those tools.

When society is changing, we want to know whether the change is good or bad, but that kind of judgment becomes meaningless with transformations this large. It's good that the kids in Belarus now have flash mobs as a tool for opposing political oppression, but for other groups, whether Voice of the Faithful or the passengers demanding better treatment from the airlines, the change looks different depending on where you sit. Loyal Catholics might regard VOTF's demands as a threat to the church they love, and union members may not want the airlines' financial position weakened by the passengers' demands.

Sorting the good from the bad is challenging in part because we're used to social disapproval making it hard for groups to form. Alcoholics Anonymous has more support from society than the Pro-Ana girls, but both groups use the language of self-help to describe what they do. The Pro-Ana movement demonstrates, along with sister movements like Pro-Mia (bulimia) and the Cutters (self-mutilation), that the

definition of self-help has suffered the same blow that journalism has. For much of the twentieth century Alcoholics Anonymous, the premier self-help organization, set the tone for social assumptions about self-help: it was a place of devotion and healing, and it promoted a generally approved goal. The shock of the Pro-Ana movement is that it seems to turn many of those aspects inside out, helping people remain sick or become sicker.

The shock turns out to be misplaced: the Pro-Ana movement is in fact a self-help movement, because the content of a self-help movement is determined by its members. The logic of self-help is affirmational—a small group bands together to defend its values against internal and external challenges. When the small group is a bunch of drunks trying to get sober, against the norms set by their drinking buddies, then society generally approves. When the small group is a bunch of teenage girls trying to get or remain dangerously thin, against the judgment of their horrified parents and friends, then we disapprove. But the basic mechanism of mutual support remains the same.

Falling transaction costs benefit all groups, not just groups we happen to approve of. The thing that kept phenomena like the Pro-Ana movement from spreading earlier was cost. The transaction costs of gathering a group of like-minded individuals, especially in an anonymous fashion, has historically been large, and self-funded and socially approved groups like AA were the only ones that could take on those costs. Once the transaction costs fell, however, the difficulties of putting such groups together disappeared; the potential members of such a group can now gather and set their own goals without needing any sort of social sponsorship or approval.

Three Kinds of Loss

Our new freedoms are not without their problems; it's not a revolution if nobody loses. Improved freedom of assembly is creating three kinds of social loss. The first and most obvious loss is to people whose jobs relied on solving a formerly hard problem. This is the effect felt by media outlets challenged by mass amateurization. The basic problem of copying and distributing information, previously an essential service of the music and newspaper industries among others, is now largely solved thanks to digital networks, undermining the commercial logic of many industries that relied on previous inefficiencies.

Andrew Keen, in *Cult of the Amateur*, describes a firm that ran a \$50,000 campaign to solicit user-generated ads. Keen notes that some professional advertising agency therefore missed out on hundreds of thousands of dollars in fees. This loss is obviously a hardship for the ad agency employees, but were they really worth the money in the first place if amateurs working in their spare time can create something the client is satisfied with? The spread of cheap and widely available creative tools is sad for people in the advertising business in the same way that movable type was sad for scribes—the loss from this kind of change is real but limited and is accompanied by a generally beneficial social change.

The second kind of loss will damage current social bargains. Many countries place restrictions on the media in the run-up to elections, but this raises the question of who “the media” is today and what controls should be put on them. Different

countries are coming up with different answers—Singapore banned blogging during the last few weeks before a 2005 election but couldn't control Singaporeans blogging overseas; the Thai government forbade blogging on all political matters, to little effect; and the U.S. election commission decided not even to try to apply its media coverage rules to blogging. The provisional and variable nature of these restrictions suggests that the old relations between the media and the state, even where they are broadly supported by the citizenry, are going to be as impossible to sustain as the old definitions for journalism, which is now less a profession than an activity.

The third kind of loss is the most serious. Networked organizations are more resilient as a result of better communications tools and more flexible social structures, but this is as true of terrorist networks or criminal gangs as of Wikipedians or student protesters. This third loss, where the harms are not merely transitional, leads to a hard question: What are we going to do about the negative effects of freedom? It's easy to tell the newspaper people to quit whining because the writing has been on the wall since the internet became publicly accessible in the early 1990s—their response has been inadequate in part because they waited so long to grapple with the change. It's harder, though, to say what we should be doing about Pro-Ana kids or about newly robust criminal networks.

It used to be hard to get people to assemble and easy for existing groups to fall apart. Now assembling latent groups is simple, and the groups, once assembled, can be quite robust in the face of indifference or even direct opposition from the larger society. (In some cases, that very opposition can *strengthen* the group's cohesion, as with the Pro-Ana girls.)

When it is hard to form groups, both potentially good and bad groups are prevented from forming; when it becomes simple to form groups, we get both the good and bad ones. This is going to force society to shift from simply preventing groups from forming to actively deciding which existing ones to try to oppose, a shift that parallels the publish-then-filter pattern generally.