

will require a sustained and imaginative effort on the part of researchers across the sciences.

Kenneth Boulding summarized science as consisting of “testable and partially tested fantasies about the real world.” The science of complex systems is not a new way of doing science but rather one in which new fantasies can be indulged.

CHAPTER 2

Complexity in Social Worlds

I adore simple pleasures. They are the last refuge of the complex.

—Oscar Wilde, *The Picture of Dorian Gray*

When a distinguished but elderly scientist states that something is possible, he is almost certainly right. When he states that something is impossible, he is very probably wrong.

—Arthur C. Clarke, *Report on Planet Three*

WE ARE SURROUNDED by complicated social worlds. These worlds are composed of multitudes of incommensurate elements, which often make them hard to navigate and, ultimately, difficult to understand. We would, however, like to make a distinction between complicated worlds and complex ones. In a complicated world, the various elements that make up the system maintain a degree of independence from one another. Thus, removing one such element (which reduces the level of complication) does not fundamentally alter the system's behavior apart from that which directly resulted from the piece that was removed. Complexity arises when the dependencies among the elements become important. In such a system, removing one such element destroys system behavior to an extent that goes well beyond what is embodied by the particular element that is removed.

Complexity is a deep property of a system, whereas complication is not. A complex system dies when an element is removed, but complicated ones continue to live on, albeit slightly compromised. Removing a seat from a car makes it less complicated; removing the timing belt makes it less complex (and useless). Complicated worlds are reducible, whereas complex ones are not.

While complex systems can be fragile, they can also exhibit an unusual degree of robustness to less radical changes in their component parts. The behavior of many complex systems emerges from the activities of lower-level components. Typically, this emergence is the result of a very powerful organizing force that can overcome a variety of changes to the lower-level components. In a garden, if we eliminate an insect the vacated niche will often be filled by another species and the ecosystem will

continue to function; in a market, we can introduce new kinds of traders and remove old traders, yet the system typically maintains its ability to set sensible prices. Of course, if we are too extreme in such changes, say, by eliminating a keystone species in the garden or all but one seller in the market, then the system's behavior as we know it collapses.

When a scientist faces a complicated world, traditional tools that rely on reducing the system to its atomic elements allow us to gain insight. Unfortunately, using these same tools to understand complex worlds fails, because it becomes impossible to reduce the system without killing it. The ability to collect and pin to a board all of the insects that live in the garden does little to lend insight into the ecosystem contained therein.

The innate features of many social systems tend to produce complexity. Social agents, whether they are bees or people or robots, find themselves enmeshed in a web of connections with one another and, through a variety of adaptive processes, they must successfully navigate through their world. Social agents interact with one another via connections. These connections can be relatively simple and stable, such as those that bind together a family, or complicated and ever changing, such as those that link traders in a marketplace. Social agents are also capable of change via thoughtful, but not necessarily brilliant, deliberations about the worlds they inhabit. Social agents must continually make choices, either by direct cognition or a reliance on stored (but not immutable) heuristics, about their actions. These themes of connections and change are ever present in all social worlds.

The remarkable thing about social worlds is how quickly such connections and change can lead to complexity. Social agents must predict and react to the actions and predictions of other agents. The various connections inherent in social systems exacerbate these actions as agents become closely coupled to one another. The result of such a system is that agent interactions become highly nonlinear, the system becomes difficult to decompose, and complexity ensues.

2.1 THE STANDING OVATION PROBLEM

To begin our exploration of complex adaptive social systems we consider a very simple social phenomenon: standing ovations (Schelling, 1978; Miller and Page, 2004). Standing ovations, in which waves of audience members stand to acknowledge a particularly moving performance, appear to arise spontaneously.¹ Although in the grand scheme of things

standing ovations may not seem all that important, they do have some important parallels in the real world that we will discuss later. Moreover, they provide a convenient starting point from which to explore some key issues in modeling complex social systems.

Suppose we want to construct a model of a standing ovation. There is no set method or means by which to do so. To model such a phenomenon we could employ a variety of mathematical, computational, or even literary devices. The actual choice of modeling approach depends on our whims, needs, and even social pressure emanating from professional fields.

Regardless of the approach, the quest of any model is to ease thinking while still retaining some ability to illuminate reality.

A typical mathematical model of a standing ovation might take the following track. Assume an audience of N people, each of whom receives a signal that depends on the actual quality of the performance, q . Let $s_i(q)$ give the signal received by person i . We might further specify the signal process by, say, assuming a functional form such as $s_i(q) = q + \epsilon_i$, where ϵ_i is a normally distributed random variable with a mean of zero and standard deviation of σ . To close the model, we might hypothesize that in response to the signal, each person stands if and only if $s_i(q) > T$, where T is some critical threshold above which people are so moved by the performance that they stand up and applaud.

Given this simple mathematical model, how much of reality can we illuminate? The model could be used to make predictions about how many people would stand. We could tie this prediction to key features of the model; thus, we can link the elements like the quality (q) of the performance, the standing threshold (T), and even the standard deviation of the signal (σ) to the likelihood of an initial ovation of a given size. Given the current form of the model, that is about the extent of what we can predict. These predictions do provide some illumination on reality, but they fail to illuminate some of the key elements that make this problem so interesting in the first place (like the waves of subsequent standing).

Given this, we might want to amend the model to shed a bit more light on the subject at hand. It is probably the case that people respond to the behavior of others in such situations. Therefore, we can add a parameter α that gives the percentage of people who must stand for others to ignore their initial signals and decide to stand up regardless. In some fields, like economics, we might even delve a bit deeper into the notion of α and see if we can tie it to some first principles, for example, perhaps people realize that their signals of the performance are imperfect and thus they update them using the information gathered by observing the behavior of others. We will avoid such complications here and just assume that α exists for whatever reason.

¹There are circumstances, such as the annual State of the Union address before the U.S. Congress, where such behavior is a bit more orchestrated.

Our elaborated model provides some new insights into the world. If the initial group of people standing exceeds α percent, then everyone will rise; if it falls short of this value, then the standing ovation will remain at its initial level. Again, we can tie the elements of the model to a prediction about the world. By knowing the likelihood of various-sized initial ovations, we can predict (given an α) the likelihood of everyone else joining the ovation.

As clean and elegant as the mathematical model may be, it still leaves us wanting some more illumination. For example, we know that real ovations do not behave in the extreme way predicted by the model; rather, they often exhibit gradual waves of participation and also form noticeable spatial patterns across the auditorium. In the model's current form, too much space exists in between what it illuminates and what we want to know about the real world.

To capture this additional illumination, we might extend the mathematical model even further by using ideas from complex systems. This approach may require us to model using a different substrate, most likely indirect computation rather than direct mathematics, but for the moment this choice is less important than the directions we wish to take the modeling. The first elaboration we could undertake is to place each person in a seat in the auditorium, rather than assuming that they attend the theater on the head of a pin. Furthermore, we might want to assume that people have connections to one another, that is, that people arrive and sit with acquaintances (see figure 2.1).²

Once we allow people to sit in a space and locate next to friends, the driving forces of the model begin to change. For example, the initial assumption of independent signals is now suspect. It is likely that people seated in one part of the theater (or "side of the aisle") receive a different set of signals than others. Locations not only determine physical factors, such as which other patrons someone can see, but also may reflect a priori preferences for the performance that is about to begin. Similarly, in an audience composed of friends and strangers, people may differentially weight the signals sent by their friends, either because of peer pressure or because the friendships were initially forged based on common traits.

