CHAPTER 1

Diverse Perspectives

HOW WE SEE THINGS

Those French, they have a different word for everything.
—Steve Martin

We all differ in how we see and interpret things. Whether considering a politician's proposal for changes in welfare policy, a new front-loading washing machine, or an antique ceramic bowl, each of us uses a different representation. Each of us sees the thing, whatever it is, in our own way. We commonly refer to the ways we encode things as *perspectives*. But if asked what a perspective is, most of us would have only a crude idea. In this chapter I provide a formal definition, but before I get to that I'll present an example of a famous perspective: the periodic table.

In the periodic table each element has a unique number. These numbers help us to organize the elements. They give structure. Compare this perspective to the perspective that uses common names such as oxygen, carbon, and copper. By convention we know what those names mean—copper is a soft brownish metal that conducts electricity—but the names don't create any meaningful structure. They are just names. We could just as well give copper the name *Kamisha*.

Mendeleyev's periodic table gave us a meaningful structure. Coming up with that perspective took hard work. To discover the structure of the elements, Mendeleyev created cards of the sixty-three known elements. Each card contained information about an element including its chemical and physical properties. Mendeleyev then spent hours studying and arranging these cards, transforming the problem into a representational puzzle. Eventually, Mendeleyev pinned the cards to the wall in seven columns, ordering the cards from lightest to heaviest. (Imagine playing solitaire on the wall using thumbtacks.) When he did this, he saw a structure that was completely understood only three decades later with the introduction of atomic numbers. Before Mendeleyev, atomic weight had been considered irrelevant. A scientist could order the elements by atomic weight from lightest to heaviest, but he could also arrange them alphabetically or by the number of letters in their name. Why bother?

As some of the elements had not been found, Mendeleyev's table had gaps. New elements were soon found that filled those gaps. Mendeleyev took information, turned it into the pieces of a puzzle, and showed us that pieces were missing.[1] Mendeleyev's representational puzzle, unlike the problem of finding the chemical composition of salt, lacks a physical analog. He was not searching for an existing structure; he was creating a structure out of thin air. That structure revealed order in the stuff of which we're made. His story is not unique. We can find stories like Mendeleyev's throughout the history of science—think of Copernicus and the heliocentric universe, or of Einstein and the construction of relatively theory. In both cases, someone saw the world differently—Einstein linked space and time—and what had been obscure, confusing, or unseen became clear.

Scholars from a variety of disciplines have studied how people and groups make breakthroughs. The common answer: *diverse perspectives*. As the philosopher of science Steven Toulmin wrote, "The heart of all major discoveries in the physical sciences is the discovery of novel methods of representation."[2] New perspectives, what Toulmin calls "novel methods of representation," are often metaphorical. The canonical model for earthquakes, for instance, involves blocks connected by springs, which can then be analyzed rigorously using mathematics.[3] Though we know perspectives lead to breakthroughs, their sources remain shrouded in mystery. The only necessary ingredient appears to be hard work and a willingness to look at things that others ignore. That's also a recurrent theme. Being diverse in a relevant way often proves hard. Being diverse and irrelevant is easy.

We can now define perspectives formally. They are representations that encode objects, events, or situations so that each gets its own unique name. No two chairs, no two people are represented in the same way. Mathematicians call these "one-to-one" mappings or bijections. That's a strong but necessary assumption. The names that perspectives assign to objects capture underlying structure. If they do not, then the perspectives are not of much use.

In this chapter, we drive home one main insight: *the right perspective can make a problem easy.* We see how most scientific breakthroughs and business innovations involve a person seeing a problem or situation differently. The germ theory of disease transformed what seemed like an intractable confusing mess of data into a coherent collection of facts. Thanks to Adam Smith, we all know the story of the pin factory and how efficient it is. But do many of us know that the first pin factory manufactured brushes with firm steel bristles? The firm began producing pins only when someone realized that the bristles could be cut off and made into pins. Diverse perspectives—seeing the world differently, seeing the brush as a forest of pins—provided the seeds of innovation.[4] Let's begin.

Putting Our Different Shoulders to the Wheel

I first provided some hints of how diverse worldviews can be useful in solving a problem such as cracking the Enigma code. We now consider a problem related to cracking a code of a different form. This example will be a bit of a teaser for the second part of the book, when we consider the application of diversity to problem solving, prediction, and choice.

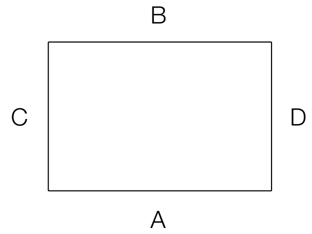

Figure 1.1 The Geometrician's Perspective

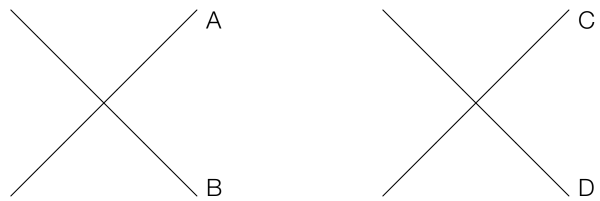

Figure 1.2 The Economist's Perspective

We start with a puzzling fact: in every object examined by a team of scientists, they find that the amount of *A* equals the amount of *B* and the amount of *C* equals the amount of *D*. For the moment, don't worry about what *A*, *B*, *C*, and *D* are. What might we do with such a fact? How might we make sense of it? Let's put on several different types of academic hats and see what we might do with it.

We might first think like geometricians and imagine that the letters represent the sides of a rectangle (see figure 1.1). If we label the top and the bottom of the rectangle *A* and *B*, and label the two sides *C* and *D*, we can then exploit the property that the opposite sides of a rectangle must be equal. We have explained our fact. Of course, we don't have any idea what to do with this rectangle, but still, we have an idea that whatever we find might well be rectangular.

