Social Technology

FROM HUNTER-GATHERERS TO MULTINATIONALS

IN 2002, WILLIAM EASTERLY of the Institute for International Economics and Ross Levine of the University of Minnesota conducted a detailed study of seventy-two rich and poor countries and asked, "What makes one country richer than another?"[1] One might assume that the major determinants of national wealth include factors such as the existence of natural resources, the competence of government policies, and the relative sophistication of a country's Physical Technologies. Easterly and Levine found that while these factors all mattered to a degree, the most significant factor was the state of a nation's Social Technology. The rule of law, the existence of property rights, a well-organized banking system, economic transparency, a lack of corruption, and other social and institutional factors played a far greater role in determining national economic success than did any other category of factors. Even countries with few resources and incompetent governments did reasonably well if they had strong, well-developed Social Technologies. On the flip side, no countries with poor Social Technologies performed well, no matter how well endowed they were with resources or how disciplined their macroeconomic policies were.

Not only do Social Technologies affect the performance of a nation-state, but they also explain differences in performance at the more granular levels of industries and companies. During the late 1990s, economists began to notice a rapid rise in the productivity of the U.S. economy. At first, researchers looked to PTs for an explanation. There had been massive investments in computing power over the previous two decades, and a leading hypothesis was that the economy was at long last seeing the payoff from that investment.

However, my colleagues at McKinsey & Company's Global Institute were skeptical and delved underneath the headline productivity figures.[2] They found that the real driver of increased productivity was changes in how companies were organizing and managing themselves—in other words, innovations in Social Technologies.

One of the industries the McKinsey team examined in depth was retail, and in particular the impact of Wal-Mart on overall sector productivity. Wal-Mart's innovations in large-store formats and highly efficient logistical systems in the 1980s and early 1990s enabled the company to be 40 percent more productive than its competitors. This challenge in turn forced its competitors to imitate Wal-Mart's organizational innovations and raise their own productivity 28 percent in the late 1990s. Meanwhile, Wal-Mart continued to increase its own productivity a further 22 percent. This particular innovation race in Social Technologies in the retail sector alone accounted for *nearly a quarter* of the growth in overall U.S. productivity during the period. Similar Social Technology innovation races in five other sectors made up virtually all the rest. Computers certainly played a vital role in this story; without them, Wal-Mart's sophisticated logistics processes would not be possible. But computer technology played an enabling rather than a primary role; it was the innovations in organization and processes that yielded the dramatic productivity gains.

Let's Get Organized

In the previous chapter, I defined Physical Technologies as methods and designs for transforming collections of matter, energy, and information from one state into another in pursuit of a goal or goals. Social Technologies can be defined similarly:[3]

> Social Technologies *(STs) are methods and designs for organizing people in pursuit of a goal or goals.*

A group of people might come together and organize themselves to start a company, to form a religion, or to create a Friday night bowling league. Such acts of organizing are always in pursuit of a goal, whether it is profits, spiritual enlightenment, or a bit of fun. Just as PTs are methods for creating order in the physical realm to meet human needs, STs are methods for creating order in the social realm also to meet human needs.

The term *Social Technologies* is a close cousin of a term used by economists: *institutions.* Nobel Prize winner Douglass North defines institutions as "the rules of the game in a society."[4] Institutions are one ingredient in organizing, but I intend for my definition of STs to be somewhat broader and include other ingredients such as structures, roles, processes, and cultural norms. Social Technologies include all the elements necessary for organizing. The STs of a soccer team include not just the rules of the game, but also the job description of the goalkeeper, the cultural norms of the team, and whether the team fields three strikers at the front or two strikers and a sweeper at the rear. While the STs of a soccer team would include a complete description of the team's organizational methods, it would *not* include the strategy used by the team. Thus, statements such as "attack on the left" or "focus on short passes" would not be included. In an economic context, such strategies are a part of a Business Plan. This is a distinction we will return to later in the book.

How Social Technologies Evolve

Given this definition of STs, it should come as no surprise that we can construct a theoretical design space for them, a Library of All Possible Social Technologies. We will follow the same path we traveled for imagining the Physical Technology design space. In the ST library are schemata that code for specific designs and instructions for creating social structures. We can imagine writing out instructions for organizing a Yanomamö hunting party, describing the organizational structure of GE, or laying out European banking regulations. These instruction sets might include natural-language text, charts, and tables that include descriptions of the organizational structure, roles, decision processes, formal rules, incentive systems, codes of behavior, and so on. As before, we can imagine these STs encoded in multivolume sets of 500-page books—a vastly-larger-than-the-universe Library of All Possible Social Technologies.

