CHAPTER ONE

Definitions and Examples of Groups

1A

From the abstract, axiomatic point of view that prevails today, one can argue that group theory is, in some sense, more primitive than most other parts of algebra and, indeed, the group axioms constitute a subset of the axiom systems that define the other algebraic objects considered in this book. Things we learn about groups, therefore, will often be relevant to our study of modules, rings, and fields. In addition, group theory has considerable indirect connection to these other areas. (The most striking example of this is probably the use of Galois groups to study fields.) It is largely for these reasons that we begin this book on algebra with an extensive study of group theory. (If the whole truth were told, the fact that the author's primary research interest and activity are in group theory would be seen as relevant, too.)

The subject we call "algebra" was not born abstract. In its youth, algebra was the study of concrete objects such as polynomials, rather than of things defined by axiom systems. In particular, early group theory was concemed with groups of mappings, known as "transformation groups." (In the early literature, for instance, the elements of a group were referred to as its "operations.")

For at least two reasons, we begin our study of group theory by (temporarily) adopting this nineteenth-century point of view. First, mappings of one kind or another are ubiquitous throughout algebra (and most of the rest of mathematics, too) and so it makes sense to begin with them. Furthermore, some of the most interesting examples of groups are best constructed and visualized as transformation groups.

We begin our study of mappings with some notation and definitions. (It is this author's belief that mathematics at its best consists of theorems and examples. Definitions are often dull, although they are a necessary evil, especially near the beginning of an expository work. We pledge that the balance of theorems and examples versus definitions will become more favorable as the reader progresses through the book.)

The notation $f : A \to B$ means that f is a mapping (that is, a function) from the set A to the set B. (If either A or B is empty, there are no mappings, and so the existence of f implies that A and B are both nonempty.) The set A is the *domain* of f and B is the *target*. The *image* or *range* of f is denoted $f(A)$. It is the subset

$$\{f(a) \mid a \in A\} \subseteq B .$$

The map f is *onto* or *surjective* if its image is all of the target B. It is *one-to-one* or *injective* if distinct elements of A map to distinct elements of B. If f is both injective and surjective, we say that it is a *bijection*.

Note that we have not specified whether our functions "act on" the right or the left, although in writing "$f(a)$" above, we seem to be implying an action on the left. Our point of view on this question is perhaps slightly unconventional, but it will, we hope, be quite comfortable for the student.

We maintain that functions do not "act" on any particular side, and so it is permissible to write $f(a)$ or $(a)f$, whichever is more intelligible in a given context. Both notations mean precisely the same thing: namely, the result of applying f to a. What is the cost of this freedom of notation? Confusion could enter when two mappings are composed. For instance, if

$$f : A \to B \quad \text{and} \quad g : B \to A ,$$

does fg mean "f then g" or does it mean "g then f"? Proponents of "action on the right" would say the former and "leftists" would choose the latter. Our convention throughout this book is that fg always means "f then g." This does not, however, constrain us to write the mappings on the right, but in a setting in which function composition is important, it will usually enhance clarity to do so, and so we shall. According to our notation, therefore, we have

$$(a)(fg) = ((a)f)g ,$$

but it would be equally correct (though more confusing) to write

$$(fg)(a) = g(f(a)) .$$

For any nonempty set A, we write i_A (or sometimes just i) to denote the identity map. Thus $i_A(a) = a$ for all $a \in A$, and it is clear that

$$i_A f = f \quad \text{and} \quad g i_A = g$$

for arbitrary maps $f : A \to B$ and $g : B \to A$. Note that the associative law for maps is a triviality. If $f : A \to B$, $g : B \to C$, and $h : C \to D$, then $f(gh)$ and $(fg)h$ are equal, since both are the map obtained by doing first f, then g, and then h.

(1.1) LEMMA. *Let $f : A \to B$.*

a. *f is injective iff there exists $h : B \to A$ such that $fh = i_A$.*
b. *f is surjective iff there exists $g : B \to A$ such that $gf = i_B$.*

c. *If f is a bijection, then the maps g and h above are uniquely determined and equal.*

Proof. Suppose f is injective. Fix an element $a \in A$ and define $h : B \to A$ by

$$(b)h = \begin{cases} a & \text{if } b \notin (A)f \\ x & \text{if } (x)f = b \text{ for some } x \in A. \end{cases}$$

Note that by the injectivity of f, there is at most one element $x \in A$ such that $(x)f = b$. Also, the mapping h is unambiguously defined, since the two cases are mutually exclusive and exhaust the possibilities.

