chapter, on external intervention and introduction of explicit barriers.

These remarks on the mixture of internal and external processes in nation building serve as a reminder that actual situations typically involve many mechanisms at once. Variety may be created by imitating foreign visitors even as it is destroyed by censoring media. Interaction may be decreased by a policy decision to reduce foreign language instruction, at the same time as interaction is increased by the need to trade with each other. An idea can be widely adopted in spite of being publicly condemned—indeed, *because* of being condemned, so that publishers will be eager to print a book "banned in Boston."

Having reviewed this collection of mechanisms for changing interaction patterns, we can now focus on the third major aspect of Complex Adaptive Systems, their mechanisms of selection. Interaction among agents shapes the creation and destruction of variety and produces the events that drive the attribution of credit. Now we can examine how selection itself works and see how it feeds back onto variety and interaction.

IV

Selection

In previous chapters, we have looked at the mechanisms that create and destroy types, and at the processes and structures that govern interaction among types. In this chapter, we turn to the fundamental question of which agents or strategies should be copied and which should be destroyed. In other words, how should selection be employed to promote adaptation?

Natural selection in evolutionary biology provides a familiar and well-studied example of how selection can work. Although selection in a Complex Adaptive System need not operate in the same way as natural selection, evolutionary biology is a good place to start our analysis. Evolution by natural selection requires three things. First, it requires a means to retain the essential character of the agent. In biological systems, genetic material preserves the key patterns. Evolution by natural selection also requires a source of variation. In the simplest biological systems, this can be achieved by mutation. In sexual reproduction, novelty is generated through recombination of characteristics from different parents, as well as by mutation. Finally, evolution requires amplification, changes in the frequencies of types. In biological systems this is the result of some individuals having many offspring while others have few or none.

If you want to design a system that is able to explore new possibilities while being able to exploit what has already been achieved, biological evolution provides an important benchmark. It demonstrates that adaptation can be achieved even without the agents (or anyone else) having any understanding of how the system works.

While natural selection provides an important paradigm for how an adaptive system can work, it also has some serious disadvantages compared with more directed methods of achieving adaptation. Whenever it is feasible to attribute success to something more specific than the entire agent, there is the possibility of selecting strategies rather than whole agents. If you find that quinine-related compounds reduce malaria, you can spread them through the world instead of waiting many generations for natural selection to breed malaria-resistant humans. This is especially valuable since the main antimalarial solution nature has so far evolved makes the carrier susceptible to sickle-cell disease, itself a debilitating condition. When attribution is sufficiently precise—and this can be far from perfectly accurate—it can pay handsomely to make numerous copies of a good strategy on a fast time scale that would be impossible if complete agents had to be reproduced.

These two approaches, selecting at the level of entire agents and selecting at the level of strategies, share the need to make copies that retain effective adaptations, to incorporate variation for further adaptation, and to amplify the success (and cull the failure) that does occur. But they differ in the level at which they operate—and selection at the two levels can work *very* differently. Selection of one advertising agency from a population of competing firms can have quite different dynamics from selecting among a population of advertising themes proposed by a single agency. Nonetheless, whether it is whole agents or strategies that are evaluated and undergo reproduction, a design for an adaptive system of selection must deal with four issues:

1. Defining criteria of success.
2. Determining whether selection is at the level of agents or strategies.

3. Attributing credit for success and failure.
4. Creating new agents or strategies.

This chapter will consider each of these elements of a selection process in turn. While these elements do not separate neatly in the everyday world, distinguishing them will help simplify our discussion without introducing too much distortion.

The chapter includes two extended examples that show how principles for harnessing complexity can be applied. The first deals with improving the criteria for success used by agents. It explores prize competitions as a mechanism that changes success criteria in a field by identifying and rewarding exemplary individuals or activities. The second example suggests approaches to attributing credit for success and failure when there is only a limited amount of relevant experience available. It explores military use of simulation as a method of generating surrogate experience that can be helpful in accelerating the attribution of credit to strategies that might be successful. The chapter concludes with a discussion of how leaders can use their visibility to help shape the selection of effective agents and strategies.

Defining riteria of Success

The importance of knowing what to count as success is the point of an Army story about the new draftee who was an operations analyst in civilian life. After standing with fellow draftees in a long line to get their dinner plates washed and rinsed, the recruit went up to the old sergeant and explained it is inefficient to use two vats for washing dishes and two vats for rinsing them. It would be faster to use three vats for washing and only one for rinsing since washing takes more time than rinsing. The old sergeant looked with disdain at the new recruit, and said, "You've got it exactly backwards. I want them to stand just as *much* as possible. I can't keep them running around all day, but the longer I can keep them on their feet, the better."

Clearly, selection of agents or strategies implies some metric of success. Agents need not attend to the measure. Animals can have many offspring out of motives far more compelling than the eventual adaptation of their species. Fashions may be copied without much big-picture reflection by those adopting a new style. In such cases, success is actually defined by outside observers as "frequently copied." Rather than specifying a success measure and copying what scores well, this approach measures success by numbers of copies. Biologists take this line when they assess fitness as number of offspring. In biology, survival defines what is fittest.

However, in most of the situations we consider, performance measures are active in the minds of designers, policy makers, and other actors, whether they are acting inside the system or contemplating it from the outside. Recall the example of Linux software development, with its thousands of volunteers proposing solutions to specific problems in a massive operating system. Being able to evaluate the effectiveness of proposed solutions using clear measures such as speed and crash-avoidance was one of the requirements for such open software development to work. Typically, however, the assessment of alternatives in a Complex Adaptive System is not easy. In fact, there is usually more than one criterion that could be used to assess results.

For a business, profit seems a natural measure of success. For a checkers player, winning games is a natural performance measure. Yet even in these examples, with success criteria that seem indisputable, complexity might be harnessed more effectively if other measures of success are used. In the business example, market share provides an additional measure that can be a useful supplement to profits. One reason is that changes in profits may reflect factors beyond the control of the company, such as improvement in the national economy. You might not want to attribute credit for increased profits to your new marketing campaign if you knew that your entire industry had prospered during a buoyant economy. An increase in your market share could provide a better indication than profits of whether you were doing something right—and what it was. We will also see below for our checkers example that there are measures of success that may be more effective than waiting for the outcome of the game.

