additional radian takes us to $2e^{2i}$. Shifting this point to the right one unit yields $1+2e^{2i}$. It follows that every point of E can be represented as a polynomial in e^i of the form $z = a_0 + a_1e^i + a_2e^{2i} + \ldots + a_ne^{ni}$ where the a_i are nonnegative integers and n is a positive integer. The set E_1 consists of points with algebraic representations such that $a_0 > 0$. The set E_2 consists of points with representations such that $a_0 = 0$. Shifting any point (z) to the right one unit adds 1 to a_0. Rotating any point (z) one radian counterclockwise increases the coefficient of i in all exponents by one unit.

The properties of E given in this chapter should be apparent from the algebraic description of E. Clearly $E = E_1 \cup E_2$. It now becomes clear why $E_1 \cap E_2 = \emptyset$. If it were the case that a point of E had two distinct polynomial representations (two distinct words in letters T and R) then e^i would be the solution to some polynomial equation. This is impossible as it has been shown that e^i is transcendental (satisfies no polynomial equation with rational coefficients).

5

Statement and Proof of the Theorem

Infinity is where things happen that don't.
—An anonymous schoolboy

Both statement and proof of the theorem of Stefan Banach and Alfred Tarski are straightforward. It is only when one contemplates the conclusion of the theorem that the magic appears. This chapter presents the proof of the theorem in as much detail as possible, without getting bogged down in specific mathematical technicalities which are better left to be investigated outside of this presentation. Mathematical notation is kept to a minimum; however it is assumed that the reader is familiar with the contents of Chapter 3. The interpretation and resolution of this theorem's stunning conclusion are presented in Chapter 6.

Statement of the Banach-Tarski Theorem

A solid ball may be separated into a finite number of pieces and reassembled in such a way as to create two solid balls, each identical in shape and volume to the original.

Formally, this duplication version of the theorem can be stated as follows:

The unit ball $B=\{(x,y,z) : x^2 + y^2 + z^2 \leq 1\}$ *can be partitioned into two sets* B_1 *and* B_2 *such that* $B \sim B_1$ *and* $B \sim B_2$. *Here "~" means "is piecewise congruent to" or "is equidecomposable to."*

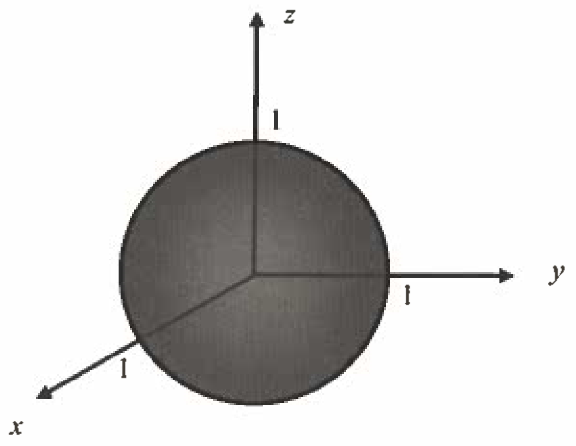

Figure 5.1. The unit ball *B*.

A different, but equally paradoxical version of the theorem asserts that a solid of any shape and volume can be decomposed and reassembled to form another solid of any specified shape and volume. Thus, the theorem is sometimes referred to as the *pea and the sun* paradox. It will be shown that this magnification version follows from the duplication version.

Proof of the Banach-Tarski Theorem

Without loss of generality, we focus on a ball of radius one (the unit ball) centered at the origin of the standard *xyz* Cartesian (rectangular) coordinate system. We can think of the ball as a set of points $B=\{(x,y,z): x^2+y^2+z^2 \leq 1\}$ (Figure 5.1).

We will show that B can be partitioned into a finite number of sets and reassembled in such a way as to form two copies of itself—B_1 and B_2.

The proof is presented in three steps. Step I creates a group of rotations of the unit sphere (surface of the ball). Step II uses these rotations to partition the sphere into subsets possessing a remarkable property known as the *Hausdorff Paradox*. Lastly, Step III extends the above mentioned property from the spherical surface to the solid ball. This will conclude the proof of the duplication version of the Banach-Tarski Theorem.