Conversely, if $h : B \to A$ and $fh = i_A$, we wish to show that f is injective. Suppose $(x)f = (y)f$. Then $x = (x)fh = (y)fh = y$, as required.

Now suppose f is surjective. For each $b \in B$, choose $a \in A$ with $(a)f = b$, and once this choice is made, define $g : B \to A$ by $(b)g = a$. Clearly $gf = i_B$. Conversely, suppose $gf = i_B$. Then

$$B = (B)i_B = (B)gf \subseteq (A)f,$$

since $(B)g \subseteq A$. It follows that $(A)f = B$, as required.

Finally, assume f is a bijection so that maps h and g as in parts (a) and (b) exist. Then

$$g = gi_A = g(fh) = (gf)h = i_Bh = h.$$

In particular, g is uniquely determined, since it must equal any valid choice for h. Similarly, h is uniquely determined. ■

In our proof that g exists when f is assumed to be surjective, we needed to make some choices; the definition of g was not forced. In fact, if B is an infinite set, we would need to make infinitely many choices. Some mathematicians feel that a definition that requires infinitely many choices is somewhat suspect, and they have created an additional axiom of set theory, called the "axiom of choice," to deal with this situation. (It is precisely this axiom to which we implicitly appealed in the preceding proof.) It has been proved that the axiom of choice is not a consequence of the rest of set theory, but that it can be assumed without introducing any contradictions into mathematics. Most mathematicians (except those working in set theory itself) freely assume and use the axiom of choice whenever it is convenient to do so, and we will follow that policy here. We shall have a little more to say about the axiom of choice in Chapter 11, when we prove Zorn's lemma.

In Lemma 1.1(a) we say that h is a *right inverse* of f and in (b), that g is a *left inverse*. In the case that f is a bijection, the unique left and right inverse of f is simply called the *inverse* of f and it is denoted f^{-1}. It is interesting to observe the striking symmetry in the statement (though not in the proof) of Lemma 1.1. The conditions that f has right and left inverses are essentially mirror images, although there is no such relationship apparent between the equivalent conditions that f be injective or surjective, respectively. We shall see more of this "duality" later.

1B

Let X be an arbitrary nonempty set. We denote by $\mathrm{Sym}(X)$ the set of all bijections from X to itself. (These bijections on X are also called *permutations*, and if X is finite, this is, in fact, the more common term.) The object $\mathrm{Sym}(X)$ is called the *symmetric group* on X. (Note that if X is finite, containing n elements, say, then $\mathrm{Sym}(X)$ consists of precisely $n!$ permutations.)

(1.2) COROLLARY. *Let* $G = \mathrm{Sym}(X)$.

a. $i_X \in G$.
b. g^{-1} *exists and lies in* G *for each* $g \in G$.
c. $gh \in G$ *for each* $g, h \in G$.

Proof. Part (a) is immediate, and (b) follows from Lemma 1.1. It is not hard to prove (c) directly, but we prefer the following argument. Given $g, h \in G$, we see that

$$(gh)(h^{-1}g^{-1}) = i = (h^{-1}g^{-1})gh\,,$$

and thus gh has a left and right inverse. It follows by Lemma 1.1 that $gh \in \mathrm{Sym}(X)$. ■

(1.3) DEFINITION. Let X be any set. A *permutation group* on X is any nonempty subset $G \subseteq \mathrm{Sym}(X)$ such that

i. $g^{-1} \in G$ for each $g \in G$ and
ii. G is closed under function composition.

Note that for $g \in G$, there is no question that the mapping g^{-1} exists and lies in $\mathrm{Sym}(X)$. The point of condition (i) is that g^{-1} actually lies in the subset G. Conditions (i) and (ii), together with the assumption that $G \neq \varnothing$, imply that $i_X \in G$, and so this need not be assumed.

Given a nonempty set G of mappings on some set X, perhaps the best strategy for showing that G is a permutation group is first to verify that each element $g \in G$ has both a left and a right inverse in G. From this it follows that $G \subseteq \mathrm{Sym}(X)$ and this condition need not be verified separately. All that remains, then, is to check the closure condition.