Our approach to harnessing complexity does not take any performance measure as "given." It does not anoint any one measure as a highest goal. Performance measures can be seen as instruments that shape what events are likely to occur. Even the preservation of life is not a goal that trumps all others, as human willingness to die for principles so dramatically reveals. Since goals are not seen as fixed, setting goals, the criteria that govern processes of selection, is one of the main interventions for those who would harness complexity. Our view leads to two important and uncommon observations about performance measures.

First, it is valuable to appreciate that *performance measures are defined within the system.* They are modified (or maintained) and applied (or disregarded) by the agents themselves. This observation is not a surprise to many experienced practitioners, who are well aware of the political work that lies behind measures later taken as givens. Unfortunately, many efforts to apply complexity concepts to social systems give little attention to how performance measures are defined within the system. To see what we mean, consider the case of profit as one such measure. What may count as a profit depends on many factors, including what the law allows individuals to own, what social norms and religions define as morally fair, whether actual practices conform to those norms, what the tax code recognizes as legitimate costs, and whether society charges for disposal of the by-products of activity, such as used motor oil or even carbon dioxide. We also regulate the scope of profit as a permissible goal. We largely removed profit from the decision making within American schools, hospitals, and prisons at the beginning of the twentieth century and are experimenting now with reintroducing it.

A further consequence of performance measures being defined by the agents themselves is that there can be more than one measure active. In addition, the measures may be inconsistent and may change over time. Change that is seen as improvement by one type of agent may be seen as a loss by others. There are issues of variety in performance measures just as there are in other characteristics of agents and their strategies. When members of an organization assess a situation from different evaluative angles, they generate a greater variety of new possibilities that, if not excessive, can have great value for the organization (Cohen, 1984). But it is clear that beyond some level, variety in performance measures can also be a source of debilitating inconsistency and conflict.

Second, *how success is defined affects the chances for effective learning*. To return to our example of checkers, consider the difficulties for learning if victory is the sole criterion of success. The central problem is that victory or defeat comes only once per game. However, getting more than one measurement of performance per game could dramatically improve the rate of adaptation. The typical way to do this is to use criteria that can be measured in the course of the game. In checkers or chess, this is possible by evaluating the current board to see who is ahead in pieces and in various aspects of position. Such evaluations allow intelligent choices in the midst of the game based on what promises to lead to a better board position in a few moves. This doesn't require seeing all the way to victory or defeat at the end of the game. Since you cannot precisely measure the consequences of early moves for victory, you introduce other metrics that are more easily predicted. In a seeming paradox, you increase the chance of winning by concentrating on a set of criteria that does not include winning.

Even better, with finer-grained measures you can actually learn to improve the criteria by which you evaluate board positions. For example, you might learn from experience that having many pieces in the center often leads to surprisingly good results a few moves later. Indeed, the Samuel checker-playing program, one of the early triumphs of artificial intelligence research, learns on its own to play better checkers by using expected results in just this way (Samuel,

1959). When it arrives at a board position that is surprisingly good or bad, it uses this information to revise its own success criteria. The program determines what changes in its evaluators would have avoided the surprise and makes the corresponding changes. When it next encounters a similar board, the program will have a better set of criteria for attributing value to board positions. This approach to learning new success criteria is very powerful. Samuel's program, running on an early computer that could not keep up with today's digital wristwatches, could learn checkers well enough to defeat a state champion. Moreover, these are techniques of very broad applicability:

- When success is measurable only rarely, new measures with a faster tempo can speed learning, even if they do not perfectly reflect the longer-term goal.
- Whenever outcomes are better or worse than expected, the experience can help to revise evaluation criteria so that, in the future, the attribution of credit will produce better outcomes (Cohen and Axelrod, 1984).

Using fine-grained and short-term measures of success can help individual learning by providing focused and rapid feedback. Such narrow and prompt measures of success can also be used by an organization to evaluate who is successful and who is not. For this reason, managers are often judged by how well their unit does each quarter, or even each month, or by very specific indicators such as cost reductions. But there is a problem. If the challenges the manager is dealing with are long-term or widespread in the organization, then using fine-grained and prompt measures of success can easily miss much of the value to the organization of any improvement the manager discovers. As we saw in Chapter II, on variation, there can be a lot of bang for the exploratory buck when advances in one domain can be applied for a long time and/or in many places. A challenge for an organization is to develop measures of success that support appropriate levels of exploratory behavior while taking into account that learning is fostered by fine-grained and rapid feedback.

Another challenge in defining measures that will support learning is that a measure may be correlated with what ultimately matters without actually being causally related. A medical example is the reduction of fever as a measure of success in fighting a disease. A fever indicates the presence of a disease, and the fever disappears when the disease does. But with the development of aspirin, one can reduce the fever without curing the disease. Therefore, using body temperature as a measure of success can be misleading for some diseases. Because the elevated temperature might even be part of the body's method of fighting the disease (Nesse and Williams, 1998), parents may learn to treat the fever with aspirin in ways that can actually be harmful. The implication is that one needs to be careful about which indirect measures of success are used to guide action and learning.

Taken together, these observations about success measures imply not only risks but also rich possibilities for harnessing complexity through shaping the criteria by which the agents or their activities are evaluated. Performance measures are not immutably given, but are subject to change, both from the outside and from within the system where they operate. What measures are used profoundly affects which agents and strategies will be copied and recombined and, therefore, what adaptation will occur. This is the logic that gives long-term power to what may seem modest changes in measures, such as introducing on-time performance into airline regulation, body counts into battle assessments (Gartner, 1997), "pawn structure" into chess (Euwe, 1968), and portfolio risk into financial management (Sharpe et al., 1998).

Example: Prize Competitions

A particularly illuminating example of changing success criteria is the method of establishing a prize competition. Consider, for example, the ancient Athenian practice of conducting annual

dramatic contests (Pickard-Cambridge, 1968). By explicitly declaring which drama was the best, the award accomplished three things. First, the author was honored for success, bringing fame and influence to individuals such as Aeschylus and Sophocles. Second, the award encouraged the production of new plays composed to meet the criteria implied by the previous awards. In our terms, the strategies of later playwriting were changed. Third, the award helped educate and shape the tastes of the audience, thereby providing future support for the criteria of excellence the award implied.