As with PTs and Business Plans, in the real world, some STs exist in writing, but many exist only in people's heads. Social Technologies don't actually have to be written down, but in principle could be written down to a sufficient degree that a qualified reader could act on them to realize the design. So a Yanomamö hunter could understand the schema for a hunting band, a GE executive could understand the GE organization description, and an appropriately experienced EU bureaucrat could understand the banking regulatory structure.

Like the other design spaces we have discussed, the library of STs has three important attributes. First, like its PT cousin, the ST design space is self-feeding and exponentially unfolding.[5] Each ST breakthrough creates more headroom for the next set of breakthroughs—the invention of money enabled the invention of accounting, which enabled the invention of the joint stock company, which enabled the invention of stock markets, and so on.

Second, STs have a modular, building-block quality to them. For example, the organizational design of a large multinational corporation is a collection of modules that includes designs for organizing its business units, designs for its accounting and control systems, designs for its committee structure, and designs for its cultural norms of behavior.[6]

Third, the fitness landscape associated with ST design space is highly likely to be rough-correlated. Small differences in ST designs tend to yield small differences in relative fitness, but occasionally, small changes will make a ST either unworkable or much better. Thus, the ST fitness landscape, like our other fitness landscapes, has an alpine, rough-correlated shape, with flat spots, Swiss cheese holes, and the occasional portal to higher ground. A prediction from this assumption is that just as we observed S-curves and disruptive technologies in PT space, we would expect to see the equivalent in ST space. History seems to bear this out.[7] For example, the jump from hunter-gatherer ST to settled agriculture can be regarded as a major shift in the S-curve of human economic organization.[8] Likewise, Henry Ford's 1914 development of a radical new way of organizing manufacturing—the production line—was a highly disruptive ST that changed the structure of the early automotive industry, as well as many other industries.[9]

Deductive-Tinkering in Social Technology Space

If the Social Technology fitness landscape is indeed rough-correlated, then a further implication is that a highly effective way to search it is our grand-champion search algorithm—evolution. Just as people use deductive-tinkering to search Physical Technology space, they use deductive-tinkering to search for fit STs in ST space. For example, when Henry Ford and his team developed the production line, they didn't just sit down and deductively theorize about it on paper.[10] Nor did they merely try random experiments. Instead, they used a bit of both. Ford was motivated by the desire to manufacture a car that "the masses" could afford. To do this, he needed to reduce the number of skilled craftsmen in his manufacturing process, thus enabling more of the work to be done by less skilled and less expensive workers. Ford was familiar with advances by the U.S. Ordnance Department's Springfield (Massachusetts) Armory in using standardized interchangeable parts in its manufacturing process, and while he and his team probably did not read much theoretical economics, they were generally familiar with views on the benefits of labor specialization. Armed with a set of deductive hypotheses, Ford began experimenting with different configurations of his plant between 1908 and 1912. After four years of tinkering, in 1913 he struck on the key insight that the car itself should move along the production line rather than the workers, and by 1914, he had implemented a fully working, moving assembly line.

In STs, the ratio of deduction to tinkering is more weighted toward the latter than is the case with PTs. Despite advances in economics and organization theory, there is still far more art than science in activities such as redesigning a company's organization or creating a central banking system, as compared with building a jet aircraft or designing a new heart drug. The pattern of exploration on the fitness landscape of ST space has thus historically featured relatively few directed fingers of deduction and more clouds of trial-and-error exploration. One of the promises of Complexity Economics is that, over time, it will push the art-science boundary in ST a step further toward science. Despite the lesser role of deduction in ST space, the process of searching for fit STs using deductive-tinkering is nonetheless an evolutionary process. People conduct experiments with various STs, and then over time, successful designs tend to persist, while less successful ones fade away. Successful designs tend to be amplified as they are copied, attract more resources, and spread. For example, Ford's innovation of the moving assembly line spread rapidly through the manufacturing sector, displacing other STs, and remains standard practice today.[11]

There are tight linkages between PTs and STs. As humans move across the fitness landscape of PTs, they cause rumblings, earthquakes, and other upheavals in the landscape of STs, and vice versa. An advance in PT such as the ox-drawn plow could only have happened after the ST innovation of village-based agriculture (try carrying a plow as a nomad). Likewise, as mentioned earlier, many management innovations in the modern era have depended heavily on advances in computing and communications technology. In fact, the agricultural, industrial, and information revolutions can each be viewed as coevolutionary merry-go-rounds of advances in PTs leading to new forms of STs, which in turn were crucial for further advances in PTs, and so on.

