1

ON DIVERSITY AND COMPLEXITY

Armageddon is not around the corner. This is only what the people of violence want us to believe. The complexity and diversity of the world is the hope for the future.
—MONTY PYTHON'S MICHAEL PALIN

In this chapter, I pose and answer some basic questions. What is diversity? What is complexity? And, why link diversity and complexity—what does one have to do with the other? First, diversity. Diversity applies to populations or collections of entities. A ball bearing cannot be diverse. Nor can a flower. Diversity requires multitudes. Cities are diverse; they contain many people, organizations, buildings, roads, etcetera. Ecosystems are diverse because they contain multiple types of flora and fauna.

When scientists speak of diversity, they can mean any of three characteristics of a population. They can mean *variation* in some attribute, such as differences in the length of finches' beaks. They can mean *diversity* of types, such as different types of stores in a mall. Or they can mean differences in *configuration*, such as different connections between atoms in a molecule.

Complexity proves to be a much more problematic concept. As mentioned in the Prelude, complexity can be loosely thought of as interesting structures and patterns that are not easily described or predicted. Systems that produce complexity consist of *diverse* rule-following entities whose behaviors are *interdependent*. Those entities interact over a *contact structure* or *network*. In addition, the entities often *adapt*. That adaptation can be learning in a social system, or natural selection in an ecological system. I find it helpful to think of complex systems as "large" in Walt Whitman's sense of containing contradictions. They tend to be robust and at the same time capable of producing large events. They can attain equilibria, both fixed points and simple patterns, as well as produce long random sequences.

To provide an example of the type of analysis that follows, I begin with an example of how diversity contributes to complexity in economics. Imagine an exchange market—a bazaar in which people bring wheelbarrows of goods to trade. This example demonstrates how diversity can reduce volatility in a system and also produce complexity. In an exchange market, diversity can enter in three ways: (1) in what the agents bring to buy and sell, their *endowments*; (2) in the agents' *preferences* for the different goods; and (3) in the ways the agents *adapt* to information, specifically prices.

If the market had no diversity, not much would happen. If everyone had identical endowments and preferences, then no one would have any reason to trade. So, we need diversity on at least one of these dimensions just to make the market come to life. Let's add diversity to both endowments and preferences so that agents bring different goods to market and desire different bundles of goods as well. In such a market, we need some mechanism for prices to form. Following standard economics, let's assume that there exists a market maker, who calls out prices with the intent of producing equilibrium trades.

Once we introduce the market maker, we have to take into account how agents respond to prices. Let's start by assuming no diversity. If all of the agents react in the same way, then prices will be volatile. They'll jump all over the place. This volatility results from everyone reacting in the same way to a price that's too low, resulting in a massive increase in demand and a similar rise in price. Gintis (2007) shows that diversity in the learning rules reduces this volatility. Later in the book, I provide a simple model involving negative and positive feedbacks that explains the stabilizing effect of variability in responses. Here, I just wish to raise the point that diversity can stabilize.

This model can be made even more complex. Kirman and Vriend (2001) add realism by dispensing with the market maker. Instead, they allow individual buyers and sellers to strike up relationships with one another. With this added realism, diversity has more subtle effects. If buyers differ in the price at which they value the goods, then buyers with relatively high values tend to pay higher prices. Furthermore, high value buyers exhibit less loyalty than buyers with low values. In this model, diversity produces complexity through the web of connections and reputations that emerge from the system. Without diversity, nothing interesting happens. With diversity, we get relatively stable market prices, but when we look at the agents and how they behave, we see a complex system.

In the remainder of this chapter, I begin with brief overviews of what is meant by diversity and complexity. I then describe how diversity contributes to complexity with some specific examples including the spatial prisoner's dilemma. I conclude the chapter by describing what I call the *assemblage* problem—the fact that many complex systems are assembled, typically from the bottom up. The fact that complex systems are assembled complicates empirical tests of the benefits of diversity.

Characterizing Diversity

There are many ways to characterize diversity. Each affects how much diversity we see in a particular situation. I may walk into a furniture store and see tremendous diversity in style. You may walk in and see no diversity at all—just a bunch of bedroom furniture. In this section, I describe several categorizations of types of diversity as well as some common measures of diversity.

One logical starting place for thinking about how to categorize diversity is to distinguish between continuous and discrete differences. The weights of the members of a murder of crows or of a parliament of owls vary. These differences in weight can take on any real value; hence we can think of them as a continuous variable. Alternatively, we can think of diversity as the number of types or as the distribution across those types. For example, to capture diversity, we might count the number and types of animals in a zoo or species in a rainforest. The two approaches, measuring variation in weight and counting the number of types, capture different types of diversity. Many of the populations that interest us will include mixtures of discrete and continuous differences. Bluebirds differ from cardinals, but among the cardinals there exist continuous differences. Some cardinals appear just a little redder than others.

Though logically clean, the continuous/discrete dichotomy approach doesn't accord with how people typically categorize diversity. Instead, people more often distinguish between differences within a type (*variation*) and differences across types (*diversity*). The notion of types will prove problematic, but people like to create types or categories. Doing so allows us to make sense of the barrage of stimuli coming at us. (It's a bird, it's a plane, it's Superman!) I should note that the within/across types categorization mostly agrees with the continuous/discrete categorization. The two disagree primarily in cases where the differences within a type are discrete, such as differences in colors of paint.