Today prize competitions are used to reward, encourage, and define excellence in a wide range of activities, from grammar school art contests to the Nobel Prizes in physics, peace, and literature. There are now prizes for beauty, for most valuable player, for best dressed, and for business quality. The effectiveness of prizes is enhanced as society develops more extensive channels to disseminate news of awards. So we should not be surprised that their use is increasing. Every increment in the reach of printing, television, or e-mail newsgroups increases the possibilities for affecting success criteria by announcing winners of awards.

Some prizes are for accomplishments that can be assessed more or less objectively, such as the winner of a solar-powered car race. For our purposes, the most interesting prizes are those that are based on subjective criteria. Indeed, for many prizes, the criteria are so indefinite that the burden of defining excellence within some realm falls heavily, if not entirely, on the subjective evaluations of a panel of judges.

From the point of view of harnessing complexity, a major advantage of prize competitions is that they can award credit to people or activities based on criteria that are different from current standards. The presumption is that a carefully selected panel of judges can make worthwhile evaluations of quality. The indirect effects are as powerful as the direct effects. Giving a prize not only rewards a winner who might not have excelled in other assessments but also provides a target for others to emulate. Emulation may take the form of superficial imitation, but it may also create innovative exemplars of just what was most valued by experts. In addition, by helping to shape the tastes of the general audience, a prize competition can also shape the criteria used by the broader public. For example, book awards not only provide guidance to writers and publishers about what is being valued but also provide guidance to readers and reviewers about what is worth reading. The promotion of a sophisticated reading public, in turn, helps provide a market for good writing.

A prize competition can also promote useful variety. Prizes sometimes serve to identify and promote things that are new and valuable. When a science or literary prize is awarded, it tends to legitimate and promote the entire field or genre of the winner. Of course, there is a tension here. Deciding who or what should receive an award involves the application of standards of excellence. The judges inevitably use standards that are shaped in part by the standards in the broader community of which they are a part. Indeed, judges are usually selected on the basis of their own standing, which in turn is often based on their adherence to current standards. And even if the judges may wish to be leaders in the identification of what is both new and worthy and are willing to take a risk on something that stretches current standards, they also need to be concerned about looking arbitrary or even foolish. The judges are also judged. Therefore, they face the familiar trade-off between exploitation and exploration in making their selections. The trade-off creates a tension between making a safe choice that reflects current standards and making a bold choice that can help transform those very standards.

To the extent that prize committees are willing to go beyond the orthodoxy of the moment, they represent a valuable potential for increasing useful variety. This potential is not al-

ways fully seized. Of the first 85 winners of the Nobel Prize for literature, all but one wrote in a European language (Espmark, 1986).

Prizes can even stifle variety. It is now very hard for a young pianist to establish a successful recording or concert career without having won one of the major competitions. The reason is that producers rely on the competitions to screen pianists. Young pianists therefore train to win these competitions, go to teachers who have won or whose students have won, choose repertoire suited to winning, and so on. Thus there is some truth to the criticism that competitions can reduce the variety of piano expression exactly because the competitions can become the dominant focus for young players. It can take a long time for the weak signals of public taste or music reviews to counter the now strengthened signals of prize jury standards.

While each prize sets up a competition among those aspiring to win it, there is also competition among the prizes themselves. The sponsors and judges of each prize seek attention and prestige for their award. Within each domain there is competition for how much credit will be garnered by the winners of a particular award. Is a Pulitzer Prize for fiction better than a National Book Award? Prize competitions themselves interact, as when getting one prize makes a winner more likely to get another prize. Moreover, a lesser-known prize can gain prestige if its winners often go on to receive some better-known prize. Thus there is an intricate set of interactions within and between four populations of agents: prize seekers, members of their audience, judges on awards panels, and the various prize competitions themselves. Together they function to alter the criteria that define success in their respective domains.

Determining the Level of Selection

Two basic processes amplify success: selection of agents and selection of strategies. The natural selection of biological agents works by making an entirely new agent without the need to determine the cause of the success of the parent or parents. The selection of strategies, on the other hand, creates new strategies for an existing agent. It often involves some explicit decision about what strategy or part of the agent was responsible for the success.

Selection of Agents

Biological systems are not the only ones that select entire agents. Elections are another such method. If a congressional representative is defeated in an election, another person gets the job. The voters are not able to pick and choose among the features they like in the incumbent and a challenger. They simply have to pick one candidate or the other. This provides an easy answer to the question of what should be given credit for success (or failure). The answer is the whole candidate. As much as a voter might want to give credit and blame separately for some good and bad policy positions or character traits, the vote requires selection at the level of the whole agent. One agent will occupy the office for the coming term; all others will be cast aside.

Elections offer a nice example of several coevolving complex systems. While voters are selecting at the level of agents, active politicians are selecting at the level of strategies. They observe carefully what positions were taken by recent victors around the country. Many will adopt those more successful strategies in future campaigns.

The economy also can select at the level of agents. Companies that go bankrupt and are liquidated are thereafter not present in the

population. On the constructive side, imagine a decentralized firm that has a highly successful branch office. The firm might use its earnings to "clone" the successful branch office by setting up another branch that, insofar as possible, duplicates the entire operation of the successful one. If the branches operate fairly autonomously, this would amount to creating a new agent. The central office would have given credit to an entire branch (rather than to any of its particular strategies or characteristics) and tried to amplify success by producing a duplicate agent.

As we noted, biological evolution works by selecting agents. The success of an organism leads to reproduction. This does not entail any determination of which of the genes "deserve" credit for the reproductive success. Instead, all the genes in the reproducing organism get a roughly equal chance to be passed on to the offspring. This fact is the root of the phenomenon called hitchhiking, in which nonproductive, even mildly deleterious, genes are carried into subsequent generations by the success of the overall agent package to which they belong (Maynard Smith, 1978).

In all these examples of agent selection, there must be fairly substantial accumulations of resources to create a new agent, whether that agent is an infant organism, a political candidate, or a branch office. The need to accumulate sufficient resources to embody a new agent operates as an important limiting factor in agent-level selection. It contrasts with the situation we will see below for strategy-level selection, where what is copied can often be merely the abstract pattern of the strategy. The extreme example of this, a process that is profoundly reshaping our era, is the copying of computer algorithms. Here, the marginal costs of assembling a new copy may hover just above zero, allowing low-cost software to run on millions of computers.

