UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Higher Dimensional Trichotomy Conjectures in Model Theory

Permalink
https://escholarship.org/uc/item/0wb761q8

Author
Castle, Benjamin

Publication Date
2021

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Higher Dimensional Trichotomy Conjectures in Model Theory

by

Benjamin Castle

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thomas Scanlon, Chair
Associate Professor Pierre Simon
Associate Professor Adityanand Guntuboyina

Spring 2021

Higher Dimensional Trichotomy Conjectures in Model Theory

Copyright 2021
by
Benjamin Castle

Abstract

Higher Dimensional Trichotomy Conjectures in Model Theory

by

Benjamin Castle

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Thomas Scanlon, Chair

In this thesis we study the Restricted Trichotomy Conjectures for algebraically closed and o-minimal fields. These conjectures predict a classification of all sufficiently complex, that is, non-locally modular, strongly minimal structures which can be interpreted from such fields. Such problems have been historically divided into 'lower dimensional' and 'higher dimensional' cases; this thesis is devoted to a number of partial results in the higher dimensional cases. In particular, in ACF_0 and over o-minimal fields, we prove that all higher dimensional strongly minimal structures whose definable sets satisfy certain geometric restrictions are locally modular. We also make progress toward verifying these geometric restrictions in any counterexample. Finally, in the last chapter we give a full proof of local modularity for strongly minimal expansions of higher dimensional groups in ACF_0.

To Katrina and Justin

Contents

Acknowledgments

I would first like to thank my advisor, Tom Scanlon, for mentoring me throughout my time at Berkeley. I feel very lucky to have been one of his students, and this thesis would not have been possible without his endless patience, support, and mathematical wisdom.

I also thank the other members of the committee, Pierre Simon and Adityanand Guntuboyina, for their time and effort. I especially thank Pierre for his contributions to the model theory group at Berkeley, including running our seminar. Similarly, I am thankful to Carol Wood and Ward Henson for their presence in the Berkeley logic community.

This thesis was inspired by the work of too many mathematicians to name. However, I would like to mention Assaf Hasson and Dmitry Sustretov, whose ideas formed the starting point for this project; and especially Assaf, for a number of helpful conversations on the material.

Next, I thank all of my peers within the logic community at Berkeley, for making graduate school a fun and memorable experience. These include, but are by no means limited to, fellow model theorists Nick Ramsey, Reid Dale, Mariana Vicaria, Adele Padgett, David Casey, Diego Bejarano Rayo, Scott Mutchnik, and Sebastian Eterović.

I also thank my former housemates, Ian Gleason, Ravi Fernando, and Patrick Lutz, for a myriad of stimulating conversations, both within and outside of mathematics; I enjoyed developing friendships with each of them during my time at Berkeley. I especially thank Ravi for patiently fielding a number of probably silly questions from a model theorist trying to do algebraic geometry; and Ian for being my best friend for the last six years.

Thank you as well to all those who mentored me mathematically before my Berkeley years: in particular, Adrienne Stanley, Douglas Mupasiri, and TJ Hitchman, my undergraduate professors; Péter Juhász, my teacher, mentor, and friend from my time in Budapest; and Brian Reynolds, my high school math coach.

Last but not least, I thank my entire family – in particular my parents, and my brother Tommy – for their unwavering support. My parents recognized my love of math early on, and I wouldn't have made it to this point without their constant encouragement.

I especially thank my wife, Katrina, for supporting me in so many ways, and helping me to stay sane during these last few months of writing. And finally, I thank my five month old son, Justin. Much of this thesis was typed with him sitting or sleeping in my lap; and even though he did everything he could to keep me fully distracted the whole time, having him around truly makes pushing through the hardest moments worth it.

Chapter 1

Introduction

1.1 Historical Background

The class of strongly minimal sets has long been of importance to Model Theory, originating in 1971 with its use in Baldwin and Lachlan's treatment of Morley's Categoricity Theorem [2]. The study of such sets, however, found a new and exciting avenue in the work of Boris Zilber in the 1970s and 1980s, which proposed a classification according to familiar algebraic objects. Zilber (see, for example, [53] and [54]) established a trichotomy which divided strongly minimal sets according to the complexity of their induced strongly minimal structures. He proceeded to classify those of 'middle' complexity – the non-trivial locally modular sets – as those whose 'geometries' arise from linear algebra over division rings (a more precise form of this identification was given later by Hrushovski [23] – see Fact 2.4.2 of this thesis). Following this line of thought, Zilber conjectured that the most complex level – the non-locally modular strongly minimal sets – should be classifiable as those arising from algebraic geometry over algebraically closed fields. Precisely, he showed that non-locally modular strongly minimal sets are exactly those which contain certain combinatorial objects called 'pseudoplanes' (see [53], Theorem 3.1 or [54], Chapter II, Theorem 3.5); he then conjectured the following (see [53], Conjecture B):

Conjecture 1.1.1 (Zilber's Trichotomy Conjecture)**.** *Every uncountably categorical pseudoplane is mutually interpretable with an algebraically closed field. In particular, every non-locally modular strongly minimal set interprets an algebraically closed field.*

This conjecture has an enticing flavor: indeed, it would suggest that algebraically closed fields are characterizable (up to mutual interpretability) according to purely model theoretic principles. Alas, Zilber's conjecture turned out to be false, as proven in 1993 by Ehud Hrushovski [21]:

Theorem 1.1.2 (Hrushovski Construction)**.** *There is a strongly minimal set which is not locally modular and does not interpret any infinite group, thus in particular cannot interpret an algebraically closed field.*

Nevertheless, Zilber's original prediction has retained significant importance to model theory – in particular, by guiding much of the development of geometric stability theory over the last 35 years. Since Hrushovski's refutation of Zilber's original conjecture, model theorists have proceeded to discover a multitude of instances in which similar trichotomy phenomena do hold. These instances include, for example:

- strongly minimal sets definable in differentially closed fields (see [26] and [41]).
- strongly minimal sets which carry a 'Zariski-like' topology [27].
- o-minimal structures viewed locally near a point [34].

In each of the above examples, a classification is achieved along Zilber's original lines, dividing the specified objects into the 'trivial,' 'vector space,' and 'field' cases. Thus it has long been apparent that various geometric situations in Model Theory witness the same surprising pattern: the emergence 'out of thin air' of the same familiar algebraic objects (vector spaces and fields).

The common guiding principle which fuels these classification results is the presence of an underlying geometry on definable sets. For example, one may choose to work only with definable sets that are 'closed' in a meaningful sense. A more difficult challenge, then, is to establish trichotomy phenomena without such an available notion of geometry. Of course, as noted above, the trichotomy fails in full generality even for strongly minimal sets. However, an interesting 'in between' question is to ask about structures for which (i) there is a geometry controlling the definable sets, but (ii) that geometry can only be seen from a certain 'outside' object, not necessarily definable from the structure itself. In this case, for example, one is able to take the closure of a definable set, but must do so 'undefinably,' resulting in a set which only the outside object can see.

The type of problem described above could be termed a 'reduct' trichotomy problem: one starts with a full structure of geometric nature, then restricts to an interpreted or 'reduct' structure carrying strong model-theoretic assumptions; the goal is to use the external geometry to show that this interpreted structure satisfies a trichotomy principle. The success of such a reduct problem carries interesting philosophical implications for the governing geometry: namely, a positive answer implies that we are unable to 'forget' the geometry while preserving certain Model Theoretic properties of its structure; indeed, in finding a vector space or field one 'recovers' the underlying geometry from scratch. The conclusion, then, is that sufficiently complex and model-theoretically tame structures, if able to be generated from within a certain geometry, must actually 'be' the original geometry.

One of the oldest such 'restricted' instances of the trichotomy conjecture concerns strongly minimal reduct structures in algebraically closed fields – that is, reducts of the full induced structure on a particular constructible set. Formally, Zilber conjectured the following (for example, see [51], Remark 1.4):

Conjecture 1.1.3 (Zilber's Restricted Trichotomy Conjecture)**.** *Let $\mathcal{M}$ be a strongly minimal structure which is interpreted in an algebraically closed field K. If $\mathcal{M}$ is not locally modular then $\mathcal{M}$ interprets a field isomorphic to K. In particular, $\mathcal{M}$ satisfies Zilber's Trichotomy Conjecture.*

This Restricted Trichotomy Conjecture remains open today, though there have been major pieces of progress. Of particular note was Eugenia Rabinovich's proof for strongly minimal structures whose universe is itself an algebraically closed field [44]. Rabinovich's work laid out the foundation for an argument in general – and this argument was later shortened and adapted by Assaf Hasson and Dmitry Sustretov, whose current preprint [20] proves the conjecture for all strongly minimal structures with universe of dimension 1. Sustretov further claimed a proof for structures with higher dimensional universes, but his argument was subsequently retracted due to a cricial error. Thus the full proof of the Restricted Trichotomy Conjecture rests on the case of structures whose universes have dimesion greater than 1. It is this problem which occupies the majority of the present thesis.

There is also a related conjecture for strongly minimal structures interpretable from o-minimal fields. This time such a reduct is never able to capture the full field structure, for example by the non-stability of the field. Instead, we expect the 'most complex' strongly minimal reduct structures to precisely capture the (stable) algebraic closure of the field. We get the following conjecture of Peterzil:

Conjecture 1.1.4 (O-minimal Restricted Trichotomy Conjecture)**.** *Let $\mathcal{M}$ be a strongly minimal structure which is interpreted in an o-minimal field $\mathcal{R} = (R, +, \cdot, <, ...)$. If $\mathcal{M}$ is not locally modular, then $\mathcal{M}$ interprets a field isomorphic to $R[i]$.*

As in Zilber's original Restricted Trichotomy Conjecture, progress on the above problem has been restricted to so-called 'lower dimensional' cases – in this case, structures with universes of dimension 1 or 2 (the o-minimal dimension of $R[i]$). Notably, Hasson, Onshuus, and Peterzil gave a full proof for structures with universe of dimension 1 [19], and more recently a paper of Eleftheriou, Hasson, and Peterzil [12] gives a proof for expansions of 2-dimensional groups.

Meanwhile, no progress has yet been made for universes of dimension greater than 2. In Chapter 6 we give a partial result in this direction, for strongly minimal structures whose definable sets satisfy certain geometric restrictions arising from o-minimal geometry.

1.2 The Challenge of Working in Higher Dimensions

The strategy of past authors in 'low' dimensional cases reflects a crucial difference from the higher dimensional cases: namely, that non-locally modular strongly minimal sets which interpret the relevant field actually exist. In the works of Rabinovich, Hasson and Sustretov, for example, one starts with a given non-locally modular strongly minimal structure, say $\mathcal{M}$; on the other hand, one has in mind another strongly minimal structure which should be mutually interpretable with $\mathcal{M}$ – namely, the field K over which $\mathcal{M}$ is defined. The argument, then, involves showing that there is enough resemblance between $\mathcal{M}$ and K that one eventually 'recognizes' K inside of $\mathcal{M}$. This is done by developing a theory of slopes of $\mathcal{M}$-definable 'curves' – which from the onset must resemble the K-points of actual curves up to finite error. One then uses the intersection theory of curves over K to identify the tangency relation on pairs of curves in $\mathcal{M}$, again up to finite error. This identification roughly enables one to interpret the 'family of slopes' of a given family of curves – a set which a priori agrees up to finite error with a one-dimensional algebraic group over K. An actual one-dimensional group, and subsequently a field, are then interpreted using Hrushovski's group and field configuration results [22].

In contrast, in the higher dimensional case of the Restricted Trichotomy Conjecture, it is known that no strongly minimal reduct structure can interpret an infinite field (see Corollary 3.1.16 of this thesis, though the result is not new). Thus the goal of the argument changes, as we are instead asked to directly establish local modularity in an arbitrary such reduct structure. Of course one could try to interpret a field anyway, proving the conjecture by way of contradiction; however, various issues arise with such an argument. Most notably, $\mathcal{M}$-definable curves need not agree with the K-points of any variety up to finite error: indeed, while two distinct $\mathcal{M}$-definable curves have finite intersection, their 'closest resembling varieties' may have infinite intersection. Thus one would require the general intersection theory of higher dimensional varieties to try to interpret a 'group of slopes' – and in this context there are not strong enough theorems to be able to satisfy the strict hypotheses of Hrushovski's group and field configurations.

Compounding the above difficulties is the nature of the tangency relation on $\mathcal{M}$-definable curves. In the one-dimensional case, we can fix a one-dimensional family of such curves whose tangency relation is a finite-to-finite correspondence, and whose set of slopes is a one-dimensional group up to finitely many points. This allows for the construction of a tuple of parameters forming a group configuration, in which certain elements are algebraic over specific subtuples. In higher dimensions, the tangency relation need only be finite-to-finite if (i) the family has rank 1, and (ii) we have succeeded in 'almost' defining tangency in the reduct (which as mentioned above does not go smoothly). Moreover, the set of slopes of a rank one family need not resemble a group; this could be fixed by applying a series of composition operations, but the result would be a larger family with a more complicated tangency relation. Such issues, in combination, significantly complicate the task of building a valid group configuration.

To resolve these problems, as stated above, we will entirely avoid the interpretation of a field. Instead we will establish local modularity directly under various assumptions. In essence, this entails directly bounding the size of families of $\mathcal{M}$-definable curves. One of the most successful past instances in which such bounds have been attained is the case of *unimodular* theories [25] – and indeed in some instances we are able to proceed directly to local modularity through unimodularity. This method cannot work universally, however, as there are many examples of locally modular, non-unimodular strongly minimal sets which are interpretable in algebraically closed fields. Thus in the most general setting, a new method is needed for bounding families of curves; such an argument is carried out in Chapter 5.

Finally, we mention here the difficulty caused by lack of agreement between a reduct structure and the governing geometry of the field. In dimension one, past authors needed to deal with the fact that $\mathcal{M}$-definable curves need not be actual curves in the geometric sense; though after adding and removing finitely many points this identification is indeed possible. In higher dimensions, an $\mathcal{M}$-definable curve can be identified with a pure dimensional Zariski closed set up to 'error of smaller dimension.' So, while in one dimension we may need to deal with a finite number of isolated points that complicate the geometry of such a set, in higher dimensions we may have infinitely many 'superfluous' points – and these points can easily break the various counting arguments we would like to use to prove the conjecture.

Our main strategy to deal with this most recent issue is to study almost everything 'up to codimension 2.' That is, we will only worry about extra or missing regions of an $\mathcal{M}$-definable set which are of pure codimension 1 – disregarding all other such points. This is largely successful, because we employ a number of counting and geometric arguments which are unphased by a codimension 2 error. Then, when trying to deal with these 'bad' points, we have the strong assumption that they form, say, a variety of codimension 1. We attempt, and in some cases succeed, to show that such a variety cannot be found in too many of the sets the reduct structure is able to define.

1.3 Outline of Results

Chapters 2 and 3 contain most of the background information needed to implement our strategies. Chapter 2 outlines the standard facts about strongly minimal sets and Zilber's trichotomy, including the characterization of local modularity in terms of 'plane curves' that we use throughout. Chapter 3, meanwhile, contains several geometric and topological preliminaries – notably introducing the notions of 'almost purity' and 'almost closedness' that we will exploit throughout our work in ACF_0.

In Chapter 4 we present our first main result: namely, we establish the local modularity of higher dimensional strongly minimal reduct structures in ACF_0 in which (i) the universe is a smooth quasiprojective variety of finite fundamental group, and (ii) any 'generic' plane curve is almost pure. We do this in two independent ways. First, we analyze degrees of covering maps to directly reduce to unimodularity, which as previously mentioned implies local modularity. Second, given a sufficiently large family of plane curves, we show how to modify the curves, definably over the full field structure, until they become pairwise disjoint; this provides a contradiction for large enough families since all such curves lie inside the 'plane' – which is a space of bounded size. In a sense this second argument gives a more geometrically satisfying interpretation of the implication between unimodularity and local modularity, provided there is a sufficiently strong background geometry which can be used to perform the necessary modifications.

Chapter 5 generalizes the results of Chapter 4 to varieties with arbitrary fundamental group. Namely, we show that if the universe of a higher dimensional strongly minimal reduct structure $\mathcal{M}$ is a smooth quasiprojective variety in characteristic zero, then there are no rank ≥ 2 families of plane curves whose 'generic' members are almost pure. The proof involves approximating $\mathcal{M}$-definable sets with covering spaces to show that, in a certain sense, any two generic plane curves in a rank ≥ 2 family have finite intersection. The interesting thing about this statement is that, as worded, it makes no reference to the relationship between the two curves: they can even be the same curve. Thus by intersecting an infinite curve with itself, we obtain a contradiction – and thereby conclude that such families of curves do not exist in the first place.

Chapter 6 is an orthogonal generalization of Chapter 4: instead of relaxing the fundamental group, we relax the background geometry. Namely, we consider strongly minimal reduct structures on affine space over o-minimal fields. We conclude a similar result to Chapter 4: after adopting a generalized notion of almost purity, any such structure in higher dimensions, with every 'generic' plane curve almost pure, is unimodular. The overarching strategy is identical to the work in Chapter 4; however in the o-minimal setting we lack the availability of canonical irreducible components of closed sets. The main point of the chapter, then, is that under a suitable generalization of almost purity we can actually recover a canonical component decomposition from scratch. To this end, we define a notion of 'components' for definable sets, and show that for almost pure sets, various properties transfer from the algebraic case. We then use these components to carry out the unimodularity argument from Chapter 4.

Chapter 7 addresses the extent to which we can establish the almost purity of enough plane curves to carry out the work in the previous chapters. The main result of this chapter is that, if the universe is a smooth variety in characteristic zero, we can assume that all 'sufficiently generic' plane curves are 'generically almost closed' (i.e. any large frontier regions are constrained to non-generic parts of the plane). We then outline a general strategy to try to establish the almost purity of enough plane curves in compact universes by exploiting this generic almost closedness. This is not successful, but to an extent we can quantify where the argument can fail, and we hope that a similar approach might be completed to a full proof in the future.

Finally, Chapter 8 applies the results of Chapter 7 to cover the general case of strongly minimal groups in characteristic zero. The main result is that all higher dimensional strongly minimal groups interpreted in ACF_0 are locally modular; thus to prove the full conjecture in characteristic zero one only needs to interpret a group, and not a field. The proof is to an extent modeled after a similar argument given in [12]. The case of compact groups (that is, abelian varieties) follows quite easily from Chapter 7. For non-compact groups, as in [12], the main step is to show that plane curves can have only finitely many 'poles,' and thus in a sense we can treat the universe as if it actually is compact. We then deduce almost purity using a straightforward adaptation of a result from Chapter 7.

Chapter 2

Preliminaries on Strongly Minimal Sets

In this chapter we review the basic facts about strongly minimal sets, and develop the tools used in studying the Zilber trichotomy. The facts stated are not original, and indeed most of the material can be found in essentially equivalent forms in e.g. [30] or [38]. However, we have chosen to give an exposition of the basic theory for the sake of completeness and continuity.

We assume the reader is familiar with basic model theory – see, for example, chapters 1-4 of [30], or chapters 1-5 of [4]. Note that we will almost always view structures as given relationally, i.e. as a universe along with select subsets of its cartesian powers. This does not affect the study of trichotomy problems, as we are only interested in studying definable sets of certain structures – and not language-specific topics such as substructures or quantifier complexity.

2.1 Strong Minimality and Dimension

Definition and Examples

Definition 2.1.1. Let X be a definable set in a structure $\mathcal{M}$. Then X is *strongly minimal* if:

1. X is infinite.

2. Every definable subset of X is finite or cofinite.

3. Item (2) also holds in every elementary extension $\mathcal{N}$ of $\mathcal{M}$ – that is, the set defined in $\mathcal{N}$ by the same formula used to define X in $\mathcal{M}$, satisfies (2) in $\mathcal{N}$.

If only conditions (1) and (2) are satisfied, then X is called *minimal*. The point of condition (3) is to prevent growing families of uniformly defined finite sets, which would result in an infinite coinfinite set in an elementary extension. In fact, the definition of strong minimality can equivalently be expressed without (3) if we assume uniform bounds on finite sets:

Fact 2.1.2. *The following are equivalent for an infinite definable set X in a structure $\mathcal{M}$:*

1. *X is strongly minimal.*
2. *For any uniformly definable family $\mathcal{D}$ of subsets of X, there is a positive integer n such that for all $D \in \mathcal{D}$, either D or $X - D$ has cardinality at most n.*

We will work mainly with the subtly different notion of a strongly minimal structure:

Definition 2.1.3. A *strongly minimal structure* is a structure $\mathcal{M}$ whose universe is a strongly minimal set. Equivalently, $\mathcal{M}$ is a strongly minimal structure if it is an infinite structure in which every one-variable definable set, even in elementary extensions, is finite or cofinite.

There is no real harm in simplifying our discussion of strongly minimal sets to the special case of strongly minimal structures, as any strongly minimal set can be viewed as a strongly minimal structure:

Fact 2.1.4. *If X is a strongly minimal set in a structure $\mathcal{M}$, then the structure $\mathcal{X}$, with universe X and equipped with exactly those subsets of powers of X which are definable in $\mathcal{M}$, is a strongly minimal structure.*

The strongly minimal structure $\mathcal{X}$ is the *induced* structure on the strongly minimal set X. It carries all of the same properties (rank and degree functions, pregeometry structure, etc.) as the original strongly minimal set X. For this reason, we will from now on assume to work only with strongly minimal structures.

Example 2.1.5. Let K be an algebraically closed field, viewed as a structure in the language of rings. Then by quantifier elimination every definable subset of K is defined (even uniformly in families) by a Boolean combination of polynomial equalities; the strong minimality of K then follows from the fact that a nonzero polynomial of degree d can have at most d roots.

Example 2.1.6. Other key examples of strongly minimal structures are pure sets, infinite vector spaces (with unary function symbols representing field elements), and all models of $\mathrm{Th}(\mathbb{Z}, s)$, the theory of the integers with the successor function. In each case the argument is nearly identical to the one above for algebraically closed fields: one uses quantifier elimination to reduce to the case of equations, then uses the fact that nontrivial equations have a (uniformly bounded in families) finite number of solutions.

We will keep the above examples in mind while discussing various properties of strongly minimal structures throughout this chapter. The case of algebraically closed fields is particularly important, as it serves as the background setting for much of the rest of the thesis.

Dimension and Degree

We now turn to the basic properties of strongly minimal structures. To start, one of the main consequences of Fact 2.1.2 is the availability of a definable dimension theory for definable sets and types – which, to the familiar reader, is just a specific case of Morley rank.

Fact 2.1.7. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure. Then there are functions RM and DM, each with domain the collection of non-empty definable subsets of cartesian powers of M, with the following properties:*

1. *For each such D, $RM(D)$ is a natural number and $DM(D)$ is a positive integer.*
2. *RM and DM are automorphism invariant.*
3. *A finite set D of size $n > 0$ has $RM(D) = 0$ and $DM(D) = n$; M itself satisfies $RM(M) = 1$ and $DM(M) = 1$.*
4. *$RM(D \cup E) = \max(RM(D), RM(E))$. If $RM(D) = RM(E)$ and $D \cap E = \emptyset$ then $DM(D \cup E) = DM(D) + DM(E)$.*
5. *If $\{D_a\}$ is a uniformly definable family of pairwise disjoint sets, each with a constant value n under RM, and indexed by a definable set A, then $RM\left(\bigcup D_a\right) = n + RM(A)$. In particular, $RM(D \times E) = RM(D) + RM(E)$ for all D and E.*
6. *RM is definable in families: that is, if $\{D_a\}$ is a uniformly definable family of sets, indexed by a definable set A, then $\{a \in A : RM(D_a) = n\}$ is definable for each n.*

We will refer to the funciton RM in various settings as either Morley rank, rank, or dimension; the funciton DM will always be called degree or Morley degree.

Remark 2.1.8. Some basic properties of these functions can quickly be inferred from Fact 2.1.7. For example, a non-empty finite set is exactly a set of rank 0; a definable set D is strongly minimal if and only if it has rank and degree both equal to 1; and RM is preserved under definable bijections, and more generally definable finite-to-finite correspondences.

Remark 2.1.9. Note, also, that the functions RM and DM naturally extend to *interpretable* sets, by moving to the structure $\mathcal{M}^{\text{eq}}$. Moreover, the basic properties given above all transfer to this more general context.

Example 2.1.10. Let D be a definable set over an algebraically closed field K – that is, a subset of K^n for some n which is a Boolean combination of affine algebraic sets. Then the Zariski closure $\overline{D}$ of D is a finite union of irreducible components, say $\overline{D} = C_1 \cup ... \cup C_m$, with each C_i of dimension d_i as an algebraic variety. Then we can interpret the RM and DM functions in terms of these dimensions d_i: $RM(D) = RM(\overline{D})$ is the maximum of the d_i, and $DM(D) = DM(\overline{D})$ is the number of the d_i which attain this maximum.

In other words, the RM function agrees with the notion of dimension on affine algebraic sets; and the DM function counts the number of top-dimensional irreducible components.

Convention 2.1.11. Throughout, when working with an algebraically closed field K, we interpret the term *variety*, or *variety over* K, as referring to the set of K-points of an irreducible quasiprojective variety over K. We will make this identification throughout the thesis, but will try to remind the reader as we go.

Remark 2.1.12. Given the above discussion, the reader may wonder about varieties which are not affine. Indeed, it follows from elimination of imaginaries [43] that any variety V over an algebraically closed field K may be identified with a definable subset of K^n for some n. This identification need not preserve the Zariski topology, but it does preserve the class of constructible sets. In light of this fact, when studying definable subsets of cartesian powers of V, one may choose to work with the Zariski and analytic topologies on powers of V rather than K, and to interpret 'definable' sets as Boolean combinations of Zariski closed sets. It is important to note that this is harmless: indeed, all of the facts stated in this section about dimension and genericity over algebraically closed fields extend to definable sets in this more general context. We will make heavy use of this generalization throughout the thesis.

Dimension of Tuples and Types

The function RM also extends to complete types, as follows:

Definition 2.1.13. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, a a tuple of elements of M, and A a set of parameters. Then by $RM(a/A)$ we mean the smallest natural number n such that there is an A-definable set D with $a \in D$. If p is a type over A, then by $RM(p)$ we mean $RM(a/A)$ for some (any) realization of p.

As with sets, Definition 2.1.13 extends naturally to tuples in $\mathcal{M}^{\text{eq}}$, and carries the same basic properties.

Example 2.1.14. Over an algebraically closed field, $RM(a/A)$ is the transcendence degree, over the field generated by A, of the extension generated by the elements of the tuple a. In a vector space V, $RM(a/A)$ is the dimension of the span of a in V after quotienting by the subspace generated by A. In $\text{Th}(\mathbb{Z}, s)$ any model is a disjoint union of copies of $\mathbb{Z}$; in this light, the function $RM(a/A)$ is just the number of such copies which both (i) contain an element of a and (ii) do not contain an element of A.

As with sets, we will refer to the function RM on tuples and types as either Morley rank, rank, or dimension. Essentially everything we need to know about this function is contained in the 'additivity' property given below, which roughly follows from the last two clauses of Fact 2.1.7:

Fact 2.1.15. *Let a and b be tuples, and A a set of parameters. Then*

$$RM(a, b/A) = RM(a/A) + RM(b/A, a).$$

Remark 2.1.16. For example, it follows that if $a_1, .., a_n$ are tuples, and A is a set of parameters, then we always have

$$RM(a_1, ..., a_n/A) \leq RM(a_1/A) + ... + RM(a_n/A).$$

If equality holds, then $a_1, ..., a_n$ are said to be *independent over* A.

We will frequently exploit Fact 2.1.15 in performing rank computations throughout this thesis. As a caution to the reader, we note that Fact 2.1.15 does not hold of Morley rank in general theories; rather, it follows in our setting from the hypothesis of strong minimality.

Genericity

We conclude this section with a discussion of genericity in strongly minimal structures.

Definition 2.1.17. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, and D, E definable sets with $D \subset E$. Then we say D is *generic in* E if $RM(D) = RM(E)$.

Remark 2.1.18. Note that Definition 6.4.1 is not completely standard, but is useful when working in strongly minimal structures. We have borrowed the terminology form the theory of stable groups.

Example 2.1.19. The generic subsets of M are precisely the cofinite sets. In an algebraically closed field, D is a generic subset of E if and only if $\overline{D}$ contains at least one top dimensional component of $\overline{E}$.

We can use Definition 6.4.1 to make the following intuitive notions precise:

Definition 2.1.20. If D and E are definable sets, then D is *almost contained in* E if $D - E$ is not generic in D. If D and E are almost contained in each other, then we say D and E are *almost equal*, denoted $D \sim E$.

It follows, for example, that almost equal sets have the same Morley rank and degree. A helpful example is that every definable set over an algebraically closed field is almost equal to its Zariski closure.

A common feature of Morley rank is the (essentially) unique decomposition of any definable set into degree one pieces:

Fact 2.1.21. *Let D be a non-empty definable set with $DM(D) = n$. Then there are pairwise disjoint generic definable subsets $D_1, ..., D_n$ of D of Morley degree 1 such that $D = D_1 \cup ... \cup D_n$. Moreover the D_i are unique up to almost equality: if $E_1 \cup ... \cup E_m = D$ is any other such decomposition then $m = n$, and the almost equality relation gives a bijection between the D_i and the E_j.*

For example, if D is definable in an algebraically closed field, then each D_i almost contains exactly one of the top-dimensional components of $\overline{D}$.

Another helpful consequence of Definition 6.4.1 is the following, which follows from clause (4) of Fact 2.1.7:

Fact 2.1.22. *If $D_1, ..., D_n, E$ are definable sets and $E = D_1 \cup ... \cup D_n$, then D_i is generic in E for at least one i. If E has Morley degree 1 and the D_i are pairwise disjoint, then D_i is generic in E for exactly one i.*

Finally, we note that genericity extends naturally to points and types:

Definition 2.1.23. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, a a tuple of elements of M, A a set of parameters, and D an A-definable set with $a \in D$. Then a is *generic in D over A* if $RM(a/A) = RM(D)$. If p is a complete type over A which contains a formula defining D, then p is a *generic type of D over A* if some (any) realization a of p is generic in D over A.

Thus generic tuples and types are those with the full Morley rank of the set we consider them in. A crucial fact about these notions is:

Fact 2.1.24. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, A a set of parameters, and D a non-empty A-definable set of Morley degree n. Then there are precisely n generic types of D over A. If the language of $\mathcal{M}$ is countable and $|A| + \aleph_0 < |M|$, then each of these types has a realization in M. In particular, in this situation every definable set has a generic point.*

Remark 2.1.25. The reader may recognize that there is nothing special about working with generic types in Fact 2.1.24: the assertion is just a reflection of the fact that uncountable strongly minimal structures are saturated. However, the most common application for us will be taking generic points.

Note also that if D is infinite, we can actually find $|M|$-many generic points, by iteratively adding each one to the parameter set A before finding a new one.

Example 2.1.26. In an algebraically closed field, $a \in D$ is generic in D over A if and only if it does not belong to any Zariski closed defined over A of dimension smaller than $RM(D)$. In particular, (i) a generic element of an irreducible variety V over A is one which does not belong to any proper closed A-definable subset of V, and (ii) an n-tuple is generic in affine n-space over A if and only if its coordinates are algebraically independent over A.

Example 2.1.27. Similarly, in a vector space V an n-tuple is generic in V^n over A if and only if its coordinates are linearly independent after quotienting V by the subspace generated by A. In a model M of $\mathrm{Th}(\mathbb{Z}, s)$, an n-tuple is generic in M^n over A if and only if its coordinates reside in n distinct copies of $\mathbb{Z}$, each of which is disjoint from A.

By Fact 2.1.24, a definable set D of Morley degree 1 has exactly one generic type over any parameter set. Sets with this property are also called *stationary*. The main combinatorial advantage of stationary sets is the following:

Fact 2.1.28. *If D is stationary, then the generic subsets of D are closed under finite intersection. In particular, they give a complete type over any set A, which is the unique generic type of D over A.*

For example, a variety V over an algebraically closed field is stationary; moreover, if V is affine, its unique generic type over A describes an element whose annihilator ideal over the field generated by A is exactly the annihilator ideal of V over the same field.

Finally, we note that the generic types of a definable set D (over any parameter set) correspond bijectively to the decomposition of D into degree 1 generic sets in Fact 2.1.21: namely, the unique generic type of a degree one generic definable subset of D is also a generic type of D; and two such degree one sets are almost equal if and only if they have the same generic type. For this reason, we will often refer to a degree 1 generic subset of D as a *stationary component* of D. This is especially useful to keep in mind in the case of algebraically closed fields, as the stationary components of a definable set correspond to the top-dimensional irreducible components of its Zariski closure.

2.2 Pregeometries and Local Modularity

In this section we give an overview of the theory of pregeometries and its connection to strongly minimal structures. The main goal for the reader is to motivate the next section, where we present the equivalent description of local modularity in terms of abstract plane curves. Again, for more details, including proofs, the reader could consult [38].

Definition and Examples

Definition 2.2.1. A *pregeometry* is a set X along with a 'closure' operator $\mathrm{cl} : \mathcal{P}(X) \to \mathcal{P}(X)$ with the following properties:

1. $\mathrm{cl}(A) \supset A$ for all $A \subset X$.
2. If $X \subset Y$ then $\mathrm{cl}(X) \subset \mathrm{cl}(Y)$.
3. $\mathrm{cl}(\mathrm{cl}(X)) = \mathrm{cl(X)}$ for all X.
4. If $b \in \mathrm{cl}(A)$ then there is a finite $A' \subset A$ with $b \in \mathrm{cl}(A')$.
5. If $c \in \mathrm{cl}(A \cup \{b\}) - \mathrm{cl}(A)$ then $b \in \mathrm{cl}(A \cup \{c\})$.

Items (1)-(3) in Definition 2.2.1 are standard axioms for closure operators. Item (4) is called *finite character*; properties of this nature appear throughout model theory as reflections of the fact that each formula has finite length. Item (5) is called *Steinitz Exchange*, and generalizes a property used to show that the dimension of a vector space is well defined.

A trivial example is given by letting X be any set and defining $\mathrm{cl}(A) = A$ for all $A \subset X$. We call this the *trivial* pregeometry on X (not to be confused with our definition of triviality for strongly minimal structures, which is broader). Below we give two more interesting examples.

Example 2.2.2. The motivating example of a pregeometry is any vector space V, with the map $A \mapsto \mathrm{Span}(A)$. In this case finite character refers to the fact that span consists only of finite linear combinations of the initial vectors; Steinitz Exchange, roughly, means that any nontrivial dependence $c_1v_1 + \cdots + c_nv_n = 0$ can be solved to express any v_i in terms of the other v_j, provided $c_i \neq 0$. As stated above, this property is crucially used in showing both that (i) bases exist and (ii) any two bases have the same cardinality.

Example 2.2.3. Any algebraically closed field forms a pregeometry with closure given by (field theoretic) algebraic closure. In this light, we think of algebraic closure in fields as a generalization of linear span in vector spaces. Indeed, using this intuition one can prove analogous statements on the existence and unique cardinality of *transcendence* bases in algebraically closed fields. These analogous results are, in essence, the main motivation to study pregeometries in the first place: indeed, the observation that algebraically closed fields form pregeometries is essentially equivalent to Steinitz's characterization of them by characteristic and transcendence degree [46].

Strongly Minimal Structures as Pregeometries

The reader will likely notice that the examples above agree with the examples of strongly minimal structures encountered in the previous section. This is no accident, and is the main reason we study these concepts together. Formally, we define the following:

Definition 2.2.4. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, $a \in M$, and A a set of parameters. Then a is *algebraic over* A if it belongs to a finite A-definable subset of M. Otherwise a is *transcendental over* A. The set of elements algebraic over A is called the *algebraic closure of* A, denoted $\mathrm{acl}(A)$.

In other words, recall that there is exactly one generic type of M over A; we define realizations of this type as transcendentals, and non-realizations – those of rank zero over A – as algebraic.

Remark 2.2.5. Note that algebraic closure does not depend on the choice of model. That is, the algebraic closure of A in $\mathcal{M}$ is the same as the algebraic closure of A in any elementary extension of $\mathcal{M}$.

We note also the following, which follows in each case from quantifier elimination:

Fact 2.2.6. *The notion of algebraic closure in Definition 2.2.4 coincides with linear span in vector spaces, and with field theoretic algebraic closure in algebraically closed fields.*

Thus, as promised, pregeometries provide a common generalization of the closure operator in vector spaces and algebraically closed fields. Generalizing this point further, we are now ready to state the connection between strongly minimal structures and pregeometries:

Fact 2.2.7. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure. Then the map $A \mapsto \mathrm{acl}(A)$ defines a pregeometry on M.*

The verification of items (1)-(4) in Definition 2.2.1 does not use strong minimality, and only depends on working in first-order logic. The main point is that strongly minimal structures satisfy the Steinitz Exchange property.

Bases and Dimension

We now turn toward properties of pregeometries. To start, one can use the intuition of linear algebra to extend various concepts to this more general setting:

Definition 2.2.8. Let (X, cl) be a pregeometry.

1. A set $A \subset X$ is *closed* if cl(A)=A.
2. If $B \subset X$ is closed and $A \subset B$, then A *spans* B if $\mathrm{cl}(A) = B$. In particular, a *spanning set for* X is a set whose closure is all of X.
3. A set $A \subset X$ is *independent* if no $a \in A$ belongs to $\mathrm{cl}(A - \{a\})$.
4. If $B \subset X$ is closed then $A \subset B$ is a *basis for* B if it is independent and spans B.

Using the axioms in Definition 2.2.1, and mimicking the corresponding proof for vector spaces, one can show:

Fact 2.2.9. *Let (X, cl) be a pregeometry, and $B \subset X$ a closed set. Then there is a basis for B. Further, any two bases have the same cardinality.*

In light of Fact 2.2.9, we thus define:

Definition 2.2.10. Let (X, cl) be a pregeometry, and $A \subset X$ a subset. Then the *dimension of* A, denoted $\dim A$, is the cardinality of any basis for $\mathrm{cl}(A)$.

Then, as expected, the dimension defined in Definition 2.2.10 agrees with linear dimension in vector spaces, and with transcendence degree in algebraically closed fields.

In the strongly minimal setting, Fact 2.2.9 has the following structural consequence:

Fact 2.2.11. *Let T be a complete strongly minimal theory – that is, a theory whose models are strongly minimal. Then each model of T is determined up to isomorphism by its dimension as a pregeometry. In particular, if T is countabe then it is uncountably categorical.*

Fact 2.2.11 generalizes the corresponding statements for vector spaces and algebraically closed fields. Indeed, the work of Baldwin and Lachlan on uncountably categorical theories [2] characterizes such theories as those which are 'close enough' to being strongly minimal. This work can be seen as the original reason for studying strongly minimal structures in detail.

Localization and Projectivization

We now describe two constructions which create new pregeometries from old ones. These constructions each have meaning both in the classical example of vector spaces, and in the model theoretic context.

First we define the 'localization' of a pregeometry:

Definition 2.2.12. Let (X, cl) be a pregeometry, and $S \subset X$ any subset. Then the *localization of X at S*, denoted X_S, is another pregeometry on X, with closure operation $\mathrm{cl}_S(A) = \mathrm{cl}(A \cup S)$.

Example 2.2.13. If V is a K-vector space, the 'affine' structure $\mathcal{V}_{\mathrm{aff}}$ on V is the reduct of the full vector space structure $\mathcal{V}$ with language given by the maps $(v_1, v_2) \mapsto cv_1 + (1-c)v_2$ for each $c \in K$. $\mathcal{V}_{\mathrm{aff}}$ is also strongly minimal, as a reduct of $\mathcal{V}$. In fact $\mathcal{V}$ and $\mathcal{V}_{\mathrm{aff}}$ have exactly the same definable sets – one just needs more parameters to define sets in $\mathcal{V}_{\mathrm{aff}}$.

As a pregeometry, $\mathcal{V}$ is the localization of $\mathcal{V}_{\mathrm{aff}}$ at the singleton $\{0\}$ – which just says that, since vector subspaces all include 0, the vector space span of $A \subset V$ is the same as the affine span of $A \cup \{0\}$. More generally, the localization of $\mathcal{V}_{\mathrm{aff}}$ at any singleton is isomorphic as a pregeometry to $\mathcal{V}$, via a map identifying that singleton with 0.

Example 2.2.14. We can generalize the above example to the more abstract model theoretic context: let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, and $A \subset M$ any subset. Let $\mathcal{M}_A$ be the expansion of $\mathcal{M}$ by adding the elements of A to the language as constant symbols. Then $\mathcal{M}_A$ is also strongly minimal, and its associated pregeometry is the localization of $\mathcal{M}$ at A.

Thus localization corresponds in strongly minimal structures to expanding the language by constants.

We now turn toward geometries, and the geometry associated to a pregeometry.

Definition 2.2.15. A pregeometry (X, cl) is a *geometry* if $\mathrm{cl}(A) = A$ whenever $|A| \leq 1$.

Example 2.2.16. The trivial pregeometry on any set is a geometry by definition. Moreover the affine structure on a vector space V gives a geometry, called the *affine geometry* on V;

this is because zero and one element sets form affine subspaces. In contrast, the standard pregeometry on a vector space is not a geometry, nor is the pregeometry on an algebraically closed field: in each case, for example, the closure of the empty set includes the element 0.

Example 2.2.17. Projective Space: If V is a vector space, we define the projectivization of V, V_{proj}, to be the space of all one-dimensional vector subspaces of V. Then V_{proj} is a pregeometry: for a collection $\mathcal{W}$ of one-dimensional subspaces of V, we define the closure of $\mathcal{W}$ to be the projectivization of the vector space spanned by $\bigcup \mathcal{W}$. It is now easy to check that the pregeometry on V_{prof} is a geometry, called the projective geometry associated to V.

An important fact about pregeometries is that any pregeometry carries with it an associated geometry, built in the same way as the projectivization of a vector space:

Definition 2.2.18. Let (X, cl) be a pregeometry. The *geometry associated to* X, also called the *projectivization of* X, is a geometry on the set X_{proj} of all closed subsets of X of dimension 1. Given a set $\mathcal{W}$ of such subsets, the closure of $\mathcal{W}$ is defined as the projectivization of $\mathrm{cl}\,(\bigcup \mathcal{W})$.

Equivalently, the relation $\mathrm{cl}(a) = \mathrm{cl}(b)$ gives an equivalence relation $\sim$ on $X - \mathrm{cl}(\emptyset)$; the underlying set of X_{proj} is the set of $\sim$-classes. For example, if $\mathcal{M} = (M, ...)$ is strongly minimal, then M_{proj} consists of transcendental elements of M up to interalgebraicity.

Example 2.2.19. Let $\mathcal{M}$ be a model of $\mathrm{Th}(\mathbb{Z}, s)$, so $\mathcal{M}$ is stronly minimal. Then the projectivization of $\mathcal{M}$ consists of the set of $\mathbb{Z}$-copies in $\mathcal{M}$, equipped with the trivial closure operator. In the next subsection we will conclude from this phenomenon that $\mathcal{M}$ is 'trivial' as a strongly minimal structure. Roughly, this corresponds to the fact that relations between points in $\mathcal{M}$ which cannot be expressed without parameters are determined 'element-by-element,' rather than by legitimate relationships between several elements.

The Main Classes of Pregeometries

We turn now to the main classes of pregeometries which serve as the motivation for trichotomy problems in model theory: namely, we define degenerate, modular, locally modular, and non-locally modular pregeometries. We start with:

Definition 2.2.20. The pregeometry (X, cl) is *degenerate* if it projectivizes to the trivial pregeometry.

Unwrapping the definition of X_{proj}, an equivalent characterization is the following: for all $A \subset X$ and $b \in X$, if $b \in \mathrm{cl}(A)$ then there is some $a \in A$ with $b \in \mathrm{cl}(\{a\})$.

We also define:

Definition 2.2.21. The saturated strongly minimal structure $\mathcal{M}$ is *trivial* if it is degenerate as a pregeometry – that is, if its associated geometry is trivial. A general strongly minimal structure is trivial if some (any) of its saturated elementary extensions is trivial.

For example, models of $\mathrm{Th}(\mathbb{Z}, s)$ are trivial, as are pure sets; while all structures arising from vector spaces and fields are nontrivial. The intuition behind triviality is that such structures are built *combinatorially* rather than *algebraically* – indeed, trivial structures can be characterized as those not admitting 'non-trivial' two-dimensional subsets of three-space (e.g. graphs of binary operations are excluded):

Fact 2.2.22. *The following are equivalent for a strongly minimal structure $\mathcal{M} = (M, ...)$:*

1. *$\mathcal{M}$ is trivial.*
2. *Every rank 2 definable set $D \subset M^3$ is almost equal to a finite union of products (in some order) of subsets of M and M^2.*
3. *There do not exist transcendentals a, b, c in an elementary extension of $\mathcal{M}$ such that any two of a, b, c are independent but the full triple (a, b, c) is not.*

Remark 2.2.23. The motivation for item (3) in Fact 2.2.22 is the following: if $\mathcal{G} = (G, ...)$ is a saturated strongly minimal group, let a and b be independent transcendentals, and $c = ab$. Then one can check that a, b, c are transcendentals, any two are independent, and the whole triple is not. It follows from this that trivial strongly minimal structures cannot interpret infinite groups, motivating their intuitive status as 'non-algebraic.'

We next define modularity:

Definition 2.2.24. The pregeometry (X, cl) is *modular* if whenever $A, B \subset X$ are closed and $c \in \mathrm{cl}(A \cup B)$, there exist $a \in A, b \in B$ with $c \in \mathrm{cl}(\{a, b\})$. A strongly minimal structure $\mathcal{M}$ is modular if some (any) of its saturated elementary extensions has modular pregeometry.

Example 2.2.25. It is obvious from the definitions that degenerate pregeometries are modular. The motivating example of a modular pregeometry is a vector space: indeed, any element in the span of two subspaces $A, B \subset V$ is a sum of one element from A and one element from B. One can also check that the projectivization of a vector space is modular.

Modularity is also equivalent to the following familiar dimension formula from vector spaces – indeed, many take this formula as the definition of modularity:

Fact 2.2.26. *If (X, cl) is a pregeometry, the following are equivalent:*

1. *X is modular.*
2. *For any two finite-dimensional closed sets $A, B \subset X$ we have*

$$\dim(\mathrm{cl}(A \cup B)) = \dim A + \dim B - \dim(A \cap B).$$

Finally we define 'local modularity,' a slight generalization of modularity:

Definition 2.2.27. A pregeometry (X, cl) is *locally modular* if there is some $s \in X$ such that the localization $X_{\{s\}}$ is modular. The strongly minimal structure $\mathcal{M}$ is locally modular if some (any) of its saturated elementary extensions has locally modular pregeometry.

It is easy to check that modularity is preserved under localization. Thus the reader may wonder why we insist on only localizing by one point in Definition 2.2.27. In fact this does not truly matter – the only issue is that we can't allow too big of a localization since, for example, X_X is modular for any pregeometry (X, cl). To clarify things in the model theoretic context, we note the following:

Fact 2.2.28. *Let $\mathcal{M} = (M, ...)$ be a saturated strongly minimal structure. Then the following are equivalent:*

1. *For some $S \subset M$ with $|S| < |M|$ the localized pregeometry M_S is modular.*
2. *For all $S \subset M$ with $|S| < |M|$ and $S \not\subset acl(\emptyset)$, the localized pregeometry M_S is modular.*
3. *$\mathcal{M}$ is locally modular.*

In other words, if $\mathcal{M}$ becomes modular upon localizing in an interesting way, then we can take the localization set to be any single transcendental – or indeed any 'small' set containing a transcendental.

Example 2.2.29. It is obvious from the definitions that modular pregeometries are locally modular. The standard example of a locally modular pregeometry which is not modular is the affine geometry on a vector space. Of course we have seen that such a geometry becomes modular after localizing at 0. To see why it is not already modular, let A and B be parallel lines, and b, c distinct points on B. Then $c \in \mathrm{cl}(A \cup \{b\})$; however this fails if A is replaced by any $a \in A$. In terms of the above dimension formula we have $\dim A = \dim B = 2$, $\dim(A \cap B) = 0$, but $\dim(\mathrm{cl}(A \cup B)) = 3$.

One of the main motivations for local modularity is that it generalizes vector spaces but is preserved under localization and projectivization. Model theoretically, we have the following:

Fact 2.2.30. *Let $\mathcal{M}$ be a saturated strongly minimal structure. Then $\mathcal{M}$ is locally modular if and only if its projectivization is. Further, if $\mathcal{M}$ is locally modular then any strongly mninimal structure interpreted in $\mathcal{M}$ is also locally modular.*

Local modularity is thus seen as a broad notion capturing the 'simple' strongly minimal structures. There is only one easy known way to build a strongly minimal which does not reside in this class (though others do exist). In a sense this example forms the basis for trichotomy problems:

Example 2.2.31. No algebraically closed field K is locally modular. To see why, suppose we add any set S of constants to the language. Then, replacing K with an elementary extension if necessary, we take a solution to the equation $y = ax + b$ with a, b, x transcendental and independent over S. Then, in the localization K_S, y is in the closure of $\text{cl}(x) \cup \text{cl}(\{a, b\})$. However, one can check that this fails if $\text{cl}(\{a, b\})$ is replaced with any single one of its elements.

Geometrically, the above example says that the family of lines in the plane is two dimensional – and correspondingly the given dependence between x and y can only be described with two parameters.

2.3 Plane Curves

This section is devoted to the reinterpretation of trivial, locally modular, and non-locally modular strongly minimal structures in the way we will consider them – by using families of abstract plane curves.

Definition and Examples

Throughout this section, assume $\mathcal{M} = (M, ...)$ is a saturated strongly minimal structure.

Definition 2.3.1. We define abstract curves as follows:

1. A *curve in* $\mathcal{M}$ is a definable set $C \subset M^n$ with $\text{RM}(C) = 1$.
2. An *irreducible curve* is a curve which is strongly minimal – that is, of degree 1.
3. A *plane curve* (resp. *irreducible plane curve*) is a curve (resp. irreducible plane curve) $C \subset M^2$.

Note, then, that every curve is a finite union of irreducible curves.

By stationarity, an irreducible curve has exactly one generic type over any set. Irreducible curves are often identified with their generic types, which amounts to thinking of them 'up to finite error,' or equivalently up to almost equality. We now make this precise. To start, the following is immediate from the definition of strong minimality:

Fact 2.3.2. *If $C_1, C_2 \subset M^n$ are irreducible curves, then either $C_1 \Delta C_2$ or $C_1 \cap C_2$ is finite.*

So any two irreducible curves are either equal or disjoint after ignoring finitely many points. We thus define:

Definition 2.3.3. The irreducible curves $C_1, C_2 \subset M^n$ are *equivalent*, denoted $C_1 \sim C_2$, if $C_1 \Delta C_2$ is finite.

The following properties are easy to check:

Fact 2.3.4. *The following hold of the relation $\sim$ defined above.*

1. *$\sim$ is an equivalence relation on irreducible curves.*
2. *Two irreducible curves are equivalent if and only if they are almost equal, or equivalently they have the same generic type.*
3. *$\sim$ is definable in families.*

We now illustrate these concepts by describing the $\sim$-classes of irreducible plane curves in canonical examples of trivial, locally modular, and non-locally modular strongly minimal structures:

Example 2.3.5. Trivial: If $\mathcal{M}$ is a pure set, the classes of irreducible plane curves are represented by the following, which of course appear in every strongly minimal structure:

- $\{m\} \times M$ for each $m \in M$.
- $M \times \{m\}$ for each $m \in M$.
- $\{(m, m) : m \in M\}$.

Furthermore, if $\mathcal{M}$ is a model of $\mathrm{Th}(\mathbb{Z}, s)$, the additional classes present are all of the form $\{(m, s^n(m)) : m \in M\}$ for $n \in \mathbb{Z}$.

Remark 2.3.6. Plane curves equivalent to one of the first two types above have been called *straight lines* by [12]. To avoid confusion with actual lines in other examples, we will call them *trivial* curves.

Example 2.3.7. Locally Modular: If $\mathcal{M}$ is a K-vector space, the classes of irreducible plane curves are all of the form $\{(x, y) : ax + by = v\}$ for some triple $(a, b, v) \in K^2 \times M$. Of course the triple (a, b, v) is only defined up to scaling, and there are certain exceptional triples which do not define curves. Note that this presentation inlcudes all irreducible plane curves which appear in the pure set, by assigning a and b to certain values among $\{-1, 0, 1\}$.

Example 2.3.8. Non-Locally Modular: If $\mathcal{M}$ is an algebraically closed field, then every class of irreducible plane curves is identifiable as the roots of a nonzero irreducible polynomial $p(x, y)$. That is, the model theoretic notion correctly captures the geometric notion – an abstract irreducible plane curve is the same thing as an actual plane curve plus or minus finitely many points.

Families

The main goal of this section is to study *families* of plane curves. We will be able to use such families to characterize triviality and local modularity in a way that will be more applicable to the rest of the thesis.

Definition 2.3.9. A *family* $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves is a definable set $F \subset M^2 \times A$, for some interpretable set A, such that F_a is a plane curve for each $a \in A$. We further say $\mathcal{F}$ is a *generically irreducible family* if for every generic $a \in A$ the curve F_a is irreducible.

Remark 2.3.10. To clarify, we use the scripts $\mathcal{F}$ and F for subtly different purposes: namely, we reserve F for when we specifically want to view $\mathcal{F}$ itself as a single interpretable set. For example, below we will develop a notion of rank for families; the notation $\text{rk } \mathcal{F}$ will refer to the rank of the family of curves, as given below, while $\text{rk } F$ will refer to the Morley rank of the interpretable set F representing the family $\mathcal{F}$.

Remark 2.3.11. Note that Morley rank is definable in strongly minimal structures; while in general irreducibility is only type definable. For this reason we can assume all fibers in a family are in fact curves, but we can only assume the generic ones are irreducible.

We now define the rank of a family of plane curves. In full generality this requires the theory of canonical bases; however we will quickly turn our attention toward special types of families whose ranks can be understood more easily.

Definition 2.3.12. We define each of the following:

1. If D is a stationary definable set in $\mathcal{M}$, then we define the *code of* D to be the canonical base of the generic type of D over some (equivalently any) set capable of defining D.

2. If D is a stationary definable set in $\mathcal{M}$, and A is a set of parameters, then we define the *complexity of* D *over* A to be the Morley rank of the code of D over A.

3. If D is any non-empty definable set in $\mathcal{M}$, and A is a set of parameters, then we define the *complexity of* D *over* A to be the maximum complexity of a stationary component of D over A.

4. Let $\mathcal{F}$ be a family of plane curves parametrized by the set A, and definable over a set $B \subset M$. We define the *rank of* $\mathcal{F}$ to be the maximum complexity of a curve F_a over B, among those $a \in A$ which are generic in A over B.

Remark 2.3.13. Some comments on this definition are needed:

1. Most importantly, we note that the rank of $\mathcal{F}$ is a well-defined natural number, which is not clear from the definition. Indeed, since A has only finitely many generic types over B, and complexity is automorphism invariant over B, it follows that the rank is the maximum of finitely many natural numbers, and thus is itself a natural number.

Additionally, the precise set B does not matter, as long as it is capable of defining $\mathcal{F}$; roughly, this follows since the generic extensions of a stationary type over different parameter sets are compatible.

2. Our use of the word 'code' is slightly non-standard; indeed, the experienced reader may have expected to see the similar notion of a 'canonical parameter.' We have chosen the terminology and definition above because we want to only distinguish sets up to almost equality.

3. As hinted at in the previous item, note that the code of a stationary definable set is determined by its generic type; so almost equal sets have the same code.

4. The code of a stationary definable set is unique up to interdefinability; thus we will use 'the code' and 'a code' interchangeably. Up to interalgebraicity, one can always take the code to be a tuple of elements of M; one may need imaginaries to give the precise code, but this is in general harmless.

5. It is not in general true that a stationary definable set D is definable over its code c; but this is true 'up to almost equality' – namely, there is always a set D' which is almost equal to D and definable over c.

6. Finally, we emphasize that in computing rk $\mathcal{F}$ we are only interested in the *generic* curves in $\mathcal{F}$. For example, if almost every curve in $\mathcal{F}$ is 0-definable, while a small subset of the indices contain positive complexity curves, we still say the rank of F is 0.

Faithfulness

We next introduce two conditions on a family which make this rank calculation simpler.

Definition 2.3.14. A family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves is *faithful* if for distinct $a, b \in A$ the sets $F_a, F_b \subset M^2$ have finite intersection. The family is *almost faithful* if for each $a \in A$ there are only finitely many $b \in A$ such that $F_a \cap F_b$ is infinite.

The motivating example of a faithful family is a family of irreducible, pairwise inequivalent plane curves. Given any generically irreducible family, one can use imaginaries to 'normalize' it into a faithful family:

Definition 2.3.15. Two families $F \subset M^2 \times A$ and $G \subset M^2 \times B$ are *equivalent* if for all generic $a \in A$ there is a generic $b \in B$ such that $F_a \Delta G_b$ is finite, and for all generic $b \in B$ there is a generic $a \in A$ such that $F_a \Delta G_b$ is finite.

Fact 2.3.16. *If $\mathcal{F}$ is any generically irreducible family of plane curves, indexed by the interpretable set A, then F is equivalent to a faithful family.*

A brief proof sketch for Fact 2.3.16 is as follows: using the Compactness Theorem, there is an interpretable set $A' \subset A$ containing all generics of A such that any two curves in A' either have finite intersection or finite symmetric difference (even though they might not all be irreducible). One then quotients A' by the equivalence relation of the corresponding curves having finite symmetric difference, and in the resulting family any two distinct indices correspond to curves with finite intersection.

In light of Fact 2.3.16, one may wonder why we additionally define almost faithfulness. The main point is to try to avoid quotienting operations, so the parameter family remains a definable (rather than interpretable) set. We get the following result; while it is not original, it is quite central to our future work, so we include a proof outline:

Fact 2.3.17. *If $\mathcal{F}$ is any generically irreducible family of plane curves, indexed by the definable set A, then $\mathcal{F}$ is equivalent to an almost faithful family, possibly defined over additional parameters, which is still indexed by a definable set.*

Proof. Using compactness, we may replace A with one of its generic subsets and assume the relation $a \sim b$, defined by the assertion that $F_a \Delta F_b$ is finite, is an equivalence relation on A. We may also assume $\operatorname{acl}(\emptyset)$ is infinite, since we are allowed additional parameters.

Now in general strongly minimal structures do not eliminate imaginaries. However, they do 'up to finite error' if the empty set has infinite algebraic closure (see for example [24], page 137). That is, there is a definable set B, and an interpretable finite-to-one surjective map $f : B \to A/\sim$. We now reparametrize by B: for $b \in B$, set G_b to be the set of all $\overline{m} \in M^2$ which belong to F_a for all generic $a \in f(b)$. Then $\mathcal{G} = \{G_b\}_{b \in B}$ consists of the same curves as $\mathcal{F}$ up to equivalence; further, since f is finite-to-one it follows immediately that $\mathcal{G}$ is almost faithful. □

Now, as promised, we note the following consequence of almost faithfulness:

Fact 2.3.18. *Let $\mathcal{F}$ be an almost faithful family of curves, parametrized by A. Then*

$$rk\ \mathcal{F} = RM(A).$$

The proof boils down to the observation that, if $(m_1, m_2, a) \in F$ is generic, then the canonical base for $\operatorname{stp}(m_1 m_2/a)$ is, by almost faithfulness, interalgebraic with a.

We conclude this subsection with some examples of families in the structures we have stressed repeatedly:

Example 2.3.19. Trivial: If $\mathcal{M}$ is a model of $\operatorname{Th}(\mathbb{Z}, s)$, then every definable, generically irreducible family of plane curves of positive rank is equivalent to either $\mathcal{F} = \{\{m\} \times M\}_{m \in M}$, $\mathcal{G} = \{M \times \{m\}\}_{m \in M}$, or the union of these two. In other words, we can only take families of 'trivial' curves. Each of these two families is faithful of rank 1.

Example 2.3.20. Locally Modular: If $\mathcal{M}$ is a K-vector space, we can take the family $\mathcal{F}_{ab}$ (for each $(a,b) \in K^2 - \{(0,0)\}$) defined by the equations $ax + by = v$ for each $v \in M$. Note that we cannot vary a and b since they are function symbols in the language, not elements of the model. Then each $\mathcal{F}_{ab}$ is a faithful family of rank 1.

Example 2.3.21. Non-Locally Modular: If $\mathcal{M}$ is an algebraically closed field, we have much more freedom in choosing families of plane curves. For example, the family defined by $y = ax + b$ for each $(a,b) \in M^2$ is a faithful family of rank 2. Indeed, by varying the coefficients in polynomial equations of high degree, one can obtain families of plane curves of any desired rank.

Characterization of Local Modularity

We now turn toward the classical application of families of curves in characterizing local modularity. We begin by formally defining 'triviality' of plane curves – as we have identified in our examples above – and subsequently discussing the composition operation on non-trivial plane curves:

Definition 2.3.22. A plane curve $C \subset M^2$ is *trivial* if it almost contains $\{m\} \times M$ or $M \times \{m\}$ for some $m \in M$; otherwise C is *non-trivial*.

Remark 2.3.23. We make two brief comments regarding Definition 2.3.22:

1. By strong minimality, C is non-trivial if and only if both projections $C \to M$ are everywhere finite-to-one; if this is the case, it follows again by strong minimality that both projections have cofinite image in M.

2. Note that non-triviality is definable in families; so, when dealing with families of curves whose generic members are non-trivial, there is no harm in restricting to a generic subfamily and assuming all curves are non-trivial.

Definition 2.3.24. Let $\mathcal{F} = \{F_a\}_{a \in A}$ and $\mathcal{G} = \{G_b\}_{b \in B}$ be families of non-trivial plane curves. We define the *composition family* $\mathcal{F} \circ \mathcal{G}$, indexed by $A \times B$, as follows: for each $a \in A$ and $b \in B$, set

$$F_a \circ G_b = \{(m_1, m_3) : \exists m_2((m_1, m_2) \in G_b \wedge (m_2, m_3) \in F_a)\}.$$

That is, we can form a new family by 'composing' each curve from one family with each curve from another. A priori, we would expect the rank of $\mathcal{F} \circ \mathcal{G}$ to be the sum of the ranks of $\mathcal{F}$ and $\mathcal{G}$. If this is not the case, one can try to glean relationships between the curves in the two families. One is led to the following, which goes back to the work of Zilber (see for example [53], Theorem 3.1 for a weaker statement) and Hrushovski (see [22], Theorems 3.4.2 and 4.1.1), and can be viewed as an original, and actualy true, 'trichotomy' phenomenon:

Fact 2.3.25. *Let* $S = \sup\{k : \text{there is a family of non-trivial plane curves of rank } k\}$*. Then* S *is either 0, 1 or* ∞*.*

The proof of Fact 2.3.25 follows the intuition of the preceding paragraph, and roughly proceeds as follows: assuming $1 < S < \infty$, one takes a family $\mathcal{F}$ of maximal rank k, and notes that $\mathcal{F} \circ \mathcal{F}$ must also have rank k. One then converts this data into a 'group configuration' (see [38], Chapter 5), obtaining a definable k-dimensional group acting transitively on a strongly minimal set. One then applies a classification of such actions ([22], Theorem 4.1.1) and concludes that the given group is one which interprets an agebraically closed field. This is a contradiction since, assuming the presence of a field, it has been noted above that $S = \infty$.

By Fact 2.3.25, we obtain three clearly distinguished 'complexity classes' of strongly minimal structures – forming the basis for trichotomy conjectures on such structures. In fact, we have just recovered the same classes that we have studied in the previous section:

Fact 2.3.26. *Let* S *be as in Fact 2.3.25. Then* $S = 0$ *if and only if* $\mathcal{M}$ *is trivial,* $S = 1$ *if and only if* $\mathcal{M}$ *is non-trivial and locally modular, and* $S = \infty$ *if and only if* $\mathcal{M}$ *is non-locally modular.*

Thus $\mathcal{M}$ is non-locally modular if and only if there exists a rank 2 family of curves, if and only if there exist families of arbitrarily high rank. It is easy to show that we can take such families to be generically irreducible; thus, combining with Fact 2.3.17, we conclude:

Fact 2.3.27. *If* $\mathcal{M}$ *is non-locally modular then we can find almost faithful, generically irreducible families of plane curves of arbitrarily high rank, indexed by definable sets in* $\mathcal{M}$*.*

The main objective in establishing the local modularity of $\mathcal{M}$, then, is to show that one only has such families of bounded rank. We develop new methods to do this in Chapters 4 and 5, and also reduce to preexisting methods in Chapters 4 and 6.

2.4 Trichotomy Conjectures and Our Setting

So far in this chapter we have presented the results that motivate the study of 'trichotomy' phenomena for strongly minimal structures. The landmark idea of Zilber was to characterize the three types of strongly minimal structures not just by abstract properties, but by the presence of familiar algebraic structures. In this short section we present the original and restricted forms of the trichotomy conjecture, and conclude with a description of the setting we will work with throughout most of the thesis.

Locally Modular Structures Seen Algebraically

The three motivating examples for strongly minimal structures are pure sets, vector spaces, and algebraically closed fields – which seems to match the trichotomy we have now seen in two different ways for strongly minimal structures. One may ask whether these three main examples are 'intrinsic' to their respective complexity classes. For the first two, there is a sense in which this holds:

Fact 2.4.1 (Zilber, e.g. [54]). *The strongly minimal structure $\mathcal{M}$ is trivial if and only if its associated geometry is isomorphic to that of a pure set. $\mathcal{M}$ is non-trivial and locally modular if and only if its associated geometry is isomorphic to an affine or projective space over a division ring.*

So, at the level of associated geometries, the main examples are the 'only' examples. In the locally modular case, we can be more precise:

Fact 2.4.2 (Hrushovski [22]). *Let $\mathcal{M}$ be a non-trivial locally modular strongly minimal structure. Then in $\mathcal{M}$ there is an interpretable strongly minimal group G, such that all definable subsets of G^n, for all n, in the induced structure on G are Boolean combinations of cosets of definable subgroups.*

Remark 2.4.3. Fact 2.4.2 is often thought of as saying '$\mathcal{M}$ is essentially a vector space.' Indeed, the group G obtained is either a vector space over a finite field, or divisible abelian with 'small' torsion (i.e. finitely many elements of each finite order). The point about definable sets in powers of G is interpreted as saying that G only carries its 'lineear' structure. Note that after quotienting G by its torsion subgroup, one obtains a vector space over the division ring of definable quasi-endomorphisms of G; the associated geometry of G as a strongly minimal set is precisely the projectivization of this vector space.

So, under the assumption of non-triviality and local modularity one is able to 'recover' a vector space. Zilber [53] predicted that, given the plethora of plane curves provided by non-local modularity, one should be able to 'recover' a field:

Conjecture 2.4.4 (Zilber). *If $\mathcal{M}$ is strongly minimal and non-locally modular, then $\mathcal{M}$ interprets an algebraically closed field.*

As stated in the introduction, this conjecture is false. One of the oldest attempts to salvage it is:

Conjecture 2.4.5 (Zilber's Restricted Trichotomy Conjecture). *If $\mathcal{M}$ is a strongly minimal, non-locally modular structure which is interpretable in the algebraically closed field K, then $\mathcal{M}$ interprets a field isomorphic to K.*

Remark 2.4.6. We make a few comments:

1. By [42] (Theorem 4.15), if $\mathcal{M}$ interprets any infinite field, then that field is isomorphic to K. So the challenge is really to interpret a field.