Or we might suppose that we are economists (see figure 1.2). If so, we might reason that *A* is the amount of some good supplied and *B* is the amount demanded. Or we might reason that *A* is the

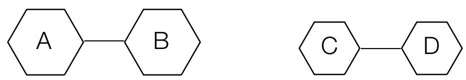

Figure 1.3 The Chemist's Perspective

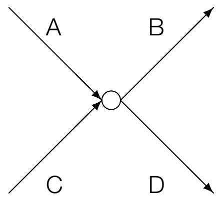

Figure 1.4 The Physicist's Perspective

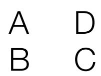

Figure 1.5 The Fashion Designer's Perspective

price of the good and *B* is the marginal cost. Or we might think that *A* is the price of an input and *B* is its marginal product. In each of these examples, some force equilibrates the amounts.

Or we might imagine ourselves chemists (see figure 1.3). We might reason that some chemical reaction results in molecules containing equal amounts of *A* and *B*, and of *C* and *D*.

Or we might imagine ourselves physicists (see figure 1.4). We might then reason that *A*, *B*, *C*, and *D* are properties that must be conserved such as energy or matter that can be neither lost nor gained.

Finally, we might imagine ourselves fashion designers with nifty, small glasses (see figure 1.5). If so, we might think of *A* as the number of people wearing argyle sweaters and *B* as the number of people wearing blue jeans. We might think of *C* as the number of people wearing cutoffs and *D* as the number wearing dingy

T-shirts. If everyone wearing an argyle sweater wears jeans and everyone wearing a dingy T-shirt wears cutoffs, then $A = B$ and $C = D$.

Each of these possible perspectives embeds an intelligence based on training and experience. And each of these ways of looking at the problem would prove useful in some cases and not so useful in others.

To see diverse perspectives at work, let's look at the real story: the discovery of the structure of DNA. Francis Crick and James Watson's piecing together of the double helix structure involved many hours of hard work and lots of dead ends. Overcoming these dead ends (what we will later call "local optima": perspectives that look pretty good until you search further) required that they develop new perspectives on their problem. One fact at their disposal, a fact ignored by many others working on the problem, was Chargaff's rules, which stated that a cell's nucleus contained equal amounts of adenine (A) and thymine (T) and equal amounts of guanine (G) and cytosine (C). The natural way to look at these rules was chemical and to think that they happened to be produced in equal amounts by some chemical reaction in the cell. But Crick one day channeled his inner fashion designer and thought to pair the adenine and thymine and the guanine and cytosine. These pairings formed the rungs in the helical ladder that is DNA. In retrospect, the idea to create these pairings seems obvious. At the time, it was a breakthrough.

The most amazing aspect of Crick's pairings idea was not that it led to the solution, but that it solved a second puzzle: the pairings revealed how DNA can be the building block of life. To see how this worked, we use our fashion designer's perspective but replace our blue jeans (B) with trousers (T) and make all of our dingy T-shirts green (G) so that we are using the appropriate letters (A=T and C=G).

Imagine a long line of people, some wearing argyle sweaters and trousers and some wearing green dingy T-shirts and cutoff shorts. We can identify each person with a two-letter code. AT is someone wearing an argyle sweater and trousers, and GC is someone wearing a green shirt and cutoffs. To describe the outfits of the entire line of people, we make a long list: *AT*, *AT*, *GC*, *GC*, *GC*, *AT*, and so on. In making this list, we've wasted a lot of ink. We don't need two letters to identify each person. If we know what pants a person wears, we know what shirt he wears (and vice versa). Thus, we need only identify one piece of clothing on each person, and we can use logic to fill in the rest. The single piece of clothing list *A*, *G*, *T*, *C*, *T*, *T*, *A*, *C* provides enough information to reproduce the full list *AT*, *GC*, *AT*, *GC*, *AT*, *AT*, *AT*, *GC* because the sweaters are always matched with the trousers and the T-shirts with the cutoffs.

This insight explains how our cells reproduce. (Hard to believe, but true.) When the helix splits, each half contains sufficient information to reproduce the missing other half. Every *A* can be matched with a *T*, every *C* can be matched with a *G*, and so on. The single strands of DNA can be thought of as half dressed, but fortunately our cell chemistry completes their outfits by matching appropriate bottoms with tops and tops with bottoms.

With the aid of modern microscopes, any one of us could now uncover the structure of DNA. But Crick and Watson didn't have modern microscopes. All they had were fuzzy pictures taken by Rosalind Franklin. (Franklin didn't share in the Nobel Prize, not because she was a woman, but because she happened to be dead at the time they awarded it.) Crick and Watson's construction was a monumental scientific achievement that required diverse thinking. Crick and Watson may be two white boys, but viewed cognitively, they're a diverse pair. If a person can be diverse, Crick was. Crick's training spanned physics, chemistry, and biology. He did not have Ph.D.s in these areas. In fact, he had no Ph.D. at all. James Watson did have a Ph.D., in zoology. He had been a wunderkind ornithologist, but became obsessed with DNA after studying viruses. Could Crick have unraveled the puzzle without Watson or Watson without Crick? Most doubt it. Historians of science assign credit to their hard work and their diverse skills. They leveraged their differences and together achieved far more than either could have alone. In Robert Wright's brief popular account, he writes, "here is a case where one plus one equals twelve."[5]

To understand this new mathematics, to see how one plus one can be twelve, we need more formalism. We need to begin the harder (and more satisfying) work of building frameworks and models.