When using selection of agents to harness complexity, a key question is how strong the **selection pressure** should be. If the best agent in a population gets many copies while the others get few or none, the selection pressure is very high. In effect, strong selection pressure greatly amplifies the success of the best agent in the population but gives very little amplification to the slightly less successful. In an era where franchising can provide strong selection pressure, the best ideas for a hardware store or a bookstore will be extensively copied, while independent competitors will languish. Conversely, weak selection pressure produces only a slight tendency for the better agents to have more copies and thus provides more uniform amplification to the relatively successful agents. The advantage of strong selection pressure is that it exploits success by quickly spreading copies of the best-performing agents. The disadvantage is that it can quickly destroy the variety in the population that is needed to explore for even better outcomes in the future. Thus the trade-off between strong and weak selection pressure raises the familiar issues of choosing the balance between exploiting the best current outcomes and using variety to explore for possible future improvements.

Managers and designers often have opportunities to change selection pressure. Among other things, they can increase rewards and visibility for top performers and set severe punishments for flaws. For example, "zero tolerance" deletion of agents or artifacts with small deficiencies has the effect of reducing variety. It thereby favors exploitation over exploration. In the short run, strong selection pressure converts existing variety to new exploitation, but in the long run exploration may suffer. Harnessing complexity requires taking advantage of variety rather than trying to ignore or eliminate it.

An instructive issue in biological reproduction is the **founder effect.** An example would be an island populated by long-beaked birds descended from a long-beaked pair that were among the first to reach the locale. In its early history, the population is small, and an outstandingly fit individual has offspring that form a large portion of the next generation. Over subsequent generations, many traits of that "founder" are carried widely through the population. Whether or not they make their own functional contribution, the traits that made the founder effective co-occurred with traits that do not have high value.

Both kinds are amplified. A nonbiological example can be seen in the Carnegie libraries that proliferated in the United States in the early twentieth century. Many different communities established libraries starting from the same plans. Overall, the favored plans were good ones and carried financial subsidies. The practice of using them was beneficial on the whole but did result in libraries with specific services that were arbitrary or even unwanted in some communities in which they were instituted (Van Slyck, 1995).

Selection of Strategies

An alternative to selecting entire agents as the basis for the amplification of success is to make copies or recombinations at the level of particular strategies. In this section, we examine three common points of difference in selection at the agent and strategy levels: cost, waiting time, and difficulty of inference. We show how the different strengths of the two levels of selection are sometimes complementary, and that there can be substantial advantages to a hybrid system in which selection goes on simultaneously at both levels.

If success can be assessed at the strategy level rather than the agent level, one difference that often occurs is a lowered *cost* of copying. To assemble or acquire a whole new agent (a new person, a new business, a new governmental unit) is typically more costly than to copy a strategy employed by a successful agent. It takes years to grow several Pacific yew trees for bark that provides cancer-fighting compounds for a single patient. A laboratory synthesis of the active chemical makes it available quickly to many thousands of patients. An owner of a baseball team can try to buy a star pitcher from another team. If the reason for success is that the pitcher is winning by throwing the forkball, it might be cheaper to teach the other pitchers that strategy during the off-season. Whether this will be promising or not depends on how easily the forkball can be copied. Is there a pitching coach for hire with success in teaching it? Or perhaps success depends on the uncanny similarity of the star's forkball and fastball motions. Then it may be necessary to pay the cost of acquiring the whole agent, with the entire complement of strategies, or of searching for another pitcher with a comparable package of skills.

A second difference that often occurs between the strategy and agent levels is *waiting time.* One could just think of this as a special case of higher costs, but it deserves a brief discussion of its own. Because assembling copies of agents is generally a larger task involving more resources, it typically takes more time than copying or recombining strategies. Even if the direct costs of agent copying were affordable, the indirect costs of delay might not be. For example, another company may have a proprietary process for manufacturing a part that goes into a product you are developing. It might be quite valuable to invent your own process for making the needed component, and plausible to create a division within your company to do it. It would lower your costs and let you tailor the part to your particular needs. But competitors are racing to the market for your own product. The delay while you create a capacity to make the part means falling behind in competition with them. So you license the existing process from its owner, copying that strategy not because of lower monetary costs but because of the value of elapsed time.

To highlight the speed at which strategies can change, consider a stock market. Agents watch changes in prices for information about what other agents believe. Thus the market has a recursive nature in which agents' expectations are formed on the basis of their anticipation of other agents' expectations. The result can be rapid bubbles and crashes. Simulations of markets as Complex Adaptive Systems demonstrate how high rates of exploration can generate these bandwagon effects and "market psychology" (Arthur et al., 1997).

Social mobilization is another arena in which agents' expectations are formed by watching each other's behavior. Again, the result can be very rapid change once a bandwagon begins. The fall of the Berlin Wall occurred with amazing speed once the initial demonstrations showed what was possible. As in a market, people formed their expectations on the basis of their anticipation of others' expectations. Once begun, a series of demonstrations set off a cascade of revised beliefs leading to irresistible levels of protest (Lohmann, 1994).

Markets and demonstrations illustrate how strategies can be selected very quickly. Typically, selection at the strategy level is faster and less costly than selection at the agent level. Nevertheless, these differences are tendencies rather than inevitable consequences. So, by way of counterexample, large corporations are often faced with new products from start-up competitors. They sometimes find it quicker to create new divisions or small spin-off firms to make a comparable product rather than modifying existing lines of activity to produce it. In effect, this is a case where agent creation may be faster than strategy copying. Although the differences we have mentioned are only tendencies, they are rooted in the added difficulty that is typical for creating full agents. Hence it is often important to compare possibilities for selection at the agent and strategy levels.

