

MULTIPLICATIVE RELATIONS AMONG
SPECIAL POINTS OF MODULAR
AND SHIMURA CURVES

GUY FOWLER

St John's College
University of Oxford

A thesis submitted for the degree of
DOCTOR OF PHILOSOPHY
Hilary Term 2021

For my family

Contents

Chapter 0. Conventions	11
Chapter 1. Introduction	13
1.1. Elliptic curves	13
1.2. The j -invariant	17
1.3. Modular functions	19
1.4. Modular curves	22
1.5. Imaginary quadratic fields	24
1.6. Complex multiplication of elliptic curves	25
1.7. Singular moduli	28
1.8. Shimura curves	30
1.9. Plan of this thesis	33
Chapter 2. Images of special points in \mathbb{G}_m	37
2.1. The mixed Shimura variety setting	37
2.2. Exemplary components	41
2.3. Ax–Schanuel	43
2.4. A finiteness result for strongly special subvarieties	46
2.5. Arithmetic ingredients	49
2.6. Finiteness of exemplary components	52
Chapter 3. Multiplicative independence of modular functions	61
3.1. Special points of modular functions	61

3.2.	Functional multiplicative independence	65
3.3.	Borcherds products for modular functions	68
3.4.	Finiteness of multiplicatively dependent tuples	74
3.5.	Describing the exemplary components	95
3.6.	An extension to finite rank	98
Chapter 4.	Atypical intersections and multiplicative dependence	107
4.1.	The Zilber–Pink conjecture	107
4.2.	Multiplicatively dependent images of special points	110
4.3.	Multiplicatively dependent f -special points	116
4.4.	The finite rank case	130
Chapter 5.	Effective results on products of singular moduli	145
5.1.	The André–Oort conjecture	146
5.2.	The multiplicative case	148
5.3.	Background on singular moduli	151
5.4.	An effective bound	158
5.5.	Eliminating non-trivial cases	170
Appendix A.	O-minimality	177
A.1.	O-minimal structures	177
A.2.	The tameness of o-minimal structures	179
A.3.	The Pila–Wilkie Counting Theorem	183
A.4.	The o-minimal structure $\mathbb{R}_{\text{an},\text{exp}}$	187
References		191

Abstract

We consider a number of Diophantine problems for (mixed) Shimura varieties. Specifically, we look at the multiplicative relations satisfied by special points of modular and Shimura curves. These problems are closely related to the André–Oort and Zilber–Pink conjectures, and we will resolve some special cases of these conjectures.

Let Y be a modular or Shimura curve. Let $V \subset Y \times \mathbb{G}_m$ be an algebraic correspondence defined over $\overline{\mathbb{Q}}$. For each $n \geq 1$, we prove that all multiplicative dependencies among n V -images of special points of Y belong to one of finitely many components of V^n of a particular special kind.

We then strengthen this result in the case that Y is the modular curve $Y(1) = \mathrm{SL}_2(\mathbb{Z}) \backslash \mathbb{H}$ and V is the graph of a (suitably generic) rational function. In this case, we prove that, for each $n \geq 1$, there are only finitely many n -tuples of distinct V -images of special points which are multiplicatively dependent and minimal for this property. The key ingredient in this strengthening is a functional multiplicative independence statement for modular functions, which we prove.

Finally, we give a complete classification of the triples of singular moduli which have rational product. We show that all such triples are “trivial” in a suitable sense. This establishes a completely explicit André–Oort statement for the family of cubic surfaces defined (in \mathbb{C}^3) by an equation $x_1x_2x_3 = \alpha \in \mathbb{Q}$.

Acknowledgements

First, I would like to thank my supervisor, Jonathan Pila. His guidance and encouragement have been invaluable over the past four years, and I am tremendously grateful to have had the benefit of learning from him.

I also owe much to many other mathematicians, who have all been generous with their time and advice. In particular, I would like to thank Vahagn Aslanyan, Yuri Bilu, Laura Capuano, Chris Daw, Gabriel Dill, Ehud Hrushovski, Gareth Jones, Minhyong Kim, Jonathan Kirby, Alan Lauder, Alex Wilkie, and Boris Zilber. I am also very grateful to the referees who reviewed the two papers which I based upon parts of this thesis; my work has benefited greatly from their comments and corrections.

I thank my fellow Logic group students John Armitage, Oliviero Casani, Alexis Chevalier, Philip Dittmann, Sebastian Eterović, Arturo Rodriguez Fanlo, Haden Spence, and Brian Tyrrell, who all provided interesting company and patiently answered many of my questions.

Cornelia Drutu and Zhongmin Qian were excellent tutors and, later, colleagues. Ellen Luckins and Rob Rockwood, my fellow Exeter mathematicians, have provided friendship and support for the past eight years. From the distant past, I would also like to thank my teachers, Ray Healy and Craig Holloway, for first introducing me to mathematics.

Away from mathematics, I thank my friends, both in Oxford and elsewhere, who have made these years so enjoyable. They may not have contributed to the contents of this thesis, but their contribution to my life has been immeasurable.

Finally, I must also thank my family. It is their support at every stage of my life and education which made this possible. I could not have done any of it without them.

I gratefully acknowledge the financial support of the Engineering and Physical Sciences Research Council, the Mathematical Institute, and St John's College.

CHAPTER 0

Conventions

We collect here some of the notation and conventions which we will adopt throughout this thesis.

Set inclusions written $A \subset B$ are not necessarily proper. We write $|A|$ for the cardinality of a set A . Given a map $f: A \rightarrow B$, we will also denote by f arbitrary Cartesian products $f^n: A^n \rightarrow B^n$. Which map is meant should always be clear from context.

All the fields considered in this thesis will be subfields of the field \mathbb{C} of complex numbers. In particular, a field is always of characteristic 0. We let \mathbb{Q} denote the field of rational numbers and $\overline{\mathbb{Q}}$ denote the field of algebraic numbers.

We identify varieties with their sets of complex points. They are not necessarily irreducible. We denote by \mathbb{G}_m the multiplicative group of complex numbers $\mathbb{C} \setminus \{0\}$. Given a field K , we will sometimes write $\mathbb{G}_m(K)$ for the multiplicative group $K \setminus \{0\}$.

We let \mathbb{H} denote the complex upper half plane. The group $\mathrm{GL}_2^+(\mathbb{R})$ of 2×2 matrices with real entries and positive determinant will always act on \mathbb{H} by Möbius transformations. Any subgroup of $\mathrm{GL}_2^+(\mathbb{R})$ will be understood to act on \mathbb{H} in the same way.

Given $\alpha \in \overline{\mathbb{Q}}^\times$, we write $h(\alpha)$ for the absolute logarithmic Weil height of α and $H(\alpha) = \exp h(\alpha)$ for the multiplicative height. We refer to $H(\alpha)$ as the height of α and $h(\alpha)$ as the logarithmic height of α . Note that for

$\alpha \in \mathbb{Q}^\times$ such that $\alpha = a/b$ for $a, b \in \mathbb{Z}$ with $\gcd(a, b) = 1$ we have that $H(\alpha) = \max\{|a|, |b|\}$. For $\alpha = (\alpha_1, \dots, \alpha_n) \in (\overline{\mathbb{Q}}^\times)^n$, we let

$$H(\alpha) = \max\{H(\alpha_i) : i = 1, \dots, n\}$$

and

$$h(\alpha) = \max\{h(\alpha_i) : i = 1, \dots, n\}.$$

We again call these respectively the height of α and the logarithmic height of α . For basic properties of the height, see e.g. [14].

We use some results from model theory, in particular o-minimality. The necessary material is contained in Appendix A. By “definable”, we will always mean definable (with parameters) in the structure $\mathbb{R}_{\text{an,exp}}$. This structure is defined in Section A.4. In particular, note that $\mathbb{R}_{\text{an,exp}}$ is an o-minimal structure. We use real and imaginary parts to identify subsets of \mathbb{C}^n with the corresponding subsets of \mathbb{R}^{2n} in the natural way. Given $z \in \mathbb{C}$, we write respectively $\operatorname{Re} z$ and $\operatorname{Im} z$ for the real and imaginary parts of z .

CHAPTER 1

Introduction

In this introduction, we recall some of the classical material which forms the setting for our work. We then introduce the main question of this thesis and outline the plan for the rest of the document.

1.1. Elliptic curves

Definition 1.1. Let K be a field. An elliptic curve over K is a smooth projective curve over K which is of genus 1 and has a distinguished K -rational point.

For an elliptic curve E/K , the K -rational points of E form an abelian group $E(K)$. We will write the corresponding group law additively. In particular, the elliptic curves over K are precisely the one-dimensional abelian varieties over K .

Every equation

$$y^2 = x^3 + Ax + B$$

with $A, B \in K$ such that $4A^3 + 27B^2 \neq 0$ defines an elliptic curve over K . Conversely, every elliptic curve over K is isomorphic to one defined by such an equation (here we use the fact that the characteristic of K is 0). An equation $y^2 = x^3 + Ax + B$ is called a Weierstrass equation.

Elliptic curves over \mathbb{C} may be thought of in a particularly nice way. An elliptic curve over \mathbb{C} is a one-dimensional abelian variety over \mathbb{C} . There

is an equivalence of categories between the category of abelian varieties over \mathbb{C} and the category of polarisable complex tori [51, Theorem 2.9]. In particular, there is a one-to-one correspondence between the elliptic curves over \mathbb{C} and the one-dimensional complex tori. A one-dimensional complex torus is a quotient \mathbb{C}/Λ , where $\Lambda \subset \mathbb{C}$ is a lattice. A lattice in \mathbb{C} is an additive subgroup $\langle \tau_1, \tau_2 \rangle \subset \mathbb{C}$, where $\tau_1, \tau_2 \in \mathbb{C}$ are linearly independent over \mathbb{R} .

Definition 1.2. Let $\Lambda \subset \mathbb{C}$ be a lattice. An elliptic function (for Λ) is a meromorphic function $f: \mathbb{C} \rightarrow \mathbb{C}$ such that $f(z + \omega) = f(z)$ for every $z \in \mathbb{C}$ and $\omega \in \Lambda$.

Let $\Lambda \subset \mathbb{C}$ be a lattice. The prototypical example of an elliptic function for Λ is the corresponding Weierstrass \wp -function

$$\wp_{\Lambda}(z) = \frac{1}{z^2} + \sum_{\substack{\omega \in \Lambda \\ \omega \neq 0}} \left(\frac{1}{(z - \omega)^2} - \frac{1}{\omega^2} \right).$$

It is straightforward to see that \wp_{Λ} is an elliptic function for Λ , which is holomorphic away from its double poles at the points of Λ . Differentiating, we obtain that

$$\wp'_{\Lambda}(z) = \sum_{\omega \in \Lambda} \frac{-2}{(z - \omega)^3},$$

which is again an elliptic function for Λ , this time with triple poles at the points of Λ and holomorphic elsewhere.

Comparing terms, one may verify that \wp_{Λ} satisfies the equation

$$(\wp'_{\Lambda})^2 = 4(\wp_{\Lambda}(z))^3 - g_2(\Lambda)\wp_{\Lambda}(z) - g_3(\Lambda),$$

where

$$g_2(\Lambda) = 60 \sum_{\substack{\omega \in \Lambda \\ \omega \neq 0}} \frac{1}{\omega^4}$$

and

$$g_3(\Lambda) = 140 \sum_{\substack{\omega \in \Lambda \\ \omega \neq 0}} \frac{1}{\omega^6}.$$

This differential equation makes explicit the link between lattices and elliptic curves, as we will now explain.

Consider the map $\mathbb{C} \rightarrow \mathbb{C}^2$ given by $z \mapsto (\wp_\Lambda(z), \wp'_\Lambda(z))$. Write $(x, y) = (\wp_\Lambda(z), \wp'_\Lambda(z))$. Then the differential equation satisfied by \wp_Λ implies that

$$y^2 = 4x^3 - g_2(\Lambda)x - g_3(\Lambda).$$

This equation defines a curve, isomorphic to one given by an equation $y^2 = x^3 + Ax + B$ for some $A, B \in \mathbb{C}$. To show that the equation defines an elliptic curve, it is thus enough to show that the curve is non-singular. If the curve were singular, then one would have that $g_2(\Lambda)^3 - 27g_3(\Lambda)^2 = 0$. For every lattice Λ , it may be verified that this is not the case [79, Proposition VI.3.6]. Every lattice $\Lambda \subset \mathbb{C}$ thus gives rise to an elliptic curve over \mathbb{C} in the following way.

Theorem 1.3 ([79, Proposition VI.3.6]). *Let $\Lambda \subset \mathbb{C}$ be a lattice. Write E_Λ for the corresponding elliptic curve defined by the equation*

$$y^2 = 4x^3 - g_2(\Lambda)x - g_3(\Lambda).$$

The map $\Phi: \mathbb{C}/\Lambda \rightarrow E_\Lambda(\mathbb{C}) \subset \mathbb{P}^2(\mathbb{C})$ induced by $z \mapsto [\wp_\Lambda(z) : \wp'_\Lambda(z) : 1]$ is an isomorphism of Riemann surfaces which is also a group homomorphism from the additive group \mathbb{C}/Λ to the elliptic curve $E_\Lambda(\mathbb{C})$.

Every elliptic curve E/\mathbb{C} arises from a lattice in this way.

Theorem 1.4 ([79, Theorem VI.5.1]). *Let E/\mathbb{C} be an elliptic curve. Then, with the above notation, there exists a lattice $\Lambda \subset \mathbb{C}$ such that $E \cong E_\Lambda$.*

The question then arises of for which lattices $\Lambda_1, \Lambda_2 \subset \mathbb{C}$ are the associated elliptic curves $E_{\Lambda_1}, E_{\Lambda_2}$ isomorphic.

Proposition 1.5 ([79, Corollary VI.4.1.1]). *Let $\Lambda_1, \Lambda_2 \subset \mathbb{C}$ be lattices and $E_{\Lambda_1}, E_{\Lambda_2}$ the associated elliptic curves. Then $E_{\Lambda_1} \cong E_{\Lambda_2}$ over \mathbb{C} if and only if Λ_1, Λ_2 are homothetic, i.e. if and only if there exists some $\alpha \in \mathbb{C}^\times$ such that $\alpha\Lambda_1 = \Lambda_2$.*

To classify elliptic curves up to isomorphism, we therefore need only classify lattices up to homothety. Clearly, every lattice $\langle \tau_1, \tau_2 \rangle$ is homothetic to a lattice $\langle 1, \tau \rangle$ for some $\tau \in \mathbb{H}$. So we may reduce to considering when lattices $\langle 1, \tau \rangle$ and $\langle 1, \tau' \rangle$ with $\tau, \tau' \in \mathbb{H}$ are homothetic. They are homothetic if and only if there are $a, b, c, d \in \mathbb{Z}$ and $\alpha \in \mathbb{C}^\times$ such that $ad - bc \neq 0$ and

$$\alpha\tau' = a\tau + b \text{ and } \alpha = c\tau + d.$$

We thus see that $\langle 1, \tau \rangle$ and $\langle 1, \tau' \rangle$ are homothetic if and only if

$$\tau' = \frac{a\tau + b}{c\tau + d}$$

for some $a, b, c, d \in \mathbb{Z}$ with $ad - bc = 1$.

The group $\mathrm{SL}_2(\mathbb{Z})$ acts on \mathbb{H} by Möbius transformations. The previous paragraph shows that there is a bijection between the corresponding quotient $\mathrm{SL}_2(\mathbb{Z}) \backslash \mathbb{H}$ and the isomorphism classes of elliptic curves over \mathbb{C} .

This bijection sends (the orbit of) $\tau \in \mathbb{H}$ to (the isomorphism class of) the elliptic curve $E_{\langle 1, \tau \rangle}$. For $\tau \in \mathbb{H}$, we write E_τ for the elliptic curve $E_{\langle 1, \tau \rangle}$.

Corollary 1.6. *Every elliptic curve E/\mathbb{C} is isomorphic to an elliptic curve E_τ for some $\tau \in \mathbb{H}$ (and τ is unique up to the action of $\mathrm{SL}_2(\mathbb{Z})$ on \mathbb{H}).*

1.2. The j -invariant

Definition 1.7. Let E/K be an elliptic curve, which we may take to be isomorphic to an elliptic curve defined by a Weierstrass equation

$$y^2 = x^3 + Ax + B,$$

where $A, B \in K$ with $4A^3 + 27B^2 \neq 0$. The j -invariant of the elliptic curve E is defined

$$j(E) = 1728 \frac{4A^3}{4A^3 + 27B^2}.$$

The quantity

$$\frac{4A^3}{4A^3 + 27B^2}$$

is independent of the choice of a Weierstrass equation $y^2 = x^3 + Ax + B$ for E . The j -invariant is thus well-defined. The j -invariant of an elliptic curve is an isomorphism invariant of elliptic curves. If the field K is algebraically closed, then the j -invariant classifies elliptic curves over K up to isomorphism.

Theorem 1.8 ([79, Proposition III.1.4]). *Let K be a field and let \overline{K} be an algebraic closure of K . Then elliptic curves $E, E'/K$ are isomorphic over \overline{K} if and only if $j(E) = j(E')$.*

The j -invariant gives a field of definition for an elliptic curve.

Theorem 1.9. *Let $j_0 \in \mathbb{C}$. Then there exists an elliptic curve $E/\mathbb{Q}(j_0)$ such that $j(E) = j_0$.*

PROOF. Let $j_0 \in \mathbb{C}$. Then ([22, p. 378]) an example of an elliptic curve $E/\mathbb{Q}(j_0)$ with $j(E) = j_0$ is the curve with Weierstrass equation:

- (1) $y^2 = x^3 - 3\alpha x + 2\alpha$, where $\alpha = j_0/(j_0 - 1728)$, if $j_0 \notin \{0, 1728\}$;
- (2) $y^2 = x^3 - 1$ if $j_0 = 0$;
- (3) $y^2 = x^3 - x$ if $j_0 = 1728$. □

We now define the j -invariant of a lattice $\Lambda \subset \mathbb{C}$. To do this, we use the uniformisation theorem for elliptic curves.

Definition 1.10. Let $\Lambda \subset \mathbb{C}$ be a lattice. The j -invariant of the lattice Λ is defined by $j(\Lambda) = j(E_\Lambda)$, where E_Λ is the elliptic curve associated to the lattice Λ as in Theorem 1.3.

Since the field \mathbb{C} is algebraically closed, the combination of Theorems 1.4 and 1.8 shows that lattices Λ_1, Λ_2 have $j(\Lambda_1) = j(\Lambda_2)$ if and only if Λ_1, Λ_2 are homothetic.

We may now define the modular j -function. The j -function will be central to much of this thesis.

Definition 1.11. The modular j -function $j: \mathbb{H} \rightarrow \mathbb{C}$ is defined by setting $j(\tau) = j(\langle 1, \tau \rangle)$, the j -invariant of the lattice $\langle 1, \tau \rangle$. Equivalently, for each $\tau \in \mathbb{H}$, set $j(\tau) = j(E_\tau)$, the j -invariant of the elliptic curve E_τ .

Theorem 1.8 and Proposition 1.5 together imply the following fact.

Theorem 1.12. *Let $\tau, \tau' \in \mathbb{H}$. Then $j(\tau) = j(\tau')$ if and only if there exists $\gamma \in \mathrm{SL}_2(\mathbb{Z})$ such that $\tau' = \gamma\tau$.*

The j -function has many other nice properties with which we will become familiar; we state two of them now.

Theorem 1.13 ([23, Theorem 11.2]). *The function $j: \mathbb{H} \rightarrow \mathbb{C}$ is holomorphic and surjective.*

The j -function thus induces a bijection from $\mathrm{SL}_2(\mathbb{Z}) \backslash \mathbb{H}$ to \mathbb{C} .

1.3. Modular functions

Suppose that a function $f: \mathbb{H} \rightarrow \mathbb{C}$ is meromorphic and satisfies the identity $f(\tau) = f(\tau + 1)$ for all $\tau \in \mathbb{H}$. The function f may then be viewed as a function of the nome $q = \exp(2\pi i\tau)$ which is holomorphic on the region $0 < |q| < 1$. The function f thus has a Laurent expansion

$$f(\tau) = \sum_{n \in \mathbb{Z}} c_n q^n$$

for some (unique) $c_n \in \mathbb{C}$. We call this the q -expansion of f . We say that the function f is meromorphic at the cusp if $c_n = 0$ for all but finitely many $n < 0$.

Definition 1.14. A modular function $f: \mathbb{H} \rightarrow \mathbb{C}$ is a meromorphic function which is invariant under the action of $\mathrm{SL}_2(\mathbb{Z})$ on \mathbb{H} and also meromorphic at the cusp.

Since

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}),$$

observe that every $\mathrm{SL}_2(\mathbb{Z})$ -invariant function $f: \mathbb{H} \rightarrow \mathbb{C}$ satisfies the identity $f(\tau) = f(\tau + 1)$ for all $\tau \in \mathbb{H}$. We know that the j -function $j: \mathbb{H} \rightarrow \mathbb{C}$

from Section 1.2 is holomorphic and $\mathrm{SL}_2(\mathbb{Z})$ -invariant; it thus has a q -expansion.

Proposition 1.15 ([23, Theorem 11.8]). *There exist $c_n \in \mathbb{Z}_{>0}$ such that*

$$j(\tau) = \frac{1}{q} + 744 + \sum_{n=1}^{\infty} c_n q^n.$$

Remark 1.16. Proposition 1.15 explains why the factor of 1728 appears in the definition of the j -invariant in Definition 1.7. The factor is chosen so that the leading coefficient in the q -expansion of the j -function is equal to 1 [23, p. 225]. Petersson [58] and Rademacher [71] independently (see [72]) derived an explicit formula for the coefficients c_n , when $n \geq 1$. Their formula is given by an infinite series. Since this series converges relatively rapidly and every $c_n \in \mathbb{Z}$, one only needs to calculate the first few terms of the series in each case in order to compute c_n .

We can see from the q -expansion in Proposition 1.15 that the j -function is meromorphic at the cusp. Combined with Theorems 1.12 and 1.13, this implies that the j -function is a modular function. It is clear that the modular functions form a field. It turns out that this field is generated by the j -function.