Competing to Cooperate

We can next ask, what drives humanity's deductive-tinkering search through Social Technology space? What spurs us to constantly seek out new and better ways of organizing ourselves? The answer lies in the magic of non-zero-sum games.

In chapter 10, I noted the distinction between zero-sum games, in which one person's gain is another person's loss, and non-zero-sum games, in which both people can be made better off by cooperating. Cooperation in non-zero-sum games has a 1 + 1 = 3 logic, whereby if you scratch my back, I'll scratch yours, and together we can do something neither can do as well on our own and we both benefit. Non-zero-sum cooperation is one of those Good Tricks of survival that has been widely employed by biological evolution. Dogs hunt in packs, termites collectively build mounds, fish swim in schools, and, like most primates, members of *Homo sapiens* live in groups.

But while the benefits of cooperation in non-zero-sum games are substantial, as the Prisoner's Dilemma showed us, there is often a tension between cooperating for the greater good and pursuing one's narrow self-interest.[12] In his thought-provoking book, *Non Zero*, the journalist and science writer Robert Wright argues that much of human history can be viewed as the outcome of this central tension between cooperation and self-interest.[13] Wright claims that the process of bootstrapping social complexity, from simple hunter-gatherer tribes to organized villages to nation-states and global corporations, has been the result of humans innovating new ways to cooperate across larger and larger scales and devising ways to play increasingly complex and profitable non-zero-sum games. He notes that in a world where resources are finite at any given moment, there are competitive pressures to cooperate. Over time, societies that are better able to organize themselves will socially, economically, and militarily dominate societies that are less successful at creating cooperative structures. Thus, it is the competition to cooperate that drives social innovation.

Recasting Wright's thesis in the language we have developed, we can view the deductive-tinkering search through the ST fitness landscape as a quest for STs that enable people to play and capture the benefits of non-zero-sum games. Social Technology fitness will therefore depend on three factors. First, the ST must provide the potential for non-zero-sum payoffs. Second, it must provide methods for allocating the payoffs in such a way that people have an incentive to play the game. And third, the ST must have mechanisms for managing the problem of defection. Let's take a closer look at each factor.

Non-Zero Magic

There are four basic sources of 1 + 1 = 3 magic in non-zero-sum games. All four have been well known to Traditional Economic theory for a long time. First is the division of labor. As discussed in chapter 2, this benefit was pointed out by Adam Smith over two centuries ago. If two people have even slightly different skill sets, mutual gain can be created by each person's focusing on what he or she does best and then trading. If Larry is a good hunter and Harry a good ax maker, then Larry is better off stalking game than futilely pounding rocks, and vice versa.

Second is the heterogeneity of people. Their different needs and tastes create opportunities to trade for mutual benefit (something we saw in Sugarscape). Charles Darwin observed the benefits of this type of trade while interacting with Fuegian Indians during his voyage on HMS *Beagle*: "both parties laughing, wondering, gaping at each other; we pitying them, for giving us good fish and crabs for rags, [etc.]; they grasping at the chance of finding people so foolish to exchange such splendid ornaments for a good supper."[14]

Third are the benefits of increasing returns to scale, a concept discussed in chapter 3. A lone hunter, for example, might invest 500 kilocalories (kcal) of energy in a few hours of hunting, and have a 20 percent chance of killing an animal worth 2,500 kcal of food. Thus, his expected return would equal 500 kcal, and he would just break even. Now imagine he joins two others to form a hunting party. The three still invest 500 kcal each, but their odds of getting a kill now jump to 90 percent. Thus, the expected value becomes 750 kcal each (90 percent of 2,500 kcal, divided by three). Simply by joining the group, our hominid hunter has made a 250-kcal profit on his investment and greatly reduced his risk.