2. By elimination of imaginaries in algebraically closed fields [43], we can assume the universe of $\mathcal{M}$ is definable, i.e. a Boolean combination of varieties over K.

3. If the universe M of $\mathcal{M}$ is a rational curve over K, the interpretability of the field was proven in [44]. More generally, the currently unpublished paper [20] generalizes Rabinovich's result to universes M having dimension 1 as a K-definable set. So the conjecture reduces to the case that $\dim M > 1$.

4. Also by [42], if K is interpretable then it follows that $\dim M = 1$ (see Corollary 3.1.16 of this thesis). So a positive answer is equivalent to the local modularity of all strongly minimal $\mathcal{M}$ with higher dimensional universes.

5. If a counterexample exists, then one exists in any saturated algebraically closed field of the same characteristic, and the language can be taken to be finite (in which case $\mathcal{M}$ is also saturated and of the same cardinality as K). We will at times add to the language, but all such extensions will remain countable. So, when assuming characteristic zero (which we do throughout), we may always assume K to be the $\aleph_1$-saturated field $\mathbb{C}$, and thereby obtain both (1) access to generic points of definable sets over countable sets of parameters in both K and $\mathcal{M}$, and (2) access to analytic reasoning in studying definable sets.

6. When working over $\mathbb{C}$, we will frequently use both the Zariski and analytic topologies on varieties. Note, importantly, that the Zariski and analytic closures agree on $\mathbb{C}$-definable sets (by [33], I.10, Corollary 1, and the fact that by quantifier elimination constructible is equivalent to definable). It follows that the same holds for related operations such as interior and frontier. We will abuse this fact liberally in the subsequent chapters.

Our Setting

The majority of this thesis is devoted to addressing the case of higher dimensional universes in characteristic zero. We now conclude this chapter by briefly summarizing what this means. This will serve as the setting for the rest of the thesis, with the sole exception being Chapter 6.

To start, we are given:

- An algebraically closed field K, which we typically assume is the complex field $\mathbb{C}$.
- A strongly minimal structure $\mathcal{M} = (M, ...)$, such that all definable subsets of powers of M are definable over K.
- The assumption that $\dim M > 1$ as a K-definable set.

Our goal is to show $\mathcal{M}$ cannot have families of plane curves of arbitrarily high rank. We will do this under various assumptions by using the background geometry of K to place restrictions on $\mathcal{M}$-definable sets.

Chapter 3

First Observations and Almost Purity

In this chapter we introduce several of the main technical observations we will need throughout the remaining chapters. Most importantly, we introduce the geometric notions of 'almost purity' and 'almost closedness,' under which we will be able to solve the Restricted Trichotomy Conjecture for certain universes. These notions represent, essentially, those definable sets which 'agree enough' with the background geometry of the field.

3.1 Dimension and Genericity

We first discuss certain issues arising from dimension and genericity. The proofs in this section are quite straightforward and elementary; however we include them in full, since they involve notions that we will use later on.

Convention 3.1.1. Throughout this section, assume K and $\mathcal{M} = (M, ...)$ satisfy the hypotheses outlined at the end of the last chapter: K is a saturated algebraically closed field, and $\mathcal{M}$ is a saturated strongly minimal structure definable in K whose universe is of dimension greater than 1.

We thus have *two* strongly minimal structures to work with, K and $\mathcal{M}$. These two structures each carry their own version of the rank and degree functions RM and DM. One of the main points of the higher dimensional version of the problem is that these functions do *not* coincide – indeed, the universe M has dimension 1 according to $\mathcal{M}$ but a higher dimension according to K.

Convention 3.1.2. For sets D definable in $\mathcal{M}$, we will use the notation rk D to refer to its dimension according to $\mathcal{M}$, and $\dim D$ to refer to its dimension according to K. Similarly for points and types, we use rk when computing according to $\mathcal{M}$ and dim when computing according to K. When discussing the genericity of points in $\mathcal{M}$-definable sets, we use the notation $\mathcal{M}$-generic when computing in $\mathcal{M}$ and K-generic when computing in K.

Following the above convention, our first goal is to describe the relationship between dim and rk for $\mathcal{M}$-definable sets. The following notion will be useful throughout:

Definition 3.1.3. Let $\mathcal{X}$ be any strongly minimal structure, and let D, E be definable sets with $D \subset E$. Then D is *fully generic in* E if $E - D$ is non-generic in E.

Remark 3.1.4. As with our use of the word 'generic,' Definition 3.1.3 is not standard. It will, however, be a useful notion to have.

The idea of 'fully generic' is that we omit, for example, E having degree ≥ 2 and D being one of its degree 1 components. The following are now straightforward to show:

Lemma 3.1.5. *Let $\mathcal{X}$ be any saturated strongly minimal structure, and D, E definable sets with $D \subset E$.*

1. *If D is fully generic in E then D is generic in E.*
2. *D is generic in E if and only if it contains at least one generic element of E, over some (any) set capable of defining D and E.*
3. *D is fully generic in E if and only if it contains all generic elements of E, over some (any) set capable of defining D and E.*
4. *If DM $(E) = 1$ and D is generic in E, then D is fully generic in E.*

Proof. Without loss of generality assume D and E are $\emptyset$-definable.

1. Since $E = D \cup (E - D)$, one of D and $E - D$ has the same rank as E. If D is fully generic in E then $E - D$ has smaller rank than E, so D must have the same rank as E. Thus D is generic in E.
2. Fitst assume D is generic in E. Let $d \in D$ be generic. Then $d \in E$, and

   $$\text{RM }(d) = \text{RM }(D) = \text{RM }(E),$$

   so d is generic in E.

   Now assume D contains a generic element $e \in E$. Then since $e \in D$,

   $$\text{RM }(E) = \text{RM }(e) \leq \text{RM }(D).$$

   On the other hand $D \subset E$, so RM $(D) \leq$ RM (E). Thus RM $(D) =$ RM (E), so D is generic in E.
3. The statement that D is fully generic in E is equivalent to the statement that $E - D$ is not generic in E – which, by (2), is equivalent to the statement that $E - D$ does not contain generic elements of E, i.e. D has all generics of E.

4. If DM $(E) = 1$ then E has exactly one generic type: hence all generics of E belong to the same $\emptyset$-definable sets – and thus, if one belongs to D, all of them do.

$\square$

We can now define:

Definition 3.1.6. Let $\mathcal{X} = (X, ...)$ be any strongly minimal structure, D and E definable sets, and $f : D \to E$ a definable function.

1. f is *almost finite-to-one* if the (definable) union of all finite fibers is fully generic in D.

2. f is *almost surjective* if its image is fully generic in E.

Remark 3.1.7. Note, for example, that almost finite-to-one is *not* saying generic target elements have finite fibers. For example, a coordinate projection of the union of two lines in the plane, one horizontal and vertical, is not almost finite-to-one.

Almost finite-to-one and almost surjective functions have the following rank preservation properties:

Lemma 3.1.8. *Let $\mathcal{X}$ be any saturated strongly minimal structure, D and E definable sets, and $f : D \to E$ a definable function.*

1. *If f is almost finite-to-one then RM $(D) \leq$ RM (E).*

2. *If f is almost surjective then RM $(D) \geq$ RM (E).*

3. *If f is almost finite-to-one and almost surjective then RM $(D) =$ RM (E).*

Proof. We may assume all relevant data is $\emptyset$-definable.

1. Let $d \in D$ be generic. Since f is almost finite-to-one, $f(d)$ has a finite preimage. Thus d is algebraic over $f(d) \in E$, so

$$\text{RM } (D) = \text{RM } (d) \leq \text{RM } (f(d)) \leq \text{RM } (E).$$

2. Let $e \in E$ be generic. Since f is almost surjective, there is some $d \in D$ with $f(d) = e$. Then e is algebraic over d, so

$$\text{RM } (D) \geq \text{RM } (d) \geq \text{RM } (e) = \text{RM } (E).$$

3. This follows from the two previous clauses.

$\square$

The main point of the above notions is to give a smooth proof of the following straightforward fact, which in turn allows for an easy characterization of the relationship between dim and rk:

Lemma 3.1.9. *Let $\mathcal{X} = (X, ...)$ be any saturated strongly minimal structure, and $D \subset X^k$ a non-empty definable set with $RM(D) = r$. Then there is a definable, almost finite-to-one, almost surjective function $f : D \to X^r$ (here we identify X^0 with a singleton). Moreover if $DM\ (D) = 1$ then f can be taken to be a coordinate projection.*

Proof. We first assume $DM(D) = 1$. We may further assume D is $\emptyset$-definable. Now let $d = (d_1, ..., d_k)$ be a generic element of D. Then d has dimension r in the pregeometry structure on X. Now take a basis for d – that is, an r-element set $B \subset \{1, ..., k\}$ so that the tuple $d_B = \{d_i : i \in B\}$ is a basis for the closed set acl (d). Thus RM $(d_B) = r$, so d_B is generic in X^r.

Let f be the restriction to D of the projection $X^k \to X^r$ to the coordinates in B. Then d_B is in the image of f. Since d_B is generic in X^r, it follows that im f is generic in X^r. Since X^r has degree 1, this implies im f is fully generic, hence f is almost surjective.

Claim 3.1.10. *There are only finitely many elements of D which project to d_B under π.*

Proof. Otherwise the set of such elements of D has dimension at least 1; let d' be such an element which is generic, so $\mathrm{RM}(d'/d_B) \geq 1$. Then by additivity it follows that

$$\mathrm{RM}\ (d') = \mathrm{RM}\ (d_B, d') \geq r + 1,$$

contradicting that $d' \in D$ and RM $(D) = r$. □

We have just deduced that $f(d)$ has a finite fiber under f. Since d is generic in D and D has degree 1, the set of elements with this property is fully generic in D – thus f is almost finite-to-one.

This concludes the proof if D has degree 1. Now if DM $(D) = m$, write D as a disjoint union of degree 1 sets, say $D_1 \cup ... \cup D_m$. Choose a projection f_i satisfying the claim for each D_i; then the function $f = \bigcup f_i$ suffices for D. □

We now arrive at the main goal of this section:

Corollary 3.1.11. *Let D be any non-empty $\mathcal{M}$-definable set. Then*

$$\dim D = \dim M \cdot rk\ D.$$

In particular, the dimension of any $\mathcal{M}$-definable set is a multiple of $\dim M$*.*

Proof. We work by induction on $r = \mathrm{RM}\ (D)$; the statement is clear if $r = 0$. So, assume $r > 0$ and the claim is true for all $r' < r$. We first apply Lemma 3.1.9 to $\mathcal{X} = \mathcal{M}$, obtaining an $\mathcal{M}$-definable function $f : D \to M^r$ which is almost finite-to-one and almost surjective according to $\mathcal{M}$. We want to conclude that f is also almost finite-to-one and almost surjective according to K.

Let A be the $\mathcal{M}$-definable set of elements of D belonging to infinite fibers in f, and let B be the $\mathcal{M}$-definable set of elements of E which are not in the image of f. Then since f is almost finite-to-one and almost surjective, both A and B have rank less than r, so the statement of the corollary holds for A and B. In particular, each of $\dim A$ and $\dim B$ is strictly less than

$$r \cdot \dim M = \dim M^r.$$

Now since $\dim B < \dim M^r$, it follows that B is not K-generic in M^r, so f is almost surjective according to K. Thus

$$\dim D \geq \dim M^r > \dim A.$$

But this further implies A is not K-generic in D, so f is also almost finite-to-one according to K. Then, since f is almost finite-to-one and almost surjective according to K, Lemma 3.1.8 implies

$$\dim D = \dim M^r = r \cdot \dim M,$$

as desired. □

Corollary 3.1.11 has the following immediate consequences, which we also use throughout:

Corollary 3.1.12. *Let D and E be $\mathcal{M}$-definable sets. Then:*

1. *If $D \subset E$ then D is $\mathcal{M}$-generic (resp. fully generic) in E if and only if D is K-generic (resp. fully generic) in E.*
2. *D is almost contained in E according to $\mathcal{M}$ if and only if D is almost contained in E according to K.*
3. *D and E are almost equal according to $\mathcal{M}$ if and only if D and E are almost equal according to K.*

Proof. 1. First assume D is $\mathcal{M}$-generic in E. Then

$$\dim D = \dim M \cdot \text{rk } D = \dim M \cdot \text{rk } E = \dim E,$$

so D is K-generic in E.

Next assume D is K-generic in E. Then

$$\text{rk } D = \frac{\dim D}{\dim M} = \frac{\dim E}{\dim M} = \text{rk } E,$$

so D is $\mathcal{M}$-generic in E.

Finally, the corresponding statements about full genericity immediately follow by replacing D with $E - D$.

2. This follows from (1) and the fact that almost containment is solely determined by whether $D \cap E$ is generic in D.

3. This follows immediately from (2).

□

So the correspondence between dim and rk on definable sets is quite well behaved. From now on we will not distinguish between $\mathcal{M}$-genericity and K-genericity of sets – we will simply use the term 'generic.'

For points and types, the relationship is not quite as desirable:

Corollary 3.1.13. *Let D be an $\mathcal{M}$-definable set, A a set of parameters defining D, and a a tuple. Then:*

1. $\dim(a/A) \leq \dim M \cdot rk\ (a/A)$.

2. *If a is K-generic in D over A then a is $\mathcal{M}$-generic in D over A.*

Proof. 1. Let E be a set of smallest rank which contains a and is $\mathcal{M}$-definable over A. Then E is K-definable, contains a, and has dimension

$$\dim M \cdot \text{rk } E = \dim M \cdot \text{rk } (a/A),$$

which means

$$\dim(a/A) \leq \dim E = \dim M \cdot \text{rk } (a/A).$$

2. If a is K-generic in D over A then

$$\dim(a/A) = \dim D = \dim M \cdot \text{rk } D.$$

At the same time, by (1) we have

$$\dim(a/A) \leq \dim M \cdot \text{rk } (a/A).$$

Thus rk $D \leq$ rk (a/A). Since $a \in D$ this implies rk $(a/A) =$ rk D, thus a is $\mathcal{M}$-generic in D over A.

□

Convention 3.1.14. Unless otherwise specified, we will reserve the word 'generic' applied to points and types for the stronger notion of K-genericity.

We end this section by pointing out that the results above extend easily to *interpretable* sets in $\mathcal{M}$. In doing so we deduce, as promised in the previous chapters, that $\mathcal{M}$ cannot interpret any infinite field.

Corollary 3.1.15. *Corollary 3.1.11 and Corollary 3.1.13 hold for interpretable sets, in the following senses:*

1. *Let D be any non-empty $\mathcal{M}$-interpretable set. Then*

$$\dim D = \dim M \cdot rk\ D.$$

2. *If A is a set of parameters and a is a tuple in $\mathcal{M}^{eq}$, then*

$$\dim(a/A) \leq \dim M \cdot rk\ (a/A).$$

Proof. 1. As in Fact 2.3.17, after adding constants to the language, D is in interpretable finite-to-finite correspondence with an $\mathcal{M}$-definable set. Since both dim and rk are preserved under finite correspondences, the result follows immediately.

2. This follows from (1) using exactly the same argument as in Corollary 3.1.13.

□

Corollary 3.1.16. *$\mathcal{M}$ does not interpret any infinite field.*

Proof. Assume $\mathcal{M}$ interprets the infinite field L. Since L is infinite, rk $L \geq 1$. Thus by Corollary 3.1.15,

$$\dim L \geq \dim M \geq 2.$$

On the other hand, the interpretation of L in $\mathcal{M}$ also gives an interpretation of L in K. By [42] (Theorem 4.15), there is then a K-definable isomorphism $f : K \to L$. Since dim respects definable bijections we conclude

$$\dim L = \dim K = 1,$$

a contradiction. □

3.2 Pure Parts of Sets

The main approach we will use toward the Restricted Trichotomy Conjecture uses a loose generalization of pure dimensionality for curves to the higher dimensional setting. In the next two sections we will introduce this notion and study its basic properties. Recall that we interpret a *variety* over an algebraically closed field K to be the set of K-points of an irreducible quasiprojective variety over K.

Definition 3.2.1. Let V be a variety over an algebraically closed field K, and $D \subset V^k$ a K-definable set. Then the *pure part of* D, denoted D^P, is the union of the top dimensional components of the Zariski closure $\overline{D}$. We say that D is *pure* if $D \subset D^P$.

Remark 3.2.2. Note that Definition 3.2.1 depends on the embedding of D into V^k. Later on we will assume to work with structures $\mathcal{M}$ whose universe is a variety; we will then work with definable sets equipped with fixed embeddings into powers of that variety.

We note the following:

Lemma 3.2.3. *The map $D \mapsto D^P$ is K-definable in families. That is, for any definable family $\mathcal{D} = \{D_a\}_{a\in A}$ of definable sets, the family $\{(D_a)^P\}_{a\in A}$ is also definable over the same parameters.*

Proof. It is a general fact in the model theory of algebraically closed fields that irreducibility, components, dimension, and Zariski closure are definable in families (for example, see Chapter 10 of [28]). So we need only note that D^P is defined as the union of the irreducible components of a certain dimension that are contained in a Zariski closure. □

We will often use the following equivalent definitions of the pure part of D:

Lemma 3.2.4. *Let V be a variety over a saturated algebraically closed field K, $D \subset V^k$ a K-definable set, and A any countable set of parameters over which D is definable. Then D^P is precisely the Zariski closure of the set of generic elements of D over A.*

Proof. First assume $a \in D^P$. Then there is a top dimensional component $C \subset \overline{D}$ such that $a \in C$. Now let W be any Zariski open set containing a. Since C is irreducible, $W \cap C$ is dense in C, and therefore generic in C. So there is some $b \in W \cap C$ which is generic in C over A. Then

$$\dim(b/A) = \dim C = \dim D = \dim \overline{D},$$

so b is also generic in $\overline{D}$ over A. But $\dim(\overline{D} - D) < \dim \overline{D}$, so $b \notin \overline{D} - D$; thus $b \in D$. In particular, since $\dim(b/A) = \dim D$, b is generic in D over A.

Now assume $a \notin D^P$. Noting that D^P is Zariski closed in V^k, set $W = V^k - D^P$. Then W is a Zariski open neighborhood of a in V^k. Moreover, $W \cap D = D - D^P$ is an A-definable set of dimension less than $\dim D$ – so it does not contain any generics in D over A. Thus a is not in the closure of the generic points in D over A. □

A similar argument shows that if $K = \mathbb{C}$ we can replace the Zariski topology with the analytic topology:

Lemma 3.2.5. *Let V be a variety over $\mathbb{C}$, $D \subset V^k$ a $\mathbb{C}$-definable set, and A any countable set of parameters over which D is definable. Then D^P is precisely the closure (in the analytic topology on V) of the set of A-generic elements of D.*

Proof. By Lemma 3.2.4 D^P is the Zariski closure of the A-generic points of D. Since the analytic topology refines the Zariski topology, it follows that D^P contains the analytic closure of the A-generic points.

Now suppose $a \in D^P$, and let W be an analytic open set containing a. Let C be a top dimensional component of $\overline{D}$ containing a. Then C is a space of local dimension $2 \cdot \dim D$

over $\mathbb{R}$ at every point (for example, treating C as definable in the o-minimal real field – see chapter 6 for more details). Moreover $W \cap C$ is an open subspace, hence also has dimension $2 \cdot \dim D$. Now there are only countably many non-generic A-definable subsets of C; the closure of each of these intersects W in a space of dimension $< 2 \cdot \dim D$, which is therefore nowhere dense in $W \cap C$. By the Baire Category Theorem [1] these countably many spaces cannot cover all of $W \cap C$, so it follows that there is an element of $W \cap C$ not belonging to any of them – that is, an A-generic element $b \in C$ which belongs to W. As in the proof of Lemma 3.2.4, b is also an A-generic element of D. □

Finally, we note that the pure part 'commutes with fibers' in a precise sense.

Notation 3.2.6. Given a projection $\pi : A \to B$, a subset $S \subset A$, and an element $b \in B$, we use the notation S_b to denote $\{s \in S : \pi(s) = b\}$.

Lemma 3.2.7. *Let V be a variety over a saturated algebraically closed field K, let $D \subset V^j$ and $E \subset V^k$ be definable, and let $\pi : V^j \to V^k$ be a projection satisfying $\pi(D) \subset E$. Assume that for all generic $e \in E$ the fiber D_e has dimension $\dim D - \dim E$. Then the equation $(D_e)^P = (D^P)_e$ holds for all generic $e \in E$.*

Proof. Without loss of generality we assume V, D, and E are $\emptyset$-definable; otherwise we add parameters to the langauge until this is the case. Now fix a generic element $e \in E$.

First assume $x \in (D_e)^P$. Then $x \in \overline{D_e} \subset (V^j)_e$, so $\pi(x) = e$. It remains to show $x \in D^P$. To do this, let W be any Zariski open neighborhood of x in V^j. We seek a generic element of D which belongs to W. Now since $x \in (D_e)^P$, there is an e-generic $x' \in D_e$ which belongs to W. Since e is generic in E we have

$$\dim e = \dim E,$$

and by the choice of x' we have

$$\dim(x'/e) = \dim D_e = \dim D - \dim E.$$

By the additivity of dimension it follows that

$$\dim x' = \dim(e, x') = \dim D.$$

Thus x' is generic in D, as desired.

Now assume $x \in (D^P)_e$, and let W be a Zariski open neighborhood of x in V^j. We seek a generic element of D_e which belongs to W.

Since $x \in D^P$, there is a top dimensional component C of $\overline{D}$ which contains x. Since e is generic in E, there is a unique component G of $\overline{E}$ containing e; moreover G is of top dimension in $\overline{E}$, and e is generic in G.

Since $\pi(D) \subset E$, it follows that $\pi(\overline{D}) \subset \overline{E}$, and so $\pi(C) \subset \overline{E}$. Since C is irreducible, $\pi(C)$ is irreducibe, so in fact $\pi(C)$ is contained in a single irreducible component of $\overline{E}$. Since $e = \pi(x) \in \pi(C)$ and G is the only component of $\overline{E}$ containing e, we get $\pi(C) \subset G$.

So, since it is a projection, π restricts to a morphism of varieties, $f : C \to G$. Since im f contains the generic element $e \in E$, f is dominant. By [48] (page 323, Theorem 11.4.1), it follows that for all generic $e' \in G$ the fiber $C_{e'}$ is of pure dimension

$$\dim C - \dim G = \dim D - \dim E.$$

In particular, C_e is of pure dimension $\dim D - \dim E$.

Now since W is a Zariski open neighborhood of x, and x belongs to the pure dimensional set C_e, we can find a generic element $x' \in C_e$ which belongs to W. Thus

$$\dim(x'/e) = \dim C_e = \dim D - \dim E,$$

and so by additivity

$$\dim x' = \dim(e, x') = \dim D = \dim C.$$

In particular, x' is generic in C. Since C is almost contained in D, we get $x' \in D$, and thus $x' \in D_e$. Finally, recalling that

$$\dim(x'/e) = \dim D - \dim E,$$

and also

$$\dim D_e = \dim D - \dim E,$$

we conclude that x' is generic in D_e. This completes the proof of the lemma.

□

3.3 Almost Purity and Almost Closedness

Intuitively, we think of the pure part of D as an 'intended' definable set in $\mathcal{M}$ – that is, a set $\mathcal{M}$ 'would' define if it could see the background geometry of the field. The problem is $\mathcal{M}$ can't see this – so we have no way of preventing $\mathcal{M}$ from 'getting it wrong' – i.e. adding or removing a sporadic set of points from D^P to make its eventual definable set D. In this light, our goal in approaching the higher dimensional case of the conjecture is two-fold: (1) prove local modularity assuming the difference between D and D^P is never 'too much' in $\mathcal{M}$-definable curves, and (2) generate curves where D is close enough to D^P in the sense of (1). Our main technical notion is thus our notion of 'close enough' between D and D^P:

Convention 3.3.1. So that the following definition makes sense, we interpret the dimension of the empty set set as $-\infty$.

Definition 3.3.2. Let V be a variety over an algebraically closed field K, $D \subset V^k$ a non-empty K-definable set, and $\overline{D}$ its Zariski closure in V^k. Then D is *almost pure* if $\dim(D - D^P) \leq \dim D - 2$, and D is *almost closed* if $\dim(\overline{D} - D) \leq \dim D - 2$.

So, D is almost pure if it doesn't have 'added' pieces in codimension 1, and D is almost closed if it doesn't have 'subtracted' pieces in codimension 1. We end this section with the following characterization of almost purity and almost closedness:

Lemma 3.3.3. *Let V be a variety over a saturted algebraically closed field K, $D \subset V^k$ a non-empty K-definable set, and $\overline{D}$ its Zariski closure in V^k. Let A be any set of parameters over which all of this data is definable. Then:*

1. *D is almost pure if and only if for all $d \in V^k$ with $\dim(d/A) \geq \dim D - 1$, if $d \in D$ then $d \in D^P$.*
2. *D is almost closed if and only if for all $d \in V^k$ with $\dim(d/A) \geq \dim D - 1$, if $d \in D^P$ then $d \in D$.*

Proof. 1. First suppose D is almost pure. Let $d \in V^k$ with $\dim(d/A) \geq \dim D - 1$, and suppose $d \in D$. By almost purity, the set $D - D^P$ is A-definable of dimension at most $\dim D - 2$. Thus, since $\dim(d/A) \geq \dim D - 1$, $d \notin D - D^P$. Since $d \in D$, this implies $d \in D^P$.

Now suppose that for all $d \in D$ with $\dim(d/A) \geq \dim D - 1$, $d \in D^P$. Let d_0 be an A-generic element of $D - D^P$. Then by assumption we have $\dim(d_0/A) \leq \dim D - 2$. On the other hand, since d_0 is A-generic in $D - D^P$, we have

$$\dim(D - D^P) = \dim(d_0/A) \leq \dim D - 2.$$

So D is almost pure.

2. First suppose D is almost closed. Let $d \in V^k$ with $\dim(d/A) \geq \dim D - 1$, and suppose $d \in D^P$. Then $d \in \overline{D}$. Note that by almost closedness, $\overline{D} - D$ is an A-definable set of dimension at most $\dim D - 2$. Since $\dim(d/A) \geq \dim D - 1$ it follows that $d \notin \overline{D} - D$. So, since $d \in \overline{D}$, we get $d \in D$.

Now suppose that for all $d \in D^P$ with $\dim(d/A) \geq \dim D - 1$, $d \in D$. Let d_0 be an A-generic element of $\overline{D} - D$. We claim that $\dim(d_0/A) \leq \dim D - 2$. To see this, let C be an irreducible component of $\overline{D}$ containing d_0. We consider two cases:

- If $\dim C = \dim D$ then $d_0 \in D^P$, so since $d_0 \notin D$ our assumptions imply $\dim(d_0/A) \leq \dim D - 2$.
- If $\dim C \leq \dim D - 1$, note that C is almost contained in D (as a component of its closure). So since $d_0 \notin D$, d_0 is not generic in C. Hence

$$\dim(d_0/A) \leq \dim C - 1 \leq \dim D - 2.$$

Now, as in (1), we conclude by noting that

$$\dim(\overline{D} - D) = \dim(d_0/A) \leq \dim D - 2.$$

So D is almost closed.

□

3.4 Relation to Covering Maps

In this section we develop various properties of definable sets over $\mathbb{C}$, with applications toward those that are almost pure. Our main goal is to prove that almost pure sets definable in our strongly minimal $\mathcal{M}$ must very closely resemble covering spaces of powers of M in a precise sense. Along the way we will develop various technical facts that we use throughout the subsequent chapters.

Convention 3.4.1. Throughout this section, assume V is a variety over $\mathbb{C}$. Unless stated otherwise, we work with the analytic topology on V and its cartesian powers. Also unless otherwise stated, the word 'definable' refers to definability in $\mathbb{C}$. For definable sets $D \subset V^k$, we endow D with the subsapce topology inherited from the analytic topology on V^k.

We first define 'smoothness' for points in arbitrary constructible sets, generalizing the corresponding notion for varieties. We are intuitively trying to capture the points $d \in D$ where D looks locally like the analytic variety $\mathbb{C}^{\dim D}$.

Definition 3.4.2. Given a definable set D and an element $d \in D$, we say that d is *smooth in* D if there is an irreducible component C of $\overline{D}$ such that the following hold:

1. C is the unique irreducible component of $\overline{D}$ containing d.
2. $\dim C = \dim D$, and d is a smooth point of C.
3. There is an analytic neighborhood of d in C which is contained in D.

If every $d \in D$ is smooth in D, we say D is *smooth*.

Remark 3.4.3. It follows easily from Definition 3.4.2 that if $d \in D$ is smooth, d also belongs to D^P. Furthermore, as predicted above, there is a neighborhood of d which is isomorphic to $\mathbb{C}^{\dim D}$; this is the main application we make of smoothness.

The following notion of 'local surjectivity' will be quite useful throughout this thesis; among other things, it helps us to count fiber sizes of maps, which in many cases our reduct structure $\mathcal{M}$ is capable of seeing:

Definition 3.4.4. Let D and E be definable sets, and $f : D \to E$ a definable function. We say that f is *locally surjective near* $d \in D$ if there are neighborhood bases, $\{X_i\}_{i \in \mathbb{N}}$ of d in D and $\{Y_i\}_{i \in \mathbb{N}}$ of $f(d)$ in E, such that $Y_i \subset f(X_i)$ for each i.

Our immediate goal is to establish the following technical fact, which gives conditions under which local surjectivity can be inferred. We will then give various applications on counting fiber sizes; the eventual conclusion will be, as promised, that projections of $\mathcal{M}$-definable almost pure sets closely resemble covering maps.

Proposition 3.4.5. *Let D and E be definable sets of the same dimension n, and $f : D \to E$ a definable function. Let $d \in D$, and set $e = f(d)$. Assume the foilowing:*

1. *f is almost surjective, almost finite-to-one, and analytically continuous.*
2. *d belongs to D^P and has a compact neighborhood in D.*
3. *e is smooth in E.*
4. *d is isolated in the fiber $f^{-1}(e)$.*

Then f is locally surjective near d.

Proof. Fix any neighborhood bases $\{P_i\}$ and $\{Q_i\}$ of d and e. We will construct $\{X_i\}$ and $\{Y_i\}$ which refine $\{P_i\}$ and $\{Q_i\}$ and satisfy $Y_i \subset f(X_i)$.

Fix any i. Let $A_i \subset D$ be an open neighborhood of d isolating it in $f^{-1}(e)$. Shrinking A_i if necessary, we may assume the closure $\overline{A_i} \subset D$ is compact and contained in P_i. Thus the boundary $\partial(A_i)$ of A_i is compact as well.

Since f is continuous and $\partial(A_i)$ is compact, $f(\partial(A_i))$ is compact. By construction of A_i, $e \notin f(\partial(A_i))$. So there is a neighborhood B_i of e in E disjoint from $f(\partial(A_i))$. Shrinking B_i if necessary, we may assume that B_i is $\mathbb{R}$-definable and contained in Q_i. Since e is smooth in E, by shrinking further we may assume B_i is isomorphic to $\mathbb{C}^n$ as an analytic manifold.

Now as an almost surjective, almost finite-to-one definable function, f is 'almost everywhere' a finite-sheeted covering (for example, by applying 'generic smoothness on the source,' Theorem 25.3.1 of [48], to each top dimensional component of $\overline{D}$). That is, there are definable relatively open fully generic sets $G \subset D$ and $H \subset E$, such that the restriction of f to G is a finite-sheeted cover of H.

Claim 3.4.6. *$B_i \cap H$ is path connected.*

Proof. Since H is fully generic in E it follows that $E - H$ is a relatively closed subset of E of smaller dimension than $\dim E = \dim B_i$. In particular, $B_i \cap (E - H) = B_i - H$ is a proper closed analytic subvariety of B_i of smaller complex dimension. So $B_i \cap H$ is the complement in $B_i \cong \mathbb{C}^n$ of a closed $\mathbb{R}$-definable analytic subvariety of real codimension at least 2. This implies that $B_i \cap H$ is path connected (this is surely common knowledge, but could for example be deduced from Lemma 6.2.10, since all relevant sets are $\mathbb{R}$-definable). □

Claim 3.4.7. *$B_i \cap H \subset f(A_i)$.*

Proof. Since f is continuous and B_i is a neighborhood of e, it follows that $f^{-1}(B_i)$ is a neighborhood of d. Thus $A_i \cap f^{-1}(B_i)$ is also a neighborhood of d. Since $d \in D^P$, this implies $A_i \cap f^{-1}(B_i)$ contains a generic element $\hat{g} \in D$ over the parameters defining G and D. Since $G \subset D$ is fully generic, $\hat{g} \in G$. Let $\hat{h} = f(\hat{g}) \in H$. Then by assumption $\hat{h} \in B_i \cap H$.

Now fix any $h \in B_i \cap H$. Since $B_i \cap H$ is path connected, there is a path in $B_i \cap H$ from $\hat{h}$ to h. Using the covering f, we can lift this path to a path γ in G, which goes from $\hat{g}$ to some $g \in f^{-1}(h)$. So $f(g) = h$; it remains to show $g \in A_i$.

But recall that $\hat{g} \in A_i$. Since B_i is disjoint from $f(\partial(A_i))$, no point along γ lies on $\partial(A_i)$. In particular, γ must remain entirely inside A_i. It follows that $g \in A_i$, as desired. □

Claim 3.4.8. $B_i \subset f(\overline{A_i})$.

Proof. Since $\overline{A_i}$ is compact, $f(\overline{A_i})$ is closed. By the previous claim, $f(\overline{A_i})$ contains $B_i \cap H$, and thus contains $\overline{B_i \cap H} \supset B_i$. □

Now set $X_i = P_i$ and $Y_i = B_i$. By Claim 3.4.8 and the fact that $\overline{A}_i \subset P_i$, the proof of the proposition is complete. □

We now give some corollaries of Proposition 3.4.5 which will form the basis for many of our more geometric arguments later on. We begin with:

Corollary 3.4.9. *Let D and E be definable sets of the same dimension n, and $f : D \to E$ a definable function. Let $e \in E$. Assume:*

1. *f is almost surjective, almost finite-to-one, and analytically continuous.*

2. *There is a positive integer l such that $|f^{-1}(e')| = l$ holds for all generic $e' \in E$.*

3. *D is pure and locally compact.*

4. *e is smooth in E.*

Then $f^{-1}(e)$ is either infinite or of size at most l.

Proof. Since D is pure and locally compact, every point of D belongs to D^P and has a compact neighborhood in D.

Now assume $f^{-1}(e)$ is finite, and let $f^{-1}(e) = \{d_1, ..., d_m\}$. Note that each d_i is isolated from the others. So all assumptions of Proposition 3.4.5 are met, and we conclude that f is locally surjective near each d_j. For each $j \leq m$, fix neighborhoods X_j of $d_j \in D$ and Y_j of $e \in E$ such that $Y_j \subset f(X_j)$. Since our choices for these sets form neighborhood bases, we may assume the X_j are pairwise disjoint.

Let $Y = \bigcap_{j=1}^{m} Y_j$. Then Y is a neighborhood of e. Since e is smooth in E we have $e \in E^P$, and thus there is an element $e' \in Y$ which is generic in E over the parameters used to define D, E, and f. Since e' is generic in E, $f^{-1}(e')$ has size l.

On the other hand, since $e' \in Y$ it has at least m preimages – one in each of the X_i. Note that these m preimages are distinct since the X_i are pairwise disjoint. We conclude that $m \leq l$, which proves the corollary. □

The next corollary imroves on Corollary 3.4.9 by noting that the finite fiber case must happen 'outside codimension 2' – an observation that will be quite helpful when dealing with almost pure sets later on:

Corollary 3.4.10. *Let D, E, n, f, l, e be as in the hypothesis of Corollary 3.4.9. Let A be a set of parameters over which all this data is definable. If $\dim(e/A) \geq n-1$ then $f^{-1}(e)$ has size at most l.*

Proof. By Corollay 3.4.9, we need only show $f^{-1}(e)$ is not infinite. So, suppose it is, and let d be an A-generic element of $f^{-1}(e)$. So $\dim(d/A, e) \geq 1$. By additivity we conclude

$$\dim(d, e/A) \geq (n-1) + 1 = n.$$

But since $f(d) = e$ and f is definable over A, $\dim(e/A, d) = 0$. By additivity again, $\dim(d/A) \geq n$. Thus d is A-generic in D. But then we have an A-generic element of D which belongs to an infinite f fiber, which contradicts that f is almost finite-to-one. □

The next corollary gives conditions under which we always have a covering map; this forms the basis for the subsequent theorem:

Corollary 3.4.11. *Let D and E be definable sets of the same dimension n, and $f : D \to E$ a definable function. Assume:*

1. *f is almost surjective, almost finite-to-one, and analytically continuous.*
2. *There is a positive integer l such that $|f^{-1}(e)| = l$ holds for all generic $e \in E$.*
3. *D is pure and locally compact.*
4. *E is smooth.*

Let H be the interior of the set of all $e \in E$ such that $|f^{-1}(e)| = l$, and let $G = f^{-1}(H)$. Then the restriction of f to G is an l-covering of H.

Proof. Let $\hat{h} \in H$, and let $\hat{g}_1, ..., \hat{g}_l \in G$ be the preimages of $\hat{h}$. As in the previous two corollaries, f is locally surjective near each $\hat{g}_j$. For each j, let X_j be a neighborhood of $\hat{g}_j$, and Y_j a neighborhood of $\hat{h}$, such that $Y_j \subset f(X_j)$. Since G and H are open we may assume $X_j \subset G$ and $Y_j \subset H$. We may further assume the X_j are compact and pairwise disjoint.

Let Y be a neighborhood of E contained in $\bigcap_{j=1}^{m} Y_j$. Note that $Y \subset H$. Since E is smooth it is also locally compact – hence we can assume Y is compact as well.

Claim 3.4.12. *For each $y \in Y$, $f^{-1}(y)$ precisely consists of exactly one preimage in each X_j.*

Proof. Let $y \in Y$. By assumption y has at least one preimage in each X_j. On the other hand, since $y \in H$ it has exactly m total preimages. Since the X_j are disjoint, the claim follows immediately. □

For each i let $Z_j = X_j \cap f^{-1}(Y)$. Since X_j and Y are compact, Z_j is also compact. Let f_j be the restriction of Z_j to Y.

Claim 3.4.13. *For each j, f_j is a homeomorphism from Z_j to Y.*

Proof. By Claim 3.4.12 f_j is a bijection. On the other hand, f_j is continuous (as f is continuous), Z_j is compact, and Y is Hausdorff. Thus f_j is a continuous bijection from a compact space to a Hausdorff space, and so is a homeomorphism. □

By Claim 3.4.12 we have $f^{-1}(Y) = \bigcup_{j=1}^{m} Z_j$. So we have found a neighborhood of $\hat{h}$ whose preimage consists of l disjoint homeomorphisms. It follows that f is an l-covering on H. □

We are now finally ready to state the main result of this section, which we will apply repeatedly later on in studying the definable sets of higher dimensional reduct structures. In what follows, we will work with a strongly minimal structure whose universe is a variety M. Thus, in the style of the present chapter up to this point, we treat all definable subsets of powers of M as having inherited the analytic topology from that of M. Now our main result is:

Theorem 3.4.14. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure definable from $\mathbb{C}$. Assume the universe M is a smooth variety of dimension $n > 1$. Let $D \subset M^k$ be an $\mathcal{M}$-definable set of rank r, and let $\pi : M^k \to M^r$ be a projection which is almost surjective and almost finite-to-one on D. If D is almost pure, then there is a $\mathbb{C}$-definable open set $W \subset M^r$ such that:*

1. $\dim(M^r - W) \leq \dim(M^r) - 2$.
2. *The restriction of π to $\pi^{-1}(W)$ is a finite covering of W.*

Proof. We may assume all of this data is $\mathcal{M}$-definable over $\emptyset$. Given $y \in M^r$ and $S \subset M^k$, we use the notation S_y to denote $S \cap \pi^{-1}(y)$.

Since M is strongly minimal, M^r has Morley degree 1. Thus there is a positive integer l such that all generic fibers D_y have size l.

Lemma 3.4.15. *If $y \in M^r$ and $\dim y \geq nr - 1$, then $D_y = (D^P)_y$ and both sides have size exactly l.*

Proof. Fix $y \in M^r$ with $\dim y \geq nr - 1 = \dim(M^r) - 1$. We deduce a sequence of facts about y and its fibers in D and D^P.

Claim 3.4.16. $|D_y| = l$.

Proof. Recall that $\dim y \leq n \cdot \operatorname{rk} y$, by Corollary 3.1.13. In our case,

$$nr - 1 \leq n \cdot \operatorname{rk} y.$$

Since $n > 1$, this implies $\operatorname{rk} y \geq r$. Thus y is $\mathcal{M}$-generic in M^r, and so necessarily satifies the ($\mathcal{M}$-definable and generic) property that $|D_y| = l$. □

Claim 3.4.17. *$D_y \subset (D^P)_y$.*

Proof. Let $x \in D_y$. Then

$$\dim(x, y) = \dim y + \dim(x/y) \geq \dim y \geq nr - 1.$$

Since $\pi(x) = y$, $\dim(y/x) = 0$, so we conclude that

$$\dim x = \dim(x, y) \geq nr - 1.$$

In particular, since $\dim D = nr$, $x \in D$ is a point of codimension at most 1. Since D is almost pure, this implies $x \in D^P$. □

Claim 3.4.18. *$|(D^P)_y| \leq l$.*

Proof. The projection $\pi : D^P \to M^r$ satisfies all the hypotheses of Corollary 3.4.10. Indeed, π is continuous, almost surjective, and almost finite-to-one; D^P is locally compact (as a closed subset of M^k); and M^r is smooth (as a power of M). Since $\dim y \geq \dim(M^r - 1)$, Corollary 3.4.10 implies $|(D^P)_y| \leq l$. □

Now the previous three claims, when combined, immediately imply $D_y = (D^P)_y$ and $|(D^P)_y| = l$, so the lemma is proven. □

Now let Z be the set of all $y \in M^r$ such that $D_y = (D^P)_y$ and both sets have size exactly l. Let W be the interior of Z. Now the above lemma implies that any codimension 1 element of M^r belongs to Z. This implies that a generic element of $M^r - Z$ has codimension at least 2 in M^r – or equivalently,

$$\dim(M^r - Z) \leq \dim(M^r) - 2.$$

But $M^r - W$ is just the closure of $M^r - Z$, so we conclude immediately that

$$\dim(M^r - W) \leq \dim(M^r) - 2.$$

Finally, we show that the restriction of π to D is an l-covering of W. To see this, note that W is smooth (as an open subset of M^r), and $\pi^{-1}(W) \cap D^P$ is locally compact (as the intersection of an open set and a closed set in M^k). By Corollary 6.7.10, the restriction of π to D^P is an l-covering of W. By definition of W, this is the same as the restriction of π to D, and the proof is complete. □

Chapter 4

Finite Fundamental Groups and Unimodularity

In this chapter we study cases in which we are able to prove local moduarity through the stronger notion of unimodularity. The first appearing proof that certain strongly minimal structures are locally modular was Zilber's work on totally categorical structures (for example [52], Theorem 6.7); the class of unimodular structures was later introduced by Hrushovski [25] as a generalization of total categoricity in which a similar proof could be carried out. Unimodularity can be thought of as an abstract common generalization of 'locally-finite' and 'torsion-free': indeed, a strongly minimal pure group is unimodular if and only if it is a vector space over a finite field or $\mathbb{Q}$ – thus either of bounded exponent or with no torsion.

In the first section we will review the basic facts about unimodularity. After doing this, we will be able to give a particularly simple proof of the conjecture for simply connected universes (or more generally, universes of finite fundamental group) with 'enough' almost pure definable sets. The idea is to use the covering maps provided at the end of the previous chapter to directly prove unimodularity, thereby inheriting local modularity from Hrushovski's work.

The reduction to unimodularity is particularly simple; however, for interest we include a more direct proof of a similar statement in section 3. While this proof is unnecessary, we include it for the interesting strategy involved. At its heart, local modularity is a statement about 'almost disjoint' sets behaving like 'actually disjoint' sets. Namely, by the basic dimension theory of strong minimality, if $\{F_a\}_{a\in A}$ is a definable family of pairwise disjoint subsets of a fixed space, the size of the family, $\dim A$, is necessarily bounded. In this light, a non-locally modular structure is one in which allowing the F_a to be 'almost disjoint' (i.e. have finite intersection) lets $\dim A$ become arbitrarily large. Following this intuition, our second proof of local modularity seems intuitively like a 'correct' reason that families of curves have bounded ranks: starting with a large enough almost pure family, we proceed to produce a second family of pairwise disjoint $\mathbb{C}$-definable sets which is 'too big' to fit inside the fixed space M^2.

Finally, in the last section we answer a natural question which might be raised after reading the first three sections – namely, whether higher dimensional strongly minimal structures on simply connected universes are always unimodular. To this end, we give a short example of a strongly minimal structure on a higher dimensional simply connected variety which is not unimodular. Of course this structure is still locally modular – indeed of trivial pregeometry – so it does not violate the general conjecture we wish to prove.

4.1 Review of Unimodularity

We begin by reviewing the basic definition and facts related to unimodularity. This is harder than one might expect – indeed, there has been considerable confusion in the literature as to the correct definition. For a correct and thorough reference, the reader could consult [13]. In any case, the definition we give is certainly sufficient for our purposes.

Definition 4.1.1. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure, and let $f : D \to E$ a definable function which is almost surjective and almost fnite-to-one. Assume E is stationary, so there is some $l \in \mathbb{Z}^+$ such that $|f^{-1}(e)| = l$ holds for all generic $e \in E$. In this situation we define the integer l to be the *degree of* f.

That is, the degree of a function is its generically occuring fiber size, where this makes sense. Note that for any non-empty definable set D, by Proposition 3.1.9 we can find a function $f : D \to M^{\text{rk } D}$ satisfying the hypotheses of Definition 4.1.1.

We now define unimodularity in terms of function degrees, as follows:

Definition 4.1.2. The strongly minimal structure $\mathcal{M} = (M, ...)$ is *unimodular* if for all definable functions $f : D \to E$ satisfying the hypotheses of Definition 4.1.1, the degree of f only depends on D and E, and not the particular function f. That is, if E is stationary, $f : D \to E$ is an almost surjective, almost finite-to-one function of degree l_1, and $g : D \to E$ is such a funciton of degree l_2, then $l_1 = l_2$.

Example 4.1.3. Any vector space is unimodular: indeed, in a vector space there are no l-to-one maps between irreducible sets unless $l = 1$; it follows that the degree of any function satisfying the hypotheses of Definition 4.1.1 is $\text{DM}(D)$.

Example 4.1.4. The canonical example of a non-unimodular structure is a group of unbounded exponent which has torsion. For example, if K is an algebraically closed field and $p \neq \text{char } K$ is a prime, then the maps $x \mapsto x$ and $x \mapsto x^p$ have different degrees from K to K, showing that K (or indeed just the multiplicative group $K^\times$) is not unimodular.

The most important fact about unimodularity is the following:

Fact 4.1.5 (Hrushovski [25]). *If $\mathcal{M} = (M, ...)$ is a strongly minimal structure which is unimodular, then $\mathcal{M}$ is locally modular.*

A proof outline of Fact 4.1.5 is as follows: assuming $\mathcal{M}$ is unimodular, one gets a well-defined 'degree' (not necessarily equal to DM) for every non-empty definable set: one takes each power of M to have degree 1, and in general defines the degree of D to be the degree of any definable function $f : D \to M^{\text{rk } D}$ which is almost surjective and almost finite-to-one.

One then obtains an 'additivity' formula for degrees: if $D \subset A \times B$ is such that all fibers $D_b \subset A$ have the same Morley rank, and almost all D_b have the same degree d, then deg $D = d \cdot \deg B$. Using this additivity formula, one solves for the degree of the intersection of a generic pair of curves in a family; the conclusion is that this degree, which equals the cardinality of the intersection, only depends on the degree of a generic curve in the family. One then obtains a contradiction (assuming the original family was large enough) by finding a different family whose curves have the same individual degrees but less intersections: indeed, one can take the subfamily of all curves through a fixed generic point in the plane, then remove this common point from each of them to decrease the generic intersection number by 1.

Our goal in the next section will be to use Fact 4.1.5 to obtain a proof of the Restricted Trichotomy Conjecture in the case that (1) the universe has finite fundamental group, and (2) 'most' plane curves are almost pure. The method of proof is to study projections of plane curves $C \subset M^2$ to M. Assuming C is almost pure, we apply Theorem 3.4.14 to conclude that such a projection essentially consists of finitely many connected covers of M – one per top dimensional component of $\overline{C}$. We are then able to calculate bounds for the degree of such a projection in terms of the number of such components and the order of $\pi_1(M)$. Our conclusion is that the degrees of the two projections $C \to M$ cannot be 'too different' – in other words, we have a weak version of unimodularity. We then conclude that this weak form implies full unimodularity via a pair of technical facts.

For now, we discuss the first of these technical facts – namely Proposition 4.1.8 below – which justifies our restriction to projections of plane curves. An essentially equivalent statement appears in [13] (Corollary 3.7), but we include the proof below for the sake of completeness and geometric intuition.

Remark 4.1.6. Recall that if $\mathcal{M} = (M, ...)$ is strongly minimal, we define a *non-trivial* plane curve $C \subset M^2$ as one for which both projections $C \to M$ are finite-to-one – or equivalently, C is non-trivial if it does not almost contain $\{m\} \times M$ or $M \times \{m\}$ for any $m \in M$. Thus if C is non-trivial, both projections $C \to M$ satisfy the hypotheses of Definition 4.1.1.

By the preceding remark, the following definition is now justified:

Definition 4.1.7. Let $\mathcal{M} = (M, ...)$ be strongly minimal, and let $C \subset M^2$ be a non-trivial plane curve. Then C is *balanced* if the two projections $C \to M$ have the same degree.

The remainder of this section will be devoted to the proof of the following:

Proposition 4.1.8. *Let $\mathcal{M} = (M, ...)$ be strongly minimal, and assume every non-trivial definable plane curve $C \subset M^2$ is balanced. Then $\mathcal{M}$ is unimodular.*

Proof. We proceed with a series of lemmas. To start, assume $\mathcal{M} = (M, ...)$ is strongly minimal with all non-trivial plane curves balanced. Our goal is to show that any two almost surjective, almost finite-to-one projections of a fixed definable set to a power of M have the same degree. The first lemma below says that we are allowed to change the coordinates we project to by one at a time.

Lemma 4.1.9. *Let $D \subset M^k$ be a stationary definable set of rank $r \geq 1$, and let $\pi : D \to M^r$ and $\tau : D \to M^r$ be two projections which are almost finite-to-one and almost surjective. Let $B_\pi, B_\tau \subset \{1, ..., k\}$ be the r-element sets of coordinates that each of π and τ project to. If $|B_\pi \cap B_\tau| = r - 1$ then $\deg(\pi) = \deg(\tau)$.*

Proof. We assume D is definable over $\emptyset$. We may further assume that $B_\pi = \{1, ..., r\}$ and $B_\tau = \{1, ..., r-1, r+1\}$, and that π and τ project to these coordinates in increasing order. Let $(m_1, ..., m_k)$ be a generic element of D. Then our assumptions imply that $(m_1, ..., m_r)$ and $(m_1, ..., m_{r-1}, m_{r+1})$ are each generic in M^r. It follows that

$$\deg(\pi) = |\pi^{-1}(\pi(\overline{m}))|$$

and

$$\deg(\tau) = |\tau^{-1}(\tau(\overline{m}))|.$$

Let $C \subset M^2$ be the set of all (x, y) such that, for some element $d \in D$, the first $r+1$ coordinates of d are $(m_1, ..., m_{r+1}, x, y)$. Then the main point is:

Claim 4.1.10. *C is a non-trivial plane curve.*

Proof. First note that $(m_r, m_{r+1}) \in C$ by definition. Since $(m_1, ..., m_r) \in M^r$ is generic and C is definable over $(m_1, ..., m_{r-1})$, it follows that (m_r, m_{r+1}) has positive rank over the parameters defining C. Thus C is infinite.

Now the claim follows from the assertion that both projections $C \to M$ are everywhere finite-to-one. To see this, assume without loss of generality that for some $x_0 \in M$, we have $(x_0, y) \in C$ for infinitely many y. Then there is such an element $y_0 \in M$ which is generic over $(m_1, ..., m_{r-1})$. Let $d \in D$ be an element whose first $r+1$ coordinates are $(m_1, ..., m_{r-1}, x_0, y_0)$. So $(m_1, ..., m_{r-1}, y_0) \in M^r$ is generic and contained in the coordinates of d; this implies $\text{rk}\,(d) \geq r$, so d is generic in D. On the other hand, by choice of x_0 it follows that d belongs to an infinite fiber under π, which contradicts that π is almost finite-to-one on D. □

Now let f and g be the two projections $C \to M$, in order. By the claim we have $\deg f = \deg g$. To prove the lemma, we proceed to use f and g to compute the degrees of π and τ.

Let k be the number of elements of D whose first $r+1$ coordinates are $(m_1, ..., m_{r+1})$. By the stationarity of D, the number k is independent of the specific tuple $(m_1, ..., m_{r+1})$, as long as it extends to a generic element of D. In particular, the same k applies to $(m_1, ..., m_r, y)$ for any y with $(m_r, y) \in C$. It thus follows that the number of extensions of $(m_1, ..., m_r)$ to an element of D is $k \cdot \deg f$ – or, in other words, we have

$$\deg \pi = k \cdot \deg f.$$

By a symmetric argument, we conclude that

$$\deg \tau = k \cdot \deg g.$$

So, since $\deg f = \deg g$, the statement of the lemma follows. □

Next we show that the above lemma is sufficient for any desired change in which coordinates we project to:

Lemma 4.1.11. *Let $D \subset M^k$ be a stationary definable set of rank r, and let $\pi : D \to M^r$ and $\tau : D \to M^r$ be any two projections which are almost surjective and almost finite-to-one. Then $\deg \pi = \deg \tau$.*

Proof. If D is finite then

$$\deg \pi = \deg \tau = |D|,$$

so we are done. Thus we assume $r \geq 1$.

Let $(m_1, ..., m_k)$ be a generic element of D. Then the closure of $\{m_1, ..., m_k\}$ has dimension r in the pregeometry structure on D. Now by the stationarity of D, a choice of r coordinates giving rise to an almost surjective, almost finite-to-one projection $D \to M^r$ corresponds to a choice of r elements among $\{m_1, ..., m_k\}$ which form a basis for the closure of the whole tuple. We are claiming that the corresponding projections for any two such bases have the same degree.

By the previous lemma, we know that changing a basis by one element does not alter degree. So it suffices to show that we can turn any basis into any other by changing one element at a time. But this is a well-known corollary of the exchange axiom for pregeometries, which is used to show that dimension is well-defined. Namely, given two finite independent sets B_1 and B_2 with the same span, and any $b \in B_2$, it is possible to substitute b for one of the elements of B_1 to obtain a new independent set with the same span. Doing this repeatedly, one may replace each element of B_1 with an element of B_2, until one is left with just B_2 after finitely many steps. □

Now by the previous lemma, the following is well-defined: for any stationary definable set D of rank r, we define the *degree of* D, denoted $\deg D$, to be the degree of some (equivalently any) almost surjective, amost finite-to-one projection $\pi : D \to M^r$. Thus $\deg D$ is always a positive integer; the following formula is then immediate:

Lemma 4.1.12. *Let D and E be stationary definable sets of the same rank r, and let $f : D \to E$ be a definable function which is almost surjective and almost finite-to-one. Then $\deg D = \deg f \cdot \deg E$.*

Proof. Let $G \subset D \times E$ be the graph of f. Then G is in definable bijection with D, so G is also stationary of rank r. Let $\pi : D \to M^r$ and $\tau : E \to M^r$ be almost surjective, almost finite-to-one projections. Then composing π with the projection $G \to D$ gives an almost surjective, almost finite-to-one projection $G \to M^r$ of degree $\deg D$, and composing τ with the projection $G \to E$ gives an almost surjective, almost finite-to-one projection $G \to M^r$ of degree $\deg f \cdot \deg E$. By the previous lemma, these two degrees are equal. □

Finally, we extend the notion of degree to all non-empty definable sets, as follows: if D has Morley degree $m \geq 1$, we define

$$\deg D = \sum_{i=1}^{m} \deg D_i,$$

where $D_1, ..., D_m$ are the stationary components of D. Note that this is well-defined: indeed, the stationary components of D are well-defined up to almost equality, and the degree of a stationary set is clearly also well-defined up to almost equality.

In light of the above definition, we have:

Lemma 4.1.13. *If D and E are definable sets of the same rank, E is stationary, and $f : D \to E$ is a definable function which is almost surjective and almost finite-to-one, then*

$$\deg D = \deg f \cdot \deg E.$$

Proof. Write D as a union $D_1 \cup ... \cup D_m$ of stationary components of the same rank. For each i let f_i be the restriction to D_i. Since the D_i are almost disjoint, it is clear that each f_i is almost surjective and almost finite-to-one, and

$$\deg f = \sum_{i=1}^{m} \deg f_i.$$

Now by the previous lemma, for each i we have

$$\deg D_i = \deg f_i \cdot \deg E.$$

Adding this equality over all i then gives the desired statement. □

Now to finish the proof of Proposition 4.1.8, simply note that in the statement of the last lemma, the quantity $\deg f$ is dependent only on D and E – indeed we have

$$\deg f = \frac{\deg D}{\deg E}.$$

Noting that degrees of sets are never 0, the proof of the proposition is complete. □

4.2 Proof I: Reducing to Unimodularity

In this section we prove the Restricted Trichotomy Conjecture for smooth varieties of finite fundamental group, assuming that 'generic' curves are almost pure. We do this by showing that plane curves are balanced, reducing to Proposition 4.1.8.

The only technicality to deal with is the order of the fundamental group of the universe. The proof would be slightly easier for simply connected universes; for finite fundamental groups, however, we only immediately get that plane curves are balanced ‘up to a constant.’ In fact, this issue could be dealt with in a simple topological way, by invoking the Galois correspondence for covering spaces: what we really need is that a single space cannot i-cover and j-cover the universe unless $i = j$, and this follows since the degree of a cover is the index of the corresponding subgroup of the fundamental group, which only depends on the domain of the cover. However, we will instead deal with this issue in a more model theoretic way, as it increases the model theoretic generality of the technique.

To begin, we will need the following lemma, which is essentially immediate from the theory of canonical bases, and will show up repeatedly in later chapters.

Note 4.2.1. Recall that we define the *code* of a stationary definable set in a strongly minimal structure $\mathcal{M} = (M, ...)$ to be the canonical base of its generic type. If C and D are non-trivial plane curves then we define the *composition* $D \circ C$ as the set of $(x, z) \in M^2$ such that for some $y \in M$ we have $(x, y) \in C$ and $(y, z) \in D$.

Lemma 4.2.2. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure.*

1. *Let C be a plane curve in $\mathcal{M}$, definable over a set A. Let S be any strongly minimal component of C. Then the code of S is algebraic over A.*

2. *Let C and D be non-trivial irreducible plane curves in $\mathcal{M}$, and let E be any strongly minimal component of $D \circ C$. Let c, d, and e be the codes of C, D, and E, respectively. Then each of c, d, and e is algebraic over the other two.*

Proof. 1. This is immediate, since C has only finitely many strongly minimal components, and S is one of them.

2. We first point out that $D \circ C$ is necessarily a non-trivial plane curve, since C and D are. It follows that E is a non-trivial plane curve.

   Now the fact that e is algebraic over (c, d) follows from (1), since $D \circ C$ is definable (at least, up to finitely many points) over (c, d). We show below that c is algebraic over (d, e); the proof that d is algebraic over (c, e) is similar.

   Consider the set $E^{-1} \circ D$, where

   $$E^{-1} = \{(z, x) : (x, z) \in E\}.$$

   Note that E^{-1} is also a non-trivial plane curve, since E is; in turn, we conclude that $E^{-1} \circ D$ is a non-trivial plane curve.

   Moving to an elementary extension if necessary, we may assume $\mathcal{M}$ is saturated. So there is a generic element $x \in M$ over (c, d, e). Now since E is non-trivial, we can find an element $z \in M$ such that (x, z) is generic in E over (c, d, e); then $(x, z) \in D \circ C$, so in turn we can find an element $y \in M$ such that $(x, y) \in C$ and $(y, z) \in D$. By definition, it follows that $(x, y) \in E^{-1} \circ D$, as witnessed by z.

On the other hand, since x is generic in M over (c, d, e), and C is strongly minimal, it follows that (x, y) realizes the unique generic type of C over (c, d, e); in particular, this implies that all generic elements of C over (c, d, e) belong to $E^{-1} \circ D$. Thus C is one of the strongly minimal components of $E^{-1} \circ D$, and so it follows by (1) that c is algebraic over (d, e).

□

We now proceed with the following notions:

Definition 4.2.3. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure. We say that *all sufficiently generic plane curves in $\mathcal{M}$ have property P* if there is a natural number N such that the following hold:

1. There is a generically irreducible family of non-trivial plane curves of rank at least N.
2. For any such family $\mathcal{F} = \{F_a\}_{a \in A}$, and for any generic $a \in A$ (over the parameters that define $\mathcal{F}$), the curve F_a has property P.

Remark 4.2.4. Note that Definition 4.2.3 makes the most sense when P is a property which is invariant under automorphisms of $\mathcal{M}$ preserving the paramters defining $\mathcal{F}$ – in particular, for example, when P is definable over those parameters. In case $\mathcal{M}$ is interpreted in the algebraically closed field K, it also makes sense to consider P which are K-definable or K-autormorphism invariant – in which case the definition is equivalent to asserting that P holds whenever $a \in A$ is K-generic.

Definition 4.2.5. Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure.

1. Let $C \subset M^2$ be a non-trivial plane curve, and let $\pi_1, \pi_2 : C \to M$ be the projections to the (respectively) left and right copies of M. Then the *ratio of C* is the rational number
$$r(C) = \frac{\deg \pi_1}{\deg \pi_2}.$$
2. We say that $\mathcal{M}$ is *almost unimodular* if there is a real number $B > 0$ such that all sufficiently generic plane curves $C \subset M^2$ satisfy $r(C) < B$.

Now our main technical goal is the following:

Proposition 4.2.6. *Let $\mathcal{M} = (M, ...)$ be a strongly minimal structure. If $\mathcal{M}$ is almost unimodular then $\mathcal{M}$ is unimodular.*

Proof. Throughout this proof, if $C \subset M^2$ is a non-trivial plane curve, we use the notation $\deg_L(C)$ and $\deg_R(C)$ to denote the degrees of the projections to the (respectively) left and right copies of M.

Now we need two lemmas about ratios of curves:

Lemma 4.2.7. *For every non-trivial plane curve C, at least one of the strongly minimal components C_i of C satisfies $r(C_i) \geq r(C)$.*

Proof. Write $C = C_1 \cup ... \cup C_m$ as a union of strongly minimal components. It is clear that

$$\deg_L(C) = \sum_{i=1}^{m} \deg_L(C_i)$$

and

$$\deg_R(C) = \sum_{i=1}^{m} \deg_R(C_i).$$

Now assume $r(C_i) < r(C)$ for each i. We can rewrite this as

$$\deg_L(C_i) < r(C) \cdot \deg_R(C_i).$$

Adding over all i, we obtain

$$\deg_L(C) < r(C) \cdot \deg_R(C),$$

which can be rearranged to $r(C) < r(C)$, a contradiction. □

Lemma 4.2.8. *If C and D are non-trivial irreducible plane curves, then at least one of the strongly minimal components E_i of $E = D \circ C$ satisfies*

$$r(E_i) \geq r(C) \cdot r(D).$$

Proof. Write $E = D \circ C$ as a union $E_1 \cup ... \cup E_m$ of strongly minimal components. Define $S \subset M^3$ by

$$S = \{(x, y, z) : (x, y) \in C \text{ and } (y, z) \in D\}.$$

So $E = D \circ C$ is the projection of S to the first and third coordinates. For each i let

$$S_i = \{(x, y, z) \in S : (x, z) \in E_i\}.$$

Then each S_i is a rank 1 set equipped with projections to C, D, and E_i. It is clear that each of these projections is almost surjective and almost finite-to-one for all i. Since C, D, and E_i are stationary, each of these projections has a degree. Let j_i, k_i, and l_i be the degrees of the projections of S_i to C, D, and E_i, respectively.

Now it is clear from the definitions of all of these values that, for each i, we have

$$\deg_L(E_i) = \frac{j_i \cdot \deg_L(C)}{l_i}$$

and

$$\deg_R(E_i) = \frac{k_i \cdot \deg_R(D)}{l_i}.$$

Now assuming $r(E_i) < r(C) \cdot r(D)$ for each i, we obtain

$$\deg_L(E_i) < r(C) \cdot r(D) \cdot \deg_R(E_i).$$

Thus

$$\frac{j_i \cdot \deg_L(C)}{l_i} < \frac{r(C) \cdot r(D) \cdot k_i \cdot \deg_R(D)}{l_i},$$

or equivalently

$$j_i \cdot \deg_L(C) < r(C) \cdot r(D) \cdot k_i \cdot \deg_R(D).$$

After expanding $r(C)$ and $r(D)$, cancelling terms, and rearranging, we get

$$j_i \cdot \deg_R(C) < k_i \cdot \deg_L(D). \qquad (*)$$

On the other hand, from the definition it is clear that

$$\sum_{i=1}^{m} j_i = \deg_L(D)$$

and

$$\sum_{i=1}^{m} k_i = \deg_R(D).$$

So, adding $(*)$ over all i, we obtain

$$\deg_L(D) \cdot \deg_R(C) < \deg_R(C) \cdot \deg_L(D),$$

a contradiction. □

Now assume $\mathcal{M}$ is almost unimodular but not unimodular. Fix N and B according to the definition of almost unimodularity. Since $\mathcal{M}$ is not unimodular, there is a non-trivial plane curve C which is not balanced – or equivalently has ratio not equal to 1.

Note that

$$r(C^{-1}) = \frac{1}{r(C)},$$

where

$$C^{-1} = \{(y,x) : (x,y) \in C\}.$$

So, replacing C by C^{-1} if necessary, we can first assume $r(C) > 1$. By Lemma 4.2.7, we can further assume C is strongly minimal. Then, by alternately applying Lemma 4.2.8 and Lemma 4.2.7, we obtain a sequence of strongly minimal plane curves whose ratios grow at least exponentially. So, without loss of generality, we may also assume $r(C) > B$.

Finally, by almost unimodularity there is a generically irreducible family of non-trivial plane curves, say $\mathcal{F} = \{F_a\}_{a \in A}$, of rank at least N. By Fact 2.3.17, we may assume $\mathcal{F}$ is almost faithful. Restricting to a stationary component if necessary, we may further assume

A is stationary. It follows that all F_a, for generic $a \in A$, have the same ratio r_0. By replacing each F_a with its inverse if necessary, we may assume $r_0 \geq 1$. At this point, we may assume C and $\mathcal{F}$ are $\emptyset$-definable.

