Notre Dame Journal of Formal Logic
Volume 45, Number 3, 2004

Linear Reducts of the Complex Field

James Loveys

Abstract A reduct of a first-order structure is another structure on the same set with perhaps fewer definable predicates. We consider reducts of the complex field which are proper (not essentially the whole field) but nontrivial in a sense to be made precise below. Our main result (the summary that is Theorem 7.1) lists seven kinds of reducts. The list is complete in the sense that every reduct is a finite cover of one of these. We also investigate when two items on our list can be the same, in a couple of natural senses.

1 A Little Basic Model Theory

Here our purpose is to outline certain basics of model theory for those unfamiliar with the subject. Those who know the territory may want to skim this section quickly for the purposes of notation, and the precise definition of reduct we use.

We will explain the notion of a (first-order, relational) structure in a little detail. Such an animal consists of a fixed set (say M) together with certain distinguished subsets of M^n for various natural numbers n. There are no a priori restrictions on which subsets we choose to distinguish. For each of these predicates on M we introduce a symbol. If $\{P_i : i \in I\}$ lists these and $R_i = P_i(\mathcal{M})$ is the corresponding subset of M^n, we will denote the structure $\mathcal{M} = (M; R_i : i \in I)$. Here and everywhere else the corresponding roman letter is used to denote the underlying set of the structure represented by a script letter. In case the structure is understood, we may use P_i for both the predicate symbol and the set it carves out of M^n. If necessary (say we are talking about two different structures at the same time), we will emphasize that we are considering P_i as applied to $\mathcal{M}$ by writing $P_i(\mathcal{M})$. When we do consider two different structures $\mathcal{M}$ and $\mathcal{N}$ for the same symbol P_i we will always assume that the number n so that $P_i(\mathcal{M}) \subseteq M^n$ is the same n so that $P_i(\mathcal{N}) \subseteq N^n$. The collection of symbols $\{P_i : i \in I\}$ is known as the *signature* or *similarity type* of the structure (actually of its language).

Printed October 26, 2004
2000 Mathematics Subject Classification: Primary, 03C65, 03C45
Keywords: complex numbers, reducts, strongly minimal sets, linearity, finite covers

We will frequently also include symbols for functions f from M^n to M; it is standard in model theory to include these as part of the basic paraphernalia of a structure, but that will not be of concern here. If we do this, take the symbol f as shorthand for its graph. The sets $P_i(\mathcal{M})$ are called *basic 0-definable sets*; the expression $P_i(x_1, \dots, x_n)$ is the *atomic formula* which defines the set. The one other basic 0-definable set is $\{(x, x) : x \in M\}$ with defining formula $x_1 = x_2$; we will never explicitly include equality among our basic symbols, but it's always there. The set M is 0-definable, by $\exists x_2(x_1 = x_2)$, for example.

Other 0-definable sets are formed (by induction) as follows. If $A \subseteq M^n$ is 0-definable via the formula $\chi(x_1, \dots, x_n)$ and π is a permutation of $\{1, \dots, n\}$, then the set $\{(a_{\pi(1)}, \dots, a_{\pi(n)}) : (a_1, \dots, a_n) \in A\}$ is also 0-definable, via the formula $\chi(x_{\pi(1)}, \dots, x_{\pi(n)})$. $A \times M$ is 0-definable via (say) $\chi(\bar{x}) \wedge x_{n+1} = x_{n+1}$. $M^n \setminus A$ is 0-definable via $\neg\chi(\bar{x})$. If B is another 0-definable subset of M^n, say by the formula $\rho(\bar{x})$, then $A \cup B$ is 0-definable via the formula $\chi(\bar{x}) \vee \rho(\bar{x})$. Finally, the set $\{(a_1, \dots, a_{n-1}) : (a_1, \dots, a_{n-1}, a_n) \in A \text{ for some } a_n \in M\}$ is 0-definable via $\exists x_n \chi(\bar{x})$. All 0-definable sets are formed by iterating these operations finitely often, starting from the basic P_is. (As are, formally, all formulas without parameters. But we will feel no compunction about writing $\chi \wedge \rho$ or $\forall x \chi$ and considering them formulas with the obvious interpretations.) If, as above, A is defined by $\chi(\bar{x})$ as above, we write $A = \chi(\mathcal{M})$ and also $\mathcal{M} \models \chi(\bar{a})$ for $\bar{a} \in A$. As above, if we have more than one structure $\mathcal{M}$ for the same signature, we will use this to emphasize that we are considering the subset of M^n picked out by the formula.

You may have been wondering about the "0" in "0-definable". If we have a formula without parameters $\chi(x_1, \dots, x_n, x_{n+1}, \dots, x_{n+m})$ as described above and $\bar{a} \in M^m$, then we have the *formula* $\chi(x_1, \dots, x_n, \bar{a})$ *with parameters* $\bar{a}$ and the $\bar{a}$-*definable set* $\{\bar{b} \in M^n : (\bar{b}, \bar{a}) \in \chi(\mathcal{M})\}$. A set is *definable* if it is definable with or without using parameters. Usually it will not matter for us. We will never again be so careful about variables (the x_i's) and formulas. For us, the crucial concept is the collection of definable sets.

We say $\mathcal{M}$ is a *substructure* of $\mathcal{N}$ if the two structures have the same signature, $M \subseteq N$, and for any P in that signature $P(\mathcal{M}) = P(\mathcal{N}) \cap M^n$ for the appropriate n. If further, for any formula without parameters $\chi(\bar{x})$, we have that $\chi(\mathcal{M}) = \chi(\mathcal{N}) \cap M^n$, then we say that $\mathcal{M}$ is an *elementary substructure* of $\mathcal{N}$. The crucial point of this last definition is in the case χ begins with an existential quantifier. If $\mathcal{N} \models \exists x \rho(x, \bar{a})$ and $\bar{a} \in M^n$, we know there is a witness for this in N; elementariness says we can choose it in M. In case $\mathcal{M}$ is an (elementary) substructure of $\mathcal{N}$ we say that $\mathcal{N}$ is an (elementary) extension of $\mathcal{M}$. To illustrate the distinction, suppose that $\mathcal{N} = (N; +)$ is an Abelian group—the only nontrivial basic 0-definable set is the graph of $+$. Any subset can be made into a substructure, but an elementary substructure must at least be a subgroup; consider the formula $\exists z(x + y = z)$. (In fact, it must be a pure subgroup, and in fact . . .)

The reader uncomfortable with the notion of elementary extension of a structure may console himself with the following. We are really interested in the complex field and its reducts (see below). Whenever we mention arbitrary elementary extensions below, the reader may feel free to consider a single algebraically closed extension field $\mathcal{K}$ of the complexes with size at least $(2^{\aleph_0})^+$. However, he must keep in mind that the parameters mentioned are in that case only permitted to come from $\mathbb{C}$ itself.

Furthermore, in practice any basic predicate $P(\mathcal{M})$ of some reduct $\mathcal{M}$ of the complex field will be presented as a Boolean combination of varieties (solutions sets of polynomial equations); in the corresponding reduct $\mathcal{N}$ of $\mathcal{K}$, $P(\mathcal{N})$ will be the same Boolean combination of the corresponding varieties in $\mathcal{K}$.

In the following we provide three nonequivalent possibilities for the definition of a reduct of a first-order structure, all of which have been used in the literature.

Definition 1.1 Suppose $\mathcal{M}$ and $\mathcal{N}$ are structures with the same underlying set $M = N$. Then $\mathcal{M}$ is a *reduct* of $\mathcal{N}$ if

1. (First Take) The signature of $\mathcal{M}$ is a subset of that of $\mathcal{N}$ and for any symbol P in the signature of $\mathcal{M}$, $P(\mathcal{M}) = P(\mathcal{N})$;
2. (Second Take) Every subset of M^n 0-definable in $\mathcal{M}$ is 0-definable in $\mathcal{N}$;
3. (Third Take) Every subset of M^n definable (allowing parameters) in $\mathcal{M}$ is definable (allowing perhaps other parameters) in $\mathcal{N}$;
4. If $\mathcal{M}$ is a reduct of $\mathcal{N}$, but not otherwise, we say that $\mathcal{M}$ is a *proper reduct* of $\mathcal{N}$. (In case $\mathcal{M}$ and $\mathcal{N}$ have the same definable sets, we will sometimes abuse notation and write $\mathcal{M} = \mathcal{N}$.)

(1) is the classical notion and certainly the strictest. We will never mention it further because for the purposes of most modern model theory it is overly language-bound. My personal feeling is that the right notion is (2). (3) yields more reducts than (2); a reduct (in the sense of (3)) is a reduct (in the sense of (2)) of the structure $(\mathcal{N} : a \in N)$ with names for the elements of the underlying set—technically, a name for $a \in M$ is a predicate whose interpretation is the singleton $\{a\}$, but this is a triviality. But notion (3) identifies reducts which are distinct in the sense of (2), thus providing a cruder analysis. However, for technical reasons *from now on, when we use the word reduct, we will mean in the sense of (3) above*.

We record here a bit of "general nonsense" that actually holds for reducts in the sense of (2) above, a fortiori also for our reducts.

Proposition 1.2 *Suppose that $\mathcal{N}$ is a proper reduct of $\mathcal{M}$. Then there is an elementary extension $\mathcal{M}'$ of $\mathcal{M}$ so that if $\mathcal{N}'$ is the corresponding reduct, there is an automorphism of $\mathcal{N}'$ which is not an automorphism of $\mathcal{M}'$.*

Our concern here is with reducts of the structure $(\mathbb{C}; +, \cdot)$, the field of complex numbers. Here are some examples.

Example 1.3

1. $(\mathbb{C}, \cdot)$.
2. Fix any subfield $F \leq \mathbb{C}$ and take the vector space structure $(\mathbb{C}; +, \lambda_a : a \in F)$ where in this notation we interpret each λ_a as (the graph of) the function $x \mapsto a \cdot x$—notice that for $F = \mathcal{Q}$, this is the "same" reduct as $(\mathbb{C}; +)$.
3. Let $\{P_i(x, y) : i \in I\}$ be any collection of polynomials in two variables and for each i, let $R_i(a, b)$ hold if and only if $P_i(a, b) = 0$—our structure is $(\mathbb{C}; R_i : i \in I)$.
4. Let g be the function $g(x) = x^2$; take $(\mathbb{C}; +, g)$.
5. $(\mathbb{C}; R)$ where $R(a, b, c)$ holds if and only if $a^2 + b^2 = c^2$.
6. As in the last item, except let $R(a, b, c)$ if and only if $a^2 + b^2 = c$.
7. As in the last two for $R(a, b, c)$ if and only if $ab(a + b) \neq 0$ and $c = \frac{ab}{a+b}$.
8. Here $R(a, b, c)$ holds exactly if $a^2 + b^2 + c^2 = 4 + abc$.

Obviously we could go on forever. It's not too hard to see that (1) and (2) are proper reducts (one can't define $+$ in the former or $\cdot$ in the latter). These are the paradigms for what we will define later as "linear" reducts. In (3) one can define neither, nor in fact any "genuinely ternary" relation. (Details later.) (4) is not a proper reduct: to see this, consider the definable function $z = \frac{1}{2}[(x+y)^2 - x^2 - y^2] = xy$ (if one has addition definable, one has subtraction and division by 2 definable). (6) is also not a proper reduct, but it takes a little more work to demonstrate this. Actually, there is a strong sense (see the Rabinovitch result below (Theorem 2.7)) in which for "nearly any" definable $R \subseteq \mathbb{C}^n$ for $n \geq 3$, the reduct $(\mathbb{C}; R)$ is not proper.

The others mentioned above are all proper, though this is not obvious. The right way to regard the Pythagorean reduct (5) is as a "finite cover" of $(\mathbb{C}; +)$. Consider the definable (on this structure $\mathcal{M}$) equivalence relation xEy if and only if $x^2 = y^2$; the map $x \mapsto x^2$ induces an identification of $\mathcal{M}/E$ and $(\mathbb{C}; +)$. As for (7), it is more or less the image of $(\mathbb{C}; +)$ under the fractional linear transformation $x \mapsto x^{-1}$ (that is, R is approximately the image of the graph of $+$ under this map). (8) is what happens to $(\mathbb{C}; \cdot)$ when you factor out the equivalence relation $x = y \vee x = y^{-1}$ and identify the quotient with $\mathbb{C}$ using the map $x \mapsto x + x^{-1}$. Details on all of the above later.

It should be mentioned right here that subsets of $\mathbb{C}^n$ which are definable in the structure $(\mathbb{C}; +, \cdot)$ are no more nor less than what are classically known as *constructible* sets. That is, they are Boolean combinations of the solution sets of polynomial equations in n (or fewer) variables. This is true because, like any algebraically closed field, the complex field admits what model theorists call "elimination of quantifiers"—any definable set is in fact definable by a formula that makes no use of quantifiers. (Actually, this is false the way we've presented things; it becomes true if we add predicates for all polynomial equations with natural number coefficients, or use function symbols for $+$ and $\cdot$. We won't worry about this here, because this terminology is mentioned here purely for the purpose of general culture.) This is not true of an arbitrary reduct of this structure, however, at least not in any language that can be found in a natural manner. We will frequently use the terminology "constructible" to emphasize that while something may be definable in the full field structure, there is no reason it must be definable in the reduct currently under consideration.

We now make our only semi-formal mention of the peculiarly-titled object $\mathcal{M}^{\mathrm{eq}}$. $\mathcal{M}$ as usual is some first-order structure. For any definable $A \subseteq M^n$ and definable $E \subseteq A \times A$ which is the graph of an equivalence relation on A, the quotient A/E inherits in a natural manner a structure from $\mathcal{M}$. Fix some parameters $\bar{a}$ so that A and E are $\bar{a}$-definable; we will suppress these in our description. For every formula $\chi(\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_m)$ with mn free variables (as indicated), we have a basic predicate $P_\chi(x_1, \ldots, x_m)$ for the structure $\mathcal{A}$ with universe A/E. We declare that $\mathcal{A} \models P_\chi(\bar{b}_1/E, \ldots, \bar{b}_m/E)$ exactly if for some $\bar{c}_1 E \bar{b}_1, \ldots, \bar{c}_m E \bar{b}_m$ (all tuples in A) we have that $\mathcal{M} \models \chi(\bar{c}_1, \ldots, \bar{c}_m)$. We call this the *structure induced on A/E by* $\mathcal{M}$, or just the *induced structure*. Any such A/E (with the induced structure) is said be *definable in* (or just in) $\mathcal{M}^{\mathrm{eq}}$.