Theorem 1.17 ([23, Theorem 11.9]). *The field of modular functions is $\mathbb{C}(j)$. Every modular function may thus be written as a rational function (with coefficients in \mathbb{C}) of j .*

The only holomorphic modular functions are thus the elements of the ring $\mathbb{C}[j]$. Of these, only the linear polynomials $aj + b$ with $a \neq 0$ induce

injective maps on the quotient $\mathrm{SL}_2(\mathbb{Z}) \backslash \mathbb{H}$. To maintain the property mentioned in Remark 1.16, of having a q -expansion with integer coefficients and leading coefficient 1, one must take $a = 1$ and $b \in \mathbb{Z}$. Given these constraints, the construction of j is thus canonical, up to modification by adding a factor $k \in \mathbb{Z}$.

It is sometimes convenient to make such a modification. For example, in [82], Zagier worked with the function $J(z) = j(z) - 744$, which has constant term 0 in its Fourier expansion. The aforementioned Petersson–Rademacher infinite series for the Fourier coefficients of j converges to 24 for $n = 0$. Therefore, one may prefer (see e.g. [22, p. 372] and [18, p. 675]) to work with the function $J_0(z) = j(z) - 720$, for which all the Fourier coefficients (except the leading coefficient, which we always normalise to 1) may be found from the Petersson–Rademacher series. In this thesis though, we will stick with convention, and so $c_0 = 744$ always.

We have seen already that the j -function is invariant under the action of $\mathrm{SL}_2(\mathbb{Z})$ on \mathbb{H} ; in fact, the j -function also behaves nicely under the action of the larger group $\mathrm{GL}_2^+(\mathbb{Q})$ on \mathbb{H} . For $g \in \mathrm{GL}_2^+(\mathbb{Q})$, write $N(g)$ for the determinant of the matrix $g' = \lambda g$, where $\lambda \in \mathbb{Q}^\times$ is chosen such that the matrix g' has coprime integer entries.

Proposition 1.18 ([23, Theorem 11.18]). *For $N \geq 1$, there exists a polynomial $\Phi_N(x, y) \in \mathbb{Z}[x, y]$ such that $\Phi_N(j(z), j(gz)) = 0$ identically in $z \in \mathbb{H}$, whenever $g \in \mathrm{GL}_2^+(\mathbb{Q})$ with $N(g) = N$. Further, $\Phi_1(x, y) = x - y$, and Φ_N is symmetric for all $N > 1$. We call Φ_N the N th modular polynomial.*

Let E_1, E_2 be elliptic curves. Then $\Phi_N(j(E_1), j(E_2)) = 0$ if and only if there is a cyclic N -isogeny between E_1 and E_2 , by [23, Theorem 11.23].

1.4. Modular curves

The group $\mathrm{SL}_2(\mathbb{Z})$ acts on \mathbb{H} by Möbius transformations. We may thus construct the corresponding quotient $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$. We have seen that each point of the quotient $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$ corresponds to an isomorphism class of elliptic curves over \mathbb{C} and every such class corresponds to a point on this quotient. Thus, $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$ is a moduli space for elliptic curves over \mathbb{C} . This quotient $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$ is the simplest example of a modular curve.

The quotient $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$ can be endowed (see e.g. [52, §2]) with the structure of a complex manifold; one thus obtains a non-compact Riemann surface, which we denote $Y(1)$. We may compactify it in the following way. Let $\mathbb{H}^* = \mathbb{H} \cup \mathbb{P}^1(\mathbb{Q})$. Observe that $\mathrm{SL}_2(\mathbb{Z})$ acts on \mathbb{H}^* in the same way as it does on \mathbb{H} , and so we may consider the quotient $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}^*$. This is equal to $(\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}) \cup \{\infty\}$, since all points of $\mathbb{P}^1(\mathbb{Q})$ belong to the same $\mathrm{SL}_2(\mathbb{Z})$ orbit. Thus, $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}^*$ may be endowed with a complex structure in the same way as for $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$. The resulting Riemann surface, denoted $X(1)$, is compact and has genus 0. Therefore, $X(1) \cong \mathbb{P}^1(\mathbb{C})$.

The function $j: \mathbb{H} \rightarrow \mathbb{C}$ induces a complex analytic isomorphism of the Riemann surfaces $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$ and \mathbb{C} . We will thus always identify the modular curve $Y(1)$ with the affine line \mathbb{C} by means of the j -function, giving us the slogan “the affine line is the j -line”. If we extend j by setting $j(\infty) = \infty$, then $j: X(1) \cong \mathbb{P}^1(\mathbb{C})$.

The field of modular functions is the function field of the compactified modular curve $X(1)$. The isomorphism $j: X(1) \cong \mathbb{P}^1(\mathbb{C})$ implies that the

j -function is a generator for this field, since the only meromorphic functions on $\mathbb{P}^1(\mathbb{C})$ are the rational functions.

In general, a modular curve is a quotient $\Gamma \backslash \mathbb{H}$, where $\Gamma \leq \mathrm{SL}_2(\mathbb{Z})$ is a congruence subgroup. A congruence subgroup of $\mathrm{SL}_2(\mathbb{Z})$ is a subgroup of $\mathrm{SL}_2(\mathbb{Z})$ which contains, for some $N \geq 1$, the group

$$\Gamma(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) : a \equiv d \equiv 1 \pmod{N}, b \equiv c \equiv 0 \pmod{N} \right\}.$$

Let Γ be a congruence subgroup of $\mathrm{SL}_2(\mathbb{Z})$. Then Γ acts on \mathbb{H} , and so one may consider the quotient $\Gamma \backslash \mathbb{H}$. This quotient may also (see e.g. [52, §2]) be given the structure of a (non-compact) Riemann surface in a natural way; we denote this as the modular curve $Y(\Gamma)$. Once again, we can compactify this by considering instead the quotient $\Gamma \backslash \mathbb{H}^*$. This may be endowed with the structure of a compact Riemann surface, denoted $X(\Gamma)$, and thus also has the structure of an algebraic curve.

We saw above that the modular curve $Y(1)$ was the moduli space for elliptic curves over \mathbb{C} ; points of $Y(1)$ correspond to isomorphism classes of elliptic curves. When Γ is a general congruence subgroup, the associated modular curve $Y(\Gamma)$ also has a moduli interpretation. Its points represent isomorphism classes of elliptic curves together with some extra data.

For example, consider the congruence subgroups of the form

$$\Gamma_1(N) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z}) : a \equiv d \equiv 1 \pmod{N}, c \equiv 0 \pmod{N} \right\}.$$

Since $\Gamma_1(N) \subset \mathrm{SL}_2(\mathbb{Z})$, if $\tau, \tau' \in \mathbb{H}$ belong to the same $\Gamma_1(N)$ orbit, then the elliptic curves E_τ and $E_{\tau'}$ are isomorphic. Observe also that the point

$1/N \in E_\tau$ has order precisely N and is invariant under the action of $\Gamma_1(N)$. Points of $\Gamma_1(N)\backslash\mathbb{H}$ thus correspond to equivalence classes of pairs (E, P) , where E is an elliptic curve over \mathbb{C} and $P \in E$ has exact order N . Two such pairs $(E, P), (E', P')$ are equivalent if there is an isomorphism from E to E' taking P to P' . The modular curve $Y_1(N) = \Gamma_1(N)\backslash\mathbb{H}$ is thus the moduli space for such pairs. See [79, Appendix C, §13] for further details.

1.5. Imaginary quadratic fields

The material in this section is based on [23, §7]. Let K be an imaginary quadratic field, so that $K = \mathbb{Q}(\sqrt{-n})$ for some unique square-free integer $n \geq 1$. The discriminant D of the imaginary quadratic field K is defined

$$D = \begin{cases} -n & \text{if } n \equiv 3 \pmod{4} \\ -4n & \text{otherwise.} \end{cases}$$

In particular, there is a unique imaginary quadratic field of a given discriminant. Note that $K = \mathbb{Q}(\sqrt{D})$.

Definition 1.19. An order \mathcal{O} in K is a subring $\mathcal{O} \subset K$ such that \mathcal{O} is a free \mathbb{Z} -module of rank 2 which contains 1.

The ring of integers \mathcal{O}_K of K is an order in K ; every other order \mathcal{O} of K is contained in \mathcal{O}_K . We thus call \mathcal{O}_K the maximal order. The order \mathcal{O}_K is equal to $\mathbb{Z} + \mathbb{Z} \cdot \omega$, where $\omega = (D + \sqrt{D})/2$. Any order \mathcal{O} in K is equal to $\mathbb{Z} + \mathbb{Z} \cdot (f \cdot \omega)$ for some unique integer $f \geq 1$. Call f the conductor of the order \mathcal{O} . One has that $f = [\mathcal{O}_K : \mathcal{O}]$.

The discriminant Δ of an order \mathcal{O} in an imaginary quadratic field K is defined

$$\Delta = \left(\det \begin{pmatrix} \alpha & \beta \\ \alpha' & \beta' \end{pmatrix} \right)^2,$$

where α, β are generators of \mathcal{O} and $x \mapsto x'$ denotes the non-trivial automorphism of K . (This is independent of the choice of α, β .) By taking the generators $1, f \cdot w$ for \mathcal{O} , one obtains that $\Delta = f^2 D$. (The discriminant of the maximal order \mathcal{O}_K is thus equal to the discriminant of the imaginary quadratic field K .) Observe that $K = \mathbb{Q}(\sqrt{\Delta}) = \mathbb{Q}(\sqrt{D})$.

The discriminant Δ of an imaginary quadratic order must satisfy $\Delta < 0$ and $\Delta \equiv 0, 1 \pmod{4}$. Conversely, for every integer $\Delta < 0$ such that $\Delta \equiv 0, 1 \pmod{4}$, there is a unique pair (K, \mathcal{O}) such that \mathcal{O} is an order of discriminant Δ in the imaginary quadratic field K .

Now let \mathcal{O} be an order of an imaginary quadratic field K . Denote by $I(\mathcal{O})$ the set of proper fractional ideals of \mathcal{O} . This is a multiplicative group. Let $P(\mathcal{O}) \subset I(\mathcal{O})$ be the subset of principal ideals. This forms a normal subgroup of $I(\mathcal{O})$. The ideal class group of the order \mathcal{O} is defined to be the quotient group $\text{cl}(\mathcal{O}) = I(\mathcal{O})/P(\mathcal{O})$. The class number $h(\mathcal{O})$ of the order \mathcal{O} is defined by $h(\mathcal{O}) = |\text{cl}(\mathcal{O})|$. Since the discriminant Δ uniquely determines the corresponding order \mathcal{O} , we will often write $\text{cl}(\Delta)$ and $h(\Delta)$ in place of $\text{cl}(\mathcal{O})$ and $h(\mathcal{O})$.

1.6. Complex multiplication of elliptic curves

Recall that modular curves parametrise isomorphism classes of elliptic curves (plus possibly some extra data). Additional structure on modular curves comes from the endomorphism type of the parametrised elliptic

curves. In particular, those elliptic curves with “extra” endomorphisms correspond to points on the modular curve with arithmetic significance.

Definition 1.20. Let E_1, E_2 be elliptic curves. An isogeny $E_1 \rightarrow E_2$ is a morphism $\phi: E_1 \rightarrow E_2$ satisfying $\phi(0) = 0$.

Denote

$$\text{Hom}(E_1, E_2) = \{\text{isogenies } \phi: E_1 \rightarrow E_2\}.$$

For $\phi_1, \phi_2 \in \text{Hom}(E_1, E_2)$, we can define $\phi_1 + \phi_2 \in \text{Hom}(E_1, E_2)$ by setting

$$(\phi_1 + \phi_2)(P) = \phi_1(P) + \phi_2(P),$$

where the addition on the right hand side is in the group law of E_2 . Since the points of an elliptic curve form an abelian group, doing this makes $\text{Hom}(E_1, E_2)$ into an additive group.

Now consider the isogenies from an elliptic curve E to itself. We denote $\text{End}(E) = \text{Hom}(E, E)$. For $\phi_1, \phi_2 \in \text{End}(E)$, we may define their product $\phi_1 \cdot \phi_2$ by setting

$$(\phi_1 \cdot \phi_2)(P) = \phi_1(\phi_2(P)).$$

Then $\text{End}(E)$ forms a ring under these definitions for addition and multiplication of isogenies $E \rightarrow E$ [79, Proposition III.4.2].

The ring $\text{End}(E)$ always contains \mathbb{Z} , where the element $n \in \mathbb{Z}$ represents the multiplication-by- n map

$$[n]: E \rightarrow E \quad P \mapsto \underbrace{P + \dots + P}_{n \text{ times}}.$$

For most elliptic curves E , $\text{End}(E) = \mathbb{Z}$. However, some elliptic curves have a strictly larger endomorphism ring.

Definition 1.21. Let E be an elliptic curve. We say that E has complex multiplication (CM) if $\text{End}(E) \supsetneq \mathbb{Z}$.

Section 1.1 established a correspondence between elliptic curves E/\mathbb{C} and lattices $\Lambda \subset \mathbb{C}$. We may use this to establish the following.

Theorem 1.22 ([79, Theorem VI.4.1]). *Let $\Lambda_1, \Lambda_2 \subset \mathbb{C}$ be lattices, and let E_1, E_2 be the associated elliptic curves. Then the map*

$$\{\alpha \in \mathbb{C} : \alpha\Lambda_1 \subset \Lambda_2\} \rightarrow \{\text{holomorphic maps } \phi: \mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2 \text{ with } \phi(0) = 0\}$$

given by $\alpha \mapsto \phi_\alpha$, where $\phi_\alpha(z) = \alpha z \pmod{\Lambda_2}$, is a bijection. Further, the natural map

$$\text{Hom}(E_1, E_2) \rightarrow \{\text{holomorphic maps } \phi: \mathbb{C}/\Lambda_1 \rightarrow \mathbb{C}/\Lambda_2 \text{ with } \phi(0) = 0\}$$

is also bijective.

Let E/\mathbb{C} be an elliptic curve. So $E \cong E_\tau$ for some $\tau \in \mathbb{H}$. By the above, we have that

$$\text{End}(E) \cong \{\alpha \in \mathbb{C} : \alpha\langle 1, \tau \rangle \subset \langle 1, \tau \rangle\}.$$

Clearly $\mathbb{Z} \subset \{\alpha \in \mathbb{C} : \alpha\langle 1, \tau \rangle \subset \langle 1, \tau \rangle\}$ and, in the generic case,

$$\mathbb{Z} = \{\alpha \in \mathbb{C} : \alpha\langle 1, \tau \rangle \subset \langle 1, \tau \rangle\}.$$

However, if $\tau \in \mathbb{H}$ is an imaginary quadratic number (i.e. $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$), then the ring $\text{End}(E)$ is strictly larger than \mathbb{Z} . In such a case, $\text{End}(E)$ is an order in the imaginary quadratic field $\mathbb{Q}(\tau)$. (If $\text{End}(E) \neq \mathbb{Z}$, then one

sees that for $\alpha \in \text{End}(E) \setminus \mathbb{Z}$ one has that $\alpha \in \mathbb{C} \setminus \mathbb{R}$. This explains the terminology “complex multiplication”.)

Proposition 1.23. *Let $\tau \in \mathbb{H}$. The elliptic curve E_τ has CM if and only if τ is imaginary quadratic. In this case, $\text{End}(E_\tau)$ is an order in the imaginary quadratic field $\mathbb{Q}(\tau)$.*

The points on a modular curve $\Gamma \backslash \mathbb{H}$ corresponding to $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$ therefore correspond to the CM elliptic curves (together with the relevant extra structure parametrised by the modular curve). We call these points the special points of the modular curve. The Diophantine properties of these special points will be the major focus of this thesis.

Definition 1.24. Let $Y = \Gamma \backslash \mathbb{H}$ be a modular curve. A special point of Y is a point of Y which corresponds to some $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$. Equivalently, a special point of Y is a point of Y corresponding to a $\tau \in \mathbb{H}$ such that the elliptic curve E_τ has CM.

1.7. Singular moduli

Definition 1.25. A singular modulus is the j -invariant of a CM elliptic curve. Equivalently, a number of the form $j(\tau)$ for some $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$.

Since we have identified the modular curve $Y(1)$ with \mathbb{C} by means of the j -function, the special points of $Y(1)$ are precisely the singular moduli.

Definition 1.26. Let $\sigma = j(\tau)$ be a singular modulus. The discriminant Δ of the singular modulus σ is defined to be equal to the discriminant of the imaginary quadratic order $\text{End}(E_\tau)$.

Proposition 1.27 ([23, Theorem 7.7]). *Let $\sigma = j(\tau)$ be a singular modulus of discriminant Δ . Suppose $a, b, c \in \mathbb{Z}$ are such that $a\tau^2 + b\tau + c = 0$ and $\gcd(a, b, c) = 1$. Then $\Delta = b^2 - 4ac$.*

Denote by F_j the standard fundamental domain for the action of $\mathrm{SL}_2(\mathbb{Z})$ on \mathbb{H} , given by (see e.g. [83, Proposition 1])

$$F_j = \left\{ z \in \mathbb{H} : -\frac{1}{2} \leq \mathrm{Re}(z) < \frac{1}{2}, |z| \geq 1, \text{ and } |z| > 1 \text{ for } 0 < \mathrm{Re}(z) < \frac{1}{2} \right\}.$$

The map j restricts to a bijection $j|_{F_j} : F_j \rightarrow \mathbb{C}$. Thus, for each singular modulus σ , there is a unique $\tau \in F_j$ such that $j(\tau) = \sigma$. We may thereby give an explicit description of the singular moduli of a given discriminant.

Let T_Δ be the set of triples $(a, b, c) \in \mathbb{Z}^3$ such that:

- (1) $\Delta = b^2 - 4ac$,
- (2) $\gcd(a, b, c) = 1$, and
- (3) either $-a < b \leq a < c$ or $0 \leq b \leq a = c$.

The set of singular moduli of discriminant Δ is then equal to

$$\left\{ j\left(\frac{-b + \sqrt{\Delta}}{2a}\right) : (a, b, c) \in T_\Delta \right\}.$$

Since $\frac{-b + \sqrt{\Delta}}{2a} \in F_j$ for every $(a, b, c) \in T_\Delta$, the numbers $j\left(\frac{-b + \sqrt{\Delta}}{2a}\right)$ are all distinct. In particular, the number of singular moduli of discriminant Δ is equal to $|T_\Delta|$. In fact, we have [83, pp. 71–72] that $|T_\Delta| = h(\Delta)$.

Now let $\Delta < 0$ satisfy $\Delta \equiv 0, 1 \pmod{4}$. Enumerate as $\sigma_1, \dots, \sigma_n$ the singular moduli of discriminant Δ . The Hilbert class polynomial of discriminant Δ is then defined

$$H_\Delta(X) = \prod_{i=1}^n (X - \sigma_i).$$

One has the following remarkable theorem. In particular, it implies that the singular moduli of discriminant Δ form a full Galois orbit over K and over \mathbb{Q} .

Theorem 1.28 ([83, Proposition 25]). *Let $\Delta < 0$ satisfy $\Delta \equiv 0, 1 \pmod{4}$. The polynomial $H_\Delta(X)$ is irreducible over K and has integer coefficients. Therefore, a singular modulus of discriminant Δ is an algebraic integer of degree $h(\Delta)$ over K (and over \mathbb{Q} also).*

Let \mathcal{O} be an order in an imaginary quadratic field K . The ring class field L of the order \mathcal{O} is the abelian extension of K which has Galois group $\text{Gal}(L/K) \cong \text{cl}(\mathcal{O})$. If $\mathcal{O} = \mathcal{O}_K$, then the ring class field of \mathcal{O} is the Hilbert class field of K , i.e. the maximal abelian unramified extension of K . The ring class field of \mathcal{O} is generated by the j -invariant of any one of the elliptic curves with CM by \mathcal{O} .

Theorem 1.29 ([23, Theorem 11.1]). *Let \mathcal{O} be an order in an imaginary quadratic field K . Let E be an elliptic curve with CM by \mathcal{O} . Then $K(j(E))$ is the ring class field of the order \mathcal{O} .*

The Kronecker–Weber theorem implies that the maximal abelian extension of \mathbb{Q} is obtained by adjoining all roots of unity. The theory of complex multiplication implies that, for imaginary quadratic fields, the singular moduli play a roughly analogous role. This illustrates the great arithmetic significance of the special points which we investigate.

1.8. Shimura curves

The group $\text{SL}_2(\mathbb{R})$ acts on \mathbb{H} by Möbius transformations. We have seen that the congruence subgroups of $\text{SL}_2(\mathbb{Z})$ are of significant arithmetic

interest, because they give rise to the modular curves. There are other arithmetically interesting subgroups of $\mathrm{SL}_2(\mathbb{R})$ though, beyond these congruence subgroups. Of particular interest to us are those which give rise to Shimura curves.

Shimura curves are examples of Shimura varieties. We will not give an introduction to the theory of these here, but instead refer to [50]. The prototypical examples of Shimura varieties are the Siegel moduli varieties (see e.g. [35]). For $g \geq 1$, denote by \mathcal{A}_g the Siegel moduli variety of dimension g . We note that \mathcal{A}_g is the moduli space for principally polarised abelian varieties of dimension g . In particular, $\mathcal{A}_1 = Y(1)$.

Definition 1.30. Let S be a Shimura variety. Then S is said to be of Hodge type if S admits a monomorphism into some \mathcal{A}_g . If S admits an isogeny to some Shimura variety S' of Hodge type, then S is said to be of abelian type.

Definition 1.31. A Shimura curve is a connected Shimura variety of dimension one which is of abelian type.¹

A connected component of a Shimura variety arises as the quotient of a Hermitian symmetric domain D by a subgroup $\Gamma \leq \mathrm{Hol}(D)$ satisfying certain conditions. Here $\mathrm{Hol}(D)$ denotes the group of holomorphic automorphisms of D . The only one-dimensional Hermitian symmetric domain is the complex upper half plane \mathbb{H} . We have that $\mathrm{Hol}(\mathbb{H}) = \mathrm{SL}_2(\mathbb{R})$. Thus, every Shimura curve arises as a quotient $\Gamma \backslash \mathbb{H}$ for some subgroup $\Gamma \leq \mathrm{SL}_2(\mathbb{R})$.

¹This definition is not completely standard, but is convenient for our purposes.

Let Y be a Shimura curve. Then Y is isogenous to some Shimura variety Y' of Hodge type, since Y is a Shimura variety of abelian type. Fix an isogeny $Y \rightarrow Y'$. Now Y' admits a monomorphism into \mathcal{A}_g , for some $g \geq 1$. Fix such a map. We may thereby associate to each point of Y a point of \mathcal{A}_g . So we may associate (the isomorphism class of) a principally polarised abelian variety of dimension g to each point of Y , since \mathcal{A}_g is the moduli space for principally polarised abelian varieties of dimension g .