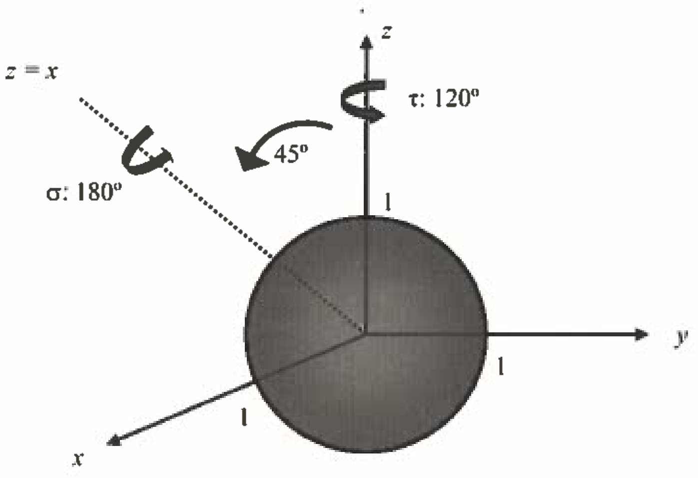

Figure 5.2. Basic rotations τ and σ.

Step I—The Group of Rotations *G*

The *rotation* of a figure (a set of points) about an axis in three-dimensional space is a rigid motion such that each point of the figure moves in a circular path about the axis in a plane perpendicular to the axis. Having said this, we may just assume that a rotation is precisely what we would expect it to be—a turning of the figure in such a way that there is no distortion of the figure.

We begin by defining two basic rotations of the unit sphere. Let τ denote a clockwise rotation of 120° about the z axis and let σ denote a clockwise rotation of 180° about the line $z = x$ in the xz plane. This line passes through the origin and makes an angle of 45° with the z axis (Figure 5.2).

These two basic rotations can be combined sequentially in a countably infinite number of ways to yield infinitely many other rotations. For example, two τ rotations, denoted as $\tau\tau$ or τ^2, corresponds to a rotation of 240° about the z axis. The rotation σ^2 would represent a rotation of 360° about the line $z = x$ and the sphere would be rotated all the way around to its initial position. We call such a rotation the identity rotation and denote it by I. So it follows that $\tau^3 = \sigma^2 = I$.

The single rotation equivalent to τ followed by σ is $\sigma\tau$. Note we use the previously adopted convention of writing the rotations from right to left with the first rotation being listed on the far right and following rotations being listed, in order, from right to left.

Rotations composed of these basic rotations should always be written in reduced, or simplified form. For example, the rotation τ^7 is best written as τ because $\tau^7 = \tau^3\tau^3\tau = II\tau = \tau$. Similarly $\tau^4\sigma^3\tau$ reduces to $\tau\sigma\tau$ because $\tau^4\sigma^3\tau = \tau^3\tau\sigma^2\sigma\tau = I\tau I\sigma\tau = \tau\sigma\tau$.

So, any rotation (except I) composed of a sequence of the basic rotations can be written, in reduced form, as a string of symbols, each of the formτ, τ^2, or σ. We define the *length* of a rotation to be the number of such symbols used to define the rotation. So, the length of the rotation $\sigma\tau^2\sigma\tau$ is four. (We define the length of the rotation I to be zero.) Again, the reader is reminded that these strings should be read from right to left. That is, the single rotation $\tau\sigma\tau^2\sigma$ should be thought of as σ, followed by τ^2, followed by σ, followed by τ.

In total, there is a countably infinite number of possible rotations formed as above and we refer to the collection of all such rotations as the group of rotations G.

The following theorem is key to our discussion. Its proof is given in this chapter's appendix.

The Uniqueness Theorem

Every rotation in G has a unique, reduced form representation.
That is, if two reduced form rotations appear different, they do,
in fact, represent different physical rotations.

Having defined this group of rotations G, we now proceed to partition G into three subsets—G_1, G_2, and G_3. We will specify a rule which assigns each rotation of G to one and only one of the specified subsets. Think of the assignment as a sorting process, placing the rotations, one at a time, into their correct subsets. The assignments are made in order of the lengths of the rotations, beginning with the identity rotation I, then the rotations of length one, two, etc. Of course, it is an infinite process with each member of G ultimately receiving its assignment to one of the three subsets. It is a recursive process in that a rotation's assignment is based, in part, on a previous rotation's assignment.

The assignment process is efficiently summarized by Figure 5.3. The process begins by assigning the identity rotation I to G_1, τ and σ to G_2, and τ^2 to G_3. At this point, all rotations up to length one have been assigned. As indicated, the assignments for all rotations of length two are made. Then it becomes possible to assign the rotations of length three, four, and so on.

	If $\alpha \in G_1$	If $\alpha \in G_2$	If $\alpha \in G_3$
If the leftmost character of α is τ or τ^2	assign $\sigma\alpha$ to G_2	assign $\sigma\alpha$ to G_1	assign $\sigma\alpha$ to G_1
If the leftmost character of α is σ	assign $\tau\alpha$ to G_2	assign $\tau\alpha$ to G_3	assign $\tau\alpha$ to G_1
	assign $\tau^2\alpha$ to G_3	assign $\tau^2\alpha$ to G_1	assign $\tau^2\alpha$ to G_2

Figure 5.3. Partitioning G into G_1, G_2, and G_3.