Assuming that individuals now have locations and friends introduces an important new source of heterogeneity. In the mathematical model, the only heterogeneity came from the different draws of ϵ_i . Now, even "identical" individuals begin to behave in quite different ways, depending on where, and with whom, they are seated.

²We once had a group of economics graduate students model the standing ovation. Not one of them allowed the possibility of people attending the theater with acquaintances. We hope this is more a reflection of how economists are trained than of how they live.




Figure 2.1. Two views of modeling the standing ovation. In its simplest form, the model requires that everyone shares the same seat in the auditorium (*left*), while the more elaborate model (*right*) allows space, friendship connections, and physical factors like vision to play a vital role in the system. While the simple model might rely on traditional tools like formal mathematics and statistics, the more elaborate model may require new techniques like computational models using agent-based objects to be fully realized.

The dynamics of the model also becomes more complicated. In the original model, we had an initial decision to stand, followed by a second decision based on how many people stood initially. After this second decision, the model reached an equilibrium where either the original group remained standing or everyone was up on their feet. The new model embodies a much more elaborate (and likely realistic) dynamics. In general, it will not be the case that the model attains an equilibrium after the first two rounds of updating. Typically, the first round of standing will induce others to stand, and this action will cause others to react; in this way, the system will display cascades of behavior that may not settle down anytime soon.

These two modeling approaches illuminate the world in very different ways. In the first model either fewer than α percent stand or everyone does; in the second it is possible to have any percentage of people left standing. In the first model the outcome is determined after two periods; in the second cascades of behavior wash over the auditorium and often reverberate for many periods. In the first model everyone's influence is equal; in the second influence depends on friendships and even seat location. Oddly, the people in the front have the most visual influence on others yet also have the least visual information, whereas those in the back with the most information have the least influence (think of the former as celebrities and the latter as academics).

The second model provides a number of new analytic possibilities. Do performances that attract more groups lead to more ovations? How does changing the design of the theater by, say, adding balconies, influence ovations? If you want to start an ovation, where should you place your shills? If people are seated based on their preferences for the performance, say, left or right side of the aisle or more expensive seats up front, do you see different patterns of ovations?

Although standing ovations per se are not the most pressing of social problems, they are related to a large class of important behaviors that is tied to social contagion. In these worlds, people get tied to, and are influenced by, other people. Thus, to understand the dynamics of a disease epidemic, we need to know not only how the disease spreads when one person contacts another but also the patterns that determine who contacts whom over time. Such contagion phenomena drive a variety of important social processes, ranging from crime to academic performance to involvement in terrorist organizations.

2.2 WHAT'S THE BUZZ?

Heterogeneity is often a key driving force in social worlds. In the Standing Ovation problem, the heterogeneity that arose from where people sat and with whom they associated resulted in a model rich in behavioral possibilities. If heterogeneity is a key feature of complex systems, then traditional social science tools—with their emphases on average behavior being representative of the whole—may be incomplete or even misleading.

In many social scenarios, differences nicely cancel one another out. For example, consider tracking the behavior of a swarm of bees. If you observe any one bee in the swarm its behavior is pretty erratic, making an exact prediction of that bee's next location nearly impossible; however, keep your eye on the center of the swarm—the average—and you can detect a fairly predictable pattern. In such worlds, assuming behavior embodied by a single representative bee who averages out the flight paths of all of the bees within the swarm both simplifies and improves our ability to predict the future.

2.2.1 Stay Cool

While differences can cancel out, making the average a good predictor of the whole, this is not always the case. In complex systems we often see differences interacting with one another, resulting in behavior that deviates remarkably from the average.

To see why, we can return to our bees. Genetic diversity in bees produces a collective benefit that plays a critical function in maintaining hive temperature (Fischer, 2004). For honey bees to reproduce and grow, they must maintain the temperature of their hive in a fairly narrow range via some unusual behavioral mechanisms. When the hive gets too cold, bees huddle together, buzz their wings, and heat it up. When the hive gets too hot, bees spread out, fan their wings, and cool things down.

Each individual bee's temperature thresholds for huddling and fanning are tied to a genetically linked trait. Thus, genetically similar bees all feel a chill at the same temperature and begin to huddle; similarly, they also overheat at the same temperature and spread out and fan in response.

Hives that lack genetic diversity in this trait experience unusually large fluctuations in internal temperatures. In these hives, when the temperature passes the cold threshold, all the bees become too cold at the same time and huddle together. This causes a rapid rise in temperature and soon the hive overheats, causing all the bees to scatter in an over ambitious attempt to bring down the temperature. Like a house with a primitive thermostat, the hive experiences large fluctuations of temperature as it continually over- and undershoots its ideals.

Hives with genetic diversity produce much more stable internal temperatures. As the temperature drops, only a few bees react and huddle together, slowly bringing up the temperature. If the temperature continues to fall, a few more bees join into the mass to help out. A similar effect happens when the hive begins to overheat. This moderate and escalating response prevents wild swings in temperature. Thus, the genetic diversity of the bees leads to relatively stable temperatures that ultimately improve the health of the hive.

In this example, considering the average behavior of the bees is very misleading. The hive that lacked genetic diversity—essentially a hive of averages—behaves in a very different way than the diverse hive. Here, average behavior leads to wide temperature fluctuations whereas heterogeneous behavior leads to stability. To understand this phenomenon, we need to view the hive as a complex adaptive system and not as a collection of individual bees whose differences cancel out one another.

2.2.2 Attack of the Killer Bees

We next wish to consider a model of bees attacking a threat to the hive.³ Some bees go through a maturation stage in which they guard the

³This is a simplified version of models of human rioting constructed by Granoveter (1978) and Lohmann (1993). Unlike the previous example, the direct applicability to bees is more speculative on our part.

entrances to the hive for a short period of time. When a threat is sensed, the guard bees initiate a defensive response (from flight, to oriented flight, to stinging) and also release chemical pheromones into the air that serve to recruit other bees into the defense.

To model such behavior, assume that there are one hundred bees numbered 1 through 100. We assume that each bee has a response threshold, R_i , that gives the number of pheromones required to be in the air before bee i joins the fray (and also releases its pheromone). Thus, a bee with $R_i = 5$ will join in once five other bees have done so. Finally, we assume that when a threat to the hive first emerges, R bees initiate the defensive response (to avoid some unnecessary complications, let these bees be separate from the one hundred bees we are watching). Note that defensive behavior is decentralized in a beehive: it is initiated by the sentry activities of the individual guard bees and perturbed by each of the remaining bees based only on local pheromone sensing.

We consider two cases. In the first case, we have a homogeneous hive with $R_i = 50.5$ for all i . In the second case, we allow for heterogeneity and let $R_i = i$ for all i . Thus, in this latter case each bee has a different response threshold ranging from one to one hundred. Given these two worlds, what will happen?

In the homogeneous case, we know that a full-scale attack occurs if and only if $R > 50$. That is, if more than fifty bees are in the initial wave, then all of the remaining one hundred will join in; otherwise the remaining bees stay put. In the heterogeneous case, a full-scale attack ensues for any $R \geq 1$. This latter result is easy to see, because once at least one bee attacks, then the bee with threshold equal to one will join the fray, and this will trigger the bee with the next highest threshold to join in, and so on.

Again, notice how average behavior is misleading. The average threshold of the heterogeneous hive is identical to that of the homogeneous hive, yet the behaviors of the two hives could not be more different. It is relatively difficult to get the homogeneous hive to react, while the heterogeneous one is on a hair trigger. Without explicitly incorporating the diversity of thresholds, it is difficult to make any kind of accurate prediction of how a given hive will behave.