Some obvious examples of permutation groups are $\mathrm{Sym}(X)$ and the singleton set $\{i_X\}$ for arbitrary nonempty X. We devote the next few pages to descriptions of several more interesting examples.

Consider a square $ABCD$ and let $X = \{A, B, C, D\}$ be its vertex set (see Figure 1.1). Within $\mathrm{Sym}(X)$, let G denote the set of permutations of X that can be realized by a physical motion of the square through 3-space. For instance, imagine the square being rotated 90° counterclockwise about its center. This brings A to the position formerly occupied by D, and so on, and the associated element of G is the map $g : A \mapsto D \mapsto C \mapsto B \mapsto A$. Similarly, a "flip" about the vertical axis v yields the map $h : A \mapsto B \mapsto A;\ C \mapsto D \mapsto C$. If we first do the rotation and then the flip, the result is the same as a flip about the diagonal axis d, and the corresponding element of G is precisely the composition $gh : A \mapsto C \mapsto A;\ B \mapsto B;\ D \mapsto D$.

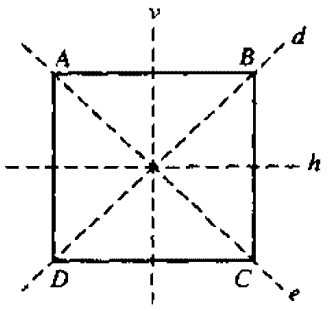

Figure 1.1

The reader should be warned of a possible source of confusion here. One might be tempted to say that the effect of the 90° counterclockwise rotation is that position A is now occupied by vertex B and so the associated map ought to take A to B. With this scheme, this rotation would yield the permutation g': $A \mapsto B \mapsto C \mapsto D \mapsto A$. The flip about axis v still yields h, but the combined operation, the flip about d, does *not* yield the composition $g'h$. (In fact, it yields hg'.) We conclude that if we want function composition, which is the group operation, to correspond to "composition of rotations," we should use the convention given earlier: The mapping g associated with a physical motion satisfies $(x)g = y$ if x goes to the position where y was.

A few moments of reflection should convince the reader that G is a group and that it contains exactly eight elements. One way to obtain the count is to focus on a particular edge, say AB. After a rotation, AB will coincide with the original position of one of the four sides, and in that position it can be in either of two orientations. This yields a total of eight alternatives for how to place AB, and each of these uniquely determines the locations of all four corners.

The eight elements of G are the permutations induced by the four "flips" (about axes h, v, e, and d)) and the four "planar rotations" of 0°, 90°, 180°, and 270°. The standard name for our group G is the *dihedral group* of order 8 and it is denoted D_8. The word "dihedral," meaning two-sided, refers to the front and back sides of the square, which are interchanged by half the elements of the group.

In general, the *order* of a group G is its cardinality (number of elements), and we write $|G|$ to denote this number. (We also write $|X|$ to denote the cardinality of any set X, although the word "order" is generally reserved for groups.)

If, instead of a square, we had started with a regular n-gon ($n \geq 3$), the resulting dihedral group would be D_{2n} of order $2n$. As with D_8, half the elements of D_{2n} correspond to flips and half (counting the identity) correspond to plane rotations. We should mention that many users of group theory write D_n to refer to the dihedral group of order $2n$, whereas most group theorists use the notation we have presented here.

As a further source of interesting examples of groups, let us move up to three dimensions. Consider the vertex set X of a regular polyhedron. The permutations of X induced by physical rotations of the object form a group called the *rotation group* of the object. A usually larger group is the *full group of symmetries*, which consists of all permutations of X realizable by geometric congruences of the polyhedron.

Consider the case of a cube. The full group of the symmetries includes the "antipodal map" τ, which reflects each vertex through the center of the cube. (Thus $(A)\tau = F$ and $(C)\tau = H$ in Figure 1.2, for instance.) The reader should check that τ does not correspond to any rotation. Note that there is no antipodal map for the regular tetrahedron, although it is true for that figure too that there are symmetries that are not rotations. In fact, in this case, the full group of symmetries is the full symmetric group on the vertex set, of order $4! = 24$.