In what follows, I will sometimes also refer to diversity of *compositions* or arrangement. A washing machine and an airplane engine may contain many of the same parts, but they differ in assembly. Putting all of this together gives three types of diversity.[3]

- *Diversity within a type, or* **variation**. *This refers to differences in the amount of some attribute or characteristic, such as the height of giraffes.*
- **Diversity** *of types and kinds, or species in biological systems. This refers to differences in kind, such as the different types of foods kept in a refrigerator.*
- *Diversity of* **composition**. *This refers to differences in how the types are arranged. Examples include recipes and molecules.*

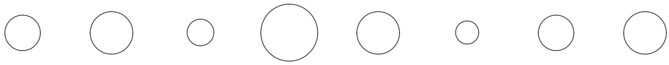

Figure 1.1. Variation: Diversity within a type.

This trichotomy will prove helpful throughout the book as I analyze the effects of diversity. Like most classifications, this one seems great if you don't think about it too deeply. Once you do, problems begin to arise. Take the length of finches' beaks. These differences would seem to fall into the category of variation. However, an ecologist will counter with the fact that finches with different sized beaks eat different types of seeds and nuts and therefore occupy different places in the food network. So, perhaps, we might also think of them as different types. In sum, this categorization won't be perfect, but it provides enough structure for us to move forward.

Variation

Diversity within a type, or variation, is often defined along dimensions, such as length, width, height, circumference, or color. Suppose that you go on a scavenger hunt and find eight marbles. If you measure the diameters of those marbles, you would probably find that they are not all the same. They exhibit *variation* in their diameters.

Variation within a type plays important roles in the adaptability and robustness of complex systems. As I just mentioned, members of the same species exhibit variation in wing size and beak length, and those differences allow them to

Figure 1.2. Diversity across types.

occupy distinct niches. Not only can the differences produce a fitness or survivability advantage for some members of that species, they also allow the species to adapt to a changing environment.

Differences of Types

When people speak of diversity, they tend to mean differences of types. Suppose that instead of asking you to gather marbles, I asked you to search your house for circular objects. You might find a frisbee, a pizza pan, a dinner plate, and a quarter. This collection would contain *diverse types* of objects even though they are all circular.

These diverse circular objects have different functions. You could eat dinner off a frisbee, and you could play catch with a dinner plate, but neither would be much fun. The functional differences between quarters and pizza pans are even more extreme. You could cook a pizza on a quarter, but it wouldn't be very filling. And, no matter how hard you tried, you couldn't load a parking meter with a pizza pan. These differences in functionalities make the world more complex, as I shall show.

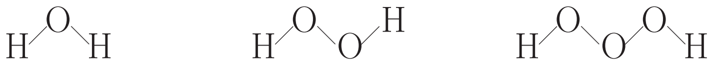

water *hydrogen peroxide* *trioxidane*

Figure 1.3. Diverse community compositions.

Differences in Community Composition

Finally, diversity can refer to differences in community or population composition. Water (H_2O) hydrogen peroxide (H_2O_2) and trioxidane (H_2O_3) all consist of combinations of hydrogen atoms and oxygen atoms, but differ in their relative amounts.

These differences in composition result in distinct emergent properties. Water has all sorts of interesting emergent properties, such as the tendency to form spheres when placed on a leaf or a freshly waxed surface, and even the ability to climb trees. Hydrogen peroxide, which differs from water by only one oxygen atom, is widely used as a disinfectant and as a whitener. It is also unstable. If exposed to sunlight it will decompose into water and oxygen, which is why it comes in brown bottles. Trioxidane, an oxidant that differs from hydrogen peroxide by only one oxygen atom, is also unstable. In the air it will decompose in a matter of minutes. If placed in water, it will decompose into a simple water molecule and an individual oxygen atom almost instantaneously.

Diversity of composition underpins much of the vast type diversity we observe in biology. The cells of all vertebrates come from only a few hundred or so types of cells. Humans, rats, and camels are comprised of muscle cells, nerve cells, glandular cells, and so on. Humans differ from rats not so much in the types of cells that we have, but in the proportions of those cells and in how those cells are arranged. That vertebrates are built from only a few cell types only moderately restricts the set of possible vertebrates. The vertebrates that presently exist are a tiny sample of what is possible (Jacob 1977).

The concept of diversity of composition provides an entrée into the concept of modularity. Many evolved and created systems are modular. Near the end of the penultimate chapter, I discuss how modularity promotes robustness. It's worth noting as well that modularity also simplifies the creation of diversity. Cars have modularized packages of extras. If you can choose from three engine modules, four stereo and communication modules, three interior models, and four trim modules, then you have a choice of one hundred and forty-four cars. The modularization is intended to guarantee that every one of those cars functions.

Complexity

Complexity has many definitions and measures. In the 1980s, Seth Lloyd began counting up definitions of complexity and stopped at forty or so (Lloyd 1988). The multitude of characterizations that Lloyd discovered reflects less a lack of agreement than an inability of any single approach to capture what scientists mean by complex. A similar problem exists for definitions of culture. Hundreds of definitions exist, and each has strengths and weaknesses. For both complexity and culture, a collection of definitions may well be needed to convey the essence of the term.

In discussing complexity, I will also devote time to describing complex systems. A complex system consists of *diverse* entities that interact in a *network* or *contact structure*—a geographic space, a computer network, or a market. These entities' actions are *interdependent*—what one protein, ant, person, or nation does materially affects others. In navigating within a complex system, entities follow rules, by which I mean prescriptions for certain behaviors in particular circumstances. These rules might be fixed: water molecules follow physical and chemical laws that are constant with respect to context.

Often, scholars distinguish between *complex systems*—systems in which the entities follow fixed rules—and *complex adaptive systems*—systems in which the entities adapt. If the entities adapt, then the system has a greater capacity to respond to changes in the environment. Adaptation occurs at the level of individuals or of types. The system itself doesn't adapt. The parts do; they alter their behaviors leading to system level adaptation.