2.2.3 Averaging Out Average Behavior

Note that the two systems we have explored, regulating temperature and providing defense, have very different behaviors linked to heterogeneity. In the temperature system, heterogeneity leads to stability. That is, increased heterogeneity improves the ability of the system to stabilize

on a given temperature. In the defense system, however, heterogeneity induces instability, with the system likely to experience wild fluctuations in response to minute stimuli.

The difference of response between the two systems is due to feedback. In the temperature system, heterogeneity introduces a negative feedback loop into the system: when one bee takes action, it makes the other bees less likely to act. In the defense system, we have a positive feedback loop: when one bee takes action, it makes the other bees more likely to act.

2.3 A TALE OF TWO CITIES

To explore further the modeling of complexity, we consider a simple world composed of two towns, each of which has three citizens. Furthermore, we assume that each town has to make a choice about an important public issue: whether to serve its citizens red or green chile at its annual picnic. Citizens possess preferences over chile and strongly prefer one type over the other.⁴ To make the analysis interesting, we assume that two of the citizens in each town prefer green to red chile while the remaining person prefers the opposite.

Though stark, this scenario builds from an extensive literature in the social sciences on the allocation of public goods and services to citizens (Samuelson, 1954; Tiebout, 1956). Public goods and services flow across all members of society without exclusion or diminution once offered. Moreover, as we will see, the model also touches on even deeper issues surrounding the decentralized sorting of agents within a complex adaptive system.

Before we can explore the behavior of the model, we need to define two further elements. The first is how does a town, given a set of citizens, select what chile to offer. The second is how do citizens react to the choices of the towns.

A town could use several mechanisms to decide what type of chile to offer. It could employ a dictator, flip a coin, or implement some other political process, such as majority rule. For the moment, we will assume that each town uses majority rule. Given this scenario, majority rule implies that each town will always offer green chile (two votes to one). Note that this outcome is not ideal, as one citizen in each town always ends up consuming her less-preferred meal (see figure 2.2).

⁴For those who enjoy both, New Mexican restaurants offer the option of ordering your chile "Christmas."




Figure 2.2. A symmetric Tiebout world. Here two towns each have three citizens, two of whom prefer green to red chili. Both towns currently offer green chili at their annual picnic. Given this scenario, the system is at an equilibrium, even though two of the citizens are not getting their favorite chili.

Now, suppose we give our citizens some mobility, that is, any citizen is free to switch towns if she so desires. We assume that citizens will move only if the alternative town is offering a better meal. If each town is serving green chili, no citizen has any incentive to relocate and everyone stays put.

Yet, something should be done. The current situation possesses a tragic symmetry that prevents the red chili lovers from ever realizing their favored outcome since they are always the minority in either town. To improve this situation, we must find a way to break the symmetry.

One way to break the symmetry is to introduce some randomness into the system. For example, we could have one citizen randomly decide to move to the other town for whatever reason. If this citizen is a red chili lover, then the town she vacated is left with two green chili lovers and her new town now has two people who like red and two who like green chili. Instead, if the citizen that relocates is a green chili lover, then the vacated town is left with one of each type, while the other town now has three green and one red chili lover. Notice that regardless of who moves, we are always left with one town that is strongly green chili and one that has equal numbers of each type.

Given this situation, we would expect that eventually the town with a split vote will offer red instead of green chili. Once this occurs, we now have one town offering red and one offering green chili. The symmetry is now broken, and the citizens in each town can immediately re-sort themselves and self-select the town that perfectly meets their chili needs. This leaves one town offering green chili populated by four green chili lovers and one town offering red chili with two red chili lovers, and all




Figure 2.3. Broken symmetry in the Tiebout world. Once the two towns offer different types of chili—perhaps due to noise in the political system—the citizens will immediately re-sort themselves. The system again attains an equilibrium, though in this case each citizen now gets her favorite type of chili. Note that this new equilibrium is much more robust to minor perturbations than the former one.

of the citizens would be worse off if they moved (see figure 2.3). This latter configuration is quite stable to random moves of individuals, as a single citizen moving will not alter the majority in either town.

An alternative way to break the symmetry is to alter slightly the behavioral rules that control our citizens. Suppose that agents are willing to relocate if they can at least maintain their level of happiness (rather than improve it). Such a change in behavior allows for what biologists call *neutral mutations*, that is, movements in the underlying structure that do not directly impact outcomes. Even though neutral mutations do not have an immediate effect, they can lead to better outcomes eventually by changing what is possible. In the initial case, any of the citizens is willing to move since both towns offer the same type of chili. Regardless of who moves, one town is always left with a split vote, and the symmetry breaking we saw previously is again possible.

The system demonstrates some key features of complex adaptive social systems. First, we have a web of connections that, in this case, results from citizens linking to one another by being resident in a given town. Second, we see change induced by choices made by all of the different types of agents in the system. Citizens must decide where to move, and towns must decide what type of chili to offer. Moreover, the system as a whole must “decide” how to sort the citizens among the towns, although this latter “choice” is not a conscious calculation of the system per se, but rather an implicit computation resulting from the decentralized choices made by each citizen and town. The model also

demonstrates how a social system can get locked into an inferior outcome and how, with the introduction of noise or different behavioral rules, it can break out of such outcomes and reconfigure itself into a better arrangement.

The model also incorporates other key themes in complex adaptive social systems: equilibria, dynamics, adaptation, and the power of decentralized interactions to organize a system. The system has multiple equilibria, some of which are inferior to others. The key dynamics that occur in the model are the choice dynamics of each town induced by the voting system and the movement dynamics of each citizen implied by her preferences and each town's offerings. Note that these dynamics imply that towns adapt to citizens, while citizens also adapt to towns. Finally, we see how the system's dynamics result in local, decentralized behaviors that ultimately organize the citizens so that their preferences align with other citizens and each town's offerings align with its residents.

2.3.1 Adding Complexity

While our model gives us some useful intuitions and insights, it is also (quite intentionally) very limited. Like all good models, it was designed to be just sufficient to tell a story that could be understood easily yet have enough substance to provide some insights into broader issues. Moving beyond the limitations of this model is going to require some compromises—namely, if we want to expand the potential for insights, we will likely need to complicate the model and, perhaps, muddy the analytic waters.

For example, suppose we wish to explore more fully Tiebout's (1956) concept of "voting with your feet." That is, can we characterize better the ability of social systems to sort citizens dynamically among towns? The simplifications in the preceding model were rather drastic; we had two towns, six citizens, a single issue (choice of chile), and a single mechanism to determine what each town offered (majority rule). If we wish to go beyond any of these constraints, we will quickly start to run into trouble in pursuing the thought experiment framework used previously.

In economics, formal modeling usually proceeds by developing mathematical models derived from first principles. This approach, when well practiced, results in very clean and stark models that yield key insights. Unfortunately, while such a framework imposes a useful discipline on the modeling, it also can be quite limiting. The formal mathematical approach works best for static, homogeneous, equilibrating worlds. Even in our very simple example, we are beginning to violate these desiderata. Thus, if we want to investigate richer, more dynamic worlds, we need

to pursue other modeling approaches. The trade-off, of course, is that we must weigh the potential to generate new insights against the cost of having less exciting analytics.