That is, we restrict $\bar{a}$-definable predicates to A, and then project onto A/E. We allow the possibility that either $A = M^n$ or that E is the identity. The choice of $\bar{a}$ affects what this structure is in a literal sense, but we shall be playing so fast and loose with parameters that this will be irrelevant for us. The elements of A/E are known as “imaginary elements” of $\mathcal{M}$, and if we create a multi-sorted structure with all such A’s, we have what model theorists know as $\mathcal{M}^{\mathrm{eq}}$.

For a classical example, consider the construction of the rational field in the structure $(\mathcal{Z}; +, \cdot)$, the ring of integers. It occurs in $\mathcal{Z}^{\mathrm{eq}}$.

In fact, if $\mathcal{M}$ is the field of complex numbers, we never need to consider equivalence relations E as above. Specifically, for any A and E as above, there is a constructible function (using the same parameters $\bar{a}$) $f : A \longrightarrow M^k$ for some (possibly large) k so that $f(\bar{b}) = f(\bar{c})$ if and only if $\bar{b}E\bar{c}$. This phenomenon is known to model theorists as “elimination of imaginaries.” It does not always hold (with “constructible” replaced by “definable in the given structure”) for reducts of the complex field, however.

Definition 1.4 Suppose that $\mathcal{M}$ is a structure and that g is a function with domain M. The *image of* $\mathcal{M}$ *under* g is the structure $g(\mathcal{M})$ with underlying set $g(M)$ and for each 0-definable subset A of M^n a predicate P_A for $g(A)$. That is, the basic 0-definable subsets of $g(\mathcal{M})$ are exactly these $g(A)$’s. For $\mathcal{N} = g(\mathcal{M})$ we say that $\mathcal{M}$ is a *cover* of $\mathcal{N}$ *via* g. If the fibers of g are finite, we call $\mathcal{M}$ a *finite cover* of $\mathcal{N}$.

In practice, we will only use these notions when all objects in question are constructible and the equivalence relation $g(x) = g(y)$ is definable in $\mathcal{M}$. We do *not* necessarily assume that g itself is definable in $\mathcal{M}$ though, even if its range is contained in M. Also, in practice, when we use the term “cover”, we will be assuming that it is a finite cover, and both M and N are the set of complex numbers. The notation $g(\mathcal{M})$ will be used in other contexts, however.

Given any g from M onto N and structure $\mathcal{N}$ on N, we have in particular the *canonical cover* of $\mathcal{N}$ by $g^{-1}(\mathcal{N}) = \mathcal{M}$ via g. The basic 0-definable sets of this structure are $g^{-1}(A)$ for every basic 0-definable $A \subseteq N^n$. We could have taken all $g^{-1}(B)$’s for all 0-definable B’s, but in this case we would get no more 0-definable (and hence no more definable) predicates on M this way. We leave this as an exercise; it uses trivial facts like $g^{-1}(B \setminus C) = g^{-1}(B) \setminus g^{-1}(C)$.

By contrast, if we just took as basic in $g(\mathcal{M})$ only the images of the basic 0-definable predicates of $\mathcal{M}$, we would in general lose a lot. For example, consider the two structures $(\mathbb{C}; \cdot)$ and $(\mathbb{C}; P)$ where P is the complement of the graph of multiplication. For our purposes, these are really the same structure. However, for $g : \mathbb{C} \longrightarrow \mathbb{C}$, where $g(x) = x + x^{-1}$ (say $g(0) = 0$), the image of the graph of multiplication is the relation R indicated in item (8) of the examples above. However, $g(P)$ is cofinite in $\mathbb{C}^3$, which makes it easy to see that R is not definable in the structure with $g(P)$ as the sole basic 0-definable set. We will see below in Section 8 that, in fact, everything 0-definable in $g((\mathbb{C}; \cdot))$ is, in fact, 0-definable in $(\mathbb{C}; R)$.

2 Strongly Minimal Sets, Linearity, etc.

In this section we will delineate precisely those reducts of the complex field that are of interest to us. They are those which are in a well-defined sense “geometrically simple,” but not too simple. We operate in a more general setting, occasionally bringing ourselves back down to earth.

Definition 2.1 Let $\mathcal{M}$ be a structure and $A \subseteq M$ be any set. For any $b \in M$, we say that b is in the *algebraic closure of* A and write $b \in acl(A)$ if there is a formula without parameters $\chi(x, \bar{y})$ and some $\bar{a}$ with each $a_i \in A$ so that the formula $\chi(x, \bar{a})$ has finitely many solutions in M, one of which is b.

That is, $\mathcal{M} \models \chi(b, \bar{a})$ and for only finitely many (possibly no) other elements c do we have $\mathcal{M} \models \chi(c, \bar{a})$. The paradigm for this definition is none other than the field of complex numbers. In this, our favorite structure, the algebraic closure of any set A is simply the smallest algebraically closed subfield containing A. If our structure is an infinite vector space over some field F (so the graph of scalar multiplication by $a \in F$ is regarded as a basic 0-definable set, as is the graph of addition), the algebraic closure of A is the subspace it generates. In the structure $(\mathbb{C}; \cdot)$, the algebraic closure of the empty set is the collection of torsion elements; the algebraic closure of any set A is the pure subgroup it generates. (That is, the smallest subgroup containing A that is closed under taking nth roots for every n.)

The following are the most basic properties of the algebraic closure operation.

Proposition 2.2 *Let $\mathcal{M}$ be any structure, and A, B be arbitrary subsets of M.*

1. $A \subseteq acl(A)$;
2. *if $A \subseteq B$, then $acl(A) \subseteq acl(B)$ (this property is called* monotonicity*);*
3. *if $a \in acl(A)$, then there is a finite set $C \subseteq A$ with $a \in acl(C)$ (*finite basis*);*
4. *if $a \in acl(A)$ and $A \subseteq acl(B)$, then $a \in acl(B)$ (*transitivity*);*
5. $acl(acl(A)) = acl(A)$.
6. *Suppose that $\mathcal{N}$ is an elementary extension of $\mathcal{M}$ and $A \subseteq M$. Then if we calculate $acl(A)$ in the two structures separately, we get exactly the same set; in particular, $M = acl(M)$ calculated in $\mathcal{N}$.*

These are all easy exercises in the definitions. For (6), we note that if $b \in acl(A)$ taken in $\mathcal{N}$, find a formula $\chi(x, \bar{a})$ witnessing this fact. If it has exactly 33 solutions in $\mathcal{N}$, there is a first-order formula true of $\bar{a}$ in $\mathcal{N}$ stating this fact. It is also true of $\bar{a}$ in $\mathcal{M}$ and the solutions in $\mathcal{M}$ are also solutions in $\mathcal{N}$. We leave the others to the reader. The reader is warned that even for finite (or definable) sets A, the set $acl(A)$ is not usually definable.

The examples mentioned above all satisfy a further property of the algebraic closure operation, namely, for any $A \subseteq M$ and $a, b \in M$,

$$\text{if } a \in acl(A \cup \{b\}) \text{ but } a \notin acl(A), \text{ then } b \in acl(A \cup \{a\}). \quad \text{(exchange)}$$

The exchange property is far from true in most structures, but it is for the mentioned examples because they are all what is known as strongly minimal structures, which we now define. The reader is cautioned that there are indeed nonstrongly minimal structures on which the algebraic closure operation satisfies exchange.

Definition 2.3 Let $\mathcal{M}$ be any infinite structure and suppose that the following is true for any elementary extension $\mathcal{N}$ of $\mathcal{M}$: every definable subset of N is either finite or cofinite. Then we call the structure $\mathcal{M}$ a *strongly minimal* structure.

Several points should be made here; first, the subsets must be of N itself, not of N^k (the diagonal is always infinite and coinfinite). We are definitely allowing parameters in the definition of these sets. In every structure the finite and cofinite sets are definable (hence the "minimal"). It is not hard to check that the property of being strongly minimal passes to reducts. And, as mentioned above, the structure $(\mathbb{C}; +, \cdot)$ is strongly minimal. (This is an easy exercise from the "quantifier elimination" result mentioned above.)

The mention of elementary extensions of $\mathcal{M}$ in this definition is a little awkward, but necessary, as we want it really to be a property of the first-order theory of the structure. The standard example of a structure with no infinite, coinfinite definable subsets, but with elementary extensions that have such subsets, is $(\mathcal{N}; <)$, the natural numbers with the usual order. But the initiated might find the following exercise amusing, and the uninitiated may find it reassuring. (As far as the author knows, it is due to Hrushovski, and unpublished.)

Proposition 2.4 *Suppose $\mathcal{M}$ is a structure and M is uncountable. Suppose also that every definable subset of M is either finite or cofinite. Then $\mathcal{M}$ is strongly minimal.*

More to the point, here is, as promised.

Proposition 2.5 *Suppose that $\mathcal{M}$ is a strongly minimal structure. Then in any elementary extension of $\mathcal{M}$, the algebraic closure operation satisfies the exchange property. (See above.)*

A pair (S, cl) where S is a set and cl an operation on the subsets of S satisfying (1)–(5) of Proposition 2.2 and the exchange property is known as an *exchange pregeometry* or a *matroid*. The "pre" refers to two things; we do not necessarily have $acl(\varnothing) = \varnothing$, nor do we have that the closure of a singleton is itself. If we did, we would have an *exchange geometry*. We can easily create one out of our pregeometry by stripping off $acl(\varnothing)$ and identifying mutually algebraic elements, but we don't really want to do that. (Imagine what the complex field would look like after such surgery!) The complexity of this geometry is a measure of the complexity of the collection of definable predicates on any strongly minimal structure.

It is routine to check from these properties that any algebraically closed set—A such that $A = acl(A)$—has a well-defined dimension, $\dim(A)$. This number is the size of a minimal subset B so that $A = acl(B)$; equivalently, it is the size of a maximal C so that for $c \in C$, $c \notin acl(C \setminus \{c\})$. The collection of closed subsets of any strongly minimal structure naturally forms a lattice. We give several versions of simplicity of the geometry in the following definition. The necessity of shifting to elementary extensions is exemplified by considering the structure $(\mathbb{C}; +, \cdot, c : c \in \mathbb{C})$, the complex field with names for all the elements. Surely this structure is no simpler than the field itself, but its geometry trivializes entirely. The complexity returns once we pass to an elementary extension, which in this case is exactly the same thing as an algebraically closed field extension, as long this extension has suitably large transcendence degree.

Definition 2.6 Let $\mathcal{M}$ be a strongly minimal structure.

1. We say that $\mathcal{M}$ is (geometrically) *trivial* if for every elementary extension $\mathcal{N}$ of $\mathcal{M}$ and $A \subseteq N$, we have that $acl(A) = \cup_{a \in A} acl(\{a\})$.
2. We say that $\mathcal{M}$ is *locally modular* if for any elementary extension $\mathcal{N}$ of $\mathcal{M}$ and any algebraically closed subsets A and B of N such that $A \cap B$ contains some element not in $acl(\varnothing)$, we have that
$$\dim(A) + \dim(B) = \dim(A \cap B) + \dim(A \cup B).$$
3. If we can remove the restriction that $A \cap B$ has a nonalgebraic element in the last clause, we call $\mathcal{M}$ *modular*.
4. If $\mathcal{M}$ is not trivial, but is locally modular (possibly modular) we call it *linear*.

Our structure $\mathcal{M}$ is modular exactly if the lattice of algebraically closed sets in any elementary extension of $\mathcal{M}$ is modular in the lattice-theoretic sense; hence the term.

It is immediate that trivial structures are modular. The distinction between those structures that are locally modular but not modular and those that are modular but nontrivial is a technical point. It can have significant model-theoretic impact (see Laskowski [3]), but will be of no concern to us here. Hrushovski calls trivial strongly minimal structures "combinatorial"; we have here used the more traditional terminology. He has a definite point; these structures do have trivial geometries but can otherwise be far from transparent. He uses the word *linear* to refer to strongly minimal structures that are locally modular but nontrivial. *From now on we will do the same.* All other strongly minimal structures he calls "geometrical".

There are so many trivial reducts of the complex field that describing them in any sensible fashion may be a hopeless cause. If we throw in as basic predicates any collection $\{P_i : i \in I\}$ whatsoever of subsets of $\mathbb{C}^2$, the result will be a trivial structure, and how do we tell them apart? We will leave them out of our discussion altogether, except for one example.

Much of the work up to the present on reducts of the complex field has been motivated by something known as "Zil'ber's Conjecture", which stated (approximately) that any nonlocally modular strongly minimal structure $\mathcal{M}$ must interpret an algebraically closed field. That is, for some definable (in the structure $\mathcal{M}$) $A \subseteq M^n$ and some definable equivalence relation E on A, the structure $\mathcal{A}$ on A/E inherited from $\mathcal{M}$ is that of an algebraically closed field. I say "approximately" because nobody, Zil'ber included, seems to know the exact content of the purported conjecture. At this level, it is of historical interest anyway, because it has been rather resoundingly refuted by Hrushovski, whatever it was. (See Hrushovski [1].)

Even since Hrushovski's counterexample machine came into being, there remains interest in the "restricted Zil'ber conjecture". This restricts itself to those strongly minimal sets which are of the form A/E as above, where $\mathcal{M}$ is the complex field (more generally, any algebraically closed field). There have been several positive results in this direction, and there is good reason to believe it true. By far the most spectacular result to date in this direction is the following, due to Rabinovitch [7]).

Theorem 2.7 *Suppose that $\mathcal{M}$ is a reduct of the complex field that is neither linear nor trivial. Then it is not a proper reduct.*

That is, if we are interested in reducts of the complex field, we may as well restrict ourselves to those that are linear or trivial. Having already dismissed the latter, we come to our (relatively) happy medium. The above is not the way Rabinovitch stated her result, but it is equivalent.

We remind the reader of the slight discussion of $\mathcal{M}^{\mathrm{eq}}$ in Section 1. The most striking general result about linear strongly minimal structures known to the author is the following, due to Hrushovski. (Actually, it holds in considerably more generality, with appropriate modifications.) We will call a structure a *strongly minimal group* if it is strongly minimal, and one of its basic definable predicates happens to be the graph of a group operation. That is, the structure has a definable group operation, and conceivably other definable predicates, and is as well strongly minimal. It is known that in this case the group structure is necessarily Abelian (see Poizat [6]).

Theorem 2.8 *Suppose that $\mathcal{M}$ is a linear strongly minimal structure. Then there is, in $\mathcal{M}^{\mathrm{eq}}$, another strongly minimal structure, which is in fact a strongly minimal group. That is, there is a definable $A \subseteq M^n$ and a definable equivalence relation E so that A/E, with its induced structure, is a strongly minimal group. It is also linear.*

The possibilities for the full structure of linear strongly minimal groups are known, but for our purposes we note only what follows below. It turns out that all the strongly minimal groups definable in $\mathbb{C}^{\mathrm{eq}}$ are of unbounded exponent. We will list the possibilities later. The next result is in Loveys [4].