Fix $a \in A$ generic. Then by Lemma 4.2.8 and Lemma 4.2.7, there is a strongly minimal component E of $F_a \circ C$ with $r(E) > B$. So, to obtain a contradiction, it suffices to realize E as a generic curve in a family of rank at least N. But this follows readily from the theory of plane curves and canonical bases.

Namely, we need only note that a is algebraic over the code e of E: indeed, it follows by Lemma 4.2.2, and the fact that C is $\emptyset$-definable, that the code c of F_a is algebraic over e; but by almost faithfulness only finitely many $F_{a'}$ have the same code, so a is in turn algebraic over c.

Now since a is algebraic over e, it follows that

$$\text{rk } e \geq \text{rk } a \geq N.$$

Thus the corresponding family $\mathcal{G}$ of plane curves parametrized by the conjugates of e (or rather, an interpretable set whose generic elements are the conjugates of e) has the same rank as $\mathcal{F}$, and all generic curves of ratio greater than B. This contradicts almost unimodularity. □

Finally, we conclude this section with the proof of our main result for the present chapter. Recall that for an algebraically closed field K, we interpret a *variety over* K to be the set of K-points of an irredubiel quasiprojective variety over K.

Theorem 4.2.9. *Let M be a smooth variety of dimension $n > 1$ over an algebraically closed field K of characteristic 0. Let $\mathcal{M} = (M, ...)$ be a strongly minimal reduct of the full K-induced structure on $\mathcal{M}$. If all sufficiently generic plane curves in $\mathcal{M}$ are almost pure, then $\mathcal{M}$ is unimodular, and thus locally modular.*

Proof. We may assume $K = \mathbb{C}$. Let m be the order of the fundamental group of M. The main observation is the following:

Lemma 4.2.10. *If $C \subset M^2$ is any non-trivial almost pure plane curve, then $r(C) \leq m^2$.*

Proof. We apply Theorem 3.4.14. Namely, let $\pi : C \to M$ be a fixed projection to either copy of M; we obtain a $\mathbb{C}$-definable open set $W \subset M$ such that:

1. $\dim(M - W) \leq \dim M - 2$.
2. The restriction of π to $\pi^{-1}(W)$ is a finite covering of W.

It is clear that the number of sheets in this covering map is $\deg \pi$. Now let $C_1, ..., C_l$ be the top dimensional components of $\overline{C}$. Then we conclude:

Claim 4.2.11. *$\deg \pi \geq l$.*

Proof. Since the projection π is almost surjective and almost finite-to-one, its restriction to each C_i is also almost surjective and almost finite-to-one. Now let $w \in W$ be generic over the parameters defining C and W; note that w is also generic in M. Then w has at least one preimage in each C_i. On the other hand, since the C_i are almost disjoint, they cannot overlap above a generic point; thus the sets $\pi^{-1}(w) \cap C_i$ are pairwise disjoint and non-empty. It follows that

$$|\pi^{-1}(w)| = \deg \pi \geq l,$$

which proves the claim. □

Claim 4.2.12. *$\deg \pi \leq l \cdot m$.*

Proof. Since $W \subset M$ is fully generic and π is almost surjective and almost finite-to-one, it follows that $\pi^{-1}(W)$ is almost equal to C, and thus in turn is almost equal to $C_1 \cup ... \cup C_l$. So for each i, $C_i \cap \pi^{-1}(W)$ is a non-empty Zariski open subset of the variety C_i. It follows that $C_i \cap \pi^{-1}(W)$ is also irreducible; in particular, $C_i \cap \pi^{-1}(W)$ is path connected, and is thus contained in one of the connected components of the covering of W.

On the other hand, since $M - W$ has complex codimension at least 2 in M, it follows that W has the same fundamental group as M (see [15], Exposé X, Corollaire 3.3 and Exposé XII, Corollaire 5.2). In particular, every connected component of $\pi^{-1}(W)$ has at most m sheets.

By the above two paragraphs, it follows that each $C_i \cap \pi^{-1}(W)$ can account for at most m sheets. Thus the total number of sheets, $\deg \pi$, is at most $l \cdot m$, as desired. □

Now by the previous two claims we have

$$l \leq \deg \pi \leq l \cdot m.$$

In particular, since this applies to both projections $C \to M$, we obtain that $r(C)$ is the quotient of two numbers in the interval $[l, l \cdot m]$. So we need only note that the quotient of any two numbers in this interval is at most m^2. □

Finally, if all sufficiently generic curves are almost pure, then by the Lemma all sufficiently generic plane curves have ratio at most m^2. This implies $\mathcal{M}$ is almost unimodular. By Proposition 4.2.6, $\mathcal{M}$ is unimodular. □

4.3 Proof II: Producing a Pairwise Disjoint Family

In this section we give a direct, and drastically different, proof of a similar result to Theorem 4.2.9. There are subtle differences: first, rather than assuming *all* sufficiently generic plane curves are almost pure, we only assume the existence of one 'large enough' family with an almost purity property. On the other hand, rather than assuming only the almost purity of the generic curves in the family, we need a stronger geometric assumption: our family itself, as a single definable set, is almost pure. Finally, since we are not passing through unimodularity, our conclusion in this section is a direct verification of local modularity: in other words, we simply assume such a family exists and derive a contradiction.

As mentioned in the introduction of this chapter, the main reason for including this proof is the overarching strategy used. Even though the argument is quite technical, the guiding intuition is quite simple: the order of the fundamental group gives us bounds on the number of intersections of components of curves; using these bounds, we are able to modify a sufficiently large family of curves until it they become pairwise disjoint – an easy contradiction since these curves live inside a fixed space of bounded dimension. The strategy for performing this modification is a more technical version of the strategy used in showing that unimodular structures are locally modular: given a large family with bounded pairwise intersections, we take the subfamily sharing a fixed intersection, then remove the common intersection to leave pairwise disjoint sets.

A more detailed outline of the strategy is as follows: we start with a 'very large' almost faithful family $\mathcal{F} = \{F_a\}_{a \in A}$ such that the associated definable set $F \subset M^2 \times A$ is almost pure. We then modify the family $\mathcal{F}$ in a field-definable way to obtain a new family $\mathcal{G}$ of subsets of M^2 with a superior intersection property: roughly, 'outside codimension 2' any two sets in G have finite intersection, and the cardinality of this intersection is uniformly bounded in terms of the group $\pi_1(M)$ (regardless of the size of $\mathcal{G}$). Assuming $\mathcal{G}$ is large enough, we move to the subfamily $\mathcal{H}$ of sets passing through a fixed finite set of points, then remove these common points; the conclusion is that, 'outside codimension 2,' any two sets in $\mathcal{H}$ are disjoint. Finally, restricting to a further subfamily, we arrive at a family $\mathcal{I}$, indexed by a set D of dimension $\dim M + 1$, such that each set I_d in $\mathcal{I}$, with finitely many exceptions, can only intersect other sets $I_{d'}$ at non-generic points of I_d. So, removing a small piece of each I_d, we may assume the family $\mathcal{I}$ is pairwise disjoint. But $\mathcal{I}$ is a $\dim M + 1$-dimensional family of $\dim M$-dimensional sets inside a space of dimension $2 \cdot \dim M$ – thus we reach our contradiction.

We now proceed with the proof. We first need some preliminaries on almost faithfulness and almost purity:

Definition 4.3.1. Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a definable family of plane curves in any strongly minimal structure $\mathcal{M} = (M, ...)$. A point $x \in M^2$ is a *common point of* $\mathcal{F}$ if the set

$$_xF = \{a \in A : x \in F_a\}$$

is generic in A.

The most important fact about common points is that, assuming almost faithfulness, there are only ever finitely many: roughly, if there were infinitely many common points then there would be a 'common curve,' which would easily contradict almost faithfulness. Because of this, we can definably remove any common points, and thereby work with an equivalent family which does not have any. Formally, we have the following:

Lemma 4.3.2. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a definable, almost faithful family of plane curves in any strongly minimal structure $\mathcal{M} = (M, ...)$. If rk $\mathcal{F} \geq 1$, then there are only finitely many common points of $\mathcal{F}$.*

Proof. We may assume $\mathcal{F}$ is $\emptyset$-definable. Let D be the set of common points, so D is also $\emptyset$-definable. Let $d \in D$ be generic. Our goal is to show rk $d = 0$. To do this, let $(a, b) \in A^2$ be generic over d. Since d is common, such a and b can be found so that $d \in F_a \cap F_b$. Now since (a, b) is generic, we in particular have

$$\text{rk } (b/a) = \text{rk F} \geq 1.$$

Since $\mathcal{F}$ is almost faithful, this implies $F_a \cap F_b$ is finite, and thus rk $(d/a, b) = 0$. But d is independent from (a, b) by definition, so rk $d = 0$. ☐

We now assume the setup of the problem. Namely:

Convention 4.3.3. For the remainder of the section, assume M is a smooth variety of dimension $n > 1$ over an algebraically closed field K of characteristic zero, and $\mathcal{M} = (M, ...)$ is a strongly minimal reduct of the full K-induced structure on M. When discussing genericity of points, the words 'generic' and 'independent' are always intended in the sense of K. The corresponding notions in the sense of $\mathcal{M}$ will be expressed as '$\mathcal{M}$-generic' and '$\mathcal{M}$-independent.'

We proceed with two lemmas on almost pure sets.

Lemma 4.3.4. *Let D be an $\mathcal{M}$-definable set of rank $r \geq 0$ which is almost pure, and let $\pi : D \to M^s$ be a projection. Assume that each generic $y \in M^s$ has a fiber of rank $r - s$. Then for each K-generic $y \in M^s$, the fiber $\pi^{-1}(y)$ is almost pure.*

Proof. We may assume D is $\emptyset$-definable in $\mathcal{M}$. Let $y \in M^s$ be generic, and let $x \in \pi^{-1}(y)$ be a point of codimension at most 1. By Lemma 3.3.3, it suffices to show that $x \in (D_y)^P$.

Now by the choice of x we have

$$\dim(x/y) \geq n(r - s) - 1.$$

Since $\dim y = ns$, additivity implies

$$\dim(x, y) \geq nr - 1.$$

But y is algebraic over x, so

$$\dim x \geq nr - 1 = \dim D - 1.$$

Since D is almost pure, we get $x \in D^P$. Now recall that by Lemma 3.2.7 we have

$$(D_y)^P = (D^P)_y.$$

Since $x \in (D^P)_y$, it follows that $x \in (D_y)^P$, as desired. ☐

Lemma 4.3.5. *If D and E are $\mathcal{M}$-definable sets which are almost equal, and D is almost pure, then E is almost pure.*

Proof. Let $r = \text{rk } D = \text{rk } E$. We wish to show $\dim(E - E^P) \leq nr - 2$. Now since D and E are almost equal they have the same top dimensional components, so in particular we have $D^P = E^P$. So

$$E - E^P \subset (D - E^P) \cup (E - D) = (D - D^P) \cup (E - D).$$

Now since D is almost pure we have $\dim(D - D^P) \leq nr - 2$. Since D and E are $\mathcal{M}$-definable and almost equal it follows that $E - D$ is an $\mathcal{M}$-definable set of rank at most $r - 1$, and thus of dimension at most

$$n(r-1) = nr - n \leq nr - 2.$$

The statement of the lemma follows immediately. □

Now our main theorem is:

Theorem 4.3.6. *Assume M has finite fundamental group. If $\mathcal{F} = \{F_a\}_{a \in A}$ is an almost faithful family of non-trivial plane curves in $\mathcal{M}$, and the associated definable set $F \subset M^2 \times A$ is almost pure, then rk $\mathcal{F}$ is bounded by a positive integer which only depends on the order $m = |\pi_1(M)|$.*

Proof. We may assume $K = \mathbb{C}$. Take such a family of rank r. We can and do assume $r \geq 2$. Our first step is:

Lemma 4.3.7. *We may assume A is a generic subset of M^r.*

Proof. Fix an $\mathcal{M}$-definable function $f : A \to M^r$ which is almost surjective and almost finite-to-one. Indeed, as in Lemma 3.1.9, we can choose f to be given piecewise by projections on the stationary components of A. Let F' be the image of F in $M^2 \times M^r$. We may assume F, F', and f are all $\emptyset$-definable in $\mathcal{M}$.

Now we claim that F' is almost pure. To see this, note first that $\dim F' = \dim F$. Let $(x, f(y)) \in F'$ be of codimension at most 1. Then $(x, f(y))$ is $\mathcal{M}$-generic in F', since otherwise its dimension would drop by at least n from $\dim F'$. Now since f is almost surjective and almost finite-to-one, this implies (x, y) is $\mathcal{M}$-interalgebraic with $(x, f(y))$ – and thus (x, y) has codimension 1 in F, and is similarly $\mathcal{M}$-generic in F. Since F is almost pure, (x, y) belongs to $F \cap F^P$, and so belongs to $C \cap C^P$ for one of the stationary components C of F. Now take a sequence $\{(x_i, y_i)\}$ of generic points of C converging to (x, y); since f is a projection on C it is continuous on C, so the sequence $\{(x_i, f(y_i))\}$ converges to $(x, f(y))$. Since f is almost surjective and almost finite-to-one, each $\{(x_i, f(y_i)\}$ is generic in F', and we are done.

Now for generic $z \in M^r$, the fiber F'_z is a finite union of generic curves in $\mathcal{F}$, and so is itself a non-trivial plane curve. Let $\mathcal{G}$ be the family of curves of the form F'_z, for $z \in M^r$ such that $f^{-1}(z)$ is finite and F'_z is a plane curve. Let $G \subset M^2 \times M^r$ be the associated definable set corresponding to $\mathcal{G}$. Note that G is also almost pure by Lemma 4.3.5, since it is almost equal to F'.

Finally, we need only verify that the family $\mathcal{G}$ is almost faithful. To see this, let $z_1, z_2 \in M^r$ be such that $G_{z_1} \cap G_{z_2}$ is infinite. Since z_1 and z_2 have finite fibers under f, it follows that $F_{y_1} \cap F_{y_2}$ is infinite for some y_1, y_2 with $f(y_i) = z_i$ for each i. Since F is almost faithful, this implies y_1 and y_2 are interalgebraic. Then, since f has finite fibers at each z_i, it is clear that z_i and y_i are interalgebraic for each i; thus z_1 and z_2 are interalgebraic, as desired.

Now replacing $\mathcal{F}$ with $\mathcal{G}$ proves the lemma. □

So, assume $A \subset M^r$ is generic.

Convention 4.3.8. For the remainder of the proof, when the context is clear, we will often omit parameters when discussing dimension and genericity. This will make the presentation cleaner and easier to follow. For example, to say that an element is generic in a set implicitly means 'generic over the relevant parameters.' When the specific parameters needed are unclear, or when the parameters are especially important, we will include them.

We are now ready to make our first modification of $\mathcal{F}$: let $\mathcal{G}$ be the family of sets of the form $F_a \cap C$, where C is an irreducible component of $(F_a)^P$. Note that $\mathcal{G}$ is indeed a $\mathbb{C}$-definable nr-dimensional family of n-dimensional subsets of M^2 – though it is far from being $\mathcal{M}$-definable. We write $\mathcal{G}$ as the $\mathbb{C}$-definable set $G \subset M^2 \times B$, for some $\mathbb{C}$-definable set B. Note that $\dim B = \dim A$. By shrinking B if necessary, we assume B is itself a variety, thus in particular irreducible. Note that each set G_b is stationary as a $\mathbb{C}$-definable set, since it is fully generic in an irreducible component of the corresponding $(F_a)^P$. At this point, we may assume $\mathcal{G}$ is $\emptyset$-definable in the field structure.

Recall that we defined m to be the order of $\pi_1(M)$. Now our main goal is prove:

Lemma 4.3.9. *If $b, c \in B$ are separately generic and the pair (b, c) has codimension at most 1 in B^2, then there are at most m^4 points $x \in M^2$ which are generic in both G_b and G_c.*

Proof. Let $a_b, a_c \in A$, C_b a component of $(F_{a_b})^P$, and C_c a component of $(F_{a_c})^P$, such that $G_b = F_{a_b} \cap C_b$ and $G_c = F_{a_c} \cap C_c$. Note that a_b and b are interalgebraic, as are a_c and c. Thus a_b and a_c are generic in A, and (a_b, a_c) has codimension at most 1 in A^2. It follows that we can find $r-1$ of the M-coordinates of a_c which are generic over a_b. That is, let

$$a_b = (b_1, ..., b_r) \in M^r$$

and

$$a_c = (c_1, ..., c_r) \in M^r;$$

then we may assume that $(a_b, c_1, ..., c_{r-1})$ is generic in M^{2r-1} and

$$\dim(c_r/a_b, c_1, ..., c_{r-1}) \geq n - 1.$$

Let $\mathcal{T}$ be the subfamily of $\mathcal{F}$ consisting of curves indexed by $a \in A$ with first $r-1$ coordinates equal to $(c_1, ..., c_{r-1})$. Then $\mathcal{T}$ is naturally parametrized by M, seen as the last coordinate of A. Let $T \subset M^2 \times M$ be the corresponding $\mathcal{M}$-definable set. By Lemma 4.3.2, $\mathcal{T}$ has only finitely many common points. Furthermore, the set of such points is $\mathcal{M}$-definable over $c_1, ..., c_{r-1}$. Let $\mathcal{T}'$ be the subfamily obtained by deleting each common point of $\mathcal{T}$ from each curve in $\mathcal{T}$, and let T' be the associated definable set.

Now let

$$R = \{(x, y, z) \in M^3 : (x, y) \in F_{a_b} \cap T'_z\}.$$

That is, R parametrizes the intersections of curves in T' with F_{a_b}. Since each generic such curve has a non-empty finite intersection with F_{a_b}, it readily follows that rk $R = 1$.

We proceed to show a number of properties of these sets:

Claim 4.3.10. *Each of the three projections $R \to M$ is almost surjective and everywhere finite-to-one.*

Proof. Since rk $R = 1$, this follows from the assertion that each projection is everywhere finite-to-one. Now for the projections to either of the leftmost two copies of M, this second assertion follows from the fact that F_{a_b} is non-trivial and $\mathcal{T}'$ has no common points – i.e. each point belongs to at most finitely many curves in $\mathcal{T}'$.

For the projection to the rightmost copy of M, the desired statement follows from the fact that all curves in $\mathcal{T}'$ have finite intersection with F_{a_b}: indeed, if not then some curve index in $\mathcal{T}'$ would (by the almost faithfulness of $\mathcal{F}$) be algebraic over a_b – thus c_1 would not be generic over (a_b), a contradiction. □

Claim 4.3.11. *F_{a_b}, T, and T' are almost pure.*

Proof. We cite Lemma 4.3.4. Namely, F_{a_b} is the fiber above the generic point a_b in the projection $F \to M^r$; and T is the fiber above the generic point $(c_1, ..., c_{r-1})$ in the projection $F \to M^{r-1}$. This shows F_{a_b} and T are almost pure. To see that T' is almost pure, we note that T and T' are almost equal, and cite Lemma 4.3.5. □

With a bit more work, we are also able to show:

Claim 4.3.12. *R is almost pure.*

Proof. Let $w = (x, y, z) \in R$ be a point of codimension at most 1. So w is $\mathcal{M}$-generic in R. By Claim 4.3.10, each coordinate of w is $\mathcal{M}$-generic, and any two are interalgebraic. It follows, then, that each coordinate has codimension at most 1 in M. In particular, (x, y) is a point of codimension at most 1 in F_{a_b}. Since F_{a_b} is almost pure, we can find a sequence of generics $\{(x_i, y_i)\} \subset F_{a_b}$ converging to (x, y). Without loss of generality, we may assume each (x_i, y_i) is generic in F_{a_b} over $(a_b, c_1, ..., c_{r-1})$.

Now we have

$$\dim(a_b, c_1, ..., c_{r-1}, w) \geq n(2r-1) + n - 1 = 2nr - 1.$$

Since (x, y) is not algebraic over $(a_b, c_1, ..., c_{r-1})$, it is in particular not algebraic over $\emptyset$, and so is not a common point of $\mathcal{F}$. Thus

$$\text{rk } (a_b/c_1, ..., c_{r-1}, w) \leq r - 1,$$

and so

$$\dim(a_b/c_1, ..., c_{r-1}, w) \leq n(r - 1).$$

By additivity,

$$\dim(c_1, ..., c_{r-1}, w) \geq n(r + 1) - 1.$$

In particular, this implies

$$\dim(w/c_1, ..., c_{r-1}) \geq 2n - 1.$$

Thus w is a point of codimension at most 1 in T, and therefore also in T'. On the other hand, since T' is a rank 1 almost faithful family with no common points, it is clear that the projection of T' to the plane is almost surjective and almost finite-to-one. We conclude:

Subclaim 4.3.13. *The projection $T' \to M^2$ to the first two coordinates is locally surjective near w.*

Proof. We verify the hypotheses of Proposition 3.4.5. The only non-obvious points to verify are:

- $w \in (T')^P$. Indeed, this follows since w has codimension at most 1 in T' and T' is almost pure.
- w is isolated in its fiber. Indeed, since w has codimension at most 1 in T', it is in fact $\mathcal{M}$-generic in T'; so, since the projection $T' \to M^2$ is almost finite-to-one, w belongs to a finite fiber.
- w has a compact neighborhood in T'. Indeed, if not then w belongs to the frontier of the frontier of T', which shows that w has codimension at least 2 in T', a contradiction.

□

By the subclaim, and after restricting to a subsequence, we can find z_i converging to z such that each $(x_i, y_i, z_i) \in T'$. Thus each $(x_i, y_i, z_i) \in R$, and by the assumption on the genericity of (x_i, y_i) in F_{a_b} it follows that each (x_i, y_i, z_i) is generic in R. This shows that R is almost pure. □

We will use the almost pureness of these sets to count fiber sizes. In a manner similar to Theorem 4.2.9, we note the following:

Claim 4.3.14. *If D is any $\mathcal{M}$-definable almost pure set of rank r, and $\pi : D \to M^r$ is an almost surjective, almost finite-to-one projection, then the restriction of D to any component of D^P has degree at most m^r.*

Proof. In an identical manner to the argument given in Theorem 4.2.9, we conclude that the projection $D \to M^r$ restricts to a covering map after removing a closed subset of complex codimension at least 2 from M^r. The removal of such a subset does not change the fundamental group ([15], Exposé X, Corollaire 3.3 and Exposé XII, Corollaire 5.2), so the resulting target of the projection has fundamental group $(\pi^1(M))^r$. We also conclude identically that each component of D^P, restricted to the domain of the cover, is contained in a single connected component – and therefore contributes at most m^r sheets. □

Now toward the proof of Lemma 4.3.9, let (x_0, y_0) be such that (x_0, y_0, b) and (x_0, y_0, c) are both generic in G. Since by assumption $\dim(c_r/a_b, c_1, ..., c_{r-1}) \geq n - 1$, it follows that (x_0, y_0, c_r) is a point of codimension at most 1 in R, and so belongs to R^P. Let W be a component of R^P containing (x_0, y_0, c_r). By the irreducibility of W, there are components Y of $(F_{a_b})^P$ and Z of $(T')^P$, such that whenever $(x, y, z) \in W$ we have $(x, y) \in Y$ and $(x, y, z) \in Z$.

On the other hand, since (x_0, y_0, b) and (x_0, y_0, c) are generic in G, it follows that (x_0, y_0) is generic in F_{a_b} and (x_0, y_0, c_r) is generic in T'. Thus there are unique components of $(F_{a_b})^P$ and $(T')^P$ containing these points. Since $(x_0, y_0, c_r) \in W$, it follows that Y is the unique component of $(F_{a_b})^P$ containing (x_0, y_0) and Z is the unique component of $(T')^P$ containing (x_0, y_0, c_r).

Now the main point is that Y and Z are determined by b and c: indeed, b itself picks out a component of $(F_{a_b})^P$, which must be Y; and c picks out a component of $(T'_{c_r})^P$ – then, since c is generic, this second component is contained in a unique component of $(T')^P$, which must be Z.

To summarize, (x_0, y_0, c_r) is contained in a component W of R^P which is itself contained in the set

$$S = \{(x, y, z) : (x, y) \in Y, (x, y, z) \in Z\}.$$

But we can bound the number of components of R^P which are contained in S: indeed, since by Claim 4.3.14 the projections $Y \to M$ and $Z \to M^2$ have degrees at most m and m^2, respectively, it follows that the projection $S \to M$ (to the leftmost copy of M) has degree at most m^3. It then follows that at most m^3 components $W' \subset R^P$ are contained in S.

Finally, note that (x_0, y_0) belongs to the fiber W_{c_r}. Since M is smooth and c_r has codimension at most 1 over the parameters defining W, Corollary 3.4.10 implies that the fiber W_{c_r} has at most the generic finite size. But by Claim 4.3.14 this finite size is at most m.

To conclude, there are at most m^3 components which could contain the points (x_0, y_0) we are interested in; and each such component can contain at most m of them. So there are at most m^4 such points in total. □

The proof of Lemma 4.3.9 is now complete. Using Lemma 4.3.9 and the Compactness Theorem, we can restrict G to a generic subset G', consisting of a generic $B' \subset B$, and a generic $G'_b \subset G_b$ for each $b \in B'$, so that for each $b \in B'$, the set of $c \in B'$ with $|G'_b \cap G'_c| > m^4$ has codimension at least 2 in B'. Without loss of generality, we assume $G = G'$ for the remainder of the proof.

For ease of notation, we use the letter k to denote the quantity m^4. We can and do assume $r > k$.

Let $b \in B$ be generic, and let $x_1, ..., x_k$ be generic independent elements of G_b. Our goal will be, roughly, to take the subfamily of curves through $x_1, ..., x_k$, and then remove $x_1, ..., x_k$ to leave a pairwise disjoint family. We first show:

Claim 4.3.15. *$x_1, ..., x_k$ are $\mathcal{M}$-generic and $\mathcal{M}$-independent in M^2.*

Proof. Let a be such that $G_b \subset F_a$. Then a and b are interalgebraic, and $(x_1, ..., x_k)$ is an $\mathcal{M}$-generic independent tuple in F_a. Let c be an independent realization of the strong type of a over $(x_1, ..., x_k)$. Note that $\text{rk}\ (a, x_1, ..., x_k) = r + k$ and $\text{rk}\ (x_1, ..., x_k) \leq 2k$. Thus

$$\text{rk}\ (a/x_1, ..., x_k) \geq r - k > 0.$$

In particular, a and c are not interalgebraic; since $\mathcal{F}$ is almost faithful, $F_a \cap F_c$ is finite, and so $\text{rk}\ (x_1, ..., x_k/a, c) = 0$. Thus $\text{rk}\ (a, c, x_1, ..., x_k) \leq 2r$. Since $\text{rk}\ (a, x_1, ..., x_k) = r + k$, this means $\text{rk}\ (c/a, x_1, ..., x_k) \leq r - k$. By the choice of c, $\text{rk}\ (c/x_1, ..., x_k) \leq r - k$, and so $\text{rk}\ (a/x_1, ..., x_k) \leq r - k$. Since $\text{rk}\ (a, x_1, ..., x_k) = r + k$, this implies $\text{rk}\ (x_1, ..., x_k) \geq 2k$. Equivalently, $(x_1, ..., x_k)$ is a generic independent tuple of elements of M^2. □

We immediately conclude:

Corollary 4.3.16. *$x_1, ..., x_k$ are generic and independent in M^2.*

Proof. Let a be as in the proof of the previous claim. Then the previous claim implies $\text{rk}\ (a/x_1, ..., x_k) = r - k$, and so $\dim(a/x_1, ..., x_k) \leq n(r-k)$. But $\dim(a, x_1, ..., x_k) = n(r+k)$, thus $\dim(x_1, ..., x_k) \geq n(2k)$, as desired. □

Now let $\mathcal{H} = \{H_c\}_{c \in C}$ be the subfamily of $\mathcal{G}$ consisting of sets which contain $(x_1, ..., x_k)$. Note that $b \in C$. Restricting C if necessary, we assume

$$\dim C = \dim(b/x_1, ..., x_k) = n(r - k).$$

Then we conclude:

Claim 4.3.17. *If $c, c' \in C$ are generic and $(c, c') \in C^2$ is a point of codimension at most 1, then $H_c \cap H_{c'} = \{x_1, ..., x_k\}$.*

Proof. Let a and a' be such that $H_c \subset F_a$ and $H_{c'} \subset F_{a'}$. So a and a' are interalgebraic, as are c and c'. Now since $c' \in C$ has codimension at most 1 over c in C, and $\dim C = n(r-k) > 1$, it follows that c and c' are not interalgebraic; thus a and a' are not interalgebraic. By almost faithfulness, this implies $F_a \cap F_{a'}$ is finite, and so $H_c \cap H_{c'}$ is finite as well. Thus $\dim(x_1, ..., x_k/c, c') = 0$. Now since $(c, c') \in C^2$ is a point of codimension at most 1, we have

$$\dim(x_1, ..., x_k, c, c') \geq n(2k) + 2n(r - k) - 1 = 2nr - 1.$$

So $\dim(c, c') \geq 2nr - 1$. In particular, (c, c') is a point of codimension at most 1 in B^2 as well. But, by choice of the family $\mathcal{G}$, this implies that

$$|G_c \cap G_{c'}| = |H_c \cap H_{c'}| \leq k.$$

Since this intersection contains $\{x_1, ..., x_k\}$, it must be equal to $\{x_1, ..., x_k\}$. □

We now replace $\mathcal{H}$ by a family $\mathcal{H}'$ as follows: using compactness, we restrict to a generic $C' \subset C$ such that for any $c \in C'$, the set of $c' \in C'$ with $H_c \cap H_{c'} \neq \{x_1, ..., x_k\}$ has codimension at least 2 in C'. Then for each $c \in C'$, let H'_c be the n-dimensional set formed by removing $x_1, ..., x_k$ from H_c, i.e.

$$H'_c = H_c - \{x_1, ..., x_k\}.$$

It follows readily that, for each $c \in C'$, the set of $c' \in C'$ such that H'_c intersects $H'_{c'}$ is of codimension at least 2 in C'. From this point forward, we assume we have replaced $\mathcal{H}$ with $\mathcal{H}'$ and C with C'.

We have one more modification to perform: in the ambient projective space containing M^r, let $\mathcal{P}$ be a finite set of generic, independent hyperplanes of the appropriate cardinality so that $D = C \cap (\bigcap \mathcal{P})$ has dimension $n + 1$. This is possible assuming $r \geq k + 2$. Let $\mathcal{I}$ be the subfamily of $\mathcal{H}$ indexed by D. We conclude:

Claim 4.3.18. *If $d \in D$ is generic over $\mathcal{P}$, then the union of all intersections $I_d \cap I_{d'}$ which are non-generic in I_d, across all $d' \in D$, is itself non-generic in I_d.*

Proof. First consider those intersections which are infinite. Such intersections come from infinite intersections of curves in the original family $\mathcal{F}$, and therefore by almost faithfulness only happen at interalgebraic parameters. So there are only finitely many such intersections, each of which is non-generic in I_d by definition; and the union of finitely many non-generic subsets of a fixed set is again non-generic.

Now consider those intersections which are finite. Recall that the set Q of $d' \in C$ which intersect $I_{d'}$ has codimension at least 2 in C. Now since d is generic in D it is independent from the collection $\mathcal{P}$; thus each $P \in \mathcal{P}$ decreases $\dim Q$ by at least 1. It follows that the set of $d' \in D$ which intersect d is of codimension at least 2 in D, and therefore dimension at most $n - 1$. So we have an $n - 1$-dimensional family of finite intersections; the union of such a family is of dimension $n - 1$, and is thus non-generic in the n-dimensional set I_d. □

Finally, we replace D by a generic subset D' so that, for each $d \in D'$, the union of all non-generic intersections of other sets in $\mathcal{I}$ with I_d is itself non-generic in I_d. We then define a family $\mathcal{I}'$, indexed by D', by setting each I'_d to be the complement in I_d of the union of all such intersections. Thus each I'_d is still of dimension n. We assume we have replaced $\mathcal{I}$ and D with $\mathcal{I}'$ and D'.

Now it follows immediately that, for any $d, d' \in D$, the intersection $I_d \cap I_{d'}$ is either empty or generic in I_d. Since the sets in $\mathcal{I}$ are stationary over $\mathbb{C}$, this equivalently means that I_d and $I_{d'}$ are either disjoint or almost equal. Note that by almost faithfulness the almost equality classes are finite: so, modding by almost equality, the dimension of the family is unchanged. That is, replace each almost equality class in $\mathcal{I}$ with the intersection of the finitely many representatives of the class. We obtain an $n+1$-dimensional family of n-dimensional sets which are pairwise disjoint. But these sets are all contained in the $2n$-dimensional set M^2, which is a contradiction.

To recap, we obtained a contradiction assuming $r \geq k+2$. Thus, we obtain the bound

$$r \leq k+1 = m^4 + 1,$$

which proves the theorem.

□

4.4 A Non-Unimodular Example with Trivial Geometry

As stated in the introduction to this thesis, our envisioned strategy for the Restricted Trichotomy Conjecture in higher dimensions consists of two main parts: (1) show that we can generate ‘enough’ almost pure definable sets, and (2) prove bounds on the complexity of such sets. The contents of the present chapter, to this point, accomplish (2) in a strong way for universes of finite fundamental group – namely showing that such structures are unimodular. So one might naturally conjecture that, if we can accomplish (1) as well, it would follow that all higher dimensional strongly minimal reducts on universes of finite fundamental group are unimodular. In this section we show that this natural conjecture fails, by constructing a non-unimodular reduct structure on the higher dimensional simply connected variety $\mathbb{C}^2$. Interpreted in the lens of the goals (1) and (2) above, what our example really shows is that (1) cannot be verified without having a reasonably complex structure to begin with. That is, our strategies in approaching (1) in subsequent sections will involve studying the interactions of failures of almost purity with a plethora of other definable sets; however we only really have access to a plethora of definable sets under the assumption of non-local modularity. Indeed, the example we construct below is at the opposite end of the spectrum, of trivial type. Roughly, we simply name a single non-balanced (and thus necessarily non-almost pure) set $C \subset (\mathbb{C}^2)^2$, and show that the generated class of definable sets gives rise to a strongly minimal structure on $\mathbb{C}^2$. Of course we expect that naming any ‘rich’ collection of such sets, as opposed to just one, would lead to a violation of strong minimality.

We now give the example. The structure we are really trying to define is $(\mathbb{C}, x \mapsto x^2)$ – the complex numbers endowed with the square function. This structure has a simply connected universe, is strongly minimal (as a reduct of the complex field), and is clearly not unimodular, as witnessed by the square function. Thus it would suffice for our purposes if $\mathbb{C}$ was not of dimension 1; so, our goal will be to axiomatize the theory of $(\mathbb{C}, x \mapsto x^2)$, and then find a constructible self-map on $\mathbb{C}^2$ satisfying the same axioms.

We start with our language. Let $P \subset \mathbb{C}$ be the set consisting of 0, and every root of unity. Note that P is closed under squaring, and under taking square roots. Then we define the following:

Definition 4.4.1. Let $\mathcal{L}$ be the language consisting of a unary function symbol f, and a constant symbol $\overline{c}$ for each $c \in P$.

Remark 4.4.2. The added constant symbols are not really necessary, but they allow for a much smoother analysis of the theory we will study.

Next we define our theory:

Definition 4.4.3. Let T be the $\mathcal{L}$-theory given by the following axiom schema:

A1 : For each $c, d \in P$ with $c \neq d$, the axiom $\overline{c} \neq \overline{d}$.

A2 : For each $c \in P$, the axiom $f(\overline{c}) = \overline{c^2}$.

A3 : The sentence stating that $\overline{0}$ has exactly one preimage under f, and every other element has exactly two preimages under f.

A4 : For all non-negative integers $j < i$, the axiom asserting that if $f^i(x) = f^j(x)$, then x is equal to either $\overline{0}$ or one of $\overline{c_1}, ..., \overline{c_k}$, where $c_1, ..., c_k$ are the distinct roots of unity in P whose orders divide $2^i - 2^j$. Note that we interpret f^n as the n-fold composition of f, and f^0 as the identity.

We note that if $M \subset \mathbb{C}$ is any set which contains P and is closed under squaring, then M is naturally identified as an $\mathcal{L}$-structure $\mathcal{M}$, by interpreting each $\overline{c}$ as c and f as the square function. If in addition M is closed under taking square roots, we have the following:

Lemma 4.4.4. *Let $M \subset \mathbb{C}$ be any set which contains P and is closed under squaring and taking square roots. Let $\mathcal{M}$ be the associated $\mathcal{L}$-structure described above. Then $\mathcal{M} \models T$.*

Proof. The axiom schema A1 and A2 are obvious, and A3 follows since M is closed under taking square roots. To verify A4, let $x \in M$ and $j < i$ with $f^i(x) = f^j(x)$. Equivalently, this says

$$x^{2^i} = x^{2^j}.$$

Now if $x \neq 0$, then dividing by x^{2^j} gives

$$x^{2^i - 2^j} = 1.$$

Thus x is one of the roots of unity whose order divides $2^i - 2^j$, as desired. □

In particular, P and $\mathbb{C}$ are models of T. By A1 and A2, the atomic diagram of the constants is determined by T – thus P canonically embeds into every model of T. Note that by A2 and A3, the copy of P in any model of T is closed under taking images and preimages of f.

Our aim is to show that T is complete, i.e. any two models of T agree on all first order sentences. We do this by showing that T has quantifier elimination. Recall the following (Fact 4.4.6 follows easily from [30], Proposition 4.3.28):

Definition 4.4.5. The theory T in the language $\mathcal{L}$ has *quantifier elimination* if every $\mathcal{L}$-formula is equivalent modulo T to one without quantifiers.

Fact 4.4.6. *Let T be an $\mathcal{L}$-theory, and assume $\mathcal{L}$ containis at least one constant symbol. Suppose that whenever $\mathcal{M} = (M, ...)$ and $\mathcal{N} = (N, ...)$ are models of T, N is $|M|^+$-saturated, and $g : A \to B$ is an isomorphism of substructures $A \subset M$ and $B \subset N$, then for any $a \in M$ we can find substructures $A' \subset M$ and $B' \subset N$, and an isomorphism $g' : A' \to B'$ extending g, such that $a \in A'$. Then T has quantifier elimination.*

Now we show:

Proposition 4.4.7. *T has quantifier elimination.*

Proof. We use Fact 4.4.6. Let $\mathcal{M}$ and $\mathcal{N}$ be models of T with universes M and N, and $g : A \to B$ an isomorphism of substructures. Assume $\mathcal{N}$ is $|M|^+$ saturated, and let $a \in M$. We seek an extension of g to a substructure containing a.

Note that if $a \in A$ there is nothing to prove, so we assume $a \notin A$. In particular, since substructures contain all constants, a is not equal to any $\bar{c}$. By the remarks after the proof of Lemma 4.4.4, no image or preimage of a under any iterate of f is equal to any $\bar{c}$. Also, by A4, it follows that all $f^i(a)$ are distinct for non-negative integers i.

Now we consider two cases:

- First suppose $f^i(a) \in A$ for some $i > 0$. Assume i is minimal with this property. Then the substructure A' generated by $A \cap \{a\}$ is just

  $$A \cup \{a, f(a), ..., f^{i-1}(a)\},$$

  where $a, f(a), ..., f^{i-1}(a)$ are all distinct and do not belong to A.

  Let $u = f^i(a) \in A$. Then $u \neq \bar{0}$, so u has exactly two preimages under f. Since $f^{i-1}(a)$ is one such preimage and does not belong to A, it follows that u has at most one preimage under f in A.

  Now let $v = g(u) \in B$. By the isomorphism g it follows that $v \neq \bar{0}$ and v has at most one preimage in B; so by A3 we can find some $z \in N - B$ with $f(z) = v$. Noting that f is surjective by A3, let b be any preimage of z under f^{i-1}.

Note that, for any $f^j(b)$ with $0 \leq j \leq i-1$, we have

$$f^k(f^j(b)) = z \notin B,$$

where

$$k = (i-1) - j \geq 0.$$

Since B is closed under f^k, it follows that $f^j(b) \notin B$. In particular, $b \notin B$. It follows, as in the case of $a \notin A$, that b is not equal to any constant symbol, and so all iterates of f applied to b are distinct. Finally, since $f^i(b) = v \in B$, we conclude that $f^j(b) \in B$ for all $j \geq i$.

Now let B' be the substructure of N generated by $B \cup \{b\}$. By the above paragraph we have

$$B' = B \cup \{b, f(b), ..., f^{i-1}(b)\},$$

with each of these i additional elements distinct and not already in B, subject to the axiom that $f(f^{i-1}(b)) = v$. In other words, we have said exactly that the map $g' : A' \to B'$, extending g and sending $f^j(a)$ to $f^j(b)$ for $0 \leq j < i$, is an isomorphism.

- Now suppose $f^i(a) \notin A$ for all i. It follows that the substructure A' generated by $A \cup \{a\}$ is just

$$A \cup \{a, f(a), f^2(a), ...\},$$

where $a, f(a), f^2(a), ...$ are distinct elements of M which form a copy of $\mathbb{N}$ (with the successor function) that is disjoint from A.

Our task is now to extend the isomorphism g by finding such a copy of $\mathbb{N}$ in N which is disjoint from B. But this can be expressed in a partial type over B: namely, let $\Phi(x)$ be the colleciton of formulas asserting that (1) any two iterates of f, applied to x, are distinct, and (2) no iterate of f applied to x is equal to any element of B. Note that by A3 and A4, each formula in $\Phi(x)$ is satisfied by all but finitely many elements of N. By A1 N is infinite, which implies that $\Phi(x)$ is finitely satisfiable. So $\Phi(x)$ extends to a complete type p over B. Note that

$$|B| = |A| \leq |M|.$$

So, since $\mathcal{N}$ is $|M|^+$-saturated, p has a realization $b \in N$. Then the iterates of f applied to b form a copy of $\mathbb{N}$ disjoint from B, so the extension g' sending $f^i(a)$ to $f^i(b)$ is the desired isomorphism.

□

Armed with quantifier elimination, we can now show:

Proposition 4.4.8. *T is complete and strongly minimal, with trivial pregeometry. Moreover, P is a prime model of T.*

Proof. By quantifier elimination, any embedding of models of T is an elementary embedding. So, since P embeds into every model of T, it is a prime model. Now since T has a prime model it is complete: indeed, P is elementarily equivalent to every model of T, which shows that any two models of T are elementarily equivalent.

To show strong minimality, we could use two approaches. First, $\mathbb{C}$ is a model of T, with all structure definable from the field structure. So by strong minimality of the complex field, it follows that T is strongly minimal. Alteratively we could more carefully study the definable sets of models of T. By quantifier elimination, every formula in one variable is a Boolean combination of those of the form $f^i(x) = f^j(x)$ or $f^i(x) = a$ (where a is a parameter). By A3 and A4, such sets are either finite or cofinite. Since this holds in all models, T is strongly minimal.

Finally, we verify that T is of trivial type. Indeed, by quantifier elimination every nontrivial irreducible plane curve in a model of T is equivalent to one of the form $f^i(x) = f^j(y)$. In particular, such sets are always $\emptyset$-definable, and so there cannot be families of positive rank. □

Finally, we show the following:

Proposition 4.4.9. *There is a model of T with universe $\mathbb{C}^2$, with f interpreted as a constructible function.*

Proof. For each $c \in P$ we interpret $\bar{c}$ as $(c, 0) \in \mathbb{C}^2$. Then we interpret f as follows: for $a \in \mathbb{C}$ set

$$f(a, 0) = (a^2, 0),$$

and for $(a, b) \in \mathbb{C} \times (\mathbb{C} - \{0\})$ set

$$f(a, b) = (a + 1, b^2).$$

Then f is constructible, so the induced model of T is definable from the field structure on $\mathbb{C}$.

For convenience we split $\mathbb{C}^2$ into the sets $Y = \mathbb{C} \times \{0\}$ and $Z = \mathbb{C} \times (\mathbb{C} - \{0\})$. Note that $f(Y) = Y$ and $f(Z) = Z$.

Now by the definition of f on Y it is clear that Y is a substructure isomorphic to $\mathbb{C} \models T$. Thus axioms A1 and A2 are immediate. It remains to verify A3 and A4.

We first verify A3. Since f preserves Y and Z and $Y \cong \mathbb{C}$, it follows that A3 holds for all elements of y. Thus we need to show that every element of Z has exactly two preimages. But this is clear: if $b \neq 0$ then the preimages of (a, b) are precisely $(a - 1, c_i)$, where c_1, c_2 are the two square roots of b.

Lastly we verify A4. Again, since f preserves Y and Z and $Y \cong \mathbb{C}$, the statement of A4 is clear for elements of Y. So, assume $b \neq 0$, and suppose $f^i(a,b) = f^j(a,b)$ with $i > j \geq 0$. Evaluating using the formula for f, we get

$$(a+i, b^{2^i}) = (a+j, b^{2^j}).$$

Thus

$$a+i = a+j,$$

and so $i = j$, a contradiction. So in the case of points in Z, A4 is vacuously true, and we are done. □

Remark 4.4.10. Note that, by quantifier elimination and the proof of the above proposition, the structure on $\mathbb{C}^2$ is an elementary extension of $Y \cong \mathbb{C}^2$.

Finally, we conclude:

Corollary 4.4.11. *There is a trivial, non-unimodular strongly minimal reduct of the full field-induced structure on a smooth, simply connected, higher dimensional complex algebraic variety.*

Proof. The structure given in Proposition 4.4.9 is such a structure. Note, in particular, that no model of T is unimodular, since the functions $\mathrm{id} : M \to M$ and $f : M \to M$ have degrees 1 and 2, respectively. □

Chapter 5

The General Case of Smooth Varieties

The purpose of this chapter is to generalize the results of the previous chapter to varieties with infinite fundamental group. Precisely, we will prove:

Theorem 5.0.1. *Let M be a smooth variety of dimension $n > 1$ over an algebraically closed field K of characteristic zero, and let $\mathcal{M} = (M, ...)$ be a strongly minimal reduct of the full K-induced structure on M. If $\mathcal{F} = \{F_a\}_{a \in A}$ is an almost faithful family of plane curves in $\mathcal{M}$, and for all generic $a \in A$ the curve F_a is almost pure, then rk $\mathcal{F} \leq 1$.*

That is, we will bound the ranks of 'generically almost pure' families – those families in which every generic curve is almost pure. Our main strategy is to consider the fiber product $F \times_{M^2} F$ of such a family: that is, we consider

$$R_{\mathcal{F}} = \{(x, a, b) \in M^2 \times A^2 : x \in F_a \cap F_b\}.$$

Using the assumption that $\mathcal{F}$ is generically almost pure, we show that $R_{\mathcal{F}}$ satisfies a relativized, or 'local,' version of almost purity. We then show that, after replacing F with a fully generic subset F', and forming the corresponding fiber product R', we can assume that R' is pure dimensional, smooth, and generic in R. Now using a local version of Theorem 3.4.14 we conclude that the projection $R' \to A^2$ is unramified, and thus quasifinite. On the other hand, this projection is easily shown to not be quasifinite: indeed, if $a \in A$ is generic, then the fiber $R'_{(a,a)}$ contains all generic elements of F_a, and is therefore infinite.

It is worth noting the resemblance of this argument, at least philosophically, to the case of unimodularity. A common interpretation of the proof that unimodular structures are locally modular is as follows: using the degree theory of definable sets mentioned in the previous chapter, one proves a 'Bezout' theorem for intersections of plane curves: the number of intersection points of two inequivalent plane curves (or rather, the number of common realizations of their generic types) is the product of the degrees of the two curves. The interesting thing about this statement is that the degree of a curve is a 'generically determined' property: any two equivalent curves have the same degree. So, one concludes that the number of intersections of two curves is already determined by knowing 'almost all' points on the two curves. Informally, when intersecting two curves, if we know *almost all* information then we know *all* information. This is easily seen to be contradictory if there are enough curves present in the plane: one simply finds two curves intersecting at a generic point, and removes this point from each curve; what remains is a pair of curves, equivalent to the original pair, with fewer intersections.

Our strategy in this chapter is at least philosophically similar. Namely, if $\mathcal{F}$ is a family of rank at least 2, then the projection $R \to A^2$ mentioned above is easily seen to be almost surjective and almost finite-to-one. Now for $(a, b) \in A^2$, the fiber $R_{(a,b)}$ is simply the intersection $F_a \cap F_b$. We conclude that for a generic pair $(a, b) \in A^2$, the intersection of F_a and F_b is non-empty and finite. Now the strategy outlined above shows that, if we restrict to the fully generic $F' \subset F$, the corresponding projection $R' \to A^2$ is *everywhere* finite-to-one: that is, for *every* $(a, b) \in A^2$, the intersection $F'_a \cap F'_b$ is finite. So, as in the unimodular case, knowing *almost all* information about intersections of curves tells us *all* information (at least in the related family F'). But this assertion, that $F'_a \cap F'_b$ is always finite, is even more obviously contradictory: we simply observe that almost all of the F'_a are infinite, and thus the intersection $F'_a \cap F'_a$ is almost always infinite.

We now proceed with the argument.

5.1 Algebro-Geometric Preliminaries

In this section we mention a couple facts that we will use in the subsequent sections.

Convention 5.1.1. Recall that we assume all varieties are irreducible and quasiprojective, and we identify them with their K-points for some algebraically closed field K, which will always be clear from context, and most of the time will just be $\mathbb{C}$.

Now we begin with the following, which says that under reasonable hypoteses we can assume fiber products are smooth and pure dimensional:

Lemma 5.1.2. *Let X, Y, and Z be smooth complex algebraic varieties of dimensions m, n, and k, respectively, and let $f : X \to Z$ and $g : Y \to Z$ be dominant smooth morphisms. Then the fiber product $X \times_Z Y$ is smooth and of pure dimension $m + n - k$.*

Proof. It suffices to work in the category of smooth manifolds. That is, by the smoothness of X, Y, and Z, we may observe that X, Y, and Z naturally form smooth manifolds of dimension $2m$, $2n$, and $2k$. Further, since they are smooth morphisms, f and g define smooth submersions in the category of smooth manifolds (for example, by [48], Theorem 25.2.2(ii)).

Now the category of smooth manifolds has fiber products along submersions, which respect the underlying sets. More precisely, the fiber product of sets is the preimage of the diagonal $\Delta \subset Z \times Z$ in the map

$$f \times g : X \times Y \to Z \times Z.$$

But since f and g are submersions, it readily follows that $f \times g$ is transverse to Δ, so that the preimage of Δ is a smooth manifold of dimension $2m + 2n - 2k$ at every point (see [16], p. 28).

Now it follows that the algebraic fiber product $X \times_Z Y$, at the level of $\mathbb{C}$-points, is a smooth manifold of dimension $2m + 2n - 2k$ at every point; equivalently, the algebraic set $X \times_Z Y$ is smooth and of pure dimension $m + n - k$. □

We also need the following 'generic smoothness' phenomenon:

Fact 5.1.3. *Let $f : X \to Y$ be a dominant morphism of complex algebraic varieties, and assume Y is smooth. Then there is a dense open $X' \subset X$ such that:*

1. *X' is smooth.*
2. *The restriction of f to X' is a smooth morphism.*

Proof. By Theorem 5.3 of Chapter 1 of [17], we can achieve (1). Then by Theorem 25.3.1 of [48], we can achieve (2). Note that applying (2) does not affect (1), since open subsets of smooth varieties are smooth (again by Theorem 5.3 of Chapter 1 of [17]). □

Finally, we disuss the notion of ramification. Recall ([48], Section 21.6):

Definition 5.1.4. Let $f : X \to Y$ be a morphism of smooth complex algebraic varieties. The *ramification locus of f* is the support of the sheaf $\Omega_{X/Y}$ of relative differentials: that is, the set of $x \in X$ at which the induced map of tangent spaces is not injective. The image of the ramification locus in Y is called the *branch locus of f*.

It is a general fact that the ramification locus is Zariski closed (see for example [48], Exercise 21.6.H) – thus in particular it is $\mathbb{C}$-definable. If x belongs to the ramification locus, we say that f is *ramified at x*; otherwise we say that f is *unramified at x*. If f is unramified at all $x \in X$ then we say f is *unramified*.

Now the main facts we need about ramification are:

Fact 5.1.5. *Let $f : X \to Y$ be a morphism of smooth complex algebraic varieties.*

1. *f is unramified at $x \in X$ if and only if (x, x) belongs to the relative interior of the diagonal in the fiber product $X \times_Y X$.*
2. *If f is unramified then f is quasifinite – that is, all fibers $f^{-1}(y)$ are finite.*
3. *If $\dim X = \dim Y$ and f is dominant, then the ramification locus is either empty or of pure codimension 1 in X.*

Proof. (3) is the well-known 'purity of the ramification locus' ([50]; in our specific context, see [48], Exercise 21.7.A). For (2), it follows from [14] (Corollaire 17.4.3) that f is locally quasifinite – that is, all fibers $f^{-1}(y)$ are discrete. By quasicompactness of X, all discrete subsets are finite; thus all fibers $f^{-1}(y)$ are finite.

For (1), note that both of the given conditions are open, local conditions; so, passing to a dense open subset of X if necessary, it suffices to show that f is unramified if and only if the diagonal is relatively open in $X \times_Y X$. But this equivalence is well-known – see for example [14], Corollaire 17.4.2. □

5.2 Relative Almost Purity

The main theorem of this chapter relies on the observation that a certain set is 'locally' almost pure in a special way: that is, there is an open subset of a special type which is almost pure. In this section, we develop this 'local' or 'relative' version of almost purity. Our goal is an analogue of Theorem 3.4.14.

Convention 5.2.1. For the remainder of this section, assume M is a smooth complex algebraic variety of dimension $n > 1$, and $\mathcal{M} = (M, ...)$ is a strongly minimal reduct of the full $\mathbb{C}$-induced structure on M.

Now our relative form of almost purity is the following:

Definition 5.2.2. Let $D \subset M^k$ be a non-empty constructible set, let $\pi : M^k \to M^r$ be any projection, and let $U \subset M^r$ be a dense Zariski open set. Then D is (π, U)-*almost pure* if $D \cap \pi^{-1}(U)$ is almost pure.

Remark 5.2.3. Under reasonable assumptions (e.g. if the top dimensional components of D^P project dominantly but the codimension 1 components do not), D is automatically (π, U)-almost pure for some U (e.g. set U to be the complement of the closure of the projections of the codimension 1 components). So this notion is only really interesting if we verify it for a particularly interesting U. As stated above, this is precisely what we will do in the next section.

Our goal in this section is to prove the following proposition. This is done essentially identically as in Theorem 3.4.14:

Proposition 5.2.4. *Let $D \subset M^k$ be an $\mathcal{M}$-definable set of rank r, and let $\pi : D \to M^r$ be a projection which is almost surjective and almost finite-to-one. Assume D is (π, U)-almost pure for some dense Zariski open set $U \subset M^r$. Then there is a Zariski open $W \subset U$ such that:*

1. $\dim(U - W) \leq \dim U - 2$.
2. *The restriction of π to $\pi^{-1}(W)$ is a finite covering of W.*

Proof. Assume all of this data is definable over $\emptyset$. Note that $\dim U = nr$. Let l be such that $|D_y| = l$ for all generic $y \in M^r$. We now follow the same steps as in Theorem 3.4.14.

Lemma 5.2.5. *If $y \in U$ and $\dim y \geq nr - 1$, then $D_y = (D^P)_y$ and both sets have size exactly l.*

Proof. As in Theorem 3.4.14, note that y has codimension at most 1 in M^r, so y is $\mathcal{M}$-generic in M^r. The equality $|D_y| = l$ then follows since it is a generic $\mathcal{M}$-definable condition on y.

Also as in Theorem 3.4.14, the projection of D^P to M^r is almost surjective and almost finite-to-one; since y has codimension at most 1 in M^r, Corollary 3.4.10 implies $|(D^P)_y| \leq l$.

So, as in Theorem 3.4.14, it suffices to observe that $D_y \subset (D^P)_y$. To see this, let $x \in D_y$. Then $\dim(y/x) = 0$, so

$$\dim x \geq nr - 1.$$

In particular, x has codimension at most 1 in $D \cap \pi^{-1}(U)$. But $D \cap \pi^{-1}(U)$ is almost pure, so we get

$$x \in (D \cap \pi^{-1}(U))^P.$$

On the other hand, note that U is fully generic in M^r. So, since π is almost surjective and almost finite-to-one on D, it follows that $D \cap \pi^{-1}(U)$ is almost equal to D. In particular, $D \cap \pi^{-1}(U)$ has the same pure part as D. Thus $x \in D^P$, as desired. □

Now let W be the interior of the set of all $y \in U$ such that $D_y = (D^P)_y$ and both sides have size exactly l. By the above lemma W contains all points of codimension at most 1 in U, so it follows that

$$\dim(U - W) \leq nr - 2.$$

Finally, as in Theorem 3.4.14, it follows by Corollary 6.7.10 that the restriction of π to $\pi^{-1}(W)$ is an l-sheeted covering of W. □

Finally, we point out the main application we make of Proposition 5.2.4:

Corollary 5.2.6. *Assume the hypotheses of Proposition 5.2.4, and let $D' \subset D$ be a smooth variety of the same dimension as D, whose projection is contained in U. Then, letting $f : D' \to U$ be the projection, the branch locus of f has codimension at least 2 in U.*

Proof. Assume all of this data is $\emptyset$-definable. Then it suffices to show that, if $y \in U$ is a point of codimension at most 1, and $x \in f^{-1}(y)$, then f is unramified at x.

So, let $y \in U$ with $\dim y \geq nr - 1$, and let $x \in f^{-1}(y)$. Let W be the set provided by Proposition 5.2.4. Then $U - W$ has codimension at least 2 in U, so it follows that $y \in W$.

Let $x_1, ..., x_l$ be the distinct elements of D_y. Without loss of generality we may assume $x = x_1$. By Proposition 5.2.4, there are disjoint open neighborhoods B_i of each x_i in D, and V of y in W, such that the restriction of π to each B_i is a homeomorphism to V. It follows that π is injective on $B_1 \subset D$. Since $D' \subset D$, it also follows that f is injective on $B_1 \cap D'$. Equivalently, the restriction of $D' \times_U D'$ to $(B_1)^2$ is contained in the diagonal. Thus (x, x) belongs to the interior of the diagonal in $D' \times_U D'$; by Fact 5.1.5, we conclude that f is unramified at x. □

5.3 Relative Almost Purity of the Fiber Product of a Family

Now that we have our relative notion of almost purity, our next aim is to verify this notion for a specific set, and apply Corollary 5.2.6.

Convention 5.3.1. Throughout this section, we retain the underlying assumptions of the previous section. That is, we assume M is a smooth complex algebraic variety of dimension $n > 1$, and $\mathcal{M} = (M, ...)$ is a strongly minimal reduct of the full $\mathbb{C}$-induced structure on M.

Given a family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves in $\mathcal{M}$, we define the *associated fiber product*

$$R_{\mathcal{F}} = \{(x, a, b) \in M^2 \times A^2 : x \in F_a \cap F_b\}.$$

So $R_{\mathcal{F}}$ is $\mathcal{M}$-definable over the same parameters as $\mathcal{F}$. We first note the following:

Lemma 5.3.2. *Assume $\mathcal{F}$ is almost faithful and of rank at least 2, with A stationary. Set $R = R_{\mathcal{F}}$. Then rk $R = 2 \cdot$ rk A, and the projection $\pi : R \to A^2$ is almost surjective and almost finite-to-one.*

Proof. First we show that rk $R \geq 2 \cdot$ rk A and π is almost surjective. Indeed, let $a \in A$ be generic, and let x be generic in F_a. It follows that rk $(x/a) = 1$, and so rk $x \geq 1$. By Lemma 4.3.2, x is not a common point of $\mathcal{F}$; it follows that

$$\text{rk } (a/x) \leq \text{rk } A - 1.$$

On the other hand we have

$$\text{rk } (a, x) = \text{rk } A + 1,$$

so this implies rk $x \geq 2$. But $x \in M^2$, so rk $x = 2$, and thus

$$\text{rk } (a/x) = \text{rk } A - 1.$$

Let b be an independent realization of the (strong) type of a over x. So $x \in F_a \cap F_b$, and

$$\text{rk } (b/a, x) = \text{rk } A - 1.$$

We are assuming rk $A \geq 2$, so this implies b is not algebraic over a. In particular, by almost faithfulness, $F_a \cap F_b$ is finite, and so

$$\text{rk } (x/a, b) = 0.$$

Now by additivity

$$\text{rk } (a, x, b) = 2 \cdot \text{rk } A,$$

which implies rk $R \geq 2 \cdot$ rk A. Further, since rk $(x/a, b) = 0$ we get

$$\text{rk } (a, b) = 2 \cdot \text{rk } A = \text{rk } A^2.$$

Since A is stationary, so is A^2, so this implies π is almost surjective.

Secondly, we show that rk $R \leq 2 \cdot$ rk A and π is almost finite-to-one. For this, let $(x, a, b) \in R$ be any element with

$$\text{rk } (x, a, b) \geq 2 \cdot \text{rk } A.$$

Note that since $x \in F_a$,

$$\text{rk } (a, x) \leq \text{rk } A + 1,$$

so

$$\text{rk } (b/a, x) \geq \text{rk } A - 1 \geq 1.$$

As in the previous paragraph, this implies $F_a \cap F_b$ is finite. So x belongs to a finite fiber of π, which shows that π is almost finite-to-one.

Furthermore, we have rk $(x/a, b) = 0$, so

$$\text{rk } (x, a, b) = \text{rk } (a, b) \leq 2 \cdot \text{rk } A.$$

This shows rk $R \leq 2 \cdot$ rk A.

□

Convention 5.3.3. For the rest of this section, fix an almost faithful family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves of rank $r \geq 2$. Let $F \subset M^2 \times A$ be the associated definable set, and $R = R_{\mathcal{F}} \subset M^2 \times A^2$ the fiber product. Let $\pi : F \to A$ and $\tau : R \to A^2$ be the projections, so that Lemma 5.3.2 implies τ is almost surjective and almost finite-to-one. To avoid confusion with the curves F_a, we use the notation

$$_xF = \{a \in A : (x, a) \in F\}$$

for $x \in M^2$; we similarly use $_x(F^P)$ for the fiber at x in F^P, etc.

Remark 5.3.4. In the next section, we will modify such a family in order to assume the parameter set A is a power of M – or at least a generic subset thereof. That is, since A has rank r, we will produce a scenario where A is a generic subset of M^r. Note that in this case we may view π and τ as projections to M^r and M^{2r}. Of course, since A and M^r are almost equal, it will still follow that τ is almost surjective and almost finite-to-one.

Now our main goal in this section is the following proposition. Roughly, it says that the relative almost purity of F over a given set is always inherited by R. The main thing to note is that R becomes relatively almost pure over a *square* of an open subset of A: later on this will allows us to 'intersect a generic curve with itself,' as described in the introduction to this chapter.

Proposition 5.3.5. *Assume A is a generic subset of M^r. If $U \subset M^r$ is a dense Zariski open set such that F is (π, U)-almost pure, then R is (τ, U^2)-almost pure.*

Proof. Assume $U \subset M^r$ is a dense Zariski open set such that F is (π, U)-almost pure. We set

$$F_U = F \cap \pi^{-1}(U)$$

and

$$R_U = R \cap \tau^{-1}(U^2).$$

Since all F_a are of the same dimension, it follows that F and F_U are almost equal. Also, since τ is almost surjective and almost finite-to-one and $U \subset M^r$ is fully generic, it follows that R_U is almost equal to R, and so $\dim R_U = 2nr$. Adding parameters if necessary, we may assume F and U are $\emptyset$-definable.

Now our goal is to show that R_U is almost pure. To do this, let $(x_0, a_0, b_0) \in R_U$ with

$$\dim(x_0, a_0, b_0) \geq 2nr - 1.$$

We wish to show $(x_0, a_0, b_0) \in (R_U)^P$. Since R and R_U are almost equal, this is equivalent to showing $(x_0, a_0, b_0) \in R^P$. To do this, let $W \subset M^2 \times M^{2r}$ be any Zariski open set containing (x_0, a_0, b_0); we may further assume W is $\emptyset$-definable. We wish to find a generic element of R which belongs to W.

First suppose x_0 is a common point of F. In this case, since A is stationary it follows that $x_0 \in F_a$ for all generic $a \in A$. Let $W_{x_0} \subset M^{2r}$ be the fiber above x_0 in W. Then W_{x_0} is open and non-empty, since it contains (a_0, b_0). Since M^r is irreducible, W_{x_0} is dense, and therefore fully generic in M^{2r}. Let $(a_1, b_1) \in A^2$ be any generic point; it thus follows that $x_0 \in F_{a_1} \cap F_{b_1}$ and $(a_1, b_1) \in W_{x_0}$. So

$$(x_0, a_1, b_1) \in R \cap W.$$

Since (a_1, b_1) is generic in A^2 we have

$$\dim(x_0, a_1, b_1) \geq \dim(a_1, b_1) = 2nr,$$

so (x_0, a_1, b_1) is generic in R, as desired.

Now suppose x_0 is not a common point of F. In this case the argument is much more complicated. To start, we note:

Claim 5.3.6. *For every generic $x \in M^2$, the fiber ${}_xF$ has rank $r - 1$.*

Proof. Since $\mathcal{F}$ has only finitely many common points, it is clear that $\text{rk}\ ({}_{x'}F) \leq r - 1$ holds generically. Now let $a \in A$ be generic, and x' generic in F_a over a. Then

$$\text{rk}\ (a, x') = r + 1,$$

and

$$\text{rk}\ (x'/a) = 1.$$

So $\text{rk}\ x' \geq 1$, and thus x' is not a common point. We conclude:

1. rk $x' \leq 2$, since $x' \in M^2$.
2. rk $(a/x') \leq r-1$, since x' is not common.

Using additivity and the fact rk $(a, x') = r+1$, it follows that each of (1) and (2) above is an equality. Thus x' is generic and has fiber of rank $r-1$. Since M^2 is stationary, the same holds of x. □

We immediately conclude:

Claim 5.3.7. *For every generic $x \in M^2$, the fiber ${}_x(F^P)$ has dimension $n(r-1)$.*

Proof. By the previous claim, the equality

$$\dim({}_xF) = n(r-1)$$

holds generically. Since F and F^P are almost equal, the present claim immediately follows. □

Next we show:

Claim 5.3.8. *Each of the following holds:*

1. $\dim x_0 \geq 2n-1$.
2. $\dim(x_0, a_0) \geq n(r+1)-1$.
3. $\dim(x_0, b_0) \geq n(r+1)-1$.

Proof. Since x_0 is not common we have

$$\text{rk } (a_0/x_0) \leq r-1,$$

and so

$$\dim(a_0/x_0, b_0) \leq \dim(a_0/x_0) \leq n(r-1).$$

Then, using additivity and the fact that

$$\dim(x_0, a_0, b_0) \geq 2nr-1,$$

we immediately conclude (3). Note that a symmetric argument gives (2).