One promising alternative approach is the development of computation-based models. In the Tiebout system, through computation we can allow multiple towns and citizens, as well as more elaborate preference and choice mechanisms. Thus, we can consider a world in which each town must make binary choices over multiple issues, such as whether to, say, serve red or green chile at the annual picnic, allow smoking in public places, and set taxes either high or low. Once we admit multiple issues, our citizens will need to have more complicated preference structures to account for the more elaborate set of choices. This will imply that, instead of just two types of citizens, we now have a much more heterogeneous population. Finally, instead of using majority rule as the sole means by which a town chooses its offerings, we can admit a variety of other possible social choice mechanisms. For example, towns might use a form of democratic referendum where, like simple majority rule, citizens get to vote on each issue and the majority wins; or perhaps the towns could rely on political parties that develop platforms (positions on each possible choice) and then vie for the votes of the populace.

Rather than fully pursuing the detailed version of the model we just outlined (interested readers should see Kollman, Müller, and Page, 1997), here we provide just an overview. Using computation, we can explore a world with multiple issues, citizens, towns, choices, and choice mechanisms. For example, consider a model where each town must make binary decisions across eleven issues. Each citizen has a preference for each issue that takes the form of a (randomly drawn) weight that is summed across all of the choices in a town's platform to determine the citizen's overall happiness. Of particular interest at the moment is the effectiveness of different public choice mechanisms in allocating citizens to towns and towns to platforms.

We will allow towns to use a variety of choice mechanisms to determine what they will offer. At one extreme we can employ *democratic referenda* (essentially majority rule on an issue-by-issue basis), while at the other we will consider a party-based political processes whereby political parties propose platforms and then compete with one another for votes. In this latter mechanism, we can consider worlds with two or more parties, either where the winning party takes all in *direct competition* (that is, the winning party's platform is what the town offers) or where, in a system of *proportional representation*, the final platform offered by the town is a blend, weighted by votes, of each individual party's platform.

Again, we impose a simple dynamic on the system: the citizens in a town, mediated by the choice mechanism, determine what the town will offer across the eleven issues and, once that is determined, citizens look around and move to their favorite town based on their own preferences and each town's current offerings. We iterate this process multiple times and ultimately investigate the final match of citizens to towns and towns to issues. For the moment, we judge each mechanism only by its effectiveness at maximizing the overall happiness of the citizens after a fixed amount of time. Thus, a good outcome will have citizens with similar preferences living in the same town, and that town offering a platform that aligns well with the preferences of its, relatively homogeneous, residents.

To get our bearings, first consider the case of a world with only a single town. In such a world the dynamic implied by citizens moving from town to town is nullified, and the only dynamic element of the model is that arising from the town altering its offerings via the choice mechanism. Thus, the best outcome will depend on the ability of the choice mechanism to come up with a platform that closely matches the preferences of the population. We find that, under these conditions, democratic referenda lead to the best outcome, followed by two political parties competing under direct competition, then multiple parties with proportional representation, and finally more than two parties using direct competition. Under democratic referenda, the system immediately locks into the median position of the voters on each issue; under the other mechanisms, party competition can result in the town's platform changing from period to period and not necessarily achieving the median on any one issue. Under the preference structure of our model, the median voter position on each issue will typically maximize the overall welfare of a fixed group of citizens confined to *a single town*. Therefore, democratic referenda are the best mechanisms for maximizing social welfare in a world consisting of only a single town.

Oddly, when we allow additional towns into the system, democratic referenda no longer lead to the highest social welfare. In fact, the effectiveness of the different choice mechanisms is completely reversed, and democratic referenda become the worst possible institution rather than the best. (See figure 2.4.)

Why does this happen? Fortunately, computational models are quite amenable to exploring such questions; in essence, we have a laboratory on the desktop and can systematically propose, test, and eliminate key hypotheses to understand better the outcomes we are observing.

To develop some needed intuition, consider the following. If we are interested in maximizing the overall happiness of our citizens with multiple towns, we must achieve two ends. First, we need to sort




Figure 2.4. Results of a computational Tiebout model. As we increase the number of towns in the system, the effectiveness of the different choice mechanisms in achieving high social welfare completely reverses.

the citizens among the available towns so that citizens with similar preferences reside in the same town. Without such a sorting, the social welfare generated by each town will be compromised given the diversity of wants. Second, each town must choose across the issues so as to maximize the happiness of its residents. As noted, democratic referenda are very effective at deriving a stable platform of choices that maximizes happiness for a given town. Given this observation, the failure of democratic referenda with multiple towns must be related to their inability to sort adequately the citizens among the towns.

A deeper investigation into the dynamics of the system confirms that the mechanisms other than democratic referenda result in far more initial movement of the citizens among the towns. Democratic referenda tend to stabilize the system quickly, freezing the citizens and platforms in place after only a few iterations. That is, after only a few rounds each town is offering a fixed platform, and no citizen wants to move. The other mechanisms are much more dynamic, in the sense that the platforms of each town keep changing during the early periods and the citizens tend to migrate much more often. Eventually, even these latter systems settle down to a state with little platform change and few migrations.

Earlier we saw how noise in the system allows it to break out of inferior sortings and to lock into superior ones. Of course, noise alone is not sufficient to guarantee a quality sorting of the citizens—to achieve high levels of social welfare, you need the noise to result in relatively homogeneous groups of citizens in each town and each town implementing platforms that approach something akin to the median issue positions across the local voters.

In fact, the choice mechanisms that work best in our more complicated model have a subtle, but key, property. These mechanisms tend to

introduce noise into the system when the local citizens' preferences are heterogeneous and to reduce this noise as the citizens become more homogeneous. Thus, if the citizens in a given town have very different preferences from one another, the more successful mechanisms will tend to induce more sorting. As the local citizens become more and more similar, these same mechanisms tend to converge on something approaching the median position on each issue. The notion that good political mechanisms should have such an inherent design is somewhat intuitive: if everyone in a district wants the same thing, the mechanism should deliver it; if, on the other hand, there is a diversity of wants, then the political process should jump around among the various options.

This "natural" annealing process turns out to be a very effective way to promote the decentralized sorting of citizens among towns. To achieve the highest social welfare, we need homogeneous collections of citizens in each town receiving roughly the median policy of the local residents. When the overall sorting of the system is poor, that is, when the mix of citizens in each town tends to be heterogeneous rather than homogeneous, then we should introduce a lot of noise into the platforms. Such noise will induce some citizens to migrate, and this migration will often cascade across other towns and result in a fairly large-scale resorting of the citizens. However, as the citizens become better sorted, that is, as each town becomes more homogeneous, the choice mechanisms should "cool" (anneal) the system by stabilizing on platforms that closely match the relatively homogeneous preferences of each town's citizens.

The notion of annealing to improve the structure of decentralized systems was first recognized a few thousand years ago in early metal-working. Heating metal tends to disrupt the alignment of (add noise to) the individual atoms contained in a metal; then, by slowly cooling the metal, the atoms can align better with one another, resulting in a more coherent structure. Kirkpatrick, Gelatt, and Vecchi (1983), based on some ideas from Metropolis et al. (1953), suggested that "simulated" annealing could be used as an effective nonlinear optimization technique. Thus, the Tiebout model shows how different institutions (here, public choice mechanisms) can become natural annealing devices that ultimately result in a decentralized complex adaptive social system seeking out global social optima.

By pursuing the more elaborate computational model, we achieved a number of useful ends. First, we were able to investigate some important new questions, such as the impact of citizen heterogeneity, multiple towns, and differing choice mechanisms on the ability of a system to achieve high social welfare. Second, the more elaborate model provided some new insights into how such systems behave, the most important

being the idea that well-structured noise can jolt a system out of inferior equilibria and lead it toward superior ones, and that choice mechanisms can be designed to introduce such noise in a decentralized way. This intuition is contrary to our usual way of thinking about such problems. Noise is usually considered to be a disruptive force in social systems, resulting in perturbations away from desirable equilibria rather than a means by which to attain them.