Theorem 2.9 *Let $\mathcal{A}$ be a strongly minimal group of unbounded exponent, which is also a linear structure. Let R be the ring of all definable (in the structure $\mathcal{A}$) endomorphisms of $(A; +)$. Then any predicate definable in the structure $\mathcal{A}$ is in fact definable in its reduct $(A; +, r : r \in R)$.*

The result in [4] is both more precise and more general than this, but for our needs this is sufficient.

Our situation is as follows; we have some given strongly minimal linear structure $\mathcal{M}$, and we know that there is a strongly minimal Abelian group $\mathcal{A}$ somewhere in $\mathcal{M}^{\mathrm{eq}}$. Under these hypotheses, the following is demonstated in Loveys [5].

Proposition 2.10 *With the notation of the last paragraph, there are definable equivalence relations E on M and $\sim$ on A with finite classes and a definable bijection between M/E and $A/{\sim}$.*

That is, M and A have what can reasonably be called a "definable common factor". The "factors" of such a structure $\mathcal{A}$ are rather constrained. Suppose that $\mathcal{A} = (A; +, r : r \in R)$ is, as above, a strongly minimal module and that $\sim$ is a definable equivalence relation on A with finite classes. Suppose that R contains all the definable endomorphisms of the group. The following is shown in [5].

Proposition 2.11 *We use the notation of the previous paragraph. Suppose that $\sim$ is defined via parameters $\bar{a}$. Then whenever $x \sim y$, there is a nonzero integer m, an $r \in R$, and an element b of A in $acl(\bar{a})$ so that $mx = ry + b$. The ring R embeds into a division ring D and in D we have that $m^{-1}r$ is a root of unity.*

Note that in this proposition we are *not* saying that $r = ms$ for some element $s \in R$.

Occasionally we will need to compare two groups both in $\mathcal{M}^{\mathrm{eq}}$ for one of our structures $\mathcal{M}$. The following is what we will need; it is a very special case of the results in Hrushovski and Pillay [2].

Proposition 2.12 *Let $\mathcal{M}$ be a strongly minimal linear structure and $\mathcal{A}_1$ and $\mathcal{A}_2$ be strongly minimal groups in $\mathcal{M}^{\mathrm{eq}}$. Then any subset of $A_1 \times A_2$ definable in $\mathcal{M}^{\mathrm{eq}}$ is a Boolean combination of cosets of definable subgroups of $A_1 \times A_2$. The parameters needed to define the subgroups in question can always be chosen algebraic over those needed to define the groups.*

As advertised, we now describe what the candidates for the Abelian group are. A strongly minimal group definable from the structure $(\mathbb{C}; +, \cdot)$ is nothing more or less than what is known to algebraic geometers as a "one-dimensional, irreducible, algebraic group". The following classical result lists all these. It can be found in Silverman [9] (Theorem IV.1.6, p. 293).

Theorem 2.13 *Let $(A; *)$ be an irreducible, one-dimensional algebraic group. Then there is a constructible isomorphism between $(A; *)$ and one of the following:*

1. $(\mathbb{C}; +)$*;*
2. $(\mathbb{C} \setminus \{0\}; \cdot)$*;*
3. *an elliptic curve with its canonical group operation.*

Now, if a group is definable from a reduct of $\mathbb{C}$, it is certainly definable from $\mathbb{C}$ itself. A slightly more subtle point is that one could conceivably define sets in $\mathbb{C}^{\mathrm{eq}}$ that are definable and strongly minimal in some reduct, but not strongly minimal in the full structure, so by restricting ourselves to the ones above, we would not get the full story. The following says that this does not happen.

Proposition 2.14 *If a group definable in a reduct of the complex field is strongly minimal in the reduct, it must also be strongly minimal in $\mathbb{C}^{\mathrm{eq}}$ with the full algebraic structure.*

Proof For such a set $S \subseteq \mathbb{C}^n$, clearly there can only be at most $a \in \mathbb{C}$ such that the set $\{\bar{x} \in \mathbb{C}^{n-1} : (a, \bar{x}) \in S\}$ is infinite. If there is such an a and we throw away the rest of S—which will be finite—we finish by induction on n by projecting. So suppose that for every $a \in \mathbb{C}$, this set is finite. Then clearly S has “Morley rank one”, or in more standard terminology, is one-dimensional.

In the full field structure, a basic result (see, e.g., [6], Lemma 2.1, p. 41) tells us that if S is a group, it has a definable subgroup S^0 of finite index which is strongly minimal in the full structure, so S^0 must be one of the three types mentioned above. They are all divisible groups so in each case S^0 is just mS for some natural number m. In particular, S^0 is definable purely from the group structure, and thus in the reduct. If it weren’t the whole of S, strong minimality in the reduct would be contradicted; consider its cosets.

The assumption that S lives in $\mathbb{C}^n$ is easily removed to complete the proof. □

It should be noted that the only place where the group operation comes in is in showing that the Morley rank one set S (in the full field) is, in fact, strongly minimal. In the trivial or bounded exponent cases, we may not get this, as the following illustrate.

Example 2.15

1. Let $R = \{(x, y) \in \mathbb{C}^2 : x^2+y^2 = 0\}$. In the trivial reduct $(\mathbb{C};\, R)$ the definable set R is strongly minimal. (This easily follows from quantifier elimination, which we get once we adjoin a constant symbol for 0, and a function symbol, the definable unary function $x \mapsto -x$, to the language.) But R split definably into $\{(x, ix) : x \in \mathbb{C}\}$ and $\{(x, -ix) : x \in \mathbb{C}\}$ over $(\mathbb{C}; +, \cdot)$—indeed these sets are definable in the linear structure $(\mathbb{C}; +, R)$.
2. Let V be an infinite-dimension vector space over the field GF_4 of four elements, and V_0 a one-dimensional subspace of V. $\mathcal{M} = (V; +, \lambda_a, v_0 : a \in GF_4, v_0 \in V_0)$. Let $\alpha \neq 0, 1$ be in GF_4 and $R(x, y) \Leftrightarrow y - \alpha x \in V_0$. Let $\mathcal{M}_0 = (V; +, R, v_0 : v_0 \in V_0)$. Then R is a strongly minimal group when regarded in $\mathcal{M}_0^{\mathrm{eq}}$; again, this can be seen using quantifier elimination for $\mathcal{M}_0$, or by using automorphisms of $\mathcal{M}_0$. But of course it has degree 4 in $\mathcal{M}$.

Theorem 2.13 mentions a constructible (therefore definable in the field structure) isomorphism between the given group and one on the list. But even if the group is definable in some reduct of $\mathbb{C}$, there is no reason the isomorphism must be definable in said reduct. It usually will not be.

About the groups themselves, nothing need be said about either the additive or multiplicative groups of $\mathbb{C}$. However, as we are assuming no knowledge of elliptic curves, our next project is to describe them in sufficient detail for our purposes.

3 A Few Words about Elliptic Curves

We summarize here for the unacquainted certain (mostly quite basic) facts about elliptic curves. Everything in this section can be found in, for example, Silverman [8], especially Chapter 3. Elliptic curves over $\mathbb{C}$ are curves (that is, defined by an equation) in projective 2-space over the complex field which have genus 1 with a distinguished base point. We will not attempt to present the definition in this form, however. It is known that any elliptic curve over $\mathbb{C}$ is definably isomorphic to one with an equation in *simplified Weierstrass form*. That is, up to a (locally rational) definable map, the curve has an equation of this form: $Y^2Z = X^3 + aXZ^2 + bZ^3$, where a and b are complex numbers. Also, the discriminant ($\Delta = -16(4a^3+27b^2)$) must be nonzero—equivalently, the equation $x^3 + ax + b = 0$ must have 3 distinct solutions. Here we have written the equation with homogeneous variables, as is customary when discussing projective space. However, as is customary when discussing elliptic curves, we will in fact write the equation of the curve as

$$y^2 = x^3 + ax + b$$

and regard it as living in $\mathbb{C}^2$. We must, of course, remember that there is a “point at infinity” on the curve, as well. It may be worthwhile to note that two elliptic curves are definably isomorphic exactly if they have the same j-invariant, a number defined as $j = 1728(4a)^3/\Delta$ for the given equation. From this it is easily seen that if E_i is the curve $y^2 = x^3 + a_ix + b_i$ for $i = 1, 2$ and E_1 and E_2 are isomorphic, then there is a complex number c so that $a_1 = c^4a_2$ and $b_1 = c^6b_2$. An isomorphism in this case is $(x, y) \mapsto (c^2x, c^3y)$; any other isomorphism is the composition of this with an automorphism of E_2.

An elliptic curve has a definable group operation on it, which we will now describe. The identity $\mathcal{O}$ is the point at infinity. Three distinct points (x_0, y_0), (x_1, y_1), and $(x_2, -y_2)$ on the curve are collinear exactly if $(x_0, y_0)\oplus(x_1, y_1)\oplus(x_2, -y_2) = \mathcal{O}$. That is, to calculate $(x_0, y_0) \oplus (x_1, y_1)$ for two distinct points on the curve, one first draws the line through them to find a (uniquely defined) third point $(x_2, -y_2)$ on the intersection of the curve and the line; then the sum is (x_2, y_2), which is the inverse of $(x_2, -y_2)$ under this operation. To find $(x_0, y_0) \oplus (x_0, y_0)$ one first takes the tangent line to the curve at the point, finds the other point $(x_2, -y_2)$ where this line intersects the curve (if there is no other point, (x_0, y_0) is an element of order 2 in the group), and defines the sum as (x_2, y_2). If one works out the details for the curve $y^2 = x^3 + ax + b$, for (x_0, y_0) and (x_1, y_1) distinct points, the sum is the point (x_2, y_2), where

$$x_2 = (\frac{y_1 - y_0}{x_1 - x_0})^2 - x_0 - x_1$$

and

$$y_2 = -(y_0 + \frac{(x_2 - x_0)(y_1 - y_0)}{x_1 - x_0}).$$

It turns out that, abstractly, all the groups arising as above from an elliptic curve over $\mathbb{C}$ are isomorphic to the direct product of two copies of the multiplicative group of $\mathbb{C}$. We will also be interested (for reasons indicated in Section 2) in the ring R of all definable endomorphisms of an elliptic curve. There are certain obvious ones; for a positive integer m, the map $[m]$ where $[m](x, y) = (x, y) \oplus (x, y) \oplus \cdots \oplus (x, y)$ (m times) is, of course, a definable endomorphism. Given that the inverse of (x, y) is $(x, -y)$ the map $[-m]$ where $[-m](x, y) = (x, -y) \oplus \cdots \oplus (x, -y)$ (m times) is also a definable endomorphism. This provides an embedding of the integers into the ring R. Usually (e.g., for generic a and b) this embedding is onto; there are no definable endomorphisms but the obvious ones. In case there are further definable endomorphisms, we say the curve has *complex multiplication*. We outline exactly the possibilities for R in this case.

Let d be a square-free positive integer. If d is not equivalent to 3 (mod 4), an *integer* in the field $\mathcal{Q}[\sqrt{-d}]$ is an element of the form $m + n\sqrt{-d}$ for standard integers m, n. If d is equivalent to 3 (mod 4), an integer in $\mathcal{Q}[\sqrt{-d}]$ is of the form $m + n(\frac{1}{2} + \frac{1}{2}\sqrt{-d})$ for standard integers m, n.

Proposition 3.1 *Let E be an elliptic curve over $\mathbb{C}$ which has complex multiplication and let R be its ring of definable endomorphisms. Then for some square-free positive integer d, R is isomorphic to a subring of the ring of all integers in $\mathcal{Q}[\sqrt{-d}]$.*

Notice (this will be relevant for us) that $\mathcal{Q}[\sqrt{-d}]$ has no roots of unity except 1 and -1 unless $d = 1$ or $d = 3$. Notice also that any such R is generated over $\mathcal{Z}$, as an Abelian group, by a single element.

Example 3.2

1. Consider the elliptic curve with equation $y^2 = x^3 + b$. Let η be a primitive 3rd root of unity. The map $[\eta]$ given by $(x, y) \mapsto (\eta x, y)$ is easily seen to be an endomorphism of the curve. The ring R mentioned above is isomorphic to $\mathcal{Z}[\eta]$ in this case.
2. Consider the curve $y^2 = x^3 + ax$. The map $[i]$ given by $(x, y) \mapsto (-x, iy)$ is a definable endomorphism of this curve. R here is isomorphic to $\mathcal{Z}[i]$.

Definition 3.3 Let E_1 and E_2 be elliptic curves. A definable (i.e., constructible) map $f : E_1 \longrightarrow E_2$ is an *isogeny* if it preserves the group operation. If there is a nonzero isogeny from E_1 to E_2, we say that the curves are *isogenous*.

It turns out that any nonzero isogeny is necessarily onto and has finite kernel. Further, it has a *dual*; that is, there is another nonzero isogeny $g : E_2 \longrightarrow E_1$ and for some natural number m, we have that $gf : E_1 \longrightarrow E_1$ is the map $[m]$ mentioned above and $fg : E_2 \longrightarrow E_2$ is the map $[m]$ for this curve.

The following is listed an exercise in [8] (6.9, p. 168).

Proposition 3.4 *Let E_1 and E_2 be elliptic curves over $\mathbb{C}$ with complex multiplication. Let R_1 and R_2 be their respective rings of definable endomorphisms. Then E_1 and E_2 are isogenous if and only if the rings of quotients of R_1 and R_2 are isomorphic.*

We finally record the following direct corollary of the classical theorem of Hurvitz on genus (see Theorem 5.9 in [8]). We will not explain the concept of "genus" mentioned in (1), as the corollary (2) is all we need.

Proposition 3.5

1. *If there is a constructible map from the curve C_1 to the curve C_2 with infinite range, then the genus of C_1 is at least that of C_2 (assuming we are working over the field $\mathbb{C}$, as usual).*
2. *There does not exist any constructible map from $\mathbb{C}$ onto an elliptic curve.*

4 In Case the Group is Multiplication

We are assuming that we are starting out with a linear reduct of the complex field. We know that there is a strongly minimal group close by, which is also a linear structure. We know it is constructibly isomorphic to one of three things: the multiplicative group of $\mathbb{C}$, the additive group of same, or an elliptic curve with addition as spelled out in Section 3. In this section, we will deal with the first of these possibilities. So we have a structure $\mathcal{M}$ which is a reduct of the complex field, and we are given that there is a definable equivalence relation E (in said structure) on M, a constructible equivalence relation $\sim$ on $(\mathbb{C}; \cdot)$ and a constructible isomorphism of $\mathbb{C}/{\sim}$ and $\mathcal{N} = \mathcal{M}/E$.