The Perspective Framework

We now get down to business and consider the perspective framework in detail. In this framework, we'll assume a large set of objects, situations, events, or solutions that must be given a representation. That set could be big. It could contain a billion billion billion things, but it is finite nonetheless.[6] Think stars, atoms, and all creatures great and small. The challenge is that they must be organized. To do this, each person possesses an *internal language* that describes these objects, situations, events, or solutions. An internal language can be written in words, in numbers or symbols, or in abstract shapes and forms. An internal language differs from a spoken or written external language. Internal languages can assign different words to the same object.

What do we mean when we say that these internal languages differ across people? To modify an example from Daniel Dennett, one person could internally represent a right triangle by the length of its two edges. Another person could internally represent the same triangle with a nonright angle and its adjacent edge. These two people can communicate the existence of the triangle by drawing it, but they don't have to translate between their internal languages; they need only translate their internal languages into reality.[7]

So we have some thing, some object or event, and we have an individual's internal language with which they represent that thing. Think of the UPC codes on products as an internal language for scanners. Each item in the supermarket gets a unique code. We call the mapping that takes reality and encodes it in the internal language a *perspective*. Perspectives assign unique names to objects. Mathematicians call them "one-to-one mappings." Each object, situation, problem, or event maps to a unique "word" in the person's internal language.

A **perspective** *is a map from reality to an internal language such that each distinct object, situation, problem, or event gets mapped to a unique word.*[8]

Although perspectives represent solutions in terms of an internal language, from here on I refer to perspectives as both the mapping to the internal language and as the representation itself. This makes the prose much easier to follow, but it comes at a small cost. When we say that two people have different perspectives we could mean one of two things. We could mean that they map reality differently into the same internal language—long strings of zeros and ones—or we could mean that they map reality into different internal languages. But in both cases, the perspectives differ and that's our main focus.

The first perspectives most of us learn are bases—not in baseball but in mathematics. When we say that the speed limit is seventy miles an hour, we use base ten. We mean seven times ten miles per hour. In my older son's kindergarten class, they use the concept of base ten to teach addition. To add seven plus five, a child grabs five straws and then seven more straws. He or she then counts out ten straws, which get traded in for a "bundle." Bundles consist of ten straws tied together. This bundle is then added to the two remaining straws to give an answer of "one ten bundle and two ones" or 12. The teacher might alternatively ask the students to create bundles of size eight. Assigned that task, the child could create one eight bundle with four straws left over. So, in base eight, the answer equals 14. The number of straws could be written in any number of bases. In base four, the answer would be 30, and in base two, the language of computer scientists, the answer would be 1100.

Though we think in base ten, we need not. It is just one of many perspectives we can use to represent numbers. Its use is not universal. The Mayans, who had greater toe awareness, used base twenty. Even more amazing, cultural anthropologists report that one former British colony still uses base sixteen for some weights and measures. Once they reach sixteen ounces, they use the term *pint*.

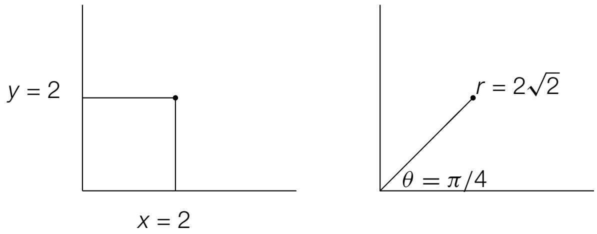

Figure 1.6 Cartesian and Polar Perspectives

The second time that we confront diverse perspectives in mathematics is usually in trigonometry, a subject loved by all. In trigonometry, we learn that a point in space can be represented using Cartesian coordinates, the familiar (x, y), or using something called polar coordinates. Polar coordinates describe a point by an *angle* θ and a distance from the origin called the *radius* r. Figure 1.6 shows a point represented in Cartesian coordinates as the point $(x, y) = (2, 2)$ and in polar coordinates as the point $(r, \theta) = (2\sqrt{2}, \frac{\pi}{4})$.

The Cartesian and polar coordinates systems are both perspectives of points in Euclidean space. Neither perspective is better than the other (despite claims made by those "Polar coordinates are cooler" T-shirts). Cartesian coordinates simplify the description of rectangles by labeling a length and a width. Describing a rectangle using polar coordinates requires a calculator, a clean sheet of paper, and an eraser. But polar coordinates show their worth when we have to describe circles and arcs. A circle is all of the points of some fixed radius. To describe a circle using Cartesian coordinates, we have to rely on complicated functions of x and y. Thus, easy problems in polar coordinates can be difficult in Cartesian, and vice versa.[9] I'll often return to this insight—*the choice of perspective contributes to a problem's difficulty*. It's a touchstone for the more general results.

A second insight is that perspectives often impose structure. For example, if our internal language assigns numbers, we create a complete ordering. One is the smallest number. Two lies between one and three, and so on. Complete and even partial orderings give language power. They simplify identification.[10] If an internal language fails to create structure, if it consists of idiosyncratic names such as *Isabella*, *Roland*, *Susan*, and *Hugo*, then, unlike Mendeleyev's periodic table, the language cannot clarify relationships. Even if we were to alphabetize the names, we learn nothing about the relatedness of Roland and Susan. They could be husband and wife. They could be father and daughter. This raises a key point: *Internal languages that fail to create structure do not aid problem solving or understanding. To be of functional value, perspectives must embed meaningful relatedness.* Just assigning names isn't enough.

Ben and Jerry

To see more concretely how a perspective creates structure, we consider an example that takes some liberties with a story about how Ben and Jerry's Homemade developed a new flavor.[11] In addition to foreshadowing future claims about the benefits of diverse perspectives, this example provides further evidence of the benefits of owning an ice cream company. Suppose that when Ben and Jerry were searching for a good recipe for New York Super Fudge Chunk Ice Cream, they created a two-dimensional array of pints. The two dimensions that they considered were the number of chunks and the size of the chunks. These dimensions impose a structure on the pints that is understandable and manipulable.