For example, we have been given the assignments of all length one rotations and now wish to assign the length two rotations to their respective subsets. The rotation $\tau\sigma$ would be assigned to G_3 because σ (having a leftmost character of σ) has already been assigned to G_2 and attaching τ in front obligates us, by the scheme, to place $\tau\sigma$ in G_3. To assign, say, $\sigma\tau^2$, we note that τ^2 (having a leftmost character of τ^2) has already been assigned to G_3. Attaching σ at the left requires we assign $\sigma\tau^2$ to G_1. When we get around to assigning rotations of length three and wish to assign $\tau\sigma\tau^2$ to its subset, we would choose to assign it to G_2 because $\sigma\tau^2$ (having a leftmost character of σ) has already been assigned to G_1 and attaching τ up front requires us to place $\tau\sigma\tau^2$ in G_2.

A few members of the subsets G_1, G_2, and G_3 are shown in Figure 5.4.

Note that as a consequence of the Uniqueness Theorem, the process continues ad infinitum listing all rotations in G, each of which is physically unique.

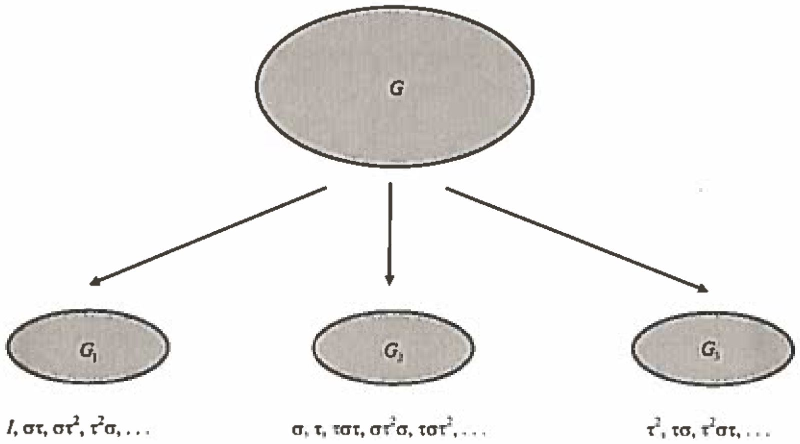

Figure 5.4. Members of the subsets.

A clever way of envisioning the creation and sorting of the rotations was devised by Robert French, currently at the University of Liege, Belgium. He created a Rube Goldberg type machine [French 88, p. 27] which constructs the rotations and sorts them into their respective sets G_1, G_2, or G_3. The schematic shows the device to consist of three hoppers, three appenders, three sorters, and three collection bags named G_1, G_2, and G_3 (Figure 5.5).

The hoppers are to receive copies of previously created and sorted rotations, sent to them by the copiers. Each copied rotation is then passed on to the appender which increases the length of the rotation by one, in all possible ways, by appending σ, τ, or τ^2 at the left end. For example, the rotation $\sigma\tau$ would be appended (on the left) yielding $\tau\sigma\tau$ and $\tau^2\sigma\tau$. There is no other way to increase the length by one (on the left) and still have the rotation represented in simplified form. The sorter then sends each newly created rotation to a specific copier, based on the leftmost character, as indicated by the schematic. The copier makes a copy of the newly created rotation, sends the copy up to the hopper as shown, and drops the newly created rotation into the collection bag below. The process continues, with all three bags shown at the bottom collecting their rotations.

To fire up the process, we begin by placing the identity rotation I in G_1 and placing a copy of I in the hopper on the far left, above G_1. The rest is automatic as the machine does it all. The identity rotation I (of length zero) is appended to form the three rotations of length one—σ, τ, and τ^2. The sorter sends σ and τ to G_2 as shown by the schematic. The rotation τ^2 is sent to G_3. The rotations are copied, and then dropped into the collection bags below. The copies are sent to the hoppers above and the process continues ad infinitum.