Elliptic curves are one-dimensional principally polarised abelian varieties. Just as for elliptic curves, there is a corresponding (though more complicated) theory of complex multiplication for general abelian varieties. For details of this, see [56].

Definition 1.32. For a Shimura curve Y , the special points of Y are those points which correspond to CM abelian varieties under the relevant map from Y to some \mathcal{A}_g described above.

We note the analogy with Definition 1.24, which defines special points of modular curves. The Diophantine properties of special points of modular and Shimura curves are the focus of this thesis.

Proposition 1.33. *Let Y be a Shimura curve with uniformising map $\pi: \mathbb{H} \rightarrow Y$. There exists $d \geq 1$ such that: for all $\tau \in \mathbb{H}$, if $\pi(\tau)$ is a special point of Y , then $[\mathbb{Q}(\tau) : \mathbb{Q}] \leq d$.*

PROOF. See [24, §1.2]. □

The discriminant of a special point of a Shimura curve Y is defined to be the discriminant of the corresponding CM abelian variety (i.e. the discriminant of the centre of the endomorphism ring of the abelian variety).

1.9. Plan of this thesis

Let Y be a modular or Shimura curve. Then we have seen that Y comes with a collection of special points, which are distinguished by their close connection to CM abelian varieties. The algebraic properties of these special points are thus an area of interest. In particular, a natural question to consider is, given an algebraic group G and a map $\phi: Y \rightarrow G$, when are images of special points of Y dependent in the group law of G .

Versions of this question have been considered for various such choices of G, ϕ . For example, one area of particular interest has been the case of modular parametrisations of elliptic curves. The Modularity Theorem implies that, for each elliptic curve E/\mathbb{Q} , there exists a modular curve Y and a non-constant morphism $\phi: Y \rightarrow E$. Rosen and Silverman [75], followed by e.g. [5, 77], considered the question of when the corresponding ϕ -images of special points are dependent in the group law of E . More widely, the case of more general maps from a modular or Shimura curve to an elliptic curve was studied by Buium and Poonen [20] and Pila and Tsimerman [63]. See also the related work of Baldi [7] and Kühne [44].

In this thesis, we consider the case where $G = \mathbb{G}_m$, the multiplicative group of complex numbers. Dependence in the group law of \mathbb{G}_m is thus multiplicative dependence: a set $\{x_1, \dots, x_n\}$ is called multiplicatively dependent if there are some $a_1, \dots, a_n \in \mathbb{Z}$, not all zero, such that $x_1^{a_1} \cdots x_n^{a_n} = 1$.

Definition 1.34. A correspondence $V \subset Y \times \mathbb{G}_m$ is an irreducible closed subvariety of $Y \times \mathbb{G}_m$ which has dimension 1 and projects dominantly to both Y and \mathbb{G}_m .

Fix a correspondence $V \subset Y \times \mathbb{G}_m$. Given $(s, x) \in Y \times \mathbb{G}_m$, call x a V -image of s if $(s, x) \in V$. Our main aim in this thesis is to investigate the question of when V -images of special points of Y are multiplicatively dependent.

One expects that this happens relatively rarely. Indeed, ultimately one would like to show that it only happens in “trivial” cases. In Chapter 5, we prove such a result for one very special case of this problem. In greater generality though, we are not able to do this.

In Chapter 2, we will prove instead that, for each $n \geq 1$, if n V -images x_1, \dots, x_n of special points s_1, \dots, s_n are multiplicatively dependent, then the point (s, x) must belong to one of finitely many components of V^n of a particular kind. This result is a finiteness result on the different possible “shapes” of multiplicative dependences among n V -images of special points.

In Chapter 3, we will show that, in a more restricted setting, one can in fact prove, for each $n \geq 1$, the finiteness of n -tuples of distinct V -images of special points which are multiplicatively dependent and minimal for this property.

The results of Chapters 2 and 3 and, more generally, the question of multiplicatively dependent V -images of special points are intimately related to the concept of so-called “atypical intersections” and, in particular, the Zilber–Pink conjecture. The Zilber–Pink conjecture is a major conjecture in Diophantine geometry. In Chapter 4, we discuss this connection and prove some special cases of the Zilber–Pink conjecture in low dimensions.

Our proofs in Chapters 2–4 will make use of techniques from model theory, in particular the o-minimal Counting Theorem of Pila and Wilkie

[66]. We include in Appendix A an introduction to the results from model theory which we use, intended for those readers who are unfamiliar with such techniques.

Our uses of o-minimality to prove Diophantine results in Chapters 2–4 will follow the so-called Pila–Zannier strategy. This strategy, using the o-minimal Counting Theorem, was suggested by Zannier and first used by Pila and Zannier [67] to re-prove the Manin–Mumford conjecture. Subsequently, this approach has been widely applied to prove Diophantine results, see e.g. [78] for a survey of some of these results. In common with most Diophantine results proved using o-minimality, the finiteness results of Chapters 2–4 are all ineffective.

In Chapter 5, we prove an effective result on triples of singular moduli with rational product. The proof of this result is based on only arithmetic ingredients and makes use of extensive computations in the package PARI [57]. We deduce from this result a completely explicit André–Oort statement for the family of cubic surfaces in $Y(1)^3$ defined by equations of the form $x_1x_2x_3 = \alpha \in \mathbb{Q}$. The André–Oort conjecture, which was proved (ineffectively) for subvarieties of products of modular curves by Pila [60], is a special case of the Zilber–Pink conjecture. We describe the André–Oort conjecture in Chapter 5.

CHAPTER 2

Images of special points in \mathbb{G}_m

2.1. The mixed Shimura variety setting

Let Y be a modular or Shimura curve. Then Y has a uniformisation $q: \mathbb{H} \rightarrow Y$. Recall that \mathbb{G}_m denotes the multiplicative group of complex numbers. For $m, n \geq 1$, let $X = X_{m,n} = Y^m \times \mathbb{G}_m^n$. Then X is an example of a mixed Shimura variety. We will not discuss these in general, but instead refer the reader to [49] for background.

Let $e: \mathbb{C} \rightarrow \mathbb{G}_m$ be given by $e(u) = \exp(2\pi i u)$. Let $U = U_{m,n} = \mathbb{H}^m \times \mathbb{C}^n$. Then the mixed Shimura variety X has a uniformisation $\pi = \pi_{m,n}: U \rightarrow X$ given by

$$\pi(z_1, \dots, z_m, u_1, \dots, u_n) = (q(z_1), \dots, q(z_m), e(u_1), \dots, e(u_n)).$$

Let F_Y denote a fundamental domain for the uniformisation $q: \mathbb{H} \rightarrow Y$. (When $Y = Y(1)$, we will always take $F_Y = F_j$, where F_j is as defined in Section 1.7.) Let $F_e = \{z \in \mathbb{C} : 0 \leq \operatorname{Re} z < 1\}$, which is a fundamental domain for the map $e: \mathbb{C} \rightarrow \mathbb{G}_m$. Then $F = F_{m,n} = F_Y^m \times F_e^n$ is a fundamental domain for the uniformisation $\pi: U \rightarrow X$. We note here the fact, which we will rely on repeatedly, that the map π , restricted to the fundamental domain F , is definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$. See Appendix A for details.

Since X is a mixed Shimura variety, it comes with a collection of “weakly special” subvarieties, included among which are the “special” subvarieties. We now describe these. To do so, we first need to describe the (weakly) special subvarieties of \mathbb{G}_m^n and Y^m . We begin with the case of \mathbb{G}_m^n ; our description is based on [84, p. 15].

2.1.1. The case of \mathbb{G}_m^n . There is a one-to-one correspondence between the set of lattices $\Lambda \subset \mathbb{Z}^n$ and the set of algebraic subgroups of \mathbb{G}_m^n . This correspondence may be given explicitly by the map

$$\Lambda \mapsto \{(x_1, \dots, x_n) \in \mathbb{G}_m^n : x_1^{a_1} \cdots x_n^{a_n} = 1 \text{ for every } (a_1, \dots, a_n) \in \Lambda\}.$$

Write G_Λ for the algebraic subgroup which corresponds to the lattice Λ in this manner. Then G_Λ is irreducible if and only if the lattice Λ is primitive.

A weakly special subvariety of \mathbb{G}_m^n is a coset of an irreducible algebraic subgroup of \mathbb{G}_m^n . We note that, for $\Lambda \subset \mathbb{Z}^n$ a lattice, any coset of G_Λ is of the form

$$\{(x_1, \dots, x_n) \in \mathbb{G}_m^n : x_1^{a_1} \cdots x_n^{a_n} = c_a \text{ for every } (a_1, \dots, a_n) \in \Lambda\}$$

for some suitable constants $c_a \in \mathbb{G}_m$.

A special point of \mathbb{G}_m^n is a torsion point. A special subvariety of \mathbb{G}_m^n is a weakly special subvariety which contains a special point. Equivalently, a weakly special subvariety is special if and only if the constants c_a in the above description are all roots of unity. In particular, a special subvariety of dimension 0 is a special point. A special subvariety of \mathbb{G}_m^n is always a component of an algebraic subgroup of \mathbb{G}_m^n of the same dimension.

2.1.2. The case of Y^m . We begin with a definition from group theory.

Let Γ be a subgroup of a group G . The commensurator of Γ in G is defined

$$\text{Comm}_G(\Gamma) = \{g \in G : \Gamma \cap g\Gamma g^{-1} \text{ has finite index in both } \Gamma \text{ and } g\Gamma g^{-1}\}.$$

It is straightforward to verify that $\text{Comm}_{\text{GL}_2(\mathbb{R})}(\Gamma) = \mathbb{R}^\times \cdot \text{GL}_2(\mathbb{Q})$ for every congruence subgroup Γ of $\text{SL}_2(\mathbb{Z})$.

Recall that the modular or Shimura curve Y arises as a quotient $\Gamma \backslash \mathbb{H}$ for some subgroup $\Gamma \leq \text{GL}_2(\mathbb{R})$. Let G denote the neutral component¹ of $\text{Comm}_{\text{GL}_2(\mathbb{R})}(\Gamma)$. (So, if Y is a modular curve, then $G = \mathbb{R}^\times \cdot \text{GL}_2^+(\mathbb{Q})$.) Let S_0, \dots, S_k be pairwise disjoint subsets of $\{1, \dots, m\}$ such that $k \geq 0$,

$$\bigcup_{i=0}^k S_i = \{1, \dots, m\},$$

and, for every $i > 0$, $S_i \neq \emptyset$. For each $i \in S_0$, fix some $h_i \in \mathbb{H}$. For each $i = 1, \dots, k$, fix some $s_i \in S_i$ and, for every $j \in S_i \setminus \{s_i\}$, fix some $g_{i,j} \in G$.

Let

$$W = \{(u_1, \dots, u_m) \in \mathbb{H}^m : \forall i \in S_0 \quad u_i = h_i \\ \forall i = 1, \dots, k \quad \forall j \in S_i \setminus \{s_i\} \quad u_j = g_{i,j} u_{s_i}\}.$$

A weakly special subvariety of \mathbb{H}^m is a set W of this kind for some choice of the associated data $k, S_0, \dots, S_k, h_i, s_i, g_{i,j}$. Note that this definition depends on Y (insofar as Y determines G).

A weakly special subvariety of Y^m is the image of a weakly special subvariety of \mathbb{H}^m under (the Cartesian product of) the map $q: \mathbb{H} \rightarrow Y$. The special points of Y were described in Chapter 1. A special point of Y^m is a

¹i.e. the connected component containing the identity element

point $(s_1, \dots, s_m) \in Y^m$ such that each s_i is a special point of Y . A special subvariety of Y^m is a weakly special subvariety which contains a special point. A strongly special subvariety of Y^m is a special subvariety which has no constant coordinates (equivalently, a strongly special subvariety is a special subvariety which projects dominantly to every coordinate).

We have the following technical fact.

Lemma 2.1 ([63, Lemma 2.1]). *Let $M \subset Y^m$ be a weakly special subvariety. Then there are pairwise disjoint sets $A, B_i \subset \{1, \dots, n\}$, where i runs over an index set I , such that*

$$A \cup \bigcup_{i \in I} B_i = \{1, \dots, n\}$$

and also:

- (1) for every $i \in I$, the set B_i is non-empty;
- (2) for every $a \in A$, the image of the projection $\pi_a: M \rightarrow Y^{\{a\}}$ is a point;
- (3) for every $i \in I$ and $a, b \in B_i$ distinct, the image of the projection $\pi_{a,b}: M \rightarrow Y^{\{a,b\}}$ is a strongly special subvariety of $Y^{\{a,b\}}$.

(Note that we do not require A or I to be non-empty.) Further, M is a special subvariety if and only if, for every $a \in A$, the image of the projection $\pi_a: M \rightarrow Y^{\{a\}}$ is a special point of Y .

For a weakly special subvariety $M \subset Y^m$, we thus have that M is a special subvariety of Y^m if and only if every constant coordinate of M is a special point of Y .

2.1.3. The case of X . A (weakly) special subvariety of $X_{m,n}$ is the product of a (weakly) special subvariety of Y^m with a (weakly) special subvariety of \mathbb{G}_m^n . A special point of X is a point (s, x) such that s is a special point of Y^m and x is a special point of \mathbb{G}_m^n . We thus see that the special subvarieties of X are precisely the weakly special subvarieties which contain a special point.

2.2. Exemplary components

In this chapter, for each $n \geq 1$, we denote by π_1 and π_2 the projections of $Y^n \times \mathbb{G}_m^n$ onto the Y^n and \mathbb{G}_m^n coordinates respectively. (Note that we will always use the same notation for these projections, even as n varies. It should always be clear from context which map is meant.)

For the remainder of this chapter, fix $V \subset Y \times \mathbb{G}_m$ an irreducible correspondence defined over $\overline{\mathbb{Q}}$. To state the main result of this chapter, we need the notion of an exemplary component, which was first formulated (in a different setting) by Pila and Tsimerman [63].

Definition 2.2. Let $n \geq 1$, so that V^n is a subvariety of the mixed Shimura variety $Y^n \times \mathbb{G}_m^n$.

- (1) A distinguished component of V^n is a (geometrically-) irreducible component $W \subset V^n \cap (S \times B)$, where $S \subset Y^n$ and $B \subset \mathbb{G}_m^n$ are special subvarieties, such that $\pi_1(W) = S$.
- (2) A distinguished component W of V^n is said to be dependent if, in the definition of a distinguished component in (1), the special subvariety B (may be taken such that it) is a proper special subvariety of \mathbb{G}_m^n .

- (3) A distinguished component W of V^n is called an exemplary component of V^n if, setting B to be the smallest special subvariety of \mathbb{G}_m^n containing $\pi_2(W)$, there is no distinguished component $W' \supsetneq W$ with $\pi_2(W') \subset B$.

Note that a dependent distinguished component is always contained in some (not necessarily unique) dependent exemplary component. Observe also that the unique non-dependent exemplary component of V^n is V^n itself.

Suppose that $(s, x) \in V^n$, where s is a special point of Y^n . If $(s, x) \in W$ for some dependent distinguished component W , then the coordinates x_1, \dots, x_n of x are multiplicatively dependent, since they lie in some proper special subvariety of \mathbb{G}_m^n . Conversely, if x_1, \dots, x_n are multiplicatively dependent, then x lies in some proper special subvariety of \mathbb{G}_m^n , and hence (s, x) is contained in some dependent distinguished component.

The main theorem of this chapter is the following finiteness result for exemplary components. The analogous result when \mathbb{G}_m is replaced by an elliptic curve was proved by Pila and Tsimerman [63]. Our proof of Theorem 2.3 is based on their approach. In particular, the proof will proceed by an o-minimal counting argument.

Theorem 2.3. *Let $n \geq 1$. Then there are only finitely many exemplary components of V^n .*

Fix some $n \geq 1$. For $A \subset \mathbb{G}_m^n$, denote by $\langle A \rangle$ the smallest special subvariety of \mathbb{G}_m^n containing A . By Theorem 2.3, there are only finitely many exemplary components of V^n , which may thus be enumerated as $W_{i,j,k}$. Here $W_{i,j,k}$ is an irreducible component of $V^n \cap (S_{i,j} \times B_i)$ such that

$S_{i,j} = \pi_1(W_{i,j,k})$ and $B_i = \langle \pi_2(W_{i,j,k}) \rangle$ are special subvarieties of Y^n and \mathbb{G}_m^n respectively and

- (1) $i = 1, \dots, r$ for some $r \geq 1$;
- (2) $j = 1, \dots, s_i$ for some $s_i \geq 1$; and
- (3) $k = 1, \dots, t_{i,j}$ for some $t_{i,j} \geq 1$.

Suppose then that $(s, x) \in V^n$ is such that s is a special point of Y^n . Let $B = \langle \{x\} \rangle$. By the definition of an exemplary component, there is then some index (i, j, k) such that $(s, x) \in W_{i,j,k}$ and $B_i = B$. If $(s', x') \in W_{i,j,k}$, then:

$$\forall a \in \mathbb{Z}^n (x_1^{a_1} \cdots x_n^{a_n} = 1 \implies (x'_1)^{a_1} \cdots (x'_n)^{a_n} = 1),$$

since $x' \in B_i = \langle \{x\} \rangle$.

For $(s, x) \in V^n$ such that s is a special point of Y^n , the possible multiplicative relations

$$x_1^{a_1} \cdots x_n^{a_n} = 1,$$

where $a \in \mathbb{Z}^n$, satisfied by the coordinates of x are therefore completely characterised by the list of exemplary components of V^n that contain (s, x) . Theorem 2.3 therefore implies the finiteness of the possible ‘‘shapes’’ of multiplicative dependencies amongst V -images of special points of Y^n . It is thus a first step towards answering the question, posed in Section 1.9, of when V -images of special points are multiplicatively dependent.

2.3. Ax–Schanuel

Recall the setting of Section 2.1. So, for $m, n \geq 1$, we let $U = \mathbb{H}^m \times \mathbb{C}^n$ and $X = Y^m \times \mathbb{G}_m^n$, and then take the uniformisation $\pi: U \rightarrow X$ given by

$$\pi(z_1, \dots, z_m, u_1, \dots, u_n) = (q(z_1), \dots, q(z_m), e(u_1), \dots, e(u_n)).$$

In this section, we formulate a functional transcendence result of “Ax–Schanuel”-type for the uniformisation $\pi: U \rightarrow X$. We will need this both for the proof of Theorem 2.3 and also in subsequent chapters.

Definition 2.4. An algebraic subvariety of U is a complex-analytically irreducible component of $Y \cap U$, for some algebraic subvariety (in the usual sense) $Y \subset \mathbb{C}^m \times \mathbb{C}^n$.

This definition is required since \mathbb{H}^m has no algebraic subvarieties.

Definition 2.5. A (weakly) special subvariety of U is an irreducible component of $\pi^{-1}(W)$, where W is a (weakly) special subvariety of X .

Observe that a weakly special subvariety of U is a special subvariety of U if it contains the preimage of a special point of X . Note also that the definition of a (weakly) special subvariety of U depends on the choice of Y , although U itself does not.

Our first Ax–Schanuel result is the following statement.

Theorem 2.6. *Let $W \subset X$ and $T \subset U$ be algebraic subvarieties. Suppose that $A \subset T \cap \pi^{-1}(W)$ is a complex-analytically irreducible component. Then*

$$\dim A = \dim W + \dim T - \dim X,$$

unless A is contained in a proper weakly special subvariety of U .

PROOF. When $Y = Y(1)$, this is [65, Theorem 3.2]. The general statement here may be proved by exactly the same argument. In this way, the result follows from the corresponding Ax–Schanuel statements for Y^n and \mathbb{G}_m^n . The former is encompassed by [64, Theorem 1.1] in the modular case

and by [53, Theorem 1.1] in the Shimura setting. The latter is included in Ax's theorem for the exponential function [6]. \square

From an Ax–Schanuel statement of the form of Theorem 2.6, one may deduce the following version (see [62, paragraph above 5.7]).

Theorem 2.7. *Let $U' \subset U$ be a weakly special subvariety. Let $X' = \pi(U')$. Let $W \subset X'$ and $T \subset U'$ be algebraic subvarieties. Suppose that A is a complex-analytically irreducible component of $T \cap \pi^{-1}(W)$. Then*

$$\dim A = \dim W + \dim T - \dim X',$$

unless A is contained in a proper weakly special subvariety of U' .

To state a third version of Ax–Schanuel, we need the following definitions from [39].

Definition 2.8. Fix a subvariety $W \subset X$.

- (1) A component A with respect to W is a complex-analytically irreducible component $A \subset T \cap \pi^{-1}(W)$ for some algebraic subvariety $T \subset U$.
- (2) The Zariski-defect with respect to W of a component A with respect to W is defined

$$\delta_{\text{Zar}}(A) = \dim \langle A \rangle_{\text{Zar}} - \dim A,$$

where $\langle A \rangle_{\text{Zar}}$ denotes the Zariski-closure of A .

- (3) A component A with respect to W is called Zariski-optimal with respect to W if there is no component B with respect to W such that $B \supsetneq A$ and $\delta_{\text{Zar}}(B) \leq \delta_{\text{Zar}}(A)$.

- (4) A component A with respect to W is called geodesic with respect to W if it is a complex-analytically irreducible component of $\langle A \rangle_{\text{Zar}} \cap \pi^{-1}(W)$ and $\langle A \rangle_{\text{Zar}}$ is a weakly special subvariety of U .

Theorem 2.7 is then equivalent to the following statement, as shown in [39, §5].

Theorem 2.9. *Let $W \subset X$ be a subvariety. A Zariski-optimal component with respect to W is geodesic.*

2.4. A finiteness result for strongly special subvarieties

In this section, we prove the following result, which we need for the proof of Theorem 2.3. The Ax–Schanuel results of the previous section are used in this proof.

Proposition 2.10. *Let $k \geq 1$. Define \mathcal{S}_k to be the set of $S \subset Y^k$ such that:*

- (1) S is a strongly special subvariety of Y^k ,
- (2) there exists a proper weakly special subvariety $T_S \subset \mathbb{G}_m^k$ such that the V -image of S is contained in T_S , and
- (3) there is no strongly special subvariety $S' \supsetneq S$ of Y^k such that the V -image of S' is contained in T_S .