A third difference between selection at the two levels involves problems that commonly occur in inferring exactly what is to be copied. There are myriad ways that selection can go awry and incorrectly reward an agent or strategy that was not responsible for a success. Such failures plague selection at both levels. However, one important difference does occur. Agents are collections of strategies. Successful agents generally use strategies that are mutually compatible. The interaction among those strategies does not have to be understood if selection is at the level of the agent, copying all its strategies. Biological selection of whole agents capitalizes mightily on this fact, but so can identical replication of franchised business units. Selection at the strategy level generally demands higher quality of inference. How many of the agent's action patterns must be copied to replicate the success? Which ones? To obtain the same low defect rate as a rival firm, which of their quality control procedures should be emulated? Selection at the agent level, on average, is more **context preserving** than at the strategy level. In a Complex

Adaptive System, where many results derive from effects that multiply other effects, context preservation can work to retain and spread synergies that are not fully understood. We made a related point in the previous chapter when we observed that the longer time horizons of those in authority create a common context for coordinating the faster actions of those they supervise. There we were examining agents' activities. Here we are examining the selection that follows from their success or failure.

We have argued that there are tendencies for selection at the agent level to be more costly, slower in elapsed time, and more context preserving. The first two effects are often not wanted, while the last one frequently is. This can set up a tension in which a designer or policy maker who has some freedom to influence the level of selection may have to trade off the various factors.

To take an example: Suppose that you want to discourage a dangerous behavior such as violating crucial safety regulations. We have usually considered selection for positive traits, but here we can look at negative selection. At the level of *strategy,* selection may correspond to punishing the action pattern. Each detected instance of rule breaking could be heavily fined, for example. On the other hand, *agents* could be negatively selected in response to their violations. An offending employee could be suspended, transferred, or even fired. These forms of removal will make the agent less likely to be copied. Taking the agent out of circulation and making the effort of replacement typically costs more and takes longer than simply changing an agent's strategy. If the safety violation is integrated with other strategies—for example if the agent's entire work style used a set of methods now considered unsafe—simply punishing the violations may not discourage the behavior, so removal may be worthwhile. If the violations are more a matter of "fashion"—for example, not wearing a hard hat in order to look fearless—punishing the action itself may be the preferred approach.

Schemes to amplify success are nearly always imperfect. Selection at the level of agent and selection at the level of strategy are families of mechanisms that have somewhat complementary strengths. Agent selection often works on longer time scales—faster is not always better—and preserves variation and context. Strategy selection isolates key patterns that can be more easily and rapidly copied.

Thus it is not surprising that there are many hybrid systems, where selection is found to be operating at both levels in a single population of agents. Many species of birds and mammals seem to select at both the agent level, by conventional natural selection, and at the strategy level, by processes of cultural diffusion, which operate at a much faster time scale (Lumsden and Wilson, 1981; Cavalli-Sforza and Feldman, 1981). In the human case, cultural evolution is so rapid and effective that we tend to ignore the continuing operation of natural selection. At the other extreme, we often do not notice cultural aspects of an animal population. But close observation reveals striking cases, such as the English birds that discovered how to peck through foil milk bottle caps. Their discovery spread across the entire country within a few years (Hinde and Fisher, 1951).

Hybrid systems such as this have tremendous advantages. Herbert Simon has argued that they are so beneficial that we could expect biological evolution to create individuals with increased susceptibility to following strategies suggested by others (Simon, 1990). Even though this "docile" quality makes it possible to take advantage of individuals who possess it, that can be outweighed by the tremendous gains of adding cultural selection of strategies to natural selection of agents.

These observations on complementary strengths and hybrid selection systems have a cumulative implication. When there is room to alter selection processes, it can be wise to look for changes to the system that could diversify it, adding fast elements if its selection processes are slow. If the fast processes are not succeeding, it can pay to add slower elements that sustain a new context. As an example of adding fast elements, organizations that rely heavily on change through personnel turnover are often ripe for improved trading of employee "war stories." A series of failures in piecemeal importing of "best practices" might suggest bringing in a new supervisor experienced in how the various routines form an interlocking system. As with many other interventions we discuss, hybridizing selection processes is not guaranteed to be better, but it is often a beneficial focus of attention.

Attributing Credit for Success and Failure

At the beginning of this chapter we outlined four interdependent aspects of selection. So far we have examined two of them, how success criteria are defined and whether selection operates at the agent or strategy level. To complete the picture we must consider two more issues:

- how an agent uses a performance criterion to increase the frequency of successful strategies or decrease the frequency of unsuccessful ones, a step we call **attribution of credit,** and
- how agents or strategies that receive credit are copied, recombined, or destroyed.

Credit attribution, though difficult and necessarily imperfect, can nonetheless be designed to help harness complexity. In the preceding section, we pointed out that context preservation could be advantageous if the cause of apparent success is not fully understood. This indicates a general problem. Since Complex Adaptive Systems are inherently difficult to understand or predict, it follows that attribution of credit in selection will often be difficult and prone to mistakes. If it were feasible, the best response would be not to make mistakes in credit attribution. Because such mistakes can be very costly, vast bodies of academic knowledge and expensive social apparatus have been created to reduce them. Systems of logic,

methods of statistics, and philosophies of science are all aimed at improving the extent to which our conclusions follow from our premises and evidence. There are public debates, professional review boards, and courts of law. All contribute to limiting the mistakes in attribution of credit that may drive selection processes.

Where these tools for improving inference are cost-effective, we certainly believe they should be used, and we applaud the work that maintains and extends them. However, despite all the effort put into these valuable resources, totally accurate attribution of credit is often infeasible. To see why, consider this list from a survey of research on the limits of rationality (Conlisk, 1996). It details the factors that make it easy to learn appropriate lessons from the experience accumulated in making a series of choices:

- clear rewards for the appropriate choices,
- repeated opportunities for observation or for practice,
- small deliberation costs at each choice so that frequent choices are easier,
- good feedback on the results of choices,
- unchanging circumstances that keep inferences valid, and
- a simple context that can be effectively analyzed.

The contrast of this list with the properties we have seen in Complex Adaptive Systems is stark. In complex systems, it is difficult to determine what should be rewarded or which choice is appropriate. Measurement of success is often infrequent, and shifting context makes few observations comparable. Deliberation costs for choices can be high, especially if they require the apparatus of formal logic or statistics, or social processes of choice such as scientific peer review—not to speak of court proceedings. Feedback is ambiguous. Circumstances, even goals, are changing. All of this follows from the fundamental premise: we are coping with systems that are complex and adaptive, not simple or static. In the short run we are not likely to have a direct approach that "gets it completely right." We will need as well the indirect methods of harnessing complexity.