Note that even if the individuals seek or are selected for better performance, we have no guarantee that the system will perform better, the Tragedy of the Commons (Hardin 1968) in which individual self-interest harms collective performance being the classic example of a disconnect between individual adaptation and community failure.[4]

Systems possessing diverse, connected, interacting and adaptive agents often prove capable of producing *emergent* phenomena as well as *complexity*. Before describing complexity, I take a moment to discuss emergence. Emergence refers to higher order structures and functionalities that arise from the interactions of the entities. Ant bridges, market crashes, national cultures, and collective wisdom are all examples of emergence. As the physicist Philip Anderson (1972) wrote in a seminal article, “more is different.”

Emergence underpins the idea of the ladder of science. Physics becomes chemistry, chemistry becomes biology, biology becomes psychology, and so on. Or, put another way: cells emerge from the interactions of atoms, organs emerge from the interactions of cells, and societies emerge from the interactions of people. Each level of emergence produces higher order functionalities. Cells divide. Hearts beat. People think. Societies mobilize.

Emergent properties can also be functional. Complex systems often prove robust to internal and external disturbances. This robustness emerges even though it was neither engineered from the top down nor an objective of the parts. Individual species don’t set out to create robust ecosystems. Yet, their individual pursuit of survival produces diverse interactions and behaviors that combine to form systems that can withstand mighty blows.

Juxtaposing emergence and complexity provides ample food for thought. Complexity refers to interesting behavior produced by the interactions of simple parts. Emergence refers to simpler higher order behavior that arises from underlying complexity. On the one hand, we have complexity from simplicity. And on the other hand, we have simplicit from complexity.

Keep in mind that complexity can be thought of as a property—something is either complex or not—or it can be conceptualized as a measure: a rainforest can then be said to be more complex than a cornfield. Wolfram (2002) considers complexity to be a matter of kind, a property. He classifies systems as producing one of four types of outcomes that can roughly be characterized as: fixed points, simple structures/periodic orbits, randomness, or complexity. In this conceptualization, complexity lies between simple structures and randomness.

Wolfram's definition aligns with casual intuition. Complex outcomes are neither simple patterns nor completely random. They are longer, interesting structures. This approach gives us an exclusionary test for complexity. Complex outcomes are those which cannot be classified as equilibria, simple patterns, or random. This approach works pretty well. Consider the interactions between members of a family. Most family dynamics do not produce fixed points. Families don't do the same thing at every moment. And, though there may be family traditions and the like, there's sufficient novelty that interactions do not produce periodic orbits, where each week mimics the previous one. Despite the novelty, family dynamics probably cannot be considered chaotic. Yes, daily actions may be unpredictable but family life produces recurrent structures and novelty, that is, complexity.

Most categorizations of complexity align with Wolfram's, with complexity residing between order and randomness. Unfortunately, constructing a useful measure of complexity proves difficult (see Lloyd 2001; Mitchell 2009). Measures that sound good in theory often don't work in practice. Let's start with a naïve definition based on description length and move on from there.

Description Length

The most basic definition of description length is simply the number of words required to describe some object, event, or sequence: the more words needed to describe an event, the more complex that event probably is. At the first level of approximation, description length seems to work pretty well. A dead branch is much less complex than a living tree. The branch can be characterized in far fewer words. Similarly, a game of marbles is less complex than a game of soccer. This is true whether we describe the rules or the play of a given game.

A moment's reflection reveals a potential problem with description length: the length depends on the language used. Fortunately, the problem of multiple languages has been overcome. To show how, I have to formalize the notion of minimal description length. It's worth working through this level of detail for two reasons. First, we get a peek at the subtleties involved in constructing a useful definition of complexity. Second, we see the extent to which the complexity depends on the encoding.

First, let me give you an idea of why choice of language matters. Suppose that I want to communicate that a period of human existence had a particular spirit. In the English language, I would have to say something awkward like "the spirit of the age," whereas in German, I could just use the word "zeitgeist."

I can take this same idea and make it more formal. Suppose that I want to describe numbers, but the only number in my language is three. To express the number five, I would have to write $3 + \frac{(3+3)}{3}$. To write six, I would only need to write $3 + 3$. Thus, six would be less complex than five in this language.

To show that the language only matters up to a point, we first need to assume that the phenomenon of interest can be written in a *description language*. We can then define *Kolmogorov complexity* as follows:

Kolmogorov complexity: *The minimum length of a program written in the description language that produces the desired sequence of symbols.*

If we have two languages, A and B, then we can write the Kolmogorov complexity of a sequence s relative to those languages as $K_A(s)$ and $K_B(s)$. The *invariance* result states that given any two languages A and B, there exists some constant C such that for any sequence s, $|K_A(s) - K_B(s)| < C$.[5] This result implies that for sequences with high Kolmogorov complexity, the choice of language does not matter.

So, at this point, we can say that if the thing we're describing is reasonably complex—if it requires a lot of words—then the choice of language doesn't matter much. This does not mean that minimal description length is a perfect measure. It has two serious shortcomings. The first is technical. Given a sequence and an alphabet, there does not necessarily exist an algorithm that will spit out the Kolmogorov complexity.

The second problem is more relevant for our discussion. Description length assigns high values to random sequences. Consider the following three sequences of 0s and 1s. If it helps to make the presentation less abstract, think of 1s as days the stock market goes up and 0s as days that the stock market goes down.

000000000111111111

010011000111000011111

011011110101100010100

The first sequence can be described as "ten zeroes followed by ten ones." That's a total of six words. The second sequence can be described as "K zeroes then K ones: K from 1 to 4." That's nine words. The third sequence can be described as "zero, two ones, zero, four ones, zero one twice, one, zero, zero one twice, zero zero." That's a lot of words. Yet, if we look at the second and third sequences, we see that the second has patterns and structure and the third is random. Therefore, rather than give a long description of the third, we'd like to just say "it's random." That's only two words.