Then the set \mathcal{S}_k is finite.

Before coming to the proof of Proposition 2.10, we need to introduce the strongly Möbius subvarieties.

Definition 2.11. Let $n \geq 1$. A strongly Möbius subvariety of \mathbb{H}^n is (the intersection with \mathbb{H}^n of) a subvariety defined by equations of the form $gz_i = z_k$ for $g \in \mathrm{GL}_2^+(\mathbb{R})$.

We define the slope g of a strongly Möbius subvariety M to be the vector of elements $g_i \in \mathrm{SL}_2(\mathbb{R})$ which appear in the defining equations of M , as explained in [39, §6.2]. A strongly Möbius subvariety is determined uniquely by its slope, so we will write M_g for the strongly Möbius subvariety of \mathbb{H}^n with slope g . The collection of all strongly Möbius subvarieties, parametrised by their slopes, is a definable set.

Note that a weakly special subvariety of \mathbb{H}^n without constant coordinates is a strongly Möbius subvariety of \mathbb{H}^n . The collection of slopes of all the weakly special subvarieties of \mathbb{H}^n without constant coordinates is a countable set.

For the remainder of this section, fix $k \geq 1$. We take the mixed Shimura variety $X = Y^k \times \mathbb{G}_m^k$. As described in Section 2.1, this comes with a uniformisation $\pi: U \rightarrow X$, which has a fundamental domain F . Here $U = \mathbb{H}^k \times \mathbb{C}^k$. We consider the subvariety $V^k \subset X$, where $V \subset Y \times \mathbb{G}_m$ is the irreducible correspondence defined over $\overline{\mathbb{Q}}$ from Section 2.2.

Definition 2.12. Let A be a component with respect to V^k . There is a smallest strongly Möbius subvariety $M \subset \mathbb{H}^k$ and a smallest linear subvariety $L \subset \mathbb{C}^k$ such that $A \subset M \times L$. Write $\langle A \rangle_M = M \times L$. The Möbius-defect of A with respect to V^k is defined

$$\delta_M(A) = \dim \langle A \rangle_M - \dim A.$$

We say that A is Möbius-optimal with respect to V^k if there is no component B with respect to V^k such that $B \supsetneq A$ and $\delta_M(B) \leq \delta_M(A)$.

Proposition 2.13. *Let A be a component with respect to V^k . Suppose that A is Möbius-optimal with respect to V^k . Then A is Zariski-optimal with respect to V^k .*

PROOF. The approach is the same as [25, Lemma 6.7]. Suppose that B is a component with respect to V^k such that $B \supset A$ and $\delta_{\text{Zar}}(B) \leq \delta_{\text{Zar}}(A)$. Without loss of generality, we may assume that B is Zariski-optimal with respect to V^k . So B is geodesic with respect to V^k , thanks to Ax–Schanuel in the form of Theorem 2.9. In particular, $\langle B \rangle_{\text{Zar}}$ is a weakly special subvariety of U , and so $\langle B \rangle_{\text{Zar}} = \langle B \rangle_M$. Hence,

$$\delta_M(B) = \delta_{\text{Zar}}(B) \leq \delta_{\text{Zar}}(A) \leq \delta_M(A),$$

since clearly $\dim \langle A \rangle_{\text{Zar}} \leq \dim \langle A \rangle_M$. But A is Möbius-optimal with respect to V^k , so we must have that $A = B$, and thus A is Zariski-optimal with respect to V^k as required. \square

Now we come to the proof of Proposition 2.10.

PROOF OF PROPOSITION 2.10. Let Mob be the definable set of all (slopes of) strongly Möbius subvarieties of \mathbb{H}^k . Let Lin be the collection of all proper linear subvarieties of \mathbb{C}^k . Note that Lin is also a definable set.

Consider the subset Mob_0 of Mob comprising those $g \in \text{Mob}$ with the property that there exists $L \in \text{Lin}$ such that $(M_g \times L) \cap \pi^{-1}(V^n) \cap F$ has a component A of dimension equal to $\dim M_g$ which is Möbius-optimal with respect to V^k and such that L is the smallest linear subvariety containing

the projection of A . All these conditions may be checked definably, see e.g. [39, Proposition 6.5] and [25, §6].

So Mob_0 is a definable set. By Proposition 2.13, we see that any such A is Zariski-optimal and $\langle A \rangle_{\text{zar}} = \langle A \rangle_M$ is weakly special. In particular, if $g \in \text{Mob}_0$, then M_g is a weakly special subvariety of \mathbb{H}^k . In particular, the set Mob_0 , must be countable, since it is a subset of the set of slopes of weakly special subvarieties, which is itself a countable set. Therefore, Mob_0 is a countable and definable set, and thus must be finite by o-minimality.

Now suppose that $S \in \mathcal{S}_k$. Then, by the definition of Mob_0 , there is some $g \in \text{Mob}_0$ such that $q(M_g) = S$. In particular, the set \mathcal{S}_k must be finite as desired. \square

2.5. Arithmetic ingredients

We collect in this section the arithmetic ingredients we need for the proof of Theorem 2.3. Fix $n \geq 1$. In this section, constants may depend implicitly on Y, V, n , and the choice of fundamental domain F_Y for the uniformisation $q: \mathbb{H} \rightarrow Y$ of Y . The dependence of a constant on any other parameters will be explicitly indicated.

Given a special point $s \in Y$, we write $\Delta(s)$ for the discriminant of s . If Y is a Shimura curve, then $\Delta(s)$ is defined to be the discriminant of the CM abelian variety associated to s by the approach described in Section 1.8. If Y is a modular curve, then $\Delta(s)$ is defined, as in Section 1.7, to be the discriminant of the quadratic order corresponding to the endomorphism ring of the CM elliptic curve associated to s .

Proposition 2.14 ([63, Proposition 5.1]). *Let $s \in Y$ be a special point and $\tau \in F_Y$ a preimage of s for the uniformisation $q: \mathbb{H} \rightarrow Y$. Then:*

- (1) $h(s) \leq c(\epsilon)|\Delta(s)|^\epsilon$ for every $\epsilon > 0$;
- (2) $H(\tau) \leq C|\Delta(s)|^c$;
- (3) $[\mathbb{Q}(s) : \mathbb{Q}] \leq c(\epsilon)|\Delta(s)|^{\frac{1}{2}+\epsilon}$ for every $\epsilon > 0$;
- (4) $[\mathbb{Q}(s) : \mathbb{Q}] \geq c(\epsilon)|\Delta(s)|^{\frac{1}{2}-\epsilon}$ for every $\epsilon > 0$.

(The constants here are all independent of s .)

For bounds on the multiplicative relations satisfied by V -images of special points in Y , we use a result of Masser [48]. For K a number field, we denote:

$$\eta = \eta(K) = \inf\{h(P) : P \in \mathbb{G}_m(K) \text{ is non-torsion}\},$$

the infimum of the logarithmic heights of the non-torsion points of $\mathbb{G}_m(K)$, and

$$\omega = \omega(K) = |\mathbb{G}_m(K)_{\text{tors}}|,$$

the cardinality of the torsion subgroup of $\mathbb{G}_m(K)$.

Theorem 2.15 ([48, Theorem \mathbb{G}_m]). *Let K be a number field. Suppose that $x_1, \dots, x_n \in \mathbb{G}_m(K)$ have logarithmic heights at most $h \geq \eta$. The set*

$$\{(m_1, \dots, m_n) \in \mathbb{Z}^n : x_1^{m_1} \cdots x_n^{m_n} = 1\}$$

forms an additive subgroup of \mathbb{Z}^n . There is a basis $\{m^{(j)} : j \in I\}$ for this group with

$$|m_i^{(j)}| \leq n^{n-1} \omega \left(\frac{h}{\eta} \right)^{n-1}$$

for every $i \in \{1, \dots, n\}$ and $j \in I$. Further, we have the following estimates on the parameters η, ω :

$$\eta \geq c^{-1} D^{-1} \left(\frac{\log L}{L} \right)^3$$

and

$$\omega \leq cD \log L,$$

where we write $D = [K : \mathbb{Q}]$ and $L = \log(D + 2)$ for convenience.

Note, in particular, that both $\frac{1}{\eta}, \omega$ can therefore be bounded above polynomially in D . We may then deduce the following.

Proposition 2.16. *There are constants C, c with the following property.*

Let $x_1, \dots, x_n \in \mathbb{G}_m$ be V -images of special points $s_1, \dots, s_n \in Y$. Write $\Delta(s_i)$ for the discriminant of s_i and set $\Delta = \Delta(s) = \max\{|\Delta(s_i)|\}$. Then there is a basis $\{m^{(j)} : j \in I\}$ for the additive group of $(m_1, \dots, m_n) \in \mathbb{Z}^n$ such that $x_1^{m_1} \cdots x_n^{m_n} = 1$ with $|m_i^{(j)}| \leq C\Delta^c$ for every $i \in \{1, \dots, n\}$ and $j \in I$.

PROOF. Let K_0 be a number field over which Y, V are defined. Suppose $x = (x_1, \dots, x_n) \in \mathbb{G}_m^n$ and $s = (s_1, \dots, s_n) \in Y^n$ are such that $(s_i, x_i) \in V$ for all i . Let $K = K_0(x_1, \dots, x_n)$ and $D = [K : \mathbb{Q}]$. Since each x_i is a V -image of s_i and V is defined over K_0 , we have that

$$[K_0(x_i) : K_0] \leq c_1 [\mathbb{Q}(s_i) : \mathbb{Q}]$$

for some constant C . Thus, $D \leq c_2 \Delta^{c_3}$ by Proposition 2.14 (3).

Let $h = \max\{h(x_1), \dots, h(x_n)\}$. Suppose for now that some x_i is non-torsion. Then $h \geq \eta$, and so by Theorem 2.15 there is a basis $m^{(j)}$ for the

multiplicative relations satisfied by the x_i , with every $m_i^{(j)}$ having

$$(2.1) \quad |m_i^{(j)}| \leq n^{n-1} \omega(K) \left(\frac{h}{\eta(K)} \right)^{n-1},$$

where $\omega(K), \eta(K)$ are as defined above.

Note that $H(x_i) \leq c_4 H(s_i)^{c_5}$ for every i , and so $h \leq c_6 \Delta^{c_7}$ by Proposition 2.14 (1). If we combine these facts with the bounds for $\omega(K), \eta(K)$ in terms of D , then we see that the right hand side of equation (2.1) is bounded polynomially in terms of Δ as required.

If every x_i is torsion, then the degrees m_i of the x_i form a generating set for the multiplicative relations satisfied by the x_i , and so the result follows from the fact that D is polynomially bounded in terms of Δ . The proof is thus complete. \square

2.6. Finiteness of exemplary components

Now we give the proof of Theorem 2.3. Recall that Y is a modular or Shimura curve and $V \subset Y \times \mathbb{G}_m$ is an irreducible correspondence defined over $\overline{\mathbb{Q}}$. Fix $n \geq 1$. Let K be a number field over which V, Y are defined.

PROOF OF THEOREM 2.3. Recall that a special subvariety $S \subset Y^n$ may be written as $S = S_1 \times \{S_2\}$ for $S_1 \subset Y^{I_1}$ a strongly special subvariety and $S_2 \in Y^{I_2}$ a special point. Here $I_1, I_2 \subset \{1, \dots, n\}$ are disjoint, possibly empty, sets of coordinates such that $I_1 \cup I_2 = \{1, \dots, n\}$.

Fix a strongly special subvariety $S_1 \subset Y^{I_1}$. We may and do assume that S_1 is defined over K . Let $W_1 = \{(x, t) \in V^{I_1} : x \in S_1\}$. Then $\pi_2(W_1)$ is contained in a proper weakly special subvariety of $\mathbb{G}_m^{I_2}$ if and only if some

equations

$$\prod_{i \in I_1} \xi_i^{a_i} = p$$

hold for all $\xi \in \pi_2(W_1)$, where $p \in \mathbb{G}_m$ and the $a_i \in \mathbb{Z}$ are not all zero. Let $p_1, \dots, p_k \in \mathbb{G}_m$ be the points corresponding to a generating set of such relations. The span of the p_i is $\text{Gal}(\overline{\mathbb{Q}}/K)$ invariant, so we may assume that the p_i are defined over K .

Fix a preimage $\nu = (\nu_1, \dots, \nu_k) \in F_e^k$ of (p_1, \dots, p_k) . Let

$$Z = \{(z, w, \alpha, \beta) \in F_Y^{I_2} \times F_e^{I_2} \times \mathbb{R}^{(k+|I_2|) \times (k+|I_2|)} \times \mathbb{R}^{(k+|I_2|)} : \\ \pi(z, w) \in V^{I_2}, \quad \alpha \cdot (vw)^t = \beta^t\}.$$

Here A^t denotes the transpose of a matrix A . Note that Z is definable.

Suppose that $S_2 \in Y^{I_2}$ is a special point. Write Δ for the complexity $\Delta(S_2)$ of S_2 , which is defined as in Proposition 2.16. We will show that if Δ is larger than some constant (depending on S_1), then there is no exemplary component $W \subset V^n$ with $\pi_1(W) = S_1 \times \{S_2\}$. In what follows, constants c_i are positive and depend possibly on Y, V, n, K, S_1 , but are independent of S_2 .

Let $\eta \in \mathbb{G}_m^{I_2}$ be such that $(S_2, \eta) \in V^{I_2}$. Then (S_2, η) has a preimage $(z, w) \in F_Y^{I_2} \times F_e^{I_2}$ such that the z coordinates are algebraic, of bounded degree (cf. Proposition 1.33), and have height bounded by $c_1 \Delta^{c_2}$ by Proposition 2.14 (2).

Any equation

$$\prod_{i \in I_1} \xi_i^{a_i} \prod_{j \in I_2} \eta_j^{b_j} = 1,$$

where $a_i, b_j \in \mathbb{Z}$, which holds for all $\xi \in \pi_2(W_1)$ implies that

$$\prod_{i \in I_1} \xi_i^{a_i}$$

is constant on $\pi_2(W_1)$. It is thus equivalent to an equation of the form

$$\prod_{i=1}^k p_i^{a'_i} \prod_{j \in I_2} \eta_j^{b'_j} = 1,$$

where $a'_i, b'_j \in \mathbb{Z}$. Conversely, any equation of the second form is equivalent to an equation of the first form.

Consider then the system of equations

$$\prod_{i=1}^k p_i^{a'_i} \prod_{j \in I_2} \eta_j^{b'_j} = 1,$$

where $a'_i, b'_j \in \mathbb{Z}$, which is equivalent to the system of all the equations

$$\prod_{i \in I_1} \xi_i^{a_i} \prod_{j \in I_2} \eta_j^{b_j} = 1,$$

where $a_i, b_j \in \mathbb{Z}$, which hold for all $\xi \in \pi_2(W_1)$. Let $H \subset \mathbb{Z}^k \times \mathbb{Z}^{I_2}$ be the corresponding set of $(a', b') \in \mathbb{Z}^k \times \mathbb{Z}^{I_2}$. Then H is an additive subgroup of $\mathbb{Z}^k \times \mathbb{Z}^{I_2}$.

Proposition 2.16 then gives a basis \mathcal{B} for the group H such that the $\leq k + |I_2|$ elements of \mathcal{B} are vectors in $\mathbb{Z}^{k+|I_2|}$ with all their entries $\leq c_3 \Delta^{c_4}$ in absolute value. Let $\alpha \in \mathbb{R}^{(k+|I_2|) \times (k+|I_2|)}$ be the matrix whose rows comprise the elements of \mathcal{B} , followed by some zero rows as necessary. Then α has integral entries and height (viewed as the vector of its entries) bounded by $c_5 \Delta^{c_6}$, since the height of a non-zero integer is equal to its absolute value.

Let $\beta \in \mathbb{R}^{(k+|I_2|)}$ be defined by $\beta^t = \alpha \cdot (\nu w)^t$. Observe that $\beta \in \mathbb{Z}^{(k+|I_2|)}$. Since all the ν_i, w_j have real part in $[0, 1]$, the absolute value (and hence height) of an entry of β is bounded by $c_7 \Delta^{c_8}$. The special point S_2 thus gives rise to the point $(z, w, \alpha, \beta) \in Z$, which is algebraic of bounded degree in its z coordinates, integral in its α and β coordinates, and has height in its z, α, β coordinates $\leq c_9 \Delta^{c_{10}}$.

Now consider a K -conjugate (S'_2, η') of (S_2, η) . Since $Y, V, S_1, p_1, \dots, p_k$ are all defined over K , the conjugate (S'_2, η') also gives rise (as above) to a point $(z', w', \alpha, \beta') \in Z$, where the z' coordinates are algebraic of bounded degree, the entries of α, β' are integral, and the height of the z', α, β' coordinates is $\leq c_9 \Delta^{c_{10}}$. Distinct conjugates must give rise to distinct points of Z , but note that the α coordinate is the same for all the resulting points. By Proposition 2.14, there are at least $c_{11} \Delta^{c_{12}}$ distinct K -conjugates of (S_2, η) . We are thus in a position to apply the o-minimal Counting Theorem, in the form of [39, Corollary 7.2] (Theorem A.21), to the set Z .

If Δ is suitably large, then we thereby obtain a continuous definable function $\gamma: [0, 1] \rightarrow Z_\alpha$, where Z_α denotes the fibre of Z over the point $\alpha \in \mathbb{R}^{(k+|I_2|) \times (k+|I_2|)}$, with the following properties:

- (1) The composition $\pi_a \circ \gamma: [0, 1] \rightarrow F_Y^{I_2} \times \mathbb{R}^{k+|I_2|}$ is semialgebraic.
- (2) The composition $\pi_b \circ \gamma: [0, 1] \rightarrow F_e^{I_2}$ is non-constant.
- (3) There is some conjugate (S'_2, η') of (S_2, η) such that $\beta(0)$ is the point of Z_α arising from (S'_2, η') .
- (4) The restriction of γ to $(0, 1)$ is real analytic.

Here π_a, π_b are the projection maps $F_Y^{I_2} \times F_e^{I_2} \times \mathbb{R}^{k+|I_2|} \rightarrow F_Y^{I_2} \times \mathbb{R}^{k+|I_2|}$ and $F_Y^{I_2} \times F_e^{I_2} \times \mathbb{R}^{k+|I_2|} \rightarrow F_e^{I_2}$ respectively. The property we state here as (3) is established in the proof of [39, Corollary 7.2], although it is not stated there in this form. Note that (4) holds since $\mathbb{R}_{\text{an,exp}}$ admits analytic cell decomposition [30]. By (2) and the definition of Z , we see that the projection of $\gamma([0, 1])$ to $F_Y^{I_2}$ is positive-dimensional.

We now show that the projection of $\gamma([0, 1])$ to $\mathbb{R}^{(k+|I_2|)}$ is constant. Fix $r \in \{1, \dots, k + |I_2|\}$. If the r th row of α is a zero row, then $\gamma([0, 1])$ must just project to 0 in the r th coordinate of $\mathbb{R}^{(k+|I_2|)}$. So we assume that $\alpha_r = (\alpha_{r,1}, \dots, \alpha_{r,k+|I_2|})$ is not a zero row. Let $\Phi: F_Y^{I_2} \times \mathbb{R}^{k+|I_2|} \rightarrow \mathbb{H}^{I_2} \times \mathbb{C}$ be given by $(z, \beta) \mapsto (z, \beta_r)$. Then $(\Phi \circ \pi_a \circ \gamma)([0, 1])$ is a positive-dimensional semialgebraic set which is contained in the set

$$A = \{(u, t) \in \mathbb{H}^{I_2} \times \mathbb{C} : \exists \lambda \in \mathbb{G}_m^{I_2} \quad (q(z), \lambda) \in V^{I_2}, \\ \prod_{i=1}^k p_i^{\alpha_{r,i}} \prod_{j=1}^{|I_2|} \lambda_j^{\alpha_{r,k+j}} = e(t)\}.$$

The set A is in fact a complex analytic set; this follows from the implicit function theorem. We may thus find an open neighbourhood $\Omega \subset \mathbb{H}^{I_2} \times \mathbb{C}$ of $\Phi(\pi_a(\gamma(0)))$ and a set P such that:

- (1) $\Omega \cap \gamma([0, 1]) \subset P \subset A$,
- (2) P may be written as a finite union of irreducible Nash subsets of Ω which each contain $\Phi(\pi_a(\gamma(0)))$.

This uses [1, Proposition 1] as well as the characterisation of Nash sets given in [1, p. 989].

If every complex-analytically irreducible component of P had constant projection to its \mathbb{H}^{I_2} coordinates, then $\gamma([0, 1])$ would have constant projection to its \mathbb{H}^{I_2} coordinates by real analytic continuation. Since $\gamma([0, 1])$ does not have constant projection to its \mathbb{H}^{I_2} coordinates, there must exist some complex-analytically irreducible component of P which has non-constant projection to its \mathbb{H}^{I_2} coordinates. Every complex-analytically irreducible component of P contains $\Phi(\pi_a(\gamma(0)))$. Thus, by complex analytic continuation, there exists an algebraic subvariety $T \subset \mathbb{C}^{I_2} \times \mathbb{C}$ and a complex-analytically irreducible component $D \subset (\mathbb{H}^{I_2} \times \mathbb{C}) \cap T$ such that $\Phi(\pi_a(\gamma(0))) \in D \subset A$ and D has non-constant projection to its \mathbb{H}^{I_2} coordinates. We now apply Ax–Schanuel. There must then exist weakly special subvarieties $M_1 \subset \mathbb{H}^{I_2}$ and $M_2 \subset \mathbb{C}$ such that $D \subset M_1 \times M_2 \subset A$. Since the projection $A \rightarrow \mathbb{H}^{I_2}$ has discrete fibres, the weakly special subvariety M_2 must just be a point. Since $\Phi(\pi_a(\gamma(0))) \in D$, we see that M_2 must in fact be equal to the projection of $\Phi(\pi_a(\gamma(0)))$ and so, in particular, be an integer, which we denote b_r . Repeating for each r in turn, we thus show that the projection of $\gamma([0, 1])$ to $\mathbb{R}^{(k+|I_2|)}$ is constant and integer-valued.