Their perspective allows them to compare and maneuver in meaningful ways. They might want more chunks, or larger ones. Given their perspective, they can do that. They have a structure they can exploit. Had they given the pints arbitrary names such as *Captain Crunchy* and *New York State of Mind* and randomly arranged them on the table, then searching among the pints would be nothing more than a random search. But having a structured perspective allows for structured search. Figure 1.7 is a representation of Ben and Jerry's perspective, which is structured. It helps us make sense of how ice cream differs in a meaningful way.

4,1	4,2	4,3	4,4
3,1	3,2	3,3	3,4
2,1	2,2	**2,3**	2,4
1,1	1,2	1,3	1,4

Figure 1.7 Ben and Jerry's Perspective: Number of Chunks (in dozens) and Size of Chunks (cm diameter)

250	252	255	256	258	262

Figure 1.8 Nelly Armstrong's Perspective (in calories)

3,2	**2,3**	**4,1**	3,3	2,4	4,3

Figure 1.9 Nelly's Perspective in Number and Size of Chunks (in calories)

The internal language consists of numbers representing the number of chocolate chunks and the size of those chunks. Their perspective maps the pints into numbers. This perspective, though useful, is just one of many. It just happens to be the one that Ben and Jerry chose. To see a different perspective, suppose that Ben and Jerry hired a consultant named Nelly Armstrong, and that she arranged the pints according to the number of calories per serving (see figure 1.8).

These two perspectives create distinct spatial relationships between the pints of ice cream. We can translate Nelly's perspective into Ben and Jerry's by replacing the calories in each pint with the number and size of chunks (see figure 1.9).

In Nelly's perspective, the pint with two dozen three centimeter diameter chunks lies adjacent to the pint with four dozen chunks of diameter one centimeter (they are typed in **bold** in her perspective). In Ben and Jerry's perspective, these two pints are far apart

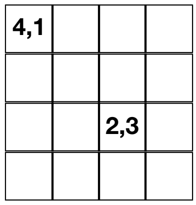

Figure 1.10 Ben and Jerry's Perspective (Number of Chunks, Size of Chunks)

(see figure 1.10). Suppose that Ben and Jerry proposed the pint with two dozen chunks of diameter three, and the consultant then suggested the pint with four dozen chunks of diameter one, and let's suppose that hers tasted better. Ben and Jerry would have perceived this as a cognitive leap, as a brilliant insight by their highly paid consultant. But it is only a giant leap because of the structure imposed by their perspective. It is a small step for Nelly. (If you would like to test these theories in practice, get out your ice cream maker and get cranking.)

Nelly's perspective leads to an improvement, but let's not confuse it with genius. Genius is Einstein linking space and time to create his theory of relativity. Nelly takes into account the fact that people count calories—a good idea but not a great one. We might even speculate on the origin of this perspective. Calories may have been salient for Nelly. Some experience, such as having a father who attended Weight Watchers, could have caused her to focus on calories and not chip size. Nellie also may have gained this perspective from her years of experience in the food industry, which taught her that calories affect taste. Regardless, it gave her a different way of seeing than Ben and Jerry's—and that's what did the trick.

Not all diverse perspectives are helpful. Just because someone brings a different perspective doesn't mean that it will lead to a better solution. To see this, consider a second consultant, Hannah, who organizes the pints of ice cream by the average length of time it takes to chew a spoonful of the ice cream. In order to rationalize being paid some stupendous fee, Hannah might give this attribute some impressive name like *masticity*. For the problem of making ice cream, this perspective would not create meaningful structure. It would be a diverse but not very useful perspective. That doesn't mean it's always a bad perspective. A food industry consultant who saw this example remarked that masticity does matter for breakfast cereals. Cereals should be neither mush nor bark and twigs.

Our Ben and Jerry's example shows a benefit of diverse perspectives in problem solving. Perspectives can also be useful in strategic situations. The difficulty of formulating a strategy depends on how the strategies are represented. To see this, we consider three games.

Three Games

These games that we consider require skill, like chess and Go, although they are not as challenging. Unlike children's board games such as Candy Land or Chutes and Ladders or adult games such as poker, bridge, and backgammon, they involve no luck. The first game is called Sum to Fifteen. Herbert Simon, a computer scientist and political scientist who was awarded the Nobel Prize in economics, created it.

Sum to Fifteen

Setup: Nine cards numbered 1 through 9 are laid out on a table face up.

Order of Play: One player is randomly chosen to go first and then players alternate taking cards.

Object: To collect three cards that sum to 15.

Sum to Fifteen creates two contradicting incentives. A player wants to build combinations that could sum to 15 while simultaneously preventing the other player from doing the same. Walking through a single play of the game reveals some of its complexity.

Suppose that the first player chooses the 5 and the second player chooses the 3. The first player would then not want to choose the 7 because $(5 + 3 + 7 = 15)$ and the 3 has already been selected.

TABLE 1.1:
First player's cards and needs

Cards	*Sum*	*Needed Card*
2,4	6	**9**
2,5	7	8
4,5	9	**6**

Suppose then the first player chooses the 2. Her cards now sum to 7, so the second player must choose the 8 to prevent the first player from winning. The second player's cards now sum to only 11, meaning that the first player must choose the 4. After this card has been chosen, the situation looks as follows:

Player 1: Holds 2, 4, 5
Player 2: Holds 3, 8
Remaining Cards: 1, 6, 7, 9

The table above shows all pairs of cards that the first player holds, their sum, and what third card must be added to it in order to sum to 15. For example, cards 2 and 4 sum to six, so card number 9 brings the total to fifteen. The numbers in bold in the last column have yet to be chosen.