As the sequences above make clear, description length conflates randomness and complexity. As Huberman and Hogg (1986) point out, complexity is not the same thing as randomness.[6] Instead, complexity lies in between order and disorder. One approach to making randomness less complex which that has been advocated by Murray Gell-Mann and Seth Lloyd, uses the concept of *effective complexity*. This concept emends the idea of minimal description length by considering only the "length of a highly compressed description of its regularities" (Gell-Mann and Lloyd 1996). In other words, Gell-Mann and Lloyd strip away randomness and count only what's left.

Alternatively, complexity can be measured as the difficulty of generating the sequence. This can be calculated either as the minimal number of steps, *logical depth* (Bennet 1988), or as the number of steps in the most plausible sequence of events, *thermodynamic depth* (Lloyd and Pagels 1988). The logical depth of a human being would be huge, but it's relatively small compared to the thermodynamic depth, which also takes into account the most likely evolutionary history that resulted in humans.

Measures that capture the difficulty of generating a phenomenon align with how many people think about complexity, but they too have shortcomings. First, they have computability issues just like description length (Crutchfield and Shalizi 1999). So, the measures work great in theory, but they're cumbersome to apply in practice. A measure that you cannot calculate quickly has limited practical value even if it's theoretically sublime.

A second problem with this whole class of measures is that they apply to a fixed sequence of events or outcomes. The system being studied must have a beginning and an end. We might instead consider a sequence that continues to grow over time and ask how difficult it is to predict that sequence as accurately as possible. This idea underpins Crutchfield and Young's (1989) concept of *statistical complexity*. With statistical complexity, a random sequence would have low complexity because a machine that generates that sequence would be relatively simple. To calculate statistical complexity, an algorithm classifies past data into categories so as to produce patterns that are statistically indistinguishable from the real data. Though not easily calculated, Crutchfield and Young's measure has been derived in some cases.

So What is Complexity?

The analysis so far reveals multiple approaches to measuring complexity. The fact that complexity, which is itself a complex idea, lacks a single definition should, thus, not be a surprise. Nor should it be seen as undermining the science of complexity. A scientific approach to complexity is possible: it just may need multiple lenses. Following Melanie Mitchell

(2009), I see the abundance of definitions as more a strength than a weakness. Diversity can be good. Nevertheless, it would be nice to have at least one or two core definitions, principles, or characteristics of complexity that we can use as a foundation. Here are two: complexity is a BOAR (a somewhat wild one), and complexity is DEEP.

BOAR *Complexity lies* **B***etween* **O***rder* **A***nd* **R***andomness.*:

DEEP *Complexity cannot be easily* **D***escribed,* **E***volved,* **E***ngineered, or* **P***redicted.*

Before moving on, one further point merits attention. Complexity has also been described as lying on the edge of chaos (Langton 1990). I could just wave my hands, claim "chaos = randomness" and move on, but unfortunately, I cannot. Chaos is not randomness. Chaos refers to extreme sensitivity to initial conditions. If we change the initial point by a little bit in a chaotic system, we end up on different paths. Many chaotic systems are completely deterministic. If you know the current state, then you know all future states. When we speak of randomness, we often mean the exact opposite: complete unpredictability. The next flip of the coin does not depend in any way on the current flip of the coin.

There's a simple resolution to this seeming contradiction. Complex systems that produce randomness, such as Wolfram's one-dimensional cellular automata, are also extremely sensitive to initial conditions. They have to be. They're using simple rules, so in order to produce a long random string they have to be adding new information that comes from points far away on the one-dimensional grid. Wolfram (2002) dedicates several pages to this insight in his book. In systems that produce simple structures, there exists very limited sensitivity to initial conditions. Class IV systems (complex systems) lie in between. They produce some sensitivity to initial conditions, but not an extreme amount. So, in a sense, we can say that complexity lies both between order and randomness and between order and chaos. The systems that produce randomness are also highly chaotic.

How can we make sense of all this? Well, we could draw a cover of the *New Yorker* showing Complexistan as a triangle bordered on one side by the Sea of Order, on the second by the Ocean of Chaos, and on the third by the River of Randomness. The interior of that land would not be easily explained. It would be complex.

Complexity and Diversity Together at Last

We can now begin to ask: if complex systems can produce complexity how does diversity contribute to making them do so? I'll begin with a system that's not complex and show that including variation within types and diversity across types doesn't have much of an effect. The setting I want to consider is building a small wooden deck. The process of building a wooden deck is pretty close to an ordered system, and it's not difficult to describe. Once built, the deck will be about as noncomplex as can be. It'll be an equilibrium system with non-adapting parts (save for shrinking and expanding with temperature). But here, we're considering the construction process. The variation in the width of the boards has no functional value. In fact, the builder hopes that whatever variation does exist cancels out due to averaging. Any individual two by four won't be exactly three and a half inches wide (the actual width of a two by four). Some will be a little wider and some a little narrower. With enough boards, these deviations from the mean will cancel one another out—or come close—so that the variation won't matter.

In this example, variation and diversity contribute to complexity, but only a little. The diversity of the parts, the mix of two by four planks, four by four posts, etc. makes the process more complex—in that it makes the structure more elaborate and the process harder to define. But, whether we use BOAR or DEEP, we don't find very much complexity. Overall, then, in noncomplex systems, variation has no real effect, and diversity's effect is minimal.

Now let's move to a system that's more complex: a cakewalk involving nine-year-olds. For those unfamiliar with a cakewalk, imagine the numbers one through twelve written on pieces of tape, which are then stuck to a gym floor, making a circle with a diameter of about eight feet. This creates a clock face. One person stands on each number. When all twelve numbers are covered, the music starts. As in musical chairs, once the music begins the participants walk clockwise around the numbers. When the music stops, each participant stops as well, and a number from one to twelve is drawn from a hat. The participant standing on the number that matches the number drawn wins a cake.