Let L be the linear subvariety of \mathbb{C}^{k+I_2} defined by the equations

$$\sum_{i=1}^k \alpha_{r,i} x_i + \sum_{j \in I_2} \alpha_{r,j} x_j = b_r$$

for $r = 1, \dots, k + |I_2|$. Let $L_\nu \subset \mathbb{C}^{I_2}$ be the fibre of L over (ν_1, \dots, ν_k) . By the above argument, there is a positive-dimensional weakly special subvariety of $M \subset \mathbb{H}^{I_2}$ such that $(M \times L_\nu) \cap \pi^{-1}(V)$ has a component U_0 which projects onto M and contains the preimage of a conjugate (S'_2, η') of (S_2, η) .

Let $S^* = q(M)$. Then S^* is a positive-dimensional weakly special subvariety of Y^{I_2} . In fact, S^* is a special subvariety since it contains a special point. An appropriate K -conjugate S^{**} of S^* is then a positive-dimensional special subvariety of Y^{I_2} which contains S_2 itself. By construction, the V -image of $S_1 \times S^{**}$ is contained in the same special subvariety of \mathbb{G}_m^n as the V -image of $S_1 \times \{S_2\}$. In particular, there can be no exemplary component W of V^n with $\pi_1(W) = S_1 \times \{S_2\}$.

Thus, there is some constant $C(S_1)$ such that, if $\Delta(S_2) \geq C(S_1)$, then there is no exemplary component $W \subset V^n$ with $\pi_1(W) = S_1 \times \{S_2\}$. So there are only finitely many exemplary components W such that $\pi_1(W) = S_1 \times \{S_2\}$ for some special point $S_2 \in Y^{I_2}$. If there exists an exemplary component W such that $\pi_1(W) = S_1 \times \{S_2\}$, then either $S_1 = Y^{I_1}$ or the V -image of S_1 is contained in a proper weakly special subvariety of $\mathbb{G}_m^{I_1}$. By Proposition 2.10, applied with $k = 1, \dots, n$, there are thus only finitely many possibilities for S_1 too. Consequently, there are only finitely many exemplary components of V^n , as desired. \square

2.6.1. The one-dimensional case. We finish this chapter by showing that, in the case $n = 1$, we may drop the condition that V is defined over $\overline{\mathbb{Q}}$.

Proposition 2.17. *Let $V \subset Y \times \mathbb{G}_m$ be a correspondence (not necessarily defined over $\overline{\mathbb{Q}}$). Then there are only finitely many exemplary components of V .*

PROOF. The only special subvarieties of \mathbb{G}_m are \mathbb{G}_m itself and those of the form $\{\zeta\}$, where ζ is a root of unity. Since V dominates \mathbb{G}_m , the only exemplary components of V are therefore V itself and all the special

points $(s, \zeta) \in V$. It therefore suffices to show the finiteness of such special points.

Since special points of Y are algebraic, if V contains infinitely many special points (s, ζ) , then we may assume that V is defined over $\overline{\mathbb{Q}}$. But then the finiteness of exemplary components (and hence also of special points) of V follows already from Theorem 2.3. \square

Remark 2.18. Since V does not contain any positive-dimensional special subvarieties, what we have really established in Proposition 2.17 is the André–Oort statement (see Section 5.1) for V .

Pila and Tsimerman [63] are able to extend their version of Theorem 2.3 to correspondences defined over \mathbb{C} for all $n \geq 1$. They do this via a specialisation argument. A similar approach should work in our setting, but we have not yet been able to carry it out.

CHAPTER 3

Multiplicative independence of modular functions

3.1. Special points of modular functions

Let $R \in \mathbb{C}(t)$ be a rational function (of one variable). Let V be the intersection with $Y(1) \times \mathbb{G}_m$ of the graph of R (viewed as a subset of $\mathbb{P}^1 \times \mathbb{P}^1$). The results of Chapter 2 imply that if R is defined over $\overline{\mathbb{Q}}$, then, for each $n \geq 1$, multiplicatively dependent n -tuples of R -images of special points of $Y(1)$ belong to one of finitely many exemplary components of $V^n \subset Y(1)^n \times \mathbb{G}_m^n$.

The results of Chapter 2 do not though imply that, for each n , there are only finitely many such n -tuples. In this chapter, we prove, for every $n \geq 1$, a finiteness result for these n -tuples, which holds for all suitably generic rational functions R . Our starting point is the following result of Pila and Tsimerman [65].

Theorem 3.1. [65, Theorem 1.2] *Let $n \geq 1$. There exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct singular moduli such that $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent.*

Observe that the multiplicative independence of proper subsets and the distinctness of the σ_i are required to avoid trivialities. Pila and Tsimerman proved Theorem 3.1 by an o-minimal counting argument; in particular, the result is ineffective. The critical new ingredient in their proof

was a functional multiplicative independence result for pairwise distinct $\mathrm{GL}_2^+(\mathbb{Q})$ -translates of the j -function. Here and throughout $\mathrm{GL}_2^+(\mathbb{Q})$ and its subgroups act on \mathbb{H} by Möbius transformations.

Functions $f_1, \dots, f_n: \mathbb{H} \rightarrow \mathbb{C}$ are called multiplicatively dependent modulo constants if some relation $\prod f_i^{a_i} = c$ holds for $a_i \in \mathbb{Z}$, not all zero, and $c \in \mathbb{C}$; if no such relation holds, then f_1, \dots, f_n are called multiplicatively independent modulo constants. The functional independence result of Pila and Tsimerman was the following.

Theorem 3.2. [65, Theorem 1.3] *Let $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$. If the functions $j(g_1z), \dots, j(g_nz)$ are pairwise distinct, then $j(g_1z), \dots, j(g_nz)$ are multiplicatively independent modulo constants.*

The proof of Theorem 3.2 in [65] is via an elaborate tree argument. In particular, the method there does not readily generalise to other modular functions. In Section 3.2, we provide a new proof of Theorem 3.2. This uses just elementary properties of j . Notably, the proof generalises in a straightforward way to a wide class of modular functions, and so we are able to prove Theorem 3.5, a generalisation of Theorem 3.2.

Recall that the map j restricts to a bijection $j|_{F_j}: F_j \rightarrow \mathbb{C}$. Therefore, any non-zero modular function $f: \mathbb{H} \rightarrow \mathbb{C}$ has only finitely many zeros and poles in F_j , and further if f is also non-constant, then f has at least one zero or pole in F_j . We may therefore make the following definition.

Definition 3.3. Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function. Enumerate the zeros and poles of f contained in F_j as w_1, \dots, w_r , where $\mathrm{Im}(w_i) \leq \mathrm{Im}(w_{i+1})$. We say that f satisfies the divisor condition if $w_r + s$ is neither a zero nor a pole of f for all $0 < s < 1$.

Remark 3.4. If $\text{Im}(w_r) \geq 1$, then the divisor condition holds if $\text{Im}(w_r) > \text{Im}(w_{r-1})$.

Our generalisation of Theorem 3.2 is the following result.

Theorem 3.5. *Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function satisfying the divisor condition. Let $g_1, \dots, g_n \in \text{GL}_2^+(\mathbb{Q})$. If the functions $f(g_1z), \dots, f(g_nz)$ are pairwise distinct, then they are multiplicatively independent modulo constants.*

The divisor condition is satisfied generically by modular functions (but not by all modular functions). In Section 3.3, we show that the divisor condition is in particular satisfied by all of the (infinitely many) non-constant functions belonging to a certain multiplicative group of modular functions with product formulae which arise as Borcherds lifts of some weakly holomorphic modular forms. This gives us a natural set of examples of modular functions satisfying the divisor condition.

Definition 3.6. Let f be a modular function. An f -special point is a complex number σ such that $\sigma = f(\tau)$ for some $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$.

The j -special points are thus precisely the singular moduli (i.e. the special points of $Y(1)$). If f is a modular function, then there exists a rational function $R \in \mathbb{C}(t)$ such that $f(z) = R(j(z))$. The f -special points are then exactly the images under R of the singular moduli. For clarity, we will use the term j -special point in place of singular moduli for the remainder of this chapter.

If f satisfies the divisor condition, then we are able to establish a finiteness result on multiplicatively dependent tuples of f -special points, analogous to Theorem 3.1. The proof is via a modification of Pila and Tsimerman's o-minimal counting argument and is ineffective.

Theorem 3.7. *Suppose f is a non-constant modular function satisfying the divisor condition. Let $n \geq 1$. There exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent.*

An n -tuple $(\sigma_1, \dots, \sigma_n)$ of f -special points is multiplicatively dependent if some product $\prod \sigma_i^{a_i}$, with the $a_i \in \mathbb{Z}$ not all zero, lies in the trivial subgroup $\{1\} \leq \mathbb{G}_m$. A natural extension of Theorem 3.7 is then to consider, given a fixed subgroup $\Gamma \leq \mathbb{G}_m$, those n -tuples $(\sigma_1, \dots, \sigma_n)$ of f -special points such that, for some $a_i \in \mathbb{Z}$ not all zero, one has $\prod \sigma_i^{a_i} \in \Gamma$.

When $\Gamma \leq \mathbb{G}_m$ is of finite rank and $f \in \overline{\mathbb{Q}}(j)$, we are able to extend Theorem 3.7 to this setting. Recall that a subgroup $\Gamma \leq \mathbb{G}_m$ is said to be of finite rank if there exists $\Gamma_0 \subset \Gamma$ such that Γ_0 is a finitely generated subgroup of \mathbb{G}_m and, for every $\gamma \in \Gamma$, there exists $m \geq 1$ such that $\gamma^m \in \Gamma_0$. We thereby establish the following result. Here a set $\{x_1, \dots, x_n\}$ is called Γ -dependent, for $\Gamma \leq \mathbb{G}_m$, if there exist $a_1, \dots, a_n \in \mathbb{Z}$, not all zero, such that $\prod x_i^{a_i} \in \Gamma$.

Theorem 3.8. *Suppose f is a non-constant modular function satisfying the divisor condition and also $f \in \overline{\mathbb{Q}}(j)$. Let $\Gamma \leq \mathbb{G}_m$ be of finite rank and $n \geq 1$. Then there exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of*

distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent.

The plan of the rest of this chapter is as follows. In Section 3.2, we give a new proof of Theorem 3.2 and then generalise it to prove Theorem 3.5. Section 3.3 establishes a natural class of modular functions, arising via Borcherds lifts, to which this theorem applies. The proof of Theorem 3.7 is contained in Section 3.4. In Section 3.5, we explore the link between the results of this chapter and those of Chapter 2. Finally, the proof of Theorem 3.8 then takes place in Section 3.6. Many of the results of this chapter first appeared in the author's article [32].

3.2. Functional multiplicative independence

In this section, we give a new proof of Theorem 3.2, and then show how this method may be extended to prove the more general Theorem 3.5.

PROOF OF THEOREM 3.2. Let $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$, and suppose that the functions $j(g_1 z), \dots, j(g_n z)$ are pairwise distinct. Then the cosets

$$[g_1], \dots, [g_n] \in \mathrm{PSL}_2(\mathbb{Z}) \backslash \mathrm{PGL}_2^+(\mathbb{Q})$$

are pairwise distinct. For $g \in \mathrm{GL}_2^+(\mathbb{Q})$, we may, as in the proof of [61, Proposition 7.1], write $g = \gamma h$, where $\gamma \in \mathrm{SL}_2(\mathbb{Z})$ and $hz = rz + s$ for $r, s \in \mathbb{Q}$ with $0 < r$ and $0 \leq s < 1$. The cosets $[g_i] \in \mathrm{PSL}_2(\mathbb{Z}) \backslash \mathrm{PGL}_2^+(\mathbb{Q})$ are pairwise distinct if and only if the corresponding linear functionals $r_i z + s_i$ are pairwise distinct. Different g_i may have associated the same r_i , so reindex them as $g_{i,k}$, associated with the functional $r_i z + s_{i,k}$, where $r_1 < r_2 < \dots < r_l$ and $s_{i,k} < s_{i,k'}$ for $k < k'$.

To prove Theorem 3.2, it is enough to find $z \in \mathbb{H}$ such that $j(r_i z + s_{i,k}) = 0$ if and only if $(i, k) = (1, 1)$. Recall that $j(\zeta_6) = 0$, where $\zeta_6 = \exp(\pi i/3)$. Therefore, setting

$$z = \frac{1}{r_1} \left(\frac{1}{2} - s_{1,1} \right) + \frac{1}{r_1} \frac{\sqrt{3}}{2} i \in \mathbb{H},$$

so that $r_1 z + s_{1,1} = \zeta_6$, gives that $j(r_1 z + s_{1,1}) = 0$.

It remains to show that $j(r_i z + s_{i,k}) \neq 0$ for $(i, k) \neq (1, 1)$. To do this we use two elementary facts about j . First, that for w with $1/2 < \operatorname{Re}(w) < 3/2$, if $j(w) = 0$, then $\operatorname{Im}(w) < \sqrt{3}/2$. Second, that $j(w) \neq 0$ whenever $\operatorname{Im}(w) > \sqrt{3}/2$. Either of these is clear by considering the tessellation of \mathbb{H} by translates of the fundamental domain F_j for the action of $\operatorname{SL}_2(\mathbb{Z})$, since the only zero of j in F_j is at ζ_6 .

For $k > 1$, note that

$$\operatorname{Re}(r_1 z + s_{1,k}) = \operatorname{Re}(r_1 z + s_{1,1}) + \operatorname{Re}(s_{1,k} - s_{1,1}) = \frac{1}{2} + (s_{1,k} - s_{1,1}).$$

Since $0 \leq s_{1,1} < s_{1,k} < 1$, we have that $1/2 < \operatorname{Re}(r_1 z + s_{1,k}) < 3/2$. Therefore, $j(r_1 z + s_{1,k}) \neq 0$ by the first fact, since $\operatorname{Im}(r_1 z + s_{1,k}) = \sqrt{3}/2$. Now for $i > 1$,

$$\operatorname{Im}(r_i z + s_{i,k}) = r_i \operatorname{Im}(z) = r_i \frac{1}{r_1} \frac{\sqrt{3}}{2} > \frac{\sqrt{3}}{2}$$

since $r_i > r_1$ for $i > 1$. Therefore, by the second fact, $j(r_i z + s_{i,k}) \neq 0$ for $i > 1$. Hence, $j(r_i z + s_{i,k}) \neq 0$ for $(i, k) \neq (1, 1)$. \square

Now we show how this method may be generalised to prove Theorem 3.5.

PROOF OF THEOREM 3.5. Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function. So f is a rational function of j . Let $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$, and suppose that $f(g_1z), \dots, f(g_nz)$ are pairwise distinct. In particular, the associated cosets $[g_i] \in \mathrm{PSL}_2(\mathbb{Z}) \backslash \mathrm{PGL}_2^+(\mathbb{Q})$ are pairwise distinct, and so, as above, we may rewrite the $f(g_i z)$ as $f(r_i z + s_{i,k})$ for $r_i, s_{i,k} \in \mathbb{Q}$ with $r_i > 0$ and $0 \leq s_{i,k} < 1$. These linear functionals $r_i z + s_{i,k}$ are again pairwise distinct and we may assume $r_1 < r_2 < \dots < r_l$ and $s_{i,k} < s_{i,k'}$ for $k < k'$.

The function f is meromorphic on \mathbb{H} . To prove the multiplicative independence modulo constants of the functions $f(r_i z + s_{i,k})$, it will be enough to find $z \in \mathbb{H}$ such that $r_i z + s_{i,k}$ is either a zero or a pole of f if and only if $(i, k) = (1, 1)$. Since f is a rational function of j , and $j|_{F_j}: F_j \rightarrow \mathbb{C}$ is bijective, the function f has only finitely many zeros and poles in F_j . Further, f has at least one zero or pole in F_j since f is non-constant.

Enumerate the zeros and poles of f in F_j as w_1, \dots, w_r , where $\mathrm{Im}(w_i) \leq \mathrm{Im}(w_{i+1})$. We may then proceed for f as we did in the above proof for j , replacing ζ_6 by w_r , provided that

- (1) $w_r + s$ is neither a zero nor a pole of f for all $0 < s < 1$, and
- (2) f has no zero or pole with imaginary part $> \mathrm{Im}(w_r)$.

The first of these is just the divisor condition. For the second, note that if f has a zero or pole in \mathbb{H} with imaginary part $> \mathrm{Im}(w_r)$, then f must have a zero or pole in F_j with imaginary part $> \mathrm{Im}(w_r)$, as may be seen by considering the tessellation of \mathbb{H} by $\mathrm{SL}_2(\mathbb{Z})$ -translates of F_j . This cannot happen by the definition of w_r , and so we are done. \square

Remark 3.9. Suppose f is a non-constant modular function which does not satisfy the divisor condition. Let $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$ be such that the functions $f(g_1z), \dots, f(g_nz)$ are pairwise distinct. For such g_i suitably generic, there will still exist some $z \in \mathbb{H}$ such that $g_i z$ is either a zero or a pole of f for exactly one i . Thus the translates $f(g_1z), \dots, f(g_nz)$ would still be multiplicatively independent modulo constants. The divisor condition is sufficient for this to be true for all possible choices of the g_i , and hence sufficient to establish Theorem 3.5, but is not obviously necessary. Indeed, there does not appear to be an obvious obstruction to a corresponding functional independence result holding for an arbitrary non-constant modular function. It seems likely that there is a weaker condition which would still suffice to prove Theorem 3.5.

3.3. Borcherds products for modular functions

In this section, we show that the divisor condition is satisfied by the non-constant elements of a natural class of modular functions, which arises as the set of Borcherds lifts of certain weakly holomorphic modular forms. We thereby establish multiplicative independence for non-constant functions belonging to this set. We introduce the following notation, after [16]. Here $q = \exp(2\pi iz)$ is the usual nome.

Let A be the set of weakly holomorphic modular forms f of weight $1/2$ on $\Gamma_0(4)$ which have a Fourier series of the form $f(z) = \sum c(n)q^n$, where $c(n) \in \mathbb{Z}$ are such that $c(n) = 0$ unless $n \equiv 0, 1 \pmod{4}$. (Weakly holomorphic means that we allow f possibly to have poles at the cusps.) It is clear that the elements of the set A form an additive group.

Let B be the set of integer weight meromorphic modular forms for some character of $\mathrm{SL}_2(\mathbb{Z})$, all of whose zeros and poles are located at either cusps or imaginary quadratic numbers, and that have Fourier expansions with integer coefficients and leading coefficient 1. The elements of B form a multiplicative group.

Denote by $H(n)$ the Hurwitz class number of discriminant $-n$ when $n > 0$ and set $H(0) = -1/12$. Then let the function \tilde{H} be defined by

$$\tilde{H}(z) = \sum_{n \geq 0} H(n)q^n = -\frac{1}{12} + \frac{q^3}{3} + \frac{q^4}{2} + q^7 + \dots$$

Definition 3.10. The discriminant Δ of an imaginary quadratic number τ is defined as $\Delta = b^2 - 4ac$, where $a, b, c \in \mathbb{Z}$ satisfy $a\tau^2 + b\tau + c = 0$ and $\gcd(a, b, c) = 1$. Equivalently, the discriminant of an imaginary quadratic number τ is equal to the discriminant (in the sense of Definition 1.26) of the j -special point $j(\tau)$.

Borcherds established the following isomorphism.

Theorem 3.11. [17, Theorem 14.1] *Let $\Psi: A \rightarrow B$ be given by*

$$f(z) = \sum a(n)q^n \mapsto \Psi(f(z)) = q^{-h} \prod_{n>0} (1 - q^n)^{a(n^2)},$$

where h is the constant coefficient in the Fourier series of $f(z)\tilde{H}(z)$. Then Ψ is an isomorphism from the additive group A to the multiplicative group B . Further, the weight of the meromorphic modular form $\Psi(f)$ is equal to $a(0)$, and the multiplicity of the zero of $\Psi(f)$ at an imaginary quadratic

number $\tau \in \mathbb{H}$ of discriminant Δ is equal to

$$\sum_{n>0} a(n^2\Delta).$$

A modular function is a weight 0 meromorphic modular form for the trivial character of $\mathrm{SL}_2(\mathbb{Z})$. The modular functions in B clearly form a subgroup of the group B , which we denote B_0 . Applying Borcherds's isomorphism, we have that $B_0 \subset \Psi(A_0)$, where A_0 is the subgroup of A comprising those functions $f \in A$ which have a Fourier expansion $f(z) = \sum a_f(n)q^n$ with $a_f(0) = 0$.

In this section, we prove the following theorem.

Theorem 3.12. *Let $f \in B_0$ be non-constant and $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$. Suppose that the functions $f(g_1z), \dots, f(g_nz)$ are pairwise distinct. Then the functions $f(g_1z), \dots, f(g_nz)$ are multiplicatively independent modulo constants.*

This theorem will follow from Theorem 3.5 once we show that non-constant functions $f \in B_0$ satisfy the divisor condition. To do this, we consider a basis for the (free abelian) group A . The group A has a basis

$$\{f_d(z) : d \geq 0 \text{ and } d \equiv 0, 3 \pmod{4}\},$$

where f_d is the unique element of A with a Fourier expansion of the form

$$f_d(z) = q^{-d} + \sum_{D>0} A(d, D)q^D.$$

See, for example, [54, Chapter 4], where the first few such functions are listed. The isomorphism of Theorem 3.11 then tells us that the functions

$\Psi(f_d)$ form a basis for the group B . Recall that the weight of an element $\Psi(f) \in B$ is equal to the constant coefficient $a_f(0)$ in the Fourier expansion of $f = \sum a_f(n)q^n \in A$. The subgroup B_0 of modular functions in B is therefore contained in $\Psi(A_0)$, which has a basis given by

$$\{\Psi(f_d) : d > 0 \text{ and } d \equiv 0, 3 \pmod{4}\}.$$

Lemma 3.13. *Let $d > 0$ satisfy $d \equiv 0, 3 \pmod{4}$. Then there exists a simple zero $\tau^* \in F_j$ of $\Psi(f_d)$ with the property that, for every $\tau \in F_j \setminus \{\tau^*\}$ which is either a zero or a pole of $\Psi(f_d)$, one has that $\text{Im}(\tau) < \text{Im}(\tau^*)$. Let $\alpha_d = \text{Im}(\tau^*)$. The resulting sequence $\{\alpha_d : d > 0, d \equiv 0, 3 \pmod{4}\}$ is strictly increasing in d .*

PROOF. Write $A(d, n)$ for the Fourier coefficients of f_d , so that

$$f_d(z) = \sum_{n \in \mathbb{Z}} A(d, n)q^n.$$

Then, in particular, $A(d, -d) = 1$ and $A(d, n) = 0$ for all $n < 0, n \neq -d$. By Theorem 3.11 all the zeros and poles of $\Psi(f_d)$ in F_j are located at imaginary quadratic numbers in F_j . Further, the multiplicity of the zero or pole of $\Psi(f_d)$ at an imaginary quadratic number $\tau \in F_j$ of discriminant Δ is equal to

$$\sum_{n > 0} A(d, n^2 \Delta).$$

In particular, this is equal to zero if $\Delta < -d$ and equal to 1 if $\Delta = -d$. Hence, $\Psi(f_d)$ has a simple zero at each imaginary quadratic number in F_j of discriminant $-d$. Further, all zeros and poles of $\Psi(f_d)$ in F_j are contained

in the set

$$\{\tau \in F_j : \tau \text{ is imaginary quadratic of discriminant } \Delta \text{ with } |\Delta| \leq d\}.$$

Recall, from Section 1.7, that the quadratic imaginary numbers in F_j of a given discriminant Δ can be explicitly described. Each corresponds to the unique (in \mathbb{H}) solution to an equation $az^2 + bz + c = 0$, where $a, b, c \in \mathbb{Z}$ satisfy $\Delta = b^2 - 4ac$, $\gcd(a, b, c) = 1$, and either $-a \leq b < a < c$ or $0 \leq b < a = c$. The imaginary quadratic τ corresponding to such a triple (a, b, c) is

$$\tau = \frac{-b + \sqrt{\Delta}}{2a},$$

which has $\text{Im}(\tau) = |\Delta|^{1/2}/2a$.