The difficulty of attributing credit in real experience can be reinforced by considering a few examples. The war in Vietnam provides a striking case. Although war usually produces large rewards (and punishments) and, in the end, provides clear feedback on the result, none of the other circumstances for effective learning obtains. For the Americans, the Vietnam War was not a victory. But exactly what lessons should have been learned from it? There are many contending lessons and no obvious way to determine which candidates are most appropriate. Despite these impediments, lessons were learned by the American military. These included the need for decisive force in any future war, the need to avoid slow escalation, and the need to avoid civilian interference in the conduct of the war (Powell, 1995). These lessons—"strategies," in our terms—were applied to the planning and conduct of the Gulf War and seemed to be effective in that application.

On the other hand, for the Soviet Union the Vietnam War was a success. The lessons drawn from the war by the Soviet Union emphasized that their Vietnamese allies won because of their great will and courage, assisted by military aid from Communist nations (Zimmerman and Axelrod, 1981). These optimistic lessons would not have warned the Soviets about the dangers of their later intervention in Afghanistan.

Biological systems also face difficulties in attributing credit. Consider birds, which determine from experience the visual characteristics of dangerous predators. Their situation illustrates one of the many interesting complications of credit attribution in a Complex Adaptive System: exploitation by others. The method is **mimicry**—as when many species of moths evolve spots on their wings that resemble the eyes of larger predators. This works because the birds develop a "prediction" of danger from appearance and rely on it to avoid predators. One presumes that the birds' capability to associate certain appearances with danger, which is a mechanism for attributing credit, serves the birds well overall. But the moths can also exploit the birds' imperfect credit attribution to avoid being eaten. Once again coevolution increases complexity and inhibits prediction.

For a nonbiological example of the limits of credit attribution, consider the person who ends the year with the highest sales volume, receives a significant bonus, and is singled out to be emulated. Years later, more careful cost accounting may show that most of the sales actually lost money for the firm because of eventual refunds or support costs. The business literature is rife with stories of performance indicators that failed to capture important aspects of a complex setting. These misattributions may occur because of causal connections that no one understands, or because some employees, like the spotted moths, come to mimic features that other employees, like the birds, have come to associate with success or failure.

The difficulties of credit attribution are endemic in Complex Adaptive Systems. Our aim in this section is not to escape them, though we recommend that when it is feasible. Instead, our aim is to suggest how the side effects of inevitable mistakes of attribution can be turned to some advantage. Each of these three categories is constructed as a composite of actual cases in which complexity makes some mistakes of attribution inevitable. They illustrate three different problems of inference that are highly characteristic of credit attribution in complex systems:

- the mistake of crediting or blaming a part when a larger ensemble is responsible,
- the mistake of attributing credit or blame to a particular ensemble of factors when in fact a different ensemble is responsible, and
- the mistake of crediting a misconstrued strategy, where the action involved produced success, but the conditions in which the action should be taken have been misunderstood.

The first type of mistake, crediting a part when a larger ensemble is responsible, is very common in Complex Adaptive Systems since they so often involve a number of entangled causal factors. It is easy to notice that a single agent or strategy is associated with a series of successes (or failures). If you are not positioned to observe the operation of other necessary forces, you reach an incorrect conclusion that it alone causes the results. Consider a manager of a department that uses project teams assembled for specific tasks. If it is the practice in the unit to reward team members whose work contributed to notable success, a manager can almost be sure that there will be some occasions where an individual receives credit for what was produced by the interplay among contributions of several team members—what is sometimes labeled the group's "chemistry."

We have stressed insufficient exploration in examples throughout this book because we so often have seen variation being undervalued by managers of Complex Adaptive Systems. But for this case, let's stipulate that the manager believes the department has a problem of insufficient exploitation. Perhaps "back channel" communication suggests that a project group has done well by ensemble effect rather than the efforts of the most prominent individual. How can the manager get "mileage" out of discovering those attribution mistakes without knowing what caused them? One approach is to make a special effort to reassemble that identical team for a later problem, retaining (and exploiting) the uncredited ensemble that may be there. Doing this has a cost, of course. It reduces the ability to mix and match individuals to the characteristics of the next task. Harnessing complexity does not always come for free.

The second type of mistake, attribution of credit to the wrong set of factors, is often made in Complex Adaptive Systems for much the same reasons. Diagnosis of causes in complex, multicausal situations is error prone. We might take as an example the problem of examining customer complaints about product malfunction in order to discover product defects or possible design improvements. Many large consumer product companies have service desks that answer thousands of calls per week about products. They frequently have systems that generate "trouble tickets" associated with each call. It is natural to ask what can be learned from the records of all this work

that would contribute to improvement of the products, but closing this loop of organizational learning has often proven quite difficult.

Working with a group of such reports, an analyst searches for patterns in the way the features and structure of the product interact with the circumstances of use reported by the customers. The hallmarks of complexity are present. The analyst may develop hypotheses such as: "All these customers reported that sound quality deteriorated when they were driving on country roads. Could it be that the audio unit is disturbed by shocks spaced at a particular frequency?"

Many hypotheses like this one are generated, but not all will be correct. In many organizations, such hypotheses are tested by checking if they are sufficient to reproduce the problem. In a complex world, many of those tests will fail. Someone from product development (not the same division as customer service) will subject the unit to low-frequency jolts and observe that it still performs well. An interesting strategy at such a moment of impasse is to bring into the process some of the frontline customer service agents who took the original calls. They may suggest something like, "These all came in last winter. Does it only happen if the unit is cold?"

Of course, this may not turn out to be the answer. But in an organization having trouble maintaining contact patterns between two divisions, the effort to correct a misattribution provides an occasion for interaction during which other useful information may flow. It functions as an episode of triggered recombination. Product people learn of other patterns noticed in customer service. The frontline agents learn about new product ideas in development and can then be alert to relevant remarks from customers.

Our third class of mistake is failure to appreciate the critical role of context. This kind of mistake is especially common when selection is at the level of strategies because strategies so often take the form of conditional action patterns: "If you encounter circumstances X, then do Y." The problem is that the actions are frequently much easier to observe than the conditions. For example, if your opponent in a chess game gives you the opportunity to take a piece, it may not be easy to determine from the context if this is a stupid blunder or a clever sacrifice.