Finally, recall from above that

$$\dim(a_0/x_0) \leq n(r-1).$$

Then, using additivity and (2), we conclude (1). □

We can now conclude the following, which is the first main step of the proof:

Corollary 5.3.9. *(x_0, a_0) and (x_0, b_0) both belong to F^P.*

Proof. Since $(x_0, a_0, b_0) \in R_U$, it follows that (x_0, a_0) and (x_0, b_0) both belong to F_U. By the previous claim they each have codimension at most 1 in F_U. So, since F_U is almost pure, each of (x_0, a_0) and (x_0, b_0) belongs to $(F_U)^P$. Since F and F_U are almost equal, $(F_U)^P = F^P$, and we are done. □

Next we note:

Claim 5.3.10. *The set ${}_{x_0}(F^P)$ has dimension $n(r-1)$.*

Proof. If $\dim x_0 = 2n$ this follows by Claim 5.3.7. So, since $\dim x_0 \geq 2n-1$, we may assume $\dim x_0 = 2n-1$. We first show ${}_{x_0}(F^P)$ has dimension at most $n(r-1)$. To see this, suppose

$$\dim({}_{x_0}(F^P)) \geq n(r-1)+1,$$

and let a be a generic element of this set. Then by additivity

$$\dim(x_0, a) \geq n(r+1) = \dim F^P,$$

so (x_0, a) is a generic element of F^P. Since F and F^P are almost equal, we get $(x_0, a) \in F$. But then, since x_0 is not common, we get

$$\text{rk } (a/x_0) \leq r-1,$$

and so

$$\dim(a/x_0) \leq n(r-1),$$

a contradiction.

Now note that, since $\dim x_0 = 2n-1$ and

$$\dim(x_0, a_0) \geq n(r+1) - 1,$$

it follows by additivity that

$$\dim(a_0/x_0) \geq n(r-1).$$

But $a_0 \in {}_{x_0}(F^P)$. So it follows that ${}_{x_0}(F^P)$ has dimension exactly $n(r-1)$, as desired. □

Now the next main step in the proof is the following conclusion:

Corollary 5.3.11. *The projection $F^P \to M^2$ is locally surjective near each of (x_0, a_0) and (x_0, b_0).*

Proof. In the ambient projective space containing $M^2 \times M^r$, let $\mathcal{H}$ be a collection of $n(r-1)$ hyperplances through (x_0, a_0), which are independent and generic over all other relevant data. Let

$$G = F^P \cap \bigcap \mathcal{H}.$$

We conclude:

Claim 5.3.12. *G is of pure dimension $2n$.*

Proof. By the genericity and independence of $\mathcal{H}$, each $H \in \mathcal{H}$ decreases the dimension of F^P by at least one, which means that each irreducible component of G has dimension at most

$$\dim F^P - n(r-1) = n(r+1) - n(r-1) = 2n.$$

On the other hand, since $M^2 \times M^r$ is smooth and F^P is pure dimensional, each irreducible component of G has codimension at most $n(r-1)$ in F^P (see [11], Theorem 0.2), and therefore has dimension at least $2n$. The claim follows. □

Claim 5.3.13. *The fiber ${}_{x_0}G$ is finite and non-empty.*

Proof. Similarly to the previous claim, the genericity and independence of $\mathcal{H}$ implies that the dimension of ${}_{x_0}(F^P)$ drops by at least $n(r-1)$ when intersected with $\bigcap \mathcal{H}$. So by Claim 5.3.10, ${}_{x_0}(F^P)$ is finite. On the other hand $a_0 \in G_{x_0}$ by definition, so G_{x_0} is non-empty. □

Now by the semicontinuity of fiber dimension ([48], Theorem 11.4.2), there is a Zariski open neighborhood V around (x_0, a_0) in G such that the projection $V \to M^2$ is finite-to-one. Then, noting that G is closed in $M^2 \times M^r$ by definition, it follows that V is locally compact. Moreover, V is of pure dimension $2n$ since G is. Then since $\dim D = \dim M^2$, the projection $D \to M^2$ is in fact almost surjective. So, by Proposition 3.4.5, we conclude that the projection $D \to M^2$ is locally surjective near (x_0, a_0). Since $D \subset F^P$, this clearly implies that the projection $F^P \to M^2$ is locally surjective near (x_0, a_0). Then by symmetry the same holds of (x_0, b_0), which completes the proof of the corollary. □

We are finally ready to prove Proposition 5.3.5. We proceed as follows:

Definition 5.3.14. We say that a point $(x, a) \in M^2 \times M^r$ is *good* if there is some $b \in M^r$ such that $(x, a, b) \in W$ and $(x, b) \in F^P$.

It is clear that (x_0, a_0) is good, as witnessed by b_0. Now the main point to observe is:

Claim 5.3.15. *There is an analytic neighborhood of (x_0, a_0) in $M^2 \times M^r$ consisting only of good points.*

Proof. Since $(x_0, a_0, b_0) \in W$ and W is open, there are analytic neighborhoods X of x_0 in M^2, Y of a_0 in M^r, and Z of b_0 in M^r, such that $X \times Y \times Z \subset W$. Shrinking X and Z if necessary, we may assume by Corollary 5.3.11 that $X \subset (F^P)_Z$, where $(F^P)_Z$ is the set of elements of F^P with M^r-coordinate belonging to Z.

Now we claim that $X \times Y$ is the desired neighborhood. Indeed, for any $(x, a) \in X \times Y$, there is some $b \in Z$ with $(x, b) \in F^P$. Since $X \times Y \times Z \subset W$, we also get $(x, a, b) \in W$, as desired. □

Now by the previous claim and the fact that $(x_0, a_0) \in F^P$, there is a generic element (x_1, a_1) of F which is good. Then x_1 is generic in M^2, so by Lemma 3.2.7 we obtain

$$_{x_1}(F^P) = (_{x_1}F)^P.$$

In particular, $_{x_1}(F^P)$ is closed and of pure dimension $n(r-1)$. Let $W_{(x_1,a_1)}$ be the fiber above (x_1, a_1) in W, viewed as a subset of M^r. Then $W_{(x_1,a_1)}$ is open. By the goodness of (x_1, a_1), W has non-empty intersection with $_{x_1}(F^P)$. Then, since $_{x_1}(F^P)$ is of pure dimension $n(r-1)$, its intersection with the open set W is dense in a component of dimension $n(r-1)$, and therefore is also of dimension $n(r-1)$. Let b_1 be a generic element of this intersection over (x_1, a_1). Then $(x_1, b_1) \in F^P$ and

$$\dim(b_1/x_1, a_1) = n(r-1).$$

In particular it follows that

$$\dim(b_1/x_1) \geq n(r-1),$$

so by additivity

$$\dim(x_1, b_1) \geq n(r+1).$$

Thus (x_1, b_1) is generic in F^P, and so belongs to F; we conclude that $(x_1, a_1, b_1) \in R$. Now again using that

$$\dim(b_1/x_1, a_1) = n(r-1),$$

additivity gives that

$$\dim(x_1, a_1, b_1) = 2nr,$$

so (x_1, a_1, b_1) is a generic element of R. Finally, by the choice of b_1 we have $(x_1, a_1, b_1) \in W$, which completes the proof of Proposition 5.3.5.

□

5.4 Proof of the Main Theorem

In this section we present a proof of Theorem 5.0.1. Our strategy is to combine Proposition 5.3.5 with Corollary 5.2.6 to conclude that the projection $R' \to A^2$ is unramified, where R' is a certain generic pure dimensional subset of R. We will conclude this projection is quasifinite by Fact 5.1.5 – a conclusion which is easily shown to be contradictory.

Convention 5.4.1. Throughout this section, assume M is a smooth variety of dimension $n > 1$ over an algebraically closed field K of characteristic 0, and $\mathcal{M} = (M, ...)$ is a strongly minimal reduct of the full K-induced structure on M. We fix $\mathcal{F} = \{F_a\}_{a \in A}$, an almost faithful family of plane curves in $\mathcal{M}$, where $A \subset M^k$ is a definable set. We assume all of this data is $\emptyset$-definable.

Now Theorem 5.0.1 equivalently states the following:

Theorem 5.4.2. *Assume that for all generic $a \in A$ the curve F_a is almost pure. Then rk $A \leq 1$.*

Proof. By saturation, we may assume $K = \mathbb{C}$. Let $r = \text{rk } A$, and assume $r \geq 2$. We begin with:

Lemma 5.4.3. *We may assume A is a generic subset of M^r.*

Proof. Since A is $\mathcal{M}$-definable and of rank r, there is an $\mathcal{M}$-definable function $f : A \to M^r$ which is almost surjective and almost finite-to-one. Adding parameters if necessary, we may assume f is $\emptyset$-definable. Let $B \subset M^r$ be the set of elements with non-empty and finite preimage under f; so B is generic in M^r.

We now define a new family $\mathcal{G} = \{G_b\}_{b \in B}$ as follows: for each $b \in B$, set

$$G_b = \bigcup_{a \in f^{-1}(b)} F_a.$$

Then each G_b is the union of finitely many plane curves, and is thus a plane curve. It remains to show the following two claims:

Claim 5.4.4. *If $b \in B$ is generic then G_b is almost pure.*

Proof. Let $a \in f^{-1}(b)$. Then

$$\dim a \leq \dim A = nr.$$

Since $f(a) = b$ we have $\dim(b/a) = 0$, so it follows that

$$\dim a \geq \dim b = \dim B = nr.$$

Thus a is generic in A, and so F_a is almost pure.

It follows that G_b is the union of finitely many almost pure curves. To show G_b is almost pure, let $x \in G_b$ be a point of codimension at most 1. So, for some $a \in f^{-1}(b)$, x is a point of codimension at most 1 in F_a. Since F_a is almost pure, this implies

$$x \in (F_a)^P \subset (G_b)^P.$$

□

Claim 5.4.5. *$\mathcal{G}$ is almost faithful.*

Proof. Let $b, b' \in B$ such that $G_b \cap G_{b'}$ is infinite. We wish to show b and b' are interalgebraic. Now since $f^{-1}(b)$ and $f^{-1}(b')$ are finite, it follows that $F_a \cap F_{a'}$ is infinite for some $a \in f^{-1}(b)$ and $a' \in f^{-1}(b')$. Since $\mathcal{F}$ is almost faithful, a and a' are interalgebraic. Again since $f^{-1}(b)$ and $f^{-1}(b')$ are finite, it follows that a and b are interalgebraic, as are a' and b'. Thus all of a, b, a', and b' are interalgebraic, so in particular b and b' are interalgebraic. □

□

Replacing $\mathcal{F}$ with the family $\mathcal{G}$ provided by the previous lemma, we assume from now on that A is a generic subset of M^r. Let $F \subset M^2 \times M^r$ be the associated definable set. We next show:

Lemma 5.4.6. *There is a dense Zariski open set $U \subset M^r$ such that F is (π, U)-almost pure, where $\pi : F \to M^r$ is the projection.*

Proof. Let W be the set of all $a \in A$ such that

1. F_a is almost pure.
2. $(F_a)^P = (F^P)_a$.

Note that each of (1) and (2) holds generically in M^r. Indeed, the genericity of (1) is our assumption on $\mathcal{F}$; for (2), note that as in Claim 5.3.6 we have

$$\text{rk } ({}_xF) = r - 1,$$

and thus

$$\dim({}_xF) = n(r-1),$$

for all generic $x \in M^2$. So by Lemma 3.2.7, (2) holds generically.

It follows that W is generic in M^r. Let U be the interior of W, so U is a dense (thus generic) open set in M^r. We claim that F is (π, U)-almost pure – or equivalently, that F_U is almost pure, where

$$F_U = \{(x, a) \in F : a \in U\}.$$

To see this, let (x, a) be a point of codimension at most 1 in F_U. So

$$\dim(x, a) = \dim a + \dim(x/a) \geq n(r+1) - 1.$$

Since $\dim a \leq nr$, additivity implies

$$\dim(x/a) \geq n - 1.$$

So x is a point of codimension at most 1 in F_a. By (1) we get $x \in (F_a)^P$; so by (2) we get $x \in (F^P)_a$, i.e. $(x, a) \in F^P$. On the other hand it is clear that F and F_U are almost equal, so $F^P = (F_U)^P$, and we are done. □

Fix the open set U provided in the lemma. We may assume U is $\emptyset$-definable. Let

$$R = R_{\mathcal{F}} \subset M^2 \times M^{2r}$$

be the fiber product studied in the previous section, and let $\tau : R \to M^{2r}$ be the projection. So τ is almost surjective and almost finite-to-one, and by Proposition 5.3.5 R is (τ, U^2)-almost pure.

Claim 5.4.7. *For each top dimensional irreducible component C of F^P, the projection $C \to M^2$ is dominant.*

Proof. Let $(x, a) \in C$ be generic. Then (x, a) is a generic element of F, so

$$\dim(x, a) = n(r + 1).$$

Since $\dim a \leq nr$, it follows that x is not algebraic over $\emptyset$, and so x is not a common point. Thus

$$\dim(a/x) \leq n(r - 1).$$

Since $\dim(x, a) = n(r - 1)$, additivity implies $\dim x = 2n$. So a generic element of M^2 is in the image of C, which is exactly what we needed to show. □

We now apply Fact 5.1.3 to each irreducible component of F^P in the projection to M^2. We obtain a fully generic open subset $F' \subset F^P$ such that each irreducible component of F' is smooth, of dimension $n(r+1)$, and projects smoothly and dominantly to M^2. Restricting to a dense open set if necessary, we may assume the components of F' are pairwise disjoint. As in the previous section we denote

$$F_U = \{(x, a) \in F : a \in U\}.$$

Note that F_U is almost equal to F, and therefore almost equal to F^P. So, shrinking F' again if necessary, we may further assume $F' \subset F_U$. At this point, we assume F' is $\emptyset$-definable.

Let

$$R' = \{(x, a, b) : x \in F'_a \cap F'_b\}$$

be the fiber produt using F' in place of F. Then we show:

Claim 5.4.8. *Let B be any irreducible component of R'. Then:*

1. *B is smooth and of dimension $2nr$.*
2. *The branch locus of the projection $B \to U^2$ has codimension at least 2 in U^2.*

Proof. Note that R' is the union of all fiber products $C \times_{M^2} D$, where C and D range over the irreducible components of F'. By Lemma 5.1.2, each such fiber product $C \times_{M^2} D$ is smooth and of pure dimension $2nr$, which is enough to prove (1).

To prove (2), we use the fact that R is (τ, U^2)-almost pure. That is, letting

$$R_U = \{(x, a, b) \in R : a, b \in U\},$$

we get that R_U is almost pure. Now since $F' \subset F_U$, it is clear that $R' \subset R_U$; in particular, $B \subset R_U$. The desired conclusion now follows immediately from (1) and Corollary 5.2.6. □

We conclude:

Lemma 5.4.9. *If B is any irreducible component of R', then the projection $B \to U^2$ is unramified.*

Proof. By the previous claim, B is smooth. Also, since the projection $R \to A^2$ is almost surjective and almost finite-to-one, and $B \subset R$ is generic, it follows that the projection $B \to U^2$ is dominant. Finally, since U^2 is open in M^{2r}, U^2 is smooth.

It follows that we are in the situation of Fact 5.1.5 (3); so it suffices to show that, if $(x, a, b) \in B$ is a point of codimension at most 1, then the projection $B \to U^2$ is unramified at (x, a, b).

Now by the previous claim, $\dim B = 2nr$. So, we assume $(x, a, b) \in B$ and

$$\dim(x, a, b) \geq 2nr - 1.$$

Then (x, a, b) is also a point of codimension at most 1 in R; in particular, we conclude that (x, a, b) is $\mathcal{M}$-generic in R. Now since the projection $R \to M^{2r}$ is almost finite-to-one, we obtain rk $(x/a, b) = 0$, and so $\dim(x/a, b) = 0$. By additivity, this implies

$$\dim(a, b) \geq 2nr - 1.$$

But by the previous claim, the branch locus of $B \to U^2$ has dimension at most $2nr - 2$; it follows that (a, b) does not belong to the branch locus, and so (x, a, b) does not belong to the ramification locus. □

By Fact 5.1.5, each projection $B \to U^2$ is quasifinite. Since there are only finitely many components, we obtain that the projection $R' \to U^2$ is quasifinite. That is, for all $a, b \in U$ the fiber $R'_{(a,b)}$ is finite. Finally, to obtain a contradiction we note:

Lemma 5.4.10. *The projection $R' \to U^2$ is not quasifinite.*

Proof. Let (x, a) be any generic element of F', so (x, a) is also generic in F. Then

$$\dim(x, a) = n(r + 1)$$

and

$$\dim a \leq nr,$$

so by additivity

$$\dim(x/a) \geq n.$$

Now clearly $(x, a, a) \in R'$, and

$$\dim(x/a, a) \geq n.$$

Thus the fiber $R'_{(a,a)}$ is infinite. □

Since we had concluded the projection $R' \to U^2$ is quasifinite, we now have a contradiction, and the proof of Theorem 5.0.1 is complete. □

5.5 Interpreting a Group

We end this chapter with an application of the main theorem – namely, we show that the presence of any generically almost pure family of positive rank implies the interpretability of a strongly minimal group. The main idea is to note that almost purity is closed under composition. Since we know by the main theorem that generically almost pure families have rank ≤ 1, the existence of any such family implies the existence of a family which does not grow under composition. Then, applying standard techniques (namely the 'group configuration'), we show that this implies the interpretability of a group.

We begin with a review of group configurations. For now, let $\mathcal{X} = (X, ...)$ be any strongly minimal structure. In this section we will use the notion of canonical bases. For more details and basic facts, see for example [38]. Recall that we define the *code* of a stationary definable set D to be the canonical base of its generic type, over any set of parameters which defines D.

Definition 5.5.1. A triple (a, b, c) of elements of X^{eq} is called a *non-trivial configuration* if

$$\text{rk } a = \text{rk } b = \text{rk } c = 1,$$

and

$$\text{rk } (a, b) = \text{rk } (a, c) = \text{rk } (b, c) = \text{rk } (a, b, c) = 2.$$

That is, each point has rank 1, any two are independent, and any one is algebraic over the other two.

The term 'non-trivial' arises from the fact that non-trivial strongly minimal structures are precisely those admitting non-trivial configurations (possibly after adding parameters). For example, we have the following:

Lemma 5.5.2. *Let $S \subset X^2$ be a non-trivial irreducible plane curve, and let c be a code of S. If rk $c = 1$ and $(x, y) \in S$ is generic in S over c, then (x, y, c) is a non-trivial configuration.*

Proof. We assumed rk $c = 1$, and since (x, y) is generic in S we have rk $(x, y/c) = 1$, thus rk $(x, y, c) = 2$. Since S is non-trivial, both projections $S \to M$ are finite-to-one everywhere. This implies that x and y are interalgebraic over c. In particuar,

$$\text{rk } x = \text{rk } y = \text{rk } (x/c) = \text{rk } (y/c) = 1,$$

so by additivity

$$\text{rk } (x, c) = \text{rk } (y, c) = 2.$$

It remains only to check that rk $(x, y) = 2$, i.e. that rk $(x, y) \neq 1$. But if rk $(x, y) = 1$ then there is a $\emptyset$-definable plane curve C containing (x, y). Since (x, y) is generic in S, it follows that every generic element of S belongs to C; but then S is one of the strongly minimal components of C, which means that rk $c = 0$, a contradiction. □

We now define:

Definition 5.5.3. A tuple (a, b, c, x, y, z) of elements of X^{eq} is a *one-dimensional group configuration* if:

1. The triples (a, b, c), (a, x, y), (b, y, z), and (c, x, z) form non-trivial configurations.
2. (x, y, z) is independent – that is, rk $(x, y, z) = 3$.

Remark 5.5.4. To make sense of this definition, it helps to consider the most natural example. Suppose G is a strongly minimal group definable in X, and G acts definably and transitively on X. Consider any generic element $(a, b, x) \in G^2 \times X$. Set $c = ba$, $y = ax$, and $z = by = cx$. Then (a, b, c, x, y, z) forms a one-dimensional group configuration.

In light of the above remark, one often thinks of group configurations as sets of points which 'look like' they come from groups. Indeed, a well-known result of Hrushovski ([23]; or, see chapter 5 of [38]) states that in certain theories, such configurations only arise from actual groups. In particular, restricting to our needs, we have the following:

Fact 5.5.5. *If there is a one-dimensional group configuration in the strongly minimal structure $\mathcal{X}$, then $\mathcal{X}$ interprets a strongly minimal group.*

We are now ready to give the result of this section. We return to our algebraic setting:

Convention 5.5.6. For the remainder of the section, assume M is a smooth variety of dimension $n > 1$ over an algebraically closed field K of characteristic 0, and $\mathcal{M} = (M, ...)$ is a strongly minimal reduct of the full K-induced structure on M.

Definition 5.5.7. A family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves in $\mathcal{M}$ is *generically almost pure* if for every generic $a \in A$ the plane curve $F_a \subset M^2$ is almost pure.

Now we show:

Theorem 5.5.8. *Assume $\mathcal{M}$ admits an almost faithful, generically almost pure family of non-trivial plane curves of positive rank. Then $\mathcal{M}$ interprets a strongly minimal group.*

Proof. We will construct a one-dimensional group configuration. While the process is straightforward and standard, we will need to develop a series of facts. We begin with:

Lemma 5.5.9. *If C and D are non-trivial almost pure plane curves, then $D \circ C$ is also a non-trivial almost pure plane curve.*

Proof. We may assume C and D are $\emptyset$-definable. That $D \circ C$ is a non-trivial plane curve is clear. To show $D \circ C$ is almost pure, let $(x, z) \in D \circ C$ with $\dim(x, z) \geq n - 1$. By definition there is some $y \in M$ such that $(x, y) \in C$ and $(y, z) \in D$. Now since C and D are non-trivial, each projection of either set to M is finite-to-one everywhere. It follows that x and y are interalgebraic, as are y and z. In particular,

$$\dim x = \dim y = \dim z \geq n - 1.$$

By the almost purity of C and D, this implies $(x, y) \in C^P$ and $(y, z) \in D^P$.

Let π be the projection of C^P to the left copy of M, and τ the projection of D^P to the left copy of M. Since

$$\dim x = \dim y \geq n - 1,$$

each of the fibers $\pi^{-1}(x)$ and $\tau^{-1}(y)$ is finite (for example, see Corollary 3.4.10). In particular, by Proposition 3.4.5, π is locally surjective near (x, y) and τ is locally surjective near (y, z).

Fix any analytic neighborhood $X \times Z$ of (x, z). Shrinking if necessary, we can take a neighborhood Y of y such that

$$X \subset \pi(C^P \cap (X \times Y))$$

and

$$Y \subset \tau(D^P \cap (Y \times Z)).$$

Now let $x' \in X$ be any generic element of M. Then there is some $y' \in Y$ with $(x', y') \in C^P$, and so there is some $z' \in Z$ with $(y', z') \in D^P$. Since x' is generic, it follows that $\pi^{-1}(x')$ is finite and contained in C. In particular, $(x', y') \in C$, and so x' and y' are interalgebraic. Thus y' is also generic. By similar reasoning, we conclude that $(y', z') \in D$. So $(x', z') \in D \circ C$. Since x' is generic, rk $(x', z') \geq 1$. Thus $(x', z') \in X \times Z$ is a generic element of D.

We have shown that every analytic neighborhood of (x, z) contains a generic element of $D \circ C$. It follows that $(x, z) \in (D \circ C)^P$, as desired. □

Now let $\mathcal{F} = \{F_a\}_{a \in A}$ be an almost faithful, generically almost pure family of non-trivial plane curves, and assume rk $A \geq 1$. By Theorem 5.0.1, rk $A = 1$. Shrinking and adding parameters if necessary, we may assume A is strongly minimal and $\emptyset$-definable.

Let

$$\mathcal{G} = \{G_{(a,b)}\}_{(a,b) \in A^2}$$

be the family of compositions $\mathcal{F} \circ \mathcal{F}$ – that is, for each (a, b) set $G_{(a,b)} = F_b \circ F_a$. By Lemma 5.5.9, the family $\mathcal{G}$ is also a generically almost pure family of non-trivial plane curves.

Say that a pair $(a, b) \in A^2$ is *bad* if there are only finitely many $(a', b') \in A^2$ for which $G_{(a,b)} \cap G_{(a',b')}$ is infinite. Let $Z \subset A^2$ be the set of bad points. Note that Z is $\mathcal{M}$-definable.

Lemma 5.5.10. *rk $Z \leq 1$.*

Proof. Assume instead that rk $Z = 2$. Consider the subfamily $\mathcal{G}_Z$ of curves in $\mathcal{G}$ indexed by elements of Z. It follows by definition of badness that $\mathcal{G}_Z$ is almost faithful. Since rk $Z = 2$, Z is generic in A^2, and therefore every generic curve in $\mathcal{G}_Z$ is almost pure. This contradicts Theorem 5.0.1. □

Now let $(a, b) \in A^2$ be generic. By the lemma, (a, b) is not bad, and so there are infinitely many curves in $\mathcal{G}$ whose intersections with $G_{(a,b)}$ are infinite. Let $(a', b') \in A^2$ be a generic index among these curves: that is, let (a', b') be such that rk $(a', b'/a, b) \geq 1$ and $G_{(a,b)} \cap G_{(a',b')}$ is infinite. Let S be any strongly minimal component of the rank one set $G_{(a,b)} \cap G_{(a',b')}$, and let c be a code of S. Note, by Lemma 4.2.2, that c is algebraic over both (a, b) and (a', b'), since S is a component of both $G_{(a,b)}$ and $G_{(a',b')}$. Furthermore, since F_a and F_b are non-trivial plane curves, so is $G_{(a,b)}$, and thus so is S.

Let (x, z) be a generic element of S over (a, b, a', b', c). Then

$$(x, z) \in G_{(a,b)} \cap G_{(a',b')},$$

so there are $y, y' \in M$ such that $(x, y) \in F_a$, $(y, z) \in F_b$, $(x', y') \in F_{a'}$, and $(y', z') \in F_{b'}$.

Lemma 5.5.11. *Each of x, y, y', and z is generic in M over (a, b, a', b', c). Further, (x, y) is generic in F_a over (a, b, a', b', c), (x, y') is generic in $F_{a'}$ over (a, b, a', b', c), (y, z) is generic in F_b over (a, b, a', b', c), and (y', z) is generic in $F_{b'}$ over (a, b, a', b', c).*

Proof. Since S is a non-trivial plane curve, both projections $S \to M$ are finite-to-one everywhere. Thus x and z are interalgebraic over c. In particular, since

$$\text{rk } (x, z/a, b, a', b', c) = 1,$$

we get that

$$\text{rk } (x/a, b, a', b', c) = \text{rk } (z/a, b, a', b', c) = 1.$$

Similarly, since F_a and $F_{a'}$ are non-trivial, x and y are interalgebraic over a, and x and y' are interalgebraic over a'. We get that

$$\text{rk } (y/a, b, a', b', c) = \text{rk } (y'/a, b, a', b', c) = 1.$$

Finally, it is now clear that each of (x, y), (x, y'), (y, z), and (y', z) has rank at least 1 over (a, b, a', b', c), and so is generic in the relevant curve over (a, b, a', b', c). □

Now let D be a strongly minimal component of F_a containing (x, y), D' a strongly minimal component of $F_{a'}$ containing (x, y'), E a strongly minimal component of F_b containing (y, z), and E' a strongly minimal component of $F_{b'}$ containing (y', z). Let d, d', e, and e' be codes of each of these sets, respectively.

Lemma 5.5.12. *a and d are interalgebraic, as are a' and d', b and e, and b' and e'.*

Proof. It follows from Lemma 4.2.2 that d is algebraic over a. On the other hand, since $\mathcal{F}$ is almost faithful, D can only appear as a strongly minimal component of finitely many curves in $\mathcal{F}$. This shows that a is algebraic over d.

Thus a and d are interalgebraic. The other three statements are similar. □

It follows by the previous two lemmas that (x, y) is generic in D over (d, e, d', e', c), (x, y') is generic in D' over (d, e, d', e', c), (y, z) is generic in E over (d, e, d', e', c), and (y', z) is generic in E' over (d, e, d', e', c). Moreover, we can now show:

Lemma 5.5.13. *Each of a, b, and c is algebraic over the other two. Moreover the same holds for a', b', and c.*

Proof. By the previous lemma, it suffices to prove the analagous statements for the triples (d, e, c) and (d', e', c). But these statements follow immediately from Lemma 4.2.2. □

Now we are ready to verify:

Lemma 5.5.14. *(a, b, c) forms a non-trivial configuration.*

Proof. We know by choice of (a, b) that

$$\text{rk } a = \text{rk } b = 1 \text{ and rk } (a, b) = 2.$$

Now by the previous lemma:

- rk $(c/a, b) = 0$, and so by additivity

$$\text{rk } (a, b, c) = 2.$$

- rk $(a/b, c) = 0$, and so by additivity

$$\text{rk } (b, c) = 2.$$

- rk $(b/a, c) = 0$, and so by additivity

$$\text{rk } (a, c) = 2.$$

It remains to show that rk $c = 1$. For this, we note:

- Since rk $(a/b, c) = 0$, it follows that rk $(a, b/c) \leq 1$.
- Since rk $(a', b'/a, b) \geq 1$, it follows that

$$\text{rk } (a, b, a', b', c) \geq 3.$$

  But by the previous lemma we have rk $(c/a', b') = 0$, so rk $(a', b', c) \leq 2$, and thus

$$\text{rk } (a, b/c) \geq \text{rk } (a, b/a', c', b) \geq 1.$$

By the two points above, rk $(a, b/c) = 1$. Since rk $(a, b, c) = 2$, this implies rk $c = 1$, and we are done. □

Finally, we can now conclude:

Lemma 5.5.15. *(a, b, c, x, y, z) forms a one-dimensional group configuration.*

Proof. We have just shown that rk $c = 1$. Since a and d are interalgebraic, rk $d = 1$. Similarly, rk $e = 1$. So each of d, e, and c is a rank 1 code for a non-trivial irreducible plane curve, and the pairs (x, y), (y, z), and (x, z) are respectively generic elements of those curves. By Lemma 5.5.2, we conclude that each of (d, x, y), (e, y, z), and (c, x, z) forms a non-trivial configuration. Again since a and d are interalgebraic and b and e are interalgebraic, we conclude that each of (a, x, y), (b, y, z), and (c, x, z) forms a non-trivial configuration. Additionally, from the previous lemma we know that (a, b, c) forms a non-trivial configuration.

It remains only to show that rk $(x, y, z) = 3$. But since x is generic over (a, b, c), and rk $(a, b) = 2$, it follows that

$$\text{rk } (a, b, c, x, y, z) \geq 3.$$

On the other hand, we have just concluded that

$$\text{rk } (a/x, y) = \text{rk } (b/y, z) = \text{rk } (c/x, z) = 0,$$

and thus

$$\text{rk } (a, b, c/x, y, z) = 0.$$

It follows that rk $(x, y, z) = 3$, and we are done. □

So, since we have found a one-dimensional group configuration in the strongly minimal structure $\mathcal{M}$, it follows by Fact 5.5.5 that $\mathcal{M}$ interprets a strongly minimal group.

□

Chapter 6

Unimodularity Revisited: An O-minimal Variant

In this chapter we will generalize the results of Chapter 4 to strongly minimal structures interpretable in o-minimal fields. Namely, we will define a notion of 'almost pure' for definable sets in o-minimal fields, and prove unimodularity for higher dimensional strongly minimal reducts of affine space, assuming that all sufficiently generic plane curves are almost pure. In doing this, we obtain a partial result toward the O-minimal Restricted Trichotomy Conjecture mentioned in the introduction. We expect future research to verify this almost purity condition under the assumption of a group operation – which, if successful, would for example prove the conjecture for expansions of the n-dimensional additive group for each n.

The main challenge of the generalization presented in this chapter is that o-minimal structures do not admit a canonical notion of irreducible components with the properties we need. We would like to argue similarly to Chapter 4 that every top dimensional 'component' of a plane curve is almost contained in a connected covering, under any projection to the universe, of a generic subset with the same fundamental group as the universe. However, such an argument only works if each such 'component' projects almost surjectively – a property which fails easily for any reasonable notion of components in o-minimal structures. Additionally, since we want to compare the degrees of different projections, such an argument only works if the 'components' are canonical and intrinsic to the plane curve; in general, in o-minimal structures we only have a non-unique 'cell decomposition,' which is insufficient for our purposes.

Now the main innovation of this chapter is that, under a suitable notion of almost purity, such a desired notion of 'components' can be recovered instrinsically from a definable set. In section 8 we carry out this process, along the way proving various natural facts about the behavior of components. Our conclusion will be that the components of a plane curve behave in a similar way to the irreducible components of algebraically constructible sets; then in section 9 we are able to carry out the unimodularity argument as in Chapter 4.

6.1 O-Minimality: Review of Basic Definitions and Facts

In this section we briefly mention the important notions from o-minimality that we need. For a full account of the subject, the reader could consult [7].

Definition 6.1.1. Let $\mathcal{R} = (R, <, ...)$ be a structure in a language containing the binary relation symbol $<$, such that the pair $(R, <)$ forms a linear order. Then $\mathcal{R}$ is *o-minimal* if every definable set $D \subset R$, over any parameters, is a finite union of points and open intervals.

Remark 6.1.2. Note that intervals have specified endpoints, and this is a stronger condition than convexity. So, more precisely, for every definable $D \subset R$ there are

$$a_1, ..., a_m, b_1, ..., b_m, c_1, ..., c_n \in R \cup \{\pm\infty\}$$

such that

$$D = \bigcup_{i=1}^{m}(a_i, b_i) \cup \{c_1, ..., c_n\}.$$

Definition 6.1.3. An *o-minimal field* is an o-minimal structure $\mathcal{R} = (R, +, \cdot, <, ...)$, in a language containing the binary operations $+$ and $\cdot$, such that $(R, +, \cdot, <)$ is an ordered field.

It is a fact that every o-minimal field is real closed ([40], Theorem 2.3). However, there may be a plethora of additional structure added to a real closed field while preserving o-minimality. The most common expansions are the real field with the exponential function [49], and the real field augmented with the restrictions of all analytic functions to all compact intervals (see [6], [5], and [9]).

Note also that there are many o-minimal fields with universes other than the real numbers – in fact, by quantifier elimination [47] every real closed field is o-minimal as a pure ordered field. Now even if our field R is not the field of real numbers, one still has the order topology on R, and the inherited product topology on R^n for each n. In fact, these topologies can be generated in the usual way via the Pythagorean 'distance' formula for two points – i.e. the square root of the sum of squares of their coordinate differences – which, aside from taking values outside the reals, turns out to possess all other relevant properties of a metric. This distance function is definable, which allows one to uniformly express various analytic and topological notions – e.g. limits, continuity, and closure – using first order formulas. Thus, when working with an o-minimal field, one can perform quite a bit of topological analysis on definable sets.

Of particular importance is the notion of 'definable compactness' [36] – an abstract version of compactness that, at least in the presence of a field structure, enjoys many analagous properties to compactness in $\mathbb{R}^n$. For example, restricted to the definable subsets of the affine space R^n over an ordered field, the definably compact definable sets are precisely those which are closed and bounded; moreover, definably compact sets are closed under images of continuous definable functions.

When defining strong minimality it is crucial to require that the finite/cofinite property pass to elementary extensions; interestingly, this is not the case for o-minimal structures:

Fact 6.1.4 (Knight, Pillay, Steinhorn [29]). *Every elementary extension of an o-minimal structure is o-minimal.*

Thus, for example, when working with an o-minimal field $\mathcal{R}$, we may assume $\mathcal{R}$ is κ-saturated for any fixed κ.

O-minimality is often seen as a natural setting for 'tame geometry' – that is, real analysis and topology without pathological counterexamples. Indeed, o-minimal structures possess various 'tameness' properties for definable sets and functions with respect to the topology mentioned above (see [7]). For example, every definable function in one variable has unique left and right limits at every point, and every definable function on any definable set is continuous 'almost everywhere' in a precise sense.

The most important geometric fact about o-minimal structures is the Cell Decomposition Theorem (see [29]). Similarly to the case of strongly minimal structures, a strong restriction on definable sets in one variable leads to a good structure theory for definable sets in many variables – roughly, we get a similar statement to the definition of o-minimality, but with points and intervals replaced by 'cells.' We briefly summarize below. For convenience, we identify the 0th cartesian power of a set as a single point.

Definition 6.1.5. Given an o-minimal structure on a set R, we define the class of *d-cells in R^n*, a familly of definable subsets of R^n, inductively as follows:

1. A 0-cell in R^n, for any n, is a single point.
2. There are no d-cells in R^0 unless $d = 0$.
3. If $d, n \geq 1$ and $C \subset R^{n-1}$ is a d-cell, then the graph of any continuous definable function $f : C \to R$ is a d-cell in R^n.
4. If $d, n \geq 1$ and $C \subset R^{n-1}$ is a $(d-1)$-cell, then the following are all d-cells in R^n:
   a) $C \times R$.
   b) $\{(x, y) \in C \times R : y < f(x)\}$ and $\{(x, y) \in C \times R : y > f(x)\}$, where $f : C \to R$ is any continuous definable function.
   c) $\{(x, y) \in C \times R : f(x) < y < g(x)\}$, where $f, g : C \to R$ are continuous definable functions with $f(x) < g(x)$ for all $x \in C$.

Remark 6.1.6. It follows from the construction that the class of d-cells in R^n is definable in families; moreover, in the presence of a field structure, every d-cell is definably homeomorphic to R^n.

Now the main fact about cells is the following result of Knight, Pillay, and Steinhorn [29]:

Fact 6.1.7 (Cell Decomposition). *Let $\mathcal{R} = (R, <, ...)$ be an o-minimal structure. Then every definable subset of R^n, for any n, is a disjoint union of finitely many cells.*

One can then extrapolate a dimension theory for definable sets:

Definition 6.1.8. Given an o-minimal structure $\mathcal{R} = (R, <, ...)$, and a definable set $D \subset R^n$, the *dimension of* D is the largest d such that that some (equivalently any) cell decomposition of D includes a d-cell.

This dimension notion enjoys many of the same properties as its counterpart in strong minimality: it is automorphism invariant, definable in families, preserved under finite-to-one maps, respects finite unions and products, and so on. The only missing ingredient from the strongly minimal dimension is the degree function: indeed, a set of dimension n in n-space (for example) may be just a ball around a point – thus it is quite easy to partition R^n into infinitely many pairwise disjoint sets of dimension n.

Using the Cell Decomposition Theorem, it is possible to deduce various facts about the geometric behavior of definable sets and functions (see [7]). For example, if D is any non-empty definable set, then the frontier of D, $\overline{D} - D$, has smaller dimension than D. Similarly, the set of points at which a definable function $f : D \to E$ is not continuous is of dimension less than $\dim D$.

One also has a dimension notion for points and types, defined in the same way as for strongly minimal structures: given an element $a \in R^n$ and a set A, we define $\dim(a/A)$ to be the smallest dimension of an A-definable set containing a. If $a \in D$, D is A-definable, and $\dim(a/A) = \dim D$, then we say that a is *generic in* D *over* A. Again, theses notions enjoy many of the same properties as their counterparts from strong minimality, the notable exception being stationarity. The main property we need is additivity, which does hold. Namely (see [39], Lemma 1.2):

Fact 6.1.9. *Given tuples a and b, and a set A, we have* $\dim(a, b/A) = \dim(a/A) + \dim(b/Aa)$.

Finallly, we mention the following fact, which follows from the 'Trivialization' Theorem (see [7], Chapter 9, Theorem 1.2 and Remark 2.1). Roughly, it says that whether two definable sets are definably homeomorphic is a definable property:

Fact 6.1.10. *Let $\mathcal{R} = (R, +, \cdot, <)$ be an o-minimal field, and let $S \subset R^{m+n}$ be a definable set. Then among the fibers $S_x \subset R^n$, for $x \in R^m$, there are only finitely many occurring sets up to definable homeomorphism. Moreover, these finitely many definable homeomorphism types induce a partition of R^m into finitely many definable sets.*

For example, in any definable family of sets, the indices of sets which are definably homeomorphic to a fixed R^n is definable. This is the main application of Fact 6.1.10 that we will make.

6.2 Some Algebraic Topology

The main theorem in this chapter concerns structures definable from o-minimal fields with simply connected universes. Our strategy, similarly to Chapter 4, will be to study covering spaces in order to bound the ratios of plane curves. One might worry, then, that we are conducting this work over an abstract ordered field as opposed to the real field. The main purpose of this section is to point out that the usual theory of fundamental groups, to the extent that we need, transfers to arbitrary o-minimal fields.

We note before proceeding that the definable fundamental group (also called the o-minimal fundamental group) in o-minimal structures has been well studied – as has, to a lesser extent, the Galois correspondence with definable covers. We refer the reader to section 2 of [10] for more details. Our goal in this section will largely be to present the basic definitions, and then cite [10] for the necessary portion of the Galois correspondence. We do work out a couple proofs that seem absent in the literature, though it is poossible that this work is also well known.

Convention 6.2.1. Throughout this section, we fix $\mathcal{R} = (R, +, \cdot, <, ...)$, an o-minimal field. The notation $[0, 1]$ refers to the closed interval in the sense of $\mathcal{R}$ – that is, $\{x \in R : 0 \leq x \leq 1\}$.

We begin with some basic definitions:

Definition 6.2.2. Given a definable set D, a *definable path in* D is a continuous definable function $\gamma : [0, 1] \to D$. If $\gamma(0) = \gamma(1) = x_0$, then γ is a *definable loop in* D *at* x_0.

Definition 6.2.3. The definable set D is *definably path connected* if for all $a, b \in D$, there is a definable path γ in D with $\gamma(0) = a$ and $\gamma(1) = b$.

Remark 6.2.4. Note that R^n is definably path connected for all n: indeed, for any a and b the function $\gamma(x) = a + x \cdot (b - a)$ gives a definable path from a to b.

Definition 6.2.5. Let γ_1 and γ_2 be definable paths in D, with $\gamma_1(0) = \gamma_2(0) = a$ and $\gamma_1(0) = \gamma_2(1) = b$. Then γ_1 and γ_2 are *definably homotopic in* D if there is a continuous definable function $h : [0, 1]^2 \to D$ with $h(0, y) = \gamma_1(y)$, $h(1, y) = \gamma_2(y)$, $h(x, 0) = a$, and $h(x, 1) = b$, for all $x, y \in [0, 1]$.

It is clear that definable homotopy is an equivalence relation on definable paths, since the usual proof in algebraic topology only relies on 'reparametrizing' homotopies using the field operations – a process which is certainly definable in our field. Thus, fixing a definable set D and a point $x_0 \in D$, we can take the set of definable loops at x_0 in D, modulo definable homotopy. This defines a group under the usual concatenation operation: again, the verification only relies on reparametrizing with the field operations, which our field can handle. Thus we can define:

Definition 6.2.6. Given a definable set D and an element $x_0 \in D$, the *definable fundamental group* of D at x_0, denoted $\pi_1(D, x_0)$, is the group of definable loops in D at x_0 modulo definable homotopy.

It is clear that, if D is definably path connected, the isomorphism class of $\pi_1(D, x_0)$ does not depend on the point x_0. In this case we abuse notation and write $\pi_1(D)$. Thus the following is well defined:

Definition 6.2.7. The definable set D is *definably simply connected* if it is definably path connected and has trivial definable fundamental group.

Matching our intuition from usual fundamental groups, we note:

Lemma 6.2.8. *R^n is definably simply connected for all n.*

Proof. It was noted in Remark 6.2.4 that R^n is definably path connected, so we need only show that R^n has trivial definable fundamental group. Without loss of generality, we consider the definable fundamental group at $0 \in R^n$. Now it suffices to show that every definable loop at 0 is definably homotopic to the constant loop; but indeed, if γ is any definable loop at 0, the function $h(x, y) = x \cdot \gamma(y)$ is such a definable homotopy. □

More generally, we show the following:

Convention 6.2.9. For ease of notation, we assume the empty set has dimension $-\infty$.

Lemma 6.2.10. *Let $T \subset R^n$ be a definable set.*

1. *If $\dim T \leq n - 2$ then $R^n - T$ is definably path connected.*
2. *If $\dim T \leq n - 3$ then $R^n - T$ is definably simply connected.*

Proof. First, let $f : [0, 1] \to R$ be any function with the following properties:

- f is continuous and definable.
- $f(0) = f(1) = 0$.
- $f(x) \neq 0$ for all $x \neq 0, 1$.

For example, $f(x) = x^2 - x$ will do. Now we proceed with the proof:

1. Assume $\dim T \leq n - 2$, and let $a, b \in R^n - T$. Then since R^n is definably path connected, there is a definable path γ from a to b in R^n. Now we claim that there is an element $v \in R^n$ such that the function

$$\gamma_v(x) = \gamma(x) + v \cdot f(x)$$

defines a path from a to b in $R^n - T$. It is clear that for all v we have $\gamma_v(0) = a$, $\gamma_v(1) = b$, and γ_v is continuous and definable; so the goal is to choose v so that the image of γ_v is disjoint from T. Equivalently, since $a, b \notin T$ by assumption, we desire

$$\gamma_v((0, 1)) \cap T = \emptyset.$$

Let

$$A = \{(x, v) \in (0, 1) \times R^n : \gamma_v(x) \in T\}.$$

We note:

Claim 6.2.11. *A is in definable bijection with* $(0, 1) \times T$.

Proof. There is a natural map $A \to (0, 1) \times T$ given by

$$(x, v) \mapsto (x, \gamma_v(x)) = (x, \gamma(x) + v \cdot f(x)).$$

Its inverse is given by

$$(x, t) \mapsto \left(x, \frac{t - \gamma(x)}{f(x)}\right),$$

which is well-defined on $(0, 1) \times T$ since f is nonzero on $(0, 1)$. □

Now it follows that

$$\dim A = 1 + \dim T \leq n - 1.$$

Thus the projection $A \to R^n$ cannot be surjective. So we need only choose some $v \in R^n$ which is not in the image of A; then by definition of A it follows that $\gamma_v((0, 1))$ is disjoint from T.

2. Now assume $\dim T \leq n - 3$. By (1) $R^n - T$ is definably path connected. To show that $R^n - T$ is definably simply connected, let $x_0 \in R^n - T$ and let γ be a loop in $R^n - T$ at x_0. Then by Lemma 6.2.8, there is a definable homotopy $h : [0, 1]^2 \to R^n$ with the constant loop at x_0. We seek such a definable homotopy which is contained in $R^n - T$.

   Our strategy is essentially identical to (1). For each $v \in R^n$ let $h_v : [0, 1]^2 \to R^n$ be the map given by

   $$h_v(x, y) = h(x, y) + v \cdot f(x).$$

   Then each h_v is a definable homotopy from γ to the constant loop. We claim that there is some $v \in R^n$ so that the image of h_v is disjoint from T – or equivalently, since this is automatic on boundary points of $[0, 1]^2$,

   $$h_v((0, 1)^2) \cap T = \emptyset.$$

   To see this, let

   $$B = \{(x, y, v) \in (0, 1)^2 \times R^n : h_v(x, y) \in T\}.$$

   Then, as in (1), it suffices to find an element $v \in R^n$ which is not in the image of the projection $B \to R^n$. Now we note:

Claim 6.2.12. *B is in definable bijection with $(0,1)^2 \times T$.*

Proof. The map

$$(x, y, v) \mapsto (x, y, h_v(x, y))$$

maps B to $(0,1)^2 \times T$. Its inverse is given by

$$(x, y, t) \mapsto \left(x, y, \frac{t - h(x, y)}{f(x)}\right).$$

As in (1), this is well-defined because f is nonzero on $(0,1)$. □

Now we conclude that

$$\dim B = 2 + \dim T \leq n - 1.$$

So, as in (1), the projection $B \to R^n$ cannot be surjective, and we are done.

□

Next we discuss covering spaces.

Definition 6.2.13. Let D and E be definable sets, and $f : D \to E$ a continuous definable function which is everywhere k-to-1 for some $k \in \mathbb{Z}^+$. If $U \subset E$ is definable, we say that f *definably trivializes over* U if there are continuous definable functions $g_1, ..., g_k : U \to D$ such that $f(g_k(u)) = u$ holds for all i and u, and the sets $g_i(U)$ are pairwise disjoint.

Remark 6.2.14. Note that the above definition implies $f^{-1}(U)$ is the disjoint union of the sets $g_i(U)$. In fact, it is implied that the topology on $f^{-1}(U)$ is the disjoint union topology coming from the $g_i(U)$: indeed, since there are only finitely many g_i, it suffices to check that

$$\overline{g_i(U)} \cap g_j(U) = \emptyset$$

for $i \neq j$; but if some x belonged to this intersection, the continuity of g_i and g_j would imply $g_i(f(x)) = g_j(f(x)) = x$, contradicting Definition 6.2.13.

Definition 6.2.15. Let D and E be definable sets, and $f : D \to E$ a continuous definable function which is everywhere k-to-1 for some $k \in \mathbb{Z}^+$. Then f is a *definable k-covering* if there is a cover of E by relatively open definable subsets over which f definably trivializes.

Remark 6.2.16. Definition 6.2.15 is not to be confused with the Trivialization Theorem for o-minimal fields ([7], Chapter 9). Importantly, we insist that the sets in the cover of E be open, while in the Trivialization Theorem they need only be definable.

Definable covers were studied in [10] under the assumption that one can cover the target by *finitely many* open sets over which the map trivializes. Now the main observation of the present section is that there is no harm in generalizing this to allow *arbitrary* open covers, as suggested by Definition 6.2.15. That is, we show below that if a map definably trivializes over any open cover, it definably trivializes over a finite open cover:

Proposition 6.2.17. *Let D and E be definable sets, and $f : D \to E$ a definable k-covering. Then there is a finite cover of E by relatively open definable susbsets over which f definably trivializes.*

Proof. We begin with:

Lemma 6.2.18. *For each pair of natural numbers n and k, there is a finite set $T_1, \ldots, T_L$ of linear maps from R^n to R, such that for any k nonzero vectors $v_1, \ldots, v_k \in R^n$, there is some T_j such that $T_j(v_i) \neq 0$ for all $i \leq k$.*

Proof. This is a first order statement, so we may assume $\mathcal{R}$ is $\aleph_0$-saturated. Let

$$L = k(n-1) + 1,$$

and let A be an $L \times n$ matrix whose entries are algebraically independent over $\mathbb{Q}$. Then the desired maps are given by the L rows of A.

To prove this works, take any nonzero vectors $v_1, \ldots, v_k$. We claim that each v_i can be sent to 0 by at most $n-1$ of the rows of A. Indeed, otherwise we would obtain an $n \times n$ submatrix B of A, with v_i in the null space of B. But by the independence of all entries of A, the matrix B is invertible, meaning $v_i = 0$, a contradiction.

Now since each v_i is sent to 0 at most $n-1$ times, and since there are $L > k(n-1)$ rows of A, it follows that some row of A sends no v_i's to 0.

□

Corollary 6.2.19. *For each pair of natural numbers n and k, there is a finite set $T_1, \ldots, T_L$ of linear maps from R^n to R, such that for any k distinct vectors $v_1, \ldots, v_k \in R^n$, there is some T_j such that the k vectors $T_j(v_1), \ldots, T_j(v_k)$ are also distinct.*

Proof. Apply the above lemma to the differences of the v_i's. That is, set $k' := k(k-1)$ and apply the above lemma. □

Corollary 6.2.19 says, roughly, that we can find finitely many linear maps which can be injective with respect to any k points in R^n. Our strategy will be to use these maps to order the preimages of points in E 'continuously,' thereby obtaining the desired trivializations.

We now define:

Definition 6.2.20. Let $v_1, \ldots, v_k$ be a tuple of distinct vectors in R^n. We say that a definable linear order $<$ on R^n is *continuous at* $(v_1, \ldots, v_k)$ if there are definable open neighborhoods U_i of each v_i, such that for any other tuple $(w_1, \ldots, w_k)$ in $U_1 \times \cdots \times U_k$, the points $w_1, \ldots, w_k$ come in the same order as $v_1, \ldots, v_k$.

Remark 6.2.21. Note that Definition 6.2.20 is equivalent to the assertion that the map from $(R^n)^k$ to $(R^n)^k$, which outputs the input vectors in increasing order, is continuous at $(v_1, \ldots, v_k)$.

Remark 6.2.22. Note also that the order of $v_1, \ldots, v_k$ does not matter in Definition 6.2.20; thus it is well-defined to say that $<$ is continuous at the *k-element set* $\{v_1, \ldots, v_k\}$.

Now the main point of the proof of Proposition 6.2.17 is the following:

Lemma 6.2.23. *For each pair of natural numbers n and k, there are finitely many definable linear orders $<_1, \ldots, <_L$ on R^n, such that for any k distinct vectors $v_1, \ldots, v_k \in R^n$, one of the $<_j$ is continuous at $(v_1, \ldots, v_k)$.*

Proof. Let $T_1, \ldots, T_L$ be linear maps as in Corollary 6.2.19. For each j choose a linear order $<_j$ which orders points first by their images under T_j, and then uses the lexicographic order on R^n as a tiebreaker. Thus each $<_j$ is definable; indeed, the map T_j is given by a matrix of elements of R, and the lexicographic order is easily expressed in terms of the relation $<$.

Now for any distinct $v_1, \ldots, v_k$, we can choose T_j which is injective on $\{v_1, \ldots, v_k\}$. Then it follows immediately that $<_j$ is continuous at $(v_1, \ldots, v_k)$. □

Finally, we return to the proof of Proposition 6.2.17. Let n be such that R^n is the ambient space containing D. Let $<_1, \ldots, <_L$ be linear orders as in Lemma 6.2.23, and for each j let U_j be the set of points $e \in E$ such that $<_j$ is continuous at the k-element set $f^{-1}(e)$. By the construction above it is clear that U_j is definable and relatively open in E. To see that f definably trivializes on U_j, let $g_i^j(u)$, for each $u \in U_j$, be the ith smallest preimage of u according to $<_j$. Then g_i^j is definable. By the continuity of f, and the fact that $<_j$ is continuous at $f^{-1}(u)$ for $u \in U_j$, it is clear that g_i^j is continuous. This completes the proof. □

In light of Proposition 6.2.17, the results of [10] still hold using our definition of definable covers. That is, in particular:

- The natural generalizations of unique path lifting and unique homotopy lifting hold ([10], Propositions 2.6 and 2.7).
- A definable k-covering map $f : D \to E$ of definably path connected definable sets (with compatible base points) identifies $\pi_1(D)$ as a subgroup of $\pi_1(E)$ ([10], Corollary 2.8).
- If $f : D \to E$ is a definable k-covering of definably path connected definable sets, and $\pi_1(D)$ is normal in $\pi_1(E)$, then the quotient group $\pi_1(E)/\pi_1(D)$ is canonically isomorphic to the group $Aut(D/E)$ of 'definable deck transformations' – that is, definable homeomorphisms $\pi : D \to D$ satisfying $f \circ \phi = f$ ([10], Proposition 2.10).

For our purposes, we can now conclude the main result of this section:

Proposition 6.2.24. *Let D and E be definably path connected definable sets, and $f : D \to E$ a definable k-covering. If E is definably simply connected, then $k = 1$.*

Proof. Since $\pi_1(E)$ is trivial, it is clear that $\pi_1(D)$ is normal and of index 1 in $\pi_1(E)$. Thus, by the above remarks, the group $Aut(D/E)$ is trivial. In other words, there are no definable deck transformations other than the identity map on D.

Now suppose toward a contradiction that $k > 1$. Fix $d_1 \neq d_2 \in D$ with

$$f(d_1) = f(d_2) = e_0.$$

Then by the proof of [10], Proposition 2.10, there is a definable deck transformation sending d_1 to d_2. That is, since D is definably path connected, there is a definable path γ from d_1 to d_2 in D. Thus the image $f \circ \gamma$ of γ is a definable loop at e_0 in E. Now using e_0 as the basepoint in E, the image of $f \circ \gamma$ in $Aut(D/E)$ (first quotienting by homotopy to get to $\pi_1(E)$, then quotienting by $\pi_1(D)$, then passing through the isomorphism to $Aut(D/E)$) still sends d_1 to d_2 by construction.

So if $k > 1$ then we can find non-trivial deck transformations, which gives a contradiction.

□

6.3 Our Setting

In this short section we summarize the setting we will assume for the rest of the chapter. Our aim is to show that certain higher dimensional strongly minimal reduct structures interpretable in o-minimal fields are unimodular. Note that, similarly to algebraically closed fields, o-minimal fields eliminate imaginaries (by [7], Chapter 6, Proposition 1.2); thus we need only study *definable* structures – those whose universe and definable sets are definable from the original o-minimal field. We thus make the following definition:

Convention 6.3.1. For the rest of this chapter, we fix an o-minimal field $\mathcal{R} = (R, +, \cdot, <, ...)$.

Definition 6.3.2. A *strongly minimal reduct structure* of $\mathcal{R}$ is a strongly minimal structure $\mathcal{M} = (M, ...)$, whose universe M is an $\mathcal{R}$-definable subset of R^n for some n, such that every $\mathcal{M}$-definable subset of any cartesian power of M is already definable in $\mathcal{R}$.

Convention 6.3.3. For the rest of this chapter, unless stated otherwise, we assume that $\mathcal{M} = (M, ...)$ is a fixed strongly minimal reduct structure of the o-minimal field $\mathcal{R}$. Throughout, the term 'definable' will refer to definability in the structure $\mathcal{R}$, whereas definability in $\mathcal{M}$ will be referred to as '$\mathcal{M}$-definable.' As in the previous chapters, we have two dimension notions for points and types; we will use the notation dim for the o-minimal dimension of $\mathcal{R}$, and rk for the strongly minimal dimension of $\mathcal{M}$. The word 'generic,' when applied to points and types, always refers to genericity in the sense of $\mathcal{R}$; we will write '$\mathcal{M}$-generic' for genericity in the sense of $\mathcal{M}$.

As stated in the introduction, it was conjectured by Peterzil that, in the situation outlined above, $\mathcal{M}$ must either be locally modular or interpret the algebraically closed field $R[i]$. It is known (see Corollary 6.4.10, though this result is not new) that $\mathcal{M}$ can only interpret $R[i]$ if

$$\dim M = \dim R[i] = 2,$$

i.e. the universe of M is an $\mathcal{R}$-definable set of dimension 2. Past authors have verified the conjecture in the case $\dim M = 1$ [19], for expansions of the group $(\mathbb{C}, +)$ by a function symbol [18], and more recently for all expansions of 2-dimensional groups [12]. Our aim in this chapter is to present a partial result in the case that $\dim M \geq 3$. More precisely, we will define a generalization of almost purity to this setting, and show that if $M = R^n$ for some $n \geq 3$, and all sufficiently generic plane curves are almost pure, then $\mathcal{M}$ is unimodular.

Convention 6.3.4. For the rest of this chapter, we assume the universe M is equal to R^n for some $n \geq 3$.

6.4 Dimension and Genericity

We proceed by making a few initial observations analogous to those made in Chapter 3. To start, we discuss dimension and genericity, with the aim of showing a direct analogy to Corollary 3.1.11. The following notions are all essentially identical to their counterparts in Chapter 3.

Definition 6.4.1. Let D and E be definable sets with $D \subset E$. Then D is *$\mathcal{R}$-generic in E* if $\dim D = \dim E$. If in addition $\dim(E - D) < \dim E$, we say that D is *$\mathcal{R}$-fully generic in E*.

The following is immediate and useful:

Lemma 6.4.2. *Let C, D, and E be definable sets with $C, D \subset E$. If C is $\mathcal{R}$-generic in E and D is $\mathcal{R}$-fully generic in E, Then $C \cap D$ is $\mathcal{R}$-generic in E.*

Proof. It suffices to note that

$$C = (C \cap D) \cup (C - D) \subset (C \cap D) \cap (E - D).$$

So, since $\dim C = \dim E$, one of $C \cap D$ and $E - D$ has the same dimension as E. But our assumption implies $\dim(E - D) < \dim E$, so it follows that $\dim(C \cap D) = \dim E$. □

We will often say that a property holds for 'almost all' elements of a set D. This will mean that the solution set of that property in D, is a fully generic definable subset of D.

The following equivalence is straightforward and useful:

Lemma 6.4.3. *Let $A \subset R$, $a \in R^m$ for some m, and $E \subset R^m$ a definable set containing a. Then a is generic in E over A if and only if for every fully generic A-definable set $D \subset E$, we have $a \in D$.*

Proof. Rewriting the definitions shows that a set $D \subset E$ is fully generic if and only if $E - D$ is not generic. The lemma is now immediate. □

We now define the following analogously to Chapter 3:

Definition 6.4.4. Let D and E be definable sets.

1. D is $\mathcal{R}$-*almost contained in* E if $D \cap E$ is $\mathcal{R}$-fully generic in D.
2. D and E are $\mathcal{R}$-*almost equal* if they are $\mathcal{R}$-almost contained in each other.
3. D and E are $\mathcal{R}$-*almost disjoint* if $D \cap E$ is $\mathcal{R}$-non-generic in both D and E.

Definition 6.4.5. Let D and E be definable sets, and let $f : D \to E$ be a definable function. Then:

1. f is $\mathcal{R}$-*almost surjective* if $f(D)$ is $\mathcal{R}$-fully generic in E.
2. f is $\mathcal{R}$-*almost finite-to-one* if the union of all finite fibers of f is $\mathcal{R}$-fully generic in D (or equivalently, the union of all infinite fibers is $\mathcal{R}$-non-generic in D).
3. f is $\mathcal{R}$-*almost injective* if the union of all fibers of size > 1 is $\mathcal{R}$-non-generic in D.
4. f is $\mathcal{R}$-*almost bijective* if f is both $\mathcal{R}$-almost injective and $\mathcal{R}$-almost surjective.

By an identical argument to that presented in Lemma 3.1.8, we have:

Lemma 6.4.6. *Let D and E be definable sets, and let $f : D \to E$ be a definable function.*

1. *If f is $\mathcal{R}$-almost surjective, then* $\dim D \geq \dim E$.
2. *If f is $\mathcal{R}$-almost finite-to-one, then* $\dim D \leq \dim E$.
3. *If f is $\mathcal{R}$-almost surjective and $\mathcal{R}$-almost finite-to-one, then* $\dim D = \dim E$.

We conclude the following, which is the main goal of this section:

Lemma 6.4.7. *If D is any non-empty definable set, then*

$$\dim D = n \cdot rk\ D.$$

Proof. We induct on rk D. The case rk $D = 0$ is clear, since then D is finite and non-empty, and so $\dim D$ is also 0.

Now assume rk $D = r > 0$. Using Lemma 3.1.9, choose any $\mathcal{M}$-definable function $f : D \to M^r$ which is almost surjective and almost finite-to-one in the sense of $\mathcal{M}$. As in the proof of Corollary 3.1.11, let A be the $\mathcal{M}$-definable union of all infinite fibers in f, and let B be the $\mathcal{M}$-definable complement of the image of f. By the choice of f we have

$$\text{rk } A, \text{rk } B < r,$$

and so by the inductive hypothesis

$$\dim A, \dim B < nr = \dim M^r.$$

Since $\dim B < \dim M^r$, B is not generic in M^r, and so f is $\mathcal{R}$-almost surjective. So by Lemma 6.4.6 we get that $\dim D \geq \dim M^r$. In particular, since $\dim A < \dim M^r$, we conclude that $\dim A < \dim D$. Thus A is not generic in M^r, and so f is $\mathcal{R}$-almost finite-to-one. Then again by Lemma 6.4.6, this implies that

$$\dim D = \dim(M^r) = r \cdot \dim M,$$

as desired. □

Remark 6.4.8. In light of Lemma 6.4.7, there is no need to distinguish between $\mathcal{M}$-genericity and $\mathcal{R}$-genericity of $\mathcal{M}$-definable sets. By extension, the same holds for the notions defined in Definitions 6.4.4 and 6.4.5. So, from now on, we will drop the clarifying $\mathcal{R}$ or $\mathcal{M}$ in these notions.

As in Chapter 3, we do still need to distinguish between the two notions of dimension and genericity for points and types. In general, we only have the following, which is immediate from Lemma 6.4.7 using the same argument as in Chapter 3:

Lemma 6.4.9. *Let $a \in M^k$ and $A \subset M$.*

1. $\dim(a/A) \leq n \cdot rk\,(a/A)$.
2. *If D is a set which is $\mathcal{M}$-definable over A, and a is $\mathcal{R}$-generic in D over A, then a is $\mathcal{M}$-generic in D over A.*

Though it is certainly not new, we end this section by deducing the following analog of Corollary 3.1.16:

Corollary 6.4.10. $\mathcal{M}$ *does not interpret any infinite field.*

Proof. First note that, by the exact same argument as in Corollary 3.1.15, the statement of Lemma 6.4.7 holds also for $\mathcal{M}$-interpretable sets: indeed, one only needs to note that $\mathcal{M}$-interpretable sets can be put in finite-to-finite correspondence with $\mathcal{M}$-definable sets, and such correspondences preserve both dimension notions.

Now assume $\mathcal{M}$ interprets the infinite field L. Since L is infinite, we get $\text{rk } L \geq 1$, and therefore

$$\dim L \geq n \geq 3.$$

On the other hand, by [36], Theorem 4.1, the dimension of any field interpreted in an o-minimal structure is at most 2, a contradiction. □

6.5 Geometric Preliminaries

Next we develop some geometric notions that will allow us to fomulate an almost purity hypothesis for $\mathcal{M}$-definable plane curves. The main goal is to find analogs of pure and smooth points for definable sets.

Convention 6.5.1. Throughout this chapter, we use the o-minimal topology on cartesian powers of R. Closures of sets are denoted with overlines (e.g. $\overline{X}$). If $D \subset R^m$ is any definable set, then we endow it with the subspace topology inherited from R^m.

We first define:

Definition 6.5.2. A definable set D is a *quasi-cell* if it is definably homeomorphic to R^d for some d. In this case, we say that D is a *quasi-cell of dimension* d.

Remark 6.5.3. Note that the class of quasi-cells of any given dimension is definable in families, by Fact 6.1.10.

We proceed to define pureness and smoothness:

Definition 6.5.4. Let $D \subset R^m$ be a definable set of dimension d, and let $a \in R^m$ (so not neccessarily $a \in D$).

1. a is a *pure point of* D if for every definable neighborhood N of a in R^m, $D \cap N$ is generic in D.

2. a is a *smooth point of* D if $a \in D$, and there is an definable neighborhood N of a in R^m such that $D \cap N$ is a quasi-cell of dimension d.

Remark 6.5.5. We make some comments:

1. Note that every smooth point of D is also a pure point.

2. Note also that this is an unconventionally weak notion of smoothness, as we do not require any level of differentiability in the identification with R^d.

3. Finally, note that pure and smooth points are definable in families. This follows from Remark 6.5.3.

We next define:

Definition 6.5.6. Let D be a definable set.

1. The pure part of D, denoted D^P, is the set of pure points of D.

2. The smooth part of D, denoted D^S, is the set of smooth points of D.

3. If $D \subset D^P$ then we say D is pure.

4. If $D = D^S$ then we say D is smooth.

Thus D^P and D^S are also definable. The following facts will be useful:

Lemma 6.5.7. *Let D be any non-empty definable set.*

1. *D^S is fully generic in D, and thus is almost equal to D.*
2. *$D^P = \overline{D^S}$, and thus D^P is also almost equal to D.*

Proof. 1. Take any cell decomposition of D. Then the set of non-smooth points is contained in the closures of the non-top dimensional cells, which are non-generic in D.