Looking carefully at this table reveals that the first player wins. If the second player chooses 6, the first player chooses 9. If the second player chooses 9, the first player chooses 6. In addition, the second player has no card that he can choose that allows him to win. Game over.

The second game is more complicated. In this game, each of nine picnic baskets contains a unique combination of food items. These items are Nachos (N), Eggs (E), Sausage (S), Water (W), Hot dogs (H), Vinegar (V), Lemons (L), and Raisins (R). Across them, the nine baskets contain three of each food item. The goal of this game is to collect baskets that contain all three copies of some food item. If either player gets all three lemons, or nachos, or eggs, and so on, that player wins.

TABLE 1.2:
The Unpacking Game

Basket	*Contents*
1	H, W
2	S, E, R
3	N, V
4	N, E, L
5	H, V, L, R
6	S, W, L
7	S, V
8	N, W, R
9	H, E

THE UNPACKING GAME

Setup: Nine baskets containing the items as shown in table 1.2 are placed on the table.

Order of Play: One player is randomly chosen to go first and then players alternate choosing baskets.

Object: To collect all three copies of one of the food items.

In the interests of brevity, we won't walk through a play of this game. You can try it by yourself. It's complicated and requires lots of mental accounting. As in Sum to Fifteen, players in the Unpacking Game must balance offense against defense.

The third game is familiar to all of us. In America we call it Tic Tac Toe, but in England, where it originated centuries ago, they call it Noughts and Crosses. Tic Tac Toe is an easy game. The fun of playing it wears off quickly.

TIC TAC TOE

Setup: Play begins with an empty three-by-three array of boxes. (see figure 1.11)

Order of Play: The player randomly chosen to go first places an X in an empty box. The other player places an O in an empty box. Players alternate placing Xs and Os.

Object: To get three Xs or Os in a row.

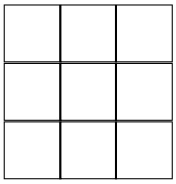

Figure 1.11 Tic Tac Toe

8	3	4
1	5	9
6	7	2

Figure 1.12 A Magic Square

I won't walk you through the play of Tic Tac Toe. It's a simple game. Now, the other games, they were challenging. Or were they?

Suppose someone said that the other games are no harder than Tic Tac Toe. You'd think they were mathematicians, crazy, or both. But, in fact, not only are they no harder than Tic Tac Toe, they *are* Tic Tac Toe—in different perspectives. To see the equivalence of Tic Tac Toe and Sum to Fifteen, we have to learn about magic squares. Some readers may remember these from seventh-grade math. In a magic square, every row and every column sums to fifteen, as do both diagonals (see figure 1.12). Take the top row: $8 + 4 + 3 = 15$. Take the second column: $3 + 5 + 7 = 15$. Take the upward sloping diagonal: $6 + 5 + 4 = 15$. Every row, column, and diagonal sums to fifteen.

Let's play Tic Tac Toe on the magic square. When a player places an X or an O, they must erase the number in the box—which makes the game Sum to Fifteen. Reconsider the play of our game of Sum to Fifteen that we walked through earlier. The first player chose the 5 (put an X in the center box), and the second player then chose the 3 (put an O there). The first player then chose the 2 and the second player the 8. Now when the first player chooses the 4, he has two ways to win, as shown in figure 1.13.

O	O	X
1	X	9
6	7	X

Figure 1.13 Tic Tac Toe on a Magic Square

N,W,R	N,V	N,E,L
H,W	H,V,L,R	H,E
S,W,L	S,V	S,E,R

Figure 1.14 Unpacking Tic Tac Toe

In Sum to Fifteen, the second player's choice of three didn't seem so bad. Seen in the Tic Tac Toe perspective, it looks dumb. Proving that Sum to Fifteen is the same game as Tic Tac Toe requires a little work. The number of sets of three numbers that sum to 15 that include the number 8 equals three. These sets are $\{8, 2, 5\}$, $\{8, 1, 6\}$, and $\{8, 3, 4\}$. All of these are possible in Tic Tac Toe on the magic square. The number 5 belongs to four sets of three numbers that sum to 15 and all four of those also exist on the magic square. Similar arguments can be made for the other seven numbers, establishing the result.

Let's next look at the Unpacking Game. We can see the nine baskets as the nine boxes in Tic Tac Toe. The contents of a box are the paths that contain that box. Instead of having *N* denoting Nachos, let it denote the Northern path—the three boxes in the top row—and instead of eggs, let *E* denote the Eastern path. We can similarly define Southern *S*, Western *W*, Horizontal *H*, Vertical *V*, ladder (upward sloping) *L*, and ramp (downward sloping) *R* paths. The eight food items in the Unpacking Game can be seen as equivalent to the eight ways to win at Tic Tac Toe as shown in figure 1.14.

If we place the basket numbers on the Tic Tac Toe board, we again get the magic square (see figure 1.15).

8	3	4
1	5	9
6	7	2

Figure 1.15 Baskets on the Board

This example shows how diverse perspectives cut both ways; for every brilliant perspective that changes a difficult situation into an easy one, there may be a multitude of perspectives that muck up our understanding of even strategic contexts. Different perspectives can simplify, but they can also muddle.

This example also demonstrates the value of being open-minded and listening to new ways of thinking. A person concentrating on playing Sum to Fifteen might ignore some other person prattling on about magic squares and Tic Tac Toe. This other person knows a better way to play the game, but the better way requires adopting a new perspective, which isn't easy.