For each discriminant $\Delta < 0$, there exists a unique such triple (a, b, c) with $a = 1$. See [10, Proposition 2.6]. Write τ_Δ for the element of F_j corresponding to this triple. Then τ_Δ is a simple zero of $\Psi(f_{-\Delta})$. Now let $\tau \in F_j \setminus \{\tau_\Delta\}$ be either a zero or a pole of $f_{-\Delta}$. Then τ is an imaginary quadratic number of discriminant Δ' , where $|\Delta'| \leq |\Delta|$. Then clearly $\text{Im}(\tau_\Delta) > \text{Im}(\tau)$. So for $d > 0$ satisfying $d \equiv 0, 3 \pmod{4}$, let $\Delta = -d$ and take $\tau^* = \tau_\Delta$. This proves the first part of the lemma, and the second part then follows immediately. \square

Now we come to the proof of Theorem 3.12.

PROOF OF THEOREM 3.12. Let $f \in B_0$ be non-constant. By Theorem 3.5, it suffices to show that f satisfies the divisor condition. Since f is not the identity element of B_0 , there are $0 < d_1 < \dots < d_k$, satisfying

$d_i \equiv 0, 3 \pmod{4}$, and $a_1, \dots, a_k \in \mathbb{Z} \setminus \{0\}$ such that

$$f = \Psi(f_{d_1})^{a_1} \cdots \Psi(f_{d_k})^{a_k}.$$

Therefore,

$$\{\text{zeros and poles of } f\} \subset \bigcup_{i=1}^k \{\text{zeros and poles of } \Psi(f_{d_i})\}.$$

By Lemma 3.13, there is a simple zero τ^* of $\Psi(f_{d_k})$ which is the unique zero/pole in F_j of maximum imaginary part for $\Psi(f_{d_k})$. Further τ^* has imaginary part greater than that of every zero or pole in F_j of a function $\Psi(f_{d_i})$, for $i < k$. Since $a_k \neq 0$, the function f therefore has a unique zero/pole (depending on the sign of a_k) in F_j of maximum imaginary part.

Enumerate the zeros and poles of f in F_j as w_1, \dots, w_r , where $\text{Im}(w_i) \leq \text{Im}(w_{i+1})$. We then have that $w_r = \tau^*$. By the previous paragraph and Remark 3.4, f satisfies the divisor condition if $\text{Im}(\tau^*) \geq 1$, which is true whenever $d_k \geq 4$.

If $d_k = 3$, then $f = \Psi(f_3)^a$ for some $a \in \mathbb{Z} \setminus \{0\}$. Thus $f = j^{a/3}$ since $j = \Psi(f_3)^3$ by [17, §14, Example 2]. Hence the divisor condition holds for f since it holds for j , as was shown already in the proof of Theorem 3.2 in Section 3.2. Hence, for all d_k , the divisor condition holds for f , and so we are done. \square

Remark 3.14. If f is a modular function with rational Fourier coefficients, then $f \in \mathbb{Q}(j)$, see e.g. [23, Proposition 12.7]. So in particular $f \in \overline{\mathbb{Q}}(j)$ for all those functions covered by Theorem 3.12. Therefore the additional hypothesis in Theorem 3.8 that $f \in \overline{\mathbb{Q}}(j)$ is not in fact a further restriction in this case. (On the other hand, the divisor condition alone does not imply

that $f \in \overline{\mathbb{Q}}(j)$, so in the fully general case Theorem 3.8 does impose an additional restriction on f .)

3.4. Finiteness of multiplicatively dependent tuples

We now prove Theorem 3.7. Our proof will be in two stages. First we will prove the theorem for the case of modular functions $f \in \overline{\mathbb{Q}}(j)$ via an o-minimal counting argument. We then extend the result to modular functions $f \in \mathbb{C}(j)$.

For the first step, we prove the conditional result Proposition 3.16, which covers multiplicatively dependent f -special points for a modular function $f \in \overline{\mathbb{Q}}(j)$ (which does not necessarily satisfy the divisor condition). This result is conditional on a functional independence statement for distinct $\mathrm{GL}_2^+(\mathbb{Q})$ -translates of f , which we formulate below as Condition 3.15. Theorem 3.5 establishes this statement for modular functions satisfying the divisor condition. We may therefore deduce Proposition 3.17, which is Theorem 3.7 for functions $f \in \overline{\mathbb{Q}}(j)$.

We adopt this approach in order to emphasise that, for the case of $f \in \overline{\mathbb{Q}}(j)$, our proof of Theorem 3.7 does not depend on any particular facts about the divisor condition. As we mentioned in Remark 3.9, it may be possible to establish Condition 3.15 under a weaker hypothesis than the divisor condition. If one could do this, then Proposition 3.16 shows that, for $f \in \overline{\mathbb{Q}}(j)$, Theorem 3.7 would hold under this weaker hypothesis.

Condition 3.15. *Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function. If the functions $f(g_1z), \dots, f(g_nz)$ are pairwise distinct for $g_1, \dots, g_n \in \mathrm{GL}_2^+(\mathbb{Q})$, then they are multiplicatively independent modulo constants.*

Proposition 3.16. *Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function such that $f \in \overline{\mathbb{Q}}(j)$. Suppose that f satisfies Condition 3.15. Let $n \geq 1$. Then there exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent.*

Proposition 3.17. *Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function such that $f \in \overline{\mathbb{Q}}(j)$. Suppose that f satisfies the divisor condition. Let $n \geq 1$. Then there exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent.*

PROOF OF PROPOSITION 3.17. Let f be a non-constant modular function satisfying the divisor condition. Then Condition 3.15 holds for f by Theorem 3.5. Proposition 3.17 thus follows from Proposition 3.16. \square

Now we come to the proof of Proposition 3.16. We collect first those arithmetic estimates on f -special points which we will require.

3.4.1. Arithmetic estimates. Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function such that $f \in \overline{\mathbb{Q}}(j)$. Then $f(z) = R(j(z))$ for some rational function R with algebraic coefficients. Suppose that σ is an f -special point. So $\sigma = f(\tau)$ for some $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$. The function f is $\mathrm{SL}_2(\mathbb{Z})$ -invariant, the restricted function $j|_{F_j}: F_j \rightarrow \mathbb{C}$ is injective, and the (non-constant) rational map $R: \mathbb{C} \rightarrow \mathbb{C}$ is finite-to-one on its domain of definition. Thus there exist $\tau_1, \dots, \tau_k \in F_j$, with $k \geq 1$ and $[\mathbb{Q}(\tau_i) : \mathbb{Q}] = 2$ for every i , with the property that, for every $\tau \in F_j$ with

$[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$, one has $f(\tau) = \sigma$ if and only if $\tau \in \{\tau_1, \dots, \tau_k\}$. We may thus make the following definition.

Definition 3.18. Let σ be an f -special point. Let $\tau_1, \dots, \tau_k \in F_j$ be such that, for every $\tau \in F_j$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$, we have that $f(\tau) = \sigma$ if and only if $\tau \in \{\tau_1, \dots, \tau_k\}$. Then the discriminant $\Delta(\sigma)$ of the f -special point σ is defined

$$\Delta(\sigma) = \min\{\Delta(\tau_1), \dots, \Delta(\tau_k)\},$$

where $\Delta(\tau_i)$ is the discriminant of the imaginary quadratic number τ_i (as defined in Definition 3.10).

Observe that, for the j -function, this definition agrees with Definition 1.26.

In this subsection, constants will be positive and with only the indicated dependencies, but in general not effective. Let K be a number field containing the coefficients of R . The arithmetic estimates on f -special points we need are the following.

Lemma 3.19. *For all $\epsilon > 0$, there exists a constant $c(\epsilon, f)$ such that $h(\sigma) \leq c(\epsilon, f)|\Delta|^\epsilon$ for every f -special point σ of discriminant Δ .*

PROOF. Fix $\epsilon > 0$. Let $\tau_0 \in F_j$ be such that $\sigma = f(\tau_0)$ and $[\mathbb{Q}(\tau_0) : \mathbb{Q}] = 2$. Note that $|\Delta(\tau_0)| \leq |\Delta|$. There exists [65, (5.4)] a constant $c(\epsilon)$ such that

$$h(j(\tau_0)) \leq c(\epsilon)|\Delta(\tau_0)|^\epsilon.$$

Recall that $\sigma = R(j(\tau_0))$, where R is a rational function defined over K . The result then follows thanks to elementary properties of the logarithmic

height:

$$h(xy) \leq h(x) + h(y) \text{ for } x, y \in \overline{\mathbb{Q}},$$

$$h(x_1 + \dots + x_r) \leq h(x_1) + \dots + h(x_r) + \log r \text{ for } x_1, \dots, x_r \in \overline{\mathbb{Q}},$$

$$h(x^r) = |r|h(x) \text{ for } r \in \mathbb{Q}, x \in \overline{\mathbb{Q}}^\times,$$

see e.g. [14, Chapter 1]. □

Lemma 3.20. *Let σ be an f -special point of discriminant Δ and suppose $\sigma = f(\tau_0)$ for $\tau_0 \in F_j$ such that $[\mathbb{Q}(\tau_0) : \mathbb{Q}] = 2$. Then $H(\tau_0) \leq 2|\Delta|$.*

PROOF. This follows immediately from [65, (5.5)]. □

Lemma 3.21. *For any $\epsilon > 0$ there exist constants $c_1(\epsilon, f), c_2(\epsilon, f)$ such that*

$$c_1(\epsilon, f)|\Delta|^{\frac{1}{2}-\epsilon} \leq [\mathbb{Q}(\sigma) : \mathbb{Q}] \leq c_2(\epsilon, f)|\Delta|^{\frac{1}{2}+\epsilon}$$

for all f -special points σ of discriminant Δ . Hence, there are also constants $c'_1(\epsilon, f), c'_2(\epsilon, f)$ such that

$$c'_1(\epsilon, f)|\Delta|^{\frac{1}{2}-\epsilon} \leq [K(\sigma) : K] \leq c'_2(\epsilon, f)|\Delta|^{\frac{1}{2}+\epsilon}.$$

PROOF. Fix $\epsilon > 0$. There exists $c(\epsilon)$ such that

$$c(\epsilon)|\Delta(\tau)|^{\frac{1}{2}-\epsilon} \leq [\mathbb{Q}(j(\tau)) : \mathbb{Q}] \leq c(\epsilon)|\Delta(\tau)|^{\frac{1}{2}+\epsilon},$$

for every $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$, see [65, (5.6), (5.7)]. We will show that there exists a constant $M(f) > 0$ such that

$$(3.1) \quad \frac{1}{M(f)}[\mathbb{Q}(j(\tau)) : \mathbb{Q}] \leq [\mathbb{Q}(f(\tau)) : \mathbb{Q}] \leq M(f)[\mathbb{Q}(j(\tau)) : \mathbb{Q}]$$

holds for all $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$. The first part of the lemma then follows by combining these two inequalities, each applied to some $\tau_1 \in F_j$ with the property that $\sigma = f(\tau_1)$, $[\mathbb{Q}(\tau_1) : \mathbb{Q}] = 2$, and $\Delta(\tau_1) = \Delta$.

For $f \in \mathbb{Q}(j)$ non-constant, the necessary inequality (3.1) was proved by Spence in [80] (which treats actually the broader case $f \in \mathbb{Q}(j, \chi^*)$, where χ^* is a certain almost holomorphic modular function). The extension to $f \in \overline{\mathbb{Q}}(j)$ was given by Spence in private communication with the author; we include a proof below.

Let $d = [K : \mathbb{Q}]$ and take $\tau \in \mathbb{H}$ with $[\mathbb{Q}(\tau) : \mathbb{Q}] = 2$. For the upper bound, observe that

$$\begin{aligned} [\mathbb{Q}(f(\tau)) : \mathbb{Q}] &\leq [K(j(\tau)) : \mathbb{Q}] \\ &= [K(j(\tau)) : \mathbb{Q}(j(\tau))] [\mathbb{Q}(j(\tau)) : \mathbb{Q}] \\ &\leq d \cdot [\mathbb{Q}(j(\tau)) : \mathbb{Q}]. \end{aligned}$$

Now for the lower bound. Suppose that $[\mathbb{Q}(j(\tau), f(\tau)) : \mathbb{Q}(f(\tau))] > M$ for some M . Then

$$\begin{aligned} &[K(j(\tau), f(\tau)) : K(f(\tau))] [K(f(\tau)) : \mathbb{Q}(f(\tau))] \\ &= [K(j(\tau), f(\tau)) : \mathbb{Q}(f(\tau))] \\ &= [K(j(\tau), f(\tau)) : \mathbb{Q}(j(\tau), f(\tau))] [\mathbb{Q}(j(\tau), f(\tau)) : \mathbb{Q}(f(\tau))] \\ &> [K(j(\tau), f(\tau)) : \mathbb{Q}(j(\tau), f(\tau))] \cdot M. \end{aligned}$$

Hence, using that $K(j(\tau)) = K(j(\tau), f(\tau))$ since $f(\tau)$ is K -rational in $j(\tau)$, we have that

$$[K(j(\tau)) : K(f(\tau))] > \frac{[K(j(\tau), f(\tau)) : \mathbb{Q}(j(\tau), f(\tau))]}{[K(f(\tau)) : \mathbb{Q}(f(\tau))]} \cdot M \geq \frac{M}{d}.$$

Taking Galois conjugates of $j(\tau)$ over $K(f(\tau))$, we then obtain at least M/d distinct quadratic points τ' in the standard fundamental domain F_j with $f(\tau') = f(\tau)$. For suitably large M though, this is impossible since f is meromorphic and, restricted to F_j , definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$, so uniform finiteness applies (Theorem A.4).

For M large enough (depending on f), we therefore have that

$$[\mathbb{Q}(j(\tau), f(\tau)) : \mathbb{Q}(f(\tau))] \leq M.$$

Fix such an M . Then

$$[\mathbb{Q}(j(\tau), f(\tau)) : \mathbb{Q}] \leq M \cdot [\mathbb{Q}(f(\tau)) : \mathbb{Q}],$$

and so

$$[\mathbb{Q}(f(\tau)) : \mathbb{Q}] \geq \frac{1}{M} \cdot [\mathbb{Q}(j(\tau), f(\tau)) : \mathbb{Q}] \geq \frac{1}{M} \cdot [\mathbb{Q}(j(\tau)) : \mathbb{Q}].$$

Therefore, taking $M(f) = \max\{M, d\}$ completes the proof of (3.1).

For the second part of the lemma, the upper bound follows since

$$[K(\sigma) : K] \leq [\mathbb{Q}(\sigma) : \mathbb{Q}],$$

while the lower bound comes from the fact that

$$[K(\sigma) : K][K : \mathbb{Q}] = [K(\sigma) : \mathbb{Q}] \geq [\mathbb{Q}(\sigma) : \mathbb{Q}],$$

and thus

$$[K(\sigma) : K] \geq \frac{1}{d} \cdot [\mathbb{Q}(\sigma) : \mathbb{Q}]. \quad \square$$

We will also require the following result, which allows us to bound the “complexity” of a multiplicative dependence.

Lemma 3.22. *Let $\alpha_1, \dots, \alpha_n$ be multiplicatively dependent non-zero elements of a number field L of degree $d \geq 2$. Suppose that any proper subset of the α_i is multiplicatively independent. Then there exists a constant $c(n)$ and $b_1, \dots, b_n \in \mathbb{Z}$, not all zero, such that $\alpha_1^{b_1} \cdots \alpha_n^{b_n} = 1$ and*

$$|b_i| \leq c(n)d^n \log d \prod_{\substack{j=1 \\ j \neq i}}^n h(\alpha_j), \quad i = 1, \dots, n.$$

PROOF. See [45, Corollary 3.2] for the case $n \geq 2$. If $n = 1$, then α_1 is a root of unity of degree $\leq d$, and hence the desired result follows from elementary bounds for the Euler ϕ -function. \square

3.4.2. Conjugates of f -special points. As in Section 1.7, we write T_Δ for the set of integer triples (a, b, c) such that: $\gcd(a, b, c) = 1$, $\Delta = b^2 - 4ac$, and either $-a < b \leq a < c$ or $0 \leq b \leq a = c$. Recall that there is a bijection between T_Δ and the j -special points of discriminant Δ given by $(a, b, c) \mapsto j((b + \sqrt{\Delta})/2a)$. Further, the j -special points of discriminant Δ form a full Galois orbit over \mathbb{Q} .

Lemma 3.23. *Let K be a number field, $R \in K(t)$ a non-constant rational function, and $f: \mathbb{H} \rightarrow \mathbb{C}$ the modular function defined by $f(z) = R(j(z))$. Let σ be an f -special point of discriminant Δ . Then the set*

$$\left\{ f\left(\frac{b + \sqrt{\Delta}}{2a}\right) : (a, b, c) \in T_\Delta \right\}$$

contains all K -conjugates of σ , which are also f -special points of discriminant Δ .

PROOF. We may write $\sigma = R(j)$, where j is a j -special point of discriminant Δ . Let σ' be a K -conjugate of σ . Since R is a rational function over K , one has that $\sigma' = R(j')$ where j' is a K -conjugate of j . But we have a complete description of the conjugates of j , they arise as the j -special points $j((b + \sqrt{\Delta})/2a)$ for $(a, b, c) \in T_\Delta$. So σ' is f -special and belongs to the set in the lemma.

It remains to show that the discriminant of the f -special point σ' is equal to Δ . Since $\sigma' = R(j')$ and j' is a j -special point of discriminant Δ , this will be true unless $\sigma' = R(j_0)$ for some j -special point j_0 of discriminant $\Delta_0 < \Delta$. If this were to happen, then one would have that $\sigma = R(j'_0)$ for some K -conjugate j'_0 of j_0 since σ, σ' are K -conjugate. But j'_0 would be a j -special point of discriminant Δ_0 , and so the discriminant of σ would have to be $\leq \Delta_0$, implying that $\Delta \leq \Delta_0 < \Delta$, a contradiction. \square

The proof of Proposition 3.16 will be by an o-minimal counting argument. We first establish a lower bound for the number of multiplicatively dependent tuples of f -special points of a given discriminant. We will then apply the Pila–Wilkie o-minimal Counting Theorem [66] (and in particular its extensions to algebraic points [59, 39]) to count the preimages of such tuples. The lower bound we use comes from a lower bound for the size of the Galois orbit of an f -special point. We will thus require that conjugates of f -special points are again f -special. This is why the o-minimal argument in Subsection 3.4.3 treats only the case of $f \in \overline{\mathbb{Q}}(j)$.

3.4.3. Proof of Proposition 3.16. Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function such that Condition 3.15 holds for f . Suppose also that $f \in \overline{\mathbb{Q}}(j)$, so that $f(z) = R(j(z))$ for some rational function $R \in \overline{\mathbb{Q}}(t)$. Let K be a number field containing the coefficients of R . Let V be the intersection with $Y(1) \times \mathbb{G}_m$ of the graph of R (viewed as a subset of $\mathbb{P}^1 \times \mathbb{P}^1$). Fix $n \geq 1$. Let $X = X_{n,n} = Y(1)^n \times \mathbb{G}_m^n$, so that

$$V^n = \{(x_1, \dots, x_n, t_1, \dots, t_n) \in X : t_i = R(x_i) \text{ for } i = 1, \dots, n\} \subset X.$$

Recall the definition of a (weakly) special subvariety of X from Section 2.1.

The proof of Proposition 3.16 is similar to the proof of [65, Theorem 1.2].

PROOF OF PROPOSITION 3.16. In the following, constants c, c' are positive and depend only on our choice of f and n , but will vary between occurrences.

The complexity $\Delta(\sigma)$ of an n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of f -special points is defined to be $\max\{|\Delta(\sigma_1)|, \dots, |\Delta(\sigma_n)|\}$. We define an f -dependent tuple to be an n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of distinct f -special points satisfying a non-trivial multiplicative relation and minimal for this property.

Let

$$Y = \left\{ (z, u, r, s) \in F_j^n \times F_e^n \times \mathbb{R}^n \times \mathbb{R} : R(j(z)) = e(u), r \cdot u = s \right\}$$

and

$$Z = \left\{ (z, r, s) \in F_j^n \times \mathbb{R}^n \times \mathbb{R} : \exists u(z, u, r, s) \in Y \right\}.$$

Then $(j(z), \exp(u)) \in V^n$ for $(z, u, r, s) \in Y$, and Y, Z are definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$.