2. First note that D^P is closed: indeed, if $a \in \overline{D^P}$ and N is a definable neighborhood of a, then $N \cap D$ contains an element $b \in D^P$; thus N is also a definable neighborhood of b, and since $b \in D^P$ it follows that $N \cap D$ is generic in D. This shows that $a \in D^P$.

   So, since $D^S \subset D^P$, we conclude that $\overline{D^S} \subset D^P$. For the other direction, suppose $a \in D^P$; we show that $a \in \overline{D^S}$. Indeed, let N be any definable neighborhood of a. Then $N \cap D$ is generic in D; since D^S is fully generic in D, Lemma 6.4.2 implies that $(N \cap D \cap D^S)$ is generic in D. In particular, N contains an element of D^S.

□

Armed with Lemma 6.5.7, we can make the following helpful reformulation of Definition 6.5.6:

Lemma 6.5.8. *Let D be any non-empty definable set.*

1. *For any fully generic set $E \subset D$ we have $D^P \subset \overline{E}$.*
2. *D is pure if and only if for every fully generic $E \subset D$ we have $D \subset \overline{E}$.*
3. *D is pure and closed if and only if for every fully generic $E \subset D$ we have $D = \overline{E}$.*

Proof. 1. If $a \in D^P$, let N be a definable neighborhood of a. Then $N \cap D$ is generic in D, so by Lemma 6.4.2 $(N \cap D) \cap E$ is generic in D. Thus $N \cap E$ is non-empty, which shows $a \in \overline{E}$.

2. Suppose D is pure and E is fully generic in D. Then by (1) we have

$$D \subset D^P \subset \overline{E}.$$

   Now suppose $D \subset \overline{E}$ for all fully generic $E \subset D$. In particular, for $E = D^S$ we have

$$D \subset \overline{D^S} = D^P,$$

   so D is pure.

3. If D is pure and closed, and $E \subset D$ is fully generic, then by (2) we have $D \subset \overline{E}$. On the other hand, since $E \subset D$ we have
$$\overline{E} \subset \overline{D} = D,$$
and thus $D = \overline{E}$.

   Now suppose $D = \overline{E}$ for every fully generic $E \subset D$. In particular, for $E = D^S$ we have
$$D = \overline{D^S} = D^P,$$
so that D is pure and closed.

□

Finally, we note that pure points also have certain useful preservation properties:

Lemma 6.5.9. *Let D and E be definable sets. If D and E are almost equal then $D^P = E^P$.*

Proof. Note that since $D \cap E \subset D$ we have $(D \cap E)^P \subset D^P$. Conversely, suppose $a \in D^P$, and let N be a definable neighborhood of a. Then $N \cap D$ is generic in D. By almost equality, $D \cap E$ is fully generic in D. So by Lemma 6.4.2, the set
$$(N \cap D) \cap (D \cap E) = N \cap (D \cap E)$$
is generic in D. Now since $\dim(D \cap E) = \dim D$, it follows that $N \cap (D \cap E)$ is also generic in $D \cap E$. This shows that $a \in (D \cap E)^P$.

We have shown that $D^P = (D \cap E)^P$. An identical argument shows that $E^P = (D \cap E)^P$, and thus $D^P = E^P$. □

Lemma 6.5.10. *Let D and E be definable sets, and let $f : D \to E$ be a definable function which is continuous and almost finite-to-one. If $a \in D$ is a pure point of D, then $f(a)$ is a pure point of $f(D)$.*

Proof. Let N be a definable neighborhood of $f(a)$ in E. Then $f^{-1}(N)$ is a definable neighborhood of a in D. So $f^{-1}(N)$ is generic in D. Since f is almost finite-to-one, its restriction to the generic subset $f^{-1}(N)$ is also almost finite-to-one. Now applying Lemma 6.4.6 to both f and its restriction to $f^{-1}(N)$, we conclude that
$$\dim f(f^{-1}(N)) = \dim f^{-1}(N) = \dim D = \dim f(D).$$
But $f(f^{-1}(N))$ is just $N \cap f(D)$. We conclude that
$$\dim(N \cap f(D)) = \dim f(D),$$
which shows that $f(a) \in (f(D))^P$. □

6.6 Almost Purity

In this short section we state the almost purity hypothesis we will require for $\mathcal{M}$. The basic idea is to require that $\mathcal{M}$-definable plane curves are 'well-approximated' by their smooth points, and $\mathcal{M}$-definable functions are 'well-approximated' by covering maps.

First, a quick notational convention:

Notation 6.6.1. If D and E are definable sets, and $f : D \to E$ is a definable function, then for a given set $W \subset E$, we denote by f_W the restricted map $f : f^{-1}(W) \to W$.

Now we define the following:

Definition 6.6.2. Let D and E be definable sets, and let $f : D \to E$ be a definable function. Then given a definable set $W \subset E$, we say that f is a *covering on* W if for some $k \in \mathbb{Z}^+$ the map f_W is a definable k-covering. If f is a covering on $E - T$ for some definable set T, we say that f is a *covering outside* T. If f is a covering outside a definable set T of codimension at least i in E, we say that f is a *covering outside codimension* i.

The following fact is then a straightforward consequence of the Cell Decomposition Theorem. In fact, (i) was given in the previous section.

Lemma 6.6.3. *The following hold for all definable sets:*

1. *For each non-empty definable set D, the set $D - D^S$ has dimension at most* $\dim D - 1$.

2. *Let D and E be definable sets, and let $f : D \to E$ a definable function which is almost surjective and almost finite-to-one. Assume that almost all fibers $f^{-1}(e)$ have size k for some $k \in \mathbb{Z}^+$. Then f is a covering outside codimension 1.*

Proof. (1) follows since D^S is fully generic in D. For (2), let $E' \subset E$ be the fully generic definable set of points with fibers of size exactly k. Let $g_1, ..., g_k : E' \to D$ be defined as follows: for a fixed definable linear order on the ambient affine space containing D, let $g_i(e)$ be the i-th smallest preimage of e. Then, since they are definable, the functions $g_1, ..., g_k$ are continuous on a fully generic definable set $E'' \subset E'$. It follows that f definably trivializes on E'', and so f is a covering on E''. Since E'' is fully generic in E, it follows that f is a covering outside the non-generic set $E - E''$. □

Our formulation of almost purity asks for an improvement on Lemma 6.6.3:

Definition 6.6.4. Let $D \subset M^2$ be a non-trivial plane curve. Then D is *almost pure* if:

1. The set $D - D^S$ has codimension at least 3 in D.

2. Each of the projections $D \to M$ is a covering outside codimension 2.

Remark 6.6.5. For example, every plane curve definable in the vector space structure on R^n is almost pure. In general, if $R = \mathbb{R}$ then any definable set D which is a Boolean combination of complex analytic varieties satisfies condition (2) of Definition 6.6.4; condition (1) is analogous in this case to banning components of complex codimension 1.

Remark 6.6.6. The reader may wonder how Definition 6.6.4 arose, and if there is any merit in assuming it. Our motivation is, in part, to find a hypothesis which can be verified in the presence of a group operation. Among the most successful partial results toward the O-minimal Restricted Trichotomy Conjecture have been the papers [18] and [12], each of which takes place in the presence of a group operation. In each case, the authors use the group operation to recover strong geometric properties of plane curves: namely, they show that plane curves have finite frontier and finitely many 'poles.' In higher dimensions, we expect progress to rely on a similar strategy, showing that frontier points and poles can only happen in an appropriate codimension. Now it is not too hard to see, using the techniques we develop in the next chapter, that Definition 6.6.4 would follow if frontier points and poles can be restricted to codimension 3. Moreover, we in fact successfully carry out this restriction in Chapter 8 for complex algebraic groups. We thus hope that future work will achieve the same restriction in the o-minimal setting; if so, one could prove the local modularity of expansions of e.g. $(R^m, +)$ for all $m \geq 3$.

6.7 Coverings

Our main goal in the next two sections will be to study the behavior of almost pure plane curves under projection maps. In particular we will obtain a strong geometric structure for plane curves – notably a theory of 'components' as referenced in the introduction. The purpose of the present section is to establish an analog of Theorem 3.4.14; we will then apply this analog in the next section to show, assuming almost purity, that each component of a plane curve projects almost bijectively to each copy of M.

Convention 6.7.1. For this section, we fix a non-trivial, $\mathcal{M}$-definable plane cuve $D \subset M^2$, and a projection $\pi : D \to M$. Thus by Lemma 6.4.7, we have $\dim D = \dim M = n$. Note that since M is stationary and D is non-trivial, there is a positive integer l such that almost all fibers $\pi^{-1}(y)$ have size exactly l. We fix this value l for the remainder of the section.

We will work with the following notion:

Definition 6.7.2. A point $y \in M$ is π-*smooth* if $\pi^{-1}(y)$ consists of precisely l points, each of which is smooth in D.

We let W be the set of π-smooth points. Note that W is definable, by Remark 6.5.5. Our goal is to show that π is a covering over W, using an analagous argument to that given in Section 3.4. We will then use the covering on W to show that, assuming almost purity, π is a covering outside codimension 3. We now proceed with the argument:

Proposition 6.7.3. *If π is a covering outside codimension 2, then π is a covering on W.*

Proof. Assuming π is a covering outside codimension 2, we obtain a definable set $U \subset M$ such that:

1. $M - U$ has codimension at least 2 in M.
2. π is a covering on U.

We proceed to show that the covering on U can be extended to W. The proof will be done via a long sequence of claims. To begin, fix $w \in W$. Then w has l preimages in D, say $d_1, ..., d_l$. Our goal is to isolate small enough neighborhoods of the d_i in order to see disjoint homeomorphic mappings to a common base in M. We will do this by developing a gradually improving sequence of approximations to the desired scenario.

Claim 6.7.4. *For each i we can find a definable set $B_i \subset D$ such that:*

1. *B_i is a quasi-cell of dimension n.*
2. *B_i is a relatively open neighborhood of d_i in D.*
3. *The closure $\overline{B_i}$ is definably compact and contained in D.*
4. *For $i \neq j$ we have $\overline{B}_i \cap \overline{B}_j = \emptyset$.*

Proof. That we can satisfy conditions (1) and (2) is immediate from the smoothness of d_i in D. We can further satisfy conditions (3) and (4) by shrinking each B_i as needed. □

Fix $B_1, ..., B_l$ as in Claim 6.7.4. Let $B = B_1 \cup ... \cup B_l$, let $\overline{B}$ be the closure of B, and let ∂B be the boundary of B.

Claim 6.7.5. *There is a definable neighborhood V of w in M such that:*

1. *V is open in M.*
2. *Both V and $U \cap V$ are definably path connected.*
3. *$\pi(\partial B) \cap V = \emptyset$.*

Proof. First we show that there is a definable neighborhood satisfying (3). To do this, note that since each $\overline{B}_i$ is definably compact, the set $\overline{B}$ is also definably compact. Thus the boundary ∂B of B with respect to D is definably compact, and by continuity $\pi(\partial B)$ is definably compact as well. In particular, $\pi(\partial B)$ is relatively closed in M. Now since the sets B_i are relatively open in D, and each d_i belongs to B_i, it follows that $w \notin \pi(\partial B)$. So there is a definable neighborhood of w which is disjoint from $\pi(\partial B)$, as desired.

We have now shown condition (3) to be satisfiable. Fix V satisfying condition (3). Note that since $M = R^n$, we may assume that V is an n-cell, and is thus generic in M, open, and definably path connected.

Finally, recall that $M - U$ has codimension at least 2 in M. Since V is generic in M, it follows that $V - U$ has codimension at least 2 in V. Since V is an n-cell it is definably homeomorphic to R^n, so Lemma 6.2.10 implies that

$$U \cap V = V - (V - U)$$

is definably path connected. □

We now fix the neigborhood V from Claim 6.7.5. We will show that π definably trivializes on $V \cap W$. We first proceed to approximate inverse maps $V \cap W \to B_i$.

Claim 6.7.6. *For each i there is at least one $x \in B_i$ such that $\pi(x) \in U \cap V$.*

Proof. Note that d_i is smooth, and thus pure, in B_i. Since B_i is generic in D by definition, the restricted function $\pi \restriction_{B_i}$ is almost finite-to-one. Thus Lemma 6.5.10 applies to $\pi \restriction_{B_i}$, implying that $\pi(d_i) = w$ is pure in $\pi(B_i)$. Moreover, since π is almost finite-to-one on B_i, Lemma 6.4.6 gives that

$$\dim \pi(B_i) = \dim B_i = n.$$

Now since V is a neighborhood of w, the pureness of w implies that $\pi(B_i) \cap V$ is generic in $\pi(B_i)$, and therefore has dimension n. But V also has dimension n, so it follows that $\pi(B_i) \cap V$ is also generic in V. Finally, since U is fully generic in M it is also fully generic in V, and so by Lemma 6.4.2 it follows that $\pi(B_i) \cap V \cap U$ is generic in V. In particular, $\pi(B_i) \cap V \cap U$ is non-empty. □

We proceed to immediately improve upon Claim 6.7.6:

Claim 6.7.7. *For each i, $U \cap V \subset \pi(B_i)$.*

Proof. Using Claim 6.7.6, fix $\hat{x} \in B_i$ with $\pi(\hat{x}) \in U \cap V$. Since $U \cap V$ is definably path connected, for any $u \in U \cap V$ we can find a definable path from $\pi(\hat{x})$ to u in $U \cap V$. Now since π is a covering on U, by unique path lifting we can lift this path to a definable path in D from $\hat{x}$ to some point $x \in \pi^{-1}(u)$. Every point in this new path lies in $\pi^{-1}(V)$, and hence by Claim 6.7.5 cannot not belong to the boundary of B_i. Since the initial point of the path is in the interior of B_i (namely it is $\hat{x}$), so is the final point x. Thus $x \in B_i$ and $\pi(x) = u$, as desired. □

And now we can immediately improve upon Claim 6.7.7:

Claim 6.7.8. *For each i, $\pi(\overline{B}_i)$ contains V.*

Proof. By Claim 6.7.7, $\pi(\overline{B_i})$ contains $U \cap V$. But $\overline{B_i}$ is definably compact, thus so is $\pi(\overline{B_i})$. So $\pi(\overline{B_i})$ contains $\overline{U \cap V}$. But $U \cap V$ is fully generic in V, so by Lemma 6.5.8, $\overline{U \cap V}$ contains V^P. Since V is a cell it is pure, and so V^P in turn contains V. □

We can now conclude:

Claim 6.7.9. *For each $v \in V \cap W$, the l preimages of v in D consist of exactly one point in each $\overline{B_i}$.*

Proof. Since $v \in W$, v has exactly l preimages in D. Since $v \in V$, it has at least one preimage in each $\overline{B_i}$. Since the $\overline{B_i}$ are disjoint, there must be exactly one preimage in each. □

And now we finally conclude:

Claim 6.7.10. *π definably trivializes on $V \cap W$.*

Proof. For each $i \leq l$, define $g_i : V \cap W \to D$ by setting $g_i(v)$ to be the unique preimage of v in $\overline{B_i}$. It is then clear that $\pi \circ g_i$ is the identity on $V \cap W$ for each i. Note also that since the B_i are definable, each g_i is definable. Moreover, since the $\overline{B_i}$ are pairwise disjoint, the images of the g_i are pairwise disjoint. It remains to show each g_i is continuous. To show this, let $\gamma : (0,1) \to V \cap W$ be a definable curve converging to an element $v \in V \cap W$. Then the composite curve $g_i \circ \gamma$ is contained inside the definably compact set $\overline{B_i}$, and so has a limit $x \in \overline{B_i}$. Since π is continuous, it follows that $\pi(x) = v$. But the only element of $\overline{B_i}$ which maps to v is $g_i(v)$, so it follows that $x = g_i(v)$. That is, the image of a definable curve converging to v is a definable curve converging to $g_i(v)$. In o-minimal fields, this condition is equivalent to continuity (see [7], Chapter 6, Lemma 4.2). □

It follows that there is an open cover of W by definable sets on which π definably trivializes, and so π is a covering on W. □

Our main application of Proposition 6.7.3 is outlined in the next two facts:

Lemma 6.7.11. *If D is almost pure, then* $\dim(M - W) \leq \dim M - 3$.

Proof. We observe that (1) to (3) below each hold for all $y \in M$ outside a set of codimension 3. Of course, this implies that the conjunction of (1) to (3) also holds outside a set of codimension 3. On the other hand, the conjunction of (1) to (3) is precisely the set W, which finishes the proof.

1. The fiber $\pi^{-1}(y)$ has size l. This follows by strong minimality: the generic fiber size must happen outside a set of corank 1, and therefore codimension $n \geq 3$.

2. y is smooth in M. This is immediate since $M = R^n$ is smooth.

3. Every point in $\pi^{-1}(y)$ is smooth in D. This follows from almost purity, since

$$\dim(D - D^S) \leq \dim D - 3 = \dim M - 3,$$

and thus in particular

$$\dim(\pi(D - D^S)) \leq \dim(D - D^S) \leq \dim M - 3.$$

□

Corollary 6.7.12. *If D is almost pure, then π definably trivializes on W.*

Proof. By Proposition 6.7.3, π is a covering on W. By Lemma 6.7.11, W is a the complement in R^n of a definable subset of codimension 3; so by Lemma 6.2.10, W is definably simply connected. By Proposition 6.2.24, each definably path connected component of $\pi^{-1}(W)$ maps bijectively to W. We thus obtain disjoint maps $g_i : W \to D$, with $\pi \circ g_i = \mathrm{id}_W$, by letting each g_i take images in a fixed definably path connected component of $\pi^{-1}(W)$. By Proposition 6.2.17, and the definition of definable trivialization, each g_i is given by a definable function on each of finitely many definable sets which cover W; it follows that each g_i is definable. Finally, note that each g_i is continuous: indeed, continuity can be checked locally, and g_i is continuous on any definable open set on which π definably trivializes. □

6.8 Components

One of the main challenges in attacking strongly minimal reducts of o-minimal fields, as opposed to algebraically closed fields, is the lack of resemblance between the Cell Decomposition Theorem and the irreducible component decomposition of algebraic sets. Our work in Chapter 4 relied heavily on the unique decomposition of any closed definable set into irrecucible components, whose images under certain projections are either 'almost all' or 'almost none' of the target. Meanwhile cell decomposition has neither of these properties – it is non-unique, and a cell can easily cover e.g. 'half' of the target of a projection. However, as it turns out, our almost purity hypothesis implies the existence of a well-behaved 'component decomposition' for almost pure plane curves, with many of the same properties as in the algebraic case. In particular, the definition is intrinsic to the definable set, and we will show in Proposition 6.8.19 that almost finite-to-one projections of components are almost surjective – in fact, in our simply connected case, almost bijective.

The definintion of the component decomposition makes sense for any definable set; however, certain natural properties only hold under stronger assumptions – assumptions which in particular hold for $\mathcal{M}$-definable almost pure sets. We do not actually need all of these properties, but for interest we will outline them after introducing components.

Finally, to avoid confustion for the reader, we note that the components we will define are only 'top dimensional' components – that is, the union of the components of a closed set consists of its pure part, but does not include smaller-dimensional regions. There is no harm here, as we only need to capture the components of highest dimension.

We start with the following definitions. Intuitively, Definition 6.8.2 is to be read as the 'same component' relation – or at least an initial approximation of such a relation.

Definition 6.8.1. Let D and T be definable sets, and let $a, b \in D$. We will say that T *separates a and b in D* if:

1. $a, b \notin T$.

2. There are no definable paths from a to b in $D - T$.

Definition 6.8.2. Let D be a definable set, and let $a, b \in D$. We will say that a and b are *very connected in* D, denoted $a \sim_D b$, if no definable set of dimension at most $\dim D - 2$ separates a and b in D.

Remark 6.8.3. So that Definition 6.8.2 makes sense for all D, one should take the empty set to have dimension $-\infty$; it is then easy to see that $\sim_D$ is just the definable path connectedness relation whenever $\dim D \geq 1$. This does not actually matter for us, though, because we will only use the notion in the case that $\dim D \geq 3$.

Thus two points are very connected in D if it is 'very difficult' to disconnect them in D. Our ideal scenario would be to define the components of D to be the $\sim_D$ classes. However, this method has subtle issues: namely, (i) $\sim_D$ is not always an equivalence relation, (ii) $\sim_D$ is not obviously definable, and (iii) we might want to allow for some points of $\overline{D} - D$ to belong to the components of D.

The solution to the above problems, as it turns out, is to take $\sim_D$ classes on the smooth part of D; in fact, we show below that on D^S the relation $\sim_D$ gives a definable equivalence relation with finitely many classes, each of which is generic in D^S. It would actually be possible to stop at this point and define the components as the $\sim_D$ classes on D^S; however, due to our aesthetic desire to match the components of algebraic sets – see for example item (iii) above – we choose to take the closures of these classes as components. We then conclude easily that D has only finitely many components, and each is definable and of the same dimension as D.

The main advantage of working with smooth points is the following technical lemma:

Lemma 6.8.4. *Let D be a definable set, and let $a, c \in D$. Suppose T is a definable set which separates a and c in D, and satisfies $\dim T \leq \dim D - 2$. Then, for any smooth point $b \in D$, the set $T - \{b\}$ also separates a and c in D.*

Proof. Suppose not. Then there is a definable path γ connecting a to c in $(D - T) \cup \{b\}$. We wish to find a definable path connecting a to c in $D - T$, thus contradicting that T separates a and c in D.

If our current path γ does not pass through b, we are done. Also, if $b \notin T$ then we are done. So, we assume that $b \in T$ and b lies somewhere on γ. Since $a, c \notin T$, we know that b is on the interior of γ.

Now since b is smooth in D, there is a definable neighborhood N of b such that $N \cap D$ is a quasi-cell of the same dimension as D. Choose a point $s \in N \cap D$ which occurs on γ before all occurences of b, and a point $t \in N \cap D$ which occurs on γ after all occurences of b. Then, since $N \cap D$ is definably homeomorphic to $R^{\dim D}$, and $\dim T \leq \dim D - 2$, we get by Lemma 6.2.10 that the set

$$(N \cap D) - (N \cap D \cap T) = (N \cap D) - T$$

is definably path connected. So, we can connect a to s using γ, then s to t using the definable path connectedness of $(N \cap D) - T$, and finally t to c using γ again. The result connects a to c in $D - T$ while avoiding b, as desired. $\square$

And now with Lemma 6.8.4, we are able to show that $\sim_D$ is an equivalence relation on smooth points:

Proposition 6.8.5. *Let D be any non-empty definable set. Then on the set D^S, the relation $\sim_D$ forms a definable equivalence relation with finitely many classes, each of which is generic in D^S.*

Proof. First we show:

Claim 6.8.6. *$\sim_D$ forms an equivalence relation on D^S.*

Proof. The only nontrivial thing to prove is transitivity. So, suppose $a, b, c \in D^S$ with $a \sim_D b$ and $b \sim_D c$. Let T be a definable set which separates a and c in D. We wish to show that $\dim T \geq \dim D - 1$.

So, assume that $\dim T \leq \dim D - 2$. Then Lemma 6.8.4 applies, so we conclude that $T - \{b\}$ also separates a and c in D. That is, a and c are not definably path connected in $(D-T)\cup\{b\}$. It follows, then, that one of the pairs $(a, b), (b, c)$ is not definably path connected in $(D - T) \cup \{b\}$. On other words, $T - \{b\}$ separates at least one of the pairs $(a, b), (b, c)$ in D. So, using either that $a \sim_D b$ or $b \sim_D c$, we conclude that $\dim(T - \{b\}) \geq \dim D - 1$, and therefore also that $\dim T \geq \dim D - 1$, a contradiction. □

Now the main observation to make for the remainder of the proof is:

Claim 6.8.7. *Let $A \subset D$ be a quasi-cell with $\dim A = \dim D$. Then any two elements of $A \cap D^S$ are very connected in D.*

Proof. Let $a, b \in A \cap D^S$, and let T be a definable set of dimension at most $\dim D - 2$ which does not contain a or b. Since $\dim A = \dim D$, we have $\dim T \leq \dim A - 2$. So by Lemma 6.2.10, using that A is definably homeomorphic to $R^{\dim A}$, there is a definable path from a to b in $A - T \subset D - T$. Thus T does not separate a and b in D. □

We also note:

Claim 6.8.8. *The $\sim_D$ classes in D^S are relatively closed in D^S.*

Proof. Let E be a $\sim_D$ class in D^S, and let $a \in D^S \cap \overline{E}$. Since $a \in D^S$, there is a definable neighborhood N of a whose intersection with D is a quasi-cell of the same dimension as D. Since $a \in \overline{E}$, there is an element $b \in N \cap E$. Then by the previous claim $a \sim_D b$, so $a \in E$. □

Now let

$$D^S = (A_1 \cup ... \cup A_j) \cup (B_1 \cup ... \cup B_k)$$

be a cell decomposition of D^S, where each A_i is generic in D^S, and each B_i is non-generic in D^S. Let $A = A_1 \cup ... \cup A_j$. By Claim 6.8.7, each A_i is contained in a single $\sim_D$ class. So, restricted to A, the relation $\sim_D$ is induced by a partition of the A_i. Since there are only finitely many A_i, it follows that $\sim_D$ is definable with finitely many classes on A. Let $Q_1, ..., Q_m$ be the $\sim_D$ classes in A. Then we show:

Claim 6.8.9. *Each $\sim_D$ class in D^S is a union of sets among $\overline{Q_1} \cap D^S, ..., \overline{Q_m} \cap D^S$.*

Proof. First note that D^S is pure, so by Lemma 6.5.8 is contained in the closure of any fully generic subset. In particular, since each B_i is non-generic in D^S, it follows that A is fully generic in D^S, and so $D^S \subset \overline{A}$. Also note that

$$\overline{A} = \overline{Q_1} \cup ...\overline{Q_m},$$

since closures commute with finite unions. It follows that

$$D^S = \bigcup_{i=1}^{m} \overline{Q_1} \cap D^S.$$

It now suffices to show that each $\overline{Q_i} \cap D^S$ is contained in a single $\sim_D$ class D^S. But this follows immediately by Claim 6.8.8 and the fact that Q_i is a $\sim_D$ class in A. □

Now by the last claim, $\sim_D$ has at most m classes in D^S. Furthermore, each class is a finite union of sets of the form $\overline{Q_i} \cap D^S$, each of which is definable because Q_i and D^S are definable. Additionally, since each Q_i contains at least one of the A_i, and $\dim A_i = \dim D^S$, we obtain that each $\overline{Q_i} \cap D^S$ is generic in D^S. Thus each $\sim_D$ class in D_S is definable and generic in D^S. Then, since $\sim_D$ has finitely many classes, each of which is definable, $\sim_D$ is itself a definable equivalence relation on D_S. □

We can now define:

Definition 6.8.10. Let D be a definable set. We define the *components of* D to be the closures of the $\sim_D$-classes in the set D^S.

Definition 6.8.11. Let D be a definable set. We define the *$\mathcal{R}$-degree of* D, denoted $\deg_{\mathcal{R}}(D)$, to be the number of components of D.

Remark 6.8.12. It is not too hard to see that Definition 6.8.11 coincides with Morley Degree in algebraically closed fields. Namely, for D definable in the algebraically closed field $R[i]$, the $\mathcal{R}$-degree of D is the number of top dimensional irreducible components of $\overline{D}$, or in other words the Morley Degree of D as computed from the full field structure.

We now proceed to generate some basic properties of components.

Corollary 6.8.13. *Let D be an $\mathcal{R}$-definable set.*

1. *D has only finitely many components, and each one is definable.*

2. *For each component C of D we have $\dim C = \dim D$, and C is almost contained in D.*

3. *Any two components of D are almost disjoint.*

4. *The union of the components of D is D^P, the set of pure points of D. In particular, D is almost equal to the union of its components.*

Proof. (i) This is immediate from Proposition 6.8.5 and the definition of components.

(ii) It is clear that C is almost contained in D, since C is the closure of a definable subset of D. By Proposition 6.8.5, each component is the closure of a definable set of dimension $\dim D$, so it follows immediately that each component has dimension $\dim D$.

(iii) Let C_1 and C_2 be distinct components of D. Then we get $C_1 = \overline{S_1}$ and $C_2 = \overline{S_2}$ for two distinct $\sim_D$ classes S_1, S_2 in D^S. By (i) we have $\dim C_1 = \dim C_2$, and so $\dim S_1 = \dim S_2$. Then S_1 and S_2 are disjoint sets of the same dimension, so it is clear that their closures are almost disjoint.

(iv) The union of the components is the union of the closures of the $\sim_D$ classes in D^S. As there are only finitely many such classes, and closures respect finite unions, we are equivalentlly taking the closure of the union of these classes. In other words, the union of the components is $\overline{D^S}$, which by Lemma 6.5.7 equals D^P. □

Corollary 6.8.13 is all we really need to verify unimodularity. Out of interest, however, we include a couple extra properties of components below, which hold under conditions reminiscent of restricted trichotomy problems; it is possible that these properties will be useful for future work on this problem, e.g. for non-simply connected groups. The reader who wishes only to see the proof of unimodularity could skip to Proposition 6.8.19 if desired.

Lemma 6.8.14. *Let D and E be definable sets of the same dimension, and assume*

$$\dim(D \triangle E) \leq \dim D - 2.$$

Then D and E have the same components.

Proof. The main observation to make is:

Claim 6.8.15. *Suppose $a, b \in D^S \cap E^S$. Then $a \sim_D b$ if and only if $a \sim_E b$.*

Proof. It is enough to prove one direction. So, let $a \sim_D b$. Let T be a set which separates a and b in E, and let $T' = T \cup (D \triangle E)$. It is then clear that T' separates a and b in D, as $D - T' \subset E - T$. It follows that

$$\dim T' \geq \dim D - 1 = \dim E - 1.$$

But since

$$\dim(D \triangle E) \leq \dim D - 2 = \dim E - 2,$$

this implies that $\dim T \geq \dim E - 1$ as well. Thus we conclude that $a \sim_E b$. □

Now to show that D and E have the same components, it suffices to show that every component of D is a component of E (as the other direction is similar). So, let C be a component of D. Write $C = \overline{A}$ for some $\sim_D$ class A of D^S, and let $A' = A \cap E^S$. By Claim 6.8.15, A' is a $\sim_E$-class in $D^S \cap E^S$. Then by Claim 6.8.15 again, there is a $\sim_E$ class in E^S, say B, such that $A' = B \cap D^S$. But D^S and E^S are almost equal, so that A' is a fully generic subset of both A and B. In particular, since A and B are smooth sets, we obtain that

$$\overline{A} = A^P = \overline{A}' = B^P = \overline{B}.$$

But $\overline{B}$ is a component of E. Thus $C = \overline{A} = \overline{B}$ is a component of E, as desired. □

Lemma 6.8.16. *Let D and E be almost disjoint definable sets of the same dimension. Assume D, E, and $D \cup E$ are all smooth. Then:*

1. *$D \cap \overline{E} = \overline{D} \cap E = \emptyset$.*
2. *The components of $D \cup E$ are just the individual components of D and E.*

Proof. 1. It is enough to prove that $D \cap \overline{E} = \emptyset$, since the other equality is similar. So, fix $d \in D$. Then d is smooth in both D and $D \cup E$, so we may fix quasi-cells A_1 and A_2, each of dimension $\dim D$, which are relative definable neighborhoods of d in D and $D \cup E$, respectively. Since $D \subset D \cup E$, by shrinking if necessary we may assume that $A_1 \subset A_2$. But then A_1 is a definably homeomorphic copy of $R^{\dim D}$ inside another homeomorphic copy of $R^{\dim D}$. By the O-minimal Invariance of Domain Theorem [37], it follows that A_1 is relatively open in A_2. We conclude that A_1 is a relative definable neighborhood of d in A_2, and therefore also in $D \cup E$. That is, we have a definable neighborhood A_1 of d in $D \cup E$, which is contained in D. It follows that $A_1 \cap E$ is a definable neighborhood of d in E. On the other hand, since E is smooth it is pure, so we obtain that $\dim(A_1 \cap E) = \dim E$. Since $A_1 \subset D$, this implies that $D \cap E$ is generic in E, which contradicts that D and E are almost disjoint.

2. Since all three sets are smooth, it suffices to show that the equivalence relation $\sim_{D\cup E}$ is just the union of the equivalence relations $\sim_D$ and $\sim_E$.

   Since
   $$\dim D = \dim E = \dim(D \cup E)$$
   and $D, E \subset D \cup E$, it is clear that
   $$\sim_D, \sim_E \subset \sim_{D\cup E} .$$

   So it suffices to show that if $a, b \in D \cup E$ and $a \sim_{D\cup E} b$, then either $a \sim_D b$ or $a \sim_E b$.

   Without loss of generality, let us assume $a \in D$; we will show that $b \in D$ and $a \sim_D b$. Let T be any definable set of dimension at most $\dim D - 2$ which does not contain a or b. Then there is a definable path γ from a to b in $(D \cup E) - T$. If γ stays inside D then we are done. If not, then there is necessarily a point c on γ which belongs to $\overline{D} \cap \overline{E}$. But since c lies on γ it belongs to $D \cup E$, and thus either to $D \cap \overline{E}$ or $\overline{D} \cap E$; in either case we contradict (1).

□

For the next corollary, we will say that a definable set D is *quasi-smooth* if

$$\dim(D - D^S) \leq \dim D - 2.$$

Corollary 6.8.17. *Let D and E be almost disjoint definable sets of the same dimension m. Assume that D, E, and $D \cup E$ are all quasi-smooth. Then the components of $D \cup E$ are just the individual components of D and E.*

Proof. Let

$$Z = (D - D^S) \cup (E - E^S) \cup ((D \cup E) - (D \cup E)^S).$$

That is, Z is the combined set of non-smooth points in D, E, and $D \cup E$. By the quasi-smoothness of all three sets, $\dim Z \leq m - 2$. Thus we can apply Lemma 6.8.14, and conclude that $D \cup E$ has the same components as $(D \cup E) - Z$. On the other hand, $D \cup E$ is a smooth set which is the almost disjoint union of the smooth sets $D - Z$ and $E - Z$. So by Lemma 6.8.16, the components of $(D \cup E) - Z$ are the individual components of $D - Z$ and $E - Z$. But again by Lemma 6.8.14, the components of $D - Z$ are the same as the components of D, and the components of $E - Z$ are the same as the components of E. □

Remark 6.8.18. Before we move on, we point out that Lemma 6.8.14 applies whenever D and E are almot disjoint $\mathcal{M}$-definable sets of the same dimension, since then $D \Delta E$ has corank at least 1, and thus codimension at least $n \geq 3$, in each of D and E. We also point out that Corollary 6.8.17 rather trivially holds if all three sets are almost pure, non-trivial plane curves, since quasi-smoothness is contained in the definition of almost purity.

We have now finished giving extra properties of components. At this point, we return to the task at hand. Our main application of components is the following propositition, which is the last main step needed before we can prove the main theorem:

Proposition 6.8.19. *Suppose M is a cell, and D is a non-trivial, irreducible, $\mathcal{M}$-definable plane curve which is almost pure. Let $\pi : D \to M$ be either projection. Then for each component C of D, the restriction $\pi \restriction_{D \cap C}: D \cap C \to M$ is almost bijective.*

Proof. Let l be the generic fiber size of π. Recalling Proposition 6.7.3, Lemma 6.7.11, and Corollary 6.7.12, let W be the set of π-smooth points in M, so that

$$\dim(M - W) \leq \dim M - 3,$$

and π definably trivializes on W. Let $g_1, ..., g_l$ be the definable inverse maps provided by Definition 6.2.13, and let S_i be the image of g_i for each i. So the S_i are definable, and form the 'sheets' of the covering π_W.

Let A be a $\sim_D$ class in D^S such that $C = \overline{A}$. Then A is a generic subset of D^S, and therefore a generic subet of D. Now since W is fully generic in M, and D is non-trivial, it follows that

$$\pi^{-1}(W) = S_1 \cup ... \cup S_l$$

is fully generic in D. So, by Lemma 6.4.2, $A \cap \pi^{-1}(W)$ is generic in D, and so in particular is non-empty. Let us fix $a \in A \cap \pi^{-1}(W)$. Since

$$a \in \pi^{-1}(W) = S_1 \cup ... \cup S_l,$$

we also fix i so that $a \in S_i$. Our goal will be to show that S_i and A are almost equal.

Lemma 6.8.20. *A is almost contained in S_i.*

Proof. We show that if $b \in A \cap \pi^{-1}(W)$ then $b \in S_i$. This is enough because $\pi^{-1}(W)$ is fully generic in D, hence $A \cap \pi^{-1}(W)$ is fully generic in A.

So, let $b \in A \cap \pi^{-1}(W)$. Let $T = D - \pi^{-1}(W)$. Thus $a, b \notin T$. In addition, we can show the following:

Claim 6.8.21. $\dim T \leq \dim D - 3$.

Proof. Since D is non-trivial, π is finite-to-one on D. So, since $T \subset D$, π is also finite-to-one on T. It follows that $\dim T = \dim \pi(T)$. Since $\pi(T) \subset M - W$, we get that

$$\dim T \leq \dim(M - W) \leq \dim D - 3.$$

□

Now since $b \in A$, there is a definable path γ connecting a to b in $D - T = \pi^{-1}(W)$. Then $\pi \circ \gamma$ is a definable path from $\pi(a)$ to $\pi(b)$ in W. Now by unique path lifting, the definable path $\pi \circ \gamma$ has exactly one lift to $\pi^{-1}(W)$ starting at a. On the other hand, we have two such lifts: one is the original path γ, and the other is $g_i \circ \pi \circ \gamma$, the lift provided by the definable trivialization of π over γ. It follows that these two lifts are equal. In particular, γ must be contained in the image of g_i. This shows that $b \in S_i$, as desired. □

Lemma 6.8.22. *$S_i \subset A$.*

Proof. Let $b \in S_i$. We wish to show that $a \sim_D b$. So, let T be a definable set not containing a or b and satisfying $\dim T \leq \dim D - 2$. Without loss of generality, we may assume $T \subset D$ (otherwise we replace T with $T \cap D$).

Now since $\dim T \leq \dim D - 2$, we have

$$\dim \pi(T) \leq \dim D - 2 = \dim M - 2.$$

In particular, since

$$M - (W - \pi(T)) = (M - W) \cup \pi(T),$$

we get that

$$\dim(M - (W - \pi(T))) \leq \dim M - 2.$$

By Lemma 6.2.10, we conclude that $W - \pi(T)$ is definably path connected. Additionally, by the choice of a, b, and T, note that $\pi(a)$ and $\pi(b)$ both belong to $W - \pi(T)$.

Let γ be a definable path from $\pi(a)$ to $\pi(b)$ in $W - \pi(T)$. Then $g_i \circ \gamma$ is a definable path from $g_i(\pi(a))$ to $g_i(\pi(b))$ in $\pi^{-1}(W) - T$. Since $a, b \in S_i$, we have $g_i(\pi(a)) = a$ and $g_i(\pi(b)) = b$. So we have shown that a and b are definably path connected in $\pi^{-1}(W) - T$, and thus also in $D - T$.

It follows that $a \sim_D b$. Now by definition of W we have $\pi^{-1}(W) \subset D^S$, so it follows that a and b belong to the same $\sim_D$ class in D^S. Since $a \in A$, this implies that $b \in A$, which completes the proof that $S_i \subset A$. □

By the previous two lemmas, S_i is almost equal to A. In turn, this that implies S_i is almost equal to $D \cap C$. But the restriction of π to S_i is a bijection with W, since it is inverted by g_i. Since S_i and $D \cap C$ are almost equal, this implies that the restriction of π to $D \cap C$ is almost bijective.

□

6.9 The Main Theorem

In this section, we finally present the main result of the present chapter. The relevant pieces have essentially all been assembled, so the proof will now be quick. The idea is straightforward: under suitable hypotheses, Proposition 6.8.19 implies that the degree of either projection of an almost pure plane curve to the universe is equal to the number of components of the curve, and thus does not depend on the projection. It will follow, as in Chapter 4, that under an appropriate hypothesis $\mathcal{M}$ is unimodular.

Before presenting the theorem, we need to extend the notion of 'all sufficiently generic plane curves' to the o-minimal setting:

Convention 6.9.1. Recall that we are under the assumption that $\mathcal{M} = (M, ...)$ is a strongly minimal structure definable in the o-minimal field $\mathcal{R} = (R, ...)$, whose universe M is equal to R^n for some $n \geq 3$.

Definition 6.9.2. We say that *all sufficiently generic plane curves in* $\mathcal{M}$ *have property* P if there is a natural number N such that the following hold:

1. There is a generically irreducible family of non-trivial plane curves in $\mathcal{M}$ of rank at least N.

2. For any such family $\mathcal{F} = \{F_a\}_{a \in A}$, and for any generic $a \in A$ (over the parameters that define $\mathcal{F}$), the curve F_a has property P.

The above definition is identical to the one presented in Chapter 4 – the only difference is that the meaning of the word 'generic' now comes from the o-minimal dimension.

We now conclude this chapter with the proof of our theorem:

Theorem 6.9.3. *Assume all sufficiently generic plane curves in $\mathcal{M}$ are almost pure. Then $\mathcal{M}$ is unimodular, and thus locally modular.*

Proof. Moving to an elemetary extension if necessary, we may assume $\mathcal{R}$ is $\aleph_0$-saturated, and thus that we have access to generic elements over finite sets. We also assume, adding parameters if necessary, that $\mathcal{M}$ is $\emptyset$-definable in $\mathcal{R}$ – that is, that all atomic relations of $\mathcal{M}$ are $\emptyset$-definable. We will assume familiarity with the notions presented in Sections 4.1 and 4.2 – in particular with degrees of functions, balanced plane curves, and almost unimodularity. Now the main observation to make is:

Lemma 6.9.4. *Let D be a non-trivial, almost pure plane curve in $\mathcal{M}$, and let $\pi : D \to M$ be either projection. Then the degree of π is equal to $\deg_{\mathcal{R}}(D)$, the number of components of D.*

Proof. Let $C_1, ..., C_k$ be the components of D. Adding parameters if necessary, we may assume that each of $D, C_1, ..., C_k$ is $\emptyset$-definable in $\mathcal{R}$, and moreover that D is $\emptyset$-definable in $\mathcal{M}$.

Let $l = \deg \pi$, and let $y \in M$ be any generic element. Then y has exactly l preimages in D, say $d_1, ..., d_l$. Note that y is definable over each d_i, so we have

$$\dim d_i \geq \dim y = \dim M = \dim D,$$

and thus each d_i is generic in D. By Corollary 6.8.13 (4), the set

$$D - (C_1 \cup ... \cup C_k)$$

is non-generic in D, so it follows that

$$\pi(D - (C_1 \cup ... \cup C_k))$$

is non-generic in M. Thus each d_i belongs to some C_j.

On the other hand, by Corollary 6.8.13 (3), the intersection of any two distinct C_j has dimension less than $\dim D$. It follows each d_i belongs to at most one of the C_j, and therefore belongs to exactly one of the C_j.

Now the key point to make is that by Proposition 6.8.19, the fiber above y in each set $C_j \cap D$ has size exactly one. In other words, each C_j contains exactly one of the d_i.

Thus the relation '$d_i \in C_j$' gives a bijective correspondence between $d_1, ..., d_l$ and $C_1, ..., C_k$, from which we conclude that $l = k$. ☐

We conclude:

Corollary 6.9.5. *Every non-trivial, almost pure plane curve in $\mathcal{M}$ is balanced.*

Proof. By Lemma 6.9.4, the degree of either projection $D \to M$ is intrinsic to the set D, and thus does not depend on the choice of projection. ☐

And finally, we conclude:

Corollary 6.9.6. $\mathcal{M}$ *is almost unimodular.*

Proof. By assumption all sufficiently generic plane curves are almost pure. By Corollary 6.9.5, we conclude that all sufficiently generic plane curves are balanced. So, taking $B = 2$ in Definition 4.2.5, we conclude directly that $\mathcal{M}$ is almost unimodular. □

Finally, by Proposition 4.2.6, we conclude that $\mathcal{M}$ is unimodular, and thus also locally modular by Fact 4.1.5. □

To recap, we have shown:

Theorem 6.9.7. *Let* $\mathcal{R} = (R, +, \cdot, <, ...)$ *be an o-minimal field, and let* $\mathcal{M}$ *be a strongly minimal reduct of the full* $\mathcal{R}$*-induced structure on* R^n *for some* $n \geq 3$*. If all sufficiently generic plane curves in* $\mathcal{M}$ *are almost pure, then* $\mathcal{M}$ *is unimodular, and thus locally modular.*

Finally, we point out:

Remark 6.9.8. The assumption that $M = R^n$ is equivalent to the assumption that M is definably homeomorphic to R^n – indeed, the definition of almost purity for plane curves is clearly invariant under definable homeomorphisms of the universe. Thus, the statement of Theorem 6.9.7 also holds in the case that M is a quasi-cell of dimension at least 3.

Chapter 7

Toward Guaranteeing Almost Purity

Until this point, the main results of this thesis have established the local modularity of certain structures with enough almost pure plane curves. In this chapter, we investigate the extent to which almost pure curves must be present. This is, in a sense, the most challenging aspect of restricted trichotomy problems, as one is forced to confront the full complexity of field definable sets. Indeed, the most successful partial result to date [20] does not fully deal with the possibility of codimension 1 components in plane curves – rather these authors show that such components, though they may always exist, cannot disrupt the argument at hand. Such a simplification relies heavily on the universe being of dimension 1, so that codimension 1 components are finite – thus such an argument is not transferrable to higher dimensions.

In the first section of this chapter, we review the notions of very ampleness and semi-indistinguishable points. The potentially complicated behavior of these points is one of the most daunting challenges in future work on this material. The second section presents a simple technical lemma that will simplify various arguments throughout. In the third section we present our first main result, proving that all sufficiently generic plane curves are 'generically almost closed' – almost closed with respect to generic points in the plane. If there is a rank 2 very ample family, we get a much stronger statement – albeit still only relevant for generic points.

In Section 4 we generalize the result of section 3 to plane curves of lower complexity. This goes through fairly smoothly if there is a rank 2 very ample family, but produces a weaker statement otherwise.

Finally, in Section 5 we investigate the potential implications between almost closedness and almost purity in compact universes. Following an argument in [18] and [12], we easily deduce that the almost closedness of *all* plane curves implies the almost purity of all plane curves. We then outline a strategy which could work more generally to show a similar statement while only using the generic almost closedness of *generic* plane curves. This approach is not successful, but to an extent gives information about what must happen in a counterexample. We hope that some version of this argument could be made to work in the future for smooth projective varieties, or at for least smooth projective surfaces.

Before proceeding, we adopt some conventions that we will follow for the rest of the chapter. In particular, as stated below, we fix once and for all a set of parameters capable of defining our reduct structure, and add it to the language. We also add infinitely many constants to the language of the reduct, so that we can always take $\emptyset$-definable almost faithful families of plane curves. We obtain:

Convention 7.0.1. For the rest of this chapter, we assume that $\mathcal{K}$ is the field of complex numbers, expanded by a countable set of constant symbols. We fix M, a smooth complex algebraic variety of dimension $n > 1$, definable over $\emptyset$ in the structure $\mathcal{K}$. We also fix $\mathcal{M} = (M, ...)$, a strongly minimal reduct of the full $\mathcal{K}$-induced structure on M, whose atomic relations are $\emptyset$-definable in $\mathcal{K}$. We assume the language of $\mathcal{M}$ contains countably infinitely many constant symbols, whose interpretations in $\mathcal{M}$ are distinct. We also assume that $\mathcal{M}$ is non-locally modular, and therefore admits $\emptyset$-definable, almost faithful families of plane curves of arbitrarily high rank. As in the previous chapters, we use dim to refer to dimension computed in $\mathcal{K}$, and rk for dimension computed in $\mathcal{M}$. Unless otherwise stated, the terms generic and independent are interpreted according to the structure $\mathcal{K}$. Finally, we tacitly assume that all sets of parameters are countable, so that generic points always exist.

7.1 Common and Semi-indistinguishable Points

In this section we review the notions of very ampleness and semi-indistinguishability. This will in large part echo previous work on related subjects; the interested reader could consult e.g. [12] for more details.

Notation 7.1.1. If $\mathcal{F} = \{F_a : a \in A\}$ is a definable family of plane curves in $\mathcal{M}$, and $x \in M^2$, the notation ${}_xF$ referes to

$$\{a \in A : x \in F_a\}.$$

Similarly if $x, y \in M^2$, the notation ${}_{xy}F$ refers to

$$\{a \in A : x \in F_a \wedge y \in F_a\}.$$

We first reformulate the definition of common points (see Definition 4.3.1) in the above language:

Definition 7.1.2. If $\mathcal{F} = \{F_a\}_{a \in A}$ is a definable family of plane curves in $\mathcal{M}$, then a point $x \in M^2$ is *$\mathcal{F}$-common* if ${}_xF$ is generic in A.

As in Section 4.1, an almost faithful family can have only finitely many common points. This in fact holds for families in general, assuming any two generic curves in the family have finite intersection. Namely, we point out the following:

Lemma 7.1.3. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a non-empty family of plane curves, definable in M over a set E. Assume that for any generic pair $(a, b) \in A^2$ over E, the intersection $F_a \cap F_b$ is finite. Then:*

1. *There are only finitely many $\mathcal{F}$-common points.*
2. *rk $A \geq 1$.*
3. *For all generic $a \in A$ over E, and generic $x \in F_a$ over (E, a), x is $\mathcal{M}$-generic in M^2 over E and rk $(a/x) =$ rk $A - 1$.*
4. *For all generic $x \in M^2$ over E, the set ${}_xF$ has rank rk $A - 1$.*

Proof. Adding parameters if necessary, we may assume $E = \emptyset$.

1. Let x be $\mathcal{F}$-common, and let a and b be independent generic elements of ${}_xF$. So

   $$\text{rk }(a, b/x) = 2 \cdot \text{rk } A,$$

   which means that a and b are independent generics in A as well. By assumption $F_a \cap F_b$ is thus finite, so rk $(x/a, b) = 0$, and thus

   $$\text{rk }(x, a, b) = 2 \cdot \text{rk } A.$$

   By additivity, we conclude that rk $x = 0$, which implies the desired statement.

2. Assume rk $A = 0$. Let $a \in A$, and let x be generic in F_a over a. Thus $a \in {}_xF$, so

   $$\text{rk } {}_xF \geq 0 = \text{rk } A.$$

   So x is $\mathcal{F}$-common; by (1) we conclude that rk $x = 0$, which contradicts that x is generic in F_a over a.

3. Let $a \in A$ be generic, and let x be a generic element of F_a. So

   $$\text{rk }(x, a) = \text{rk } A + 1.$$

   Since rk $(x/a) = 1$, (1) implies that x is not $\mathcal{F}$-common, so

   $$\text{rk }(a/x) \leq \text{rk } A - 1.$$

   By additivity we conclude that rk $x \geq 2$; but $x \in M^2$, so this is only possible if rk $x = 2$ and rk $(a/x) =$ rk $A - 1$, as desired.

4. Since M^2 is stationary in $\mathcal{M}$, suffices to find a single generic point x with the desired property. Indeed, let (x, a) be as in (3), so that x is $\mathcal{M}$-generic in M^2. Then by (1) it follows that x is not common, so

$$\text{rk }_xF \leq \text{rk } A - 1.$$

On the other hand, by (3) we have

$$\text{rk }(a/x) = \text{rk } A - 1,$$

so it follows that rk $_xF = \text{rk } A - 1$, as desired.

□

From the fact that there are only finitely many common points, we make the following conclusion:

Lemma 7.1.4. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a non-empty $\mathcal{M}$-definable family of plane curves, and assume that for any generic pair $(a, b) \in A^2$ the intersection $F_a \cap F_b$ is finite. Then there is a family $\mathcal{F}'$ with the following properties:*

1. *$\mathcal{F}'$ is also indexed by A, and is defined in $\mathcal{M}$ over the same parameters as $\mathcal{F}$.*
2. *For each $a \in A$, F'_a is a cofinite subet of F_a.*
3. *$\mathcal{F}'$ has no common points.*

Proof. Observe that the condition 'x is $\mathcal{F}$-common' is $\mathcal{M}$-definable in x over the parameters defining $\mathcal{F}$. So, we define $\mathcal{F}'$ as follows: For $(x, a) \in M^2 \times A$, we let $x \in F'_a$ if $x \in F_a$ and x is not $\mathcal{F}$-common. By Lemma 7.1.3, $\mathcal{F}$ has only finitely many common points; so all we have done is remove a fixed finite set from each F_a, and the lemma follows. □

Remark 7.1.5. In most of our applications, we will assume that we have replaced any given family $\mathcal{F}$ with the family $\mathcal{F}'$ from Lemma 7.1.4. In particular, we are able to do this whenever $\mathcal{F}$ is almost faithful of positive rank. Moreover, once we define very ample families below, we will see that such families never have common points, even without almost faithfulness. So, when employing almost faithful or very ample families, we will generally assume there are no common points. Of course, the vast majority of properties we wish to study are transferred from a family $\mathcal{F}$ to the family $\mathcal{F}'$ constructed above, so this assumption is typically harmless.

We next study the notion of semi-indistinguishability.

Definition 7.1.6. Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a definable family of plane curves in $\mathcal{M}$, and let $x, y \in M^2$. Then x and y are *$\mathcal{F}$-semi-indistinguishable* if the set $_{xy}F$ has rank at least rk $A - 1$.

Remark 7.1.7. We briefly explain the choice of terminology. In an ideal world we would like to say that x and y are '$\mathcal{F}$-indistinguishable' if the sets ${}_xF$ and ${}_yF$ are almost equal – thus it is very hard to find a curve which 'distinguishes' them by including one but not the other. Now as we will soon see, there are a plethora of technical difficulties that arise from the possibility of indistinguishable points which are not equal, and it would be quite convenient to pretend such pairs don't exist. Indeed, in our ideal world such a dream would have hope: $\mathcal{F}$-indistinguishability is easily seen to be an $\mathcal{M}$-definable equivalence relation, and in reasonable situations turns out to be induced from an equivalence relation in one variable. So one could simply move to imaginaries by taking a quotient of the universe, thereby eliminating the annoyance entirely. Unfortunately, it turns out that the difficulties caused by these 'bad' pairs of points really stem from the condition given in Definition 7.1.6, rather than true indistinguishability; so in general we really need to focus on $\mathcal{F}$-semi-indistinguishability instead. Sadly, there is no reason for these two notions to coincide – and, more to the point, there is no reason for $\mathcal{F}$-semi-indistinguishability to be an equivalence relation. Indeed, one could imagine that for generic $x \in M^2$ the set ${}_xF$ has e.g. Morley degree 2, so that for various (x, y) the sets ${}_xF$ and ${}_yF$ could share a single component, and thus transitivity could easily fail. We therefore make no assumptions about transitivity, and so in general treat semi-indistinguishability as a legitimage possibility.

At the least, we have the following:

Lemma 7.1.8. *Let $\mathcal{F} = \{F_a\}_{a\in A}$ be a definable almost faithful family of plane curves in $\mathcal{M}$, and assume rk $\mathcal{F} \geq 2$. Then:*

1. *For each $x \in M^2$ which is not $\mathcal{F}$-common, there are only finitely many $y \in M^2$ such that x and y are $\mathcal{F}$-semi-indistinguishable.*

2. *If there are no $\mathcal{F}$-common points, then any two $\mathcal{F}$-semi-indistinguishable points are $\mathcal{M}$-interalgebraic over the parameters defining $\mathcal{F}$.*

Proof. It is clear that (1) implies (2), so we need only show (1). Now assume that $x, y \in M^2$ are semi-indistinguishable and x is not common. Adding constants if necessary, we may assume that $\mathcal{F}$ is $\emptyset$-definable in $\mathcal{M}$; so it will suffice to show that rk $(y/x) = 0$.

Now our assumptions on x and y give that rk ${}_{xy}F \geq$ rk $A - 1$ and rk ${}_xF \leq$ rk $A - 1$. Since $\mathcal{F}_{xy} \subset {}_xF$, it follows that rk ${}_xF =$ rk $A - 1$.

Let $\mathcal{G}$ be the subfamily consisting of those curves in $\mathcal{F}$ indexed by elements of ${}_xF$; so $\mathcal{G}$ is $\mathcal{M}$-definable over x. Note that since $\mathcal{F}$ is almost faithful, so is $\mathcal{G}$. Also note that since ${}_xF$ has rank rk $A - 1$, $\mathcal{G}$ has rank rk $A - 1$. In particular, since rk $A \geq 2$, $\mathcal{G}$ is an almost faithful family of positive rank. By Lemma 4.3.2, $\mathcal{G}$ has only finitely many common points. On the other hand, the definition of semi-indistinguishability implies that y is one of these $\mathcal{G}$-common points. It follows that rk $(y/x) = 0$, as desired.

□

We next discuss the related notion of very ampleness:

Definition 7.1.9. Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a definable family of plane curves in $\mathcal{M}$. Then $\mathcal{F}$ is *very ample* if for all $x, y \in M^2$ with $x \neq y$, x and y are not $\mathcal{F}$-semi-indistinguishable.

So, a very ample family is one with no unequal semi-indistinguishable pairs. Note that such a condition really only applies for families of rank at least 2.

It is possible (e.g. see [31]) that in a non-locally modular strongly minimal structure there are no rank ≥ 2 very ample families of plane curves. However in natural examples such families do seem to exist, at least after potentially moving to imaginaries. For example, such families can be arranged from the full induced structure on an algebraic curve, or in the presence of a group operation (see Chapter 8), and the main theorem of [27] takes place under the assumption of very ampleness.

In this chapter we will at times work both with and without the assumption of very ampleness. We will point out those times where very ampleness is assumed; in these times we will generally have simpler and stronger results.

As mentioned above, very ample families cannot have common points. We prove this now:

Lemma 7.1.10. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be a very ample definable family of plane curves in $\mathcal{M}$. Then there are no $\mathcal{F}$-common points.*

Proof. We may assume $\mathcal{F}$ is $\emptyset$-definable in $\mathcal{M}$. Now assume toward a contradiction that $x \in M^2$ is $\mathcal{F}$-common. Then we can find a generic element $a \in A$ over x such that $x \in F_a$. Let y be a generic element of F_a over (a, x). Then it follows that

$$\text{rk } (a, y/x) = \text{rk } A + 1.$$

Since $y \in M^2$ we have rk $(y/x) \leq 2$, so by additivity it follows that

$$\text{rk } (a/x, y) \geq \text{rk } A - 1.$$

Thus x and y are $\mathcal{F}$-semi-indistinguishable. By very ampleness we conclude that $x = y$, which contradicts that y is generic in F_a over x. □

Perhaps the most important point to make about very ampleness is the following lemma. An essentially equivalent statement appears in [12] (Lemma 3.16).

Lemma 7.1.11. *If $\mathcal{F} = \{F_a\}_{a \in A}$ and $\mathcal{G} = \{G_b\}_{b \in B}$ are two definable, very ample families of non-trivial plane curves in $\mathcal{M}$, then the composition $\mathcal{G} \circ \mathcal{F}$ is very ample.*

Proof. Let $(x_1, z_1), (x_2, z_2) \in M^2$ be distinct. We seek to show that the set of $(a, b) \in A \times B$ with $(x_1, z_1), (x_2, z_2) \in G_b \circ F_a$ has rank at most rk $A +$ rk $B - 2$.

So, let $(a, b) \in A \times B$ such that each (x_i, z_i) belongs to $G_b \circ F_a$. Then there are $y_1, y_2 \in M$ such that each $(x_i, y_i) \in F_a$ and each $(y_i, z_i) \in G_b$. Without loss of generality assume $x_1 \neq x_2$. Then by the very ampleness of $\mathcal{F}$, we have

$$\text{rk } (a/x_1, y_1, x_2, y_2) \leq \text{rk } A - 2.$$

By the non-triviality of each curve, we also have $\text{rk } (y_i/b, z_i) = 0$ for each i. Thus

$$\text{rk } (a, b/x_1, z_1, x_2, z_2) = \text{rk } (a, b, y_1, y_2/x_1, z_1, x_2, z_2)$$
$$= \text{rk } (b/x_1, z_1, x_2, z_2) + \text{rk } (y_1, y_2/b, x_1, z_1, x_2, z_2) + \text{rk } (a/b, x_1, y_1, z_1, x_2, y_2, z_2)$$
$$\leq \text{rk } B + 0 + \text{rk } A - 2,$$

as desired. □

Thus, if we have one very ample family, then by composing repeatedly we can take a much larger very ample family. Note that there is no reason for such a composition to preserve almost faithfulness, so we in general do not assume the existence of arbitrarily large very ample, almost faithful families. However, we can assume that the ranks of these compositions become arbitrarily large. Namely, we note the following, which is also essentially given in [12] (Lemma 3.20):

Fact 7.1.12. *If $\mathcal{F}$ is a family of non-trivial plane curves in $\mathcal{M}$ of rank at least 2, and the composite family $\mathcal{F} \circ \mathcal{F}$ has rank at most rk $\mathcal{F}$, then $\mathcal{M}$ interprets an algebraically closed field.*

Proof. Adjusting the parameter set if necessary, we may assume that $\mathcal{F}$ is almost faithful. We may further assume that $\mathcal{F}$ is $\emptyset$-definable in $\mathcal{M}$. Let a and b be independent generics in A, and let c be a code of any strongly minimal component S of $F_b \circ F_a$. Let $(x, z) \in S$ be generic over (a, b, c), and let y be such that $(x, y) \in F_a$ and $(y, z) \in F_b$. Then, using an analogous argument to that presented in Section 5.5, and using the fact that $\mathcal{F} \circ \mathcal{F}$ has rank at most rk $\mathcal{F}$, one can now check the following:

1. $\text{rk } a = \text{rk } b = \text{rk } c = \text{rk } \mathcal{F}$.

2. $\text{rk } (a, b) = \text{rk } (a, c) = \text{rk } (b, c) = \text{rk } (a, b, c) = 2 \cdot \text{rk } \mathcal{F}$.

3. (x, y, z) is generic in M^3.

4. x and y are interalgebraic over a, y and z are interalgebraic over b, and x and z are interalgebraic over c.

5. Any three points among (a, b, c, x, y, z) not discussed above are independent.

The data above, for $\text{rk } \mathcal{F} \geq 2$, is known as a 'Field Configuration' – and, indeed, by results of Hrushovski (see [3], Main Theorem, Proposition 2), implies the interpretability of an algebraically closed field. □

We thus have the following conclusion:

Corollary 7.1.13. *If there is a very ample family of non-trivial plane curves in $\mathcal{M}$ of rank at least 2, then there are such families of arbtitrarily large rank.*

Proof. Assume there is a bound on the ranks of very ample families, and let $\mathcal{F}$ be such a family of maximal rank. Then by assumption and Lemma 7.1.11, $\mathcal{F}$ satisfies the hypotheses of Fact 7.1.12; we conclude that $\mathcal{M}$ interprets an algebraically closed field. But this contradicts Corollary 3.1.16. □

We close this section with an additional definition, made in light of Corollary 7.1.13. The only subtlety to point out is that we assume we have added enough constants to the language in order to define all of the families that we need.

Definition 7.1.14. We say that the structure $\mathcal{M}$ is *very ample* if the following hold:

1. There is a $\emptyset$-definable, almost faithful, very ample family of non-trivial plane curves of rank 2, indexed by a stationary definable set in $\mathcal{M}$.

2. For each $k \in \mathbb{Z}^+$, there is a $\emptyset$-definable, very ample family of non-trivial plane curves of rank at least k, indexed by a stationary definable set in $\mathcal{M}$.

7.2 Sweeping Sets

The purpose of this section is to establish a definition and lemma which will help to simplify the majority of the arguments we make in the later sections. We first introduce the notion of one set 'k-sweeping' another – which approximately refers to an independence statement over certain parameters remaining independent over fewer parameters. The goal is to establish a situation when this always happens; while elementary to verify, the statement that we show will be quite helpful moving forward.

We first define:

Definition 7.2.1. Let $\mathcal{X} = (X, ...)$ be any strongly minimal structure, and let $A \subset B$ be sets of parameters.

1. Let E be an A-definable set, and $D \subset E$ a non-empty B-definable subset. Given a positive integer k, we say that D *k-sweeps E over A* if every B-generic tuple $(x_1, ..., x_k) \in D^k$ is A-generic in E^k.

2. If D is a non-empty B-definable set which k-sweeps some A-definable set over A, then we say that D has *k-sweeping over A*.

That is, Definition 7.2.1 roughly states that the automorphism conjugates of D over A 'fill out,' our 'sweep,' the set E to 'order k.' One can also think of this as an 'independence preservation' property: we are basically saying that the independence of $x_1, ..., x_k$ over B transfers to independence over A.

We proceed to note the following easy properties:

Lemma 7.2.2. *Let $\mathcal{X}, A, B, D, k$ be as in Definition 7.2.1.*

1. *The statement 'D k-sweeps E over A' does not depend on the set B, as long as it is capable of defining B. Thus the phrases 'D k-sweeps E over A' and 'D has k-sweeping over A' are well-defined.*

2. *If D k-sweeps E over A and $j < k$, then D j-sweeps E over A. Thus, if D has k-sweeping over A and $j < k$, then D has j-sweeping over A.*

Proof. 1. Assume D k-sweeps E over A, and let B' be another set containing A which is capable of defining D. Then every generic type of D^k over B' extends to a generic type of D^k over $B \cup B'$. So, if $(x_1, ..., x_k) \in D^k$ is B'-generic, then we may assume without loss of generality that it is in fact $(B \cup B')$-generic, and therefore also B-generic; then by assumption $(x_1, ..., x_k)$ is A-generic in E^k.

2. Without loss of generality we may assume that $\mathcal{X}$ is saturated. Let $(x_1, ..., x_j) \in D^j$ be B-generic. Then by saturation we can extend to a B-generic $(x_1, ..., x_k) \in D^k$. Then by assuption, $(x_1, ..., x_k)$ is A-generic in E^k, and so in particular $(x_1, ..., x_j)$ is A-generic in E^j.

□

It is easy to build examples where Definition 7.2.1 fails. Indeed, for generic $x_0 \in X$ the x_0-definable set $\{(x_0, y) : y \in X\}$ does note have 2-sweeping over $\emptyset$: a generic pair has the form $(x_0, y), (x_0, z)$, which are not $\emptyset$-independent since they have the same first coordinate.

On the other hand, returning to the setting of Convention 7.0.1, the following states that sweeping always holds for generic plane curves:

Lemma 7.2.3. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be an almost faithful family of plane curves in $\mathcal{M}$, definable over a set E. Assume that rk $\mathcal{F} = k > 0$. Then for all generic $a \in A$ over E, the set F_a k-sweeps M^2 over E, as viewed in the structure $\mathcal{K}$.*

Proof. The argument we give can actually establish k-sweeping as viewed in either $\mathcal{M}$ or $\mathcal{K}$, but we will only need the statement in $\mathcal{K}$.