Building New Perspectives

Novel perspectives on problems do not come from the ether. We often construct them from other perspectives. In this section, we see how that can be done and see evidence of the superadditivity of diversity: how one plus one equals twelve. To be precise, five perspectives create ten pairs of perspectives. Adding a sixth perspective creates, for free, five new pairs of perspectives. These paired perspectives can enable us to solve hard problems. The problems we consider here come from an intelligence test. Spatial and analytic intelligence tests partly attempt to measure the ability to generate perspectives. Thus, many standard questions on IQ tests ask "fill in the missing number" questions. Solving these problems requires encoding the given numbers in such a way that they make sense. This requires finding a perspective that makes the sequence into a pattern. Consider the following three problems taken from actual IQ tests.

IQ Test Questions: *In each sequence, replace the X with the unique number that makes the sequence logically consistent.*

Sequence 1: 1 4 9 16 X 36

Sequence 2: 1 2 3 5 X 13

Sequence 3: 1 2 6 X 1,806

The first example is the easiest. It is a sequence of squares. The square of 1 equals 1, the square of 2 equals 4, and so on. The missing number is 25. The second sequence appears to be the prime numbers, but that is not correct. One is not considered a prime. Even if it were, the sequence of primes, 1, 2, 3, 5, 7, 11, and 13 requires space for two numbers between 5 and 13. Answering 7 to that question is what is often called a "good wrong answer."

The perspective that makes sense of this sequence is to recognize each number as the difference of the two that follow it. The first number equals the third number minus the second ($1 = 3 - 2$), the second number equals the fourth minus the third ($2 = 5 - 3$), and so on. It follows that the fifth should be such that it minus the fourth number, 5, equals the third number, 3. Therefore, the missing number is 8. The number 8 makes sense of the entire sequence because $13 - 8 = 5$, a sequence known as the Fibonacci sequence. It has many nice properties and is taught in many math classes.

People who happened to have seen this sequence in high school math probably get this question right on an IQ test and therefore appear smarter than they are. People who have never seen this sequence before and have had to develop this perspective on the fly can find this problem rather hard. It is not easy to generate this perspective under the stress of an exam.

Our third example is one of the more difficult mathematical sequence problems found on an IQ test. It differentiates those people with extraordinarily high mathematical and logical skill levels from those people who are just good. In this example, the last number, 1,806 seems out of place. It is too large. How could a logical sequence jump from 6 to 1,806? We can find the answer by combining the perspectives developed to solve the first two sequences. In doing so, we'll see the superadditivity of cognitive tools, building on existing tools to make ever more tools. Recall that solving the first sequence requires seeing the numbers as squares. The perspective that makes sense of the second sequence takes differences between successive numbers. Neither of these perspective works on this third sequence. But if they are combined, they reveal the pattern.

First, apply the perspective used in the second sequence: Look at the differences between numbers. The difference between the first two numbers equals 1 ($2 - 1 = 1$). The difference between the second two numbers is 4 ($6 - 2 = 4$). This suggests a pattern. That pattern is the perspective used to solve the first sequence: squares. Each number differs from the number after it by an amount equal to its square $1 = 2 - 1^2$, and $2 = 6 - 2^2$. This idea seems cute, but it doesn't seem as though it will get us to 1,806. And yet it does. Using this rule, the next number would be 42, $6 = 42 - 6^2$, and the number after 42 would be (guess what) 1,806: $42 = 1{,}806 - 42^2$ ($42^2 = 1{,}764$). Combining our two perspectives, we can make sense of the third sequence.

Intelligence tests that rely on sequencing problems assume some correlation between intelligence and the ability to create novel perspectives. This may not be a bad assumption. We think of someone as intelligent if they do well in school with little effort and if they have the ability to solve a range of problems. Succeeding at those tasks requires the ability to retain, generate, and combine perspectives. Having a high IQ is not the same thing as being able to use and develop perspectives, but the two are related. We return to this point later.

The ability to solve any one sequencing problem depends on some mixture of ability, experience, and serendipity. All three of these effects can be seen by reconsidering the last, and hardest, sequencing problem. In his cult classic, *The Hitchhiker's Guide to the Galaxy,* the late Douglas Adams wrote that the number 42 was the answer to the Ultimate Question of Life, the Universe, and Everything.[12] Science fiction buffs would be far more likely to guess the number 42 than would people who had never read Douglas Adams. Science fiction buffs who like math might even have seen the pattern once they inserted 42 in the sequence because $42 - 6 = 36$, a perfect square.

In these examples, once the pattern is revealed, the problem seems relatively easy. Herein lies the paradox of novel and useful perspectives. Because they make sense of a problem or situation, because they organize knowledge, they seem obvious after the fact. Of course force equals mass times acceleration. Of course the earth revolves around the sun. Of course we evolved from single-cell organisms. Of course we're composed of little vibrating strings with lots of hidden dimensions—okay, maybe not all perspectives are obvious.

Mount Fuji and Ice Cream Summits

To formalize some of these ideas, consider Alexander Pope's famous epigraph about Isaac Newton: "Nature and Nature's laws lay hid in night / God said, let Newton be! and all was light."

Now extend it to the more general claim that how people see problems determines how hard those problems are. Newton saw physical phenomena clearly. He saw white light as composed of all of the colors. He saw orbiting bodies as attracting. The mysterious made clear. The difficult made easy. We can make the following claim:

The Difficult Eye of the Beholder: *How hard a problem is to solve depends on the perspective used to encode it.*

What does this mean? Formally speaking, it means that we can encode problems so that they are easy. We can encode them so that they are hard. To demonstrate this more precisely, I will introduce the concept of a rugged landscape. Imagine that we want a perspective on houses to help us understand their prices. We consider fourteen houses for sale in West Branch, Iowa (the boyhood home of Herbert Hoover) on www.realtor.com in the fall

