

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 90 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, and the effect gets considerably worse as the group grows even moderately large.




Figure 2-I: 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.

By the time you want to go to a movie in a group often, 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 phenomenon 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 wider

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 shells. 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 news-papers, 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 Rich 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 "merrmaidparade," 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 is 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 steppingstone 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 are to say groups that existed only in potential, 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 news gathering. 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 adds 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 simplifies 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 layer 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 remains 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 the 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 absolute 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 non-institutional 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. Sonic groups are perfdy 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.