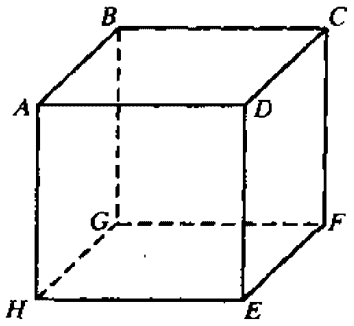

Figure 1.2

Let us compute the order of the rotation group R of a cube. After a rotation, face $ABCD$ can coincide with any of the six faces of the original cube, and in each location, it can have any of the four rotational orientations. It follows that $|R| = 6 \cdot 4 = 24$. The full group of symmetries S, on the other hand, has order 48. (We leave this as an exercise.) What are the 24 symmetries that are not rotations? Among these are the reflections in the nine planes of symmetry of the cube. These planes of symmetry are of two types: six that contain four vertices (for instance, the planes determined by B, D, G, E or by A, E, B, F) and three that are parallel to faces of the cube. A tenth nonrotational symmetry is the antipodal map τ. The remaining 14 nonrotational symmetries are rather hard to visualize and we shall not discuss them further now. The product (composition) of each of the nine reflections with τ yields a rotation of order 2. (The *order* of an element g of a group, denoted $o(g)$, is the least positive integer n, if it exists, such that g^n is the identity. If there is no such n, we say that g has *infinite order* and write $o(g) = \infty$. Elements of order 2 are usually called *involutions*.) A good exercise is to count how many elements of each order there are in the rotation group of a cube.

We shall briefly mention three more examples before proceeding with our study of groups in general. The first example is the "general linear" group $GL(V)$, where V is a vector space. This is the group of all nonsingular (invertible) linear transformations of V. It should be obvious that $GL(V) \subseteq \mathrm{Sym}(V)$ is, in fact, a group.

Next we consider the "affine group" of the line. This is the set of all mappings on the real numbers $\mathbb{R}$ that are of the form $x \mapsto ax + b$, where $a, b \in \mathbb{R}$ and $a \neq 0$. The reader should check that this really is a group.

Our final example is the group associated with the Rubik cube puzzle. (We assume that the reader has some familiarity with this object.) Of the 54 colored squares on the surface of the cube, six may be viewed as never moving from their

initial positions (although they do rotate). In other words, if we start with the red face on top and the green face in front, then all interesting cube moves can be made while keeping the red center square on top and the green center square in front. (Of course, this prohibits rotations of the entire cube, but such rotations are not strictly necessary for solving the puzzle.) Now let G be the group of those permutations on the $54 - 6 = 48$ colored squares that can be realized by some sequence of cube twists. How does one compute $|G|$?

As a first approximation, consider disassembling the cube. When this is done, one obtains eight small "corner" cubes having three colored faces each, and 12 small "edge" cubes having two colored faces each. The six face-centers remain attached to one another, and we view them as being fixed in space. To reassemble the cube, we can permute the corner cubes in 8! ways and the edge cubes in 12!. In addition, each corner cube can occur in three different orientations and each edge cube in two different orientations. This yields a total of $8! \cdot 12! \cdot 3^8 \cdot 2^{12}$ ways to reassemble the cube. It turns out (although it is not trivial to prove) that only one-twelfth of these are attainable via legal moves without doing violence to the puzzle. The order of the Rubik cube group, then, is given by

$$|G| = \left(\tfrac{1}{12}\right) 8! \cdot 12! \cdot 3^8 \cdot 2^{12} = 43{,}252{,}003{,}274{,}489{,}856{,}000.$$

1C

Throughout most of the nineteenth century, the word "group" meant "permutation group." We are now ready to give the modern definition, attributed to the English mathematician Arthur Cayley. Recall that a *binary operation* on a set G is a rule that assigns to each ordered pair of elements $x, y \in G$ another element of G. If $\circ$ is a binary operation, we write $x \circ y$ to denote the result of applying this rule to x and y, and we say that $\circ$ is *associative* if $x \circ (y \circ z) = (x \circ y) \circ z$ for all elements $x, y, z \in G$.

(1.4) DEFINITION. A *group* is a set G together with an associative binary operation $\circ$ defined on G such that there exists $e \in G$ with the following properties:

i. For each $x \in G$, $x \circ e = x = e \circ x$.
ii. For each $x \in G$, there exists $y \in G$ such that $x \circ y = e = y \circ x$.