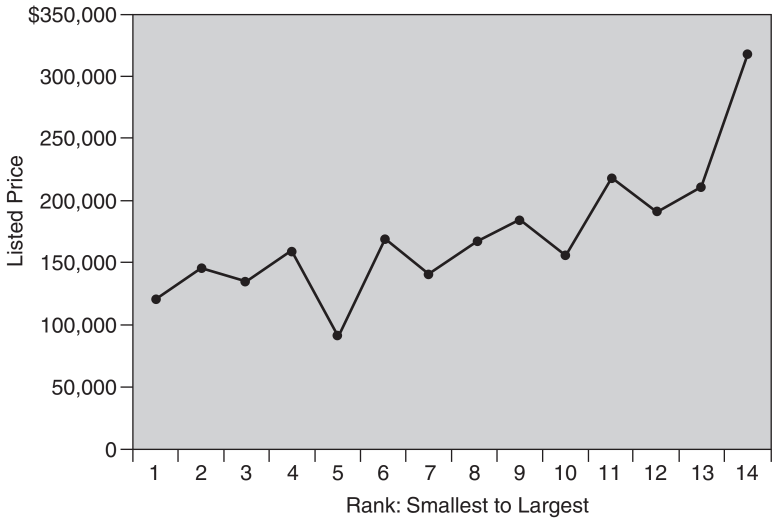

Figure 1.16 West Branch Houses in the Square Foot Perspective

of 2005. The prices of these houses range from a low of $88,500 to a high of $319,000. We can use their square footage to create a perspective on these homes. We can arrange the houses from smallest to largest and plot the prices of the fourteen homes as a function of their ranks. House number 1 is the smallest house, house number two is the next smallest, and so on. We can think of this graph as a landscape: the higher the elevation, the higher the price of the home. This landscape has several *local peaks*. A local peak is a point on the landscape from which any direction you move lowers your elevation. It is, literally, the peak of a hill. The *global peak* is the point of highest value. It's the peak of the biggest hill. On Earth, Mount Washington is a local peak, a windy one at that, but Everest is the global peak.

If square footage were a great perspective for the price of houses, this landscape would have a single local peak, which would be the global peak. But as we see in figure 1.16, our landscape is pretty rugged.[13] Instead of leading gracefully to a single peak (that looks like Mount Fuji), this perspective creates lots of ups and downs. Moving along it would be like hiking in the Adirondacks. The many local peaks imply that square footage doesn't organize the information well. If we wanted to find the least expensive house (assuming that we couldn't just get the information from the Web site) and used the square footage perspective, we'd probably start with the smallest house. If we checked the second smallest, we would find that it is more expensive. We might then think that the smallest house is the least expensive, but it's not; the fifth smallest house is. This perspective isn't helpful.

The best perspectives create landscapes with a single peak—they organize the information in such a way that a single, obvious solution becomes clear. It's easy to imagine a hiker following the various flash marks on a trail (the data points) and getting to the top. But using the square footage perspective on the West Branch houses creates lots of little peaks where our hiker could get stuck (there are five to be exact). Or if we think back to our first consultant's caloric perspective on the ice cream, a person could walk along the pints with a spoon—not a bad day's work—to reach a peak. To use the more flowery language of the late British poet laureate Ted Hughes, that peak would be an "ice cream summit," when neither the pint to the left nor that to the right tasted better.

If we restrict ourselves to one-dimensional perspectives on the fourteen houses (ordering them by their distance from the road, for instance), we can construct 8,192 distinct single-peaked perspectives. How do we know that there are 8,192 such perspectives? We arrive at this number using a perspective on the number of possible perspectives. Here's the trick: assign each house a number from 1 to 14 so that the house numbers increase with price. House number 14 has the highest price and house number 1 the lowest. Note that this is a different numbering than our square footage ranking. A one-dimensional perspective can be thought of as a list of these fourteen numbers. If the perspective creates a single-peaked landscape, the house numbers in the list must increase up to 14 and decrease after it. The perspective 1, 4, 5, 6, 9, 11, 14, 13, 12, 10, 8, 7, 3, 2 is single-peaked, but the perspective 1, 2, 3, 4, 13, 5, 6, 7, 8, 9, 14, 12, 11, 10, is not. It has a second peak at 13. Each distinct single-peaked landscape has a unique subset of houses to the left of the house number 14. So, to compute the number of perspectives that create single-peaked landscapes, we need only compute the number of unique subsets of the numbers from 1 to 13. This equals two raised to the thirteenth power.[14] We can extend this logic to the more general case with any number of objects.[15]

Thus, for any problem, Mount Fuji exists. In fact, many Mount Fujis exist. We can state this as a formal claim. We call it the *Savant Existence Theorem*.[16]

The Savant Existence Theorem: *For any problem, there exist many perspectives that create Mount Fuji landscapes.*

The outline of the proof of this theorem mimics our example. We arrange the solutions from best to worst and then assign them numbers that increase up to the best solution and decrease thereafter. This creates a Mount Fuji landscape. This approach may seem too good to be true, and it is. To use the approach, we need to know the value of every possible solution. If we had all that information, we'd have no need for a perspective. We'd just choose the solution with the highest value.

This does not mean that the theorem isn't important. It is important. It gives hope. For any problem, some Newton-like savant could show up and make everything clear as a bell, creating Mount Fuji out of the Badlands of South Dakota. Nevertheless, for some problems, we don't have good perspectives and we may never have good perspectives. We're waiting for Godot.

The Savant Existence Theorem has a flip side: the number of perspectives is enormous. So, it may not be likely that we find one that creates Mount Fuji. For example, if we return to the West Branch real estate market and place no constraints on the ordering of the houses (we can choose any one of the fourteen to place first, any one of the remaining thirteen to go second, and so on), then the total number of orderings equals 14 times 13 times 12 times 10 times 9 times 8 and so on. If we multiply all these numbers together, we get 87,178,291,200. So, for each perspective that creates a Mount Fuji landscape there exist *ten million* perspectives that do not. Ouch!