Without loss of generality we may assume that $E = \emptyset$. Fix a generic element $a \in A$, and let $(x_1, ..., x_k) \in (F_a)^k$ be generic over a. Note, then, that

$$\text{rk } (a, x_1, ..., x_k) = 2k,$$

and similarly

$$\dim(a, x_1, ..., x_k) = 2nk.$$

Let b be an independent realization of $\text{tp}_{\mathcal{K}}(a/x_1, ..., x_k)$ over $(x_1, ..., x_k)$.

Claim 7.2.4. *rk* $(a, b, x_1, ..., x_k) = 2k$.

Proof. That rk $(a, b, x_1, ..., x_k) \geq 2k$ is clear. Now if $F_a \cap F_b$ is finite then

$$\text{rk } (x_1, ..., x_k/a, b) = 0,$$

so

$$\text{rk } (a, b, x_1, ..., x_k) = \text{rk } (a, b) \leq 2k.$$

If $F_a \cap F_b$ is infinite, then by almost faithfulness rk $(b/a) = 0$, and thus

$$\text{rk } (a, b, x_1, ..., x_k) = \text{rk } (a, x_1, ..., x_k) = 2k.$$

□

By the claim, and the fact that rk $(a, x_1, ..., x_k) = 2k$, we obtain rk $(b/a, x_1, ..., x_k) = 0$, and so $\dim(b/a, x_1, ..., x_k) = 0$. By definition of b this implies that $\dim(a/x_1, ..., x_k) = 0$. So, since $\dim(a, x_1, ..., x_k) = 2nk$, we conclude by additivity that $\dim(x_1, ..., x_k) = 2nk$. So $x_1, ..., x_k$ are generic independent elements of M^2, and we are done. □

Now the main goal of this section is the following lemma and corollary:

Lemma 7.2.5. *Let $\mathcal{X}$ be any strongly minimal structure, let a, b, and c be tuples in $\mathcal{X}$, and let A be a set of parameters. Assume that b and c are independent over (A, a). If b is (A, a)-generic in an (A, a)-definable set D, and for some $k >$ rk (c/A) the set D has k-sweeping over A, then b and c are independent over A.*

Proof. Without loss of generality we may assume that $\mathcal{X}$ is saturated and $A = \emptyset$. So b is generic in the a-definable set D, and is independent from c over a; it follows that b is also generic in D over (a, c).

Let $(b_1, ..., b_k) \in D^k$ be generic over (a, c). By the above remarks, we may assume that each b_i realizes tp $(b/a, c)$, and thus in particular realizes tp (b/c).

Now since D has k-sweeping over $\emptyset$, we get that $b_1, ..., b_k$ are independent generics in some $\emptyset$-definable set; thus in particular $b_1, ..., b_k$ are $\emptyset$-independent. It follows that

$$\text{rk } (b_1, ..., b_k) = k \cdot \text{rk } b,$$

and so

$$\text{rk } (b_1, ..., b_k, c) \geq k \cdot \text{rk } b.$$

Since rk $c < k$, we conclude by additivity that

$$\text{rk } (b_1, ..., b_k/c) > k \cdot \text{rk } b - k = k \cdot (\text{rk } b - 1).$$

It follows that for some i we have rk $(b_i/c) >$ rk $b - 1$. But since b_i and b realize the same type over c, this implies that rk $(b/c) >$ rk $b - 1$, and thus rk $(b/c) =$ rk b. This is equivalent to the assertion that b and c are independent over $\emptyset$. □

Returning to our setting, we conclude:

Corollary 7.2.6. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be an almost faithful family of plane curves in $\mathcal{M}$, definable over a set E. Assume that $\operatorname{rk} \mathcal{F} = k > 2n$. Let $a \in A$ be generic over E, and let x be generic in F_a over (E, a). If $y \in M^2$ is independent from x over (E, a), then y is independent from x over E.*

Proof. By Lemma 7.2.3, F_a has k-sweeping over E. Since $y \in M^2$, we have

$$\dim y \leq 2n < k.$$

So by Lemma 7.2.5, applied in the structure $\mathcal{K}$, we conclude that x and y are independent over E. □

Remark 7.2.7. In the future we will assume that any family of curves we work with is $\emptyset$-definable. In this context, the essential content of Corollary 7.2.6 is the following: If $\mathcal{F}$ is a large enough almost faithful family of curves, F_a is a generic curve in $\mathcal{F}$, and x is generic in F_a, then for every $y \in M^2$, independence with x over a implies independence with x over $\emptyset$.

7.3 Generic Almost Closedness

In this section we truly embark on the journey of trying to establish the almost purity of plane curves in $\mathcal{M}$. We will fall short, but are able to give some partial results, and to describe to an extent what must go wrong in a counterexample. The techniques we use are geometric in nature, but are among the most technically demanding of this thesis. We will try to summarize the main geometric idea of each proof as we go.

Notable past papers on restricted trichotomy problems (e.g. [18] and [12]) have used similar strategies, in the presence of a group operation, to rule out components of intermediate dimension in plane curves. These authors first show that plane curves have finite frontier, using a geometric idea from [35]; they then establish other geometric properties of plane curves by reducing to the frontier case. It therefore seems reasonable to start our work by investigating the frontiers of plane curves. This investigation will occupy the next two sections. In short, while we cannot prove that plane curves have finite frontier, we will deduce a sort of 'generic almost closedness' statement that applies to 'most' curves.

In this section we will prove two similar statements bounding the extent to which generic points can belong to the frontiers of 'generic' plane curves. These statements only differ in the assumption of very ampleness, which allows for a stronger conclusion if assumed. Namely, we show in general that the generic plane points occuring on the frontier of a generic curve from a large enough amost faithful family have codimension at least 2 in the closure of the curve; this statement notably does not use very ampleness – something which had been critically relied on by previous authors. Moreover, under the assumption of very ampleness we arrive at the stronger conclusion that such points are $\mathcal{M}$-algebraic over the parameters defining the curve – so in parrticular there are only finitely many. Our conclusion can be intuitively summarized in the following statement: sufficiently generic plane curves become almost closed after restricting to a fixed generic subset of the plane.

Before proceeding, we briefly summarize the argument in both cases. The geometric idea goes back to [18]: given a generic point x_0 on the frontier of a generic plane curve F_{a_0}, we study the intersections of F_{a_0} with generically chosen plane curves through x_0. The intuitive conclusion is that such intersections are 'smaller than usual,' since they are 'missing' x_0. The vast majority of the proof below involves formally verifying that this intuition is correct. Once this is done, our path is simpler: we conclude easily that $\mathcal{M}$ can 'detect' a non-generic property of the point x_0 with respect to F_{a_0}; further analysis of the relevant dimension computations then gives the statement of the theorem.

We now give the theorem:

Theorem 7.3.1. *Let $\mathcal{F} = \{F_a\}_{a\in A}$ be an almost faithful family of plane curves of rank $k > 2n$, definable in $\mathcal{M}$ over a set E of parameters. Let a_0 be generic in A over E, and let x_0 be generic in M^2 over E. Assume that $x_0 \in Fr\ (F_{a_0})$. Then* $\dim(x_0/E, a_0) \leq n - 2$. *Moreover, if $\mathcal{M}$ is very ample, then x_0 is $\mathcal{M}$-algebraic over (E, a_0).*

Proof. The proof will be quite long, so to aid the reader in following the details, we begin with a glossary of the main objects that appear, including where each one is defined. Namely, we will use the following:

- The family $\mathcal{F} = \{F_a\}_{a\in A}$ of rank k, the generic index $a_0 \in A$, and the generic point x_0 which belongs to Fr (F_{a_0}), each given in the statement of the theorem; moreover, the associated definable set $F \subset M^2 \times A$. Note that, as in Assumption 7.3.3, we assume that either $\mathcal{M}$ is very ample or $\dim(x_0/a_0) = n - 1$.

- A second almost faithful family $\mathcal{G} = \{G_b\}_{b\in B}$ of plane curves of rank r, along with its associated definable set $G \subset M^2 \times B$, introduced in Definition 7.3.2, and which we use to intersect with the curves in $\mathcal{F}$.

- A generic index $b_0 \in B$ of a curve through x_0, introduced in Lemma 7.3.5.

- A subfamily $\mathcal{H}$ of $\mathcal{G}$, indexed by a strongly minimal set $S \subset B$, and defined over a parameter $z \in \operatorname{acl}_{\mathcal{M}}(b_0)$ of rank $r - 1$, each introduced in Lemma 7.3.8; moreover, the associated definable set $H \subset M^2 \times S$.

- The fiber product R of F and H over M^2, defined after the proof of Lemma 7.3.8, which is used to count intersections of the curves in $\mathcal{F}$ with the curves in $\mathcal{H}$.

- The positive integer l, which is the generic number of intersections between curves in $\mathcal{F}$ and $\mathcal{H}$, introduced after the proof of Lemma 7.3.12.

Our main goal, then, will be to show that a_0 and b_0 are not $\mathcal{M}$-independent, by analyzing fibers in the projection $R \to A \times S$. We now proceed with the argument:

We begin with some preliminary remarks. By adding constants to the language if necessary, we may assume that $E = \emptyset$. We may also assume that A is stationary in $\mathcal{M}$, since otherwise we could restrict to a stationary component of A containing a_0. As in Lemma 7.1.4, we may assume $\mathcal{F}$ has no common points: indeed, we need only observe that x_0 still belongs to the frontier of any cofinite subset of F_{a_0}. Finally, adopting the notation from Section 7.1, we set

$$_{x_0}F = \{a \in A : x_0 \in F_a\}.$$

Since x_0 is generic in M^2, Lemma 7.1.3 gives that rk $_{x_0}F = k - 1$.

We proceed to define a second family $\mathcal{G}$ of plane curves; these curves will be used to intersect with the curves in $\mathcal{F}$. Our desired conclusion will be that the curves in $\mathcal{G}$ containing x_0 tend to intersect F_{a_0} less often than the other curves in $\mathcal{G}$ – a condition which $\mathcal{M}$ will be capable of detecting.

The definition of $\mathcal{G}$ importantly depends on whether $\mathcal{M}$ is very ample, so we emphasize it below:

Definition 7.3.2. We define a family $\mathcal{G} = \{G_b\}_{b \in B}$ of plane curves as follows:

- If $\mathcal{M}$ is very ample, let $\mathcal{G}$ be a $\emptyset$-definable, almost faithful, rank 2 family of non-trivial plane curves. As in Definition 7.1.14, we assume that B is stationary in $\mathcal{M}$.
- If $\mathcal{M}$ is not very ample, let $\mathcal{G} = \mathcal{F}$.

Thus, in either case, $\mathcal{G}$ has no common points – either since $\mathcal{G} = \mathcal{F}$ or by Lemma 7.1.10. Moreover, for exactly the same reason, any generic pair of curves in $\mathcal{G}$ have finite intersection. Let $r = \text{rk } \mathcal{G}$, and let

$$_{x_0}G = \{b \in B : x_0 \in G_b\}.$$

Then, since x_0 is generic, Lemma 7.1.3 implies that $_{x_0}G$ has rank $r - 1$.

Assumption 7.3.3. We have one more clarifying remark to make: since x_0 belongs to the frontier of the n-dimensional set F_{a_0}, we have a priori that $\dim(x_0/a_0) \leq n - 1$. We will assume throughout that either $\mathcal{M}$ is very ample, or $\dim(x_0/a_0) = n - 1$. Note, by method of contradiction in the non-very ample case, that it will suffice to prove the theorem under this assumption. Our eventual conclusion will be that x_0 is $\mathcal{M}$-algebraic over a_0. If $\mathcal{M}$ is very ample then this is just the conclusion of the theorem; in case $\mathcal{M}$ is not very ample, the proper interpretation is that we have found a contradiction with the assumption $\dim(x_0/a_0) = n-1$. That is, importantly, the proof that x_0 is $\mathcal{M}$-algebraic over a_0 in the non-very ample case seems to only work under the initial assumption that $\dim(x_0/a_0) = n - 1$, so in general we only obtain that $\dim(x_0/a_0) \leq n - 2$, as in the theorem statement.

Proceeding with the proof, and working from now on under Assumption 7.3.3, we conclude the following. In essense, this is the main point of making Assumption 7.3.3 in the first place:

Claim 7.3.4. *If $\mathcal{M}$ is not very ample then $x_0 \in (F_{a_0})^P$.*

Proof. By Assumption 7.3.3 we have $\dim(x_0/a_0) = n-1$. Now if $x_0 \notin (F_{a_0})^P$, then $x_0 \in C$ for some component $C \subset \overline{F_{a_0}}$ of dimension at most $n-1$. Since $x_0 \in \operatorname{Fr}(F_{a_0})$, it follows that x_0 is not $\operatorname{acl}(a_0)$-generic in C, which implies that $\dim(x_0/a_0) \leq n-2$, a contradiction. □

Now arguably the main innovation of this theorem, which allows us to prove the non-very ample case, is the following; essentially, it states that most curves through x_0 only intersect points which are semi-indistinguishable from x_0 on the pure part of F_{a_0}:

Lemma 7.3.5. *There is a generic element $b_0 \in {}_{x_0}G$ over (x_0, a_0), such that for every element $x \in F_{a_0} \cap G_{b_0}$, if x and x_0 are $\mathcal{G}$-semi-indistinguishable then $x \in (F_{a_0})^P$.*

Proof. First assume that $\mathcal{M}$ is very ample, and so $\mathcal{G}$ is very ample by Definition 7.3.2. Then let b_0 be any generic element of ${}_{x_0}G$ over (x_0, a_0). If $x \in F_{a_0} \cap G_{b_0}$ then, since $x_0 \notin F_{a_0}$, we in particular have that $x \neq x_0$. Since $\mathcal{G}$ is very ample, we conclude that x and x_0 are not $\mathcal{G}$-semi-indistinguishable, so the desired statement is true vacuously.

Now assume that $\mathcal{M}$ is not very ample. Then by Definition 7.3.2 we have $\mathcal{G} = \mathcal{F}$, and by Claim 7.3.4 we have $x_0 \in (F_{a_0})^P$. Let F be the definable set associated to $\mathcal{F}$. Note that since a_0 is generic in A, Lemma 3.2.7 implies that $(F^P)_{a_0} = (F_{a_0})^P$.

Let T be the set of elements of $F_{a_0} - (F_{a_0})^P$ which are $\mathcal{F}$-semi-indistinguishable from x_0. Then for any $x \in T$, the above paragraph implies that $(x, a_0) \notin F^P$; so for each such x there is an analytic relatively open neighborhood of (x, a_0), say $U_x \times V_x \subset M^2 \times A$, which is disjoint from F^P. Now by Lemma 7.1.8, the set T is finite. Let V be the intersection of the sets V_x for each of the finitely many $x \in T$. Then, since each V_x is an analytic relatively open neighborhood of a_0 in A, so is V.

Since x_0 is generic in M^2, Lemma 3.2.7 implies that the fiber above x_0 in F^P is just $({}_{x_0}F)^P$; it then follows that $a_0 \in ({}_{x_0}F)^P$. So, since V is an analytic open neighborhood of a_0, we can choose $b_0 \in V$ which is a generic element of ${}_{x_0}F$ over (x_0, a_0). We conclude that

$$\dim(x_0, b_0) = n(k+1) = \dim F,$$

and thus that (x_0, b_0) is generic in F.

Now assume that $x \in F_{a_0} \cap F_{b_0}$ and x and x_0 are $\mathcal{F}$-semi-indistinguishable. Then by Lemma 7.1.8, x and x_0 are interalgebraic over $\emptyset$. In particular,

$$\dim(x, b_0) = \dim(x_0, b_0) = \dim F,$$

so that (x, b_0) is also generic in F. Of course, this implies that $(x, b_0) \in F^P$. On the other hand, by definition of V it follows that each of the sets $U_{x'} \times \{b_0\}$, for $x' \in T$, is disjoint from F^P. It follows that x cannot belong to any such $U_{x'}$, and thus in particular cannot belong to T. So $x \in (F_{a_0})^P$, as desired. □

Fix b_0 from Lemma 7.3.5. So $\dim(b_0/x_0) = n(r-1)$, and thus $\dim(x_0, b_0) = n(r+1)$. So, letting G be the definable set associated to $\mathcal{G}$, (x_0, b_0) is generic in G. In particular, b_0 is generic in B, and x_0 is generic in G_{b_0} over b_0.

Now, as hinted at above, our goal is to study the intersection points of F_{a_0} and G_{b_0}. To this end, the main observation is the following:

Lemma 7.3.6. *Let x be any element of $F_{a_0} \cap G_{b_0}$. Then at least one of the following holds:*

1. *x and x_0 are $\mathcal{G}$-semi-indistinguishable.*

2. *(x, a_0) is generic in F and (x, b_0) is generic in G.*

Proof. Assume that x and x_0 are not $\mathcal{G}$-semi-indistinguishable. So, since both x and x_0 belong to G_{b_0}, it follows that

$$\text{rk } (b_0/x_0, x) \leq r-2,$$

and so

$$\dim(b_0/x_0, x) \leq n(r-2).$$

Note also that $\dim(x/a_0, x_0) \leq n$, since $x_0 \in F_{a_0}$. We conclude:

$$\begin{aligned}\dim(x, b_0/x_0, a_0) &= \dim(x/a_0, x_0) + \dim(b_0/a_0, x_0, x) \\ &\leq n + n(r-2) = n(r-1).\end{aligned}$$

On the other hand, by the choice of b_0 we have

$$\dim(x, b_0/x_0, a_0) \geq \dim(b_0/x_0, a_0) = n(r-1).$$

The only way the last two inequalities can hold simultaneously is if

$$\dim(x/x_0, a_0) = n$$

and

$$\dim(b_0/a_0, x_0, x) = n(r-2).$$

In particular, x must be a generic element of F_{a_0} over (x_0, a_0).

It follows immediately that x is generic in F_{a_0} over a_0, so that (x, a_0) is generic in F. It also follows immediately that x and x_0 are independent over a_0. But x is generic in the generic curve F_{a_0}, from an almost faithful family of rank $k > 2n$; so by Corollary 7.2.6, we conclude that x and x_0 are independent over $\emptyset$. Now since x is generic in F_{a_0} over a_0, it is also generic in M^2 over $\emptyset$ – indeed, this follows by setting $k = 1$ in the statement of Lemma 7.2.3. We conclude that each of x and x_0 is generic in M^2, and so by independence it follows that $\dim(x_0, x) = 4n$.

Finally, we concluded above both that

$$\dim(b_0/a_0, x_0, x) = n(r-2),$$

and that

$$\dim(b_0/x_0, x) \leq n(r-2).$$

It follows imediately that

$$\dim(b_0/x_0, x) = n(r-2).$$

Thus we have

$$\begin{aligned}\dim(b_0, x_0, x) &= \dim(x_0, x) + \dim(b_0/x_0, x)\\ &= 4n + n(r-2) = n(r+2).\end{aligned}$$

Since $x_0 \in G_{b_0}$, we have $\dim(x_0/b_0, x) \leq n$; so it follows that

$$\dim(b_0, x) \geq n(r+1) = \dim G,$$

and thus that (x, b_0) is generic in G.

So, assuming that x and x_0 are not $\mathcal{G}$-semi-indistinguishable, we have concuded that (x, a_0) is generic in F and (x, b_0) is generic in G; this completes the proof of Lemma 7.3.6. □

Corollary 7.3.7. *Let x be any element of $F_{a_0} \cap G_{b_0}$. Then $x \in (F_{a_0})^P$ and (x, b_0) is generic in G.*

Proof. We apply the two cases of Lemma 7.3.6. If x and x_0 are $\mathcal{G}$-semi-indistinguishable, then by the choice of b_0 from Lemma 7.3.5 we get $x \in (F_{a_0})^P$. Furthermore, by Lemma 7.1.8 x and x_0 are $\mathcal{M}$-interalgebraic over $\emptyset$, so $\dim(x, b_0) = \dim(x_0, b_0)$. Since (x_0, b_0) is generic in G, it follows that (x, b_0) is as well.

Now assume that x and x_0 are not $\mathcal{G}$-semi-indistinguishable. Then by Lemma 7.3.6, we get that (x, a_0) is generic in F and (x, b_0) is generic in G. In particular, x is generic in F_{a_0} over a_0, so $x \in (F_{a_0})^P$, and we are done. □

For reasons that will soon become clear, it will be convenient to work with a subfamily of $\mathcal{G}$ indexed by a strongly minimal set. We do this now:

Lemma 7.3.8. *There are a subset $S \subset B$, and a tuple z, with the following properties:*

1. *S is strongly minimal and z-definable in $\mathcal{M}$.*

2. *z is $\mathcal{M}$-algebraic over any element of S.*

3. *b_0 is generic in S over z.*

Proof. Recall that B has rank r and is stationary in $\mathcal{M}$. So, by Lemma 3.1.9, there is a projection $\pi : B \to M^r$ which is almost surjective and almost finite-to-one. Now since b_0 is generic in B, it follows that b_0 is $\mathcal{M}$-interalgebraic over $\emptyset$ with the element

$$\pi(b_0) = (y_1, ..., y_r) \in M^r,$$

which is thus generic in M^r. We conclude that

$$\text{rk } (b_0/y_1, ..., y_{r-1}) = 1$$

and

$$\dim(b_0/y_1, ..., y_{r-1}) = n,$$

so that there a rank one set D, $\mathcal{M}$-definable over $(y_1, ..., y_{r-1})$, which contains b_0 as a generic element. Since $(y_1, ..., y_{r-1})$ is $\mathcal{M}$-algebraic over b_0, we may assume by shrinking D that $(y_1, ..., y_{r-1})$ is $\mathcal{M}$-algebraic over any element of D – indeed as witnessed by the same formula.

Finally, since b_0 is generic in D, there is a strongly minimal component S of D in which b_0 is generic. Since D has only finitely many components, S is definable over $\text{acl}_{\mathcal{M}}(y_1, ..., y_{r-1})$. Let $z \in \text{acl}_{\mathcal{M}}(y_1, ..., y_{r-1})$ be a tuple over which S is definable. Note that since b_0 is generic in D over $(y_1, ..., y_r)$, it follows that b_0 is also generic in S over z. Additionally, since $(y_1, ..., y_{r-1})$ is $\mathcal{M}$-algebraic over any element of D, it follows that z is $\mathcal{M}$-algebraic over any element of S, as desired. □

Fix S and z from Lemma 7.3.8. So z is $\mathcal{M}$-algebraic over b_0, and b_0 has rank 1 and dimension n over z. It follows that z has rank $r-1$ and dimension $n(r-1)$.

Let $\mathcal{H}$ be the subfamily of $\mathcal{G}$ indexed by S – that is, the family with associated definable set

$$H = \{(x, b) \in G : b \in S\}.$$

Let R be the fiber product of $\mathcal{F}$ and $\mathcal{H}$ over the plane – that is,

$$R = \{(x, a, b) \in M^2 \times A \times S : x \in F_a \cap G_b\}.$$

So H and R are $\mathcal{M}$-definable over z. We give some basic properties of these sets:

Lemma 7.3.9. *H has rank 2, and the projection $H \to M^2$ is almost surjective and almost finite-to-one.*

Proof. $\mathcal{H}$ is a family of plane curves indexed by a rank 1 set, so it is clear that H has rank 2. Further note that $\mathcal{H}$ is almost faithful, since $\mathcal{G}$ is. It now follows from Lemma 7.1.3 that for generic $(x, b) \in H$ over z, x is $\mathcal{M}$-generic in M^2 over z and b is $\mathcal{M}$-algebraic over (x, z). In other words, the projection $H \to M^2$ is almost surjective and almost finite-to-one. □

Lemma 7.3.10. *R has rank $k+1$, and the projection $R \to A \times S$ is almost surjective and almost finite-to-one.*

Proof. The proof is similar to Lemma 5.3.2. Namely, it suffices to show that (1) R has elements of rank at least $k+1$ over z, and (2) for all such elements $(x, a, b) \in R$, $F_a \cap G_b$ is finite.

For (1), let a be generic in ${}_{x_0}F$ over b_0. So $(x_0, a, b_0) \in R$. Also,

$$\text{rk } (x_0, a, b_0) = \text{rk } (x_0, b_0) + \text{rk } (a/x_0, b_0) = (r+1) + (k-1) = r + k.$$

Since rk $z = r - 1$ and z is $\mathcal{M}$-algebraic over b_0, it follows that

$$\text{rk } (x_0, a, b_0/z) = (r + k) - (r - 1) = k + 1.$$

For (2), let $(x, a, b) \in R$ be such that rk $(x, a, b/z) \geq k+1$. Since $b \in S$, rk $(z/b) = 0$. So

$$\text{rk } (x, a, b) = \text{rk } (x, a, b, z)$$

$$= \text{rk } z + \text{rk } (x, a, b/z) \geq (r - 1) + (k + 1) = k + r.$$

Now assume $F_a \cap G_b$ is infinite. Then there is a strongly minimal set C which is almost contained in both F_a and G_b. Let c be a code for C, so that by Lemma 4.2.2 c is $\mathcal{M}$-algebraic over each of a and b. Thus rk $(x, a, b, c) = k + r$.

Let a' be an independent realization of $\text{tp}_{\mathcal{M}}(a/b, c)$ over a. Then, in particular, C is almost contained in $F_{a'}$. It follows that C is almost contained in $F_a \cap F_{a'}$, so in particular $F_a \cap F_{a'}$ is infinite. Since $\mathcal{F}$ is almost faithful, we get that rk $(a'/a) = 0$. Thus rk $(a'/a, b, c) = 0$. By definition of a', this implies that rk $(a/b, c) = 0$. Then, recalling that c is $\mathcal{M}$-algebraic over b, we conclude that

$$\text{rk } (x, a, b, c) = \text{rk } (x, b, c) = \text{rk } (x, b) \leq r + 1.$$

But we previously concluded that rk $(x, a, b, c) \geq k + r$, so it follows that

$$k + r \leq r + 1,$$

and thus $k \leq 1$. This is a contradiction since by assumption $k > 2n$. □

We conclude:

Lemma 7.3.11. *Let (x, b) be generic in H over z.*

1. *The projection $H \to M^2$ is locally surjective near (x, b).*

2. *If a is generic in A and $x \in (F_a)^P$, then $(x, a, b) \in R^P$.*

3. *If $a \in A$ such that (x, a) is generic in F, and $F_a \cap G_b$ is finite, then the projection $R \to A \times S$ is locally surjective near (x, a, b).*

Proof.

1. It follows easily from Proposition 3.4.5 that an almost surjective, almost finite-to-one projection is locally surjective near any generic point of its domain. So the desired statement follows from Lemma 7.3.9.

2. Let U, V, and W be analytic neighborhoods of x, a, and b, respectively, in the relevant powers of M. It will suffice to find a generic element of R over z in $U \times V \times W$. Now by (1) the projection $H \to M^2$ is locally surjective near (x, b). So, shrinking U and W if necessary, we may assume that the projection of $H \cap (U \times W)$ to M^2 covers U.

Since a is generic in A, Lemma 3.2.7 gives that $(F^P)_a = (F_a)^P$. So, since $x \in (F_a)^P$, we conclude that $(x, a) \in F^P$. Then, since $U \times V$ is an analytic open neighborhood of (x, a), we can find a generic element $(x', a') \in F$ over z which belongs to $U \times V$. Then by assumption there is some $b' \in W$ such that $(x', b') \in H$. We conclude that $(x', a', b') \in R$, and

$$\dim(x', a', b'/z) \geq \dim(x', a'/z) = n(k + 1) = \dim R.$$

So (x', a', b') is generic in R over z, as desired.

3. We verify the hypotheses of Proposition 3.4.5. We already know that the projection $R \to A \times S$ is almost surjective, almost finite-to-one, and analytically continuous. By (2), we know that $(x, a, b) \in R^P$. Since $F_a \cap G_b$ is finite, the fiber $R_{(a,b)}$ is finite, and so (x, a, b) is isolated in its fiber.

   Now since a is generic in A, it is smooth in A. Furthermore, since (x, b) is generic in H over z, it follows that b is generic in S over z, and is thus smooth in S. So we get that (a, b) is smooth in $A \times S$.

   It remains to show that (x, a, b) has a compact neighborhood in R. But since (x, a) is generic in F, it has a compact neighborhood in F, which without loss of generality has the form $(U \times V) \cap F$ for some $U \subset M^2$ and $V \subset A$. Similarly, since (x, b) is generic in H, it has a compact neighborhood in H, say of the form $(U' \times W) \cap H$ for some $U' \subset M^2$ and $W \subset S$. Shrinking if necessary, we may assume that $U = U'$. Then

$$N = (U \times V \times W) \cap R$$

   is a compact neighborhood of (x, a, b) in R: indeed, the set

$$P = ((U \times V) \cap F) \times ((U \times W) \cap H)$$

   is a product of compact sets, thus compact; and N is isomorphic to the closed subset of P defined by equating the M^2-coordinates. It follows that N is isomorphic to a compact set, and is thus compact.

□

Before proceeding with the main geometric portion of the argument, we note:

Lemma 7.3.12. *If $\mathcal{M}$ is not very ample, then*

$$\dim(a_0, b_0) = n(k + r) - 1$$

and

$$\dim(a_0, b_0/z) = n(k + 1) - 1.$$

Proof. Let x_1 be an independent realization of $\text{tp}_{\mathcal{K}}(x_0/a_0, b_0)$ over x_0. So $\dim(x_1/a_0) = n-1$, and so $\dim(x_1/a_0, x_0) \leq n-1$. On the other hand we have

$$\dim(b_0, x_1/a_0, x_0) \geq \dim(b_0/a_0, x_0) = n(r-1).$$

It follows that

$$\dim(b_0/a_0, x_0, x_1) \geq n(r-1) - (n-1) > n(r-2).$$

In particular, this implies that $\text{rk}\,(b_0/x_0, x_1) > r-2$, so $\text{rk}\,(b_0/x_0, x_1) \geq r-1$. Since $x_0, x_1 \in G_{b_0}$, this further implies that x_0 and x_1 are $\mathcal{G}$-semi-indistinguishable. So by Lemma 7.1.8, x_0 and x_1 are $\mathcal{M}$-interalgebraic over $\emptyset$.

In particular, we may conclude that $\dim(x_1/a_0, b_0, x_0) = 0$. By definition of x_1, this implies that $\dim(x_0/a_0, b_0) = 0$. Thus

$$\dim(a_0, b_0) = \dim(x_0, a_0, b_0) = \dim(a_0) + \dim(x_0/a_0) + \dim(b_0/a_0, x_0)$$

$$= nk + (n-1) + n(r-1) = n(k+r) - 1.$$

This proves the first assertion of the lemma. For the second assertion, we simply use that z is algebraic over b_0 and of dimension $n(r-1)$. It follows immediately that

$$\dim(a_0, b_0) = n(r+k) - 1 - n(r-1) = n(k+1) - 1.$$

□

We are finally ready to state our geometrically inspired conclusion. By Lemma 7.3.10, and the fact that A and S are both stationary, there is a positive integer l such that $F_a \cap G_b$ has size l for all generic $(a, b) \in A \times S$. Then our goal is to show:

Lemma 7.3.13. *Either $F_{a_0} \cap G_{b_0}$ is not of size l, or there are infinitely many $b \in S$ such that $F_{a_0} \cap G_b$ is not of size l.*

Proof. This is obvious if $F_{a_0} \cap G_{b_0}$ is infinite, so we assume that it is finite. Let $x_1, ..., x_m$ be the distinct elements of $F_{a_0} \cap G_{b_0}$. Note that x_0 is not among $\{x_1, ..., x_m\}$, since $x_0 \notin F_a$.

We will prove a series of claims about the behavior of H and R at $x_0, x_1, ..., x_m$.

Claim 7.3.14. *If x is any of $x_1, ..., x_m$, then $(x, a_0, b_0) \in R^P$. If moreoever $\mathcal{M}$ is not very ample, then $(x_0, a_0, b_0) \in R^P$.*

Proof. The first statement follows directly from Corollary 7.3.7 and Lemma 7.3.11. For the second statement, recall by Claim 7.3.4 that if $\mathcal{M}$ is not very ample, we have $x_0 \in (F_a)^P$. So we may again apply Lemma 7.3.11. □

Claim 7.3.15. *If $\mathcal{M}$ is not very ample, then the fiber $(R^P)_{(a_0,b_0)}$ has size at most l.*

Proof. Recall that a_0 is generic, thus smooth, in A, and b_0 is generic, thus smooth, in B; it follows that (a_0, b_0) is smooth in $A \times S$. It is now easy to see, using Lemma 7.3.12, that the hypotheses of Corollary 3.4.10 are satisfied for the projection of R^P, and so the conclusion follows. □

We conclude:

Corollary 7.3.16. *If $\mathcal{M}$ is not very ample, then $m < l$.*

Proof. By the previous two claims, the fiber $(R^P)_{(a_0,b_0)}$ is of size at most l and contains the $m+1$-element set $\{x_0, x_1, ..., x_m\}$. □

We have now concluded the statement of Lemma 7.3.13 if $\mathcal{M}$ is not very ample. We seek to prove the same statement if $\mathcal{M}$ is very ample.

Claim 7.3.17. *Let x be any of $x_0, x_1, ..., x_m$. Then (x, b_0) is generic in H over z.*

Proof. First note that (x, b_0) is generic in G. Indeed this was already given for $x = x_0$, and for all other values of x it follows by Corollary 7.3.7.

So we have $\dim(x, b_0) = n(r+1)$. Since z is algebraic over b_0 and of dimension $n(r-1)$, it follows that

$$\dim(x, b_0/z) = n(r+1) - n(r-1) = 2n = \dim H,$$

and we are done. □

Claim 7.3.18. *If $\mathcal{M}$ is very ample, and x is any of $x_1, ..., x_m$, then the projection $R \to A \times S$ is locally surjective near (x, a_0, b_0).*

Proof. Assuming $\mathcal{M}$ is very ample, we chose $\mathcal{G}$ to be very ample. So, noting that $x \neq x_0$, it follows that x and x_0 are not $\mathcal{G}$-semi-indistinguishable. So by Lemma 7.3.6, (x, a_0) is generic in F. Then the conclusion now follows directly from Lemma 7.3.11. □

Finally, we conclude:

Claim 7.3.19. *Assume that $\mathcal{M}$ is very ample and $m = l$. Then there are infinitely many $b \in S$ such that $F_{a_0} \cap G_b$ is not of size l.*

Proof. This is obvious if there are infinitely many $b \in S$ such that $F_{a_0} \cap G_b$ is infinite; so, since S is strongly minimal, we assume that there are only finitely many such b.

For $i = 0, 1, ..., m$, let U_i be an analytic open neighborhood of x_i in M^2. We may assume by shrinking that the U_i are pairwise disjoint. By Claim 7.3.17 and Lemma 7.3.11, the projection $H \to M^2$ is locally surjective near (x_0, b_0). So, shrinking if necessary, we can find an analytic relatively open neighborhood W of b_0 in S so that the projection of $H \cap (U_0 \times W)$ to M^2 covers U_0.

By the Claim 7.3.18, the projection $R \to A \times S$ is locally surjective near (x_i, a_0, b_0) for $i = 1, ..., n$. So, shrinking if necessary, we can find an analytic relatively open neighborhood V of a_0 in A so that for each $i = 1, ..., n$, the projection of $R \cap (U_i \times V \times W)$ to $A \times S$ covers $V \times W$.

Since by assumption only finitely many G_b have infinite intersection with F_{a_0}, and also by assumption b_0 is not one of them, we may further assume after shrinking that $F_{a_0} \cap G_b$ is finite for each $b \in W$.

Finally, since $x_0 \in \text{Fr}\ (F_{a_0})$, and U is a neighborhood of x_0, there are infinitely many $y \in U_0 \cap F_{a_0}$. Fix any such y, and set $y_0 = y$. Then there is some $b \in W$ such that $(y_0, b) \in H$, and in turn for each $i = 1, ..., n$ there is some $y_i \in U_i$ such that $(y_i, a_0, b) \in R$. It follows that $F_{a_0} \cap G_b$ contains each of $y_0, y_1, ..., y_m$. Since the U_i are pairwise disjoint, and each $y_i \in U_i$, we get that $|F_{a_0} \cap G_b| \geq m + 1$.

We have thus shown that each of the infinitely many $y \in U_0 \cap F_{a_0}$ belongs to G_b for some $b \in W$ such that $|F_{a_0} \cap G_b| \geq m + 1$. By our assumption on W, each such G_b contains only finitely many such y. So there are infinitely many $b \in W$ such that $|F_{a_0} \cap G_b| \geq m + 1$. Thus, if $m = l$, then there are infinitely many $b \in S$ such that $|F_{a_0} \cap G_b| \geq l + 1$. The proof of Claim 7.3.19 is now complete. □

By Corollary 7.3.16 and Claim 7.3.19, the proof of Lemma 7.3.13 is also now complete. □

Our geometric conclusion is now complete: namely, by Lemma 7.3.13 we have concluded that the generically chosen curve G_{b_0} through x_0 has an atypical intersection property with respect to F_{a_0}. Our remaining task is much more straightforward: we seek to show that $\mathcal{M}$ can detect this atypical intersection correspondence between F_{a_0} and G_{b_0}, and that in turn $\mathcal{M}$ can detect x_0 from F_{a_0}. Thus, we finish the proof of Theorem 7.3.1 with the following two corollaries:

Corollary 7.3.20. *The pair (a_0, b_0) is not $\mathcal{M}$-generic in $A \times B$.*

Proof. Using either of the two cases from Lemma 7.3.13, we first conclude that (a_0, b_0) is not $\mathcal{M}$-generic in $A \times S$ over z. Indeed this is clear if $F_{a_0} \cap G_{b_0}$ is not of size l. If instead there are infinitely many $b \in S$ such that $F_{a_0} \cap G_b$ is not of size l, then there is such an element $b \in S$ which is $\mathcal{M}$-generic in S over (a, z). Thus

$$\text{rk}\ (b/a_0, z) = 1.$$

But since $F_{a_0} \cap G_b$ is not of size l, we also have that

$$\text{rk}\ (a_0, b/z) \leq \text{rk}\ (A \times S) - 1 = k.$$

It follows by additivity that

$$\text{rk}\ (a_0/z) \leq k - 1.$$

So a_0 is not $\mathcal{M}$-generic in A over z, and thus (a_0, b_0) cannot be $\mathcal{M}$-generic in $A \times S$ over z.

Now by the above paragraph, we have

$$\text{rk}\ (a_0, b_0/z) \leq \text{rk}\ (A \times S) - 1 = k.$$

Since z is $\mathcal{M}$-algebraic over b_0 and of rank $r - 1$, it follows that

$$\text{rk}\ (a_0, b_0) = \text{rk}\ (a_0, b_0, z)$$

$$= \text{rk}\ z + \text{rk}\ (a_0, b_0/z) \leq (r - 1) + k = (k + r) - 1.$$

Since $A \times B$ has rank $k + r$, the statement of the corollary follows. □

Corollary 7.3.21. *rk* $(x_0/a_0) = 0$.

Proof. Recall that rk $(b_0/x_0, a_0) = r-1$. Let b_1 be an independent realization of $\text{tp}_{\mathcal{M}}(b_0/x_0, a_0)$ over b_0. So rk $(b_1/x_0, a_0, b_0) = r-1$. Since r is either 2 or $k > 2n$, it follows that $r-1 \geq 1$; so, since $\mathcal{G}$ is almost faithful, $G_{b_0} \cap G_{b_1}$ is finite. But by definition of b_1 we have $x_0 \in G_{b_0} \cap G_{b_1}$. It follows that rk $(x_0/b_0, b_1) = 0$.

Recall that a_0 is generic in A, and so rk $a_0 = k$. So by Corollary 7.3.20, it follows that rk $(b_0/a_0) \leq r-1$. Since b_0 and b_1 realize the same $\mathcal{M}$-type over a_0, it follows that rk $(b_1/a_0, b_1) \leq r-1$. We conclude:

$$\text{rk } (x_0, b_1/a_0, b_0) = \text{rk } (b_1/a_0, b_0) + \text{rk } (x_0/a_0, b_0, b_1) \leq (r-1) + 0 = r-1.$$

On the other hand,

$$\text{rk } (x_0, b_1/a_0, b_0) = \text{rk } (x_0/a_0, b_0) + \text{rk } (b_1/x_0, a_0, b_0) = \text{rk } (x_0/a_0, b_0) + r - 1.$$

These two expressions for rk $(x_0, b_1/a_0, b_0)$ are only consistent if rk $(x_0/a_0, b_0) = 0$. But then

$$\text{rk } (x_0, b_0/a_0) = \text{rk } (b_0/a_0) + \text{rk } (x_0/a_0, b_0) \leq (r-1) + 0 = r-1.$$

On the other hand,

$$\text{rk } (x_0, b_0/a_0) = \text{rk } (x_0/a_0) + \text{rk } (b_0/x_0, a_0) = \text{rk } (x_0/a_0) + r - 1.$$

Then, similarly to above, these two expressions for rk $(x_0, b_0/a_0)$ are only consistent if rk $(x_0/a_0) = 0$. □

Finally, by Corollary 7.3.21, the proof of Theorem 7.3.1 is now complete. □

7.4 Frontiers of Plane Curves in General

Theorem 7.3.1 establishes a restriction on frontiers of 'very generic' plane curves – those which appear as generic members of almost faithful families of rank $> 2n$. So the reader may wonder what can be said about other plane curves; in this section we deduce a similar type of statement for curves which are 'less generic.' As in the previous section, our strategy is similar to that followed by [18] and [12].

The idea of the proof is simple enough to explain, but technical to execute: given any plane curve, we try to realize it as a generic member of *some* family – even if a very small one. We then compose this family with a very large family and take an almost faithful reparametrization of the composition; the result is an almost faithful family which is large enough to apply Theorem 7.3.1. Now given a generic frontier point x_0 on our original curve, we try to track x_0 through a composition to a frontier point y_0 of a generic curve in our composite family. We then apply Theorem 7.3.1 to achieve a dependence of y_0 over the relevant composite curve, and finally use this dependence to create a similar dependence of x_0 over our original curve.

There are two subtle issues where this strategy could fail, and a brief discussion will help to illuminate both the theorem statement and its proof. First, our initial frontier point might not 'remain' a frontier point through a composition. We will show that this can only happen if certain points are semi-indistinguishable – so in particular this part of the argument works if $\mathcal{M}$ is very ample. If $\mathcal{M}$ is not very ample, we conclude via the $\mathcal{M}$-interalgebraicity of the semi-indistinguishable points that x_0 is at least 'partially detectable' from $\mathcal{M}$, in the sense that it cannot be $\mathcal{M}$-generic in M^2 over the parameters defining the initial curve.

The second issue is that our original curve might not be realizable as a generic curve in *any* almost faithful family, at least without adding parameters that we aren't in a position to add. The problem is, roughly, that the parameters defining our curve may be $\mathcal{M}$-generic but not $\mathcal{K}$-generic in some $\emptyset$-definable set. To deal with this, we develop the notion of 'optimal' tuples below – roughly, an optimal tuple is both $\mathcal{M}$-generic and $\mathcal{K}$-generic in the same set. Then our theorem will apply to plane curves with 'optimal definitions' – definitions using only optimal parameters. Of course, non-optimal tuples are by their very nature non-generic, so our theorem still applies to 'most' plane curves.

Finally, an interesting feature of the theorem below is that we do not require our frontier point x_0 to be generic in M^2 – rather we only require both of its coordinates to be generic in M. This works, roughly, because a generic composition will turn a point with generic coordinates into a generic point.

We proceed to define optimal tuples:

Definition 7.4.1. Let A be a set, and a a tuple. Then a is *optimal over* A if

$$\dim(a/A) = n \cdot \text{rk } (a/A).$$

If a is optimal over $\emptyset$ then we say that a is *optimal*.

The following basic properties are routine but quite useful to remember:

Lemma 7.4.2. *Let A be a set, and a and b tuples.*

1. *If a is optimal over A, and b is optimal over (A, a), then (a, b) is optimal over A.*
2. *If a is optimal over A, and b is $\mathcal{M}$-algebraic over (A, a), then b is optimal over A, and a is optimal over (A, b).*

Proof. We may assume that $A = \emptyset$.

1. We compute:

$$\begin{aligned}\dim(a, b) &= \dim a + \dim(b/a)\\ &= n \cdot \text{rk } a + n \cdot \text{rk } (b/a)\\ &= n \cdot \text{rk } (a, b).\end{aligned}$$

So (a, b) is optimal.

2. Since b is $\mathcal{M}$-algebraic over a, we have rk $(a, b) =$ rk a and $\dim(a, b) = \dim a$. Now a is optimal, so $\dim a = n \cdot$ rk a, and thus $\dim(a, b) = n \cdot$ rk (a, b). Then we compute:

$$\dim(a,b) = \dim b + \dim(a/b)$$

$$\leq n \cdot \text{rk } b + n \cdot \text{rk } (b/a)$$

$$= n \cdot (\text{rk } b + \text{rk } (a/b))$$

$$= n \cdot \text{rk } (a,b) = \dim(a,b).$$

This is only possible if $\dim b = n \cdot$ rk b and $\dim(a/b) = n \cdot$ rk $a/b)$. In other words, b is optimal, and a is optimal over b.

□

We now state the main result of this section:

Theorem 7.4.3. *Let E be a set of parameters, and let e be a tuple which is optimal over E. Let $S \subset M^2$ be a plane curve which is $\mathcal{M}$-definable over (E, e), and let $x_0 = (m_1, m_2) \in M^2$ be such that m_1 and m_2 are each generic in M over E. Suppose that $x_0 \in$ Fr (S). Then either $\dim(x_0/E, e) \leq n - 2$, or rk $(x_0/E, e) \leq 1$. If moreover $\mathcal{M}$ is very ample, then rk $(x_0/E, e) = 0$.*

Proof. We begin by making a few simplifications. First, as usual, may assume that $E = \emptyset$. Next we note:

Claim 7.4.4. *We may assume that S is strongly minimal.*

Proof. Decompose S into a finite union $C_1 \cup ... \cup C_j$ of strongly minimal components. Since $x_0 \in$ Fr (S), it follows that $x_0 \in$ Fr (C_i) for some i. Now since C_i is a component of S, it is definable over some tuple c which is $\mathcal{M}$-algebraic over e. By Lemma 7.4.2, c is optimal. Then, if we conclude the desired level of dependence of x_0 over c, the $\mathcal{M}$-algebraicity of c over e will in turn imply the same level of dependence of x_0 over e. □

So we assume that S is strongly minimal. Thus we can show:

Claim 7.4.5. *We may assume that S is non-trivial.*

Proof. If S is trivial and strongly minimal, then without loss of generality we have

$$S \sim M \times \{m\}$$

for some $m \in M$. Note, then, that m is $\mathcal{M}$-definable over e. Now in this case, the fact that $x_0 \in$ Fr (S) implies that x_0 belongs to the finite set $M \times \{m\} - S$. We conclude that x_0 is $\mathcal{M}$-algebraic over (e, m), and thus over e. This conclusion is strong enough to give our desired result in any case. □

Now let s be a code for S. Note that s is $\mathcal{M}$-definable over e, and is thus optimal by Lemma 7.4.2.

We next define a family of curves, which we will use to compose with S. As in the proof of Theorem 7.3.1, the definition of this family depends on whether $\mathcal{M}$ is very ample:

Definition 7.4.6. We define a family $\mathcal{G} = \{G_b\}_{b \in B}$ of non-trivial plane curves in $\mathcal{M}$ as follows:

- If $\mathcal{M}$ is very ample, we choose $\mathcal{G}$ to be a $\emptyset$-definable, very ample family of non-trivial plane curves of rank at least $2n+1$, indexed by a stationary definable set in $\mathcal{M}$.
- If $\mathcal{M}$ is not very ample, we choose $\mathcal{G}$ to be a $\emptyset$-definable, almost faithful family of non-trivial plane curves of rank at least $2n+1$, indexed by a stationary definable set in $\mathcal{M}$. By Lemma 7.1.4, we may assume that $\mathcal{G}$ has no common points.

Note, then, that in either case $\mathcal{G}$ has no common points, either by assumption or by Lemma 7.1.10.

Let $b \in B$ be generic over (e, x_0), and thus also over (s, x_0). Let T be a strongly minimal component of G_b, and let t be a code of T. So t is $\mathcal{M}$-algebraic over b. Since $\text{rk } \mathcal{G} \geq 2n+1$, we may assume that $\text{rk } (t/s, x_0) \geq 2n+1$.

Note that b is optimal over s by definition. Since t is $\mathcal{M}$-algebraic over b, Lemma 7.4.2 gives that t is also optimal over s. Thus, again by Lemma 7.4.2, the pair (s, t) is optimal.

Now let $W = T \circ S$. By the choice of b, m_2 is generic in M over b, and thus also over t; so, noting that G_b is a non-trivial plane curve, and thus so is T, we find some $m_3 \in M$ such that (m_2, m_3) is generic in T over t.

The next two claims verify useful properties of the point (m_1, m_3); indeed, (m_1, m_3) is the 'y_0' described in the introduction to this section.

Claim 7.4.7. *(m_1, m_3) is generic in M^2.*

Proof. By definition $\dim(b/x_0) = \dim B$. But since $\mathcal{G}$ has no common points we have

$$\text{rk } (b/m_2, m_3) \leq \text{rk } B - 1,$$

and so

$$\dim(b/x_0, m_3) \leq n \cdot (\text{rk } B - 1) = \dim B - n.$$

Thus by additivity we have $\dim(m_3/x_0) = n$. In particular, m_3 is generic in M over m_1. Since m_1 is generic in M by assumption, the claim follows. □

Claim 7.4.8. *$(m_1, m_3) \in \overline{W}$.*

Proof. It follows by the genericity of (m_2, m_3) over t that the projection of T to the left copy of M is locally surjective near (m_2, m_3). Now since $(m_1, m_2) \in \text{Fr } (S)$, we can find $(m_1', m_2') \in S$ arbitrarily close to (m_1, m_2). Using local surjectivity, for such points we can in turn find m_3' arbitrarily close to m_3 so that $(m_2', m_3') \in T$, and therefore $(m_1', m_3') \in W$. □

It follows that there is a strongly minimal component U of W such that $(m_1, m_3) \in \overline{U}$. Let u be a code for U. Then by Lemma 4.2.2, each of (s, t, u) is $\mathcal{M}$-algebraic over the other two. We conclude:

Claim 7.4.9. *u is optimal and of rank at least $2n+1$.*

Proof. Recall that (s, t) is optimal. As noted above, u is $\mathcal{M}$-algebraic over (s, t), so by Lemma 7.4.2 we get that u is optimal.

Now as noted above, rk $(t/s, u) = 0$. Recall by the choice of t that rk $(t/s) \geq 2n+1$. It follows that rk $(u/s) \geq 2n+1$, and therefore rk $u \geq 2n+1$. □

As hinted at in the introduction to this section, we now have two cases for how to proceed, depending on whether $(m_1, m_3) \in U$:

- First suppose that $(m_1, m_3) \notin U$. Since U is strongly minimal, there is a $\emptyset$-definable almost faithful family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves in $\mathcal{M}$, of rank rk $u \geq 2n+1$, and an $\mathcal{M}$-generic $a \in A$, $\mathcal{M}$-interalgebraic with u over $\emptyset$, such that $F_a \sim U$. Thus

  $$\text{rk } \mathcal{F} = \text{rk } u \geq 2n+1.$$

  Also, since u is optimal, Lemma 7.4.2 implies that a is optimal, and is thus actually generic in A.

  Now we show:

  Claim 7.4.10. $\dim(m_1, m_3/u) \leq n-2$, *and moreover if $\mathcal{M}$ is very ample then rk $(m_1, m_3/u) = 0$.*

  Proof. First suppose that $(m_1, m_3) \in F_a$. Then since $F_a \sim U$ and $(m_1, m_3) \notin U$, we get that (m_1, m_3) is not $\mathcal{M}$-generic in F_a over (a, u); so (m_1, m_3) is $\mathcal{M}$-algebraic over (a, u), and thus also over u.

  Now suppose that $(m_1, m_3) \notin F_a$. Since $F_a \sim U$, it follows that $(m_1, m_3) \in \text{Fr}(F_a)$. Moreover, as noted above, (m_1, m_3) is generic in M^2. So we are in a position to apply Theorem 7.3.1. We conclude that (m_1, m_3) satisfies the desired level of dependence over a, and thus also over u. □

  By the claim, and the fact that u is $\mathcal{M}$-algebraic over (s, t), and thus in turn over (e, b), we conclude that (m_1, m_3) satisfies the desired level of dependence over (e, b). But m_2 and m_3 are $\mathcal{M}$-interalgebraic over b, since $(m_2, m_3) \in G_b$ and G_b is non-trivial. It follows that x_0 and (m_1, m_3) are $\mathcal{M}$-interalgebraic over (e, b), and so x_0 satisfies the desired level of dependence over (e, b). But since b is generic in B over (e, x_0), it is in particular independent from x_0 over e; so in fact our dependence of x_0 over (e, b) transfers to x_0 over e, and we prove the theorem.

- Now suppose that $(m_1, m_3) \in U$. So there is some m_4 such that $(m_1, m_4) \in S$ and $(m_4, m_3) \in T \subset G_b$. Let $\overline{m}$ denote (m_1, m_2, m_3, m_4). Then we compute:

  $$\text{rk }(b, m_3, m_4/e, x_0) \geq \text{rk }(b/e, x_0) = \text{rk } B.$$

  On the other hand, note that m_4 is $\mathcal{M}$-algebraic over (e, m_1), since $(m_1, m_4) \in S$ and S is non-trivial. So we have

  $$\begin{aligned}\text{rk }(b, m_3, m_4/e, x_0) &\leq \text{rk }(m_4/e, x_0) + \text{rk }(m_3/e, x_0, m_4) + \text{rk }(b/e, x_0, m_3, m_4) \\ &\leq 0 + 1 + \text{rk }(b/\overline{m}).\end{aligned}$$

  It follows that rk $(b/\overline{m}) \geq$ rk $B - 1$. Since G_b contains both (m_2, m_3) and (m_4, m_3), these two points are thus $\mathcal{G}$-semi-indistinguishable. But $m_2 \neq m_4$ since S contains (m_1, m_4) but not (m_1, m_2); thus $\mathcal{G}$ is not very ample, and so by the choice of $\mathcal{G}$ it follows that $\mathcal{M}$ is not very ample.

  We conclude by definition that $\mathcal{G}$ is almost faithful and has no common points; so by Lemma 7.1.8, (m_2, m_3) and (m_4, m_3) are $\mathcal{M}$-interalgebraic over $\emptyset$. Also, looking back at our computations and using that there are no common points, it follows that rk $(b/\overline{m})$ = rk $B - 1$, and so rk $(m_3/e, x_0, m_4) = 1$. In particular, m_3 is $\mathcal{M}$-generic in M over (m_2, m_4). This implies that m_2 and m_3 are $\mathcal{M}$-independent over m_4; but as noted above, m_2 is $\mathcal{M}$-algebraic over (m_4, m_3), so it follows that m_2 is in fact $\mathcal{M}$-algebraic over m_4. So, recalling that m_4 is $\mathcal{M}$-algebraic over (e, m_1), it follows that m_2 is $\mathcal{M}$-algebraic over (e, m_1). This implies that $x_0 = (m_1, m_2)$ has rank at most 1 over e, which is enough to complete the proof.

□

7.5 Toward Showing Almost Closedness Implies Almost Purity

In this final section we investigate the potential for implications of the previous two sections toward our goal of almost purity. We will assume throughout that M is projective – or equivalently compact in the analytic topology. Indeed, developing similar results for non-compact M would seem to necessitate the study of 'poles' of $\mathcal{M}$-definable sets – a task which, as we will see in the next chapter, has proved daunting enough even for strongly minimal groups.

There are two main parts of the current section. First, we essentially copy an argument from [18] and [12] to show that, if every plane curve is almost closed, then every plane curve is almost pure; we thus easily deduce local modularity in this case. Next, we propose an alternate strategy for approaching almost purity under the weaker hypotheses given in the previous two sections; this is not successful, but sheds some light on what would have to go wrong in a counterexample. We hope that this argument can be completed to a proof in future work.

Convention 7.5.1. In addition to the underlying assumptions given in Convention 7.0.1 at begining of this chapter, we further assume throughout this section that the variety M is projective, so is compact in the analytic topology.

We begin with the following argument, which goes back to [18], and serves as the main justification for our intensive study of almost closedness in the previous sections:

Lemma 7.5.2. *If every plane curve in $\mathcal{M}$ is almost closed, then every plane curve in $\mathcal{M}$ is almost pure.*

Proof. Assume that every plane curve in $\mathcal{M}$ is almost closed, and let S be any plane curve. We may assume that S is $\emptyset$-definable in $\mathcal{M}$. We first note some easy reductions:

Claim 7.5.3. *We may assume that S is strongly minimal.*

Proof. Let $S_1, ..., S_m$ be a decomposition of S into strongly minimal sets. If S is not almost pure, then there is a codimension 1 irreducible component C of $\overline{S}$. Since C is irreducible, it is contained in some $\overline{S_i}$. Then since $\overline{S_i} \subset \overline{S}$, C cannot be contained in any top-dimensional component of $\overline{S_i}$ – otherwise the same would hold in $\overline{S}$. So, C is a codimension 1 irreducible component of $\overline{S_i}$, which shows that S_i is not almost pure. In other words, if S is not amost pure then one of the S_i is not almost pure, which is enough to prove the claim. □

So, assume that S is strongly minimal. We can now show:

Claim 7.5.4. *We may assume that S is non-trivial.*

Proof. If S is both strongly minimal and trivial, then without loss of generality we have

$$S \sim M \times \{m\}$$

for some $m \in M$. Then S has finite symmetric difference with the irreducible set $M \times \{m\}$, so in particular every component of $\overline{S}$ has dimension n or 0. Thus S is almost pure. □

So, assume further that S is non-trivial. Then the projection $\pi : S \to M$ to the first coordinate has all but finitely many fibers of size l for some $l \in \mathbb{Z}^+$. Then there is a dense Zariski open set $U \subset M$ such that each $u \in U$ has exactly l preimages in S, each of which belongs to S^P.

To show that S is almost pure, let $(x_0, y_0) \in S$ be such that $\dim(x_0, y_0) \geq n-1$. We will show that $(x_0, y_0) \in S^P$. To do this, we will mildly use the o-minimal structure on $\mathbb{R}$; this is not truly necessary, but makes the argument easier to follow. Note that since $\mathbb{R}$ interprets $\mathbb{C}$, all of the sets we are working with are in fact $\mathbb{R}$-definable.

Now since U is dense, $x_0 \in \overline{U}$. So by curve selection ([7], Chapter 6, Corollary 1.5), there is an $\mathbb{R}$-definable function $\gamma : (0,1) \to U$ such that $\lim_{t\to 0^+} \gamma(t) = x_0$. Since each $\gamma(t) \in U$, and using definable choice ([7], Chapter 6, Proposition 1.2), we obtain $\mathbb{R}$-definable functions $\eta_1, ..., \eta_l : (0,1) \to M$ such that for each t, the elements $\eta_1(t), ..., \eta_l(t)$ are precisely those $y \in M$ with $(\gamma(t), y) \in S$. Since M is compact, and using o-minimality, it follows that each $\lim_{t\to 0^+} \eta_i(t)$ is equal to some $y_i \in M$. Then we note:

Claim 7.5.5. *For each $i = 1, ..., l$, we have $(x_0, y_i) \in S \cap S^P$.*

Proof. Note that $(\gamma(t), \eta_i(t)) \in S^P$ for each t by definition. Since $\gamma(t) \to x_0$ and $\eta_i(t) \to y_i$, and S^P is closed, we get that $(x_0, y_i) \in S^P$.

Now if $(x_0, y_i) \notin S$, then $(x_0, y_i) \in \text{Fr}\ (S)$. But by assumption S is almost closed, so we get that $\dim(x_0, y_i) \leq n-2$, and thus $\dim x_0 \leq n-2$. But since S is non-trivial, the two coordinates of any point in S are $\mathcal{M}$-interalgebraic; in particular this implies $\dim(x_0, y_0) \leq n-2$, which contradicts our assumption that $\dim(x_0, y_0) \geq n-1$. □

Our goal is now to show that y_0 is equal to one of $y_1, ..., y_l$. To do this, we show:

Claim 7.5.6. *The elements $y_1, ..., y_l$ are distinct.*

Proof. Toward a contradiction, assume that $y_i = y_j$ for some $i \neq j$. We work with the set

$$T = S \circ S^{-1} - \Delta.$$

That is, T is the set of all $y \neq y'$ such that for some $x \in M$ we have $(x, y), (x, y') \in S$. Note that since S is non-trivial, $S \circ S^{-1}$ is a non-trivial plane curve. Thus T is either finite or a non-trivial plane curve.

But by definition we have $(\eta_i(t), \eta_j(t)) \in T$ for each t, and thus $(y_i, y_j) \in \overline{T}$. So if $y_i = y_j = y$ for some y, it follows that $(y, y) \in \text{Fr}\ (T)$. In particular T is infinite, and is thus a plane curve. By our assumption T is then almost closed, so we get that

$$\dim(y, y) = \dim y \leq n-2.$$

But by the previous claim we have $(x_0, y) \in S$; moreover, as in the proof of the previous claim, the two coordinates of any point in S are interalgebraic: we conclude that $\dim(x_0, y) \leq n-2$, so $\dim x_0 \leq n-2$, and so $\dim(x_0, y_0) \leq n-2$, again a contradiction. □

Now again by the non-triviality of S, and the fact that $\dim(x_0, y_0) \geq n-1$, we have that $\dim x_0 \geq n-1$; in particular x_0 is $\mathcal{M}$-generic in M, so that x_0 has exactly l extensions to elements of S. By the previous two claims, these extensions must be precisely the (x_0, y_i) for $i = 1, ..., l$. In particular, since $(x_0, y_0) \in S$, we conclude that $y_0 = y_i$ for some $i = 1, ..., l$. But by the first claim above we have $(x_0, y_i) \in S^P$, so we conclude that $(x_0, y_0) \in S^P$, as desired. □

As stated in the introduction to this section, we can now immediately deduce local modularity if all plane curves are almost closed:

Theorem 7.5.7. *If every plane curve in $\mathcal{M}$ is almost closed, then $\mathcal{M}$ is locally modular.*

Proof. If $\mathcal{M}$ is not locally modular, then there is a rank 2 almost faithful family $\mathcal{F}$ of plane curves in $\mathcal{M}$. But then by Lemma 7.5.2, every curve in $\mathcal{F}$ is almost pure; this contradicts Theorem 5.0.1. □

The biggest downside of Lemma 7.5.2 is that it requires *all* plane curves to be *actually* almost closed – indeed, the set T considered in the proof is not guaranteed to be of high enough complexity for Theorem 7.3.1 to apply; and even if it was, Theorem 7.3.1 only applies to generic points in the plane, so clearly not to the point (y, y) under consideration. One could try to instead apply Theorem 7.4.3, but doing so would not glean much since (1) the point (y, y) automatically satisfies the conclusion of Theorem 7.4.3 anyway for non-very ample $\mathcal{M}$, and (2) Theorem 7.4.3 only applies to points with generic coordinates, while there is no reason for the coordinate y under consideration to be generic.

For the remainder of this section, we outline a potential strategy for generating an implication between generic almost closedness and almost purity 'one family at a time' – thus not invoking the set T needed in Lemma 7.5.2. Our strategy is as follows: we assume that we have a large almost faithful family $\mathcal{F} = \{F_a\}_{a \in A}$ of plane curves, each of which has a codimension 1 component, so that these codimension 1 components generically 'sweep' the whole space M^2 to a high order, as in Definition 7.2.1. It then follows easily that these 'sweeping' codimension 1 components intersect along a codimension 2 locus in A^2. We take two such curves F_{a_0} and F_{b_0} which are generic subject to the requirement that their codimension 1 components intersect. Now by assumption, we see an 'extra' intersection point between F_{a_0} and F_{b_0} – and thus, by strong minimality, we should correspondingly see an interection point 'lost' somehow. Now unless this comes from an abnormally strong collision between the frontiers of F_{a_0} and F_{b_0}, we conclude that the family of intersections $R = F \times_{M^2} F$ ramifies at (a_0, b_0).

On the other hand, ramification is a codimension 1 phenomenon, not a codimension 2 phenomenon. We conclude that the locus of 'intersection loss due to extra components intersecting,' which contains (a_0, b_0), is actually contained in a *larger* locus of intersection loss caused by a *different* phenomenon. Now again by strong minimality, this larger locus must see a corresponding intersection *gain* – which again must happen in a different way as that already seen in (a_0, b_0). Our conclusion will be, roughly, that F_{a_0} and F_{b_0} must also have a *second* extra intersection of a *different* type, coming from this larger locus. We thus have *two* special intersection phenomena happening at F_{a_0} and F_{b_0} – and running these phenomena through a dimension computation produces a contradiction. The caveat, of course, is that our original 'lost' intersection could have come through frontier points, which is why we do not have a complete proof; indeed, we present our argument below – see Proposition 7.5.10 – in the form of a proof that certain frontier intersections must happen. However, in the process we will deduce that such frontier intersections can only happen at certain very special points in the plane; we thus hope that future work will obtain a contradiction from the configuration of frontier intersections that we will describe.

Finally, we will note at the end that under certain assumptions, a generically non-almost pure family always gives rise to a family of codimension 1 components which sweep M^2 in the way that we need. It is conceivable that the same holds in full generality, but this did not seem obvious; at the very least, the failure of the assumptions used can be seen as another unusual phenomenon that one can assume in a counterexample.

We now proceed with the argument. As a first illlustration of how to apply Theorem 7.3.1, we begin with the following lemma; roughly, it says that a semi-indistinguishable pair should typically agree on whether they belong to the pure part of a curve containing them:

Lemma 7.5.8. *Let $\mathcal{F} = \{F_a\}_{a \in A}$ be an almost faithful, generically irreducible family of plane curves in $\mathcal{M}$, which has rank $k > 2n$, has no common points, and is definable in $\mathcal{M}$ over a set E. Let $a \in A$ be generic over E, let $x_0 \in M^2$ be generic over E, and let $x_1 \in M^2$ be $\mathcal{F}$-semi-indistinguishable from x_0. Assume that $x_0, x_1 \in F_a$, and*

$$\dim(x_0/E, a) = \dim(x_1/E, a) = n - 1.$$

If $x_0 \in (F_a)^P$, then $x_1 \in (F_a)^P$.

Proof. We may assume that $E = \emptyset$. Since $\mathcal{F}$ is almost faithful with no common points, each point is semi-indistinguishable with finitely many others. Let $I \subset M^2 \times M^2$ be the set of semi-indistinguishable pairs – so $\mathcal{I}$ is $\mathcal{M}$-definable over $\emptyset$, of rank 2, and both projections $I \to M^2$ are finite-to-one. Since x_0 is generic, there is an analytic neighborhood U of x_0 in M^2 on which the projection $I \to M^2$ to the first coordinate is a finite covering.

Now since $\dim(x_0/a) \geq n-1$, x_0 is $\mathcal{M}$-generic in F_a over a. Since F_a is strongly minimal, there is a positive integer k such that for all but finitely many $x \in F_a$, F_a contains exactly k points which are $\mathcal{F}$-semi-indistinguishable from x – note that k is indeed positive since every point is semi-indistinguishable from itself. So, by shrinking U if necessary, we may assume that F_a contains exactly k points which are $\mathcal{F}$-semi-indistinguishable from any given $x \in F_a \cap U$.