An f -dependent tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of complexity Δ gives rise to a point $(x_1, \dots, x_n, \sigma_1, \dots, \sigma_n) \in V^n$, where each x_i is a j -special point of discriminant $\Delta(\sigma_i)$ satisfying $R(x_i) = \sigma_i$. This point in V^n has a preimage

$$\tau = (z_1, \dots, z_n, u_1, \dots, u_n) \in F_j^n \times F_{\text{exp}}^n.$$

Now τ gives rise to the point

$$(z_1, \dots, z_n, u_1, \dots, u_n, b_1, \dots, b_n, b) \in Y,$$

where the b_i, b are rational integers recording the multiplicative dependence of $\sigma_1, \dots, \sigma_n$, i.e. such that

$$\sum_{i=1}^n b_i u_i = b.$$

Since x_i is a j -special point of discriminant $\Delta(\sigma_i)$, so z_i is a quadratic point with height bounded by $2|\Delta(\sigma_i)|$ by Lemma 3.20. The σ_i are f -special points of discriminants $\Delta(\sigma_i)$ such that the set $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent and minimal for this property. We may therefore use the bounds on the logarithmic height (Lemma 3.19) and degree (Lemma 3.21) of the σ_i , together with the result in Lemma 3.22, to see that the integers b_i may be chosen to have absolute value bounded by $c\Delta^{n^2}$. Since

$$\sum_{i=1}^n b_i u_i = b$$

and $0 \leq \text{Re}(u_i) < 1$, we obtain also a bound on $|b|$. Thus, the height of the point

$$(z_1, \dots, z_n, b_1, \dots, b_n, b) \in Z$$

is bounded by $c\Delta^{n^2}$ since $H(k) = |k|$ for $k \in \mathbb{Z} \setminus \{0\}$. Further, this point is quadratic in the z_i coordinates and rational (even integral) in the b_i, b coordinates.

Suppose then that there are infinitely many f -dependent tuples. Then, in particular, there are f -dependent tuples of arbitrarily large complexity Δ . By Lemma 3.21 such a tuple σ has at least $c\Delta^{1/4}$ conjugates over K , and by Lemma 3.23 each of these conjugates σ' is itself an f -dependent tuple of complexity Δ . Further, all these conjugate tuples satisfy the same multiplicative relation. Consequently, each of these σ' gives rise to a point

$$(z'_1, \dots, z'_n, u'_1, \dots, u'_n, b_1, \dots, b_n, b') \in Y,$$

which is quadratic in the z'_i coordinates and rational (even integral) in the b_i, b' coordinates. Further, the height of the corresponding Z -point is bounded by $c\Delta^{n^2}$. Note that the (b_1, \dots, b_n) coordinates are the same for all these conjugates, but the (z'_1, \dots, z'_n) and (u'_1, \dots, u'_n) coordinates must vary between distinct conjugates.

View Y as a definable family parametrised by \mathbb{R}^n . Let $Y_{(b_1, \dots, b_n)}$ denote the fibre of Y over the point $(b_1, \dots, b_n) \in \mathbb{R}^n$. Let $\Sigma \subset Y_{(b_1, \dots, b_n)}$ denote the set of points of $Y_{(b_1, \dots, b_n)}$ arising from the conjugates of σ . Let π_1, π_2 denote the projections of $Y_{(b_1, \dots, b_n)}$ to $F_j^n \times \mathbb{R}$ and F_{exp}^n respectively.

We are now in a position to apply the o-minimal Counting Theorem in the form of [39, Corollary 7.2] (see Theorem A.21). Provided that Δ is suitably large (which we may always assume), we thereby obtain a continuous, definable function $\beta: [0, 1] \rightarrow Y_{(b_1, \dots, b_n)}$ such that:

- (1) The composition $\pi_1 \circ \beta: [0, 1] \rightarrow F_j^n \times \mathbb{R}$ is semialgebraic and its restriction to $(0, 1)$ is real analytic.
- (2) The composition $\pi_2 \circ \beta: [0, 1] \rightarrow F_e^n$ is non-constant.
- (3) $\beta(0) \in \Sigma$.

(In [39, Corollary 7.2], it is only stated that $\pi_2(\beta(0)) \in \pi_2(\Sigma)$, but in fact the authors prove there that $\beta(0) \in \Sigma$.) Note also that (2) implies that the composition of β with projection to the F_j^n coordinates is non-constant. Denote by $Z_{(b_1, \dots, b_n)}$ the fibre of Z over the point (b_1, \dots, b_n) . Note that $(\pi_1 \circ \beta)([0, 1]) \subset Z_{(b_1, \dots, b_n)}$.

There thus exists a continuous, semialgebraic map $\gamma: [0, 1] \rightarrow Z_{(b_1, \dots, b_n)}$ which projects non-constantly to the F_j^n coordinates, maps 0 to a point of $Z_{(b_1, \dots, b_n)}$ corresponding to a conjugate of σ , and whose restriction to $(0, 1)$ is real analytic.

Observe that the set $\gamma([0, 1])$ is contained in the set

$$\tilde{V} = \{(z, t) \in \mathbb{H}^n \times \mathbb{C} : \prod_{i=1}^n f(z_i)^{b_i} = e(t)\}.$$

There then exists an open neighbourhood $\Omega \subset \mathbb{H}^n \times \mathbb{C}$ of $\gamma(0)$ and a set P such that:

- (1) $\Omega \cap \gamma([0, 1]) \subset P \subset \tilde{V}$,
- (2) P may be written as a finite union of irreducible Nash subsets of Ω which each contain $\gamma(0)$.

This follows from [1, Proposition 1] and the characterisation of Nash sets given in [1, p. 989].

If every complex-analytically irreducible component of P had constant projection to its \mathbb{H}^n coordinates, then $\gamma([0, 1])$ would have constant projection to its \mathbb{H}^n coordinates by real analytic continuation. But $\gamma([0, 1])$ does not have constant projection to its \mathbb{H}^n coordinates. Hence, there must exist some complex-analytically irreducible component of P which has non-constant projection to its \mathbb{H}^n coordinates. Observe that every complex-analytically irreducible component of P contains $\gamma(0)$. Hence, by complex analytic continuation, there exists a complex algebraic subvariety $W \subset \mathbb{C}^{n+1}$ and a complex-analytically irreducible component $A \subset (\mathbb{H}^n \times \mathbb{C}) \cap W$ such that $\gamma(0) \in A \subset \tilde{V}$ and A has non-constant projection to its \mathbb{H}^n coordinates.

By the Ax–Schanuel results of Section 2.3, there exist weakly special subvarieties $W_1 \subset \mathbb{H}^n$ and $W_2 \subset \mathbb{C}$ such that $A \subset W_1 \times W_2 \subset \tilde{V}$. The projection $\tilde{V} \rightarrow \mathbb{H}^n$ has discrete fibres, and hence W_2 must just be a point, which therefore has to be equal to the projection of $\gamma(0)$. So $W_2 = \{b'\}$ for some $b' \in \mathbb{Z}$.

Therefore, W_1 is a positive-dimensional weakly special subvariety of \mathbb{H}^n which is contained in the set

$$\{z \in \mathbb{H}^n : \prod_{i=1}^n f(z_i)^{b_i} = 1\}.$$

In addition, W_1 contains a preimage $(\tau_1, \dots, \tau_n) \in \mathbb{H}^n$ of some f -dependent tuple $(f(\tau_1), \dots, f(\tau_n))$. In particular, $f(W_1)$ has no two identically equal coordinates, since the $f(\tau_i)$ are pairwise distinct. Note that all the b_i are

non-zero, since the set $\{f(\tau_1), \dots, f(\tau_n)\}$ is minimally multiplicatively dependent. Taking the image of W_1 under f , we therefore obtain a multiplicative dependence modulo constants among some pairwise distinct $\mathrm{GL}_2^+(\mathbb{Q})$ -translates of f . This contradicts Condition 3.15, and so we are done. \square

3.4.4. Proof of Theorem 3.7. Theorem 3.7 follows from Proposition 3.17 by a specialisation argument. Before coming to this proof, we give the following elementary lemma.

Lemma 3.24. *Let $x_1, \dots, x_n \in \mathbb{G}_m$ be pairwise distinct. Suppose that*

$$x_1^{a_1} \cdots x_n^{a_n} = 1$$

for some $a_i \in \mathbb{Z}$ with $a_1 \neq 0$. Then there exists a set $S \subset \{x_1, \dots, x_n\}$ such that $x_1 \in S$ and S is minimally multiplicatively dependent.

PROOF. We proceed by induction on n . If $n = 1$, then the result is trivial. Now let $n > 1$. Suppose that $x_1, \dots, x_n \in \mathbb{G}_m$ are pairwise distinct and such that

$$x_1^{a_1} \cdots x_n^{a_n} = 1$$

for some $a_i \in \mathbb{Z}$ with $a_1 \neq 0$. Suppose for contradiction that there is no minimally multiplicatively dependent subset of $\{x_1, \dots, x_n\}$ which contains x_1 . Then, in particular, the set $\{x_2, \dots, x_n\}$ must be multiplicatively dependent. So

$$x_2^{b_2} \cdots x_n^{b_n} = 1$$

for some $b_2, \dots, b_n \in \mathbb{Z}$ not all zero. Without loss of generality, we assume that $b_2 \neq 0$. We may then eliminate x_2 from the first equality, to obtain

that

$$x_1^{a_1 b_2} x_3^{a_3 b_2 - a_2 b_3} \cdots x_n^{a_n b_2 - a_2 b_n} = 1.$$

Note that $a_1 b_2 \neq 0$. By induction, there exists a set $S \subset \{x_1, x_3, \dots, x_n\}$ such that $x_1 \in S$ and S is minimally multiplicatively dependent, and so we are done. \square

Now we may proceed with the proof of Theorem 3.7. Fix $f \in \mathbb{C}(j)$ to be a non-constant modular function satisfying the divisor condition. Write $f(z) = R(j(z))$ for some rational function $R(t) \in \mathbb{C}(t)$. We may assume that $R(t) \notin \overline{\mathbb{Q}}(t)$, else we would be in the case of Proposition 3.17. We may write

$$R(t) = \frac{p(a, t)}{q(a, t)},$$

where the polynomials p, q have integral coefficients and $a \in \mathbb{C}^N$ for some $N \geq 1$. We let $a = (a_0, \dots, a_{N-1})$ such that

$$p(a, t) = a_0 \prod_{i=1}^k (t - a_i)$$

and

$$q(a, t) = \prod_{i=r+1}^{N-1} (t - a_i).$$

Note that some coordinate a_i must be transcendental since we have assumed that $R(t) \notin \overline{\mathbb{Q}}(t)$. We may also assume that the polynomials $p(a, t), q(a, t) \in \mathbb{C}[t]$ are relatively prime in $\mathbb{C}[t]$.

We need the following fact about the divisor condition.

Lemma 3.25. *Let $W \subset \mathbb{C}^N$ be the smallest affine variety over $\overline{\mathbb{Q}}$ such that $a \in W$. Suppose the divisor condition holds for f . Then there exists an open neighbourhood U of a such that the divisor condition holds for the*

modular function

$$f_{a'}(z) = \frac{p(a', j(z))}{q(a', j(z))},$$

for every $a' \in W \cap U$.

PROOF. Let U be an open neighbourhood of a . Take $a' \in U$ and write $a' = (a'_0, \dots, a'_n)$. Provided that U is chosen suitably small, the divisor condition will hold for the modular function $f_{a'}(z)$, unless possibly one of the following happens:

- (1) $a_i = a_k$, but $a'_i \neq a'_k$ for some i, k ;
- (2) $a_i \in \{0, 1728\}$ and $a'_i \neq a_i$ for some i .

Both these possibilities are ruled out though if $a' \in W$. \square

We now complete the proof of Theorem 3.7. The approach taken is based on a suggestion of Ehud Hrushovski.

PROOF OF THEOREM 3.7. Fix $n \geq 1$. Let Y be the set of j -special points. Let Z be the set of $y = (y_1, \dots, y_n) \in Y^n$ such that the f -special points $R(y_1), \dots, R(y_n)$ are pairwise distinct and the set $\{R(y_1), \dots, R(y_n)\}$ is minimally multiplicatively dependent. We need to show that the set Z is finite. Suppose then, for contradiction, that Z is infinite. Then, in particular, the set

$$\tilde{Z} = \{\sigma \in Y : \exists y \in Z \text{ such that } \sigma = \pi_i(y) \text{ for some } i = 1, \dots, n\}$$

is infinite, where $\pi_i: Y^n \rightarrow Y$ denotes the projection map to the i th coordinate.

If $y = (y_1, \dots, y_n) \in Z$, then there are $m_1, \dots, m_n \in \mathbb{Z} \setminus \{0\}$ such that

$$p(a, y_1)^{m_1} \cdots p(a, y_n)^{m_n} = q(a, y_1)^{m_1} \cdots q(a, y_n)^{m_n}.$$

Multiplying by those factors with $m_i < 0$, we get a corresponding polynomial equality, which we denote (*). (Note that the equality (*) depends on y , but is always of the same form.) In addition, one has that

$$(**) \quad p(a, y_i) \neq 0 \text{ and } q(a, y_i) \neq 0 \text{ for } i = 1, \dots, n$$

for all $y \in Z$. Since $R(y_1), \dots, R(y_n)$ are pairwise distinct, one also has that

$$(***) \quad p(a, y_i)q(a, y_k) - p(a, y_k)q(a, y_i) \neq 0$$

for all $1 \leq i < k \leq n$ and $y \in Z$.

Let $W \subset \mathbb{C}^N$ be the smallest affine variety over $\overline{\mathbb{Q}}$ such that $a \in W$. Let $K \subset \mathbb{C}$ be a number field over which W is defined. If $\dim W = 0$, then we immediately contradict Proposition 3.17. So we may assume that $\dim W \geq 1$. Note that, by Lemma 3.25, there is some open neighbourhood $U \subset \mathbb{C}^N$ such that $a \in U$ and, for every $a' \in U \cap W$, the divisor condition holds for the modular function $f_{a'}$.

Suppose now that $\dim W = 1$. Given some $y = (y_1, \dots, y_n) \in Z$, a condition $p(x, y_i) = 0$ then defines a finite subset of W , of size at most d say. Now y_i is a j -special point, so is contained in a finite extension of \mathbb{Q} which is “dihedral” (see e.g. [36, p. 191]) and hence, in particular, solvable. So any a' satisfying the equation $p(a', y_i) = 0$ lies in an extension of K which has an index $\leq d$ solvable subextension (and similarly for an equation $q(a', y_i) = 0$ or an equation $p(a', y_i)q(a', y_k) - p(a', y_k)q(a', y_i) = 0$ where $i \neq k$).

Recall that a has at least one transcendental coordinate, say the first coordinate. Let $\pi_1: W \rightarrow \mathbb{C}$ be the projection map to the first coordinate. Then π_1 is an open mapping at a , so $\pi_1(W \cap U)$ contains an open neighbourhood $U' \subset \mathbb{C}$ such that $\pi_1(a) \in U'$. If we find $\alpha \in U'$ with Galois group G (over K), then any $a' \in W \cap U$ with $\pi_1(a') = \alpha$ will be algebraic and have a Galois group which contains G as a subgroup.

Fix some $b \in U' \cap K(i)$. Let c be any algebraic number whose Galois group over $K(i)$ is isomorphic to the alternating group A_m (for some suitably large m); such a c exists by [19, Corollary 12]. Fix some $r \in \mathbb{Q}$ with $r > 0$ and set $\alpha = b + rc$. Provided r is chosen suitably small, we have that $\alpha \in U'$. The Galois group of α over $K(i)$ is isomorphic to A_m , and hence the Galois group of α over K contains a subgroup isomorphic to A_m .

Now fix an algebraic $a' \in W \cap U$ with $\pi_1(a') = \alpha$. Then the Galois group of a' over K contains a subgroup isomorphic to A_m . In particular, for every $y = (y_1, \dots, y_n) \in Z$, one has that

$$p(a', y_i)q(a', y_i) \neq 0$$

for all $1 \leq i \leq n$ and

$$p(a', y_i)q(a', y_k) - p(a', y_k)q(a', y_i) \neq 0$$

for all $1 \leq i < k \leq n$. Note also that, for each $y = (y_1, \dots, y_n) \in Z$, the respective dependence $(*)$ holds with a' in place of a , since $a' \in W$.

Define the rational function $R_{a'} \in \overline{\mathbb{Q}}(t)$ by

$$R_{a'}(t) = \frac{p(a', t)}{q(a', t)},$$

so that $f_{a'}(z) = R_{a'}(j(z))$. For each $1 \leq k \leq n$, let A_k be the set of $y = (y_1, \dots, y_k) \in Y^k$ such that the $f_{a'}$ -special points $R_{a'}(y_1), \dots, R_{a'}(y_k)$ are pairwise distinct and the set $\{R_{a'}(y_1), \dots, R_{a'}(y_k)\}$ is minimally multiplicatively dependent. Then let

$$\tilde{A}_k = \{\sigma \in Y : \exists y \in A_k \text{ such that } \sigma = \pi_i(y) \text{ for some } i = 1, \dots, k\},$$

where $\pi_i: Y^k \rightarrow Y$ is the projection map to the i th coordinate as usual.

Let $y = (y_1, \dots, y_n) \in Z$. Then the above argument shows that $R_{a'}(y_1), \dots, R_{a'}(y_n)$ are pairwise distinct $f_{a'}$ -special points and

$$R_{a'}(y_1)^{m_1} \cdots R_{a'}(y_n)^{m_n} = 1,$$

for some $m_1, \dots, m_n \in \mathbb{Z} \setminus \{0\}$. The set $\{R_{a'}(y_1), \dots, R_{a'}(y_n)\}$ is not necessarily minimally multiplicatively dependent, but we may apply Lemma 3.24 since $R_{a'}(y_1), \dots, R_{a'}(y_n)$ are pairwise distinct and $m_1, \dots, m_n \neq 0$. We thereby obtain that, for each $i = 1, \dots, n$, there exists a minimally multiplicatively dependent subset of $\{R_{a'}(y_1), \dots, R_{a'}(y_n)\}$ which contains $R_{a'}(y_i)$. In particular, we see that

$$y_1, \dots, y_n \in \bigcup_{k=1}^n \tilde{A}_k.$$

Consequently, $\tilde{Z} \subset \bigcup_{k=1}^n \tilde{A}_k$, and so the set $\bigcup_{k=1}^n \tilde{A}_k$ must be infinite.

This though cannot happen. Since $a' \in U \cap W$, the divisor condition holds for the modular function $f_{a'}$ by Lemma 3.25. Also, $f_{a'} \in \overline{\mathbb{Q}}(j)$ since a' is algebraic. Therefore, by Proposition 3.17, for each $1 \leq k \leq n$, there are only finitely many k -tuples $(\sigma_1, \dots, \sigma_k)$ of pairwise distinct $f_{a'}$ -special

points such that the set $\{\sigma_1, \dots, \sigma_k\}$ is minimally multiplicatively dependent. Consequently, the sets A_1, \dots, A_n must all be finite, since the rational function $R_{a'}$ is finite-to-one. So the sets $\tilde{A}_1, \dots, \tilde{A}_n$ are all finite too. Thus, the set $\bigcup_{k=1}^n \tilde{A}_k$ is also finite, and this gives us the desired contradiction.

So we are left with the case where $\dim W > 1$. For $y \in Z$, define

$$W_y = \{w \in W : \text{either } p(w, y_i)q(w, y_i) = 0 \text{ for some } i = 1, \dots, n, \text{ or} \\ p(w, y_i)q(w, y_k) - p(w, y_k)q(w, y_i) = 0 \text{ for some } 1 \leq i < k \leq n\}.$$

Then $a \notin W_y$ by $(**)$ and $(***)$, and so each W_y is a proper subvariety of W . The polynomials p, q are of bounded degree, so all the W_y have bounded degree.

Let W' be an irreducible subvariety of W of codimension 1, defined over $\overline{\mathbb{Q}}$, and of degree higher than any of the W_y . By the theorem of Bertini, such a W' arises as the intersection of W with a sufficiently general hypersurface H in \mathbb{C}^N of large enough degree, defined over $\overline{\mathbb{Q}}$. The subvariety W' is not contained in any W_y . Given $y \in Z$, the corresponding equation $(*)$ holds with any $a' \in W'$ in place of a . For a generic $a' \in W'$, the inequalities $(**)$ and $(***)$ hold for every $y \in Z$. The union of all $H \cap W$ for H as in Bertini's theorem is a constructible, Zariski dense subset of W , and therefore contains a Zariski open subset of W . We may thus assume that H is such that $H \cap W \cap U$ is non-empty and of codimension 1 in $W \cap U$.

We now continue this process inductively. Eventually, we thereby obtain a variety W^* which is defined over $\overline{\mathbb{Q}}$, has dimension 1, and also satisfies the following two conditions. First, for every $y \in Z$, the corresponding equation $(*)$ holds with any $a' \in W^*$ in place of a . Second, for a generic

$a' \in W^*$, the inequalities (**) and (***) hold for every $y \in Z$. One thereby reduces to the above dimension 1 case, and the proof is complete. \square

Remark 3.26. Observe that the above proof depends crucially on the specific properties of the divisor condition used in Lemma 3.25. In particular, one could not use the argument in this section to prove a result corresponding to Proposition 3.16 for arbitrary non-constant modular functions $f \in \mathbb{C}(j)$ satisfying Condition 3.15 (but not necessarily the divisor condition).

We end this section by showing that the divisor condition is not necessary for the $n = 1$ case of Theorem 3.7, i.e. the case of f -special points which are also roots of unity. (Note it is well-known that no j -special point is a root of unity, see e.g. [65, p. 1365].)

Proposition 3.27. *Let $f \in \mathbb{C}(j)$ be non-constant. Then there are only finitely many f -special points σ such that σ is also a root of unity.*

PROOF. We treat first the algebraic case. Let $f \in \overline{\mathbb{Q}}(j)$ be non-constant. Then we may write $f(z) = R(j(z))$ for some rational function R with algebraic coefficients. Suppose x is a j -special point such that $R(x)$ is a root of unity. By Kronecker's theorem we have that $h(R(x)) = 0$. Viewing the rational function R as a morphism $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ defined over $\overline{\mathbb{Q}}$, one may use [40, Theorem B.2.5] to obtain that $h(x) \leq c$, where c is a constant depending only on R . Write Δ for the discriminant of the j -special point x . Then [37, Lemma 3] gives that

$$h(x) \geq c_1 \log|\Delta| - c_2,$$

for some absolute constants $c_1, c_2 > 0$. Putting these two inequalities together, we obtain that $|\Delta|$ must be bounded above by a constant depending only on R . There are thus only finitely many j -special points x such that $R(x)$ is a root of unity. Hence, there are only finitely many f -special points which are also roots of unity.