In general, the number of one-dimensional perspectives of N objects equals N factorial—N times $(N - 1)$ times $(N - 2)$ and so on, or what mathematicians write as $N!$. This number dwarfs the 2^{N-1} perspectives that create Mount Fuji landscapes. These calculations do get a bit technical, but they have a purpose. They demonstrate that having a unique perspective may not be that hard—we could arrange the houses by their distance from the road or by their house numbers. Both would be unique. They also show that for any problem, many perspectives make the problem simple. They create a Mount Fuji landscape that we can faithfully follow to the top. Unfortunately, relative to the total number of perspectives, these simplifying perspectives are few and far between. Precious few perspectives make problems simple. The vast majority of perspectives create no meaningful structure. People who discover these Mount Fuji perspectives, like Newton, get remembered in poems. The more perspectives we collectively possess, the better our chances of finding a Mount Fuji.

Identical Perspectives, Communication, and Groupthink

An implication of what we have covered so far is that when people see a problem the same way, they're likely all to get stuck at the same solutions—if we look at a problem with the same perspective, we're all likely to get stuck at the same local peaks. As we saw in the Ben and Jerry's example, someone who represents the problem differently (and not as well) probably has different local optima—a different peak. This person can help the group get unstuck. We might ask, Why, other than lack of imagination, would people rely on the same perspective? People may share a perspective because it's useful. If someone has a better perspective on a problem, copying it would seem to make sense. As counterintuitive as this advice sounds, copying better perspectives may not be such a good idea. Collectively, we may be better off if some of us continue to use less effective but diverse perspectives.

Another reason for common perspectives is that they allow for quick and error-free communication. An experiment by Colin Camerer and Roberto Weber shows how the incentive to communicate leads people to see the world in the same way, to use the same perspective.[17] In their experiment, two players are shown an identical picture of a nondescript business meeting on a computer screen. On one player's screen, a person in the picture is identified; let's say it's a man in a green Izod shirt. The player then sends a message to the other player, who in turn must click on the identified person. Each time the second player clicks on the correct person in the photograph, both players get a small cash reward. In the experiment, the players were given a fixed length of time to accumulate as many correct identifications, and as much money, as possible.

As you might expect, players developed situational patois, crude languages that helped to identify people quickly and accurately. The man in the green Izod shirt became "gator." The tall woman with the beauty mark above her lip became "Cindy" (after the supermodel Cindy Crawford), and so on. Each pair developed its own unique language.

Camerer and Weber then created havoc by switching the pairings of players. Initially, these new pairs did poorly, much worse, in fact, than the original pairings did. When a player sent the message "gator," his new partner just got confused. Sending the message "green gator" wasn't much help either. The message "the blond guy in his fifties in the green Izod shirt, you idiot!" did work, but that took time (and cost money) to send. Even though the perspectives had been developed only a few minutes earlier, they had resilience.

Similar scenarios play out in business, the academy, and the world of politics. Recall the Steve Martin quotation that begins this chapter. The same might be said of marketers, engineers, accountants, physicists, sociologists, and biologists. We all speak in jargon. We may not be aware of it, but we do. We construct communication shortcuts, often with the help of acronyms. Acronyms are so abundant that people who attend a seminar or meeting outside their own organization or field of specialization often have no idea what people are saying. Acronym overlap makes understanding even more difficult. The acronym ABM represents everything from activity-based management to agent-based models, anti-ballistic missiles, and the Association of Breastfeeding Mothers. This leads many of us to espouse AFC. (That's acronym-free communication.)

Although common perspectives arise because of imitation and the need to communicate, they also arise for less productive reasons. People are social, and insecure, animals. Members of a group sometimes lock into a common perspective because they feel more comfortable thinking about the world the same way that other people do. These common perspectives can be a type of groupthink.[18] The logic of groupthink rests on our desire to conform. If a majority of people thinks of a problem one way, they often compel others to do so. That way could be a good perspective and, if so, the group will do well. Groupthink need not be bad. But it could mean that everyone has adopted an unproductive perspective, and this can lead the group to make bad decisions. Most relevant for our investigation, groupthink—whether good or bad—reduces perspective diversity and stifles the collective ability of the group to find good solutions.[19]

The Last Bite of Ice Cream

In this chapter, we've seen how perspectives organize knowledge. The right perspective can make a difficult problem simple. Hence, we associate clarifying perspectives with genius. They make clear what had been opaque. Perspectives also create superadditive effects. They can be combined to form ever more perspectives.

Thus, if we hope to continue to innovate and reach new understandings, we must encourage the creation of new and diverse perspectives. We should invite physicists into chemistry departments, psychologists into economic departments, and political scientists into business schools. We should include engineers in marketing meetings and marketers in engineering meetings.

And when forming committees and teams, we should choose people who come from different backgrounds and have diverse identities. If not, we're shutting out perspectives. We're slamming the door on potential savants.

Of course, many diverse perspectives won't be useful; they'll make rugged landscapes, not Mount Fujis. We have no guarantee that adding someone different will turn our game of Sum to Fifteen into Tic Tac Toe. A new perspective could just as well transform it into the Unpacking Game. Thomas Edison once optimistically said in the face of the failure of a perspective "we now know a thousand ways not to make the light bulb." Edison rightly saw these failed perspective as a cost worth bearing. So should we, provided we sometimes stumble on a sublime way of seeing a problem.

In this chapter, we've focused on the pragmatic benefits of new perspectives—the scientific understandings, the engineering breakthroughs, the new ways of organizing knowledge. Although those merit our attention and appreciation, they should not cast such a large shadow that we forget the aesthetic joy that diverse perspectives bring. When we look at the periodic table and see the structure of the elements or even when we see how Sum to Fifteen is equivalent to Tic Tac Toe, we experience something sublime.