To take another example, suppose you are building a collection of rare books. Bidding at book auctions may allow you to observe the buying actions of your colleagues. But if there is competition among the bidders, they may not be willing to fully, or accurately, disclose why they bought what they did, when they did. Competitive barriers to observation are often a serious impediment to strategy-level selection. Moreover, the ultimate effect of buying decisions may not be clear for some time. It can take a while to appreciate the effect on a collection of new additions, and the market for particular kinds of holdings may grow or decline.

In such an environment, learning will go slowly. Efforts to emulate apparently successful buying strategies will involve mistakes because so many factors determine the ultimate success of a purchase, and because inferences about the conditional part of the strategies are so constrained. As we noted earlier, it could be advantageous in such a situation, as in chess or checkers, to develop shorter-range measures of factors correlated with long-run success.

Again, we look for ways that the inevitable mistakes of credit attribution can provide opportunities to harness complexity. In this case, it may be possible to gradually identify signals observable in the short run that can foretell the long-term performance that is the ultimate goal. One good approach follows Arthur Samuel's insight into learning to play championship checkers. Surprises are actions that came out better, or worse, than expected. Either kind can fuel improvement. The essential thing is to see what factors were observable or predictable in the short run that were correlated with the surprise. This is a powerful idea that has been found to work not only in artificial intelligence systems but also in the neuropyschology of human learning (Cohen and Axelrod, 1984).

To return to our rare books example, we might ask what other copies of a target book were recently in the market? Are there details of its condition that might add to its value? Is the market for this type

of book cyclical or sensitive to economic conditions? Are new categories of buyers entering the market who might prefer books of this type? There are hundreds of these factors, which is why it is very hard to learn to buy well for a collection. But the harnessing complexity approach does suggest an important shift in question, asking, "What observable criteria were often high or low when you did better or worse than expected?" The search is not for what predicted the outcome but for what predicted the surprise, the deviation of your expectations from what occurred. Those are the factors to which you should give increasing credit if you want to speed the process of learning which factors to credit.

Example: Military Simulation

The problems of inferring proper lessons (attributing credit) based on limited experience occur in almost every sphere of human activity. Because military organizations only rarely obtain feedback from actual combat, their circumstances make adaptation especially difficult. Since credit attribution has long been so problematic in warfare, military organizations have a rich history of refining various forms of simulation, including many forms of gaming and field exercises. The techniques used by the military to cope with the problems of credit attribution when feedback is scarce are therefore particularly illuminating.

For these organizations the problem of determining what works well is especially vexing. Large-scale fighting is infrequent—and much work goes into keeping this true. That means that opportunities to try new weapon systems or tactical concepts, or to test officer capabilities, come rarely. Learning only from real combat experience is an unacceptably slow strategy for improvement. This is a price society happily pays for peace, but it leaves military organizations facing a difficult learning problem.

Where a firm might have several different versions of a consumer product tested in the field within a few months, a military organization might not accumulate the equivalent amount of useful experience in several decades. For an extreme example, there has never been any full combat experience for our intercontinental ballistic missile hardware, operational concepts, and crews. (Robert Powell, who creates mathematical models of nuclear deterrence, says that this field is the only branch of science where success is achieved by never having any data.)

A large portion of what military organizations learn about new technology and operational concepts must come from various forms of simulated experience. These may be war games, field exercises, small-scale engagements, mental experiments, computer models, or even imaginative reconstructions of military history.

The Information Revolution is providing computer tools that dramatically expand simulation possibilities. The United States military now routinely employs simulated aircraft, tanks, ships, and soldiers in its investigations of combat possibilities. Mobilizations of large forces for field exercises incur substantial resource costs, and even without live ammunition, there are inevitable injuries and deaths from the risky movements of personnel and heavy equipment. Such exercises cannot be repeated many times in minor variations, although exactly this capability is extremely useful in exploring a Complex Adaptive System, where deliberate variation of multiple factors may reveal large consequences.

The value of these new possibilities is also becoming evident in the business world. Although useful experience is not as scarce as in the military case, there are many situations in which exploratory trials with the real system are not possible. Major reorganizations or changes of corporate strategy are like this. They often have huge costs, and if they don't work, they

risk the bankruptcy of the entire firm. In response to this need, simulation tools for business decision making are beginning to appear. Firms are arising that specialize in building such simulation models. Some are spin-offs of computer gaming companies, while others have arisen from consulting practices (Farrell, 1998).

There are limitations, of course. One shortcoming is that simulations often place sharp and arbitrary limits on improvisation. While it is an extremely important source of military and business innovation, improvisation is generally not realistically supported in computer simulations, which often insist that the "players" obey rules and constraints that in real activities they might decide to violate. Although they may fall short of realism in significant ways, computer simulations provide the kind of rapidly assessed measure of success we have discussed previously. They generate only surrogate experience, but they can improve learning in an experience-poor domain if they are used wisely, with clear attention to their limitations.

Creating New Agents or Strategies

We have now examined three of the four aspects of selection processes set out at the beginning of this chapter: the definition of success criteria; the focus of selection on agents or strategies; and the attribution of credit that connects a measured success (or failure) to an agent or strategy. The fourth part of the process is the actual destruction of an existing agent or strategy, or the creation of a new one through copying or recombination.

We have already developed many of the key points on this topic in the course of our earlier chapter on variation, which focused on the closely related issues of creating and destroying variety. There we analyzed processes of copying and recombination, the occurrence of mistakes in those processes, and their contributions to the variation in populations. We also considered destruction of the instances of a type, up to and including extinction.

In this section, we reiterate the key role played in our framework by making, recombining, and destroying instances of agents or strategies. We add to our prior analysis by considering the consequences of the differences in detail among the many processes we have grouped together under this heading.

The Key Role of Copying

Notions of copying are central both to biology and to computer science, two disciplines that have contributed enormously to complex systems research. These two traditions do not have identical notions of copying, and the differences between them are reflected in our framework. The biological approach to making copies is much closer to our discussions of selecting at the level of agents. For most agent copying, material resources have to be assembled, and copies are made using the same materials that constitute the copied agents.

By contrast, copies as conceived in computer science concentrate on preservation of abstract form. This view corresponds more closely to our discussions of selection at the level of strategies. This alternative view of copies reaches an impressive level of abstraction in binary-encoded information that preserves its essential character across arbitrary embodiments. A digital recording of a Bach fugue is a series of "ones and zeros" that can be represented as spots of magnetism, pits in an optical disk, or a series of voltage pulses or light waves.