Fourth, and finally, cooperation helps smooth out uncertainties over time. If one hunting band has a successful day and another does not, the successful band can share its bounty with the unsuccessful group under the proviso that the others do the same when the situations are reversed (and perhaps the payback will include some interest). Cooperation is thus a Good Trick for mitigating risks. If you are on your own and have a run of bad luck, you starve. But if you are in a cooperative group, your colleagues can tide you over until you can pay them back.

Dividing the Spoils

These four sources of non-zero-sum gains can be mixed and matched in various contexts to create a near-infinite number of ways that people can cooperate for their mutual benefit. But for people to have an incentive to cooperate, they must receive some share of the spoils. How the gains of cooperation are divided up is therefore a crucial question. If the rewards are distributed in the wrong way, then cooperation collapses and the non-zero-sum gains evaporate.

Allocating the payoffs from cooperation is where John Nash (profiled in the popular book and film *A Beautiful Mind*) first made his mark with a brilliant paper in 1950 titled "The Bargaining Problem."[15] In the paper, Nash asked the simple question, how will two bargainers come to agreement? How much meat will Harry give Larry for the hand ax? As simple as it sounds, the problem stumped economists for generations. Nash's elegant solution was to say that how two or more bargainers split up the gains from exchange depends on how much each values the benefits of the deal, and what the parties' alternatives are. Each looks for his or her best deal assuming everyone else is looking for the best deal, too, and the trade is made at the point at which *no one has any incentive to change position, given the actions of the other.*

This point became known as the *Nash equilibrium*. Thus, as Harry and Larry haggle over the ax, they eventually find a point at which both are happy to trade and neither can improve his position without blowing the deal—they make the trade. Both walk away better off than they would have been had they not traded at all, thus capturing the non-zero-sum gains of cooperation.

But the existence of a Nash equilibrium does not guarantee a happy, cooperative result. In the single-round Prisoner's Dilemma, the Nash equilibrium is the solution in which both prisoners rat on each other and go to jail. This is because, if you are one of the prisoners and you don't know whether your colleague will talk or stay silent, you are better off talking. As your colleague in crime faces the same incentives, you can assume that he will talk, too. Nash's theorem tells us that for non-zero-sum games to result in cooperation, either the payoffs need to be structured such that everyone's best response is to cooperate, or the players need some mechanism for coordinating their responses. For example, let's say the prisoners have a mafia boss who promises to kill anyone who testifies and to reward anyone who keeps quiet after being released from jail. That would change the payoff structure and move the Nash equilibrium to a point at which both clam up and go free. Likewise, if the prisoners were allowed to communicate and knew what the other had been offered, they could coordinate their responses and avoid the defection trap. This last solution, however, would still have the possibility that one prisoner could nevertheless sell out the other for gain; after all, there is no honor among thieves. This leads to our third critical factor in ST fitness; for STs to be fit, they must have mechanisms for dealing with those who don't play nice.

Cheaters (Mostly) Never Win and Winners (Mostly) Never Cheat

The incentive to cheat means that cooperation is inherently difficult to achieve and potentially unstable even once attained. One hunter might run just a bit slower than his friends and expend only 400 kcal, and yet, so long as his colleagues don't notice his free riding, he will still get his 750-kcal meal. Likewise, the meat I give you for that nice hand ax might be rather old and tough, or I might give you only five ounces instead of the agreed-upon six. The selfish logic of biological evolution says that if cheaters cheat and get away with it, they improve their chances of passing their cheating genes onto their offspring. Thus, cheating confers an evolutionary advantage.

But if biological selection gives cheating genes an advantage over sucker genes, how, then, does cooperation get a foothold in a population? The answer is that the gains from cooperation are so powerful that cooperating genes have an advantage over cheating genes—but only if the genes aren't naive, don't let themselves be suckers, and ensure that cheaters get punished.[16] To survive, cooperating genes need some sophisticated defense mechanisms. Recall our earlier discussion of Kristian Lindgren's model. In the Prisoner's Dilemma, if the players play for just one round, both players have an incentive to rat on each other. When the game is repeated, however, and no one knows when it will end, the dynamics become much more complex. The types of robust, successful strategies that evolve usually have a logic along the lines of this: "I will begin on the assumption of mutual cooperation. If you cheat on me, however, not only will I refuse to cooperate, but I will punish you, even to my own near-term detriment. After some time, though, I might forgive you and try cooperating again, just in case your cheating was an error or a miscommunication, or on the chance you have reformed your ways. If, however, you cheat again, the probability of my forgiving you again will become lower and my punishment even more terrible."[17]