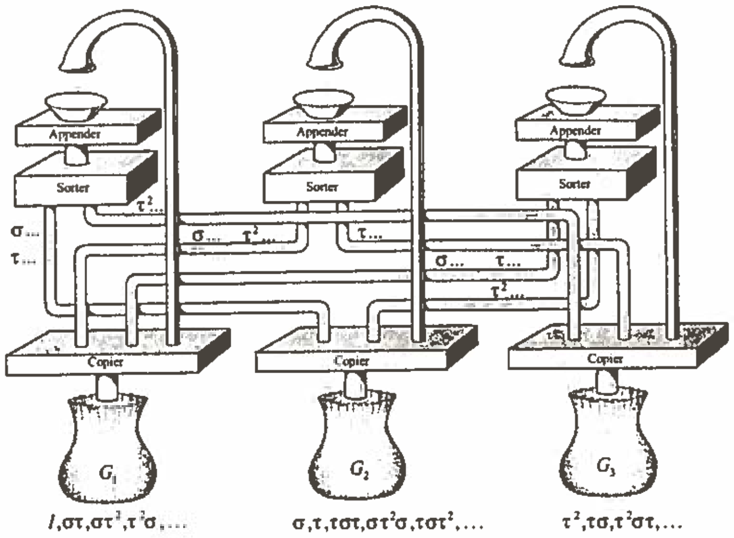

Figure 5.5. Robert French's rotation maker [French 87].
(Reprinted with the permission of Springer-Verlag.)

The three subsets thus created are related in a peculiar yet significant way. Any given rotation in G_1, if immediately followed by the basic rotation τ (denoted by writing τ immediately to the left of the given rotation), yields a rotation in G_2. For example, the rotation $\sigma\tau$ is in G_1 and the rotation $\tau\sigma\tau$ is in G_2. Furthermore, *each and every* rotation in G_2 is related to a specific rotation in G_1. In fact, we could think of G_2 as being the set of all rotations of G_1, each followed by τ (with τ attached on the left).

Mathematically we write $\tau G_1 = G_2$. Similarly, we can verify $\tau^2 G_1 = G_3$. And finally, it could be verified that following each rotation in G_1 by σ (attaching σ on the left) would give us every rotation in $G_2 \cup G_3$. We write $\sigma G_1 = G_2 \cup G_3$.

Summarizing, the following relations hold among the subsets:

$$\tau G_1 = G_2$$
$$\tau^2 G_1 = G_3$$
$$\sigma G_1 = G_2 \cup G_3$$

A formal proof of the first of these equations is presented in this chapter's appendix. Proofs of the second and third equations are omitted as they are analogous to the proof of the first.

To the reader, these relations may seem insignificant; yet, we carry these properties into Step II of the proof. They will induce a remarkable partition of the unit sphere, the Hausdorff Paradox, which some would consider every bit as remarkable as the Banach-Tarski Theorem itself. In Step III of the proof, we use the partition of the sphere established in Step II to induce the partition of the ball as specified by the Banach-Tarski Theorem.

Step II—Partitioning the Unit Sphere S into Two Copies of Itself

In this step of the proof, we will take the unit sphere, the surface or *skin* of the unit ball, separate and reassemble it into two sets, each of which is identical in shape to the original unit sphere. To say the least, this is paradoxical and would be considered by some as stunning as the Banach-Tarski Theorem itself. In fact, once we have the unit sphere partitioned in Step II, it becomes a rather small step to extend the process to the ball, as will be done in Step III.

There are two essential mathematical tools we need in Step II. First, we need the three relations among the subsets of G, as previously established. Second, we will need to invoke the Axiom of Choice. The arguments presented in Step II are straightforward and take us through the Hausdorff Paradox to finally arrive at the *two spheres from one* conclusion.

We begin by noting that every rotation in G of the unit sphere has two *poles*. A pole is a point which remains fixed for a given rotation. For example, the rotation τ which rotates the unit sphere 120° about the z axis, has the two poles $(0, 0, 1)$ and $(0, 0, -1)$, much like the north and south poles of the earth. So every rotation in G has its axis and at the ends of the axis we find the two poles associated with the rotation. In that every rotation in G has two poles and there is a countable infinity

of rotations in G, we know there are countably many poles on the unit sphere associated with G. Let P denote the set of all poles associated with G on the unit sphere. Let $S-P$ denote all other points on the sphere. Clearly there are infinitely many points in both P and $S-P$. Yet, there are far more points in $S-P$ than P. Recall from Chapter 3 that there are *orders of infinity* and, in a sense, some infinities are larger than others. To put it another way, if a point were randomly chosen from the surface of the sphere, it would almost certainly belong to $S-P$, and not to P. The fractional part of the sphere corresponding to P is infinitesimally small.

Every point in $S-P$ can be thought of as being connected to a countable infinity of other points in $S-P$ via the rotations in G. If any two points in $S-P$ are so connected, we say they belong to the same *orbit*. There is an uncountable infinity of such orbits which make up $S-P$. We now use the Axiom of Choice to select one point from each of these orbits and collectively define these points so chosen as the set C. There is no other way to do this, other than using the Axiom of Choice. These orbits have no quality which allows us to define one member of each. We must simply assume that such a choice is possible and we create the set C by choosing exactly one point from each of these orbits. The set C of points is called the *choice set* as it was formed using the Axiom of Choice.