The complex-systems approach also allows us to explore the system's robustness. The system autonomously responds to all kinds of changes. We can randomly change the preference profiles for some of the citizens, introduce or remove issues, and so on. In each case, the system will adapt to these changes by presenting new platforms and inducing new migrations. Depending on the rate of change, we may see the system slowly moving through a sequence of equilibria or find ourselves with a world constantly in flux.

Although we have focused our discussion on a political system allocating public goods, the basic ideas embodied in the model are much broader. Decentralized sorting arises across a variety of domains. For example, workers seek jobs, traders march with trading partners, individuals form social groups and clubs, and industries sort out standards and geographic locations. All of these scenarios could be cast as decentralized sorting problems similar to the one just discussed. Moreover, we could use the ideas developed here to formulate new kinds of decentralized sorting algorithms that could be used to, say, sort computer users across resources (like servers) or on-line communities (like bulletin boards or tagging).

The Tiebout world we have explored is a nice example of a much broader quest. There is nothing that is unique about the Tiebout world in terms of its complexity. Like most social systems, it displays some dynamics, heterogeneity, and agent interactions that, even in vastly simplified models, can easily introduce complexity. Even a little bit of complexity implies that the conventional tools we often employ to investigate the world will be limited in their ability to yield insights and prescriptions. We are not claiming that these more conventional tools are useless; indeed, they are an important complement in any quest to understand the world.⁵ The computational approach pursued here provided a number of new directions and insights that both enhanced, and was enhanced by, more conventional techniques.

⁵In the example presented, the investigation of the system first began with the more elaborate computational model. Based on that experience, we were able to develop the thought experiment with which we opened this section.

2.4 NEW DIRECTIONS

The notion that real social systems often result in complex worlds is nothing new. More than two hundred years ago Adam Smith described a world where the self-interested social behavior of butchers, brewers, bakers, and the like resulted in the emergence of a well-defined order. While social science has been able to develop tools that can help us decipher some parts of this system, we have yet to understand fully the inner workings of the world around us. Unfortunately, we are at the mercy of a world characterized by change and connections, and thus our ability to make sense of our world is often undermined by the same characteristics that make it so fascinating and important.

The application of computational models to the understanding of complex adaptive social systems opens up new frontiers for exploration. The usual bounds imposed by our typical tools, such as a need to keep the entire model mathematically tractable, are easily surmounted using computational modeling, and we can let our imagination and interests drive our work rather than our traditional tools. Computational models allow us to consider rich environments with greater fidelity than existing techniques permit, ultimately enlarging the set of questions that we can productively explore. They allow us to keep a broad perspective on the multiple, interconnecting factors that are needed to understand social life fully. Finally, they give us a way to grow worlds from the ground up and, in so doing, provide a viable means by which to explore the origins of social worlds.

As we move into new territory, new insights begin to spill forth. Sometimes these insights are strong enough to stand on their own; at other times, they provide enough of a purchase on the problem that we can employ time-tested older techniques to help us verify and illuminate the newly acquired insights. On occasion, of course, computational models leave us with a jumbled mess that may be of no help whatsoever, though, with apologies to Tennyson, 'tis better to have explored and lost than never to have explored at all.

Social science has failed to answer, or simply ignored, some important questions. Sometimes important questions fall through the cracks, either because they are considered to be in the domain of other fields (which may or may not be true) or because they lie on the boundaries between two fields and subsequently get lost in both. More often than not, though, questions are just too hard and therefore either get ignored or (via some convoluted reasoning) are considered unimportant. The difficulty of answering any particular scientific question is often tied to the tools we have at hand. A given set of tools quickly sorts problems into those

we could possibly answer and those we perceive as too difficult to ever sort out. As tools change, so does the set of available questions.

Throughout this book, we pursue the exploration of complex systems using a variety of tools. We often emphasize the use of computational models as a primary means for exploring these worlds for a number of reasons. First, such tools are naturally suited to these problems, as they easily embrace systems characterized by dynamics, heterogeneity, and interacting components. Second, these tools are relatively new to the practice of social science, so we take this as an opportunity to help clarify their nature, to avoid misunderstandings, and generally to advance their use. Finally, given various trends in terms of the speed and ease of use of computation and diminishing returns with other tools, we feel that computation will become a predominant means by which to explore the world, and ultimately it will become a hallmark of twenty-first-century science.

2.5 COMPLEX SOCIAL WORLDS REDUX

We see complicated social worlds all around us. That being said, is there something more to this complication? In traditional social science, the usual proposition is that by reducing complicated systems to their constituent parts, and fully understanding each part, we will then be able to understand the world. While it sounds obvious, is this really correct? Is it the case that understanding the parts of the world will give us insight into the whole? If parts are really independent from one another, then even when we aggregate them we should be able to predict and understand such "complicated" systems. As the parts begin to connect with one another and interact more, however, the scientific underpinnings of this approach begin to fail, and we move from the realm of complication to complexity, and reduction no longer gives us insight into construction.

2.5.1 Questioning Complexity

Thus, a very basic question we must consider is how complex, versus complicated, are social worlds. We suspect that the types of connections and interactions inherent in social agents often result in a complex system. Agents in social systems typically interact in highly nonlinear ways. Of course, there are examples, such as when people call one another during the course of a normal day, where agent behavior aggregates in ways that are easily described via simply statistical processes.

Nonetheless, a lot of social behavior, especially with adaptive agents, generates much more complex patterns of interaction. Sometimes this is an inevitable feature of the nature of social agents as they actively seek connections with one another and alter their behavior in ways that imply couplings among previously disparate parts of the system. Other times, this is a consequence of the goal-oriented behavior of social agents. Like bees regulating the temperature of the hive, we turn away from crowded restaurants and highways, smoothing demand. We exploit the profit opportunities arising from patterns generated by a stock market and, in so doing, dissipate their very existence. Like bees defending the hive, we respond to signals in the media and market, creating booms, busts, and fads.

If social worlds are truly complex, then we might need to recast our various attempts at understanding, predicting, and manipulating their behavior. In some cases, this recasting may require a radical revision of the various approaches that we traditionally employ to meet these ends. At the very least, we need to find ways to separate easily complex systems from merely complicated ones. Can simple tests determine a system's complexity? We would like to understand what features of a system move it from simple to complex or vice versa. If we ultimately want to control such systems, we either need to eliminate such forces or embrace them by productively shaping the complexity of a system to achieve our desired ends.

Another important question is how robust are social systems. Take a typical organization, whether it be a local bar or a multinational corporation. More often than not, the essential culture of that organization retains a remarkable amount of consistency over long periods of time, even though the underlying cast of characters is constantly changing and new outside forces are continually introduced. We see a similar effect in the human body: typical cells are replaced on scales of months, yet individuals retain a very consistent and coherent form across decades. Despite a wide variety of both internal and external forces, somehow the decentralized system controlling the trillions of ever-changing cells in your body allows you to be easily recognized by someone you have not seen in twenty years. What is it that allows these systems to sustain such productive, aggregate patterns through so much change?