We present the first of several diagrams to help indicate what is going on. We introduce here and throughout certain conventions for our diagrams. First, every structure and map (arrow) is constructible. Second, every map is a surjection. Third, all the strictly vertical and curved maps are definable in the structure at the top of them; thus if one structure $\mathcal{S}$ is directly above another $\mathcal{T}$, then $\mathcal{T}$ exists in $\mathcal{S}^{\mathrm{eq}}$.

$$\begin{tikzcd}
\mathcal{M} \arrow[dd, bend right] & \\
\mathcal{A} \arrow[r, "\cong"] \arrow[d] & (\mathbb{C}; \cdot) \arrow[d] \\
\mathcal{N} \arrow[r] & \mathbb{C}/{\sim}
\end{tikzcd}$$

We will be repeatedly using the following simple observation: Suppose that $\mathcal{M}_1$ and $\mathcal{M}_2$ are two structures definable in $\mathbb{C}^{\mathrm{eq}}$ and that f is a constructible function mapping M_1 onto M_2. Suppose that E is a definable equivalence relation on M_1 and that whenever $f(x) = f(y)$ we have that xEy. Then the image E' of E under f is a constructible equivalence relation on M_2 and f induces a definable bijection of M_1/E and M_2/E'.

We must first investigate the possibilities for $\sim$. We remind the reader that any reduct of the complex field which has (the graph of) multiplication among its definable predicates and is also linear must be really nothing more than $(\mathbb{C}; \cdot)$ itself. That is, anything definable in such a reduct is already definable in the pure multiplicative group structure. (This follows easily from Theorem 2.9 above, and the well-known fact that the only constructible maps from $\mathbb{C}$ to itself which are multiplicative group homomorphisms are the maps $x \mapsto x^n$ for n an integer.)

In the notation of Section 2, the ring R here is just the integers. As pointed out in Proposition 2.11, if $\sim$ is a definable equivalence relation on $(\mathbb{C}; \cdot)$, then for any $x, y \in \mathbb{C}$ if $x \sim y$ we must then have $x^n = y^m a$ for some a algebraic in the parameters $\bar{a}$ used to define $\sim$. We also must have $n = m$ or $n = -m$, again by 2.11, and we may assume that n is positive. Fix any x_1 not algebraic in $\bar{a}$. Let $x_1, \ldots, x_k$ list the $\sim$-class of x. For $j = 1, \ldots, k$ choose n_k, m_k, a_k with n_k positive, a_k algebraic in $\bar{a}$, and $x_k^{n_k} = x_1^{m_k} a_k$; further, choose such a triple with n_k minimal possible. Then it's easy to see that, in fact, every $n_k = 1$ and every m_k is then either 1 or -1.

So the class of x_1 consists of $x_1, x_1a_2, \dots, x_1a_\ell$ and $x_1^{-1}a_{\ell+1}, \dots, x_1^{-1}a_k$. (The second part of this list may be empty. We may have reordered the x_is, and we have set $a_1 = 1$.) Notice now that $\{a_1, \dots a_\ell\}$ must be a multiplicative subgroup of $\mathbb{C}$. To see this, notice that nothing in the class is algebraic in $\bar{a}$, so the elements all realize the formula $xa_j \sim x$ for each $j = 1, \dots, \ell$. For similar reasons, the set $\{a_{\ell+1}, \dots, a_k\}$ is either empty or a coset of said group. We can thus choose a generator γ for the group and a single element a so that the class of x_1 is either $\{x_1, \gamma x_1, \dots, \gamma^{\ell-1}x_1\}$ or this together with $\{x^{-1}a, x_1^{-1}\gamma a, \dots, x_1^{-1}\gamma^{\ell-1}a\}$.

We assume we are in the first case for the moment, so $k = \ell$. By strong minimality, for all but finitely many elements x, the $\sim$-class of x consists precisely of the elements $\{x, \gamma x, \dots, \gamma^{\ell-1}x\}$. For our purposes, there is no loss of generality in assuming that *all* the classes have exactly this form. That is, we are interested in the structure that $\mathbb{C}/\sim$ inherits, which will have the same definable sets as $\mathbb{C}/\sim'$ where we adjust finitely many of the $\sim$-classes. Consider $f : \mathbb{C} \longrightarrow \mathbb{C}$ defined by $x \mapsto x^n$. The image of $\sim$ under this map is simply the identity. Again, without loss of generality, we may assume that $\sim$ itself was the identity. What we are doing here is adding the following diagram to the one above, and then ignoring the top line by considering $\mathcal{A}/S$ instead of $\mathcal{A}$ and so on. Here $H = \{x \in \mathbb{C} : x^n = 1\}$ and S is the corresponding subgroup of $\mathcal{A}$.

$$\begin{tikzcd}
\mathcal{A} \arrow[r, "\cong"] \arrow[d] & (\mathbb{C};\cdot) \arrow[d] \arrow[dr, "f"] & \\
\mathcal{A}/S \arrow[r] & \mathbb{C}/H \arrow[r] & (\mathbb{C};\cdot)
\end{tikzcd}$$

Thus we are in the following situation. There is a constructible function $g : \mathbb{C} \longrightarrow \mathbb{C}$ such that if our reduct is $\mathcal{M}$, the structure $g(\mathcal{M})$ is nothing more or less than the multiplicative group of the complex numbers. Put another way, in this case our reduct is just the multiplicative group, up to a finite definable cover. (We remind the reader that g may not be definable in the reduct $\mathcal{M}$; the relation $xEy \Leftrightarrow g(x) = g(y)$ is, however.)

So assume we are in the other case. Just as in the last paragraph, we may assume that $\ell = 1$ and that every class is of the form $\{x, x^{-1}a\}$ for some fixed a. (We put 0 in a class by itself.) Choose b so that $b^2 = a$ and consider the map $h : \mathbb{C} \longrightarrow \mathbb{C}$ where $h(x) = xb^{-1} + bx^{-1}$ (and $h(0) = 0$). We have $h(x) = h(y)$ if and only if $x \sim y$, so that h induces a constructible bijection of $\mathbb{C}/\sim$ and $\mathbb{C}$.

$$\begin{tikzcd}
\mathcal{M} \arrow[dd, bend right] & & \\
\mathcal{A} \arrow[r, "\cong"] \arrow[d] & (\mathbb{C};\cdot) \arrow[d] \arrow[dr, "h"] & \\
\mathcal{N} \arrow[r] & \mathbb{C}/\sim \arrow[r] & (\mathbb{C};\dots)
\end{tikzcd}$$

Clearly, the binary function $(x, y) \mapsto x \cdot^* y = xyb^{-1}$ is definable in $(\mathbb{C};\cdot)$, and the image of this function under the map h is the graph of the relation

$P(x, y, z) \Leftrightarrow x^2 + y^2 + z^2 = 4 + xyz$. (This is left as an exercise, basically in high-school algebra, as are similar claims below; we don't really need the specific polynomial defining P, but it was fun to find it. Also, here we have essentially "moved the identity" of the group to b.)

It turns out (see Section 8) that anything definable in the structure $h((\mathbb{C}; \cdot))$ is in fact definable in the structure $(\mathbb{C}; P)$. Thus up to a finite definable cover, our reduct is precisely $(\mathbb{C}; P)$.

5 In Case the Group is Addition

In this section we do the same as in the last, except for the case where the group we have is constructibly isomorphic to the additive group of the complex field. One factor that makes this case a little more involved than that of the previous section is that there are reducts of $(\mathbb{C}; +, \cdot)$ strictly between $(\mathbb{C}; +)$ and the whole field. But they are easily described. Indeed, let F be any subfield of $\mathbb{C}$. The reduct $(\mathbb{C}; +, \lambda_a : a \in F)$ is linear, and every reduct of the field which contains $+$ and is linear has this form. (Again, use 2.9 for this.) Once more, we use λ_a to represent the function $x \mapsto ax$.

Now suppose that we have an equivalence relation $\sim$ definable (using parameters $\bar{a}$) in the vector space $(\mathbb{C}; +, \lambda_a : a \in F)$. By 2.11 again, if $x \sim y$, we must have a natural number m, an element $a \in F$ and some $b \in acl(\bar{a})$ so that $my = ax + b$. In fact, it is easy to see that one can again always choose $m = 1$; in that case, a must be a root of unity. Fix some particular x_1 not in the algebraic closure of $\bar{a}$ (we may have to shift to an elementary extension of the vector space to do this) and let $x_1, \ldots, x_k$ list the $\sim$-class of x_1. For each $j = 1, \ldots, k$ choose a root of unity $\gamma_j \in F$ and an element b_j in the algebraic closure of $\bar{a}$ with $x_j = \gamma_j x_1 + b_j$. It is not hard to see that $\{\gamma_j : j = 1, \ldots, k\}$ is a multiplicative subgroup of F. Also, we may check that for $j \neq \ell$, that $\gamma_j \neq \gamma_\ell$. Let γ be a generator for this group. Without loss of generality, then, $x_j = \gamma^{j-1} x_1 + b_j$. (We reorder the x_j's if necessary.)

Write b for b_2. Our next claim is that, in fact, for each $j = 2, \ldots, k$ we have that $b_j = (1 + \gamma + \cdots + \gamma^{j-2})b$. Indeed, suppose that this is true for j. We have that x_j is also not algebraic in $\bar{a}$, so it also satisfies the formula $\gamma x + b \sim x$; thus, one of the x_ℓ's is $\gamma(x_j + b)$. It is apparent that the only possibility is x_{j+1}. Thus, our equivalence class is

$$\{x_1, \ldots, \gamma^j x_1 + \frac{\gamma^j - 1}{\gamma - 1}(b), \ldots, \gamma^{k-1} x_1 + \frac{\gamma^{k-1} - 1}{\gamma - 1}(b)\}.$$

Here we have assumed that $\gamma \neq 1$, of course; the possibility also exists that the equivalence relation is the identity. By altering a finite number of classes, we may assume that every equivalence class has the above form. If we set $c = \frac{1}{\gamma - 1} b$, it is easily seen that we then have $x \sim y$ if and only if $(x + c)^k = (y + c)^k$.

Our next objective is to observe that we may assume that $c = 0$. Indeed, if we replace $(\mathbb{C}; +, \lambda_a : a \in F)$ by its image under the map $x \mapsto x + d$ for $d = \frac{1}{\gamma - 1} c$, the graph of $+$ goes to the graph of the function $(x, y) \mapsto x +^* y = x + y - d$ and the graph of λ_a goes to the graph of $\lambda_a(x - d) + d$, which function we call λ_a^*. In particular, the graph of the equivalence relation

$$x + c = y + c \vee x + c = \gamma(y + c) \vee \cdots \vee x + c = \gamma^{k-1}(y + c)$$

is taken to the graph of the equivalence relation

$$x = y \vee x = \gamma^* y \vee \cdots \vee x = (\gamma^*)^{k-1} y.$$

$$\begin{tikzcd}
\mathcal{M} \arrow[dd, bend right] & & & \\
\mathcal{A} \arrow[r, "\cong"] \arrow[d] & (\mathbb{C}; +) \arrow[r] \arrow[d] & (\mathbb{C}; +^*) \arrow[d] \arrow[dr, "h"] & \\
\mathcal{N} \arrow[r] & \mathbb{C}/{\sim} \arrow[r] & \mathbb{C}/{\sim^*} \arrow[r] & (\mathbb{C}; \ldots)
\end{tikzcd}$$

Under this replacement, then, we effectively have $c = 0$. (This trick, again, is "moving the identity" of the group.) Of course then, the function $h : \mathbb{C} \longrightarrow \mathbb{C}$ where $h(x) = x^n$ induces a definable bijection between $\mathbb{C}/{\sim}$ and $\mathbb{C}$.

The upshot of our observations to date is that, under the assumption that the group we have in $\mathcal{M}^{\text{eq}}$ is constructibly isomorphic to the additive group of the complex numbers, we have the following characterization up to a finite definable cover of the reduct: it is the image under the function $h(x) = x^n$ of a structure of the form $(\mathbb{C}; +, \lambda_a : a \in F)$ as described above. The nth roots of unity must be in F.

It should be mentioned that the image under h of the graph of λ_a is the graph of λ_{a^n}. The image of the graph of $+$ is the zero set of the following homogeneous polynomial of degree n:

$$\Pi_{i=0}^{n-1}(z - (x^{\frac{1}{n}} + \gamma^i y^{\frac{1}{n}})^n).$$

For $n = 2$, this polynomial is $x^2 + y^2 + z^2 - 2(xy + xz + yz)$. Finally, it is worth noting here that including the image of the graph of $+$ and of the λ_as is sufficient to give all the structure on the image of the vector space. (Section 8, again.)

6 In Case the Group is an Elliptic Curve

Here our aim is to give a description of those linear reducts $\mathcal{M}$ of the complex field in which the group we find in $\mathcal{M}^{\text{eq}}$ is constructibly isomorphic to an elliptic curve with its group operation. As always, our description will be complete only up to a finite constructible cover.

Our first observation is that to get a complete list of all such reducts (up to a finite definable cover), it is not necessary to consider every elliptic curve. One representative from each isogeny class is sufficient, and we can choose whichever representative we like. To see this, notice that if E_1 and E_2 are isogenous then there is a definable bijection between E_1/A_1 and E_2 for some finite subgroup A_1 of E_1. Thus for any quotient of E_2 by an equivalence relation definable in some linear expansion of the group structure, there is a constructibly isomorphic quotient of some linear expansion of E_1. This remark will be of use to us later. See the following picture.

$$\begin{tikzcd}
 & \mathcal{E}_1 \arrow[dl, shift left] \arrow[d] \\
\mathcal{E}_2 \arrow[ur, shift left] \arrow[r] \arrow[d] & \mathcal{E}_1/A_1 \arrow[d] \\
\mathcal{E}_2/\sim_2 \arrow[r] & \mathcal{E}_1/\sim_1
\end{tikzcd}$$

Notice that if E is an elliptic curve, given by the equation $y^2 = x^3 + ax + b$, then the equivalence relation $\sim$ on E given by $(x, y) \sim (z, w)$ if and only if $x = z$ is definable in the pure group structure. This is because (x, y) and $(x, -y)$ are inverses of one another. Obviously, there is a constructible isomorphism between $E/\sim$ and $\mathbb{C}$ given by the projection π_x of E onto the x-coordinate. Therefore $\pi_x(E; \oplus)$ is a linear structure on $\mathbb{C}$.

$$\begin{tikzcd}
\mathcal{M} \arrow[dd, bend right] & & \\
\mathcal{A} \arrow[r, "\cong"] \arrow[d] & (\mathcal{E}; \oplus) \arrow[d] \arrow[dr, "\pi_x"] & \\
\mathcal{N} \arrow[r] & \mathcal{E}/\sim \arrow[r] & \pi_x(\mathcal{E})
\end{tikzcd}$$

Similarly, if E has complex multiplication, say its ring of definable endomorphisms is R and R' is some subring of R, then $\pi_x(E; \oplus, r : r \in R')$ is also linear. We will see that up to a finite cover, this gives a nearly complete list of the reducts arising in this way.