Just as evolution produced this kind of logic in the world of Lindgren's computer model, so too did evolution produce this kind of logic in the minds and instincts of our ancestors. Earlier in the book, we discussed a series of experiments called the ultimatum game in which researchers gave two subjects a pool of money and asked one person to decide on how to split it, and if the other agreed the split was fair, each could keep his or her share. If the partner rejected the split, however, neither got any money. The results of these reciprocity experiments were striking. Economic logic says that people should accept any offer of a split, no matter how small, because some money is always better than no money. In test after test, however, subjects rejected offers that were perceived as unfair, even to their own detriment. The results were consistent across cultures around the world, including hunter-gatherer cultures. Other games and experiments confirm the consistent and deep-rooted nature of human cooperative-reciprocity behavior.[18] Evolution has steered us in a direction whereby we are naturally inclined to be cooperative to capture the riches of non-zero-sum gains. Nevertheless, it has also equipped us with a sensitivity to cheating, expectations of fairness, and a willingness to mete out punishment to those we believe have crossed the line. In effect, evolution has programmed into our mental software sophisticated, intuitive "Nash equilibrium finders" and "fairness detectors" that enable groups of humans to form coalitions that are at least reasonably stable and resistant to attack by free riders and cheaters.[19]

Our reciprocity software, however, is not hardwired—it can adapt to local circumstances. When we are in an environment in which most of our experience is of other people's cooperation and reciprocation and in which social norms give us signals that people can be trusted (e.g., people tell admiring stories about self-sacrificing, trustworthy types), then our mental cooperation software will tend to be biased toward cooperating. It also will be more surprised and more forgiving when it encounters an example of defection or cheating. In essence, our minds statistically sample the population around us, and if people are usually cooperative, then when we encounter a cheater, we will tend to assume that the person's behavior is probably the result of an error or misunderstanding. In contrast, in a low-cooperation, high-cheating environment with social norms that don't support cooperation (e.g., the stories are all about thieves, and people tell you to "watch your back"), our cognitive cooperation software biases us toward being suspicious. We react harshly to the first signs of cheating, forgive only slowly if at all, and are likely to resist cooperating until given a sign of cooperation from the other party first.

The local tuning of reciprocity norms can create very complex dynamics at the level of populations. High-cooperation societies can see collapses in cooperation if cheating reaches a critical mass; low-cooperation societies can get stuck in uncooperative, economically impoverished dead ends; and when people from different cooperative traditions mix, it can lead to misunderstanding and turmoil.[20] The same issues can arise in organizations too. As we will discuss in chapter 16, high-cooperation, high-trust cultures in companies tend to lead to higher economic performance, and mergers between companies can run into significant problems when populations that have evolved different trust parameters are suddenly thrown together.

Thus, Social Technologies that are better at tapping into sources of non-zero-sum gains, finding cooperative Nash equilibriums for allocating those gains, and managing the defection problem will be higher on the fitness landscape than those that do not. As people have deductively tinkered their way across the landscape in search of fit STs, humankind has evolved increasingly complex and sophisticated social structures for addressing these three issues.

From Family Units to Business Units

The journey of Social Technology evolution began with our genetic predisposition to cooperate most closely with near kin.[21] Family members, after all, share some of our genes, and helping them helps increase the odds that those genes are passed on to the next generation. Our ancestors' earliest cooperative social structure was thus the family, with hominid family habits somewhere between profligate chimpanzees and monogamous apes. In most hunter-gatherer tribes, polygynous men took multiple wives if they could, and the higher the status of the male, the more mates.[22] Early humans tended to be polygynous, but unlike chimps, the males did generally stick around and invest in their mates and offspring, creating relatively stable family units.

Some societies never advanced beyond this most basic of social structures. Robert Wright notes, for example, that the Nunamiut Eskimos and Shoshone Indians were until quite recently organized around family units, with little social structure beyond that.[23]

The first step up the social structure ladder was cooperative hunting bands. The basic caloric logic of hunting bands—you get more to eat—is so compelling that most present-day hunter-gatherer societies and, by implication, most early societies discovered this form of cooperation. However, cooperative hunting bands in most early societies were relatively small scale and were made up of mostly kin or near kin.