Our modeling of social agents tends toward extremes: we either consider worlds composed of remarkably prescient and skilled agents or worlds populated by morons. Yet, we know that real agents exist somewhere in between these two extremes. How can we best explore this middle ground? A key issue in exploring this new territory is figuring out the commonalities among adaptive agents. While it is easy to specify behavior at the extremes, as we move into the middle ground, we are

suddenly surrounded by a vast zoo of curious adaptive creatures. If we are stuck having to study every creature individually, it will be difficult to make much progress, so our underlying hope is that we can find some way to distill this marvelous collection of behaviors down to just a few prototypical ones. Once this is done, we can begin to make progress on a science of adaptive behaviors.

We know that adaptive agents alter the world in which they live. What we do not know is how much agent sophistication is required to do so effectively and what other conditions are necessary for this to happen. In general, the link between agent sophistication and system outcome is poorly understood. Theoretical work in economics suggests that optimizing agents out for their own benefit can, without intention, lead a market system toward efficiency under the right conditions. Moreover, experimental and computational work suggests that such outcomes are possible even with nonoptimizing agents. Ultimately, it would be nice to have a full characterization of the interplay between adaptation and optimality in social systems.

Another realm where we have a limited understanding is the role of heterogeneity in systems. We know that in, say, ecological systems homogeneity can be problematic. For example, using a few genetic lines of corn maximizes short-term output but subjects the entire crop to a high risk of destruction if an appropriate disease vector arises. Homogeneity in social systems may have similar effects. A homogeneous group of agents in, say, a market might result in a well-functioning institution most of the time, but leave the possibility that these behaviors could synchronize in such a way that on occasion the market will crash. By introducing an ecology of heterogeneous traders, such fluctuations might be mitigated. Perhaps heterogeneity is an important means by which to improve the robustness of systems. If so, does this work via complexifying the system or via some other mechanism?

The idea of social niche construction is also important. Agents, by their activities, can often alter the world they inhabit and, by so doing, form new niches. For example, the development of membranes early in the history of life on Earth allowed various biological components to bind together and isolate themselves from the external world. This fundamentally altered their local environment creating new opportunities for interacting with the world. Similarly, the formation of merchant guilds, corporations, and political organizations fundamentally altered both the internal world faced by agents and the external world in which these new entities operated. We would like to know when and how agents construct such niches.

The role of control on social worlds is also of interest. The ability to direct the global behavior of a system via local control is perhaps one

of the most impressive, yet mysterious, features of many social systems. In the natural world, tens of thousands of swarm-raiding army ants can form cohesive fronts fifty feet across and six feet deep that can sweep through the forest for prey. This entire operation is controlled via locally deposited chemical signals. At a grander scale, a vast decentralized system of human markets of all types orchestrate the activities of billions of individuals across the span of continents and centuries. Fully understanding how such decentralized systems can so effectively organize global behavior is an enduring mystery of social science. We do have some hints about how this can happen. For example, adding noise to the system (as we saw in our Tiebout model) may actually enhance the ability of a system to find superior outcomes. We also know that some simple heuristics that arise in some contexts, such as the notion that in a market new offers must better existing ones, result in powerful driving forces that enhance the ability of the system to form useful global patterns.

Every social agent receives information about the world, processes it, and acts. For example, in our Tiebout model, the behavior of the citizens was very straightforward (get information about the offerings of the various towns, process this via your preferences, and act by moving to your favorite town), while that of each town was a bit more elaborate (get information about the preferences of the citizens across the issues, process this via either exact or adaptive mechanisms to develop a new platform, and act by implementing this platform).

Traditional economic modeling tends to have a fairly narrow view of the issues that arise in acquiring information, processing it, and acting. In these models, agents tend to have access to all available information, process it with good fidelity and exacting logic directed toward optimization, and act accordingly. Where traditional economics gains its power is that these restrictions make for relatively easily modeling across a broad spectrum of social activity. Notwithstanding the apparent success of this approach in some domains, one wonders whether such a restricted view of these three elements is appropriate. While clearly these restrictions give us leverage from which to generate insights across a variety of social realms, we also know that in many cases the core tenets driving the approach are misplaced (though it is still an open issue whether this matters in the end). For example, the recent wave of work in behavioral economics is based on the notion that the processing of information by humans may take place in ways that dramatically diverge from the traditional view.

Much of the work we discuss throughout this book relaxes the traditional assumptions about information acquisition, processing, and acting. We want to consider models in which information is selectively

acquired across restricted channels of communication. We want to look at agents that process information via adaptive mechanisms or restricted rules rather than exacting logic. We want to explore models in which actions are often limited and localized. How do all of these factors embody social complexity and what does this mean for the practice of social science?

On Emergence

He intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention.

—Adam Smith, *Wealth of Nations*

Any sufficiently advanced technology is indistinguishable from magic.

—Arthur C. Clarke, *Profiles of the Future*

MUCH OF THE FOCUS of complex systems is on how systems of interacting agents can lead to emergent phenomena. Unfortunately, emergence is one of those complex systems ideas that exists in a well-trodden, but relatively untraveled, bog of discussion. The usual notion put forth underlying emergence is that individual, localized behavior aggregates into global behavior that is, in some sense, disconnected from its origins. Such a disconnection implies that, within limits, the details of the local behavior do not matter to the aggregate outcome. Clearly such notions are important when considering the decentralized systems that are key to the study of complex systems. Here we discuss emergence from both an intuitive and a theoretical perspective.

The notion of emergence has deep intuitive appeal. Consider for the moment standing up close to a pixelated picture (see figure 4.1).¹ While each individual pixel can be easily understood in terms of its shape, color, hue, and other properties, it is typically impossible to figure out the entire image by simply scanning across the pixels at close range. As the observer moves back, there is some point at which the overall image begins to resolve, and the pixels become indistinguishable. Once the image has resolved, we can typically make many possible alterations to individual pixels and still not impact the overall image. Indeed, depending on the image, certain types of global pixel properties, such as color, may not even be needed to have a good sense of the final image.

We may see emergence at many levels. For example, instead of having each pixel composed of a single solid color, we could replace it with a tiny picture whose overall properties can approximate the key




Figure 4.1. Emergence from a mosaic. While the properties of each tile are easy to understand at close range, the true nature of the full image is impossible to comprehend from such information. It is only when you view the mosaic from far away that emergence allows the entire image to become viable.

visual attributes of the previous pixel.² Of course, each of these tiny photographs, each emerging from its own set of pixels, could stand by itself. Thus, there is the possibility of multiple layers of emergence, where pictures become pixels that become pictures that become pixels and so on. This may lead us to a “Horton-Hears-A-Who” theory of the universe, in which the world is composed of stacked layers of emergence.

The notion of emergence at many levels is an important one, as each level of emergence can serve as a convenient point at which to dissect the larger system and attempt to understand some of its secrets. Indeed, the boundaries of modern science rely on this property—for example, physics resolves into chemistry, which resolves into biology, which resolves into psychology, which resolves into economics, and so on. Each new science is able to exploit the emergence that is attained by the previous level.

While this metaphor of emergence is very appealing, it leaves open the question of how it should fit into scientific discourse. Part of the innate appeal of emergence is the surprise it engenders on the part of the observer. Many of our most profound experiences of emergence come from those systems in which the local behavior seems so entirely disconnected from the resulting aggregate as to have arisen by magic, echoing Clarke’s observation about advanced technology. Examples of such dramatic disconnections include photomosaic pictures, the order and persistence of beehives and forging ant colonies through simple sets of localized signals, and the stability of a market price generated by the often chaotic and heterogeneous efforts of traders.

¹For the more romantic among you, assume a stained glass window.

²The technique of photomosaic pictures exploits this idea.