Let $T \subset F_a$ be the set of all $x \in F_a$ such that every element of F_a which is $\mathcal{F}$-semi-indistinguishable from x belongs to $(F_a)^P$. Note that T includes every generic element of F_a over a, and so is a fully generic $\mathcal{K}$-definable subset of F_a.

Now if $x_0 \in (F_a)^P$, then since $T \subset F_a$ is fully generic it follows easily that $x_0 \in \overline{T}$; so there is an $\mathbb{R}$-definable curve $\gamma : (0, 1) \to T$ which converges to x_0. Since U is a neighborhood of x_0 in M^2, we may assume that $\gamma(t) \in U$ for each t. We thus obtain $\mathbb{R}$-definable functions

$$\eta_1, ..., \eta_k : (0, 1) \to F_a \cap (F_a)^P$$

which output the distinct elements of F_a that are $\mathcal{F}$-semi-indistinguishable from each $\gamma(t)$.

Since $I \to M^2$ is a finite covering on U, the η_i have distinct limits $y_1, ..., y_k$ in M^2, such that each $(x_0, y_i) \in I$. Note that each $y_i \in (F_a)^P$, since each $\eta_i(t) \in (F_a)^P$ and $(F_a)^P$ is closed.

Now by semi-indistinguishability, each y_i is $\mathcal{M}$-interalgebraic with x_0 over $\emptyset$; in particular, each y_i is generic in M^2 over $\emptyset$, and $\dim(y_i/a) \geq n-1$ for each i. We now make our promised application of Theorem 7.3.1: by the previous remarks, Theorem 7.3.1 implies that no y_i can belong to Fr (F_a). So, since each $y_i \in (F_a)^P$, we conclude that each $y_i \in F_a$.

Finally, recall that F_a contains exactly k points which are $\mathcal{F}$-semi-indistinguishable from x_0; by the above remarks, these k-points are exactly $y_1, ..., y_k$. On the other hand, x_1 is also one of these points; we conclude that $x_1 = y_i$ for some i. Then, since $y_i \in (F_a)^P$, we get that $x_1 \in (F_a)^P$, as desired. □

We also need the following notion:

Definition 7.5.9. If $\mathcal{F} = \{F_a\}_{a\in A}$ is a family of plane curves, then a point $x \in M^2$ is *frontier persistent* if the set $\{a \in A : x \in \text{Fr } (F_a)\}$ has dimension at least $\dim A - n$.

That is, a frontier persistent point is one which belongs to the frontiers of at least as many curves as most points *actually* belong to. Intuitively, we think of such a point as having been 'universally removed' from the family – a process which is only undetected by $\mathcal{M}$ if we then 'universally add' the point in smaller components.

Now the main result of the remainder of this section is:

Proposition 7.5.10. *Let $\mathcal{F} = \{F_a\}_{a\in A}$ be an almost faithful, generically irreducible family of plane curves in $\mathcal{M}$, which has rank $r > 2n$, has no common points, and is definable in $\mathcal{M}$ over a set E. Let $R = \{(x, a, b) \in M^2 \times A^2 : x \in F_a \cap F_b\}$. Assume that for all generic $a \in A$, there is a codimension 1 irreducible component of $\overline{F_a}$ which k-sweeps M^2 over E for some $k > 2n$. Then there are frontier persistent points in $\mathcal{F}$; moreover, one can find separately generic $a, b \in A$ over E, with* $\dim(a, b/E) = 2 \cdot \dim A - 2$, *and points $x, y \in M^2$, so that x belongs to the aforementioned codimension 1 components of both F_a and F_b, y is a frontier persistent point belonging to the frontiers of both F_a and F_b, and $(y, a, b) \in \overline{R}$.*

We make some comments about the statement of the proposition.

- First, the assertion $\dim(a, b/E) = 2\cdot\dim A-2$ says that a and b are 'generic' subject to their codimension 1 components intersecting; that is, we are really saying that 'most of the times' two curves' codimension 1 components intersect, their frontiers also intersect at a frontier persistent point.
- Second, the condition $(y, a, b) \in \overline{R}$ is necessary for the statement to have meaning. Consider, for example, a family which had a common pure point P removed from each of its curves; the point P now belongs to the frontier of almost every curve in the family. However, it is easy to see that a tripe (P, a, b) usually does not belong to $\overline{R}$ in this case. So we are really saying that the intersection of frontiers at y is 'special' in that it can be approximated by actual intersections on nearby pairs of curves.

We now proceed with the argument.

Proof. First, we may assume that $E = \emptyset$. We may also assume that A is stationary in $\mathcal{M}$ – note that restricting to a stationary component does not affect the hypotheses of the proposition, and it is sufficient to prove the conclusion for such a restricted family.

We will proceed to build a tuple (x_0, y_0, a_0, b_0) as described in the statement of the proposition. Before doing this, we give some easy but useful properties of the set R:

Lemma 7.5.11. *Let F be the definable set associated to $\mathcal{F}$, and let R be the set defined in the statement of the proposition. Then:*

1. *R has rank $2r$.*
2. *The projection $R \to A^2$ is almost surjective and almost finite-to-one.*
3. *If $(x, a, b) \in R$ is $\mathcal{M}$-generic, then (x, a) and (x, b) are both $\mathcal{M}$-generic in F.*
4. *If $(x, a, b) \in R$ is generic, then (x, a) and (x, b) are both generic in F.*
5. *If $(x, a, b) \in R^P$, then (x, a) and (x, b) both belong to F^P.*
6. *If (x, a) and (x, b) both belong to F^P, and x is generic in M^2, then $(x, a, b) \in R^P$.*

Proof.

1. Recall that $\mathcal{F}$ has no common points. So if $(x, a, b) \in R$, then rk (a/x) and rk (b/x) are each at most $r - 1$. Combined with the fact that rk $x \leq 2$, we obtain by additivity that rk $(x, a, b) \leq 2r$.

   On the other hand, if $x \in M^2$ is $\mathcal{M}$-generic, then the set ${}_xF = \{a \in A : x \in F_a\}$ has rank $r - 1$. So, choosing (a, b) to be $\mathcal{M}$-generic in $({}_xF)^2$ over x, we obtain a triple (x, a, b) of rank $2r$. Thus rk $R = 2r$.

2. Given (1) and the fact that A is stationary, it suffices to show that rk $(x/a, b) = 0$ whenever $(x, a, b) \in R$ is $\mathcal{M}$-generic. So, let (x, a, b) be such a triple. Then we have rk $(x, a, b) = 2r$ and rk $(x, a) \leq r + 1$, which by additivity implies that

   $$\text{rk } (b/x, a) \geq r - 1 > 0.$$

   In particular, by almost faithfulness, $F_a \cap F_b$ is finite, so that rk $(x/a, b) = 0$, as desired.

3. Let $(x, a, b) \in R$ be $\mathcal{M}$-generic. Since there are no common points in $\mathcal{F}$, we have rk $(b/a, x) \leq r - 1$, so that by additivity

   $$\text{rk } (a, x) \geq 2r - (r - 1) = r + 1 = \text{rk } F,$$

   as desired. By symmetry, the same holds for (x, b).

4. If (x, a, b) is generic in R, then it is in particular optimal. Since (x, a) and (x, b) are $\mathcal{M}$-algebraic over (x, a, b), both pairs are also optimal by Lemma 7.4.2. But by (3) these pairs are both $\mathcal{M}$-generic in F, so by optimality they are both generic in F.

5. If $(x, a, b) \in R^P$, then we can find generic elements $(x', a', b') \in R$ arbitrarily close to (x, a, b). Now for such (x', a', b'), (4) gives that (x', a') and (x', b') are generic in F; this shows that $(x, a), (x, b) \in F^P$.

6. Since x is generic and $(x, a) \in F^P$, Lemma 3.2.7 gives that $a \in ({}_xF)^P$. Similarly, $b \in ({}_xF)^P$. We can thus find a' arbitrarily close to a which is generic in ${}_xF$ over x, and b' arbitrarily close to b which is generic in ${}_xF$ over (x, a'). Then (x, a', b') is generic in R, which shows that $(x, a, b) \in R^P$.

□

We now begin building the desired tuple, starting with a_0, x_0, and b_0. First, let $a_0 \in A$ be generic. Then by assumption there is a codimension 1 component C of $\overline{F_{a_0}}$ which k-sweeps M^2 over $\emptyset$. Let $x_0 \in C$ be generic over a_0. Then let $b_0 \in A$ be an independent realization of $\mathrm{stp}_{\mathcal{K}}(x_0/a_0)$ over a_0 – that is, an independent realization of the $\mathcal{K}$-type of x_0 over $\mathrm{acl}_{\mathcal{K}}(a_0)$.

We note:

- $\dim a_0 = nr$, $\dim(x_0/a_0) = n - 1$, and $\dim(x_0) = 2n$. Indeed, the first two statements follow by definition, and the third follows by k-sweeping.
- From the previous point and additivity, $\dim(a_0/x_0) = n(r - 1) - 1$.
- By the independence of b_0, $\dim(b_0/a_0, x_0) = n(r-1)-1$. In particular, $\dim(b_0/a_0) > 0$, so by almost faithfulness $\dim(x_0/a_0, b_0) = 0$.
- By addivity, $\dim(a_0, b_0) = \dim(a_0, b_0, x_0)$ can now be seen to be $2nr - 2 = 2 \cdot \dim A - 2$.

Since A is stationary, and by (2) of Lemma 7.5.11, there is a positive integer l such that $|F_a \cap F_b| = l$ holds for all generic $(a, b) \in A^2$. We next note:

Lemma 7.5.12. *(a_0, b_0) is $\mathcal{M}$-generic in A^2. In particular, $|F_{a_0} \cap F_{b_0}| = l$.*

Proof. If not then $\mathrm{rk}\ (a_0, b_0) \leq 2r - 1$, so

$$\dim(a_0, b_0) \leq 2nr - n \leq 2nr - 2.$$

But we know that $\dim(a_0, b_0) = 2nr - 2$. We conclude that $n = 2$, $\mathrm{rk}\ (a_0, b_0) = 2r - 1$, and (a_0, b_0) is optimal. But as above, x_0 is $\mathcal{M}$-algebraic over (a_0, b_0), and thus so is (a_0, x_0). So by Lemma 7.4.2, (a_0, x_0) is also optimal. But $\dim(a_0, x_0) = n(r - 1) - 1$ is not a multiple of n, so (a_0, x_0) cannot be optimal, a contradiction. □

Now in a similar manner to the proof of Theorem 7.3.1, our main goal is to study the intersections of F_{a_0} and F_{b_0}. The reader may recall that the proof of Theorem involved a crucial dimension computation, which showed that any intersection point of the two curves at hand was either semi-indistinguishable from a given point or was generic in each curve. The following lemma can be seen as the analogous computation in the current proposition. In essence, it is the main technical point of the proof:

Lemma 7.5.13. *For all $x \in F_{a_0}$, either x is $\mathcal{F}$-semi-indistinguishable from x_0 or we have* $\dim(x/b_0) \geq n$. *Similarly, for all $x \in F_{b_0}$, either x is $\mathcal{F}$-semi-indistinguishable from x_0 or we have* $\dim(x/a_0) \geq n$.

Proof. By construction the tuples (x_0, a_0, b_0) and (x_0, b_0, a_0) realize the same $\mathcal{K}$-type, so the two statements are symmetric. We will thus only prove the second. We do this by way of contradiction: namely, let $x \in F_{b_0}$, and suppose that x and x_0 are not $\mathcal{F}$-semi-indistinguishable and $\dim(x/a_0) \leq n-1$.

Note that

$$\dim(a_0, b_0, x_0, x) \geq \dim(a_0, b_0) = 2nr - 2.$$

Now since x and x_0 are not $\mathcal{F}$-semi-indistinguishable, and both lie on F_{b_0}, we conclude that rk $(b_0/x_0, x) \leq r-2$, and so in particular $\dim(b_0/x_0, x) \leq n(r-2)$. By additivity, it follows that

$$\dim(a_0, x_0, x) \geq (2nr-2) - n(r-2) = nr + 2n - 2. \quad (1)$$

Since $\dim a_0 = nr$, we further conclude by additivity that

$$\dim(x_0, x/a_0) \geq 2n - 2 = 2(n-1).$$

But by assumption each of $\dim(x_0/a_0)$ and $\dim(x/a_0)$ is at most $n-1$: for x_0 this is because $x_0 \in C$, and for x this is part of the statement of the lemma; so, since $\dim(x_0, x/a) \geq 2(n-1)$, we conclude that x_0 and x are independent over a_0. By Lemma 7.2.5, and the fact that C k-sweeps M^2 over $\emptyset$ for some $k > 2n \geq \dim x$, we conclude that x_0 and x are independent over $\emptyset$.

Note also that (1) above is only consistent if $\dim(b_0/x_0, x) = n(r-2)$. Indeed, otherwise the same argument would yield that $\dim(a_0, x_0, x) \geq nr + 2n - 1$, which is impossible given that $\dim(x_0/a_0)$ and $\dim(x/a_0)$ are both at most $n-1$.

Thus we have $\dim(b_0/x_0, x) = n(r-2)$. Combined with the independence of x_0 and x, and the genericity of x_0, we have

$$\dim(b_0, x_0, x) = \dim(x_0, x) + \dim(b_0/x_0, x)$$

$$= 2n + \dim x + n(r-2) = \dim x + nr. \quad (2)$$

On the other hand, the fact that $x \in F_{b_0}$ and there are no $\mathcal{F}$-common points gives us that rk $(b_0/x) \leq r-1$, and thus $\dim(b_0/x) \leq n(r-1)$; so we conclude that

$$\dim(b_0, x_0, x) = \dim x + \dim(b_0/x) + \dim(x_0/b_0, x)$$

$$\leq \dim x + \dim(b_0/x) + \dim(x_0/b_0)$$

$$\leq \dim x + n(r-1) + n - 1 = \dim x + nr - 1. \quad (3)$$

And now combining (2) and (3), we have a contradiction. □

We now immediately conclude:

Lemma 7.5.14. *If $x \in F_{a_0} \cap F_{b_0}$, then exactly one of the following happens:*

1. *x and x_0 are $\mathcal{F}$-semi-indistinguishable, $x \notin (F_{a_0})^P$, and $x \notin (F_{b_0})^P$.*

2. *x is both generic in F_{a_0} over a_0 and generic in F_{b_0} over b_0.*

In particular, x is generic in M^2 over $\emptyset$.

Proof. If x and x_0 are $\mathcal{F}$-semi-indistinguishable, then we get (1) by Lemma 7.5.8. We in addition get by semi-indistinguishability that x and x_0 are $\mathcal{M}$-interalgebraic over $\emptyset$; thus x is generic in M^2 because x_0 is.

If x and x_0 are not $\mathcal{F}$-semi-indistinguishable, then by Lemma 7.5.13 x is generic in each of F_{a_0} and F_{b_0} over the relevant parameter. Now a_0 is generic in A and x is generic in F_{a_0} over a_0; so, for example by Lemma 7.2.3, it follows that x_0 is generic in M^2 over $\emptyset$. □

Now the main geometric content of the proof is the following:

Lemma 7.5.15. *There is at least one $x \in M^2$ such that $(x, a_0, b_0) \in \overline{R} - R$.*

Proof. Suppose no such x exists – that is, whenever $(x, a_0, b_0) \in \overline{R}$ we have $(x, a_0, b_0) \in R$. Since $|F_{a_0} \cap F_{b_0}| = l$, it follows that the fiber $\overline{R}_{(a_0,b_0)}$ has size exactly l. Now by (5) of Lemma 7.5.11, we have $(x_0, a_0, b_0) \in R - R^P$. Thus the fiber $(R^P)_{(a_0,b_0)}$ has size strictly less than l. We next note:

Claim 7.5.16. *If $(x, a_0, b_0) \in R^P$, then (x, a_0, b_0) is smooth in R^P and (a_0, b_0) is smooth in A^2.*

Proof. Since a_0 and b_0 are generic in A, they are each smooth in A; thus the pair (a_0, b_0) is smooth in A^2. Now if $(x, a_0, b_0) \in R^P$, then we get the following:

- $x \in F_{a_0} \cap F_{b_0}$, by assumption – indeed we have $(x, a_0, b_0) \in \overline{R}$, and so $(x, a_0, b_0) \in R$.
- $(x, a_0), (x, b_0) \in F^P$, by (5) of Lemma 7.5.11.
- x is generic in each of F_{a_0} and F_{b_0} over the relevant parameter, by the previous item and Lemma 7.5.14.

Thus (x, a_0) and (x, b_0) are each generic in F, and so also in F^P. It follows that the projection $F^P \to M^2$ is a submersion of smooth manifolds in a neighborhood of each of (x, a_0) and (x, b_0) (see for example Facts 5.1.2 and 5.1.3). Recall that fiber products of smooth manifolds exist along submersions (again see Fact 5.1.2); it follows that R^P is a smooth manifold in a neighborhood of (x, a_0, b_0). That is, (x, a_0, b_0) is smooth in R^P. □

Now by the claim, the fact that M is compact, and the fact that $(R^P)_{(a_0,b_0)}$ has size strictly less than l, it follows that the projection $R^P \to A^2$ (or rather of some component of R^P to $\overline{A^2}$) ramifies at some point (w_0, a_0, b_0): indeed, this follows by a similar argument to the proof of Lemma 7.5.2. First note that $R^P \to \overline{A^2}$ is an l-covering almost everywhere, by

(2) of Lemma 7.5.11. Now since we have assumed there are no frontier points of R at (a_0, b_0), and by the compactness of M there cannot be points at infinity, it follows that two of the sheets must collide above (a_0, b_0). Then, using the claim and (1) of Fact 5.1.5, it follows that there is indeed a ramification point above (a_0, b_0).

So, fix such a ramification point (w_0, a_0, b_0). Then, by the purity of the ramification locus (see Fact 5.1.5), there is an irreducible closed subset W of the ramification locus in R^P which has codimension 1 in R^P and contains (w_0, a_0, b_0). Since W is irreducible, there is a sequence $\{(a_i, b_i, w_i)\}_{i \in \mathbb{Z}^+}$ of generic elements of W which converges to (w_0, a_0, b_0).

Recall that the fiber $(R^P)_{(a_0,b_0)}$ is finite. By the semi-continuity of fiber dimension ([48], Theorem 11.4.2), we may assume that each (w_i, a_i, b_i) belongs to a finite fiber in R^P; in particular, for each i we conclude that

$$\dim(a_i, b_i) = \dim W = 2nr - 1 = \dim(A^2) - 1.$$

We next note:

Claim 7.5.17. *For each i, a_i and b_i are each generic in A.*

Proof. Without loss of generality assume that some a_i is not generic in A. Since W is irreducible and (w_i, a_i, b_i) is generic in W, we conclude that the projection $W \to \overline{A}$ to the left $\overline{A}$-coordinate is not dominant on any top dimensional component of $\overline{A}$; in particular, $W \to \overline{A}$ is not almost surjective. But this contradicts that $(w_0, a_0, b_0) \in W$ and a_0 is generic in A. □

In particular, each a_i is smooth in $\overline{A}$, and thus (a_i, b_i) is smooth in $\overline{A^2}$. So, since the fiber $R^P_{(a_i,b_i)}$ is finite, by Corollary 3.4.9 it has size at most l. On the other hand, since $R^P \to \overline{A^2}$ ramifies at (w_i, a_i, b_i), the fiber $R^P_{(a_i,b_i)}$ in fact has size strictly less than l.

Now since $\dim(a_i, b_i) = \dim(A^2) - 1$, each (a_i, b_i) is $\mathcal{M}$-generic in A^2; so the fiber $R_{(a_i,b_i)}$ has size exactly l. It follows that we can find $z_i \in F_{a_i} \cap F_{b_i}$ for each i, so that each (z_i, a_i, b_i) belongs to $R - R^P$. Note, then, that each (z_i, a_i, b_i) has dimension $2nr - 1$, and is therefore generic in $R - R^P$. Since $R - R^P$ has only finitely many generic $\mathcal{K}$-types, we may assume that each (z_i, a_i, b_i) realizes the same $\mathcal{K}$-type.

Since M is compact, we may assume after passing to a subsequence that the z_i converge to an element $z_0 \in \overline{R}_{(a_0,b_0)}$. Then, since we have assumed there are no frontier points of R above (a_0, b_0), it follows that $z_0 \in F_{a_0} \cap F_{b_0}$. Now our strategy going forward is as follows: using that $(z_i, a_i, b_i) \notin R^P$ for each $i \geq 1$, in addition to Lemma 7.5.14, we will show that z_0 is non-generic in M^2 over $\emptyset$; this is done via a sequence of approximations to the desired statement. Once we have that z_0 is non-generic, we obtain a direct contradiction with Lemma 7.5.14, thereby proving Lemma 7.5.15. We proceed with our first approximation:

Claim 7.5.18. *For each $i \geq 1$, one of (z_i, a_i) and (z_i, b_i) is non-generic in F.*

Proof. Assume that both pairs are generic in F. Then in particular z_i is generic in M^2, and both (z_i, a_i) and (z_i, b_i) belong to F^P. By (6) of Lemma 7.5.11 we thus get $(z_i, a_i, b_i) \in R^P$, contradicting the choice of z_i. □

Without loss of generality, we thus assume from now on that each (z_i, a_i) is non-generic in F. We next conclude:

Claim 7.5.19. *(z_0, a_0) is non-generic in F.*

Proof. Since each (z_i, a_i) for $i \geq 1$ is non-generic in F, and all such (z_i, a_i) realize the same $\mathcal{K}$-type, there is in fact a single non-generic subset $S \subset F$ which is $\mathcal{K}$-definable over $\emptyset$ and contains each (z_i, a_i) for $i \geq 1$. Then $(z_0, a_0) \in \overline{S}$, which shows that (z_0, a_0) is non-generic in F. □

Now by Lemma 7.5.14, we are forced to conclude that z_0 is $\mathcal{F}$-semi-indistinguishable from x_0 and does not belong to $(F_{a_0})^P$ or $(F_{b_0})^P$. Then by Lemma 3.2.7, and the fact that a_0 and b_0 are each generic in A, we get that neither of (z_0, a_0) or (z_0, b_0) belongs to F^P. Since F^P is closed, we may thus assume that $(z_i, a_i), (z_i, b_i) \notin F^P$ for all $i \geq 1$. We therefore conclude:

Claim 7.5.20. *For each $i \geq 1$, z_i is non-generic in M^2.*

Proof. If z_i is generic, then Lemma 3.2.7 gives ${}_{z_i}(F^P) = ({}_{z_i}F)^P$. In particular, neither of a_i or b_i belongs to $({}_{z_i}F)^P$. It follows that $\dim(a_i/z_i)$ and $\dim(b_i/z_i)$ are each at most $n(r-1)-1$, and so by additivity

$$\dim(z_i, a_i, b_i) \leq 2n + 2(n(r-1)-1) = 2nr - 2,$$

contradicting that $\dim(z_i, a_i, b_i) = 2nr - 1$. □

And now finally we conclude:

Claim 7.5.21. *z_0 is non-generic in M^2.*

Proof. Since z_i is non-generic in M^2 for each $i \geq 1$, and all such z_i realize the same $\mathcal{K}$-type, there is a single non-generic subset $T \subset M^2$ which is $\mathcal{K}$-definable over $\emptyset$ and contains z_i for all $i \geq 1$. Then $z_0 \in \overline{T}$, which shows that z_0 is non-generic in M^2. □

Now, as stated above, we have directly contradicted the final statement in Lemma 7.5.14; thus the proof of Lemma 7.5.15 is complete. □

We are finally ready to finish building our desired tuple to prove Proposition 7.5.10. Using Lemma 7.5.15, fix an element $y_0 \in M^2$ such that $(y_0, a_0, b_0) \in \overline{R} - R$. Then it remains to show that y_0 is frontier persistent and $y_0 \in \text{Fr }(F_{a_0}) \cap \text{Fr }(F_{b_0})$. Note that since $(y_0, a_0, b_0) \in \overline{R} - R$, it is clear that (y_0, a_0) and (y_0, b_0) both belong to $\overline{F}$. Now we begin with:

Lemma 7.5.22. $y_0 \in \overline{F_{a_0}} \cap \overline{F_{b_0}}$.

Proof. By symmetry it is enough to show that $y_0 \in \overline{F_{a_0}}$. Now since $(y_0, a_0) \in \overline{F}$, there is an irreducible component U of $\overline{F}$ containing (y_0, a_0). Since a_0 is generic in A, there is also a top dimensional irreducible component V of $\overline{A}$ in which a_0 is generic.

Now it follows that the projection of U to $\overline{A}$ is irreducible and almost contains V; in other words, U projects dominantly to V. Thus Lemma 3.2.7 applies to the projection

$(F \cap U) \to V$, and we conclude that $y_0 \in ((F \cap U)_{a_0})^P$. In particular, $y_0 \in \overline{F_{a_0}}$, as desired. □

Next, since $(y_0, a_0, b_0) \notin R$, it follows that either $y_0 \notin F_{a_0}$ or $y_0 \notin F_{b_0}$. Without loss of generality we assume that $y_0 \notin F_{a_0}$. Then by the previous lemma, we get that $y_0 \in \text{Fr}\ (F_{a_0})$. Then we further show:

Lemma 7.5.23. $y_0 \notin F_{b_0}$.

Proof. Suppose that $y_0 \in F_{b_0}$. Then by Lemma 7.5.13, either y_0 is $\mathcal{F}$-semi-indistinguishable from x_0 or $\dim(y_0/a_0) \geq n$. But $y_0 \in \text{Fr}\ (F_{a_0})$, so $\dim(y_0/a_0) \leq n-1$. It follows that x_0 and y_0 are $\mathcal{F}$-semi-indistinguishable. In particular, y_0 is algebraic over x_0. Now recall that by that, by construction, a_0 and b_0 have the same $\mathcal{K}$-type over $\text{acl}_{\mathcal{K}}(x_0)$. We conclude, in particular, that a_0 and b_0 have the same $\mathcal{K}$-type over y_0. Then, since $y_0 \notin F_{a_0}$, it follows that $y_0 \notin F_{b_0}$, contradicting our assumption. □

So we conclude that $y_0 \in \text{Fr}\ (F_{a_0}) \cap \text{Fr}\ (F_{b_0})$. Finally, we show:

Lemma 7.5.24. *y_0 is frontier persistent.*

Proof. If not, then we have:

- $\dim y_0 \leq 2n$, since $y_0 \in M^2$.
- $\dim(a_0/y_0) \leq n(r-1)-1$, since $y_0 \in \text{Fr}\ (F_{a_0})$.
- $\dim(b_0/y_0) \leq n(r-1)-1$, since $y_0 \in \text{Fr}\ (F_{b_0})$.

By additivity, it follows that

$$\dim(y_0, a_0, b_0) \leq 2n + 2(n(r-1)-1) = 2nr - 2 = \dim(a_0, b_0).$$

In particular, this last inequality is only consistent if all three of the itemized inequalities above are equalities. We conclude that $\dim y_0 = 2n$, and

$$\dim(a_0/y_0) = n(r-1)-1,$$

so that by additivity

$$\dim(y_0, a_0) = n(r+1)-1 = nr+n-1.$$

Since $\dim a_0 = nr$, we further conclude that $\dim(y_0/a_0) = n-1$. To recap, we have that $\dim y_0 = 2n$, so y_0 is generic in M^2; and furthermore, y_0 is a frontier point of F_{a_0} which satisfies $\dim(y_0/a_0) = n-1$. This contradicts Theorem 7.3.1, and thereby proves the lemma. □

Finally, we have shown that y_0 is a frontier persistent point which belongs to the frontiers of both F_{a_0} and F_{b_0}. The proof of Proposition 7.5.10 is now complete. □

We finish this chapter by discussing the extent to which the hypotheses of Proposition 7.5.10 can necessarily be met in a counterexample. We first note the following, which is again an application of Theorem 7.3.1 and Theorem 7.4.3; it roughly says that a generic non-pure point of a curve must stay a non-pure point under a generic composition:

Lemma 7.5.25. *Let $\mathcal{F} = \{F_a\}_{a\in A}$ be an almost faithful, generically irreducible family of non-trivial plane curves in $\mathcal{M}$, which has rank $k > 2n$, has no common points, and is definable in $\mathcal{M}$ over a set E. Let $a \in A$ be generic over E, and let $(x, y) \in F_a - (F_a)^P$ be such that $\dim(x, y/E, a) = n - 1$ and $\dim(x/E) = \dim(y/E) = n$. Let $\mathcal{G} = \{G_b\}_{b\in B}$ be another almost faithful family of non-trivial plane curves which has no common points, and which is also definable in $\mathcal{M}$ over E. Let $b \in B$ be generic over (E, a, x, y), and let $z \in M$ be such that $(y, z) \in G_b$. Then $(x, z) \in G_b \circ F_a$, $\dim(x, z) = 2n$, and $\dim(x, z/a, b) = n - 1$. Moreover, if either $\dim(x, y/E) = 2n$, or $\mathcal{M}$ and $\mathcal{F}$ are both very ample, then $(x, z) \notin (G_b \circ F_a)^P$.*

Proof. We may assume that $E = \emptyset$. Note that $(x, z) \in G_b \circ F_a$ by definition.

Since b is independent from (a, x, y), we have $\dim(x, y/a, b) = n - 1$. Moreover, since F_a and G_b are non-trivial, the elements x, y, and z are all $\mathcal{M}$-interalgebraic over (a, b). We conclude that $\dim(x, z/a, b) = n - 1$. We also note the following:

Claim 7.5.26. *z is generic in M over (a, x, y), and thus $\dim(x, z) = 2n$.*

Proof. By definition we have

$$\dim(b, z/a, x, y) \geq \dim(b/a, x, y) = \dim B.$$

Meanwhile, since $\mathcal{G}$ has no common points and $(y, z) \in G_b$, we have

$$\dim(b/y, z) \leq \dim B - n.$$

By additivity, it follows that $\dim(z/a, x, y) \geq n$, so that z is generic in M over (a, x, y). In particular, z is generic in M over x, which implies that $\dim(x, z) = 2n$. □

Now assume that either $\dim(x, y) = 2n$, or $\mathcal{M}$ and $\mathcal{F}$ are both very ample; it remains only to show in this case that $(x, z) \notin (G_b \circ F_a)^P$. So, suppose toward a contradiction that $(x, z) \in (G_b \circ F_a)^P$. We first show:

Claim 7.5.27. *There is some $y' \in M$ such that $(x, y') \in (F_a)^P$ and $(y', z) \in (G_b)^P$.*

Proof. Since $(x, z) \in (G_b \circ F_a)^P$, there is a sequence $\{(x_i, z_i)\}$ of generic elements of $G_b \circ F_a$ which converges to (x, z). Then for each i there is some $y_i \in M$ such that (x_i, y_i) is generic in F_a and (y_i, z_i) is generic in G_b. Since M is compact, the y_i have a limit point $y' \in M$. Then y' satisfies the desired property. □

Our goal will be to show that (x, y') actually belongs to F_a, and is moreover $\mathcal{F}$-semi-indistinguishable from (x, y); we will then conclude that the points (x, y) and (x, y') give a direct contradiction to Lemma 7.5.8, since $(F_a)^P$ contains one but not the other. We begin with:

Claim 7.5.28. $(y', z) \in G_b$.

Proof. Recall that $\dim y = n$ and b is generic in B over y, thus independent from y. It follows that y is generic in M over b, and so (y, z) is generic in G_b over b. Since G_b is non-trivial, this implies that z is generic in M over b. Then in fact we have

$$\dim(y', z/b) \geq \dim(z/b) = n.$$

Since $(y', z) \in (G_b)^P$, (y', z) is then generic in $(G_b)^P$ over b, and so is also generic in G_b over b. In particular, $(y', z) \in G_b$. □

We also note:

Claim 7.5.29. *y' is $\mathcal{M}$-algebraic over (a, x, y).*

Proof. Since b is generic in B over (a, x, y), it is in particular optimal over (a, x, y). Now since G_b is non-trivial and (y, z) and (y', z) both belong to G_b, it follows that y' is $\mathcal{M}$-algebraic over (b, y), and thus over (a, x, y, b). So, by Lemma 7.4.2, y' is also optimal over (a, x, y).

But since $(x, y') \in (F_a)^P$, we have $\dim(x, y'/a) \leq n$; so since $\dim(x/a) = n - 1$, it follows by additivity that $\dim(y'/a, x, y) \leq 1$. By optimality, and the fact that $n > 1$, we conclude that $\text{rk}\ (y'/a, x, y) = 0$, as desired. □

We conclude:

Claim 7.5.30. *y and y' are $\mathcal{M}$-interalgebraic.*

Proof. By the last claim, b is in fact generic in B over (a, x, y, y'). So, in particular, b is independent from (y, y'). Now since $(y, z), (y', z) \in G_b$, we know that y and y' are $\mathcal{M}$-interalgebraic over b. Then by independence, it follows that y and y' are $\mathcal{M}$-interalgebraic over $\emptyset$. □

We thus further conclude:

Claim 7.5.31. $(x, y') \in F_a$.

Proof. By the previous claim, all of our dimension theoretic assumptions on the triple (a, x, y) also hold of (a, x, y'). In particular, each of x and y' is generic in M, and either $\mathcal{M}$ is very ample or $\dim(x, y') = 2n$.

Now assume that $(x, y') \notin F_a$. Then, since $(x, y') \in (F_a)^P$, it follows that $(x, y') \in \text{Fr}\ (F_a)$. Thus, in either case, the hypotheses of either Theorem 7.3.1 or Theorem 7.4.3 are met: more precisely, we use Theorem 7.3.1 if $\mathcal{M}$ is not very ample, and Theorem 7.4.3 if $\mathcal{M}$ is very ample.

We conclude using either theorem that $\dim(x, y'/a) \leq n - 2$. On the other hand, as stated above, the tuples (x, y, a) and (x, y', a) have identical dimension theoretic properties; it follows that $\dim(x, y/a) \leq n - 2$ as well. But this contradicts the assumptions of Lemma 7.5.25, since it was given that $\dim(x, y/a) = n - 1$. □

And lastly, we conclude:

Claim 7.5.32. *(x, y) and (x, y') are $\mathcal{F}$-semi-indistinguishable, and thus (x, y) and (x, y') are both generic in M^2.*

Proof. By construction we have

$$\text{rk }(a, x, y, b, z, y') = k + 1 + \text{rk } b,$$

while

$$\text{rk }(b, y, y', z) = \text{rk } b + 1.$$

By additivity, it follows that

$$\text{rk }(a, x/b, y, y', z) = k.$$

Since rk $x \leq 1$, we conclude further that

$$\text{rk }(a/x, b, y, y', z) \geq k - 1,$$

and so

$$\text{rk }(a/x, y, y') \geq k - 1,$$

which implies that (x, y) and (x, y') are $\mathcal{F}$-semi-indistinguishable. But $y \neq y'$, since $(F_a)^P$ contains (x, y') but not (x, y). It follows that $\mathcal{F}$ is not very ample. Since we assumed that either $\dim(x, y) = 2n$ or $\mathcal{M}$ and $\mathcal{F}$ are very ample, we conclude that $\dim(x, y) = 2n$, and thus (x, y) is generic in M^2. Then, since y and y' are $\mathcal{M}$-interalgebraic, it follows that (x, y') is also generic in M^2. □

Finally, we recap what we have concluded thus far: we have two $\mathcal{F}$-semi-indistingushable points, (x, y) and (x, y'), each of which is generic in M^2 and belongs to F_a. Moreover, (x, y) has codimension 1 in F_a, and thus so does (x, y') since the two points are interalgebraic. Finally, we have $(x, y') \in (F_a)^P$ and $(x, y) \notin (F_a)^P$. This data directly contradicts Lemma 7.5.8, and so the proof of Lemma 7.5.25 is complete. □

Our main remaining observation in this chapter is now the following:

Corollary 7.5.33. *Assume that $\mathcal{F}$ is an almost faithful, generically irreducible family of non-trivial plane curves which has rank $k > 2n$, has no common points, and is $\mathcal{M}$-definable over a set E. Assume that $b \in A$ is generic over E, and (y, z) is an element of $F_b - (F_b)^P$ which satisfies $\dim(y, z/E) = 2n$ and $\dim(y, z/E, b) = n - 1$. Then there is a family of curves satisfying the hypotheses of Proposition 7.5.10.*

Proof. We may assume that $E = \emptyset$. Our goal is, roughly, to show that the threefold composition $\mathcal{F} \circ \mathcal{F} \circ \mathcal{F}$ suffices. Of course this family need not be almost faithful, so what we really do is extract an almost faithful family from it.

First, let $(a, c) \in A^2$ be generic over (b, y, z). Using non-triviality, we obtain $x, w \in M$ so that $(x, y) \in F_a$ and $(z, w) \in F_c$. Note, then, that x, y, z, w are all $\mathcal{M}$-interalgebraic over

(a, b, c). Now by the independence of (a, c) over (b, y, z), we have that $\dim(y, z/a, b, c) = n-1$; so it follows that each of x, y, z, w has dimension $n-1$ over (a, b, c).

Let $D = F_c \circ F_b \circ F_a$. So D is $\mathcal{M}$-definable over (a, b, c). Now by definition $(x, w) \in D$, and by the previous paragraph $\dim(x, w/a, b, c) = n-1$. Let C be a subset of D of which is $\mathcal{K}$-definable over (a, b, c), contains (x, w), and has dimension $n-1$. By shrinking if necessary, we may assume that the Morley degree of C in $\mathcal{K}$ is as small as possible subject to these requirements, so that any two generic elements of C over (a, b, c) realize the same $\mathcal{K}$-type over (a, b, c).

The main point is then the following:

Claim 7.5.34. *C k-sweeps M^2 over $\emptyset$.*

Proof. Let $(x_1, w_1), ..., (x_k, w_k)$ be independent generic elements of C over (a, b, c). It will suffice to show that these pairs are independent generics in M^2.

Now by the choice of C, each (x_i, w_i) realizes the same $\mathcal{K}$-type as (x, w) over (a, b, c). So by the saturation of $\mathcal{K}$, for each i we can find $y_i, z_i \in M$ such that (x_i, y_i, z_i, w_i) realizes the same $\mathcal{K}$-type as (x, y, z, w) over (a, b, c).

Note then that for each i, the elements x_i, y_i, z_i, w_i are $\mathcal{M}$-interalgbraic over (a, b, c), and so each has dimension $n-1$ and rank 1 over (a, b, c). We now use the independence of the (x_i, w_i) over (a, b, c): letting $\overline{x} = (x_1, ..., x_k)$, and similarly for the other letters, it follows by the aforementioned independence that each of $\overline{x}, \overline{y}, \overline{z}, \overline{w}$ has rank k over (a, b, c). In particular, $(\overline{x}, \overline{y})$ forms a sequence of $\mathcal{M}$-independent $\mathcal{M}$-generics in F_a. Now since $\mathcal{F}$ has rank k, Lemma 7.2.3 gives that F_a k-sweeps M^2 over $\emptyset$; it follows that $(\overline{x}, \overline{y})$ are $\mathcal{M}$-generic and $\mathcal{M}$-independent, and so $\text{rk}\ (\overline{x}, \overline{y}) = 2k$.

Now we easily compute by additivity that

$$\text{rk}\ (a, \overline{x}, \overline{y}) = 2k,$$

using that the $(\overline{x}, \overline{y})$ are $\mathcal{M}$-independent $\mathcal{M}$-generics in F_a. So since $\text{rk}\ (\overline{x}, \overline{y}) = 2k$, we conclude that

$$\text{rk}\ (a/\overline{x}, \overline{y}) = 0.$$

Now it is clear by definition that

$$\text{rk}\ (a, b, c, \overline{x}, \overline{y}, \overline{z}, \overline{w}) = 4k,$$

while

$$\text{rk}\ (b, c, \overline{y}, \overline{z}, \overline{w}) = 3k.$$

Thus we conclude that

$$\text{rk}\ (a, \overline{x}/b, c, \overline{y}, \overline{z}, \overline{w}) = k.$$

Since $\text{rk}\ (a/\overline{x}, \overline{y}) = 0$, this implies that

$$\text{rk}\ (\overline{x}/b, c, \overline{y}, \overline{z}, \overline{w}) = k.$$

That is, $\overline{x}$ is $\mathcal{M}$-generic in M^k over $(b, c, \overline{y}, \overline{z}, \overline{w})$.

Additionally, note by definition that a is generic in A over $(b, c, \overline{w})$. But $\overline{y}$, $\overline{z}$, and $\overline{w}$ are all $\mathcal{M}$-interalgebraic over (b, c), so it follows that a is in fact generic in A over $(b, c, \overline{y}, \overline{z}, \overline{w})$. In particular, a is optimal over $(b, c, \overline{y}, \overline{z}, \overline{w})$. Since $\overline{x}$ is $\mathcal{M}$-algebraic over $(a, b, c, \overline{y}, \overline{z}, \overline{w})$, Lemma 7.4.2 gives that $\overline{x}$ is also optimal over $(b, c, \overline{y}, \overline{z}, \overline{w})$. So, since $\overline{x}$ is $\mathcal{M}$-generic over $(b, c, \overline{y}, \overline{z}, \overline{w})$, it is in fact generic over $(b, c, \overline{y}, \overline{z}, \overline{w})$.

In particular, it follows that $\overline{x}$ is generic in M^k over $\overline{w}$. By a symmetric argument, $\overline{w}$ is generic in M^k over $\overline{x}$. Thus $(\overline{x}, \overline{w})$ is generic in M^{2k}, as desired. □

Now to finish the proof of the corollary, we simply apply Lemma 7.5.25 twice: we first note that $(y, w) \notin (F_c \circ F_b)^P$, by applying Lemma 7.5.25 to $\mathcal{F} \circ \mathcal{F}$ – note in particular that switching the order of the coordinates does not matter. Once this is done, we extract an almost faithful family $\mathcal{G}$ which has as a generic member a strongly minimal component of $F_c \circ F_b$ containing (y, w).

We then perform the exact same process to $\mathcal{G} \circ \mathcal{F}$, extracting an almost faithful family which has a generic curve whose non-pure part contains a generic subset of C which contains (x, w). The resulting family suffices to prove the corollary. □

We end this section with some remarks on the above results.

- If $\mathcal{M}$ and $\mathcal{F}$ are very ample, then we can weaken the hypotheses of Corollary 7.5.33 by only requiring y and z to be *separately* generic in M over E; this is due to the improved statement of Lemma 7.5.25 in the very ample case. Note that the first composition operation might not result in a very ample family, but the resulting point (y, w) will then be generic anyway by Lemma 7.5.25 – so very ampleness is not required for the second composition. It seems quite conceivable that the same weakening is possible even without very ampleness, but we did not see a simple proof.

- In general, if $(y, z) \in F_b$ is a non-pure point of dimension $n - 1$ over b, then (y, z) has codimension at most 1 in M^2; this is easy to see by double counting the dimension of the tuple (b, y, z). So if y and z are not separately generic, then without loss of generality we can assume that y and z are independent, y is generic, and $\dim z = n - 1$. In other words, if the conclusion of Proposition 7.5.10 can be exploited to reach a contradiction, then we reduce the problem of non-pure points to those with at least one coordinate non-generic in $\mathcal{M}$; moreover, we can assume that these are the only codimension 1 non-pure points of generic curves in the family.

- Thus, to briefly summarize the results of this section: we have shown that, if a generic curve in any sufficiently large family has a codimension 1 component whose generic points are generic in the plane, then we can set up the hypotheses of Proposition 7.5.10. Moreover, in this case the conjecture reduces to the scenario that when two curves' sweeping codimension 1 components intersect, generically often their frontiers also intersect at a frontier persistent point, which moreover can be approximated by actual intersection points of nearby pairs of curves. Finally, if no generic curve in a large family has a codimension 1 component whose generic points are generic in the plane, then we can assume that any sufficiently large family has generic curves containing codimension 1 components, and that all of these codimension 1 components are exclusively constrained to non-generic regions of the plane. Consequently, this restriction is preserved under composition of families: that is, no amount of composition is able to 'move' these non-generic points to generic points, without simultaneously adding them to the pure parts of all resulting curves. Of course, this seems quite unlikely, but we were unable to rule it out.

Chapter 8

The Case of Groups

In this final chapter we give an example where the methods of Chapter 7 allow for a full proof of local modularity: namely, we prove the higher dimensional case of the Restricted Trichotomy Conjecture for strongly minimal expansions of groups in characteristic zero. The main goal, of course, is to prove that plane curves are almost pure; once this is accomplished, we apply Theorem 5.0.1.

Section 1 gives a few easy reductions: namely, we show that it suffices to consider very ample expansions of irreducible, commutative algebraic groups. In Section 2 we proceed to apply the results of Chapter 7, deducing in a straightforward manner that every plane curve is almost closed, and in fact has finite frontier.

Once we have almost closedness, it is tempting to conclude almost purity using Lemma 7.5.2. Indeed, we do exactly this in the case of abelian varieties. However, the proof of Lemma 7.5.2 fails for non-compact universes, roughly due to the possibility of 'poles.' Thus our main challenge, which occupies the vast majority of the present chapter, is to place a bound on 'poles' of plane curves. Once we have this bound, we will be able to verify that the proof of Lemma 7.5.2 goes through, and in doing so deduce local modularity.

Our treatment of poles is inspired by, but not identical to, the analogous arguments in [18] and [12]. In [12] the authors work with a strongly minimal expansion of a 2-dimensional group interpreted in an o-minimal field; along the way, they prove that plane curves can have only finitely many poles. They do so, roughly, by finding a link between the poles of a plane curve and the frontier of a certain associated definable set of rank 2. Now, similarly to our situation, they only a priori have results about the frontiers of plane curves, not sets of rank 2; however, assuming an infinitude of poles they were able to build a frontier point with a generic coordinate, and subsequently 'transfer' the frontier property to that coordinate's fiber, effectively reducing to the case of curves.

Such a reduction seems hopeless in our situation: indeed, a similar strategy in higher dimensions does not produce a frontier point with a generic coordinate, but rather at best could produce one with a coordinate of dimension 2 – hence the success of the argument for 2-dimensional groups. Instead we take a different approach, by developing a general result on the frontiers of definable sets of rank 2. This is carried out in Sections 3 and 4, in direct analogy to the corresponding sections of the previous chapter. Once we have this result, we then apply it directly in Section 5 to bound the frontier of the relevant set associated to a plane curve, thus in turn bounding the poles of the curve. Finally, we conclude local modularity in Section 6, using an adaptation of the proof of Lemma 7.5.2.

8.1 Standard Reductions

We begin by reducing the general case of strongly minimal groups to a more precise and desirable scenario. Namely, we note that it suffices to consider strongly minimal expansions of irreducible commutative algebraic groups, which are very ample in a canonical way. There is really nothing new in this section, so we do not include all details. Essentially everything needed is given in [12] in the context of o-minimally definable groups; what follows is simply an adaptation of certain details to the complex algebraic setting.

Recall that an *algebraic group* over an algebraically closed field K is a group object in the category of (potentially reducible) varieties over K. Identifying a variety over K with its set of K-points, as we do throughout, one may think of an algebraic group over K as an actual group $(G, \cdot, ^{-1})$, such that G is a (potentially reducible) variety over K, and the maps $\cdot : G \times G \to G$ and $^{-1} : G \to G$ are given by morphisms of varieties. It is a fact (see, for example, [32], Proposition 1.26 and Remark B.38) that all algebraic groups over K are smooth and quasi-projective.

If G is an algebraic group which is furthermore irreducible and projective, then G is called an *abelian variety*. Among the irreducible algebraic groups over K, the abelian varieties are precisely those which are complete – indeed, this follows by quasi-projectivity. In particular, the complex abelian varieties are precisely those irreducible complex algebraic groups which are compact in the analytic topology.

Common examples of algebraic groups over K include the additive and multiplicative groups of K, $GL_n(K)$ and its Zariski closed subgroups, and any elliptic curve over K; of course one can generate further examples by taking products. The elliptic curves form the prototypical examples of abelian varieties, though we should point out that there are many higher dimensional abelian varieties which are not products of elliptic curves – indeed, Jacobian varieties give such examples.

Any algebraic group over K can of course be interpreted in K. Conversely, it is a well-known theorem of model theory that every group interpretable in an algebraically closed field K is K-definably isomorphic to an algebraic group over K [8]. Thus, when working with a group G interpretable in an algebraically closed field, there is no harm in assuming that G is algebraic.

Now our goal is to show that, in addition to algebraicity, we can assume certain additional convenient properties in our setting. We sum up the desired scenario with the following proposition:

Proposition 8.1.1. *Let $(M, \cdot)$ be a group interpreted in the field structure on $\mathbb{C}$. Assume that the underlying set M has dimension at least 2 as a $\mathbb{C}$-definbale set. Assume further that $\mathcal{M} = (M, \cdot, ...)$ is a reduct of the full $\mathbb{C}$-induced structure on M, which contains the group operation, is strongly minimal, and is non-locally modular. Then there is an expansion of a group $\mathcal{N} = (N, +, ...)$, with the following properties:*

1. *$(N, +)$ is an irreducible, commutative algebraic group of dimension at least 2 over $\mathbb{C}$.*
2. *The structure $\mathcal{N}$ is a reduct of the full $\mathbb{C}$-induced structure on N.*
3. *$\mathcal{N}$ is strongly minimal and non-locally modular.*
4. *There is a non-trivial irreducible plane curve $C \subset N^2$ such the family of translates $\{C + x : x \in N^2\}$ is faithful and very ample.*

Proof. By [45] all strongly minimal groups are commutative, so we may assume that $(M, \cdot)$ is commutative; thus, at this point, we switch to writing the operation as $+$. We also need the following well-known fact:

Fact 8.1.2. *Every infinite algebraic group G over $\mathbb{C}$ has unbounded exponent.*

Proof. It suffices to show that the map $x \mapsto x^n$ is non-constant for each $n > 0$. But this map induces multiplication by n on the tangent space at the identity; in characteristic zero we obtain a vector space isomorphism, which since $\dim G \geq 1$ cannot be the zero map. □

Continuing with the proof of Proposition 8.1.1, we first show that we can satisfy (4). This argument is largely given in detail in section 3 of [12], but we summarize briefly. First we define the *stabilizer* of a plane curve C to be $\{x \in M^2 : C + x \sim C\}$. The stabilizer is always a definable subgroup of M^2. Then we show:

Lemma 8.1.3. *There is an irreducible plane curve C with finite stabilizer.*

Proof. The main point to make it:

Claim 8.1.4. *Let C be any irreducible plane curve whose stabilizer S is infinite. Then C is determined up to almost equality by any generic pair in C^2.*

Proof. Since S is infinite, it is easy to see that C is almost equal to the union of finitely many cosets of S. Since C is irreducible, we in fact conclude that S is strongly minimal, and C is almost equal to a single coset of S. Now let $x, y \in C^2$ be generic and independent over the parameters defining C. Then it follows by the previous remarks that $x - y$ is a generic element of S. Since S is infinite, Fact 8.1.2 gives that it has unbounded exponent; so since $x - y$ is generic in S, it has infinite order. Thus $x - y$ generates an infinite subgroup $S' \leq S$. By strong minimality, S is in fact the only strongly minimal set, up to almost equality, which contains S'. In other words, S is determined up to almost equality by the element $x - y$. But then, since x is generic in C, it is clear that C is almost equal to $S + x$; so in fact C is determined up to almost equality by the pair $(x, x - y)$, and thus also by (x, y). □

By the claim, any irreducible plane curve with infinite stabilizer is determined up to finitely many points by a tuple in M^4. Since $\mathcal{M}$ is non-locally modular, it has irreducible plane curves of arbitrarily high complexity, and thus not all curves have this property. □

Now fix an irreducible plane curve C with finite stabilizer. As shown in [12], one can then pass to the quotient of M by a finite subgroup $H \leq M$, and obtain a strongly minimal group in which the image of C is a non-trivial irreducible plane curve with trivial stabilizer. Finally, it follows by definition that the translates of a plane curve with trivial stabilizer form a faithful, very ample family.

We let $N = M/H$ be the quotient group described above, and $\mathcal{N}$ the full $\mathcal{M}$-induced structure on N. So $\mathcal{N}$ is a strongly minimal expansion of the group $(N, +)$, and there is an $\mathcal{N}$-definable, non-trivial, irreducible plane curve whose translates form a faithful, very ample family. Note that this family is indexed by N^2, so by faithfulness is of rank 2. Thus the structure $\mathcal{N}$ is non-locally modular.

We have now verified (3) and (4). Since $\mathcal{N}$ is interpreted in $\mathcal{M}$, (2) is clear. Finally, we verify (1). First note that N is commutative, since M is. Furthermore, since we only quotiented by a finite subgroup, we have $\dim N = \dim M \geq 2$.

Now since $(N, +)$ is interpreted in $\mathbb{C}$, it is $\mathbb{C}$-definably isomorphic to an algebraic group over $\mathbb{C}$. So we may pass the structure on $\mathcal{N}$ through such a definable isomorphism, and thereby assume without loss of generality that $(N, +)$ is an algebraic group over $\mathbb{C}$; in particular, since the isomorphism is $\mathbb{C}$-definable, we do not lose any of the properties we have thus far guaranteed of the structure $\mathcal{N}$.

It remains only to show that N is irreducible. To this end, we prove the following:

Lemma 8.1.5. *N has no proper subgroups of finite index.*

Proof. Suppose that $S \leq N$ is a subgroup of finite index $k \in \mathbb{Z}^+$. Let $f : N \to N$ be the $\mathcal{N}$-definable map given by multiplication by k, i.e. $x \mapsto k \cdot x$. Since N is commutative, f is an endomorphism of K. Also since N is commutative, N/S is a group of order k; by Lagrange's Theorem applied to N/S, we conclude that im $f \subset S$.

Now by Fact 8.1.2, N has unbounded exponent; so there are infinitely many elements of N which do not belong to ker f. Thus ker $f \leq N$ is coinfinite, and so by strong minimality is in fact finite. It follows that im f is infinite, and so by strong minimality is in fact cofinite. In particular, S is cofinite.

Now let C be any coset of S. Then C is infinite since S is. Since S is cofinite it therefore has infinite intersection with C, and so must be equal to C.

We conclude that S has only one coset, and so $S = N$. □

By Lemma 8.1.5, it follows immediately that N is irreducible: indeed, the connected component N° of the identity is an irreducible algebraic subgroup (see [32], Proposition 1.34 and Corollary 1.35), whose cosets then form the other connected components of N. Since there are only finitely many connected components ([32], Remark 1.33(a)), it follows that N° has finite index, and by Lemma 8.1.5 is therefore equal to N. It follows that N is connected, and thus is moreover irreducible ([32], Corollary 1.35). We have now completed the proof of Proposition 8.1.1. □

In light of Proposition 8.1.1, the following is now justified:

Convention 8.1.6. For the rest of this chapter, we assume that $\mathcal{K}$ is the field of complex numbers, possibly expanded by a countable set of constant symbols. We fix $(M, +)$, an irreducible, commutative, complex algebraic group of dimension $n > 1$, definable over $\emptyset$ in the structure $\mathcal{K}$. We also fix $\mathcal{M} = (M, +, ...)$, a strongly minimal reduct of the full $\mathcal{K}$-induced structure on M, whose atomic relations are $\emptyset$-definable in $\mathcal{K}$ and contain the group operation $+$. We assume that the language of $\mathcal{M}$ contains infinitely many constant symbols, whose interpretations in $\mathcal{M}$ are distinct. We further assume that $\mathcal{M}$ is very ample; in particular, we fix a non-trivial, irreducible plane curve $C \subset M^2$, $\mathcal{M}$-definable over $\emptyset$, and assume that the family of translates of C is faithful and very ample. As in the previous chapters, we use dim to refer to dimension computed in $\mathcal{K}$, and rk for dimension computed in $\mathcal{M}$. Unless otherwise stated, the terms generic and independent are interpreted according to the structure $\mathcal{K}$. Finally, we tacitly assume that all sets of parameters are countable, so that generic points always exist.

In particular, M is smooth and quasi-projective, and $\mathcal{M}$ is non-locally modular. Thus we have all of the assumptions from Chapter 7, and therefore can use the results proved there.

Note that the assumption of non-local modularity is contradictory to our eventual conclusion of local modularity. Thus, when we deduce local modularity later on, we are not showing it directly; indeed, the proper interpretation is that we have assumed non-local modularity and arrived at a contradiction by proving local modularity. We point this out to avoid any potential confusion, because the vast majority of our arguments rely extensively on the presence of large families of plane curves.

Finally, we conclude this section by noting that the assumptions outlined above are not vacuous. Namely, the reader may wonder whether the presence of the group operation on a higher dimensional algebraic group, in conjunction with strong minimality, reduce the problem to a small number of easier cases; indeed, as mentioned above every strongly minimal group is commutative, so there can be no such structure on e.g. $GL_n(\mathbb{C})$.

In fact, we can precisely classify those algebraic groups which are strongly minimal in the pure language of groups. ln [45], Reineke showed that an infiinite pure group is strongly minimal if and only if it is either (1) elementary abelian of prime exponent, or (2) divisible abelian with finitely many elements of each finite order. The case (1) cannot arise over characteristic zero fields, by Fact 8.1.2; on the other hand, using a similar argument to Claim 8.1.5, one can classify those complex algebraic groups which satisfy (2): they are precisely the irreduble commutative algebraic groups of positive dimension. We conclude that any such group is strongly minimal as a pure group, and thus is worthy of consideration for our purposes. For example, the class of strongly minimal algebraic groups includes all abelian varieties, all irreducible one-dimensional groups, and is closed under products – thus contains a plethora of higher dimensional groups.

8.2 Almost Closedness and Abelian Varieties

In this short section we show that plane curves in $\mathcal{M}$ have finite frontier, and subsequently deduce local modularity if $\mathcal{M}$ is an abelian variety. Most of the work was done in Chapter 7, so the finiteness of the frontier is fairly straightforward:

Proposition 8.2.1. *Let $S \subset M^2$ be any plane curve, definable in $\mathcal{M}$ over a set A. If $x_0 \in Fr\ (S)$, then x_0 is $\mathcal{M}$-algebraic over A. In particular, $Fr\ (S)$ is finite.*

Proof. We may assume that $A = \emptyset$; so our goal is to show that x_0 is $\mathcal{M}$-algebraic over $\emptyset$. Let $p \in M^2$ be generic over x_0, and consider the curve $S + p$. Since translation gives an isomorphism, we have $x_0 + p \in \text{Fr}\ (S + p)$.

Now $S+p$ is $\mathcal{M}$-definable over p, which is optimal since it is generic in M^2. Additionally, x_0+p is $\mathcal{M}$-interalgebraic with p over x_0; thus x_0+p is generic in M^2 over x_0, and so is also generic in M^2 over $\emptyset$.

To recap, we have a generic point $x_0 + p \in M^2$, which belongs to the frontier of the plane curve $S + p$, which is defined over the optimal parameter p. Thus Theorem 7.4.3 applies: since $\mathcal{M}$ is very ample, we conclude that $x_0 + p$ is $\mathcal{M}$-algebraic over p. Then since $x_0 = (x_0 + p) - p$, x_0 is also $\mathcal{M}$-algebraic over p. But x_0 and p are $\mathcal{M}$-independent over $\emptyset$, since p is $\mathcal{M}$-generic over x_0. It follows that x_0 is indeed $\mathcal{M}$-algebraic over $\emptyset$, as desired. □

Before moving on to the next section, we point out that we have now solved the case of expansions of abelian varieties:

Theorem 8.2.2. *If M is an abelian variety, then $\mathcal{M}$ is locally modular.*

Proof. By Proposition 8.2.1, and the fact that $\dim M > 1$, every plane curve in $\mathcal{M}$ is almost closed. So, since abelian varieties are projective, the result now follows immediately from Theorem 7.5.7. □

8.3 Frontiers of Generic 2-Hypersurfaces

In the next two sections we study frontiers of '2-hypersurfaces' – rank 2 subsets of M^3. Our conclusion will be that any frontier point of a 'non-trivial' 2-hypersurface D, $\mathcal{M}$-definable over a set A, satisfies a certain $\mathcal{M}$-dependence over A: namely, either the point has rank at most 1 over A, or has rank 2 and each of its coordinates is $\mathcal{M}$-algebraic over the other two. This statement might seem oddly specific; in fact, we are really just trying to rule out frontier points (x, y, z) such that $\text{rk}\ (x, y/A) = 1$ but $\text{rk}\ (x, y, z/A) = 2$. Our motivation comes from the fact that, as we will soon see, frontier points of this form arise naturally from the existence of infinitely many poles in any non-trivial plane curve. Thus, the ruling out of exactly this type of frontier point on 2-hypersurfaces allows us to easily deal with the issue of poles discussed in the introduction.

Our proof strategy is just a more intricate version of the one used Sections 7.3 and 7.4 to study frontiers of plane curves. Namely, in the present section we prove the desired statement for generic points on the frontiers of generic 2-hypersurfaces in large families. In the next section we deduce the same statement for generic points on the frontier of any 2-hypersurface defined over optimal parameters; finally, we then deduce the desired statement in full generality via translation by generic points, exactly as in Proposition 8.2.1.

The structure of proof in the present section is quite similar to Section 7.3 – indeed, the reader may choose to largely skip the proof and pay attention only to the key distinctions from Section 7.3, which we point out as we go. The idea is as follows: given a large family $\mathcal{F} = \{F_a\}_{a \in A}$ of non-trivial 2-hypersurfaces, we study the intersections of the F_a with a carefully chosen family $\mathcal{G} = \{G_b\}_{b \in A}$ of curves in M^3. It will be non-trivial, but possible, to verify that these intersections $F_a \cap G_b$ are indeed generically non-empty and finite – or rather, that there is no harm in assuming so. Once we have this assumption, we show that these intersections usually become smaller when G_b contains a frontier point of F_a that is generic in M^3. We thus detect an $\mathcal{M}$-definable dependence between a generic F_a and its generic frontier points.

Before stating the main result of this section, we begin with some easier topics. To start, we show that our analysis of frontier points of plane curves extends, albeit in a less strong way, to frontier points of curves in M^3:

Lemma 8.3.1. *Let A be a set of parameters.*

1. *If $D \subset M^2$ is $\mathcal{M}$-definable over A and of rank at most 1, then there is a closed set $D' \supset D$ which is also $\mathcal{M}$-definable over A and of rank at most 1. In particular, D' contains $\overline{D}$.*

2. *If $D \subset M^3$ is $\mathcal{M}$-definable over A and of rank at most 1, then there is a closed set $D' \supset D$ which is also $\mathcal{M}$-definable over A and of rank at most 1. In particular, D' contains $\overline{D}$.*

Proof.

1. If D is finite then we take $D' = D$. Otherwise D is a plane curve, so by Proposition 8.2.1 the frontier of D is contained in $\operatorname{acl}_{\mathcal{M}}(A)$. In particular, there is a formula $\phi(w)$, with parameters contained in A and with only finitely many solutions, such that every element of Fr (D) satisfies ϕ. We thus take D' to be the union of D and the solution set of ϕ.

2. Le π be any of the three projections $\pi : M^3 \to M^2$. Since rk $D \leq 1$ it follows that rk $(\pi(D)) \leq 1$. So by (1) there is a closed set $E_\pi \supset \pi(D)$, which is $\mathcal{M}$-definable over A and of rank at most 1.

   Let D' be the set of all $w \in M^3$ such that $\pi(w) \in E_\pi$ for each of the three projections $\pi : M^3 \to M^2$. Then D' is closed, $\mathcal{M}$-definable over A, and contains D. It remains to see that rk $D' = 1$. But the fact that each E_π has rank at most 1 implies that any two coordinates of a point in D' have combined rank at most 1 over A; on the other hand, if there were a point in D' with rank at least 2 over A, then some two of its coordinates would have to be $\mathcal{M}$-generic and independent over A. Thus no such points exist, and we are done.

□

We now proceed to discuss 2-hypersurfaces:

Definition 8.3.2. A 2-*hypersurface* in $\mathcal{M}$ is an $\mathcal{M}$-definable subset $D \subset M^3$ which has rank 2. We say that D is *non-trivial* if each of the projections $D \to M^2$ is almost surjective and almost finite-to-one. A definable family $\mathcal{F} = \{F_a\}_{a \in A}$ of 2-hypersurfaces is *almost faithful* if for each $a \in A$, the inequality rk $(F_a \cap F_b) \leq 1$ holds for all but finitely many $b \in A$. In this case we define rk $\mathcal{F}$ to be the rank of the set A. The family $\mathcal{F} = \{F_a\}_{a \in A}$ of 2-hypersurfaces is *generically stationary* if for all $\mathcal{M}$-generic $a \in A$, the set F_a is stationary in $\mathcal{M}$.

As in the case of plane curves, since there are infinitely many constants in the language, any stationary 2-hypersurface can be realized up to almost equality as an $\mathcal{M}$-generic member of a $\emptyset$-definable, generically stationary, almost faithful family indexed by a definable set: one first uses the Compactness Theorem to find a faithful family over an interpretable set using a sufficient fragment of the type of the code of D, then reparametrizes almost faithfully using the same argument as in Fact 2.3.17.

We proceed to develop some basic facts about almost faithful families of 2-hypersurfaces. First, we discuss common points:

Definition 8.3.3. Let $\mathcal{F} = \{F_a\}_{a \in A}$ be an almost faithful family of 2-hypersurfaces in $\mathcal{M}$. Then a point $w \in M^3$ is $\mathcal{F}$-*common* if $\{a \in A : w \in F_a\}$ is generic in A.

It is possible for a family of 2-hypersurfaces to have infinitely many common points, if all 2-hypersurfaces in the family share a common rank 1 subset. On the other hand, for positive rank families the set of common points always has rank at most 1, as shown by the following lemma:

Lemma 8.3.4. *If $\mathcal{F} = \{F_a\}_{a \in A}$ is an almost faithful family of 2-hypersurfaces in $\mathcal{M}$, and rk $\mathcal{F} \geq 1$, then the set of $\mathcal{F}$-common points has rank at most 1.*

Proof. We may assume that $\mathcal{F}$ is $\emptyset$-definable in $\mathcal{M}$. Let w be $\mathcal{F}$-common. Then we can find a pair $(a, b) \in A^2$, generic in A^2 over w, with $w \in F_a \cap F_b$. Since rk $\mathcal{F} \geq 1$, neither of a and b is $\mathcal{M}$-algebraic over the other; in particular, by almost faithfulness, $F_a \cap F_b$ has rank at most 1. Thus rk $(w/a, b) \leq 1$. But since (a, b) is generic over w, it follows that (a, b) and w are $\mathcal{M}$-independent over $\emptyset$, so that in fact rk $w \leq 1$. □

In particular, the rank of the common points is less than the rank of each 2-hypersurface in the family; this allows us to 'remove' the common points $\mathcal{M}$-definably without changing most properties of the family – thus effectively letting us assume that such points don't exist.