The big bang in social cooperation came with the advent of settled agriculture. Physical Technologies for domesticating crops were independently discovered in various parts of the world beginning around eleven thousand years ago.[24] The increased calories and reduced risks of settled agriculture enabled settlements to become more permanent and the size of human groupings to rise significantly, which meant that cooperation began to extend beyond clans of family members. This in turn opened up a slew of new non-zero-sum games.[25] Cooperative groups could tap economies of scale such as the ability to build shelters and other structures that could not be built with only a few family members, could glean benefits from the division of labor in creating artifacts, and could pursue trade between distant villages.

However, all these innovations in cooperation created a new issue: Nash's problem of how to divide the resultant wealth. The sexual hierarchy of our primate background provided a natural answer to this, what Robert Wright and others call the "Big Man Society."[26] In our closest primate relatives, as in many other species, males compete with each other for sexual access to females. Big, strong, clever, and aggressive males push away weaker males for access to females, causing their big, strong, clever, and aggressive genes to be passed on to subsequent generations. In polygynous societies such as those of early humans, the more dominant the male, the more mates, and the more mates, the more offspring. This created a hierarchy of high-status and low-status males. As groups of early humans coalesced into societies of nonrelatives, this sexual hierarchy quite naturally transferred into a socioeconomic status hierarchy for dividing up the spoils of cooperation. The two hierarchies are, of course, two sides of the same coin: the characteristics of sexual dominance (e.g., size, intelligence, and aggressiveness) in our ancestral environment also tended to translate into economic success, and being economically well-off tended to translate into sexual status, as wealthy men could provide more resources for mates and offspring.[27] Thus, from ancient Greek myths to modern tabloid newspapers, humankind's three favorite topics of conversation—sex, money, and status—have been linked from our earliest days.

As discussed, hierarchy is not very much in vogue. No one likes dictators. Corporations are urged to flatten themselves into teams, and business leaders are supposed to resist the temptation to strut around their organizations as the alpha male or alpha female. Yet, the theory of networks shows that hierarchy plays a critical role in any information-processing system, whether it is a computer chip, the Internet, the human brain, or the economy. To tap into the benefits of a division of labor and economies of scale, someone needs to divide up tasks, coordinate their execution, bring things back together, and allocate the spoils. In early societies, that someone was usually a sexually dominant male. Wright describes the role of such a "Big Man" in an early American northwest coastal tribe:

> The chief planner was the political leader, the "Big Man." He held the allegiance of a clan, maybe a village. He orchestrated the building of salmon traps or fish cellars, and he made sure that some villagers specialized in, say making canoes that other villagers could then use. To pay for all of this he would take one-fifth, or even half, of a hunter's kill. Some of this revenue would be returned to the people in the form of chief-sponsored feasts . . .
>
> Needless to say, the Big Man skimmed a little off the top. He lived in a nicer-than-average house and owned a nicer-than-average wardrobe.[28]

One can see that it is not a terribly far distance to modern CEOs and politicians.

Once humans had the invention of hierarchy, it was then a simple step to the nested structure of hierarchy within hierarchy. We can just imagine the progression: at some point, a successful Big Man with a growing village to run does not have enough time to keep his eye on salmon trap production, so he appoints his younger brother or best friend to run that aspect of village life—and voilà, the business unit is born. The Big Man boss has reporting to him minibosses, who in turn have minibosses reporting to them. Hierarchy facilitates the division of labor and the processing of information. It is pervasive in all human social structures, ranging from hunter-gatherer tribes to neighborhood bowling leagues to big corporations.[29]

Peace, Love, and Understanding

While social and economic hierarchies have important information-processing benefits, such structures are inherently unstable. There is constant competition for the top spots, and inevitable succession battles when the Big Man loses his effectiveness or dies. Organizational turmoil is very costly, while stability has many benefits. Stable organizations have the ability to accumulate knowledge and skills over time, play non-zero games with longer-term (and potentially more lucrative) payoffs, and provide more certainty for the participants, thus attracting their cooperation at lower cost. In primate troops, the method of hierarchy management is fairly straightforward. A large, clever, aggressive male dominates the group, until another larger, cleverer, more aggressive male topples him in a violent coup. Unfortunately, this method of hierarchy management is still used today by dictatorships, organized-crime syndicates, and other unsavory social organizations.