Alas, surprise and ignorance are closely related. It could be that emergent behavior is simply reflective of scientific ignorance rather than some deeper underlying phenomenon. What may start out as a mystical emergent phenomenon, such as planetary motion prior to Kepler, may turn out to be something rather simple—in the case of Kepler, *just* an ellipse. If all such scientific conundrums can be easily resolved, then perhaps it is true that all of our world is *just* physics. Nonetheless, whether our fascination with emergent phenomena is driven by ignorance or a more profound scientific mystery matters little. Profound scientific mysteries often get resolved in such a way that our prior ignorance becomes apparent, yet it is the ignorance that drives the quest for understanding.

4.1 A THEORY OF EMERGENCE

To move forward on the scientific exploration of emergence, it is useful to consider what types of theoretical ideas are possible in this area. As we have discussed, emergence is a phenomenon whereby well-formulated aggregate behavior arises from localized, individual behavior. Moreover, such aggregate patterns should be immune to reasonable variations in the individual behavior. Ideally, what we would like to develop are theorems about such phenomena, and, fortunately, at least one such theorem has existed since the early 1700s.

The theorem, the Law of Large Numbers (and its various offshoots, including the Central Limit Theorem), was developed by statisticians over the past few hundred years. It is of interest because it provides some relatively general conditions under which a certain type of aggregate behavior can emerge from the stochastic, microlevel actions of individual agents. Suppose that each individual agent's behavior is summarized by a random variable, X . Furthermore, assume that these variables are mutually independent, have a common distribution, and a mean equal to μ . According to the Law of Large Numbers, the probability that the mean will differ from μ by less than some arbitrary amount tends to one as we increase the number of agents in the system.

Thus, in such systems there is a stable, aggregate property (here the expected value of the common distribution) that emerges from aggregating the activities of the agents. Moreover, this aggregate property is robust to many underlying assumptions about the agents. In the foregoing case of the Law of Large Numbers, the only restriction is that the common distribution has mean μ ; other than that, we can vary any of its other characteristics and still maintain the identical aggregate behavior.

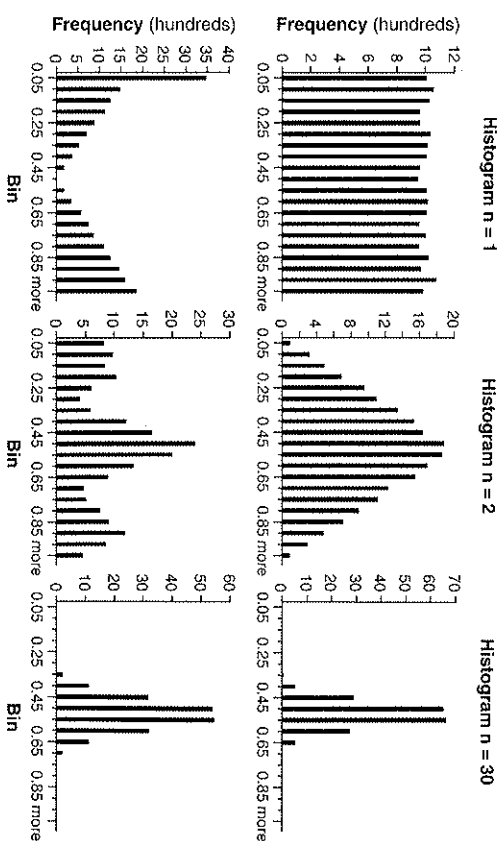


Figure 4.2. Central Limit Theorem. Notwithstanding the form of the initial distribution underlying the random process, the distribution of the mean of the variables generated by this process converges to a Normal distribution as we take larger and larger samples. In the top panel we start with a Uniform distribution, while in the bottom one we begin with a V-shaped distribution. In either case, a Normal distribution “emerges” as the sample size increases. Thus, the macrobehavior resulting from the aggregation of a remarkably diverse set of potential microbehaviors results in a very robust and predictable outcome—a hallmark of emergence.

The Central Limit Theorem provides another example of such a result (see figure 4.2). If we add the assumption that the variance of the common distribution is finite, then the distribution of the average (our aggregate property) will converge to a well-known “normal” form. The remarkable implication of this theorem is that, for an amazing variety of underlying agent behaviors, the global behavior that emerges is described by a simply specified, common form.

As Coates (1956) points out, without these laws, much of the behavior of the social worlds we live in would fall apart. Various activities, ranging from driving on the highway to enjoying the outdoors, would either be excessively crowded or notably desolate at the strangest times, stores and restaurants would run out of the oddest things, life insurance companies and telephone systems would fail, and so on.

The emergence theorems provide useful descriptions of a certain type of complexity that Weaver (1958), in a deeply prescient article, called *disorganized complexity*. The Law of Large Numbers works because as we add more and more independent agents to the world, the vagaries

of the stochastic elements, quite literally, average out. With only a few agents, these stochastic elements make it impossible to predict with any certainty the aggregate behavior because individual variation overwhelms any potential predictability, but as we increase the number of agents involved in the world, individual variations begin to cancel one another out, and systemwide predictions become possible.

The key feature of disorganized complexity is that the interactions of the local entities tend to smooth each other out. In the case of the Law of Large Numbers, an unusually high value for one random value is compensated for by an unusually low value of another. Thus, while it is difficult to predict the point at which, say, a particular rain drop meandering down a roof will fall into the gutter, it is easy to predict the activity at any point of the gutter during a rainstorm, as the various meanderings of the drops tend to disrupt one another sufficiently as they flow down the roof so as to spread out the water in a predictable way. Similarly, while predicting the motion of a planet surrounded by a few neighbors is difficult, it is easy to calculate its motion when it is among a sea of other planets, as the various gravitational forces that come into play begin to cancel one another out, and soon only the mean force becomes important. Other phenomena, ranging from population genetics to physical properties like temperature and pressure, also fall within the realm of disorganized complexity.

Thus, in cases of disorganized complexity, it should be easy to derive fairly precise emergence theorems based on fundamental concepts that are centuries old. Unfortunately, disorganized complexity accounts for only one part of our world.

4.2 BEYOND DISORGANIZED COMPLEXITY

Consider a picture of a face that is composed of black and white pixels. In such a picture, the pixels have relationships to one another that are quite important if we are to recognize the face. Some changes to the picture will not cause us to “lose” the face; for example, having a few of the pixels randomly change color, or even allowing some neighboring pixels to switch places, will preserve the “face.” Even some radical changes may not impact our ability to perceive the face, such as altering the color of related pixels (think Warhol) or just showing the important edges of the photograph (see figure 4.3). Moreover, if we are careful, we may be able to “capture” the image with just a few carefully drawn lines, as is done in caricature drawings.

Nonetheless, while we can make some slight changes to the pixels (or even some carefully designed radical changes) and still maintain the




Figure 4.3. Beyond disorganized complexity. The essence of the photograph remains robust to a variety of radical changes. All of these transformations keep intact the relationship of key parts of the original photograph (*left*).

image, doing anything more is likely to destroy the image. As we start to impact more and more pixels (by either randomly altering their colors or allowing neighbors to switch places), we quickly descend into the realm of disorganized complexity. In such a world, the photograph quickly resembles the white noise we see on the television (at least prior to the advent of late-night infomercials) when stations end their broadcast day. While it would be possible to construct the usual disorganized-complexity emergence theorems about, say, the average tone of the picture or the likelihood of an eyelike shape arising somewhere in the photograph, such theorems would fail to capture the essence of the problem of understanding how a decentralized set of pixels can emerge into a familiar face.