Both notions of copy have a place in our framework because the way copies spread through a Complex Adaptive System does not always conform to the patterns seen in natural selection. There can be adaptation, but through patterns that are not necessarily like those

seen in biology. A computer virus can spread much faster than a successful physical virus. Within hours it can clog thousands of computers all over the world with copies of itself. Thus a computer virus is different with respect to both time and space. Being immaterial, it can spread incredibly rapidly, and it can spread through a space in which "nearby" machines are physically far away. A Complex Adaptive Systems framework needs to encompass much more than the biological cases, even if those have provided much of its inspiration.

Detailed Differences Among Generic Copying Processes

Just as the difference between copying strategies and agents matters, so too do the detailed differences among various copying processes. Imitating someone's method for making telephone charity requests is not an identical process to passing along a photocopy of a fundraising letter. Both involve copying, but the former involves far more integration of a pattern into one's own behavior. Setting an example that triggers imitation is very different when the population comprises nation-states than it is when the population is made up of schoolyard playmates.

By calling many different processes "copying," it has not been our intention to deny the important differences of detail. Indeed, as we showed in Chapter II, the details have to be studied very closely. Errors and recombining processes depend on those details. And the character of the variation in the system is shaped by them in turn. Making fund-raising calls using your friend's method is much more of a recombination of strategies than is photocopying and forwarding of a funding request.

While the detailed character of copying processes is of great significance, it is also important to discuss copying processes in the aggregate. That makes clear the deep similarities among Complex Adaptive Systems. Our aim in discussing "copying" in general is to guide designers and policy makers to ask questions about how copies are made, and how destruction happens, for the agents and strategies in the systems they work with. We want to stimulate the recognition of many different kinds of processes as "copying," from duplicating computer files to replicating fast-food franchises. Once copying mechanisms are identified, the questioner will have knowledge of the important details that we cannot have. In this way, the framework aims to suggest fruitful questions.

Exercising Visible Leadership

We have given many examples of what managers, designers, and policy makers might do in a Complex Adaptive System. Virtually everything we've said about how to harness complexity can be regarded as advice about leadership. In this section, we focus on one particular aspect of leadership that deserves special attention: that what a leader does is especially likely to be copied by others.

Why would someone want to copy the visible behavior of a leader? In the ambiguous and hard-to-predict world of a Complex Adaptive System, agents often don't know what criteria of success they should use or how to evaluate the strategies they could select. This is especially important in an age of uncertainty and rapid change. When adaptive agents live in a rapidly changing environment, they tend to look to other agents to see which performance measures tend to work and which ones tend to fail. When agents are not able to predict the effects of various possible courses of action, they may resort to imitating the observable behavior of agents who seem to be successful, or who at least have more experience with the new environment (Cialdini, 1984). Imitating others who are successful or experienced is a form of implicit attribution of credit that certainly has its disadvantages. When features that are copied are only superficially relevant, the results can be wasteful or even comical. Nevertheless, following the practices of those with more experience or success is often a good strategy in an uncertain world.

There are three basic reasons a leader in a formal organization or other social system is especially likely to be copied. First, a leader can sometimes set standards that provide incentives for others to copy. Second, a leader's actions or performance measures are typically seen to be successful and hence worth emulating. Third, a leader may set an example that helps establish beneficial norms in a community.

Leadership in setting a standard can cause others to go along for their own reasons. Consider the case of Norway as a country that writes much of the world's maritime insurance. When the standards body in Norway set certain regulations for insuring oil platforms, the makers of oil platforms had an incentive to build in ways that met those standards. Thereafter other marine insurers tended to gravitate toward similar regulations (Stinchcombe and Heimer, 1985). Norway's regulations helped shape the industry in ways that led other maritime insurers to copy their visible behavior.

The emulation of a leader need not be based on a full understanding of how the emulation will help. Other agents may wish to emulate the actions or performance measures of a visibly successful leader in the hopes that what worked for the leader will work for them. A business leader who wishes to promote environmentally friendly production can, of course, make decisions that give high weight to environmental concerns. But if the firm is highly visible and is able to show that it becomes more successful because of its environmental practices and reputation, then a much more powerful dynamic comes into play. Imitation of the firm's performance measures by other firms creates a cascade that can transform an industry. Many forms of inspirational leadership work in this same fashion. For example, Gandhi's criterion of nonviolence was advanced throughout the world by the success its practitioners achieved in winning India's independence from Britain. Gandhi's leadership was successful in large part because he visibly embodied the very values he was advocating. This led others to emulate not only his tactics but also his values (Gardner, 1995).

Visible leadership can also be exercised by setting an example that helps establish beneficial norms in a community. In Complex Adaptive Systems, norms are often important regulatory mechanisms. Central monitoring and control can be difficult when many agent interactions are widely distributed across physical or social spaces. Criteria that the agents themselves apply are a very attractive alternative. Especially when they become internalized, norms regulate not through fear of consequences but through the belief that some actions are right and others wrong. This is extremely important when monitoring by central authorities is costly or intrusive. Moreover, once established, a norm can be reinforced and spread by dispersed agents who accept the norm and are willing to punish others who deviate from it (Axelrod, 1986). The Internet is a vast example of opportunities for one agent to exploit another from afar. The eventual character of its culture will be established in large measure by decisions made in the next few years, as significant and highly visible leaders promote the norms they will exemplify and expect others to enforce. The major providers of e-mail and chat facilities provoke widespread debates when they announce or modify positions on how they will handle unwanted advertisements or offensive language. The dialogs that occur build communities of users who may well enforce standards among each other more effectively than central authorities could hope to do.

We have examined five major aspects of selection: criteria of success, focus on selection at the level of agents or strategies, attribution of credit, mechanisms for creating new agents and strategies, and the exercising of visible leadership. In doing so, we have seen that each aspect has dense connections to issues of variation and interaction. Our central question in considering selection has been, "Which agents or strategies should be copied or destroyed?" But the

answer to that question is clearly intertwined with the two major questions of our earlier chapters, "What is the right balance of between variety and uniformity?" and "What should interact with what, and when?" In our concluding chapter we bring together these three elements of our framework for harnessing complexity.