Let's start with a cakewalk with no diversity or variation. Each person walks at exactly the same speed (no variation), and each person moves to the nearest number in front once the music stops (no diversity). Given these assumptions, the cakewalk is not complex. It's not between order and randomness. It's ordered, with a random winner. Nor is it difficult to describe. People walk in a circle. Music stops. A winner is drawn.

Now let's add some variation in walking speed. Some people walk quickly and some people walk slowly. With variation in walking speed, participants bunch up while the music plays. When the music stops, some numbers won't have any people near them, and other numbers will have too many people. If four participants are bunched together near one number, then they will have to scurry to open spots. If they all follow the same rule: *move clockwise to the first open space*, they would all keep circling forward and sequentially fill in the gaps in the circle. If so, depending on the number of fast walkers and how they are spaced, the resulting pattern could be rather beautiful. For instance, suppose eleven fast walkers all bunch up behind a single slow walker. When the music stops, the fast walkers will all walk past the slow walker and one by one peel off onto open numbers. (Think through this pattern; it's pretty cool.)

Next, let's add in some diversity in the behavioral rules, so that when the fast walkers get bunched up they don't necessarily move in an orderly fashion around the circle. Instead, some people might look backward. Others might look across the circle and run toward the first open spot they see. This diversity would destroy order. If the rules people follow are chosen without much thinking, the system will become more random. Again, it won't be entirely random—it will lie *between order and randomness*. As a result, describing the process and predicting where people will end up proves more difficult. The cakewalk becomes DEEP.

Notice the relative magnitude of the effects of variation and diversity. The set of imaginable outcomes in the scenario where we vary walking speed is vastly smaller than the set of outcomes that are possible when we allow diversity in behavioral rules. It will often be the case that when we ramp up the amount of diversity, the set of possible outcomes grows enormous.

To complete the logic, suppose that we have a totally random process, such as throwing a bunch of coins on the floor. If we increase variation by adding coins of different sizes, we don't fundamentally alter the outcomes, which will be random sequences of heads and tails. If we add diversity by adding dice, then we increase the set of possible outcomes. However, we could always convert the numbers on the die using binary notation to heads and tails. If we did, our results would not be any different from the initial experiment. To summarize, increases in variation and diversity can increase complexity but they typically won't have big effects on ordered systems unless they tip them out of the ordered region, or on random systems.

An Example of How Diversity Produces Complexity

To further demonstrate how diversity can produce complexity, I present a model by Nowak and May (1993). This model considers the evolution of cooperation in a spatial setting. In their model, agents play a Prisoners' Dilemma (PD) game. The PD captures a variety of strategic contexts including firm level price competition and arms races. In both cases, each of two players decides whether to cooperate (C) or defect (D). For each player, defecting always gives a higher payoff, but if both players cooperate, they are better off than if they both defect. For example, both firms make higher profits if they keep prices high (if they cooperate), but each does better by cutting prices and getting a bigger market share. The payoffs can be written as follows:

The Prisoners' Dilemma: General Case

		Column Player	
		C	D
Row Player	*C*	**R**,R	**S**,T
	D	**T**,S	**P**,P

The payoffs in the PD satisfy two conditions. First, $T > R > P > S$. This condition guarantees that defecting gives the higher payoff regardless of the other player's action. Second, $2R > S + T$. This condition guarantees that joint cooperation produces the highest average payoff. The PD game has been the focus of thousands of papers in fields ranging from economics, to political science, to ecology. It commands so much interest because individual incentives promote defection but collectively everyone is better off cooperating. A numerical example makes these incentives clearer:

The Prisoners' Dilemma: Numerical Example

		Column Player	
		C	D
Row Player	*C*	**4**,4	**1**,5
	D	**5**,1	**2**,2

Given the optimality of cooperation, two questions arise. First, can cooperation be induced? Second, will cooperation emerge? By the latter, I mean—will players cooperate even though it is in no one's self-interest to cooperate? In the repeated game setting, individuals play many times in a row. This enables individuals to use more sophisticated strategies. Both questions can be answered in the affirmative. Cooperation can be induced and it also emerges. Consider the strategy called *Tit for Tat*. In *Tit for Tat*, an individual cooperates the first time the game is played and continues to cooperate so long as her opponent cooperates. It's straightforward to see that if both individuals play *Tit for Tat*, then both cooperate. In a famous experiment in which individuals submitted strategies with the goal of earning the highest payoff, *Tit for Tat* proved to be the winning strategy (Axelrod 1984). So, not only can cooperation be supported, it can emerge with real players.

Recall that our focus here is on how diversity can create complexity. In the one-shot PD, there appears to be little complexity—everyone should defect. But suppose that we put a complexity spin on the PD. Recall the core attributes of a complex system: *diverse*, *interdependent*, *networked* entities that can *adapt*. By definition, a game assumes interdependent payoffs, and if the initial population includes both defectors and cooperators, then it is diverse. To create a complex system, we need only add dynamics and a network. This is what Nowak and May did. And they did so in the simplest possible way.

Their network was a two-dimensional lattice with N rows and N columns. In other words, they put agents on a checkerboard. Rather than assume that agents play a single agent or the entire population of agents, they assume that each agent plays its four neighbors to the North, South, East, and

West, as shown in the figure below:

		N_1		
	N_4	A	N_2	
		N_3		

Figure 1.4. Neighbors on a lattice.

To avoid boundary issues, connect the left edge of the lattice to the right edge and the top to the bottom. The first connection creates a tube. The second turns the tube into a donut, what mathematicians call a torus. Many of the canonical models in complex systems take place on tori.