Thus, disorganized complexity, while often useful, leaves out a lot of interesting complexity-related phenomena. Disorganized-complexity emergence theorems can be used to calculate the vanishingly small probability that a room full of apes randomly banging away at typewriters will come up with *Hamlet*. Of course, a close relative of an ape did write *Hamlet*, but obviously not by randomly placing words to parchment and hoping for the best. Similarly, while disorganized-complexity theorems can be used to predict the life-span of a human body or a beehive composed of individual agents (cells or bees, respectively), they do not provide any insight into how the various communication and behavioral pathways among the individual agents are able to aggregate into these larger-scale organizations that survive and have behaviors on scales that are completely different than their constituent parts.

Explorations of complex systems have begun to identify the emergent properties of interacting agents—for want of a better term, *organized complexity*. We often see unanticipated statistical regularities emerging in complex systems. These regularities go beyond the usual bounds covered

by Central Limit Theorems and such. In chapter 9 we explore a model of sand piles in which we randomly drop grains of sand onto a table. A pile forms as the sand falls, and eventually grains begin to run off the edges of the table in avalanches of various sizes. The distribution of avalanche sizes follows a power law that implies behavior that is quite different from that arising from a normal distribution.

Agent intention can also alter the patterns that emerge in complex systems. In the case of the Sand Pile model, if we give the falling grains of sand a bit of control on where they land and some desires (like maximizing the size of the resulting avalanche), the system is no longer governed by a power law and instead enters a bizarre periodic cycle. As we give agents even more strategic ability, we often see elaborate dances of strategies, with good and bad epochs, cycles, and crashes.

In systems characterized by the Central Limit Theorem, interactions cancel one another out and result in a smooth bell curve. In complex systems, interactions reinforce one another and result in behavior that is very different from the norm.³ The complex phenomena that arise in physical systems (like earthquakes, floods, and fires) and social ones (like stock market crashes, riots, and traffic jams) are decidedly not “normal,” nor are the patterns that emerge as we see birds flock, fish school, and pedestrians follow sidewalks demarcated by invisible traffic lanes.

4.2.1 *Feedback and Organized Complexity*

When interactions are not independent, feedback can enter the system. Feedback fundamentally alters the dynamics of a system. In a system with negative feedback, changes get quickly absorbed and the system gains stability. With positive feedback, changes get amplified leading to instability.

For example, consider a world in which we have one hundred consumers, each of whom must choose to shop at one of two identical grocery stores. In a world ruled by the Central Limit Theorem, a consumer would choose a store with probability one-half. Thus, each store could expect to see fifty customers on average, though the actual number that showed up would be subject to random variation. In fact, given the underlying process just described, we know that a given store will have, say, more than sixty customers only about 2 percent of the time.

Now, allow customers to act more purposefully and interact with one another. Suppose that customers prefer to be in less-crowded stores. Such

an assumption introduces a feedback into the system, whereby customers who find themselves in the crowded store begin to shop at the other store. To avoid some odd system-level behavior, we allow only a single customer per period to make such a decision.⁴ Given this assumption, in very short order the number of shoppers in each store equilibrates at fifty. Even if we include small external shocks to the system, for example, two customers from different stores take a liking to one another and begin to shop together, the system as a whole will quickly resettle back to the stable configuration with exactly fifty people in each store. Thus, the desire to avoid crowding by each individual induces a negative feedback on the system, resulting in a very stable and predictable outcome.

Agent interactions can also introduce positive feedback into the system. Suppose the same group of one hundred people, also does some banking each day. Imagine that each person has some chance, say 50 percent, of going to the bank and withdrawing money. The bank has limited reserves to cover withdrawals, and thus if too many people withdraw their money, the bank will be unable to cover the claims and become insolvent, causing depositors to panic and demand their money. If the bank has 60 percent reserves, then, as we saw earlier, around 2 percent of the time the bank will go insolvent, resulting in an unfortunate “large event” and an all-out run on the bank.

In our three worlds we see very different behavior. In the first, customers act independently and ignore one another, so the resulting number of customers shopping at a given store is nicely approximated by a normal distribution with a mean of fifty and standard deviation of five. In the second, where customers avoid crowding, we get a degenerate distribution with each store having exactly fifty customers each day. Finally, with the potential for panic, the number of customers arriving at the bank looks identical to the normal distribution we saw in the first case when we have less than sixty customers, but once we hit sixty, all of the remaining weight of the distribution shifts to the right, and we get a fat upper tail.

The contrasts between these images are startling. The world would be a lot easier to understand if we could restrict our attention to the first two scenarios, that is, if agents either avoid direct interaction with one another or interact in such a way that strong negative feedback results in a stable equilibrium. Alas, the vast majority of social science theory focuses on exactly these two types of outcomes. Nonetheless, there are many canonical examples of “large events” that arise in social

³One only needs to look at the failure of Long-Term Capital Management in 1998 to realize the practical importance of this distinction. The world in which Long-Term Capital Management played was one governed by fat-tailed distributions, not the Central Limit Theorem.

⁴Without this, there are some dynamics where we can get large swings in the number of customers as they overreact to crowding.




Figure 4.4. Gliders in the Game of Life. A glider in the Game of Life is a configuration of live cells that “moves” across the space. During each successive time step (*left to right*), the set of live cells is altered based on the simple, local rules (see text) of the game. After four time steps, the original configuration of live cells is regenerated, only displaced down and to the right by one cell. If left undisturbed, this structure will continue to “glide” across the space. A more elaborate configuration of live cells, known as a glider gun, is capable of generating such gliders.

systems, such as stock market crashes, riots, outbreaks of war and peace, political movements, and traffic jams. These events are driven by positive feedback, arising from perhaps externalities driven by the behavior of others that change each individual’s costs or benefits from acting (for example, as rioting breaks out, your chance of going to jail decreases, and the social benefit of joining in increases) or physical constraints on behavior (such as when the car in front of you on the highway slows down, forcing you to slow down as well to avoid a crash).

Thinking about positive and negative feedback provides only a crude window into the set of possibilities that can emerge in a complex social system. Many complex systems contain both types of feedback. For example, consider Conway’s Game of Life. In this game, the world moves in lockstep and is arrayed on a two-dimensional grid, each cell of which can either be dead or alive. A dead cell with exactly three live neighbors is “born” and becomes a live cell next period; otherwise, it remains dead. A live cell with two or three live neighbors “survives” into the next period; otherwise, it dies (either out of “loneliness” or “overcrowding”). Thus, in this system an intermediate amount of life begets life (a positive feedback), while too much or too little life leads to death (a negative feedback). Ultimately, this results in a remarkable set of global patterns in both space and time that can emerge from this simple set of microlevel rules. These patterns are so coherent at times that we can ignore the underlying microlevel rules that generated them and instead rely on the resulting global structures to predict systemwide behavior (see, for example, figure 4.4).

As discussed previously, we have access to some useful “emergence” theorems for systems that display disorganized complexity. However, to fully understand emergence, we need to go beyond these disorganized systems with their interrelated, helter-skelter agents and begin to develop theories for those systems that entail organized complexity. Under organized complexity, the relationships among the agents are such that through various feedbacks and structural contingencies, agent variations no longer cancel one another out but, rather, become reinforcing. In such a world, we leave the realm of the Law of Large Numbers and instead embark down paths unknown. While we have ample evidence, both empirical and experimental, that under organized complexity, systems can exhibit aggregate properties that are not directly tied to agent details, a sound theoretical foothold from which to leverage this observation is only now being constructed.