Summarizing, the surface S consists of a set of poles P, and the remaining points $S-P$. These remaining points, $S-P$, naturally separate into infinitely many orbits. We invoke the Axiom of Choice to create a set C by choosing one point from each of these orbits.

We make the following observations with respect to C:

1. C is uncountably infinite.
2. C and P have no points in common.
3. No point in C can be rotated to any other point in C by one of the rotations in G.
4. If every point in C were to be rotated by every rotation in G, we would ultimately get every point in $S-P$.

Keep in mind that the set $S-P$ represents almost every point on the sphere S, with the exception of the poles P. We will now focus on the set $S-P$, and deal with the poles later.

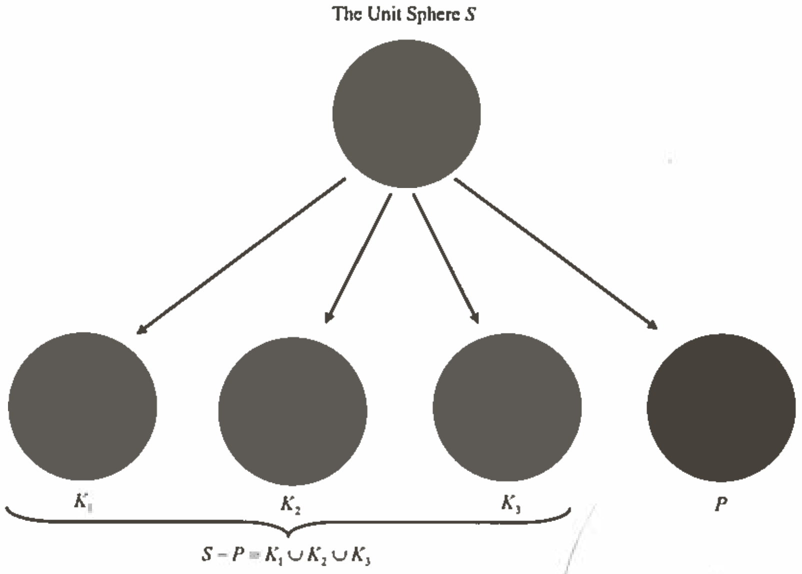

Figure 5.6. The Hausdorff partition of S.

If we apply every rotation of G_1 to C, we will get a set of points on the sphere which we can denote as G_1C. Similarly, if we apply every rotation of G_2 to C, we get another set of points on the sphere which we denote as G_2C. Lastly, we apply the rotations of G_3 to C obtaining G_3C. These three sets of points, G_1C, G_2C, and G_3C form a partition of $S - P$. That is, these three sets are disjoint (have no points in common) and include all points of $S - P$. With this in mind, we have now partitioned the entire sphere into four disjoint sets—G_1C, G_2C, G_3C, and the set of poles P. To make the notation a bit clearer, let $K_1 = G_1C$, $K_2 = G_2C$, and $K_3 = G_3C$. Then we have the sphere, S, partitioned into the four disjoint sets of points K_1, K_2, K_3, and P. We can write $S = K_1 \cup K_2 \cup K_3 \cup P$.

Figure 5.6 illustrates how we have the sphere partitioned. Remember, at this point we are partitioning the sphere S, and not the ball B. We are partitioning the surface of the ball.

By definition of K_1 and K_2, we see that if all points of K_1 are rotated by the basic rotation τ, we get K_2. So K_1 and K_2 are congruent and we write $K_1 \cong K_2$. We also note that K_3 is obtained from K_1 by rotating K_1 by τ^2. So $K_1 \cong K_3$. And finally, rotating K_1 by σ gives us $K_2 \cup K_3$.

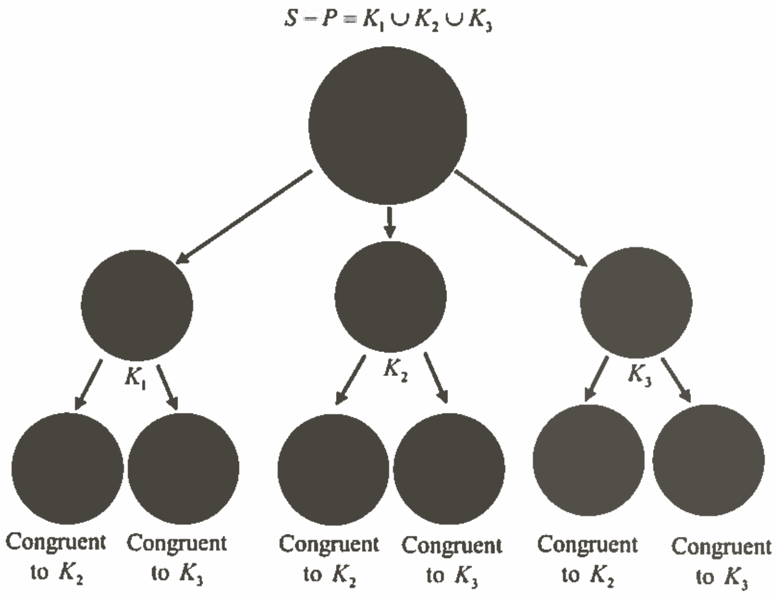

Figure 5.7. Six subsets of $S - P = K_1 \cup K_2 \cup K_3$.