Note that the "closure" condition that $x \circ y \in G$ whenever $x, y \in G$ need not be stated explicitly, since it is subsumed in the assumption that $\circ$ is a binary operation on G. Observe also that any permutation group is a group with respect to the operation of function composition. In addition to permutation groups, Definition 1.4 allows such objects as the additive group of the integers, the multiplicative group of the positive rationals, and the groups of $n \times n$ nonsingular matrices over fields (with respect to matrix multiplication).

We shall usually follow the custom of suppressing the symbol "$\circ$" and writing xy in place of $x \circ y$. The operation is usually called "multiplication," and xy is referred to as the "product" of x and y.

(1.5) LEMMA. *Let G be a group. Then, for $a, b \in G$, there exist unique elements $x, y \in G$ such that*

$$ax = b \quad \text{and} \quad ya = b.$$

In particular, the element e is unique, and for each $x \in G$, the element y of Definition 1.4(ii) is unique.

Proof. Choose z such that $az = e = za$. Now

$$a(zb) = eb = b,$$

and so we can take $x = zb$.

For uniqueness, if $ax = ax'$, we have

$$x = ex = zax = zax' = ex' = x',$$

as required. The existence and uniqueness of y are proved similarly. ■

In a permutation group, the unique element satisfying condition (ii) of Definition 1.4 is, of course, the identity map i. By analogy, this special element in an abstract group is called the *identity element* of the group and it is customarily denoted 1. The reader should note that the identity of a permutation group is defined by what it *is* (a particular mapping), whereas the identity of an abstract group is defined by how it *behaves* with respect to the group operation. Similarly, the element y of a permutation group that satisfies condition (ii) with respect to x is the inverse map, x^{-1}, and by analogy, in an abstract group, y is said to be the *inverse element* of x and the notation x^{-1} is used in this case, too.

In fact, the conditions of Definition 1.4 are more stringent than they really need to be.

(1.6) THEOREM. *Let G be a set with an associative multiplication and suppose there exists $e \in G$ with the following properties:*

i. *$xe = x$ for all $x \in G$ and*
ii. *for each $x \in G$, there exists $y \in G$ with $xy = e$.*

Then G is a group.

Proof. Let $x \in G$ and choose y according to property (ii). It suffices to show that $ex = x$ and $yx = e$.

Use property (ii) to find $z \in G$ with $yz = e$. We have

$$x = xe = x(yz) = (xy)z = ez,$$

and so

$$yx = y(ez) = (ye)z = yz = e,$$

as required. Now

$$ex = (xy)x = x(yx) = xe = x,$$

and the proof is complete. ■

We should mention that the "elementwise" calculations in the preceding proof are not typical of most of algebra. The proof of Theorem 1.6, in fact, could almost serve as a model of what algebra is not, or at least should not be, in the opinion of the author.

One way to describe the operation (multiplication) in an abstract group G is via a *multiplication table*. This is a square array, with rows and columns labeled by the elements of G and where the position in row x and column y is occupied by the element xy. Generally, it is neither useful nor practical to actually write down a multiplication table for G, but we can think of G as being defined by such a table.

One of the advantages of thinking about groups abstractly, as in Definition 1.4, is that it allows us to see that certain groups, perhaps defined very differently, are essentially "the same." Suppose, for example, that we rename all the elements of some group G, and that we use these new names to relabel the rows and columns of the multiplication table of G and also to replace the entries in the table. The result will be the multiplication table of a group that is not, in any essential respect, different from G. We can make this notion of "essential sameness" more precise, as follows.

(1.7) DEFINITION. Let G and H be two groups and suppose $\theta : G \to H$ is a bijection. We say that θ is an *isomorphism* if

$$\theta(xy) = \theta(x)\theta(y)$$

for all $x, y \in G$. We say that G and H are *isomorphic*, and we write $G \cong H$ if an isomorphism between them exists.

If θ is an isomorphism from G to H, then θ induces a match-up of the elements of G with the elements of H that causes their multiplication tables to coincide. To the extent that we view groups as being defined by their multiplication tables, we see that isomorphic groups are essentially "the same." All "group theoretic" questions will have the same answers in G and H. For example, each of G and H will have equal numbers of elements of any given order, and G will be abelian iff H is abelian. (A group is said to be *abelian* if all of its elements commute, if $xy = yx$ for all elements x, y.)