With the advent of agriculture and larger, more permanent settlements, a series of innovations were developed for managing power transfers in hierarchies without (or at least with less) costly violence. These innovations included primogeniture (e.g., a prince becoming a king) and the selection of leaders by "elders" (e.g., an Afghani *loya jirga* or a papal election by cardinals). Relatively more recent innovations include democratic elections with universal suffrage, and shareholder governance of corporations. Interestingly, despite the layers of civility in the modern versions of these processes, the threat of physical force always lurks below the surface. If a U.S. president were ever to refuse to leave the White House after losing a valid election, some rather large men in dark sunglasses would presumably cart him or her away. Nonetheless, Social Technologies for managing changes in hierarchical structure have been crucial in enabling organizations to maintain stability and endure over time.

Not only does competition within hierarchies need to be managed, but competition *between* members of different hierarchies also presents both threats and opportunities. Human groups need STs that enable cooperation between complete strangers. The first problem that strangers must overcome when they meet is that they don't know whether to trust each other. The parameters for their norms of cooperation and reciprocity behaviors might be very different, and one party might take advantage of the other in a transaction. Thus, people need STs for figuring out whom they can and cannot trust beyond their immediate kin. The first such ST extending beyond individual villages was undoubtedly tribal identity. By identifying "your people" versus "outsiders," you could efficiently find people with whom you were more likely to share social norms. And since your interactions in a tribe were likely to be repeated over time, you were less likely to get ripped off or to have destructive miscommunications.

Historically, trading networks have tended to develop first and most strongly within tribal, ethnic, and religious groups. Consider the Mayan trading network that covered a large swath of Mexico, Guatemala, and Belize from 250 to 900 AD; the extensive Muslim trading networks that stretched from North Africa through the Middle East and into Central Asia over the past thousand years; and the Ivy League–dominated Wall Street firms. The ugly side of this tagging of "your people" versus "outsiders" is discrimination, which is inherently self-limiting. In the absence of other information, tagging may be a strategy that reduces risk, but it also excludes a larger world of potentially beneficial relationships.

In a sense, exclusionary tagging is like having a closed computing environment in which only computers of the same brand can talk to each other (e.g., IBMs to IBMs, Dells to Dells)—there are benefits to standard protocols for communications and behavior, but at the cost of scalability. Thus, a major breakthrough in ST was the development of an open protocol to enable strangers to cooperate: the rule of law. Laws enable complete strangers, with different backgrounds, histories, ethnicities, and social norms, to conduct business with each other with greatly reduced risk. For example, a major investment such as buying a home can be intimidating even with all the protection of property law, building codes, and insurance. Imagine what it would be like if such transactions were completely unregulated. There is a powerful correlation between the wealth of a society and the existence of written laws with mechanisms for enforcement and adjudication.[30] Establishing the rule of law is considered a critical hurdle by development economists when they are trying to stimulate growth in poor countries, and countries without strong legal systems inevitably fall back on tagging as a less efficient and socially divisive substitute. Laws, of course, don't completely replace trust, and societies can become dysfunctional if social trust breaks down and people over rely on legal institutions (witness the increasing litigiousness of U.S. society). Nevertheless, complex, large-scale cooperation is impossible without a well-functioning legal and regulatory system to provide protocols for cooperation.

Communications is also critical in engendering cooperative behavior. Thus, the development of language increased dramatically the potential for social and economic cooperation, unlocking a host of new non-zero-sum gains. There is much debate on when language developed—the admittedly wide range of estimates is from 30,000 years ago to around 1 million years ago.[31] Regardless, language probably developed well after toolmaking, for which there is evidence from 2.5 million years ago. Consequently, there was a phase of human economic activity that was prelinguistic and thus inherently limited in its level of complexity. Just as language transformed humankind's ability to create Physical Technologies, so too did it transform the creation of STs. One can imagine the advantage that language genes must have carried with them; as vocabulary rises, the space of beneficial, cooperative social games opens up exponentially. Imagine trying to negotiate a complex deal with someone if you could use only grunts, gestures, and facial expressions. Now imagine trying it with fifty words of bad, tourist-quality French. Now imagine it with the fluent language skills of a native speaker.