Nowak and May populate their model randomly with cooperators and defectors. They then calculate the *fitness* or *payoff* of each agent as the sum of its payoffs in all four of its games. Consider the following two agents:

	C				D	
C	D	C		C	C	D
	C				D	
	Agent 1				Agent 2	

Using the payoffs from the numerical example of the PD, Agent 1 earns a payoff of twenty, the highest possible payoff, and Agent 2 earns a payoff of only seven (one each in the three games with defectors and four in the game with the cooperator).

To add dynamics to the model, Nowak and May assume that each agent looks to its neighbors. If one of those neighbors gets a higher payoff, then the agent matches the neighbor's action. It's important to note that the agents are not employing sophisticated strategies like *Tit for Tat.* Instead, each agent is either a cooperator or a defector. For the purposes of this discussion, I will refer to the action of an agent as its *type* to align with the earlier definition of diversity of types.

The extended PD can now be thought of as a complex system. It has diverse types with interdependent payoffs situated in a network and we've characterized the dynamics: agents switch their types if one of their neighbors is of a different type and earns a higher payoff. So, what happens? First, note that any agent of the same type as all its neighbors will maintain that type. Therefore, the interior of a region of all cooperators (or defectors) will not change in the next period. What's left to figure out is what happens at the boundary between regions of defectors and cooperators. Surprisingly, defection doesn't necessarily take over. To see why, consider the following configuration:

C	D	C	D	C
D	D	D	D	C
D	D	D	D	D
C	C	C	C	D
C	C	C	C	C

Figure 1.5. Evolution of types.

The agent in the center of the grid (denoted by the circle) is of type D. The agent's payoff equals eleven. To determine what the agent's type will be in the next period, we calculate the payoff of each of the agent's neighbors. The neighbors to the East, North, and West all also earn payoffs of eleven, but the neighbor to the South earns a payoff of thirteen. Therefore, in the next period, the agent in the center changes its type to C. This suggests that the system might support cooperation in the long term. And in fact, it can.

The dynamics of this model prove to be complex. They have a long description length, so they are DEEP. And if we apply the BOAR criterion, we see that they are neither ordered nor random. They lie between the two. Early on, the model produces dynamic structures that are consistent with Wolfram's complexity class. Over time, the model produces relatively stable regions of cooperators and defectors with chaotic boundaries (Nowak and May 1992, 1993; Schweitzer et al. 2002). When we have multiple regions, they need not all be in the same state. Some can be fixed, others in periodic orbits, and others chaotically moving about. It makes sense to characterize the overall system as complex.

Emergent Levels: How Simple Parts Create Complex Wholes

In the informal thought experiment of the cakewalk and in the formal PD model, increasing diversity ramps up complexity. These examples align with the common intuition that diversity and complexity go hand in hand. It's natural to think that diversity is required for complexity. Like many intuitive causal relationships, this one proves subtle. A system need not have diverse parts to produce complexity. In fact, two of the most famous complex systems models, Conway's Game of Life and Wolfram's two-dimensional cellular automata model, possess minimal diversity: cells (or agents) assume one of two states—on or off.

Let's consider the Game of Life. As in Nowak and May's PD game, in the Game of Life, agents exist on a torus. Each agent (or cell) can be alive (type 1) or dead (type 0). In the Game of Life, agents have eight neighbors. These include the neighbors to the North, South, East, and West and the four neighbors on the diagonal.

	N_3	N_5	N_8	
	N_2	A	N_7	
	N_1	N_4	N_6	

Figure 1.6. Neighbors in the Game of Life.

The rules for the Game of Life are deceptively simple. A dead agent comes to life if it has exactly three live neighbors; otherwise it stays dead. A live agent remains alive if and only if it has two or three live neighbors. Otherwise, it dies, either of boredom (fewer than two live neighbors) or of suffocation (more than three live neighbors). The Game of Life can produce complex patterns including blinkers that flip back and forth, gliders that float across the torus, and even pulsing glider guns that spit out gliders at regular intervals.

If the Game of Life doesn't include much diversity, how can it produce complexity? Two answers: large numbers of parts and interdependence that produces emergent structures. First, the large number answer: a long string of zeros and ones proves a sufficiently rich space to support complexity, just as a long string of DNA can contain the instructions for life. Complexity requires little diversity in the parts, provided there are enough parts. With enough zeros and ones it's possible to say anything. Second, the parts are interdependent. The sum, the component consisting of the parts, can be more complex than the parts themselves. What is a brain but a collection of spatially situated simple parts that interact according to rules? In the brain, the rules depend on chemistry and physics, whereas in the Game of Life, the rules depend on logic, but in both cases simple parts produce complexity.

To see this logic in greater detail, return to the idea of emergence. In the Game of Life, the fundamental objects are networked and interact according to rules. Given that additional structure, the two types can produce higher order structures. Those higher order structures produce the complexity. Let me make this more precise. Even though each agent can be only one of two types, each three by three square of nine cells takes on 512 possible types.[7] Larger configurations of agents can support structures like the aforementioned glider guns. It is those diverse structures that then produce the complexity. Thus, not very diverse fundamental building blocks prove capable of generating complexity because they produce diverse higher order parts.

Here, then, is the take away: *fundamental diversity* is not required for complexity. *Emergent diversity* is. The Game of

Life produces complexity through the interactions of diverse interacting parts, but those parts are not the cells. The relevant parts are the emergent structures, like the gliders. These exist on a higher level. Recall the third type of diversity—community level diversity. Even though the entities themselves are not diverse, they form diverse communities. Just as hydrogen and oxygen can combine to produce water and hydrogen peroxide, live and dead agents can combine to produce blinkers, gliders, and glider guns.