Summarizing: $K_1 \cong K_2 \cong K_3 \cong K_2 \cup K_3$

This last statement is the Hausdorff Paradox. It may not at first appear paradoxical, but consider that the sphere S is made up, almost entirely, of points in $K_1 \cup K_2 \cup K_3$. Since $K_1 \cong K_2 \cong K_3$, we can conclude that each K is roughly *one third* of the entire sphere (or at least one third of $K_1 \cup K_2 \cup K_3$). But if each K is congruent to $K_2 \cup K_3$, we could also conclude that any given K, say K_1, is *one half* of the entire sphere. So which is it, one third or one half? This *half-third* dilemma is the Hausdorff Paradox.

As we shall soon see, the Hausdorff Paradox allows us to separate the sphere S into a finite number of pieces which can be reassembled to form two spheres, each identical in shape to the original.

We will split $S - P = K_1 \cup K_2 \cup K_3$ into two copies of itself and deal with the problem of the poles afterward (Figure 5.7). The technique suggested by Robert French uses $K_2 \cup K_3$ as a *cutting template* which can be placed directly onto each K, partitioning each into two sets, one of which being congruent to K_2 and the other being congruent to K_3. This template will fit perfectly as $K_2 \cup K_3$ is congruent to each K.

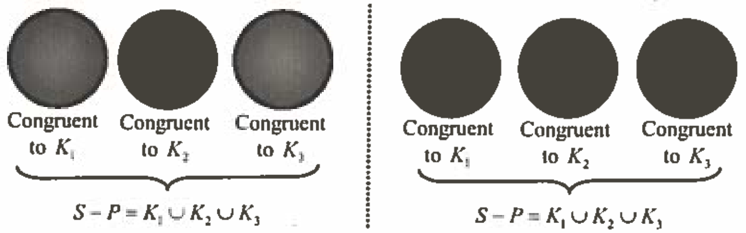

Figure 5.8. Two copies of $S - P = K_1 \cup K_2 \cup K_3$.

We have now partitioned the sphere S, except for a countable and relatively small set of points P, into six disjoint subsets, each of which is congruent to either K_2 or K_3, as illustrated. But since $K_1 \cong K_2 \cong K_3$, we can say that any given K is congruent to each of the other two. So, we can rearrange these six subsets and declare the congruencies as shown in Figure 5.8.

We've now done a most remarkable thing. We have partitioned and reassembled the set of points $S - P = K_1 \cup K_2 \cup K_3$, which represents almost all points on the unit sphere, in such a manner as to form two identical copies of itself. In that $K_1 \cup K_2 \cup K_3$ represents virtually the entire unit sphere S (except for the set of poles P), we are almost done. The fly in the ointment is P.

As it turns out, the problem of the poles is easily disposed of. We can use the set of poles P from the original sphere S to plug the holes representing the missing poles in one of the two copies of $S-P$. This gives us one complete copy of S. But how can we fill the holes in the other copy of $S-P$ to give us the second copy of the complete sphere? Recall the discussion of piecewise congruence (equidecomposability) given in Chapters 3 and 4. In Chapter 4 we showed that the surface of a sphere missing a countable number of points (holes) is equidecomposable to the complete sphere. We apply this technique, that of shifting from infinity, to the point set $S-P$ to give us the second complete copy of the sphere S (Figure 5.9).

This completes the second sphere and we now have two complete copies of the sphere S made from S alone. Call these copies S_1 and S_2 (Figure 5.10).

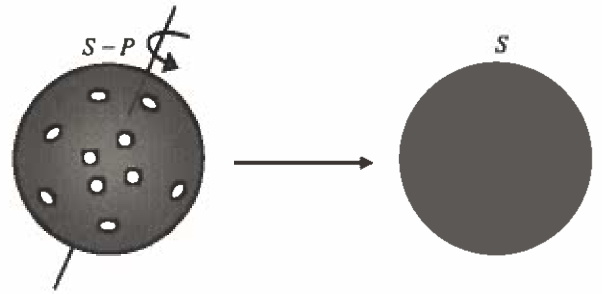

Figure 5.9. Shifting from infinity to show $S - P \sim S$.