Recall that the ring R embeds into $\mathcal{Q}[\sqrt{-d}]$ for some square-free positive integer d. In case $d \neq 1, 3$ (or in case E does not have complex multiplication) the only elements of R which are natural number multiples of roots of unity are the (ordinary) integers. So, as above for the case where the group was the multiplicative group of the field, if $\sim$ is a definable equivalence relation with finite classes on E and (x_1, y_1) is an element of E not algebraic over the parameters used in the definition of $\sim$, we can write the $\sim$-class of (x_1, y_1) as a union of two sets. One of them will be the coset $(x_1, y_1) \oplus S$ for some finite subgroup S of E. The other is either empty or has the form $(x_1, -y_1) \oplus (c, d) \oplus S$ for a constant (c, d) in the algebraic closure of parameters mentioned.

First notice that we may assume that S is trivial. In case E doesn't have complex multiplication, it is immediate that the groups E and E/S are definably isomorphic; since the group is the full structure, we just pass to E/S.

If we do have complex multiplication, the full structure on E may also involve the scalar multiplications from some R' contained in the ring R of definable endomorphisms. S may not be an R'-submodule, but $S + [r]S$ is, if R' is generated by r as an Abelian group over $\mathcal{Z}$. Suppose that we first "fatten" the equivalence relation so that a typical class is $(x_1, y_1) \oplus (S + [r]S) \cup (x_1, -y_1) \oplus (c, d) \oplus (S + [r]S)$. We then replace $\mathcal{N}$ by $\mathcal{N}/\cong$ for some definable equivalence relation $\cong$ on $\mathcal{N}$. Then pass to $E/(S + [r]S)$. In either case, we lose no generality in assuming $S = \{0\}$.

Now we see that the first of the two possibilities (that the equivalence classes are exactly the S-cosets) is out. There is no constructible bijection between a definable factor of $\mathbb{C}$ and an elliptic curve. This follows from the bowdlerized result (3.5.2.) of Hurvitz quoted above.

Now we try a trick similar to one we used in the multiplicative case (and in the additive case, too). That is, we find an element $(c', d') \in E$ so that $(c', d') \oplus (c', d') = (c, d)$ and consider the addition with (c', d') as the new zero. That is, we look at the image of our structure $\mathcal{E}$ under the map $(x, y) \mapsto (x, y) \oplus (c', d')$. (Actually, because this map is definable in the group $(E; \oplus)$, we haven't really altered our structure at all, but this is the right way to regard it.) The effect of this is to allow us to assume without loss of generality that (c, d) is zero. With these two assumptions, we have indeed that our equivalence relation is just the one which identifies (x, y) and $(x, -y)$ after all, just as in the previous diagram.

It may be worth noting that the image of the graph of $\oplus$ under the projection π_x is the relation $R(x, y, z)$ which holds if and only if

$$(xy + xz + yz - a)^2 = 4(xyz + b)(x + y + z).$$

(This for the curve $y^2 = x^3 + ax + b$, as usual.)

We are left with the cases when we have complex multiplication and the ring embeds into $\mathcal{Q}[i]$ or into $\mathcal{Q}[\sqrt{-3}]$. For each of these two possibilities, there is a single isogeny class, so we may choose our favorite representative for each to investigate what happens for these situations. A good candidate in the first case is the curve $y^2 = x^3 + x$, for the second $y^2 = x^3 + 1$.

Now suppose that $\sim$ is an equivalence relation definable on the structure $(E; +, r : r \in R)$ where E is the elliptic curve $y^2 = x^3 + x$ and R its ring of all definable endomorphisms. Note that R is generated as an additive group by the identity and the map $[i]$ where $[i](x, y) = (-x, iy)$. Then we can find a fixed finite set of elements

$$\{(c_0, d_0), \dots, (c_{j-1}, d_{j-1}), (c_j, d_j), \dots, (c_{k-1}, d_{k-1}),$$
$$(c_k, d_k), \dots, (c_{\ell-1}, d_{\ell-1}), (c_\ell, d_\ell), \dots, (c_{m-1}, d_{m-1})\}$$

so that for almost every $(x_0, y_0) \in E$, the $\sim$-class of (x_0, y_0) is the union of the following four sets:

1. $\{(x_0, y_0) \oplus (c_n, d_n) : n < j\}$,
2. $\{[-1](x_0, y_0) \oplus (c_n, d_n) : j \le n < k\}$,
3. $\{[i](x_0, y_0) \oplus (c_n, d_n) : k \le n < \ell\}$, and
4. $\{[-i](x_0, y_0) \oplus (c_n, d_n) : \ell \le n < m\}$.

We have dealt with the cases where the sets in (3) and (4) are empty, so assume that the set in (3) is not, or the one in (4) is not; this implies that all are nonempty. As usual, we must have that $\{(c_n, d_n) : n < j\}$ is a subgroup of E and the other three collections of (c_n, d_n)'s are cosets of this group.

Altering notation slightly, we have that there is a finite subgroup S of E and $e \in E$ so that for (almost all) $x, y \in E$, we have $x \sim y$ if and only if

$$y \oplus [-1]x \in S \vee y \oplus [-i]x \in e \oplus S \vee y \oplus x \in ([i]+[1])e \oplus S \vee y \oplus [i]x \in [i]e \oplus S.$$

(For x generic, suppose that the coset in (3) above is $[i]x \oplus e \oplus S$; then $x \sim [i]x \oplus e \sim [i]([i]x \oplus e) \oplus e = [-1]x \oplus ([1]+[i])e \sim [i]([-1]x \oplus ([1]+[i])e) \oplus e = [-i]x \oplus [i]e$.)

I claim that in fact S must be a submodule of E; that is, S is closed under $[i]$. For fix $x \in E$ not algebraic in e. If $s \in S$ and $y = [i]x \oplus e \oplus s$, then $y \sim x$ and also $x \oplus [i]y = [i]e \oplus [i]s$. Further, y is not algebraic in e. As $x \sim y$, one of $x \oplus [-1]y \in S$, $x \oplus [-i]y \in e \oplus S$, $x \oplus y \in ([1]+[i])e \oplus S$, or $x \oplus [i]y \in [i]e \oplus S$ must hold. We cannot have $x \oplus [-1]y \in S$ because that would imply that $([1]+[i])y = [i]e + [i]s + s'$ for $s' \in S$, which contradicts our choice of x; y must also be independent from e. Similarly it is impossible that $x \in [-i]y \oplus e \oplus S$ or $x \in [-1]y \oplus ([1]+[i])e + S$. This tells us that we must have $[i]s \in S$, establishing our claim.

By the usual trick, we may now assume that S is trivial. Again, altering a finite number of $\sim$-classes and moving the identity, we get that (without loss of generality) $(x_0, y_0) \sim (x_1, y_1)$ if and only if (x_1, y_1) is $[r](x_0, y_0)$ for $r \in \{1, -1, i, -i\}$.

How to describe $\mathcal{E}/\sim$? The easiest way is as follows; it is a quotient of the structure $\mathcal{M} = \pi_x(\mathcal{E})$ as described above via the identification of x and $-x$ for every x. Thus, if g is the map (on the complex numbers) $x \mapsto x^2$, we see that $\mathcal{E}/\sim$ is constructibly isomorphic to $g(\mathcal{M})$.

Now consider the curve E defined by $y^2 = x^3 + 1$. The ring R here is generated by $[1]$ and $[\eta]$, where $\eta \in \mathbb{C}$ is a primitive cube root of unity and $[\eta](x, y) = (\eta x, y)$. Upon factoring out a subgroup (as above it is an R-submodule), moving the identity, and altering finitely many classes, we get $(x, y) \sim (z, w)$ if and only if $(z, w) = [r](x, y)$ for $r \in R_0$, where R_0 is either $\{1\}$, $\{1, -1\}$, $\{1, \eta, \eta^2\}$, or $\{1, -1, \eta, -\eta, \eta^2, -\eta^2\}$. We have dealt with the first two cases. Case 3 is most easily described as $\pi_y(\mathcal{E})$; this is the only curve (up to isogeny) where when we project to the y-coordinate we get a linear structure. In the fourth case, we can describe the structure as $g(\pi_y(\mathcal{E}))$ for the function $g(x) = x^2$.

Notice that in all cases we have identified our quotient structure with a linear structure on $\mathbb{C}$ itself.

7 Summary, So Far

Here we give the bottom line, so far as we have it. The mention of the function f at the beginning of the statement means that our description is complete only up to a finite cover.

Theorem 7.1 *Let $\mathcal{M}$ be a linear reduct of the complex field. Then there is a constructible function $f : \mathbb{C} \longrightarrow \mathbb{C}$ such that $f(\mathcal{M})$ is one of the following structures:*

1. $(\mathbb{C}; \cdot)$.
2. $g((\mathbb{C}; \cdot))$, *where $g : \mathbb{C} \longrightarrow \mathbb{C}$ is the function $g(x) = x + x^{-1}$.*
3. *Let F be a subfield of $\mathbb{C}$ containing the nth roots of unity and g_n be the function $g_n(x) = x^n$. The structure (obviously depending on F and n) is $g_n((\mathbb{C}; +, \lambda_a : a \in F))$.*

4. *Let $\mathcal{E} = (E; \oplus, r : r \in R')$ be an elliptic curve with the usual addition, where R' is some subring of the ring of all definable endomorphisms of $(E; \oplus)$. Let π_x be the projection onto the first coordinate. Our structure is $\pi_x(\mathcal{E})$.*
5. *Let E be the curve with equation $y^2 = x^3 + x$ and $g : \mathbb{C} \longrightarrow \mathbb{C}$ the function $g(x) = x^2$. The structure is $g(\pi_x(\mathcal{E}))$ for $\mathcal{E}$ as in the last item.*
6. *Let E be the elliptic curve with equation $y^2 = x^3 + 1$ and $\mathcal{E}$ as in (4). The structure is $\pi_y(\mathcal{E})$, the projection onto the second coordinate.*
7. *Let $g(x) = x^2$ and $\mathcal{E}$ as (6). The structure here is $g(\pi_y(\mathcal{E}))$.*

We remind the reader of several things; the structure in (2) is $(\mathbb{C}; P)$ where $P(x, y, z)$ holds if and only if $x^2 + y^2 + z^2 = 4 + xyz$. In (3), for $n = 1$, we get just the vector space. In general in (3), we need only include predicates for the image of the graph of $+$ and the λ_a's to give the full structure here. Always g_n(graph of λ_a) is the graph of λ_{a^n}; we have no nice description of the image of the graph of $+$ in general, but for $n = 2$, it is the relation P where $P(x, y, z)$ holds if and only if $x^2 + y^2 + z^2 = 2(xy + xz + yz)$.

The structure in case (4) is just $(\mathbb{C}; P_{a,b})$ for the curve $y^2 = x^3 + ax + b$ if it doesn't have complex multiplication. Here $P_{a,b}(x, y, z)$ if and only if $(xy + xz + yz - a)^2 = 4(xyz + b)(x + y + z)$. To get a full list of these structures, we need only one E from each isogeny class; if E_1 and E_2 are isogenous, there is a constructible isomorphism between $\pi_x(\mathcal{E}_1)$ and $\pi_x(\mathcal{E}_2)$. It is of course given by a fractional linear transformation and can be absorbed into the function f mentioned at the beginning.

There are several as yet unproved claims in the past two paragraphs. They will be settled in Section 8.

There is clearly some redundancy on this list. (1) is a finite cover of (2), so could have been left off; for the same reason, we could have omitted (6), a finite cover of (7). Also, several special cases of the others could be omitted. We chose not to do so, for the reason that the list is minimally complete in the following sense. Suppose that $\mathcal{A}$ is a constructible linear group and $\sim$ a definable equivalence relation on $\mathcal{A}$. We don't necessarily have $\mathcal{A}/\sim$ on our list, but our constructions tell us there is something constructibly isomorphic to it on this list. So if we have a reduct $\mathcal{M}$ of $\mathbb{C}$ which happens to be such an $\mathcal{A}/\sim$, there is a fractional linear transformation σ such that $\sigma(\mathcal{M})$ is on this list.

There are two major shortcomings of this list, as I see it. The most obvious is this business of the cover by some structure $\mathcal{M}$ via f. Given any constructible function $f : \mathbb{C} \longrightarrow \mathbb{C}$, we can always endow the second copy of $\mathbb{C}$ with any of the structures here and lift it to get the canonical (minimal) cover via f. But there are in general surely many other covers via f, at least in some cases; what are the possibilities? I think this may be a very difficult question indeed, but perhaps there is something sensible to be said about it.

Another flaw is that we have not made clear when two differently presented reducts on our list are in fact the same. A related question is when one of them is a reduct of another. And of course these questions can be asked of the covers as well. For the moment, we will mention a few simple facts about this situation. See Section 9 for more information.

We note first that it makes sense to ask these questions on three different levels: the abstract, the constructible, and as actual reducts of the field. For example, for

a, b independent transcendentals, it is quite clear that all the structures mentioned in item (4) above are abstractly isomorphic. They are not constructibly isomorphic, however, unless the corresponding curves are isogenous; in that case, they are. When they give exactly the same reduct, that is, when $P_{c,d}$ is definable in $(\mathbb{C}; P_{a,b})$ and vice versa, will be explored below; in fact, it almost never happens.

Observe that the group that one finds in $\mathcal{M}^{\mathrm{eq}}$ of a linear reduct is essentially unique in the constructible sense. That is, if there were to be a constructible isomorphism between $\mathcal{M}_1$ and (even a reduct of) $\mathcal{M}_2$, both as in the statement of the theorem, they would both have to be covers of the structures in (1) or (2), or both of structures in (3), or both of one of the later items on the list; further, in the latter case, the elliptic curves in question would have to be isogenous.