In characterizing the relationship between diversity and complexity, it is also important to remember that diversity is not sufficient on its own to produce complexity. If you visit a landfill, you'll see diverse consumer waste products, but you won't see much complexity. That's because pieces of trash don't interact in interesting ways. As the saying goes: when you put garbage in, you cannot get anything but garbage out. At the same time, a system cannot just have any interacting rules and produce complexity. The rules must interact in particular ways. Flip open your computer and write a computer model with a large number of types of agents. Make each of those types meaningfully diverse; that is assign each a distinct behavioral rule. When you're finished, you're more likely to get a mangle, stasis, or randomness than complexity.

For some classes of models, it has been shown that in the space of all possible interacting rules, complexity is a low probability event (Wolfram 2002). Creating a complex system from scratch takes skill (or evolution). Therefore, when we see diverse complex systems in the real world, we should not assume that they've been assembled from whole cloth. Far more likely, they've been constructed bit by bit.

Each step along the way, complex systems rely on their rule-based parts to adapt to changing surroundings. London began with a few individuals carving out a life along a bend in the river. Over time, as London expanded, the parts—people, businesses, governments, and social organizations—adapted. In turn, they enabled ever more diversity to emerge and the city to become ever more complex.

Measuring the Effects of Diversity in Complex Systems

I've shown some examples in which diversity helps produce of complexity. Later in the book, I'll focus on the effects of complexity. That analysis will be made more difficult because complex systems are assembled through selection. Only entities that function well within the system survive. This process of selection of the entities in a complex system can be thought of as a form of assembly. By that I mean, the system does not contain agents with foresight or intention who put the pieces in place. Ecosystems self-assemble. Other systems, such as organizations, are assembled deliberately with a specific purpose in mind. Market economies and political systems lie in between. They self-assemble to some extent, and they're also partly assembled. In social system, actors—governments and other large organizations—structure interactions and incentives so as to alter the assembly process.

Assembly implies that the level of diversity in a system has survived some winnowing process. This winnowing creates three problems for empirical tests of the effects of diversity. I refer to these as the *problem of multiple causes*, *the sample problem*, and *selection (squared) bias*. I cover each in turn.

The Problem of Multiple Causes

In most complex systems, multiple factors influence the independent variable of interest. Economic growth, ecosystem robustness, etc. depend on more than one variable. We cannot therefore just run a regression of economic growth on human diversity or of robustness on the diversity of flora and fauna. We must also account for other variables. Determining the independent effect of diversity may be problematic because the independent variables may not be sufficiently random.

An example makes this clearer. Suppose that the social networks that survive can be classified as one of two types. The first type consists of diverse people with few connections, and the second type consists of homogeneous people who have lots of connections.[8] To ground the analysis, suppose that networks of the first type prove to be more innovative. We cannot discern whether that innovativeness was due to the diversity of the individuals in the network or to the looseness of their connections. To unpack the true cause, we would also need either strongly connected networks of diverse people or weakly connected networks of homogeneous people. Those might not exist. Thus, we might not be able to determine the effect of diversity.

Therefore, stylized facts that more diverse societies are more innovative or less trusting do not necessarily imply that diversity causes those phenomena (Putnam 2007). The effect (innovation or trust) might be caused by lower connectivity between people or by something else. If diversity correlates with some other characteristic of systems, we cannot be certain whether an empirical phenomenon is caused by diversity, the other characteristic, or some combination of the two.

The Sample Problem

The second difficulty for empirical analysis concerns defining the sample. Do we take all possible combinations of entities or only those that have survived some form of selection? Consider recipes. Recipes exist in a complex system. They adapt to tastes (or more accurately, they're adapted). And they face selective pressure. If a recipe's no good, it won't survive. The idea that recipes evolve is supported by data. Physicists who examined cookbooks from four cultures found a scale-invariant distribution of ingredients that's consistent with a "mutate and copy" model of recipe evolution in which new recipes make small modifications in existing recipes (Kinouchi et al. 2008).

The recipes that we find in books, then, are not random combinations of ingredients. They're the result of a long process of combining and winnowing. To see why this causes a problem for empirical research, suppose that we want to learn whether a greater diversity of ingredients improves the quality of the recipe. Making this a well-defined problem requires measures of diversity and quality. Let's use the number of ingredients as a diversity measure. To measure quality, assume that there exists a panel of food tasters who rate each dish between one and ten.

Here's the key question: how do we define the sample? Over what set of recipes do we test whether quality improves with diversity? Do we consider all possible recipes? Do we consider only those recipes that currently exist? Or do we attempt to define a set of theoretically plausible recipes and randomly select from that distribution? If I choose all recipes, then I will find that quality decreases with higher diversity.

Diversity does not always increase quality within this domain. An arbitrary combination of ingredients probably produces an inedible mélange. Peanut butter-watermelon-squash enchiladas drenched in chocolate sauce doesn't sound tasty.

Selection (Squared) Bias

Suppose that we define our sample to be those recipes whose qualities lie above some survivability threshold—perhaps recipes that appear in a cookbook. This selection could bias empirical tests. For example, suppose that recipes with more ingredients require more time and effort. If the person writing the cookbook assumed that people trade off the quality of a recipe against the costs of making it, then the author would place a higher quality bar on recipes with more ingredients. A mediocre recipe with two ingredients might make the book, but a mediocre recipe with twenty ingredients wouldn't. However, a great recipe with lots of ingredients would be included. With this rule of thumb for choosing recipes, the result would be a collection whose quality would improve with diversity.

Let me put numbers to this logic to make it more precise. Assume recipes have qualities between zero and twenty and that the cookbook's author requires a quality greater than or equal to the number of ingredients to include the recipe in the book. If so, recipes with two ingredients will have qualities between two and twenty, recipes with five ingredients will have qualities between five and twenty, and recipes with fifteen ingredients will have qualities between fifteen and twenty. An empirical test that regresses the quality of a recipe on the number of ingredients will find that quality increases with the number of ingredients.