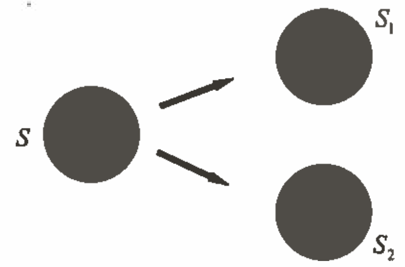

Figure 5.10. Two spheres from one.

Step III—Partitioning the Unit Ball B into Two Copies of Itself

The difficult part of the proof is behind us and it is a rather simple matter to extend the duplication of the sphere S to the duplication of the ball B.

To each subset of S, say K_1 as defined in the previous section, we can associate an inward *thickening*, extending from the surface of the sphere up to, but not including, the center (0, 0, 0). We denote this thickened section of the ball as $\overline{K}_1$. Putting it another way, we can think of K_1 as the projection on S of all points belonging to $\overline{K}_1$. Of course we can thicken all subsets of S in similar fashion.

K_1 thickens to $\overline{K}_1$
K_2 thickens to $\overline{K}_2$
K_3 thickens to $\overline{K}_3$
P thickens to $\overline{P}$

Figure 5.11. Thickening S to form $\overline{S}$.

Similarly, S_1 thickens to $\overline{S_1}$ and S_2 thickens to $\overline{S_2}$. In fact, we could thicken the entire sphere S to $\overline{S}$, the *punctured ball*, as illustrated in Figure 5.11.

In the previous section, we established various relations among the above mentioned subsets of S. Most notably, we established the Hausdorff Paradox—$K_1 \cong K_2 \cong K_3 \cong K_2 \cup K_3$, and the creation of the two spheres, S_1 and S_2, from one sphere S. But, intuitively, all of these congruencies hold as well for the thickened subsets. So we can just as easily claim $\overline{K_1} \cong \overline{K_2} \cong \overline{K_3} \cong \overline{K_2 \cup K_3}$ and that $\overline{S}$ can be separated and rearranged into $\overline{S}_1$ and $\overline{S}_2$, each of which is piecewise congruent to $\overline{S}$. Since $\overline{S}$ is the punctured unit ball, and represents all points of the unit ball B except the center, we have virtually proven the Banach-Tarski Theorem, except for the matter of dealing with the center (0, 0, 0). If we can plug the center holes in each of the two punctured balls $\overline{S}_1$ and $\overline{S}_2$, then we will have created two solid balls from one and our proof will be complete.

We use the center (0, 0, 0) of the original ball B to plug the center hole in $\overline{S}_1$ giving us one complete copy of the ball. We plug the hole in the second punctured ball $\overline{S}_2$ by the now familiar technique of shifting from infinity. That is, we consider the center of the punctured ball $\overline{S}_2$ as one of countably many carefully chosen points on a circle completely contained in $\overline{S}_2$. Shifting the points from infinity plugs the center hole and we now have our second complete copy of B. Call the two copies B_1 and B_2 (Figure 5.12). This completes our proof of the Banach-Tarski Theorem.

We have just proven the *duplication* version of the Banach-Tarski Theorem. The *magnification* or *strong version* of the theorem may be even more striking.

> *If A and B are any two bounded three-dimensional sets with nonempty interiors, then* $A \sim B$.

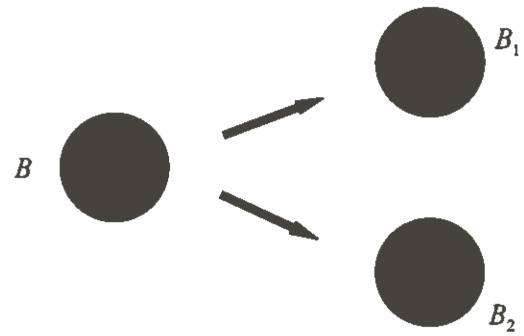

Figure 5.12. Two balls from one.

If, say, A is a ball the size of a pea and B is a ball the size of the sun, then this version of the Banach-Tarski Theorem asserts that the pea (A) can be decomposed into a finite number of pieces and reassembled to form a ball the size of the sun (B) (Figure 5.13).

Interestingly, this version of the theorem does not require that A and B be spherical in shape. In fact, their shapes need not be identical. So, we could just as easily call this version of the theorem the *mosquito and the elephant* paradox!

The magnification version of the theorem follows directly from the duplication version. Before giving the proof, we must state a preliminary theorem—the *Banach-Schröder-Bernstein Theorem*. Its proof is given in this chapter's appendix.