8 On Squashed Modules

All the reducts identified above are of the following form, up to a constructible isomorphism and a finite cover: We start with a strongly minimal module $\mathcal{M} = (M; +, r : r \in R)$ such that the underlying group has unbounded exponent. We let $\gamma \in R$ be a primitive nth root of unity in the center of R. Let $\sim$ be the relation on M defined by

$$x \sim y \Leftrightarrow y = x \vee y = \gamma x \vee y = \gamma^2 x \vee \cdots \vee y = \gamma^{n-1} x.$$

Our structure is then $\mathcal{M}/{\sim}$. We will find a natural language for such a structure and investigate it somewhat.

First we recall a few facts about the module itself. We may assume that R acts faithfully on M, so that for any nonzero $r \in R$, r is onto and has finite kernel. (This uses strong minimality.) R will embed into a division ring D. For any nonzero $r \in R$, there is, definable in the structure $\mathcal{M}$, an endomorphism s such that $rs = sr = m$ for some positive integer m; call s a *semi-inverse* for r. We will assume, with no loss of generality, that $s \in R$ for some such s. This implies that if $ra = sb$ for some $r, s \in R$ with s nonzero and $a, b \in M$, then $r'a = mb$ for some positive integer m and some $r' \in R$. Note that the collection of torsion elements of M is either just $\{0\}$ or is countably infinite, as M must be divisible as an Abelian group.

Let P be the image of the graph of $+$ under the map $x \mapsto x/{\sim}$ from M to $M/{\sim}$. So $(a/{\sim}, b/{\sim}, c/{\sim}) \in P$ exactly if $c = \gamma^i a + \gamma^j b$ for some natural numbers i and j. For any $r \in R$, let r^* be the image of the graph of r; note that r^* is the graph of a function on $M/{\sim}$. We call the structure $\mathcal{N} = (M/{\sim}; P, r^* : r \in R)$ a *squashed module*. It is clearly a reduct of $\mathcal{M}/{\sim}$ and so any automorphism of the module induces an automorphism of $\mathcal{N}$. In fact, the reverse is also true, and this fact easily implies that the squashed module structure is the entire structure of $\mathcal{M}/{\sim}$.

Proposition 8.1 *Let $\mathcal{M}$ be a module as specified above, and $\mathcal{N}$ the corresponding squashed module for the root of unity γ. Let α be an automorphism of $\mathcal{N}$. Then α lifts to an automorphism $\bar{\alpha}$ of $\mathcal{M}$; that is, $\bar{\alpha}$ is an automorphism of the module, and $\bar{\alpha}(a) \in \alpha(a/{\sim})$ for every $a \in M$.*

Proof Actually, there will be exactly n possible choices for $\bar{\alpha}$, as will soon be clear. We may assume that there is some nontorsion element $a \in M$ by passing if necessary to an elementary extension of $\mathcal{M}$; fix such an a. Let $\bar{\alpha}(a)$ be any element of $\alpha(a/{\sim})$. We define $\bar{\alpha}(b)$ in one of two ways, depending on whether or not a and b are related in the module structure.

First suppose that there is a positive integer m and some nonzero $r \in R$ so that $mb = ra$; note that b must be nontorsion, too. Choose such an m, r with m minimal. Now we choose $\bar{\alpha}(b)$ in $\alpha(b/\sim)$ so that $m\bar{\alpha}(b) = r\bar{\alpha}(a)$. To see that this is possible, note that because we have $m^*(b/\sim) = r^*(a/\sim)$ we have also $m^*(\alpha(b/\sim)) = r^*(\alpha(a/\sim)) = \alpha(ra/\sim)$. Letting $c, \gamma c, \ldots, \gamma^{n-1}c$ list $\alpha(b/\sim)$, we must have $\gamma^i r\bar{\alpha}(a) = m\gamma^j c$ for some i, j. But then $r\bar{\alpha}(a) = m\gamma^{j-i}c$. On the other hand, $m\gamma^i c \neq m\gamma^j c$ for i, j different mod n, because otherwise c is torsion, which easily implies that b is torsion. Thus $\bar{\alpha}(b)$ is well defined. Also, if $\gamma^i \neq 1$, we will have $m\bar{\alpha}(\gamma^i b) = r\gamma^i\bar{\alpha}(a)$; if we had $\bar{\alpha}(\gamma^i b) = \bar{\alpha}(b)$, we would also have $m\bar{\alpha}(\gamma^i b) = r\bar{\alpha}(a)$, so that $\bar{\alpha}(a)$ would be torsion, which it isn't. Thus, our definition of $\bar{\alpha}(b)$ is, so far, one-to-one.

Now suppose that b does not satisfy the condition above. We choose $\bar{\alpha}(b) \in \alpha(b/\sim)$ so that $\bar{\alpha}(b) + \bar{\alpha}(a) \in \alpha(a + b/\sim)$. This is possible because we have $P(a/\sim, b/\sim, a + b/\sim)$ and hence $P(\alpha(a/\sim), \alpha(b/\sim), \alpha(a + b/\sim))$. Choose $c \in \alpha(b/\sim)$ and $d \in \alpha(a + b/\sim)$; we must have $\gamma^i\bar{\alpha}(a) + \gamma^j c = d$ for some i, j; $\bar{\alpha}(b)$ is then $\gamma^{j-i}c$. To see that this is well defined, note that if we have $\bar{\alpha}(a) + \gamma^i c$ and $\bar{\alpha}(a) + \gamma^j c$ in the same $\sim$-class, we would have $\bar{\alpha}(a)+\gamma^i c = \gamma^k(\bar{\alpha}(a)+\gamma^j c)$. If $\gamma^k \neq 1$, we get that $(1-\gamma^k)\bar{\alpha}(a) = (\gamma^{j+k}-\gamma^i)c$ so that $(1-\gamma^k)^*\alpha(a/\sim) = (\gamma^{j+k}-\gamma^i)^*\alpha(b)$ and hence $(1-\gamma^k)a/\sim = (\gamma^{j+k}-\gamma^i)b/\sim$. This implies that $(1-\gamma^k)a = rb$ for some $r \in R$ and this contradicts our assumption on b. So $\gamma^k = 1$ and $\gamma^i c = \gamma^j c$.

The demonstration that $\bar{\alpha}$ remains one-to-one is quite similar to this proof that it is well defined. This shows that it is onto, as well. Every class is $\alpha(b/\sim)$ for some b and we have $\bar{\alpha}(b) \in \alpha(b/\sim)$. Also, distinct elements of $b/\sim$ go to distinct elements of $\alpha(b/\sim)$. One final point: for those elements b such that $\gamma^i b = \gamma^j b$ where $\gamma^i \neq \gamma^j$, the class $\alpha(b/\sim)$ is easily seen to have fewer than n elements as well.

Assume for the moment that there are infinitely many torsion elements. Next we show that $\bar{\alpha}$, when restricted to the torsion elements, is an endomorphism. It should be noted first that it permutes the torsion elements, because α must permute the set $\{b/\sim\colon mb = 0\}$ for every natural number m. (In fact α permutes the set $\{b/\sim\colon rb = 0\}$ for any $r \in R$.) Suppose that $mb = 0$ and m is minimal (b is fixed for now). Then $ma = m(a + b)$ so that by definition, we have that $m\bar{\alpha}(a) = m\bar{\alpha}(a + b)$. Also $m(2a + b) = 2ma$, so that $m\bar{\alpha}(2a + b) = 2m\bar{\alpha}(a)$; put these together and we get that $m[\bar{\alpha}(2a + b) - \bar{\alpha}(a) - \bar{\alpha}(a + b)] = 0$. Now because we have $P(\alpha(a/\sim), \alpha(b/\sim), \alpha(a + b/\sim))$ and $\bar{\alpha}(x) \in \alpha(x/\sim)$ for every x, we must have that $\bar{\alpha}(2a + b) = \gamma^i\bar{\alpha}(2a) + \gamma^j\bar{\alpha}(b)$ for some i, j. Because $\bar{\alpha}(a) + \bar{\alpha}(b) \in \alpha(a + b/\sim)$ we have that $\bar{\alpha}(a + b) = \gamma^k[\bar{\alpha}(a) + \bar{\alpha}(b)]$ for some k. Now we have $\bar{\alpha}(2a) = 2\bar{\alpha}(a)$, too. Putting all these together yields

$$m[(2\gamma^i - 1 - \gamma^k)\bar{\alpha}(a) + (\gamma^j - \gamma^k)\bar{\alpha}(b)] = 0.$$

As b is torsion but a is not, we must have $2\gamma^i - 1 - \gamma^k = 0$. Embedding the subring $\mathcal{Z}[\gamma]$ of R into the complexes is the easiest way to see that this can only happen if $\gamma^k = \gamma^i = 1$ and thus that $\bar{\alpha}(a + b) = \bar{\alpha}(a) + \bar{\alpha}(b)$.

Now let b, c be any torsion elements; we know that

$$\bar{\alpha}(a + b + c) = \bar{\alpha}(a) + \bar{\alpha}(b + c) = \gamma^i\bar{\alpha}(a + b) + \gamma^j\bar{\alpha}(c)$$

$$= \gamma^i[\bar{\alpha}(a) + \bar{\alpha}(b)] + \gamma^j\bar{\alpha}(c).$$

We must have $\gamma^i = 1$ as a is not torsion. If $\gamma^j = 1$, too, then of course $\bar{\alpha}(b+c) = \bar{\alpha}(b)+\bar{\alpha}(c)$. So suppose not. Similarly, we find k so that $\bar{\alpha}(a+b+c) = \bar{\alpha}(a)+\bar{\alpha}(c)+\gamma^k\bar{\alpha}(b)$ and we may assume that $\gamma^k \neq 1$. Canceling the $\bar{\alpha}(a)$s we see that for any torsion element b, the set

$$S_b =^{\text{def}} \{c : c \text{ torsion and } \bar{\alpha}(b+c) \neq \bar{\alpha}(b) + \bar{\alpha}(c)\}$$

is finite. For our given b and c, choose a torsion element d so that $d \notin S_c \cup S_{b+c}$ and $c + d \notin S_b$. Then

$$\bar{\alpha}(b+c+d) = \bar{\alpha}(b+c) + \bar{\alpha}(d) = \bar{\alpha}(b) + \bar{\alpha}(c+d) = \bar{\alpha}(b) + \bar{\alpha}(c) + \bar{\alpha}(d)$$

and so $\bar{\alpha}$ is an additive automorphism of the torsion subgroup.

Now for any $r \in R$ and torsion b, we must have that $\bar{\alpha}(rb) = \gamma^i r\bar{\alpha}(b)$ because we have $r^*(\alpha(b/\sim)) = \alpha(rb/\sim)$. Similarly, $\bar{\alpha}((r-1)b) = \gamma^j(r-1)\bar{\alpha}(b)$. So $\gamma^i r\bar{\alpha}(b) = [\gamma^j(r-1)+1]\bar{\alpha}(b)$. That is, $[(\gamma^i - \gamma^j)r + 1 - \gamma^i]\bar{\alpha}(b) = 0$. Suppose that r is not one of the finitely many elements of R so that $(\gamma^i - \gamma^j)r + 1 - \gamma^i = 0$ for some i, j. Then the set of torsion elements c so that $\bar{\alpha}(rc) \neq r\bar{\alpha}(c)$ is finite. Choosing some d such that neither d nor $b + d$ is in this set, we see that in fact it is empty by calculating $\bar{\alpha}(rb + rd)$ in two ways. So for all but finitely many elements $r \in R$, we have that $\bar{\alpha}$ respects multiplication by r. So for arbitrary $r \in R$ find nonzero $s \in R$ so that neither s nor rs is in the set of exceptions. For torsion b find c so that $sc = b$. Calculating $\bar{\alpha}(rsc) = rs\bar{\alpha}(c) = r\bar{\alpha}(sc)$ shows that, indeed, $\bar{\alpha}$ restricted to the torsion elements is a module automorphism.

Now we "alter" the definition of $\bar{\alpha}$ on the nontorsion elements. (In fact, the "new" $\bar{\alpha}$ will turn out to be the same, after the fact.) We temporarily rename the "old" purported automorphism $\bar{\alpha}'$ and define a new one $\bar{\alpha}$ as follows: it agrees with $\bar{\alpha}'$ on the torsion elements. Fix any torsion element b not in the kernel of any nonzero $\gamma^i - 1$, nor in the kernel of $2-\gamma^i-\gamma^j$ unless it is zero. For any nontorsion c, choose $\bar{\alpha}(c) \in \alpha(c/\sim)$ so that $\bar{\alpha}(c) + \bar{\alpha}(b) \in \alpha(b + c/\sim)$. As above, this is possible, and our assumptions on b assure us that it is well defined.

Simply because every $\bar{\alpha}(c)$ is in the right class, we have for any c, d that $\bar{\alpha}(c + d) = \gamma^i\bar{\alpha}(c) + \gamma^j\bar{\alpha}(d)$ for some i, j. Similarly, $\bar{\alpha}(rc) = \gamma^k r\bar{\alpha}(c)$. Now we show that for any c, we have that $\bar{\alpha}(b + c) = \bar{\alpha}(b) + \bar{\alpha}(c)$. We may assume that c is nontorsion. We have that $\bar{\alpha}(b + c) = \gamma^\ell[\bar{\alpha}(b) + \bar{\alpha}(c)]$; also

$$\bar{\alpha}(2b+c) = 2\gamma^i\bar{\alpha}(b)+\gamma^j\bar{\alpha}(c) = \gamma^k[\bar{\alpha}(b)+\bar{\alpha}(b+c)] = \gamma^k[\bar{\alpha}(b)+\gamma^\ell(\bar{\alpha}(b)+\bar{\alpha}(c))].$$

We must have $\gamma^j = \gamma^{k+\ell}$ and then that $\bar{\alpha}(b)$ is in the kernel of $2\gamma^i - \gamma^k - \gamma^{k+\ell}$, so $\gamma^i = \gamma^k = \gamma^{k+\ell}$ and $\gamma^\ell = 1$.