We next discuss sweeping; similarly to Corollary 7.2.6, we show:

Lemma 8.3.5. *Let $\mathcal{F} = \{F_a\}_{a\in A}$ be an almost faithful family of 2-hypersurfaces in $\mathcal{M}$. Assume that rk $\mathcal{F} = k > 0$, and $\mathcal{F}$ is definable in $\mathcal{M}$ over a set E. Then for all generic $a \in A$ over E, the set F_a k-sweeps M^3 over E.*

Proof. We may assume that $E = \emptyset$. Let $a \in A$ be generic, and let $(w_1, ..., w_k)$ be generic in $(F_a)^k$ over a. Then rk $(a, w_1, ..., w_k) = 3k$ and $\dim(a, w_1, ..., w_k) = 3nk$. Let b be an independent realization of $\text{tp}(a/w_1, ..., w_k)$, where the type is taken over the full field structure.

Claim 8.3.6. *rk $(a, b, w_1, ..., w_k) = 3k$.*

Proof. That rk $(a, b, w_1, ..., w_k) \geq 3k$ is clear. Now if rk $(F_a \cap F_b) \leq 1$ then

$$\text{rk } (a, b, w_1, ..., w_k) = \text{rk } (a, b) + \text{rk } (w_1, ..., w_k/a, b)$$

$$\leq 2k + k(1) = 3k.$$

Otherwise rk $(F_a \cap F_b) = 2$. Then by almost faithfulness we get rk $(b/a, w_1, ..., w_k) = 0$, and the statement of the claim is clear. □

By the claim, we have $\dim(b/a, w_1, ..., w_k) = 0$. Now by definition of b, it follows that in fact $\dim(a/w_1, ..., w_k) = 0$, so that by additivity $\dim(w_1, ..., w_k) = 3nk$. Thus $x_1, ..., x_k$ are independent generics in M^3. □

In particular, if rk $\mathcal{F} \geq 1$, then a generic point on a generic curve from $\mathcal{F}$ is always generic in M^3, which proves the lemma.

We now state the main result of this section:

Proposition 8.3.7. *Suppose that $\mathcal{F} = \{F_a\}_{a\in A}$ is an almost faithful, generically stationary family of non-trivial 2-hypersurfaces, which is definable in $\mathcal{M}$ over a set E, and has rank $k > 3n$. Let $a_0 \in A$ be generic over E, and let $w_0 = (x_0, y_0, z_0) \in Fr\ (F_{a_0})$. Assume that w_0 is generic in M^3 over E. If rk $(x_0, y_0/E, a_0) \geq 1$, then rk $(z_0/E, a_0, x_0, y_0) = 0$.*

Proof. Similarly to the proof of Theorem 7.3.1, the proof will be quite long. Thus, as in the proof of Theorem 7.3.1, we begin with a glossary of the objects that will appear, along with where each one is introduced. We will use:

- The family $\mathcal{F} = \{F_a\}_{a\in A}$ of rank k, the index $a_0 \in A$, and the point $w_0 = (x_0, y_0, z_0)$ which belongs to the frontier of F_{a_0}, all given in the statement of the proposition. Moreover, the associated definable set $F \subset M^3 \times A$.

- A fixed non-trivial, irreducible plane curve $C \subset M^2$, whose translates form a faithful and very ample family. Recall that C was introduced in Convention 8.1.6.

- A family $\mathcal{G} = \{G_b\}_{b\in B}$ of curves in M^3, along with the definable set $G \subset M^3 \times B$ associated to $\mathcal{G}$, which are introduced in Definition 8.3.10. In fact, the set B is just M^4.

- An index $b_0 = (p_0, q_0, r_0, s_0) \in B$ of a generic curve in $\mathcal{G}$ containing w_0, introduced after the proof of Lemma 8.3.13.
- The fiber product R of F and G over M^3, introduced in Definition 8.3.18.
- A set $D \subset M^4$ which is $\mathcal{M}$-definable over (p_0, r_0) and isomorphic to M^2; moreover, the subfamily $\mathcal{H}$ of $\mathcal{G}$ indexed by D, and the associated definable set $H \subset M^3 \times D$. These are all introduced in Definition 8.3.18.
- The fiber product R_H of F and H over M^3, introduced in Definition 8.3.18.
- The set D_{a_0} of all $d \in D$ such that $F_{a_0} \cap G_d$ is non-empty, also introduced in Definition 8.3.18.
- The positive integer l, which is the generic number of intersection points between sets in $\mathcal{F}$ and $\mathcal{H}$, introduced after the proof of Lemma 8.3.24.
- The points $w_1, ..., w_l \in M^3$, which are the distinct intersection points of F_{a_0} and G_{b_0}, introduced after the proof of Lemma 8.3.25.

As in the proof of Theorem 7.3.1, our strategy will be to show that a_0 and b_0 cannot be $\mathcal{M}$-independent, by analyzing fiber sizes in the projections $R \to A \times B$ and $R_H \to A \times D$. We will then transfer the resulting dependence to a dependence of w_0 over a_0 to prove the proposition. We now proceed with the argument:

First, we may assume that $E = \emptyset$. We may also assume that A is stationary in $\mathcal{M}$, since otherwise we could replace it with a stationary component in which a_0 is generic. In particular, the resulting subfamily would be $\mathcal{M}$-definable over $\text{acl}_{\mathcal{M}}(\emptyset)$, so such an assumption is harmless.

We also note the following two facts:

Lemma 8.3.8. *We may assume that rk $(w_0/a_0) \geq 2$. In particular, we may assume that w_0 does not belong to the closure of any set of rank at most 1 which is $\mathcal{M}$-definable over a_0.*

Proof. The statement of Proposition 8.3.7 is only false if rk $(x_0, y_0/a_0)$ and rk $(z_0/a_0, x_0, y_0)$ are both positive, in which case by additivity we have rk $(w_0/a_0) \geq 2$. So this assumption is justified. The second statement in the lemma now follows immediately from Lemma 8.3.1. □

Corollary 8.3.9. *We may assume that $\mathcal{F}$ has no common points.*

Proof. We argue similarly as in the case of plane curves. In any case, the set Q of $\mathcal{F}$-common points is $\emptyset$-definable in $\mathcal{M}$, and by Lemma 8.3.4 it is non-generic in each F_a; so we simply define a new family by replacing each F_a with $F_a - Q$.

To see that this is harmless, the only non-trivial thing to verify is that $w_0 \in \text{Fr}\ (F_{a_0} - Q)$. Indeed, otherwise we would have $w_0 \in \overline{Q}$, contradicting Lemma 8.3.8. □

We next define the family $\mathcal{G}$ of curves that we will use to intersect with the F_a. The definition of $\mathcal{G}$ is specific and important, so we emphasize it below. Recall that we have a fixed $\emptyset$-definable, non-trivial, irreducible plane curve C whose translates form a faithful, very ample family.

Definition 8.3.10. Let $B = M^4$. We define a family of curves $\mathcal{G} = \{G_b\}_{b \in B}$ as follows: for $b = (p, q, r, s)$, let

$$G_b = \{(x, y, z) \in M^3 : (x, y) \in C + (p, q) \wedge (y, z) \in C + (r, s)\}.$$

As stated above, our goal is to study the intersections of the sets in the families $\mathcal{F}$ and $\mathcal{G}$. However, because $\mathcal{G}$ is not a family of plane curves, we do not have a body of results at our disposal concerning the behavior of the G_b. Thus, before proceeding, we need to develop some some basic information about the family $\mathcal{G}$. First, and most obviously, we check that each G_b is in fact a curve:

Lemma 8.3.11. *Let $\mathcal{G}$ be defined as above.*

1. *For each $b \in B$, the projection $G_b \to M$ to the middle coordinate is finite-to-one.*

2. *For each $b \in B$, the projection $G_b \to M$ to the middle coordinate is almost surjective.*

3. *For each $b \in B$, rk $G_b = 1$.*

Proof.

1. Since C is a non-trivial plane curve, so is each of its translates. Now fix any $b = (p, q, r, s) \in B$ and any $y \in M$. Then $C + (p, q)$ is non-trivial, so there are only finitely many x with $(x, y) \in C + (p, q)$. Similarly, since $C + (r, s)$ is non-trivial, there are only finitely many z with $(y, z) \in C + (r, s)$. Thus in total there are only finitely many (x, z) with $(x, y, z) \in G_b$.

2. Let $y \in M$ be generic over b. Since $C + (p, q)$ is a non-trivial plane curve, there is some $x \in M$ with $(x, y) \in C + (p, q)$. Similarly, since $C + (r, s)$ is non-trivial, there some $z \in M$ with $(y, z) \in C + (r, s)$. Thus $(x, y, z) \in G_b$, which shows that y is in the image of G_b. Since y is generic in M over b, it follows that the projection $G_b \to M$ is almost surjective.

3. Immediate from (1) and (2).

□

So, $\mathcal{G}$ is indeed a family of curves in M^3. We next fix the following notation, in analogy to the same notation for families of plane curves:

Notation 8.3.12. For $w \in M^3$, we let ${}_wG$ denote the set $\{b \in B : w \in G_b\}$.

Now we proceed to note the following basic properties of the sets ${}_wG$:

Lemma 8.3.13. *$\mathcal{G}$ has no common points and is 'almost very ample,' in the following senses:*

1. *For each $w = (x, y, z) \in M^3$, ${}_wG$ is a stationary set of rank 2, and in fact equals*
$$((x, y) - C) \times ((y, z) - C).$$
2. *Let $w_1 = (x_1, y_1, z_1)$ and $w_2 = (x_2, y_2, z_2)$ be any two distinct points in M^3. Then ${}_{w_1}G \cap {}_{w_2}G$ has rank at most 1. Moreover, if w_1 and w_2 share at most one common coordinate, then ${}_{w_1}G \cap {}_{w_2}G$ is finite.*

Proof.
1. By definition we have $(x, y) \in C + (p, q)$ if and only if $(p, q) \in (x, y) - C$. Similarly, $(y, z) \in C + (r, s)$ if and only if $(r, s) \in (y, z) - C$. So ${}_wG$ is indeed just
$$((x, y) - C) \times ((y, z) - C).$$
In particular, ${}_wG$ is isomorphic to $C \times C$, and so is stationary and of rank 2.
2. By definition ${}_xG \cap {}_yG$ consists of those (p, q, r, s) such that each $(x_i, y_i) \in C + (p, q)$ and each $(y_i, z_i) \in C + (r, s)$. Now we use the fact that the translates of C are very ample: given any two points $u_1, u_2 \in M^2$, the set of $v \in M^2$ with each $u_i \in C + v$ is either finite if $u_1 \neq u_2$, or equal to the rank 1 set $u - C$ if $u_1 = u_2 = u$. So it suffices to note that, if $w_1 \neq w_2$, then at least one of the pairs $(x_1, y_1), (x_2, y_2)$ and $(y_1, z_1), (y_2, z_2)$ consists of two distinct points; and moreover, if w_1 and w_2 have at most one common coordinate, then both of these pairs consist of two distinct points.

□

We now turn toward the proof of Proposition 8.3.7. In contrast to the proof of Theorem 7.3.1, we will argue by way of contradiction. That is:

Convention 8.3.14. For the remainder of the proof of Proposition 8.3.7, we assume toward a contradiction that both rk $(x_0, y_0/a_0)$ and rk $(z_0/a_0, x_0, y_0)$ are at least 1.

Note, in particular, that by additivity rk $(w_0/a_0) \geq 2$: indeed, we already assumed this in Lemma 8.3.8, in order to prove Corollary 8.3.9; but it now follows anyway from Convention 8.3.14.

We now define the next major object that we will use: let $b_0 = (p_0, q_0, r_0, s_0)$ be a generic element of ${}_{w_0}G$ over (a_0, w_0). Then we immediately conclude:

- By Lemma 8.3.13, $\dim(b_0/a_0, w_0) = \dim(b_0/w_0) = 2n$, and thus $\dim(w_0, b_0) = 5n$.
- Since $w_0 \in G_{b_0}$ we have $\dim(w_0/b_0) \leq n$. So since $b_0 \in M^4$, the above item is only consistent if $\dim(w_0/b_0) = n$ and $\dim b_0 = 4n$.
- In particular, we conclude that b_0 is generic in B, and w_0 is generic in G_{b_0} over b_0.

Now moving forward, the first main challenge of the proof, which is quite distinct from Theorem 7.3.1, is the following:

Lemma 8.3.15. *The pair (a_0, b_0) is $\mathcal{M}$-generic in $A \times B$.*

Proof. We know by assumption that a_0 is generic in A; so we need to show that b_0 is $\mathcal{M}$-generic in B over a_0.

Recall, as in Convention 8.3.14, that we are assuming each of the values rk $(x_0, y_0/a_0)$ and rk $(z_0/a_0, x_0, y_0)$ is at least 1. In particular, rk $(y_0, z_0/a_0, x_0, y_0)$ is also at least 1. Thus we can find $\mathcal{M}$-definable sets $S, T \subset M^2$ with the following properties:

- S and T have positive rank.
- S is $\mathcal{M}$-definable over a_0, and T is $\mathcal{M}$-definable over (a_0, x_0, y_0).
- (x_0, y_0) is $\mathcal{M}$-generic in S over a_0, and (y_0, z_0) is $\mathcal{M}$-generic in T over (a_0, x_0, y_0).

We view S and T as index sets for families $S - C$ and $T - C$ of translates of $-C$, by taking the plane curves $s - C$ and $t - C$ for $s \in S$ and $t \in T$. Note that these families are faithful by the choice of C, and thus each family has positive rank, by the choice of S and T. Moreoever, note that $(x_0, y_0) - C$ and $(y_0, z_0) - C$ are $\mathcal{M}$-generic curves in these respective families over the relevant parameters.

Now by definition, and by Lemma 8.3.13, the tuple $b_0 = (p_0, q_0, r_0, s_0)$ is chosen so that (p_0, q_0) is generic in $(x_0, y_0) - C$ over (a_0, x_0, y_0), and (r_0, s_0) is generic in $(y_0, z_0) - C$ over $(a_0, x_0, y_0, z_0, p_0, q_0)$. Now in any faithful family of plane curves of positive rank, an $\mathcal{M}$-generic element of an $\mathcal{M}$-generically indexed curve is $\mathcal{M}$-generic in the plane: indeed, otherwise there would be a strongly minimal component common to generically many curves in the family, contradicting faithfulness (one can show this in detail, for example, by using exactly the same proof as in Lemma 7.1.3 (3)). It follows, in particular, that

$$\text{rk } (p_0, q_0/a_0) = \text{rk } (r_0, s_0/a_0, p_0, q_0) = 2,$$

and thus by additivity

$$\text{rk } (p_0, q_0, r_0, s_0/a_0) = \text{rk } (b_0/a_0) = 4,$$

as desired. □

We now return to mimicking the proof of Theorem 7.3.1. Our next task is to study the intersection points of F_{a_0} and G_{b_0}. To do this, we first let $F \subset M^3 \times A$ and $G \subset M^3 \times B$ be the definable sets associated to $\mathcal{F}$ and $\mathcal{G}$. Then we prove the following, which is in direct analogy to Lemma 7.3.6:

Lemma 8.3.16. *Let $w = (x, y, z) \in F_{a_0} \cap G_{b_0}$. Then (w, a_0) is generic in F and (w, b_0) is generic in G.*

Proof. Our task is equivalent to showing that w is generic in both F_{a_0} and G_{b_0}, in each case over the relevant parameter. Note that $w_0 \neq w$, since $w \in F_{a_0}$ and $w_0 \notin F_{a_0}$. By Lemma 8.3.13, and the fact that $w_0, w \in G_{b_0}$, we conclude that rk $(b_0/w_0, w) \leq 1$. Now the main fact we need, before we can mimic the proof of Lemma 7.3.6, is that we can replace 1 with 0 in this last inequality:

Claim 8.3.17. *rk* $(b_0/w_0, w) = 0$.

Proof. If not, then by Lemma 8.3.13 it follows that w_0 and w share two coordinates. The main observation, then, is that the non-triviality of F_{a_0} forces rk $(w/a_0, w_0) = 0$. Indeed, since all projections $F_{a_0} \to M^2$ are almost finite-to-one, it follows that only finitely many elements of M^2 can have infinite fibers in any of these projections: otherwise the union of the infinite fibers would have rank 2, contradicting the definition of almost finite-to-one.

On the other hand, it is easy to see that our assumptions on w_0 – namely, that both rk $(x_0, y_0/a_0)$ and rk $(z_0/a_0, x_0, y_0)$ are positive – imply that any two coordinates of w_0 combined have positive rank over a_0; thus, by the above paragraph, any two coordinates of w_0 can be extended to only finitely many elements of F_{a_0}. But w is one such extension, since w_0 and w share two coordinates; so we conclude that rk $(w/a_0, w_0) = 0$.

It follows, in particular, that

$$\text{rk } (b_0/a_0, w_0) = \text{rk } (b_0/a_0, w_0, w) \leq \text{rk } (b_0/w_0, w) \leq 1,$$

and thus $\dim(b_0/a_0, w_0) \leq n$. But by the choice of b_0 we have $\dim(b_0/a_0, w_0) = 2n$; so we reach a contradiction, and thus prove the claim. □

In light of Claim 8.3.17, we are now in a position to mimic our work for plane curves. Since rk $(b_0/w_0, w) = 0$, we get

$$\dim(w/a_0, w_0) = \dim(b_0, w/a_0, w_0)$$

$$\geq \dim(b_0/a_0, w_0) = 2n = \dim F_{a_0}.$$

Thus w is generic in F_{a_0} over (a_0, w_0). We immediately conclude that (a_0, w) is generic in F. By Lemma 8.3.5, and the fact that $k \geq 1$, we also conclude that w is generic in M^3 over $\emptyset$.

Additionally, since w is generic in F_{a_0} over (a_0, w_0), it is in particular independent from w_0 over a_0. But w is generic in F_{a_0}, and by Lemma 8.3.5 the set F_{a_0} has k-sweeping over $\emptyset$. Since $k > 3n \geq \dim w$, Lemma 7.2.5 applies: we conclude that w_0 and w are independent over $\emptyset$; as both points are generic in M^3, this implies that $\dim(w_0, w) = 6n$.

Finally, it now follows that $\dim(w_0, w, b_0) \geq 6n$. Since $\dim b_0 = 4n$, we get by additivity that $\dim(w_0, w/b_0) \geq 2n$. Since $w_0, w \in G_{b_0}$ and $\dim G_{b_0} = n$, this is only possible if w_0 and w are independent generics in G_{b_0}. In particular, w is generic in G_{b_0} over b_0, and so (w, b_0) is generic in G. □

We have now concluded the main property that we need of the intersection points of F_{a_0} and G_{b_0}. As in the proof of Theorem 7.3.1, we next define a subfamily $\mathcal{H}$ of $\mathcal{G}$ containing G_{b_0}, whose projection to the relevant power of M will be finite-to-one. The definition in this case is more concrete, but necessarily more specific. Out of necessity we use certain variables differently than in Theorem 7.3.1: the pair (p_0, r_0) will play the role of the 'z' from Theorem 7.3.1, and the set D defined below will play the role of the set 'S' from Theorem 7.3.1. Note that we also consider two fiber products as opposed to one; thus we distinguish the fiber product using the subfamily $\mathcal{H}$ with a subscript H.

Definition 8.3.18. We define the sets R, D, H, R_H, and D_{a_0} as follows:

1. We define $R \subset M^3 \times A \times B$ as $\{(w, a, b) : w \in F_a \cap G_b\}$.
2. We define $D \subset M^4$ as $\{(p_0, q, r_0, s) : (q, s) \in M^2\}$.
3. We define $\mathcal{H}$ to be the subfamily of $\mathcal{G}$ indexed by D. That is, $\mathcal{H}$ is the family whose associated definable set $H \subset M^3 \times D$ is given by $H = \{(w, d) : w \in G_d\}$.
4. We define $R_H \subset M^3 \times A \times D$ to be the restriction of R to $A \times D$: that is, we set
$$R_H = \{(w, a, d) \in M^3 \times A \times D : w \in F_a \cap G_d\}.$$
5. We define D_{a_0} be the set of all $d \in D$ such that $F_{a_0} \cap G_d$ is non-empty.

So D is naturally identifiable with M^2, and in fact the two sets are isomorphic. Note that R is $\mathcal{M}$-definable over $\emptyset$; D, $\mathcal{H}$, and R are $\mathcal{M}$-definable over (p_0, r_0); and D_{a_0} is $\mathcal{M}$-definable over (p_0, r_0, a_0).

As in the proof of Theorem 7.3.1, we proceed to establish some basic properties of the sets H and R_H. This is another instance where our proof is more complicated than the plane curve case: namely, it is not quite as easy to show that the fibers in R_H are generically non-empty and finite. Thankfully, for H our work is still straightforward:

Lemma 8.3.19. *H has rank 3, and the projection $H \to M^3$ is almost surjective and everywhere finite-to-one.*

Proof. It is clear that H has rank 3, since $\mathcal{H}$ is a family of rank 1 sets indexed by a rank 2 set. So it suffices to show that the projection $H \to M^3$ is everywhere finite-to-one. Indeed, fix any $w = (x, y, z) \in M^3$. Then by Lemma 8.3.13 we have
$${}_wG = ((x, y) - C) \times ((y, z) - C).$$
Now since C is non-trivial, the plane curves $(x, y) - C$ and $(y, z) - C$ are also non-trivial. In particular, each of these curves has only finitely many elements of the forms (p_0, q) and (r_0, s). Thus w belongs to only finitely many G_d with $d \in D$, which shows that the fiber above w in H is finite. □

So we have a nice analysis of the set H. We next proceed to work toward the analogous conclusions for R_H. The following straightforward verifications will be useful:

Lemma 8.3.20. *The following points are generic in the respectively given senses:*

1. *(p_0, r_0) is generic in M^2 over $\emptyset$.*
2. *(w_0, b_0) is generic in H over (p_0, r_0). In particular, w_0 is generic in M^3 over (p_0, r_0).*
3. *(a_0, b_0) is $\mathcal{M}$-generic in $A \times D$ over (p_0, r_0).*

Proof. 1. This is clear, since $b_0 = (p_0, q_0, r_0, s_0)$ is generic in M^4 over $\emptyset$.

2. Recall that $\dim(w_0, b_0) = 5n$. Then since p_0 and r_0 are among the coordinates of b_0, we get
$$\dim(w_0, b_0, p_0, r_0) = 5n.$$
Now by (1) we have $\dim(p_0, r_0) = 2n$, so by additivity we get
$$\dim(w_0, b_0/p_0, r_0) = 3n.$$
Thus (w_0, b_0) is generic in H over (p_0, r_0). It now follows by Lemma 8.3.19 that w_0 is generic in M^3 over (p_0, r_0).

3. By Lemma 8.3.15 we have rk $(a_0, b_0) = k+4$. Since p_0 and r_0 are among the coordinates of b_0, we get
$$\text{rk }(a_0, b_0, p_0, r_0) = k + 4$$
as well. Since rk $(p_0, r_0) = 2$, this implies that
$$\text{rk }(a_0, b_0/p_0, r_0) = k + 2,$$
as desired.

□

Recall that we defined D_{a_0} as the set of all $d \in D$ such that $F_{a_0} \cap G_d$ is non-empty. Now the key observation toward our goal is the following:

Lemma 8.3.21. $b_0 \in \overline{D_{a_0}}$.

Proof. By Lemma 8.3.19, the projection $H \to M^3$ is locally surjective near any generic element of H; in particular, it is locally surjective near (w_0, b_0).

Now since $w_0 \in \text{Fr }(F_{a_0})$, we can find $w' \in F_{a_0}$ arbitrarily close to w_0. By local surjectivity, as $w' \to w_0$ we can find $d' \in D$ approaching b_0 and satisfying $(w', d') \in H$. Then each such d' belongs to D_{a_0}, which shows that $b_0 \in \overline{D_{a_0}}$. □

We conclude:

Lemma 8.3.22. *rk* $D_{a_0} = 2$.

Proof. If not, then rk $(D_{a_0}) \leq 1$. Recall that $D_{a_0} \subset D$ is $\mathcal{M}$-definable over (a_0, p_0, r_0), and D is isomorphic to M^2. So, by Lemma 8.3.21 and Lemma 8.3.1, we get that
$$\text{rk }(b_0/a_0, p_0, r_0) \leq 1.$$
This contradicts (3) of Lemma 8.3.20. □

Now as it turns out, it is actually easier to analyze R_H by first analyzing R. We do that now:

Lemma 8.3.23. *R has rank $k+4$, and the projection $R \to A \times B$ is almost surjective and almost finite-to-one.*

Proof. For each $(w,a) \in F$, the set of $b \in B$ for which $(w,a,b) \in R$ is just ${}_wG$, which always has rank 2. Thus

$$\text{rk } R = \text{rk } F + 2 = k + 4.$$

Now the key observation to make is that R stationary in $\mathcal{M}$. Indeed, let $(w,a,b) \in R$ be $\mathcal{M}$-generic. Then rk $(w,a,b) = k+4$. On the other hand we have $a \in A$, $w \in F_a$, and $b \in {}_wG$. Since the ranks of these three sets add to at most $k+2+2 = k+4$, we conclude that all of these membership relations are $\mathcal{M}$-generic: that is, a is $\mathcal{M}$-generic in A, w is $\mathcal{M}$-generic in F_a over a, and b is $\mathcal{M}$-generic in ${}_wG$ over (a,w). Moreover, note that each of A, F_a, and ${}_wG$ is stationary in $\mathcal{M}$: indeed, the stationarity of F_a was assumed in the statement of the proposition, since $\mathcal{F}$ is generically stationary; the stationary of A was then assumed in the beginning of the proof of the proposition; and the stationarity of ${}_wG$ was given in Lemma 8.3.13. Now by the stationarity of each of these sets, it follows that we have completely determined $\text{tp}_{\mathcal{M}}(w,a,b)$. In particular, R has only one generic type, so is stationary.

Now using Lemma 8.3.22 and Lemma 8.3.20, and the fact that D is stationary, we conclude that $b_0 \in D_{a_0}$; so there is some $w_1 \in F_{a_0} \cap G_{b_0}$. It follows that the projection $R \to A \times B$ has (a,b) in its image. But by Lemma 8.3.15, (a,b) is $\mathcal{M}$-generic in $A \times B$. So, since $A \times B$ is stationary, the projection $R \to A \times B$ is in fact almost surjective. Then, since R is stationary and rk $R =$ rk $(A \times B)$, we conclude that this projection is also almost finite-to-one. □

Finally, we conclude the analagous statement for R_H:

Lemma 8.3.24. *R_H has rank $k+2$, and the projection $R_H \to A \times D$ is almost surjective and almost finite-to-one.*

Proof. It suffices to show that (1) R_H contains elements of rank at least $k+2$ over (p_0, r_0), and (2) for any such element (w,a,d), we have rk $(w/a,d,p_0,r_0) = 0$.

For (1), recall as in the proof of Lemma 8.3.23 that there is some $w_1 \in F_{a_0} \cap G_{b_0}$. Also recall by Lemma 8.3.20 that rk $(a_0, b_0/p_0, r_0) = k+2$. It follows that $(w_1, a_0, b_0) \in R_H$ and

$$\text{rk } (w_1, a_0, b_0/p_0, r_0) \geq k+2.$$

For (2), let $(w,a,b) \in R_H$ be such that

$$\text{rk } (w,a,b/p_0,r_0) \geq k+2.$$

Since rk $(p_0, r_0) = 2$, we get by additivity that

$$\text{rk } (w,a,b,p_0,r_0) \geq k+4.$$

But since $b \in D$, it follows that p_0 and r_0 are among the coordinates of b; thus in fact we have

$$\text{rk}\ (w, a, b) \geq k + 4.$$

In particular, (w, a, b) is an $\mathcal{M}$-generic element of R. By Lemma 8.3.23, we conclude that rk $(w/a, b) = 0$, and so

$$\text{rk}\ (w/a, b, p_0, r_0) = 0,$$

as desired. □

By Lemma 8.3.24, and the fact that A and D are stationary, there is a positive integer l such that $|F_a \cap G_d| = l$ holds for all generic $(a, d) \in A \times D$ over (p_0, r_0).

We conclude:

Lemma 8.3.25. *$F_{a_0} \cap G_{b_0}$ has size exactly l.*

Proof. This is immediate from Lemma 8.3.20, since (a_0, b_0) is $\mathcal{M}$-generic in $A \times D$ over (p_0, r_0). □

Let $w_1, ..., w_l$ be the distinct elements of $F_{a_0} \cap G_{b_0}$. Our next main goal is to show that the projection $R_H \to A \times D$ is locally surjective near each (w_i, a_0, b_0):

Lemma 8.3.26. *Let w be any of $w_1, ..., w_l$. Then:*

1. *(w, b_0) is generic in H over (p_0, r_0).*
2. *$(w, a_0, b_0) \in (R_H)^P$.*
3. *The projection $R_H \to A \times D$ is locally surjective near (w, a_0, b_0).*

Proof. 1. By Lemma 8.3.16, we have $\dim(w, b_0) = 5n$. Since p_0 and r_0 are among the coordinates of b_0, it follows that

$$\dim(w, b_0, p_0, r_0) = 5n.$$

Then since $\dim(p_0, r_0) = 2n$, we conclude that

$$\dim(w, b_0/p_0, r_0) = 3n = \dim H.$$

The claim follows.

2. By Lemma 8.3.16, we have $(w, a_0) \in F^P$. Thus we can find arbitrarily close (w', a') to (w, a_0) which are generic in F over (p_0, r_0). Now by (1) and Lemma 8.3.19, the projection $H \to M^3$ is locally surjective near (w, b_0). Thus, as $w' \to w$, we can find d' arbitrarily close to b_0 such that $(w', d') \in H$. It follows that each such (w', a', d') is an element of R_H, whose dimension over (p_0, r_0) is at least

$$\dim(w', a'/p_0, r_0) = \dim F = \dim R_H.$$

That is, we can find (w', a', d') arbitrarily close to (w, a_0, b_0) which are generic in R_H over (p_0, r_0). Equivalently, $(w, a_0, b_0) \in (R_H)^P$.

3. We verify the hypotheses of Proposition 3.4.5. By Lemma 8.3.24, $R_H \to A \times D$ is almost surjective and almost finite-to-one; as a projection, $R_H \to A \times D$ is also continuous. By (2), $(w, a_0, b_0) \in (R_H)^P$. By Lemma 8.3.25, (w, a_0, b_0) belongs to a finite fiber, and is therefore isolated in its fiber.

   By assumption a_0 is generic in A over $\emptyset$, and by (1) it follows that b_0 is generic in D over (p_0, r_0); in particular, a_0 is smooth in A and b_0 is smooth in D, so that (a_0, b_0) is smooth in $A \times D$.

   It remains to show that (w, a_0, b_0) has a compact neighborhood in R_H. But this follows exactly as in the plane curve case: each of (w, a_0) and (w, b_0) is generic, and thus has a compact neighborhood, in the relevant set – either F or H; without loss of generality these sets have the form $(X \times Y) \cap F$ and $(X \times Z) \cap H$, for some compact $X \subset M^3$, $Y \subset A$, and $Z \subset D$. Then $(X \times Y \times Z) \cap R$ is a compact neighborhood of (w, a_0, b_0) in R.

□

We have finally set up everything we need to execute the geometric portion of the argument, in analogy to Claim 7.3.19. Let T be the set of all $d \in D$ such that $|F_{a_0} \cap G_d| > l$. So T is $\mathcal{M}$-definable over (a_0, p_0, r_0). Then we conclude:

Lemma 8.3.27. $b_0 \in Fr\ (T)$.

Proof. By Lemma 8.3.25, $b_0 \notin T$. So it remains to show that $b_0 \in \overline{T}$. Now the proof that $b_0 \in \overline{T}$ is essentially identical to the proof of Claim 7.3.19 in the case of plane curves, so we summarize:

Since $w_0 \in \text{Fr}\ (F_{a_0})$, we can find $w' \in F_{a_0}$ arbitrarily close to w_0. By Lemma 8.3.20, w_0 is generic in M^3 over (p_0, r_0); since the projection $H \to M^3$ is finite-to-one, it is therefore locally surjective near each point lying above w_0. In particular, $H \to M^3$ is locally surjective near (w_0, b_0). Thus, as $w' \to w_0$, we can find $d' \in D$ arbitrarily close to b_0 such that $(w', d') \in H$, and therefore $(w', a_0, d') \in R_H$. Then, by the local surjectivity of $R_H \to A \times D$ near each (w_i, a_0, b_0), for $i = 1, ..., l$, as $(w', d') \to (w_0, b_0)$ we can in turn find $w'_1, ..., w'_l$, arbitrarily close to $w_1, ..., w_l$, with each $(w'_i, a_0, d') \in R_H$. For all such points, it follows that $F_{a_0} \cap G_{d'}$ contains w_0, as well as each w'_i for $i = 1, ..., l$. As the $w'_i \to w_i$, the points $w_0, w'_1, ..., w'_l$ become distinct; thus all such d' which are close enough to b_0 belong to T. □

We are finally ready to complete the proof of Proposition 8.3.7. Since $b_0 \in \text{Fr}\ (T)$, T is necessarily infinite. Now if rk $T = 2$, then T is generic in the stationary set D, so that b_0 belongs to the non–generic subset $D - T \subset D$ which is $\mathcal{M}$-definable over (a_0, p_0, r_0). If on the other hand rk $T = 1$, then by the identification of D with M^2, and by Proposition 8.2.1, we get that b_0 is $\mathcal{M}$-algebraic over (a_0, p_0, r_0).

In either case, we conclude that b_0 is not $\mathcal{M}$-generic in D over (a_0, p_0, r_0). Thus we obtain a contradiction with Lemma 8.3.20, and so the proof of Proposition 8.3.7 is complete. □

8.4 Frontiers of Arbitrary 2-Hypersurfaces

This section is analagous to Sections 7.4 and 8.2 – namely, we show that studying the frontiers of arbitrary non-trivial 2-hypersurfaces reduces to the main result of Section 8.3. As there are especially few differences between the present section and our corresponding work on plane curves, we will be fairly terse.

We start with:

Proposition 8.4.1. *Let E be any set of parameters, and let e be a tuple which is optimal over E. Let $S \subset M^3$ be a non-trivial 2-hypersurface which is $\mathcal{M}$-definable over (E, e). Let $w_0 = (x_0, y_0, z_0)$ be a point which is generic in M^3 over E and belongs to $Fr\,(S)$. If $rk\,(x_0, y_0/E, e) \geq 1$, then $rk\,(z_0/E, e, x_0, y_0) = 0$.*

Proof. We may assume that $E = \emptyset$. Let T be a non-trivial, irreducible plane curve in $\mathcal{M}$ whose code t has rank $k > 3n$, is optimal, and is independent from (e, w_0); for example, one can take a large $\emptyset$-definable almost faithful family $\mathcal{G}$, then take T to be a curve in $\mathcal{G}$ whose index is generic over (e, w_0). Note that, after potentially editing finitely many points, we may assume that T is $\mathcal{M}$-definable over t.

Let

$$W = \{(x, y, v) \in M^3 : \text{ for some } z \in M,\ (x, y, z) \in S \wedge (z, v) \in T\}.$$

So W is $\mathcal{M}$-definable over (e, t). Since the projection $S \to M^2$ to the first two coordinates is almost surjective and almost finite-to-one, and T is non-trivial, it is easy to see that W is a 2-hypersurface. Furthermore, any generic point of W over (e, t) is $\mathcal{M}$-interalgebraic over (e, t), coordinate by coordinate, with a corresponding $\mathcal{M}$-generic point of S over (e, t); it easily follows that W is non-trivial.

We next show that w_0 can be traced to an element of $\overline{W}$. Since t is independent from (e, w_0), it follows that w_0 is generic in M^3 over t. In particular, since T is non-trivial, there is some $v_0 \in M$ such that (z_0, v_0) is a generic element of T over t. Then we show:

Claim 8.4.2. $(x_0, y_0, v_0) \in \overline{W}$.

Proof. Since (z_0, v_0) is generic in T over t, and T is non-trivial, each projection $T \to M$ is locally surjective near (z_0, v_0). Now since $w_0 \in \text{Fr}\,(S)$, we can find $(x', y', z') \in S$ arbitrarily close to w_0. Then by local surjectivity, all sufficiently close such z' can be completed to a point $(z', v') \in T$ which is arbitrarily close to (z_0, v_0). Then $(x', y', v') \in W$ is arbitrarily close to (x_0, y_0, v_0), which shows that $(x_0, y_0, v_0) \in \overline{W}$. □

Now by the claim, we can fix a stationary component U of W with $(x_0, y_0, v_0) \in \overline{U}$. Thus U is also a non-trivial 2-hypersurface. Since W is $\mathcal{M}$-definable over (e, t), we can find a tuple $b \in \text{acl}_{\mathcal{M}}(e, t)$ such that U is $\mathcal{M}$-definable over b. Additionally, let u be a code of U. Note that u and b need not be equal, but we do have $u \in \text{acl}_{\mathcal{M}}(b)$.

We proceed to generate some properties of the triple (e, t, u); our goal is to solve for the rank and dimesion of u. Now we have:

- e and t are optimal and independent, by the choice of t – thus (e,t) is also optimal.
- Since $u \in \text{acl}_{\mathcal{M}}(b)$ and $b \in \text{acl}_{\mathcal{M}}(e,t)$, we get $u \in \text{acl}_{\mathcal{M}}(e,t)$, and thus u is also optimal.
- rk $t = k > 3n$, so by independence rk $(t/e) = k > 3n$.

We further verify:

Claim 8.4.3. *t is $\mathcal{M}$-algebraic over (e,u).*

Proof. Let $x \in M$ be generic over (e,t,u), let

$$S_x = \{(y,z) : (x,y,z) \in S\},$$

and let

$$U_x = \{(y,v) : (x,y,v) \in U\}.$$

Since S and U are non-trivial, and by the genericity of x, it follows that S_x and U_x are non-trivial plane curves; moreover, by definition we have $U_x \subset T \circ S_x$. Now since u is a code of U, it follows that U is $\mathcal{M}$-definable up to almost equality over u. In particular, since x is generic over u, it follows that U_x is $\mathcal{M}$-definable up to almost equality over (u,x).

Then using that T is strongly minimal, it follows that t is $\mathcal{M}$-algebraic over (e,u,x): indeed, exactly as in the proof of Lemma 4.2.2, T is almost equal to one of the strongly minimal components of the plane curve $U_x \circ S_x^{-1}$, which is definable up to almost equality over (e,u,x). But x is independent from (e,t,u), so we in fact conclude that t is $\mathcal{M}$-algebraic over (e,u), as desired. ☐

By the previous claim, we conclude that rk $u \geq k > 3n$: indeed, this follows since rk $(t/e) = k$ but rk $(t/e,u) = 0$.

Then, since u is optimal of rank $\geq k$, we can find a $\emptyset$-definable, almost faithful, generically stationary, rank $\geq k$ family $\mathcal{F} = \{F_a\}_{a \in A}$ of non-trivial 2-hypersurfaces in $\mathcal{M}$, and a generic $a \in A$, such that a and u are $\mathcal{M}$-interalgebraic and $F_a \sim U$. Note that since u is $\mathcal{M}$-algebraic over b, it follows that a is also $\mathcal{M}$-algebraic over b.

Recall that we are assuming rk $(x_0, y_0/e) \geq 1$. Since t is independent from (e, w_0), we get rk $(x_0, y_0/e, t) \geq 1$. In particular, since b is $\mathcal{M}$-algebraic over (e,t), and by the above remarks, it follows that rk $(x_0, y_0/a, b, e, t) \geq 1$. Then we arrive at the following:

Claim 8.4.4. *rk $(v_0/a, b, e, t, x_0, y_0) = 0$.*

Proof. If $(x_0, y_0, v_0) \in F_a$, then this follows from the above remarks and the fact that F_a is non-trivial. If $(x_0, y_0, v_0) \in \text{Fr}\,(F_a)$, then Proposition 8.3.7 gives that rk $(v_0/a, x_0, y_0) = 0$, so we are again done.

Otherwise we have $(x_0, y_0, v_0) \notin \overline{F_a}$; so since $(x_0, y_0, v_0) \in \overline{U}$, we conclude that

$$(x_0, y_0, v_0) \in \overline{U - F_a}.$$

Now since F_a is almost equal to U, we have rk $(U - F_a) \leq 1$; thus we can apply Lemma 8.3.1. Since U is $\mathcal{M}$-definable over b, we conclude that rk $(x_0, y_0, v_0/a, b) \leq 1$. But since rk $(x_0, y_0/a, b, e, t) \geq 1$, it follows easily that rk $(v_0/a, b, x_0, y_0) = 0$, so that we yet again conclude the desired statement. □

So we have

$$\text{rk } (v_0/a, b, e, t, x_0, y_0) = 0.$$

But since $(z_0, v_0) \in T$ and T is non-trivial, z_0 and v_0 are $\mathcal{M}$-interalgebraic over t; thus we replace v_0 with z_0, and conclude that

$$\text{rk } (z_0/a, b, e, t, x_0, y_0) = 0.$$

Since b is $\mathcal{M}$-algebraic over (e, t), and a is $\mathcal{M}$-algebraic over b, we can drop both the a and the b, leaving

$$\text{rk } (z_0/e, t, x_0, y_0) = 0.$$

Finally, since t is independent from (e, w_0), we can drop the t, and conclude that

$$\text{rk } (z_0/e, x_0, y_0) = 0.$$

The proves Proposition 8.4.1. □

As a corollary, we finally deduce the main goal of the previous two sections. The proof is essentially identical to the proof of Proposition 8.2.1:

Proposition 8.4.5. *Let $S \subset M^3$ be any non-trivial 2-hypersurface, $\mathcal{M}$-definable over a set A. If $w_0 = (x_0, y_0, z_0) \in Fr\ (S)$, and rk $(x_0, y_0/A) \geq 1$, then rk $(z_0/A, x_0, y_0) = 0$.*

Proof. We may assume that $A = \emptyset$. Let $P = (p, q, r) \in M^3$ be generic over w_0. Then $w_0 + P$ is generic over $\emptyset$ and belongs to Fr $(S + P)$. Now each coordinate of $w_0 + P$ is $\mathcal{M}$-interalgebraic over P with the corresponding coordinate of w_0; by the independence of w_0 and P, it follows that

$$\text{rk } (x_0 + p, y_0 + q/P) \geq 1.$$

Moreover, $S + P$ is $\mathcal{M}$-definable over P, which is optimal since it is generic.

So Proposition 8.4.1 applies, and we conclude that

$$\text{rk } (z_0 + r/P, x_0 + p, y_0 + q) = 0.$$

But recalling that w_0 and $w_0 + P$ are coordinate-wise $\mathcal{M}$-interalgebraic over P, this implies that rk $(z_0/P, x_0, y_0) = 0$. Finally, by the independence of w_0 and P, we conclude that rk $(z_0/x_0, y_0) = 0$, as desired. □

8.5 Poles

The purpose of this section is to bound the 'poles' of non-trivial plane curves. As described in the introduction, our strategy is to reduce the study of the poles of a curve S to the study of the frontier points of a certain associated 2-hypersurface T. We will show that, if S has infinitely many poles, then T must contain a direct contradiction to Proposition 8.4.5; we thus conclude that S has only finitely many poles.

We note that our strategy is inspired by [12], though in practice it is structured quite differently. For example, the authors of [12] needed a quite intricate study of the interactions between different poles, in order to construct a plethora of frontier points of T; while our argument below only analyzes a single pole, and therefore avoids much of the complexity of the argument in [12]. Of course, the reason for this is that we have spent the entirety of the last two sections developing exactly the restriction we need on frontier points of T.

We will need to use the o-minimal structure on the real field in this section, though not extensively. The main notion we need is the Peterzil-Steinhorn subgroup associated to an unbounded curve, which we proceed to outline below:

First, since $(M, +)$ is an $\mathbb{R}$-definable group, it follows by [39] that M carries a unique $\mathbb{R}$-definable manifold structure making it into a topological group. In fact, by uniqueness, this is precisely the manifold structure induced on M by its analytic topology as a complex variety. Recall further that M is quasi-projective, meaning that we can embed it into a complex projective variety V. We fix such an embedding $M \to V$, and assume after adding constants that it is $\emptyset$-definable in $\mathcal{K}$; we will liberally abuse notation by identifying M with its image in V. So, to summarize, we will from this point view M as an open subvariety of V, whose analytic topology is given by an $\mathbb{R}$-definable group manifold structure. We now define:

Definition 8.5.1. An *unbounded curve in M* is an $\mathcal{R}$-definable function $\gamma : (0,1) \to M$ such that $\lim_{t\to 0^+} \gamma(t)$ does not exist in M.

Remark 8.5.2. Viewing γ as a map to the compact set V, it follows by o-minimality that γ must have a limit in V; so Definition 8.5.1 is equivalent to the assertion that

$$\lim_{t\to 0^+} \gamma(t) \in V - M.$$

In [36], Peterzil and Steinhorn showed that every unbounded curve in an o-minimally definable group gives rise to a torsion-free definable subgroup of dimension one in a uniform way. Restricting their result to our situation, we have the following definition and fact:

Definition 8.5.3. If γ is an unbounded curve in $\mathcal{M}$, we denote by $PS(\gamma)$ the set of limit points in M of expressions $\gamma(s) - \gamma(t)$ as $s, t \to 0$. That is, $z \in PS(\gamma)$ if and only if there are $\mathbb{R}$-definable functions $\eta_1, \eta_2 : (0,1) \to (0,1)$ such that $\lim_{s\to 0^+} \eta_i(s) = 0$ for each i, and

$$\lim_{s\to 0^+} \gamma \circ \eta_1(s) - \gamma \circ \eta_2(s) = z.$$

Fact 8.5.4 (Peterzil-Steinhorn)**.** *For each unbounded curve γ in M, the set $PS(\gamma)$ is an $\mathbb{R}$-definable torsion-free subgroup of M of $\mathbb{R}$-dimension 1.*

We thus call $PS(\gamma)$ the *Peterzil-Steinhorn subgroup* associated to γ. We will not need that $PS(\gamma)$ is a group – indeed, the key point is that it is infinite.

We now proceed to discuss poles. Recall that we identify M as an open subvariety of a fixed projective variety V.

Definition 8.5.5. Let $S \subset M^2$ be a non-trivial plane curve in $\mathcal{M}$, and let $x \in M$. Then x is a *pole of* S if there is some $y \in V - M$ such that (x, y) belongs to the closure of S in $V \times V$.

For example, if M is projective then no plane curve has a pole. Note that by the definition, the set of poles of S is $\mathcal{K}$-definable over the parameters that define S.

In general, we can relate poles to unbounded curves, as follows:

Lemma 8.5.6. *Let $S \subset M^2$ be a non-trivial plane curve in $\mathcal{M}$, and let x be a pole of S. Then there is an $\mathbb{R}$-definable function $\gamma = (\gamma_1, \gamma_2) : (0, 1) \to S$ such that $\lim_{t\to 0^+} \gamma_1(t) = x$, and γ_2 is an unbounded curve in M.*

Proof. By curve selection ([7], Chapter 6, Corollary 1.5) applied to the set $V \times V$, we can construct such a pair so that $\lim_{t\to 0^+} \gamma_2(t) = y$ for some $y \in V - M$. As in Remark 8.5.2, this is equivalent to saying that γ_2 is unbounded. □

Now the main goal of this section is the following:

Proposition 8.5.7. *Let S be a non-trivial plane curve. Then the set of poles of S is finite.*

Proof. We may assume that S is $\emptyset$-definable in $\mathcal{M}$. Now the point is to work with the following set:

Definition 8.5.8. Let

$$T = \{(x, u, z) \in M^3 : \text{ for some } y, v \in M \text{ we have } (x, y) \in S,\ (u, v) \in S, \text{ and } y = v + z\}.$$

So T tracks differences between y-values in S, and records them by pairs of x-values. Note that T is also $\mathcal{M}$-definable over $\emptyset$. Using that S is a non-trivial plane curve, the following is easy to check:

Lemma 8.5.9. *T is a non-trivial 2-hypersurface. In fact, each projection $T \to M^2$ is finite-to-one.*

Proof. First we show that rk $T \geq 2$. Indeed, let $(x, u) \in M^2$ be generic. Then since S is non-trivial, there are $y, v \in M$ with $(x, y), (u, v) \in S$. Set $z = y - v$; then $(x, u, z) \in T$, and

$$\text{rk } (x, u, z) \geq \text{rk } (x, u) = 2.$$

Now we claim that each of the projections $T \to M^2$ is finite-to-one. In order to show this, it suffices to note the following:

Claim 8.5.10. *If $(x, u, z) \in T$, then each of x, u, and z is $\mathcal{M}$-algebraic over the other two.*

Proof. Since $(x, u, z) \in T$, there are $y, v \in M$ such that $(x, y), (u, v) \in S$, and $y = v + z$. Now since $y = v + z$, clearly each of y, v, and z is $\mathcal{M}$-algebraic over the other two. On the other hand, since S is non-trivial we get that x and y are $\mathcal{M}$-interalgebraic over $\emptyset$, as are u and v. So we can replace y and v with x and u, and conclude that the same property holds of the triple (x, u, z). □

In particular, each projection $T \to M^2$ is finite-to-one. We conclude that rk $T \leq 2$, and so rk $T = 2$. Thus T is indeed a 2-hypersurface; additionally, it now follows by the claim that each projection $T \to M^2$ is in addition almost surjective, so that T is non-trivial. □

Now the main point is the following:

Lemma 8.5.11. *Let x be any pole of S. Then there are infinitey many $z \in M$ such that $(x, x, z) \in \overline{T}$.*

Proof. By Lemma 8.5.6, we can find an $\mathbb{R}$-definable function $\gamma = (\gamma_1, \gamma_2) : (0, 1) \to S$ such that $\lim_{t \to 0^+} \gamma_1(t) = x$ and γ_2 is unbounded. Let $H = PS(\gamma_2)$ be the Peterzil-Steinhorn subgroup associated to γ_2. Then it suffices to show:

Claim 8.5.12. *For all $z \in H$ we have $(x, x, z) \in \overline{T}$.*

Proof. If $z \in H$, then we can find $\mathbb{R}$-definable functions $\eta_1, \eta_2 : (0, 1) \to (0, 1)$ such that $\lim_{s \to 0^+} \eta_i(s) = 0$ for each i, and

$$\lim_{s \to 0^+} \gamma_2 \circ \eta_2(s) - \gamma_2 \circ \eta_1(s) = z.$$

For each s let

$$\alpha(s) = (\gamma_1 \circ \eta_1(s), \gamma_1 \circ \eta_2(s), \gamma_2 \circ \eta_1(s) - \gamma_2 \circ \eta_2(s)) \in M^3.$$

It follows immediately that each $\alpha(s) \in T$, and $\lim_{s \to 0^+} \alpha(s) = (x, x, z)$, which proves the claim. □

To conclude, the statement of the lemma now follows from the claim and the fact that $H = PS(\gamma_2)$ is infinite. □

Finally, we are ready to prove Proposition 8.5.7. Indeed, suppose toward a contradiction that S has infinitely many poles. Since the set of poles is $\mathcal{K}$-definable over $\emptyset$, there is a pole x_0 satisfying $\dim x_0 \geq 1$ – and thus in particular rk $x_0 = 1$.

Now by the previous lemma there are infinitely many $z \in M$ such that $(x_0, x_0, z) \in \overline{T}$. Of course, the set of such z is $\mathcal{K}$-definable over x_0; it follows that there is some $z_0 \in M$ such that $(x_0, x_0, z_0) \in \overline{T}$ and $\dim(z_0/x_0) \geq 1$ – so in particular rk $(z_0/x_0) = 1$. On the other hand, since each projection $T \to M^2$ is finite-to-one, there are only finitely many z with $(x_0, x_0, z) \in T$. It follows that $(x_0, x_0, z_0) \in$ Fr (T). But now (x_0, x_0, z_0) directly contradicts Proposition 8.4.5, and we are done. □

8.6 The Main Theorem

In light of the previous section, we have everything put in place to prove local modularity; the main idea is that we can now run the proof of Lemma 7.5.2. We do this now:

Lemma 8.6.1. *Every plane curve in $\mathcal{M}$ is almost pure.*

Proof. Let S be a plane curve. As in Lemma 7.5.2, we may assume that S is irreducible, non-trivial, and $\emptyset$-definable. Let $(x_0, y_0) \in S$ with $\dim(x_0, y_0) \geq n-1$; we will show that $(x_0, y_0) \in S^P$. First note that since S is non-trivial, x_0 and y_0 are $\mathcal{M}$-interalgebraic, and thus each has dimension at least $n-1$ over $\emptyset$.

Now by strong minimality, there is a positive integer l such that the projection $S \to M$ to the first coordinate is l-to-one almost everywhere. Then, as in Lemma 7.5.2, there is a dense open set $U \subset M$ such that each $u \in U$ has exactly l preimages in S, each of which belongs to S^P.

Since U is dense we have $x_0 \in \overline{U}$; so by curve selection, there is an $\mathbb{R}$-definable function $\gamma : (0,1) \to U$ with $\lim_{t \to 0^+} \gamma(t) = x$. Then, as in Lemma 7.5.2, we obtain $\mathbb{R}$-definable functions $\eta, ..., \eta_l : (0,1) \to M$, such that for each t, the elements $\eta_i(t) \in M$, for $i = 1, ..., l$, are precisely those $y \in M$ with $(\gamma(t), y) \in S$.

Since V is compact, it follows by o-minimality that each η_i converges to an element $y_i \in V$; in particular, by the coice of U, each (x_0, y_i) belongs to the closure of $S \cap S^P$ in $V \times V$. Then we note:

- Each $y_i \in M$. Indeed, otherwise it would follow that x_0 is a pole of S; then by Proposition 8.5.7, we conclude that $\dim x_0 = 0$, contradicting that $\dim x_0 \geq n-1$. This is the only essential difference between Lemma 8.6.1 and Lemma 7.5.2.

- Each $(x_0, y_i) \in S$. Indeed, suppose not; then by the previous item, it follows that $(x_0, y_i) \in \text{Fr}\ (S) \subset M^2$; then by Proposition 8.2.1, we similarly conclude that $\dim(x_0, y_i) = 0$, contradicting that $\dim(x_0, y_0) \geq n-1$.

- Each $(x_0, y_i) \in S^P$. Indeed, by construction (x_0, y_i) belongs to the closure of S^P in $V \times V$. Since $y_i \in M^2$ and M is open in V, (x_0, y_i) in fact belongs to the closure of S^P in $M \times M$. But S^P is closed in $M \times M$, so $(x_0, y_i) \in S^P$.

- The y_i are distinct. Indeed, if $y_i = y_j = y$ for some $i \neq j$, then exactly as in Lemma 7.5.2, the point (y, y) would belong to the frontier of the plane curve $T = S \circ S^{-1} - \Delta$; thus Proposition 8.2.1 would imply that $\dim y = 0$. But by the second item above we have $(x_0, y) \in S$, so that x_0 and y are interalgebraic; so we conclude that $\dim x_0 = 0$, again contradicting that $\dim x_0 \geq n-1$.

Now since $\dim x_0 \geq n-1$, it follows that x_0 is $\mathcal{M}$-generic in M, so that the fiber $S_{x_0} \subset M$ has size exactly l. Then by the above remarks, the elements of S_{x_0} are precisely $y_1, ..., y_l$. We are thus forced to conclude that $y_0 = y_i$ for some $i = 1, ..., l$, and so $(x_0, y_0) \in S^P$. $\square$

Finally, we conclude:

Theorem 8.6.2. *$\mathcal{M}$ is locally modular.*

Proof. If not, then there is a rank 2, almost faithful family $\mathcal{F}$ of plane curves in $\mathcal{M}$. By Lemma 8.6.1, every curve in $\mathcal{F}$ is almost pure; we thus contradict Theorem 5.0.1. □

We now recap the main result of this chapter:

Theorem 8.6.3. *Let $(M, \cdot)$ be a group interpreted in an algebraically closed field K of characteristic zero. Assume that the underlying set M has dimension at least 2 as a K-interpretable set. Assume further that $\mathcal{M} = (M, \cdot, ...)$ is a strongly minimal reduct of the full K-induced structure on M. Then $\mathcal{M}$ is locally modular.*

Proof. If not, then a counterexample would exist in any saturated algebraically closed field of characteritic zero. In particular, we may assume that $K = \mathbb{C}$. In this case the theorem follows immediately from Proposition 8.1.1 and Theorem 8.6.2. □

We also note that we can deduce the full Restricted Trichotomy Conjecture for groups in characteristic zero:

Corollary 8.6.4. *Let $\mathcal{M} = (M, \cdot, ...)$ be a strongly minimal expansion of a group, which is interpreted in an algebraically closed field K of characteristic zero. Then either $\mathcal{M}$ is locally modular or $\mathcal{M}$ interprets a field isomorphic to K.*

Proof. If M has dimension at least 2 as a $\mathcal{K}$-definable set, then by Theorem 8.6.3 $\mathcal{M}$ is locally modular. Otherwise M has dimension 1 as a $\mathcal{K}$-definable set, and the desired statement follows from the main theorem of [20]. □

We end by noting one more application of Theorem 8.6.3: namely, that to prove local modularity in general, it suffices to find a single rank one generically almost pure family of non-trivial plane curves.

Corollary 8.6.5. *Let M be a smooth variety of dimension at least 2 over an algebraically closed field K of characteristic zero, and let $\mathcal{M} = (M, ...)$ be a strongly minimal reduct of the full K-induced structure on M. Assume that there is an almost faithful family $\mathcal{F} = \{F_a\}_{a \in A}$ of non-trivial plane curves in $\mathcal{M}$, such that rk $A \geq 1$, and for all generic $a \in A$ the curve F_a is almost pure. Then $\mathcal{M}$ is locally modular.*

Proof. We may assume that $K = \mathbb{C}$. In this case, Theorem 5.5.8 implies that $\mathcal{M}$ interprets a strongly minimal group $(G, \cdot)$. Let $\mathcal{G}$ denote the full structure induced on G from $\mathcal{M}$.

Recall by Corollary 3.1.16 that $\mathcal{M}$ cannot interpret an infinite field; it follows that $\mathcal{G}$ cannot interpret an infinite field either. So by Corollary 8.6.4, $\mathcal{G}$ is locally modular.

We conclude that $\mathcal{M}$ is a strongly minimal structure which interprets a locally modular strongly minimal set. It follows (see for example [54], Chapter II, Theorem 3.5) that all strongly minimal sets interpreted in $\mathcal{M}$ are locally modular. In particular, $\mathcal{M}$ itself is locally modular. □

Bibliography

[1] René-Louis Baire. "Sur les fonctions de variables réelles". In: *Ann. di Mat* 3 (1899), pp. 1–123.

[2] J. T. Baldwin and A. H. Lachlan. "On strongly minimal sets". In: *J. Symbolic Logic* 36 (1971), pp. 79–96.

[3] Elisabeth Bouscaren. "The group configuration—after E. Hrushovski". In: *The model theory of groups (Notre Dame, IN, 1985–1987)*. Vol. 11. Notre Dame Math. Lectures. Univ. Notre Dame Press, Notre Dame, IN, 1989, pp. 199–209.

[4] C. C. Chang and H. J. Keisler. *Model theory*. Third. Vol. 73. Studies in Logic and the Foundations of Mathematics. North-Holland Publishing Co., Amsterdam, 1990.

[5] J. Denef and L. van den Dries. "p-adic and real subanalytic sets". In: *Ann. of Math. (2)* 128.1 (1988), pp. 79–138.

[6] Lou van den Dries. "A generalization of the Tarski-Seidenberg theorem, and some nondefinability results". In: *Bull. Amer. Math. Soc. (N.S.)* 15.2 (1986), pp. 189–193.

[7] Lou van den Dries. *Tame topology and o-minimal structures*. Vol. 248. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge, 1998.

[8] Lou van den Dries. "Weil's group chunk theorem: a topological setting". In: *Illinois J. Math.* 34.1 (1990), pp. 127–139.

[9] Lou van den Dries, Angus Macintyre, and David Marker. "The elementary theory of restricted analytic fields with exponentiation". In: *Ann. of Math. (2)* 140.1 (1994), pp. 183–205.

[10] Mario Edmundo and Margarita Otero. "Definably compact abelian groups". In: *J. Math. Log.* 4.2 (2004), pp. 163–180.

[11] David Eisenbud and Joe Harris. *3264 and all that—a second course in algebraic geometry*. Cambridge University Press, Cambridge, 2016.

[12] Pantelis Eleftheriou, Assaf Hasson, and Ya'acov Peterzil. "Strongly Minimal Groups in O-minimal Structures". In: *J. Eur. Math. Soc.* (2021 (to appear)).

[13] Darío García and Frank Wagner. "Unimodularity unified". In: *J. Symb. Log.* 82.3 (2017), pp. 1051–1065.

[14] A. Grothendieck. "Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas IV". In: *Inst. Hautes Études Sci. Publ. Math.* 32 (1967), p. 361.

[15] Alexander Grothendieck and Michèle Raynaud. *Revêtements étales et groupe fondamental (SGA I)*. Lecture Notes in Mathematics 224, Springer, 1971.

[16] Victor Guillemin and Alan Pollack. *Differential topology*. Reprint of the 1974 original. AMS Chelsea Publishing, Providence, RI, 2010.

[17] Robin Hartshorne. *Algebraic geometry*. Graduate Texts in Mathematics, No. 52. Springer-Verlag, New York-Heidelberg, 1977.

[18] Assaf Hasson and Piotr Kowalski. "Strongly minimal expansions of $(\mathbb{C}, +)$ definable in o-minimal fields". In: *Proc. Lond. Math. Soc. (3)* 97.1 (2008), pp. 117–154.

[19] Assaf Hasson, Alf Onshuus, and Ya'acov Peterzil. "Definable one dimensional structures in o-minimal theories". In: *Israel J. Math.* 179 (2010), pp. 297–361.

[20] Assaf Hasson and Dmitry Sustretov. "Incidence Systems on Cartesian Powers of Algebraic Curves". manuscript.

[21] Ehud Hrushovski. "A new strongly minimal set". In: *Ann. Pure Appl. Logic* 62.2 (1993). Stability in model theory, III (Trento, 1991), pp. 147–166.

[22] Ehud Hrushovski. "Contributions to Stable Model Theory". PhD thesis. University of California, Berkeley, 1986.

[23] Ehud Hrushovski. "Locally modular regular types". In: *Classification theory (Chicago, IL, 1985)*. Vol. 1292. Lecture Notes in Math. Springer, Berlin, 1987, pp. 132–164.

[24] Ehud Hrushovski. "Strongly minimal expansions of algebraically closed fields". In: *Israel J. Math.* 79.2-3 (1992), pp. 129–151.

[25] Ehud Hrushovski. "Unimodular minimal structures". In: *J. London Math. Soc. (2)* 46.3 (1992), pp. 385–396.

[26] Ehud Hrushovski and Željko Sokolovic. "Minimal Subsets of Differentially Closed Fields". preprint. 1994.

[27] Ehud Hrushovski and Boris Zilber. "Zariski geometries". In: *J. Amer. Math. Soc.* 9.1 (1996), pp. 1–56.

[28] William Johnson. "Fun With Fields". PhD thesis. University of California, Berkeley, 2016.

[29] Julia F. Knight, Anand Pillay, and Charles Steinhorn. "Definable sets in ordered structures. II". In: *Trans. Amer. Math. Soc.* 295.2 (1986), pp. 593–605.

[30] David Marker. *Model theory*. Vol. 217. Graduate Texts in Mathematics. An introduction. Springer-Verlag, New York, 2002.

[31] Gary A. Martin. "Definability in reducts of algebraically closed fields". In: *J. Symbolic Logic* 53.1 (1988), pp. 188–199.

[32] J. S. Milne. *Algebraic groups*. Vol. 170. Cambridge Studies in Advanced Mathematics. The theory of group schemes of finite type over a field. Cambridge University Press, Cambridge, 2017.

[33] David Mumford. *The red book of varieties and schemes*. Vol. 1358. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1988.

[34] Ya'acov Peterzil and Sergei Starchenko. "A trichotomy theorem for o-minimal structures". In: *Proc. London Math. Soc. (3)* 77.3 (1998), pp. 481–523.

[35] Ya'acov Peterzil and Sergei Starchenko. "Expansions of algebraically closed fields in o-minimal structures". In: *Selecta Math. (N.S.)* 7.3 (2001), pp. 409–445.

[36] Ya'acov Peterzil and Charles Steinhorn. "Definable compactness and definable subgroups of o-minimal groups". In: *J. London Math. Soc. (2)* 59.3 (1999), pp. 769–786.

[37] Rafał Pierzchała. "Invariance of domain in o-minimal structures". In: *Ann. Polon. Math.* 77.3 (2001), pp. 289–294.

[38] Anand Pillay. *Geometric stability theory*. Vol. 32. Oxford Logic Guides. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1996.

[39] Anand Pillay. "On groups and fields definable in *o*-minimal structures". In: *J. Pure Appl. Algebra* 53.3 (1988), pp. 239–255.

[40] Anand Pillay and Charles Steinhorn. "Definable sets in ordered structures. I". In: *Trans. Amer. Math. Soc.* 295.2 (1986), pp. 565–592.

[41] Anand Pillay and Martin Ziegler. "Jet spaces of varieties over differential and difference fields". In: *Selecta Math. (N.S.)* 9.4 (2003), pp. 579–599.

[42] Bruno Poizat. *Stable groups*. Vol. 87. Mathematical Surveys and Monographs. Translated from the 1987 French original by Moses Gabriel Klein. American Mathematical Society, Providence, RI, 2001.

[43] Bruno Poizat. "Une théorie de Galois imaginaire". In: *J. Symbolic Logic* 48.4 (1983), 1151–1170 (1984).

[44] E. D. Rabinovich. *Definability of a field in sufficiently rich incidence systems*. Vol. 14. QMW Maths Notes. With an introduction by Wilfrid Hodges. Queen Mary and Westfield College, School of Mathematical Sciences, London, 1993.

[45] Joachim Reineke. "Minimale Gruppen". In: *Z. Math. Logik Grundlagen Math.* 21.4 (1975), pp. 357–359.

[46] Ernst Steinitz. "Algebraische Theorie der Körper". In: *J. Reine Angew. Math.* 137 (1910), pp. 167–309.

[47] Alfred Tarski. *A decision method for elementary algebra and geometry*. 2nd ed. University of California Press, Berkeley and Los Angeles, Calif., 1951.

[48] Ravi Vakil. *The Rising Sea: Foundations of Algebraic Geometry*. online notes. 2017.

[49] A. J. Wilkie. “Model completeness results for expansions of the ordered field of real numbers by restricted Pfaffian functions and the exponential function”. In: *J. Amer. Math. Soc.* 9.4 (1996), pp. 1051–1094.

[50] Oscar Zariski. “On the purity of the branch locus of algebraic functions”. In: *Proc. Nat. Acad. Sci. U.S.A.* 44 (1958), pp. 791–796.

[51] Boris Zilber. “A curve and its abstract Jacobian”. In: *Int. Math. Res. Not. IMRN* 5 (2014), pp. 1425–1439.

[52] Boris Zilber. “Strongly Minimal Countably Categorical Theories. III”. In: *Siberian Mathematics Journal* 25 (1984), pp. 559–571.

[53] Boris Zilber. “The structure of models of uncountably categorical theories”. In: *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Warsaw, 1983).* PWN, Warsaw, 1984, pp. 359–368.

[54] Boris Zilber. *Uncountably categorical theories*. Vol. 117. Translations of Mathematical Monographs. Translated from the Russian by D. Louvish. American Mathematical Society, Providence, RI, 1993.