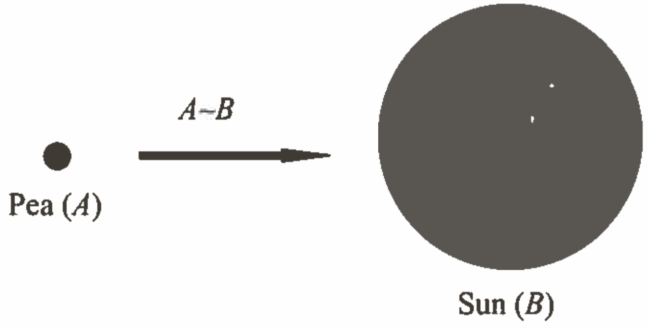

Figure 5.13. The pea and the sun paradox.

Let A and B be two bounded three-dimensional sets, each having nonempty interiors. Assume A is piecewise congruent to a subset of B. We write $A \preceq B$. *Further assume B is piecewise congruent to a subset of A and write* $B \preceq A$. *Then it follows that A is piecewise congruent to B; that is* $A \sim B$.

(This is somewhat analogous to two real numbers a and b being such that $a \le b$ and $b \le a$. It would follow that $a = b$.)

Now we prove the magnification version of the Banach-Tarski Theorem. We will show that for any two bounded three-dimensional sets A and B having nonempty interiors, $A \preceq B$. Since we could use a similar argument to show $B \preceq A$, it will follow that $A \sim B$.

Choose solid balls V and W such that A is contained in V and W is contained in B. Duplicate W (by Banach-Tarski duplication) until there are enough copies of W than can be translated and overlapped so as to completely cover V. Assume n copies are required. Note that n overlapping copies of W are equidecomposable to a subset of n disjoint copies of W. Furthermore, the Banach-Tarski Theorem guarantees that the n disjoint copies of W are equidecomposable to W.

Consequently,

$$A \subseteq V \subseteq n \text{ overlapping copies of } W \preceq n \text{ disjoint copies of } W \sim W \subseteq B,$$

which establishes $A \preceq B$. Since, by similar means, we can show $B \preceq A$, it follows by the Banach-Schröder-Bernstein Theorem that $A \sim B$. This concludes the proof of the magnification version of the Banach-Tarski Theorem.

Appendix

Proof of the Uniqueness Theorem

Every rotation in G has a unique, reduced form representation. That is, if two reduced form rotations appear different, they do, in fact, represent different physical rotations.

To begin, we note that every reduced form rotation in G (other than I, σ, τ, and τ^2) can be expressed in at least one of the four forms

$$\alpha = \tau^{P_1}\sigma\tau^{P_2}\sigma \cdots \tau^{P_n}\sigma$$

$$\beta = \sigma\tau^{P_1}\sigma\tau^{P_2} \cdots \sigma\tau^{P_n}$$

$$\gamma = \tau^{P_1}\sigma\tau^{P_2}\sigma \cdots \sigma\tau^{P_n}$$

$$\delta = \sigma\tau^{P_1}\sigma\tau^{P_2} \cdots \sigma\tau^{P_n}\sigma$$

where $n \ge 1$ and each exponent P_i is one or two. (For γ, we must have $n > 1$.)

We begin by giving the argument [Osofsky and Adams 78, p. 504] that no reduced form rotation of the form α can equal the identity matrix I. From this we will show (as given by [Stromberg 79, p. 154]) that no reduced form rotation of the form β, γ, or δ can equal I. After having established that no reduced form rotation (other than I itself) can equal I, we will show that the reduced form representation of any rotation in G is unique.

$$\text{Let } \tau = \begin{bmatrix} -\frac{1}{2} & \frac{-\sqrt{3}}{2} & 0 \\ \frac{\sqrt{3}}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \sigma = \begin{bmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

be the matrix representations of τ and σ as given in Chapter 3. Then

$$\tau^P\sigma = \frac{1}{2}\begin{bmatrix} 0 & \pm\sqrt{3} & -1 \\ 0 & 1 & \pm\sqrt{3} \\ 2 & 0 & 0 \end{bmatrix}$$

where we take $+\sqrt{3}$ if $P = 1$ and $-\sqrt{3}$ if $P = 2$. It follows by mathematical induction (or empirical verification) that

$$\alpha = \frac{1}{2^n}\begin{bmatrix} m_{1\,1} & m_{1\,2}\sqrt{3} & m_{1\,3} \\ m_{2\,1}\sqrt{3} & m_{2\,2} & m_{2\,3}\sqrt{3} \\ m_{3\,1} & m_{3\,2}\sqrt{3} & m_{3\,3} \end{bmatrix}$$