Fix c and consider the set $S_c = \{d : \bar{\alpha}(c + d) \neq \bar{\alpha}(c) + \bar{\alpha}(d)\}$. We show first that for only finitely many torsion elements c can this set be infinite. We know that $\bar{\alpha}(c + d) = \gamma^i\bar{\alpha}(c) + \gamma^j\bar{\alpha}(d)$ and also that

$$\bar{\alpha}(b+c+d)=\bar{\alpha}(b)+\bar{\alpha}(c+d)=\gamma^k\bar{\alpha}(b+d)+\gamma^\ell\bar{\alpha}(c)=\gamma^k(\bar{\alpha}(b)+\bar{\alpha}(d))+\gamma^\ell\bar{\alpha}(c).$$

The exponents, which we may assume are all between 0 and $n - 1$, depend on d, but for only finitely many d can we have $k \neq j$. Similarly,

$$\bar{\alpha}(b + c + d) = \gamma^m[\bar{\alpha}(b) + \bar{\alpha}(c)] + \gamma^p\bar{\alpha}(d)$$

and again for only finitely many d can we have $p \neq j$. Fix $d \in S_c$ so that $j = k = p$ and then cancel $\bar{\alpha}(d)$. Unless $m = \ell = i$, we have our claim. But if $m = \ell = i$, by canceling the $\bar{\alpha}(c)$'s, too, we see that we must have $m = 1$ for this choice of d, a contradiction. Let us call S the set of such c's.

Now, given arbitrary c, d we show that $\bar{\alpha}(c+d) = \bar{\alpha}(c) + \bar{\alpha}(d)$. We may assume that d is not torsion, and it suffices as above to find $e \notin S_d \cup S_{c+d}$ so that $d + e \notin S_c$. If neither c nor $c + d$ is in S, this is easy. At most one is in S. If $c \in S$, we can easily choose an element $e \notin S_c \cup S_{c+d}$ so that $d + e$ is torsion and hence not in S_c. If $c + d \in S$, we can choose e torsion so that $e \notin S_d$ and $d + e \notin S_c$. Thus $\bar{\alpha}$ is an additive automorphism of M.

Now we show that $\bar{\alpha}$ is a module automorphism of M by considering the set

$$\{r \in R : \text{ for some } c, \bar{\alpha}(rc) \neq r\bar{\alpha}(c)\}.$$

Any witness c that r is in this must be nontorsion, and we will have

$$\bar{\alpha}(rc) = \gamma^i r\bar{\alpha}(c) = \bar{\alpha}((r-1)c) + \bar{\alpha}(c) = [\gamma^j(r-1) + 1]\bar{\alpha}(c).$$

This leaves only finitely many possibilities for r; we dispose of these exactly as above. This completes the proof of the proposition in case there are nonzero torsion elements.

If 0 is the only torsion element, then our module is a vector space. The proof in this case is virtually identical to the above, but easier; we don't have to alter our original $\bar{\alpha}$. Instead we verify directly that $\bar{\alpha}$ as defined initially is a vector space automorphism of the subspace generated by a (easy) and then use elements of this subspace as we employed torsion elements in the above proof. We leave the details to the reader. The proposition is proved. □

From this proposition (applied to some elementary extension of $\mathcal{M}$) and Proposition 1.2 the following is immediate.

Corollary 8.2 *Fix notation as in Proposition 8.1. Then the squashed module structure is the entire structure of $\mathcal{M}/{\sim}$; the squashed module is not a proper reduct.*

In certain cases, we have claimed more than the above. We stated that we need only the predicate P in case $\mathcal{M}$ is the multiplicative group of the complexes, or any elliptic curve without complex multiplication. Both of these cases are dealt with in the following.

Corollary 8.3 *Fix notation as in Proposition 8.1. Suppose that the ring R is just the integers. Then the squashed module structure is the same as the structure $(M/{\sim}; P)$.*

Proof We need to check that for every integer m, the function m^* on $M/\sim$ is definable from P. For $m = 0, 1$ this is clear (1^* is just the identity). For positive m we prove this by induction, so suppose that it is true for every positive integer less than m and $m > 1$. We must have $\gamma = -1$ and then we have $P(a/{\sim}, b/{\sim}, c/{\sim})$ exactly if c is either $a + b$, $a - b$, $b - a$, or $-a - b$. We have then that $y = m^*(x)$ if and only if $P(x, (m-1)^*x, y) \wedge y \neq (m-2)^*x$. For negative m, it is now sufficient to show that -1^* is definable. But $y = (-1)^*x$ exactly if $P(x, y, 0/{\sim}) \wedge y \neq x$. This finishes the proof of the corollary. □

The next result is actually a corollary of the proof of our proposition. We will naturally be applying this in case α is a constructible map.

Corollary 8.4 *Again, we use the notation of the proposition. Suppose that our automorphism α of the squashed module is definable in $\mathcal{M}'^{\mathrm{eq}}$ for some expansion $\mathcal{M}'$ of $\mathcal{M}$. Then the lifted automorphism $\bar{\alpha}$ is also definable in $\mathcal{M}'$.*

Proof Essentially this is because there are only finitely many possibilities for $\bar{\alpha}$, but we can be more explicit. Fix some a outside the kernel of $\gamma^i - 1$ for any $i = 1, \ldots, n-1$ and let $b = \bar{\alpha}(a)$. Consider the subset G of M^2 defined by

$$(x, y) \in G \Leftrightarrow y/{\sim} = \alpha(x/{\sim}) \wedge (y+b)/{\sim} = \alpha(x + a/{\sim}).$$

It is easily checked that the symmetric difference of G and the graph of $\bar{\alpha}$ is finite. □

The following will be used at the end of Section 9.

Proposition 8.5 *Let $\mathcal{M}/\sim$ be a squashed module, P the image of the graph of addition, and c any nonalgebraic element.*

1. *Suppose that b is not algebraic in c and we have that the solutions to $P(c/{\sim}, b/{\sim}, x)$ are the same as those to $P(c'/{\sim}, b/{\sim}, x)$. Then $c/{\sim} = c'/{\sim}$.*
2. *Any automorphism of $\mathcal{M}/\sim$ that fixes every point independent from $c/\sim$ must fix $c/\sim$, too.*

Proof (1) We may obviously assume that $\sim$ is not trivial. Say it has classes of size n. Then there are n solutions to $P(c'/{\sim}, b/{\sim}, x)$, say, $d_i/{\sim}$ for $i < n$. Let γ be an nth root of unity. For each $i < n$, by perhaps altering each d_i in its $\sim$-class and reordering them, we may assume that $c' + \gamma^i b = d_i$. For each i, there are k and ℓ so that $\gamma^k c + \gamma^\ell b = d_i$. We claim first that k doesn't depend on i. Suppose for instance that $\gamma^k c + \gamma^\ell b = d_0$ and $\gamma^m c + \gamma^r b = d_1$. Then

$$\gamma^k c - c' + (\gamma^\ell - 1)b = \gamma^m c - c' + (\gamma^r - \gamma)b = 0.$$

Subtracting again gives a contradiction on our assumptions about c and b unless $\gamma^k = \gamma^m$. We can of course assume that $k = 0$ for each i, at the cost of moving c within its $\sim$-class.

So for some permutation ρ of $\{0, \ldots, n-1\}$ we have that $c + \gamma^{\rho(i)} b = d_i$. Making two additions and noting that $\Sigma_0^{n-1} \gamma^i = 0$, we see that $nc' = nc = \Sigma d_i$. So $c - c'$ is algebraic. Let $i = \rho(0)$. Then $c' + b = d_0 = c + \gamma^i b$ and $c' - c = (\gamma^i - 1)b$. As b is not algebraic, we must have $i = 0$. Thus $c = c'$, proving (1). (2) is now automatic. □

9 When Are Two of Our Reducts the Same?

Suppose that $\mathcal{M}$ and $\mathcal{N}$ are two of the items on our list and σ and τ are fractional linear transformations. We investigate here the possibility that $\sigma(\mathcal{M}) = \tau(\mathcal{N})$. That is, these two structures (both with universe $\mathbb{C}$) have the same definable sets.

We make two trivial observations that will be useful in what comes. We may assume that σ (or τ, if we prefer) is the identity. It is clear that $\sigma(\mathcal{M}) = \tau(\mathcal{N})$ if and only if $\mathcal{M} = \sigma^{-1}\tau(\mathcal{N})$. Similarly, for a particular predicate P, $\sigma(P) = \tau(P)$ if and only if $P = \sigma^{-1}\tau(P)$.

We begin with a general setup, eventually getting down to cases. Each of our structures $\mathcal{M}$ arises as (something constructibly isomorphic to) $\mathcal{A}/\sim$ where $\mathcal{A}$ is one of our particular three kinds of groups and $\sim$ identifies $x, \gamma x, \ldots, \gamma^{n-1}x$ for γ some primitive nth root of unity. The main structure on $\mathcal{M}$ is given by the image P of the graph of the group operation under the projection $\pi : \mathcal{A} \longrightarrow \mathcal{M}$. Inside the structure $(\mathbb{C}; P)^{\mathrm{eq}}$ we can find the group $\mathcal{A}$ and the projection π. Thus if $\mathcal{M}$ and $\mathcal{N}$ come from our list and $\mathcal{M} = \tau(\mathcal{N})$, we have two predicates P_1 and $P_2 = \tau(P_1)$. Inside $\mathcal{M}^{\mathrm{eq}}$ we can find Abelian groups $(\mathcal{A}_i; \oplus_i)$ and projections $\pi_i : A_i \longrightarrow \mathbb{C}$ so that the image of the graph of $\oplus_i$ is P_i. The fibers of π_i have the form $\{x, \gamma_i x, \ldots, \gamma_i^{n_i-1}x\}$ as above.

Now $A_1 \times A_2$ is an Abelian structure (being an Abelian group defined in a linear structure) and $\pi_2^{-1}\pi_1$ is a definable subset. As usual, we suppress mention of the parameters necessary to define it. Thus, according to [2], it is a Boolean combination of cosets of definable subgroups. Pick $x \in A_1$ generic and then y so that $(x, y) \in \pi_2^{-1}\pi_1$; so there is a definable subgroup $B \leq A_1 \times A_2$ with $(x, y) + B$ contained in $\pi_2^{-1}\pi_1$ except for finitely many elements. We also have $(x, \gamma_2^i y) \in \pi_2^{-1}\pi_1$ for each $i = 0, \ldots, n_2 - 1$ and so there is $B_i \leq A_1 \times A_2$ definable with $(x, \gamma_2^i y) + B_i$ contained in $\pi_2^{-1}\pi_1$ for each i. Now we cannot have $(x, \gamma_2^i y) + B_i = (x, \gamma_2^j y) + B_j$ for $i \neq j$ as this would imply that $(0, (\gamma_2^i - \gamma_2^j)y) \in B_i$, contradicting the genericity of y (hence of x).

From the above, it should be clear that the symmetric difference of $\pi_2^{-1}\pi_1$ and the union of the $(x, \gamma_2^i y) + B_i$s is finite and that each B_i is the graph of a homomorphism from A_1 to A_2. Considering the cosets $(\gamma_1^i x, y) + B_i$ for $i = 0, \ldots, n_1 - 1$ we come to two quick conclusions. First, each B_i is the graph of an isomorphism from A_1 to A_2. Second, $n_1 = n_2$; we will call it n in what follows.

We have f, a definable isomorphism of the groups A_1 and A_2, and it is the graph of B_0, where $(x, y) + B_0$ is contained in $\pi_2^{-1}\pi_1$ except for finitely many points. So $(0, y - f(x)) \in (x, y) + B_i$, too. Let $e = y - f(x)$; it is algebraic in the parameters needed to define π_1 and π_2. x is not, so e doesn't depend on the choice of x. Because $(z, \gamma_2 y) \in \pi_2^{-1}\pi_1$ exactly if $z = \gamma_1^j x$ for some j, we must have that $\gamma_2 y - f(\gamma_1^j x) = e$ for some j. Thus $(\gamma_2 - 1)e = f(\gamma_1^j x) - \gamma_2 f(x)$. Now fix an independent (from everything so far) element x'; we must also have $(\gamma_2 - 1)e = f(\gamma_1^j x') - \gamma_2 f(x')$. Subtracting, we get

$$f(\gamma_1^j(x - x')) = \gamma_2 f(x - x').$$

But $x - x'$ is just as generic as x, so $f(\gamma_1^j(x - x')) - \gamma_2(x - x') = (\gamma_2 - 1)e$. This tells us that $\gamma_2 e = e$ and the calculation tells us that $x \sim_1 z$ if and only if $f(x) \sim_2 f(z)$.

Let's take stock before proceeding to exploit this picture further. If $\mathcal{M} = \tau(\mathcal{N})$ and $\mathcal{M}$ and $\mathcal{N}$ both come from our list, then they must both come from the same *item* on our list. There is a definable isomorphism between the corresponding groups, and it takes $\sim_1$ to $\sim_2$. Thus if for instance we had $g((\mathbb{C}; \cdot)) = \tau(\mathbb{C}; \cdot)$ (items 1 and 2), we would have a bijection carrying the trivial equivalence relation to one with classes of size two, obviously nonsense. Also, if both $\mathcal{M}$ and $\mathcal{N}$ come from item 3, say $\mathcal{M} = g_m((\mathbb{C}; +, \lambda_a : a \in F_1))$ and $\mathcal{N} = g_n((\mathbb{C}; +, \lambda_a : a \in F_2))$ we must have that $m = n$ and F_1 is isomorphic to F_2. (It's hard to say here what it means for them to be the "same", as they act on different copies of the group.)

Also, if both groups are elliptic curves (item 4), they must be *isomorphic* (in the sense of algebraic geometry). That is, there is a constructible isomorphism between them. But we will see below that if E_1 and E_2 are distinct (even isomorphic) elliptic curves, the reducts corresponding to them in the sense of item 4 are distinct. (But one will be the image of the other under some f.l.t. τ, of course.)

We return to our general picture for a moment; the relation $\pi_2 f \pi_1^{-1}$ on $\mathbb{C}$ is actually the graph of a bijection. This bijection takes P_1 to P_2. Also $\pi_1(z) = x$ if and only if $\pi_2(f(z) \oplus_2 e) = x$; that is, if $g(z) = f(z) \oplus_2 e$, then the relation $\pi_2 g \pi_1^{-1}$ is the identity. This fact will allow us to compute $\pi_2 f \pi_1^{-1}$, given e. As $\gamma_2 e = e$, there are (usually) only a few possibilities for it, depending on the particular case, which we now get to.

We will actually be assuming that A_1 is one of our specified groups (the multiplicative group, the additive group, or an elliptic curve in the form $y^2 = x^3 + ax + b$) and then use the identification made above with such $A_1/\sim_1$ and the appropriate structure on our list to draw conclusions about the latter.

Suppose we are in situation 1; as $\gamma_2 = 1$, there's no constraint on e. π_1 is the identity. The function $\pi_2 f$ is easily seen to be just multiplication by $b = \pi_2(e)^{-1}$. So τ(graph of $\cdot$) is the graph of $(x, y) \mapsto xyb^{-1}$; of course this graph is definable in $(\mathbb{C}; \cdot)$. A straightforward calculation tells us that τ is either $x \mapsto bx$ or $x \mapsto bx^{-1}$. Thus by the first observation, $\sigma(\mathbb{C}; \cdot) = \tau(\mathbb{C}; \cdot)$ if and only if $\sigma(x) = \tau(bx)$ or $\sigma(x) = \tau(bx^{-1})$ for some particular b. (In the following cases, we will omit this final, trivial step.)