Notice that diversity itself isn't producing any benefits. The cookbook writer's recipe filter, and not any direct benefit from diversity, causes the diverse recipes to be better. This creates what statisticians call *selection bias*: the cases considered aren't random. In complex systems, (statistical) selection bias is caused by (evolutionary) selective pressures. The species that currently exist in the world have survived selective pressures. Therefore, the selection (the set of species existing in an ecosystem) may be biased by evolutionary selective pressures. We can think of this as selection bias caused by selection or *selection (squared) bias*. If a system exhibits selection (squared) bias, diversity may appear beneficial but not be.

Example: The Diversity-Stability Debate

As mentioned in the prelude to this book, in ecology there exists what is called the diversity-stability debate. This debate focuses on whether diverse ecosystems are more dynamically stable and more robust to perturbations. Later, I flesh out the distinction between stability and robustness in more detail. For now, think of a stable ecosystem as one that will go back to its equilibrium if the sizes of the populations are changed, and a robust ecosystem as one that can survive environmental shifts and the loss or addition of a species.

The diversity-stability debate demonstrates how the selection (squared) bias plays out in a real world example. The ecosystems we study aren't random. They've been chosen by evolution. The debate may also exhibit the problem of multiple causes if diverse ecosystems differ systematically from nondiverse ecosystems. That this basic of a question—does diversity make ecosystems more stable and more robust?—has proven difficult to prove empirically or theoretically speaks to the challenges of making definitive statements about the effects of diversity in complex systems.

This diversity-stability debate has current relevance given recent declines in the number of species. If species reductions create less stable ecosystems, then current reductions could beget future reductions and even collapses.[9]

The origins of the diversity-stability debate are empirical, but theory also plays a central role. A nonspecialist reading of the literature suggests that early on most ecologists accepted that ecosystem stability and robustness increased with species diversity. The evidence for this came first from the extremes. Rainforests were relatively stable while farms at times suffer from massive infestations and blights. Furthermore, islands with fewer species were more susceptible to invasion (Odum 1953; Elton 1958). (Bracket for the moment that rainforests have been constructed through tens of thousands of years of evolution while farms are artificially constructed.) The observation that rainforests were less prone to booms and busts led to the hypothesis that having multiple predators and prey made a system less likely to have large population fluctuations.

The logic for why diversity increased stability, as advanced by MacArthur (1955), goes as follows. If you have a single predator and its population declines, then its prey's population might explode. If swallows are the only check on mosquitos, then a sudden decline in swallows for some external reason will result in giant swarms of mosquitos. If, however, bats are also present in the ecosystem, a decline in swallows will be offset by an increase in the number of bats, thus preventing the mosquito population from getting too large.

This logic seems pretty compelling, yet it wasn't completely accepted. Part of the reason for the lack of acceptance had to do with the distinction between stability and robustness. And part had to do with the fact that the empirical evidence is messy (Woodwell and Smith 1969).

Robert May (1973) complicated the debate by constructing a mathematical model that showed the opposite conclusion, causing quite a stir. In May's model, interactions between species were assigned randomly. May found that increasing the number of species and the number of connections led to lower asymptotic stability. In retrospect, this result isn't as devastating to the MacArthur model as was first thought. In real ecosystems, interaction strengths evolve. This means that that what we see in the real world isn't a random collection. Think back to the discussion of random and assembled recipes. There the assembled recipes showed increased quality in the number of ingredients, but that would not be true for random recipes. From one perspective then, we should not be surprised that randomly created ecosystems are less stable. If we just started dumping species of all kinds into a habitat, we'd likely get an initial maelstrom, in which many species swirled to their deaths.

One response to May's result has been to replicate it with more realistic assumptions on interaction strengths. Yodzis (1981) constructed models with interaction strengths calibrated to real ecosystem data. He found that these models produce more complex *and* more stable dynamics than did

May's random interaction strengths model. This does not settle the debate so much as open more questions. Additional theoretical and empirical research has further muddied the waters. Tilman (1997) and others have constructed ecosystems in prairies and have explored how those ecosystems have fared over time. They have found a strong correlation between diversity and robustness.[10] Note that here, I use the term robustness because Tilman and his colleagues examined whether the ecosystems continue to function and not whether they return to the same equilibrium.

In his excellent book on ecological commons problems, Levin (2000) describes how May's results have been misinterpreted. Recall that May focuses on asymptotic stability. Stability and robustness differ. A system could be unstable, or continue to fluctuate, yet be very robust. Levin points to viruses. The flu virus never reaches an asymptotic equilibrium. Were it to, we could eradicate it with a vaccine. Viruses' robustness stems from the fact that they present a moving target. We're not fighting a flu virus, we're fighting a diverse, fluctuating population of flu viruses.

Schneider and Kay (1994) offer a similar critique. They view the entire debate as misguided. Diversity, they argue, could be counted in many ways. Do we use number of species? Do we count the millions of bacteria in the soil? It's difficult to make an empirical claim without a well-defined independent variable. Furthermore, they argue, stability doesn't make any sense. Here they echo Levin by noting that ecosystems are continuously in flux. Given that ecosystems are complex adaptive systems, Schneider and Kay fail to see how it makes much sense to talk about diversity creating stability when the systems being studied are not stable.

The deep unanswered question in this entire debate is why are systems diverse. Later in this book, I'll discuss how diversity contributes to robustness through particular functions. But I don't answer why robustness, a system level property, emerges in ecosystems. In designed systems, such as political systems or traffic systems, one can make an argument that the designers are mindful of robustness in their planning, that robustness can be built into a system (Bednar 2009). I agree that this is possible and only add that designers often borrow features from robust ecosystems even though the origins or emergence of those features within ecosystems remains somewhat of a mystery.