As a concrete example, consider the group R of rotations of a cube and $S = \mathrm{Sym}(4)$. (We write $\mathrm{Sym}(n)$ as a shorthand for $\mathrm{Sym}(\{1, 2, \ldots, n\})$.) We have seen that $|R| = 24$ and, of course, $|S| = 4! = 24$. In fact, we will see that $R \cong S$, and so these differently constructed objects are group theoretically identical. (Note that R permutes eight objects, the vertices of a cube, and S permutes $\{1, 2, 3, 4\}$. As permutation groups, therefore, R and S are quite different.)

In the cube of Figure 1.2, there are four "major diagonals," AF, BE, CH, and DG. Each element of R corresponds to a rotation of the cube and each such rotation induces a permutation of these four diagonals. If we fix an assignment of the numbers 1, 2, 3, and 4 to the four diagonals, then each element of R determines a particular element of $S = \mathrm{Sym}(4)$. To see that the corresponding mapping $\theta : R \to S$ is an isomorphism, we need to establish that θ is a bijection. It is not very hard to

see (although we will not write a formal proof) that θ is injective. In other words, two different rotations cannot induce the same permutation of the diagonals. Since $|R| = 24 = |S|$, it follows that θ maps onto S. Because the multiplications in both R and S come about by simply following one operation by another, it should now be fairly clear that θ is an isomorphism.

Note that if $\theta : G \to H$ is an isomorphism, then $\theta^{-1} : H \to G$ is an isomorphism also. Furthermore, if $\varphi : H \to K$ is another isomorphism, it is routine to check that $\theta\varphi : G \to K$ is an isomorphism. It follows from all this that isomorphism of groups is an equivalence relation.

Problems

1.1 A permutation group G on a set X is said to be *transitive* if for every two elements $x, y \in X$, there exists $g \in G$ with $(x)g = y$. Also, G is *regular* if it is transitive and there is a unique element that carries x to y for all $x, y \in X$. Show that a transitive abelian permutation group is necessarily regular.

1.2 Let G be any group. For $x \in G$, let r_x and l_x be the mappings $G \to G$ defined by

$$(g)r_x = gx \quad \text{and} \quad (g)l_x = xg\,,$$

or in other words, by right and left multiplication by x on G. Let $R = \{r_x \mid x \in G\}$ and $L = \{l_x \mid x \in G\}$. Show that R and L are permutation groups on G and that $R \cong G \cong L$.

NOTE: The fact that every group is isomorphic to a permutation group is known as Cayley's theorem.

1.3 Let G, R and L be as in Problem 1.2. Show that

$$L = \{f \in \mathrm{Sym}(G) \mid fr = rf \ \text{ for all } \ r \in R\}.$$

1.4 Let G be a group of mappings on a set X with respect to function composition.

a. Find an example where $G \not\subseteq \mathrm{Sym}(X)$ and $|G| \geq 2$.

b. Show that if G contains some injective function, then $G \subseteq \mathrm{Sym}(X)$.

1.5 Let G be the dihedral group D_{2n}. Let $t \in G$ correspond to a "flip" and let $r \in G$ correspond to a "plane rotation." Show that $trt = r^{-1}$. Conclude that if n is odd, then only the identity of G commutes with all elements of G.

1.6 Decide whetber or not D_{24} is isomorphic to the group of rotations of a cube. Prove your answer.

1.7 Let V be an n-dimensional vector space over a field F with (a finite number) q elements. One writes $GL(n, q)$ to denote $GL(V)$. Show that

$$|GL(n, q)| = (q^n - 1)(q^n - q)(q^n - q^2)\cdots(q^n - q^{n-1})\,.$$

1.8 Let G be a group in which every nonidentity element is an involution. Show that G is abelian.

NOTE: An abelian group in which every nonidentity element has the same prime order p is called an *elementary abelian* p-group.

1.9 Consider the eight objects ± 1, $\pm i$, $\pm j$ and $\pm k$ with multiplication rules:

$$\begin{array}{lll} ij = k & jk = i & ki = j \\ ji = -k & kj = -i & ik = -j \end{array}$$
$$i^2 = j^2 = k^2 = -1,$$

where the minus signs behave as expected and 1 and -1 multiply as expected. (For example, $(-1)j = -j$ and $(-i)(-j) = ij = k$.) Show that these objects form a group containing exactly one involution.

NOTE: This is called the *quaternion group* and is denoted Q_8.