In case 2, we have $e^2 = 1$ in the copy A_2 of $(\mathbb{C}; \cdot)$. If e is the identity, then so is $\pi_2 f \pi_1^{-1}$, so $\tau(P) = P$ where P is as usual the image of the graph of $\cdot$ under π_1. So by Corollary 8.4, τ lifts to a definable automorphism of A_1. The only such are the identity and $x \mapsto x^{-1}$, and both of these project to the identity under π_1. Using our identification $x/\sim_1 \mapsto x + x^{-1}$ of $\mathbb{C}/\sim_1$ and $\mathcal{M} = (\mathbb{C}; P)$ we see that in this case τ is the identity. If e is A_2's version of -1, then $f^{-1}g(x) = -x$ for every $x \in A_1$. So $g^{-1}f(x) = -x$, too. So $\pi_2 f(x) = \pi_2 g(-x)$ which gives us that $\pi_2 f \pi_1^{-1}(x/\sim_1) = -x/\sim_1$. Using the identification again, we see that this function remains $x \mapsto -x$, so we have $\tau(P) = \sigma(P)$ where $\sigma(x) = -x$. But this implies $\tau = \sigma$ just as above. Now $x \mapsto -x$ is definable in $(\mathbb{C}; P)$; recall that $P(x, y, z)$ if and only if $x^2 + y^2 + z^2 = 4 + xyz$, so that $P(x, y, -2) \Leftrightarrow y = -x$. Bottom line in this case: $\tau(\mathbb{C}; P) = (\mathbb{C}; P)$ if and only if τ is the identity or $x \mapsto -x$.

We move to case 3, and assume for the moment that $n > 1$, so γ_2 is not the identity. Then $e = 0$ and $\pi_2 f \pi_1^{-1}$ is the identity. So τ is a constructible isomorphism of the structure $(\mathbb{C}; +)/\sim_1$. Now τ lifts to a constructible isomorphism of $(\mathbb{C}; +)$ and these are precisely the scalar multiplications λ_a for $a \in \mathbb{C}$. Under π_1, λ_a becomes $x/\sim_1 \mapsto (ax/\sim_1)$. The identification of $\mathbb{C}/\sim_1$ with $(\mathbb{C}; P)$, where P is the image of $+$, is $x \mapsto x^n$. The image under the map $x/\sim_1 \mapsto ax/\sim_1$ under this identification is λ_a^n. Now it's clear that the ternary polynomial which is zero if and only if $P(x, y, z)$ holds is homogeneous, which implies that the map λ_b fixes the predicate P. In the language of item 3 in the summary,

$$\tau(g_n(\mathbb{C}; +, \lambda_c : c \in F')) = g_n(\mathbb{C}; +, \lambda_c : c \in F)$$

if and only if $F' = F$ and $\tau = \lambda_b$ for some b.

In case 3 when $n = 1$, what we want to know is when

$$\tau((\mathbb{C} : +, \lambda_b : b \in F')) = (\mathbb{C}; +, \lambda_b : b \in F).$$

τ applied to the graph of $+$ is a group operation defined in the vector space, so it must be the graph of $(x, y) \mapsto x + y - c$ for some c. One verifies with no difficulty that this occurs exactly if $\tau(x) = \lambda_b(x) + c$ for some b; we must also have $F' = F$. Notice that in this case as in the last paragraph, τ itself need not be definable in the structure.

Now suppose that we are in item 4, so A_1 is the elliptic curve $y^2 = x^3 + ax + b$ and $\sim_1$ identifies (x, y) with $(x, -y)$. e must be the identity or an element of order 2 in A_2. In the first case, $\pi_2 f \pi_1^{-1}$ is the identity on $A_1/\sim_1$ and τ lifts to an additive automorphism of the elliptic curve. If we suppose that $ab \neq 0$, the only such automorphisms are $[1]$ and $[-1]$ and both of these induce the identity. If $b = 0$, we have the extra automorphisms $[i]$ and $[-i]$, both of which induce the function $x \mapsto -x$. (Recall that $[i](x, y) = (-x, iy)$. Also the identification of $A_1/\sim_1$ and $(\mathbb{C}; P_{a,b})$ is so immediate in this case that we don't mention it.) If $a = 0$, we have automorphisms $[\eta]$, $[-\eta]$, $[\eta^2]$, $[-\eta^2]$ where η is a primitive cube root of unity. As $[\eta](x, y) = (\eta x, y)$ the first two induce the function $x \mapsto \eta x$ and the last two $x \mapsto \eta^2 x$.

Now suppose that e has order 2; then $f^{-1}g(x, y) = (x, y) \oplus_1 (x_0, 0)$ for some x_0 that solves $x^3 + ax + b = 0$; this is because the points $(x_0, 0)$ are the elements of order 2 in the elliptic curve. Again, this function is its own inverse, so that $g^{-1}f(x, y) = (\sigma(x), y') = (x, y) \oplus_1 (x_0, 0)$. Using the addition formula on the curve, we see that

$$\sigma(x) = \frac{x_0 x + 2x_0^2 + a}{x - x_0} .$$

From this and the fact that $\pi_2 g \pi_1^{-1}$ is the identity, we see that $\pi_2 f \pi_1^{-1} = \sigma$. It should come as no surprise that $y = \sigma(x)$ holds if and only if $P_{a,b}(x, y, x_0)$. So if $ab \neq 0$, we have that $\tau(\mathbb{C}; P_{a,b}) = (\mathbb{C}; P_{a,b})$ exactly if τ is the identity or one of the 3 σ's mentioned. Clearly such a σ never arises from an isomorphism of elliptic curves (it has the wrong form). So if $(a_1, b_1) \neq (a_2, b_2)$, then for no τ do we have $\tau(\mathbb{C}; P_{a_1,b_1}) = (\mathbb{C}; P_{a_2,b_2})$.

Two more points on item 4. We leave to the reader the straightforward calculation of possible τ's when $ab = 0$; the necessary information is contained in the two preceding paragraphs. Also, if the curve has complex multiplication, it should be clear that the rings R' mentioned must be the same for $\mathcal{M}$ and $\mathcal{N}$, when $\tau(\mathcal{M}) = \mathcal{N}$.

For item 5, we must have that $[i]e = e$. If e is the identity, then τ lifts to an additive automorphism of the curve $y^2 = x^3 + x$; the only 4 of these are $[1], [-1], [i]$, and $[-i]$, and they all project to the identity. Otherwise, recalling that $[i](x, y) = (-x, iy)$, we must have that e is A_2's copy of $(0, 0)$ and thus that $f^{-1}g(x, y) = (x, y) \oplus_1 (0, 0)$. The projection of this onto the first coordinate is x^{-1}. The map $x \mapsto x^2$ (the identification in this case) takes the map $x \mapsto x^{-1}$ to itself. In this case, then, τ must be either the identity or $x \mapsto x^{-1}$.

In item 6, we have $[\eta]e = e$. If e is the identity, τ lifts to an automorphism of $y^2 = x^3 + 1$. Of these, $[1], [\eta]$, and $[\eta^2]$ project to the identity on our structure $\pi_y(\mathcal{E})$. $[-1], [-\eta]$, and $[-\eta^2]$ all project to $y \mapsto -y$. Otherwise e is A_2's copy of either $(0, 1)$ or $(0, -1)$. In the first case, $f^{-1}g(x, y) = (x, y) \oplus_1 (0, 1)$, which projects to $y \mapsto \frac{-y+3}{y+1}$; this is its own inverse again. In the second case, $f^{-1}g(x, y) = (x, y) \oplus_1 (0, -1)$, which projects to $y \mapsto \frac{y+3}{y-1}$. Thus, in this case, if $\mathcal{M} = \tau(\mathcal{N})$, we have $\mathcal{N} = \mathcal{M}$ and τ is one of six possibilities.

Finally, in item 7, we must have that $[-\eta]e = e$; regarding A_2 as a copy of the curve $y^2 = x^3 + 1$, we would then have (for e not the identity) that

$e = (x, y) = (\eta x, -y)$ so $e = (0, 0)$. But $(0, 0)$ is not on the curve, so e must then be the identity. τ must lift to an automorphism of the curve, and all 6 of these project to the identity on $\mathcal{M}$. The only possibility in this case is that τ is the identity.

Now we turn to the possibility that, for arbitrary $\mathcal{M}$ and $\mathcal{N}$ on our list, and arbitrary definable functions f, g from $\mathbb{C}$ onto $\mathbb{C}$, $f^{-1}(\mathcal{M}) = g^{-1}(\mathcal{N})$. That is, the canonical cover of $\mathcal{M}$ via f is the same reduct as the canonical cover of $\mathcal{N}$ via g. Recall that while f need not be definable in $f^{-1}(\mathcal{M})$ the relation $x \sim_f y \Leftrightarrow f(x) = f(y)$ is; it is the pre-image of the identity on $\mathcal{M}$. Also note that if S is any generic fiber of f and α is any bijection of $\mathbb{C}$ fixing the complement of S pointwise, then α is an automorphism of $f^{-1}(\mathcal{M})$. This immediately implies that the generic fibers of f and g cannot overlap nontrivially: if x is generic, either $x/\sim_f = x/\sim_g$ or one of them is just $\{x\}$. In the second case both equivalence relations are trivial.

Thus, if $f^{-1}(\mathcal{M}) = g^{-1}(\mathcal{N})$, f and g have the same fibers (except for finitely many). This implies that gf^{-1} is a definable bijection (modulo a finite set of exceptions) and thus a fractional linear transformation τ. Making a finite number of adjustments, we have that $\tau(\mathcal{M}) = \mathcal{N}$. We summarize all our observations in this section.

Proposition 9.1 *Suppose that $\mathcal{M}$ and $\mathcal{N}$ are items from the list in 7.1. Suppose that f and g are constructible functions from $\mathbb{C}$ onto $\mathbb{C}$ so that the reducts $f^{-1}(\mathcal{M})$ and $g^{-1}(\mathcal{N})$ are the same. Then $\mathcal{M} = \mathcal{N}$ and $f(x) = \tau g(x)$ for some fractional linear transformation τ, except for finitely many x. In each case, we list the possibilities for τ:*

1. *For some fixed c, either $\tau(x) = cx$ for all x or $\tau(x) = cx^{-1}$ for all x.*
2. *Either $\tau(x) = x$ for all x or $\tau(x) = -x$ for all x.*
3. *If $n \neq 1$ then for some fixed c, $\tau(x) = cx$ for all x. If $n = 1$, there are fixed c and d so that $\tau(x) = cx + d$ for all x.*
4. *Suppose the elliptic curve in question is given by $y^2 = x^3 + ax + b$; let x_1, x_2, and x_3 be (distinct) solutions of $x^3 + ax + b = 0$. For $j = 1, 2, 3$, let $\sigma_j(x) = \frac{x_j x + 2x_j^2 + a}{x - x_j}$. In any case, τ may be the identity or one of the σ_j's.*
   (a) *If $ab \neq 0$, these four functions are the only possible τ's.*
   (b) *If $b = 0$, we have 4 more possibilities for τ: $x \mapsto -x$ and $x \mapsto \sigma_j(-x)$ for $j = 1, 2, 3$.*
   (c) *If $a = 0$, τ may be $x \mapsto \eta x$, $x \mapsto \eta^2 x$, $x \mapsto \sigma_j(\eta x)$, or $x \mapsto \sigma_j(\eta^2 x)$ for $j = 1, 2, 3$, where η is a primitive cube root of 1. (There are 12 τ's in all in this case.)*
5. *τ is either the identity or $x \mapsto x^{-1}$.*
6. *$\tau(x) = \sigma(x)$ for all x or $\tau(x) = -\sigma(x)$ for all x, where σ is either the identity, $x \mapsto \frac{-x+3}{x+1}$, or $x \mapsto \frac{x+3}{x-1}$.*
7. *τ is the identity.*

References

[1] Hrushovski, E., "A new strongly minimal set," *Annals of Pure and Applied Logic*, vol. 62 (1993), pp. 147–66. Stability in model theory, III (Trento, 1991). Zbl 0804.03020. MR 94d:03064. 168

[2] Hrushovski, U., and A. Pillay, "Weakly normal groups," pp. 233–44 in *Logic Colloquium '85 (Orsay, 1985)*, vol. 122 of *Studies in Logic and the Foundations of Mathematics*, North-Holland, Amsterdam, 1987. Zbl 0636.03028. MR 88e:03051. 169, 186

[3] Laskowski, M. C., "Uncountable theories that are categorical in a higher power," *The Journal of Symbolic Logic*, vol. 53 (1988), pp. 512–30. Zbl 0653.03019. MR 90d:03059. 168

[4] Loveys, J., "Weakly minimal groups of unbounded exponent," *The Journal of Symbolic Logic*, vol. 55 (1990), pp. 928–37. Zbl 0718.03027. MR 92e:03045. 169

[5] Loveys, J., "On locally modular, weakly minimal theories," *Archive for Mathematical Logic*, vol. 32 (1993), pp. 173–94. Zbl 0797.03035. MR 94a:03056. 169

[6] Poizat, B., *Groupes Stables*, Nur al-Mantiq wal-Ma'rifah [Light of Logic and Knowledge], 2. Bruno Poizat, Lyon, 1987. Une tentative de conciliation entre la géométrie algébrique et la logique mathématique. [An attempt at reconciling algebraic geometry and mathematical logic]. Zbl 0633.03019. MR 89b:03056. 168, 170

[7] Rabinovitch, E., *Definability in a Field with Sufficiently Rich Incidence Systems*, Queen Mary and Westfield College, School of Mathematical Sciences, London, 1993. 168

[8] Silverman, J. H., *The Arithmetic of Elliptic Curves*, vol. 106 of *Graduate Texts in Mathematics*, Springer-Verlag, New York, 1986. Zbl 0585.14026. MR 87g:11070. 171, 172

[9] Silverman, J. H., *Advanced Topics in the Arithmetic of Elliptic Curves*, vol. 151 of *Graduate Texts in Mathematics*, Springer-Verlag, New York, 1994. Zbl 0911.14015. MR 96b:11074. 169

Department of Mathematics and Statistics
McGill University
805 Sherbrooke St West
Montreal QC H3A 2K6
CANADA
loveys@math.mcgill.ca