Once language was developed, a series of PT innovations further enhanced its value in cooperative activities. Writing, which appeared around 5,000 years ago (and thus by the standards of human history is a relatively recent invention), enabled people to disseminate knowledge more widely and to preserve it more accurately over time. It is doubtful whether the highly complex societies of ancient Egypt, Greece, or Rome would have been possible without writing. The emergence of European society from its tribal Dark Ages was certainly facilitated by the printing press, and the social innovations of the industrial revolution would not have been possible without reliable mail service. Nor could the complexities of modern global corporations be managed without telephones, faxes, and e-mail.

Building Computers Out of People

We have noted that a number of key Social Technologies innovations have an information-processing angle to them and earlier that "networks of information-processing things" have the ability to compute. Once the evolution of STs reached the stage at which large numbers of people could form cooperative networks and had the means for communicating and storing significant amounts of data, human organizations took on a different character—they became capable of emergent computation.[32]

Organizations of people have the ability to process information and solve complex problems that individuals cannot process or solve on their own. British Petroleum (BP), for example, can be thought of as a computer built for solving the problem of how to extract oil and gas from locations around the world, refine it, and then distribute it to millions of energy users. At BP, there is no one who can tell you, in full detail, just how that immensely complex problem is solved. Think of the vast flows of data flowing into BP on a daily basis and of all the decisions that need to be made, decisions ranging from board-level judgments to the shift schedule for a rig in the North Sea. As much as we like to think of a CEO as being in command, it is impossible for a CEO even as capable as John Browne to be aware of more than a tiny fraction of the thousands and perhaps even millions of decisions being made at any given moment in a large organization. Yet, the hugely complex problem of finding, extracting, refining, and distributing oil is solved in a highly distributed fashion, day in and day out.

Just like an anthill, or the brain, human organizations exhibit a form of networked emergent intelligence. The University of California, San Diego, anthropologist and cognitive scientist Edwin Hutchins has studied the problem-solving capabilities of individuals versus organized groups in a variety of settings. He concluded that organizations are capable of having collective, emergent capabilities that do not exist individually within the group.[33] In essence, not only is BP smarter than any one of its people, it is also smarter than the sum of its people.

One can debate the role of large global corporations such as BP in society. But even the most hardened corporate critic would have to admit that an organization such as BP, with its 103,000 employees in over a hundred countries around the world, is a marvel of human cooperation. The vast majority of its people have never met and never will meet, but are bound together in a web of social structures, norms, protocols, legal structures, and incentives that enable them to work together for a common purpose. If one extends that web of cooperation beyond BP's immediate employees to include its 1.3 million shareholders and thousands of supplier and other partner companies, then the scale of a social structure such as BP becomes even more remarkable.

For an organization of BP's size and complexity to exist, it must sit atop a vast mountain of ST innovations that society has evolved over millennia. These innovations include such STs as money, which was first used around 2,600 years ago in Mesopotamia and provides in essence a universal utility converter—it enables one person's economic needs and wants to be translated into the same units as someone else's needs and wants.[34] Likewise, BP would have a difficult time functioning without a central nervous system of financial information provided by the ST of double-entry accounting, originally developed by Italian merchants in the thirteenth century.[35] Nor could BP even exist in its current form without the invention of the limited-liability joint stock corporation, invented by the British Parliament in a series of Acts between 1825 and 1862.[36] While BP depends on a legacy of STs to function, it is also a participant in the economic evolutionary system as its managers deductively-tinker their way to new methods of organizing and managing, and as successful STs are adopted and spread both within BP and outside it.

From the biological heritage of our primate origins, we inherited an inclination to cooperate for mutual gain and a compulsion to compete in dominance hierarchies, and eventually our developing human brains gave us language. From these humble beginnings sprang an evolutionary process of deductive-tinkering over tens of thousands of years as humans experimented with various ways of organizing their social and economic activities. The inherent non-zero-sum riches found in cooperation rewarded those STs that worked, and over time, humankind found increasingly effective Good Tricks for organizational success. As our species evolved its way through the rough-correlated landscape of ST design space, people were able to build on the tricks that worked, with each successful innovation yielding the possibility of even more future possibilities. As such innovations improved the ability of organizations to process information and solve problems, richer and richer areas of the ST fitness landscape were opened up. Simultaneously, ST space coevolved with Physical Technology space as discoveries in each sphere fed new possibilities in the other.

We have gone from hand axes to spacecraft, and from hunting bands to multinational corporations. We still have not gone the full distance in our journey from the Yanomamö to New Yorkers, but we are closing in, and in the next chapter, we will put the final pieces of economic evolution together.