Now let $f(z) = R(j(x))$ for an arbitrary non-constant rational function R . Suppose that there are infinitely many f -special points which are also roots of unity. Since j -special points are algebraic, there are then infinitely many $x \in \overline{\mathbb{Q}}$ such that $R(x) \in \overline{\mathbb{Q}}$. Hence R must be defined over $\overline{\mathbb{Q}}$, and so we contradict the above algebraic case of the proposition. \square

Remark 3.28. Proposition 5.5 is included within Proposition 2.17. The alternative proof of Proposition 5.5 we give here uses the lower bound for the height of a j -special point due to Habegger [37], in place of the appeal to Theorem 2.3 used to prove Proposition 2.17. It thus gives a proof of the result which uses only intrinsic properties of j -special points, rather than appealing to the very general Theorem 2.3 (as we had to do in the proof of Proposition 2.17).

3.5. Describing the exemplary components

Let $f \in \mathbb{C}(j)$ be a non-constant modular function. Write $f(z) = R(j(z))$ for some rational function R . Let V be the intersection with $Y(1) \times \mathbb{G}_m$ of the graph of R (viewed as a subset of $\mathbb{P}^1 \times \mathbb{P}^1$). We explain in this section how the results of Section 3.4 strengthen those of Chapter 2 in this case.

Let $n \geq 1$. Recall from Chapter 2 the definition of an exemplary component of V^n . In particular, if $(\sigma_1, \dots, \sigma_n)$ is an n -tuple of f -special points

such that the set $\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent, then there are j -special points x_1, \dots, x_n such that the point $(x_1, \dots, x_n, \sigma_1, \dots, \sigma_n)$ belongs to a dependent exemplary component of V^n .

Theorem 3.29. *Suppose that f satisfies Condition 3.15. Let $n \geq 1$. Suppose that $W \subset V^n$ is an exemplary component. Then there exist disjoint (possibly empty) sets $I, J \subset \{1, \dots, n\}$ such that $I \cup J = \{1, \dots, n\}$ and:*

- (1) *For each $i \in I$, there exists a j -special point x_i and some $a_i \in \mathbb{Z} \setminus \{0\}$ such that $\prod_{i \in I} R(x_i)^{a_i} = 1$ and the x_i are all distinct.*
- (2) *If $J \neq \emptyset$, then there exists $k \geq 1$, a partition J_1, \dots, J_k of J , and, for each $i = 1, \dots, k$, a distinguished element $j_i \in J_i$.*
- (3) *W is a geometrically-irreducible component of $V^n \cap (M \times R(M))$ which projects onto M , where*

$$M = \{(t_1, \dots, t_n) \in Y(1)^n : \forall i \in I \quad t_i = x_i \text{ and} \\ \forall i = 1, \dots, k \quad \forall j \in J_i \quad t_j = t_{j_i}\}.$$

Conversely, for every special subvariety $M \subset Y(1)^n$ of this form there exists an exemplary component W of V^n such that W projects onto M .

PROOF. This description follows immediately from Condition 3.15, the definition of an exemplary component, and the description of the special subvarieties of \mathbb{G}_m^n in Section 2.1. \square

Condition 3.15 thus imposes a very strong restriction on the form of the positive-dimensional exemplary components of V^n . It is essentially this restriction which explains why, in this chapter, we can obtain finiteness results on n -tuples of multiplicatively dependent images of special points;

results which we are not able to obtain in the more general setting of Chapter 2.

Remark 3.30. Suppose that f satisfies Condition 3.15. Let $n \geq 1$. If, for each $k \leq n$, there are only finitely many k -tuples (x_1, \dots, x_k) of distinct j -special points such that the set $\{R(x_1), \dots, R(x_n)\}$ is multiplicatively dependent, then there are only finitely many exemplary components of V^n .

Proposition 3.31. *Suppose that f satisfies Condition 3.15. Let $n \geq 1$. Suppose that, for each $k \leq n$, there are only finitely many k -tuples $(\sigma_1, \dots, \sigma_k)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_k\}$ is multiplicatively dependent and minimal for this property. Then there are only finitely many exemplary components of V^n .*

PROOF. Let $n \geq 1$. Let Y be the set of j -special points x such that there exists a j -special point $y \neq x$ such that $R(x) = R(y)$. By André's theorem [4], the set Y is finite. Let Z be the set of j -special points x such that the f -special point $R(x)$ appears as some coordinate σ_i of a k -tuple $(\sigma_1, \dots, \sigma_k)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_k\}$ is minimally multiplicatively dependent, for some $k \leq n$. Then Z is finite by assumption, since the rational function R is finite-to-one.

Now let $k \leq n$ and suppose that (x_1, \dots, x_k) is a k -tuple of distinct j -special points such that there are $a_i \in \mathbb{Z} \setminus \{0\}$ with $\prod_{i=1}^k R(x_i)^{a_i} = 1$. For $m = 1, \dots, k$, if $x_m \notin Y$, then by Lemma 3.24 there exists a subset $I \subset \{1, \dots, k\}$ such that: $x_m \in \{x_i : i \in I\}$, the f -special points $R(x_i)$ for $i \in I$ are pairwise distinct, and the set $\{R(x_i) : i \in I\}$ is minimally multiplicatively dependent. So, for every $m = 1, \dots, k$, we have that $x_m \in$

$Y \cup Z$. But the set $Y \cup Z$ is finite. So there are only finitely many possibilities for the tuple (x_1, \dots, x_k) . Hence, by Remark 3.30, there are only finitely many exemplary components of V^n . \square

Hence, if $f \in \overline{\mathbb{Q}}(j)$ satisfies Condition 3.15, then Theorem 2.3 in this case follows from Proposition 3.16. Indeed, Theorem 3.7 implies that if $f \in \mathbb{C}(j)$ satisfies the divisor condition, then there are only finitely many exemplary components of V^n (something which does not follow from Theorem 2.3, where it was required that V be defined over $\overline{\mathbb{Q}}$).

3.6. An extension to finite rank

In this section we prove Theorem 3.8. As in Section 3.4, we will prove a stronger conditional result under the assumption of Condition 3.15. Theorem 3.8 then follows from Proposition 3.32 as it corresponds to those cases where the divisor condition holds and Condition 3.15 is thus known via Theorem 3.5.

Proposition 3.32. *Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function such that $f \in \overline{\mathbb{Q}}(j)$. Assume Condition 3.15 holds for f . Let $n \geq 1$ and $\Gamma \leq \mathbb{G}_m$ be of finite rank. Then there exist only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent.*

For the remainder of this section, fix such a function f and such a group Γ . If Γ has rank 0, then the result follows immediately from Proposition 3.16. So we may assume that Γ has positive rank. As usual, we write $f(z) = R(j(z))$ for R some rational function with algebraic coefficients. Let K be a number field containing the coefficients of R .

The group Γ is of finite rank, so there exists $\Gamma_0 \leq \Gamma$ finitely generated such that for every $\gamma \in \Gamma$ there exists $m \geq 1$ such that $\gamma^m \in \Gamma_0$. Thus every Γ -dependent tuple is also Γ_0 -dependent. Therefore, we may and do assume that Γ is finitely generated.

Further, f -special points are algebraic. So if $\prod \sigma_i^{a_i} = \gamma \in \Gamma$ for some f -special points σ_i and $a_i \in \mathbb{Z}$, then $\gamma \in \Gamma \cap \overline{\mathbb{Q}}$. Now $\Gamma \cap \overline{\mathbb{Q}}$ is a subgroup of the finitely generated abelian group Γ , so is itself finitely generated. Replacing Γ with $\Gamma \cap \overline{\mathbb{Q}}$ as necessary we may and do assume as well that Γ is generated by algebraic elements.

We thus consider $\Gamma = \langle b_1, \dots, b_k \rangle$, where $k \geq 1$ and every $b_i \in \overline{\mathbb{Q}}$. A tuple of f -special points $(\sigma_1, \dots, \sigma_n)$ is then Γ -dependent if some relation

$$\prod_{i=1}^n \sigma_i^{a_i} = \prod_{i=1}^k b_i^{\alpha_i}$$

holds with $a_i, \alpha_i \in \mathbb{Z}$ and the a_i not all zero. We may also assume that the set $\{b_1, \dots, b_k\}$ is multiplicatively independent; in particular, no b_i is a root of unity. We let $K_0 = K(b_1, \dots, b_k)$.

Now fix $n \geq 1$. Let $X = X_{n,n+k} = Y(1)^n \times \mathbb{G}_m^{n+k}$, and

$$V_\Gamma = \{(x_1, \dots, x_n, t_1, \dots, t_n, b_1, \dots, b_k) \in X : t_i = R(x_i) \text{ for } i = 1, \dots, n\}.$$

We refer to Section 2.1 for the definition of a (weakly) special subvariety of X .

Now for the proof of Proposition 3.32. Constants c, c' will be positive and depend only on f, n , and Γ , but will vary between occurrences. We note that the proofs of Lemmas 3.21 and 3.23 both generalise straightforwardly when we replace K with K_0 in their statements.

PROOF OF PROPOSITION 3.32. For an n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of f -special points, define, as in the proof of Proposition 3.16, the complexity $\Delta(\sigma)$ of σ to be $\Delta(\sigma) = \max\{|\Delta(\sigma_1)|, \dots, |\Delta(\sigma_n)|\}$. Call an n -tuple $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, a Γ -tuple. We show that there are only finitely many Γ -tuples.

Suppose not. Then there are Γ -tuples σ of arbitrarily large complexity $\Delta(\sigma)$. Fix $\nu = (\nu_1, \dots, \nu_k) \in F_e^k$ a preimage of (b_1, \dots, b_k) . We let

$$Y = \{(z, u\nu, r, s) \in F_j^n \times F_e^{n+k} \times \mathbb{R}^{n+k} \times \mathbb{R} : R(j(z)) = e(u), r \cdot (u\nu) = s\}.$$

Then let Z be the projection of Y to $F_j^n \times \mathbb{R}^{n+k} \times \mathbb{R}$. Note that Y, Z are definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$.

Fix a Γ -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of complexity $\Delta = \Delta(\sigma)$. We may write $\sigma_i = R(x_i)$ for some j -special point x_i . The tuple σ satisfies some multiplicative relation

$$\prod_{i=1}^n \sigma_i^{a_i} \prod_{i=1}^k b_i^{\alpha_i} = 1$$

for $a_i, \alpha_i \in \mathbb{Z}$, with the a_i not all zero. The point

$$\hat{\sigma} = (x_1, \dots, x_n, \sigma_1, \dots, \sigma_n, b_1, \dots, b_k) \in V_\Gamma$$

has a preimage $(z, u\nu) \in F_j^n \times F_{\text{exp}}^{n+k}$. This preimage gives rise to a point $(z, u\nu, \beta, \beta', \delta) \in Y$. Here the β, β', δ coordinates are rational integers, with the β_i not all zero, which record the multiplicative dependence of the σ_i, b_i ; that is,

$$\sum_{i=1}^n \beta_i u_i + \sum_{i=1}^k \beta'_i \nu_i = \delta.$$

We may apply Lemma 3.22 to some minimally multiplicatively dependent set S , where $\{\sigma_1, \dots, \sigma_n\} \subset S \subset \{\sigma_1, \dots, \sigma_n, b_1, \dots, b_k\}$. The integers β_i, β'_i may thus be chosen such that

$$|\beta_i| \leq cd^{n+k}(\log d) \prod_{\substack{j=1 \\ j \neq i}}^n h(\sigma_j)$$

and

$$|\beta'_i| \leq cd^{n+k}(\log d) \prod_{j=1}^n h(\sigma_j).$$

Here $d \geq 2$ is the degree of a number field containing $\sigma_1, \dots, \sigma_n, b_1, \dots, b_k$. (We absorb the dependency on the logarithmic heights of the b_i into our constant.) Using Lemmas 3.19 and 3.21, we may thus bound $|\beta_i|$ and $|\beta'_i|$ by $c\Delta^{n(n+k)}$. Observe that for $u_i, \nu_i \in F_e$, the real part of u_i, ν_i is bounded by 1 in absolute value. Then, using the relation

$$\sum_{i=1}^n \beta_i u_i + \sum_{i=1}^k \beta'_i \nu_i = \delta,$$

we may also bound $|\delta|$ by $c\Delta^{n(n+k)}$. Since $\beta_i, \beta'_i, \delta$ are rational integers, their heights are also bounded by $c\Delta^{n(n+k)}$ since $H(l) = |l|$ for $l \in \mathbb{Z} \setminus \{0\}$.

The point $(z, u, \nu, \beta, \beta', \delta)$ projects to a point $(z, \beta, \beta', \delta) \in Z$, which is quadratic in the F_j coordinates and integral in the \mathbb{R} coordinates. The height of the z coordinates may be bounded by $c\Delta$ thanks to Lemma 3.20. Combining this with the height bounds from the previous paragraph, we see that the height of the point $(z, \beta, \beta', \delta)$ is thus bounded by $c\Delta^{n(n+k)}$. The Γ -tuple σ thus gives rise to a quadratic point of Z with height bounded by $c\Delta^{n(n+k)}$.

By the aforementioned generalisation of Lemma 3.21 with $\epsilon = 1/4$, a Γ -tuple σ of complexity Δ has $\geq c\Delta^{1/4}$ distinct conjugates over K_0 , provided Δ is large enough. By the similar generalisation of Lemma 3.23 and the fact that $b_1, \dots, b_k \in K_0$, any K_0 -conjugate of such a Γ -tuple is again a Γ -tuple and is also of complexity Δ . In particular, the K_0 -conjugates all satisfy the same Γ -dependence.

Therefore, each K_0 -conjugate $\tilde{\sigma}$ of σ also gives rise, in the same way as σ , to a quadratic point $(\tilde{z}, \beta, \beta', \tilde{\delta})$ of Z with height bounded by $c\Delta^{n(n+k)}$. Note that that the β, β' coordinates are the same for each conjugate. Therefore, a Γ -tuple σ of sufficiently large complexity Δ gives rise to at least $c'T^{1/4n(n+k)}$ quadratic points on Z with height at most $T = c\Delta^{n(n+k)}$, each corresponding to a distinct Γ -tuple.

The proof now proceeds as for Proposition 3.16. We will only sketch the remainder of the proof, and refer to the proof of Proposition 3.16 for details. Denote by $Y_{(\beta_1, \dots, \beta_n, \beta'_1, \dots, \beta'_k)}$ the fibre of Y over $(\beta_1, \dots, \beta_n, \beta'_1, \dots, \beta'_k)$. We will apply the Counting Theorem in the form of [39, Corollary 7.2] (Theorem A.21) to $Y_{(\beta_1, \dots, \beta_n, \beta'_1, \dots, \beta'_k)}$.

Arguing as in the proof of Proposition 3.16, if the complexity Δ of σ is suitably large, then we obtain a complex algebraic subvariety $W \subset \mathbb{C}^{n+1}$ and a complex-analytically irreducible component $A \subset (\mathbb{H}^n \times \mathbb{C}) \cap W$ such that A has non-constant projection to its \mathbb{H}^n coordinates, A contains a point of $Z_{(b_1, \dots, b_n)}$ which corresponds to a K_0 -conjugate of σ , and $A \subset \tilde{V}_\Gamma$ where

$$\tilde{V}_\Gamma = \{(z, t) \in \mathbb{H}^n \times \mathbb{C} : \prod_{i=1}^n f(z_i)^{\beta_i} \prod_{i=1}^k b_i^{\beta'_i} = e(t)\}.$$

By Ax–Schanuel, we thus obtain that there are weakly special subvarieties $W_1 \subset \mathbb{H}^n$ and $W_2 \subset \mathbb{C}$ such that A is contained in $W_1 \times W_2$ and $W_1 \times W_2 \subset \tilde{V}_\Gamma$. As in the proof of Proposition 3.16, one must have that $W_2 = \{\delta'\}$ for some $\delta' \in \mathbb{Z}$.

Therefore, W_1 is a positive-dimensional weakly special subvariety of \mathbb{H}^n which is contained in the set

$$\{z \in \mathbb{H}^n : \prod_{i=1}^n f(z_i)^{\beta_i} \prod_{i=1}^k b_i^{\beta'_i} = 1\}.$$

The weakly special subvariety W_1 contains a preimage (τ_1, \dots, τ_n) of some Γ -tuple $(f(\tau_1), \dots, f(\tau_n))$. In particular, $f(W_1)$ has no two identically equal coordinates, since the $f(\tau_i)$ are pairwise distinct. In addition, all the β_i are non-zero, since the set $\{f(\tau_1), \dots, f(\tau_n)\}$ is minimally Γ -dependent. Taking the image of W_1 under f , we therefore obtain a multiplicative dependence modulo constants among some pairwise distinct $\mathrm{GL}_2^+(\mathbb{Q})$ -translates of f . This contradicts Condition 3.15, and so we are done. \square

We now establish the analogue of Proposition 3.27, which shows that we do not need to assume that f satisfies Condition 3.15 for the $n = 1$ case.

Proposition 3.33. *Let $f \in \overline{\mathbb{Q}}(j)$ be a non-constant modular function. Let $\Gamma \leq \mathbb{G}_m$ be a finite rank subgroup. Then there are only finitely many f -special points which are Γ -dependent.*

PROOF. In contrast to roots of unity, elements of Γ are not necessarily of bounded height. We therefore cannot argue as in the proof of Proposition 3.27. Instead, we adapt the proof of Proposition 3.32.

The only point in this proof where Condition 3.15 is used is at the end, where it is used to exclude the possibility of a multiplicative dependence modulo constants among some pairwise distinct $\mathrm{GL}_2^+(\mathbb{Q})$ -translates of f . In the $n = 1$ case though, this multiplicative dependence modulo constants involves only a single $\mathrm{GL}_2^+(\mathbb{Q})$ -translate of f . Such a dependence contradicts already the fact that the function f is non-constant, and so we are done. \square

One would like to extend the results of this section to cover the case where the rational function R is not defined over $\overline{\mathbb{Q}}$. We can do this, following the same approach as in Subsection 3.4.4, provided that the finite rank subgroup $\Gamma \leq \mathbb{G}_m$ is contained in $\overline{\mathbb{Q}}$.

Theorem 3.34. *Let $f \in \mathbb{C}(j)$ be a non-constant modular function. Suppose that f satisfies the divisor condition. Let $\Gamma \leq \mathbb{G}_m$ be a finite rank subgroup such that $\Gamma \subset \overline{\mathbb{Q}}$. Then, for each $n \geq 1$, there are only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent and minimal for this property.*

PROOF. Clearly, we may assume without loss of generality that Γ is finitely generated. Suppose then that $\Gamma = \langle b_1, \dots, b_k \rangle$ for some $b_1, \dots, b_k \in \overline{\mathbb{Q}}$. The approach is then almost exactly as for the proof of Theorem 3.7 in Subsection 3.4.4. In Subsection 3.4.4, replace everywhere the words “multiplicatively dependent” by “ Γ -dependent”, and appeal to Theorem 3.8 rather than Theorem 3.7. One must also take the number field K (over which the variety W is defined) to contain b_1, \dots, b_k as well. Everything is then the same as in the previous proof, except that the polynomial equalities $(*)$ will now contain the elements b_1, \dots, b_k as well. \square

Finally, we observe that the divisor condition is not necessary in the case that $n = 1$.

Proposition 3.35. *Let $f \in \mathbb{C}(j)$ be a non-constant modular function. Let $\Gamma \leq \mathbb{G}_m$ be a finite rank subgroup such that $\Gamma \subset \overline{\mathbb{Q}}$. Then there are only finitely many f -special points which are Γ -dependent.*

PROOF. The same as for Theorem 3.34, but using Proposition 3.33 (which does not require the divisor condition), in place of Theorem 3.8. \square

Remark 3.36. We could have proved Proposition 3.27 in an analogous way to how we proved Propositions 3.33 and 3.35. That is, one may show that Condition 3.15 is not required for the proof of Proposition 3.16 when $n = 1$, and then argue that the specialisation argument to deduce Theorem 3.7 from Proposition 3.16 does not require the divisor condition for the $n = 1$ case either. We however prefer the direct approach of Proposition 3.27, which avoids all o-minimal machinery.

Remark 3.37. The assumption that $\Gamma \subset \overline{\mathbb{Q}}$ in Theorem 3.34 and Proposition 3.35 is restrictive. The difficulty arises due to the fact that if Γ contained a transcendental element γ , then a Γ -dependency of f -special points involving γ would not necessarily be preserved under specialisation of the coefficients of R .

CHAPTER 4

Atypical intersections and multiplicative dependence

4.1. The Zilber–Pink conjecture

In this chapter, we relate the results of the previous two chapters to the Zilber–Pink conjecture. This is a major conjecture about the Diophantine properties of mixed Shimura varieties. Recall that each mixed Shimura variety comes with a distinguished collection of subvarieties, called the special subvarieties.

Definition 4.1. Let X be a mixed Shimura variety and denote by \mathcal{S}_X its collection of special subvarieties. Let $V \subset X$ be an algebraic subvariety. A subvariety $A \subset V$ is called an atypical component (of V in X) if there is a special subvariety $T \in \mathcal{S}_X$ such that $A \subset V \cap T$ and

$$\dim A > \dim V + \dim T - \dim X.$$

Given such X, V, T , one would expect, by “counting conditions”, that a component of the intersection $V \cap T$ would have dimension equal to $\dim V + \dim T - \dim X$. An atypical component is therefore a component of such an intersection which has dimension strictly larger than this expected dimension.

The Zilber–Pink conjecture is the following statement. Versions of this conjecture were proposed separately by Zilber [85], by Pink [70], and by Bombieri, Masser, and Zannier [15].

Conjecture 4.2 (Zilber–Pink conjecture, version 1). *Let X be a mixed Shimura variety and $V \subset X$ an algebraic subvariety. Then there are only finitely many maximal atypical components of V .*

If X is a mixed Shimura variety and \mathcal{S}_X its collection of special subvarieties, then $X \in \mathcal{S}_X$ and \mathcal{S}_X is closed under taking irreducible components of intersections. We may therefore make the following definition.

Definition 4.3. Let X be a mixed Shimura variety and $A \subset X$ be a subvariety. Denote by $\langle A \rangle$ the smallest special subvariety of X which contains A . The defect of A is defined

$$\delta(A) = \dim \langle A \rangle - \dim A.$$

Fix a mixed Shimura variety X and a subvariety $V \subset X$. Observe that V is a special subvariety if and only if $\delta(V) = 0$. A subvariety $A \subset V$ is an atypical component (of V in X) if and only if $\delta(A) < \dim X - \dim V$. In particular, V itself is an atypical component if and only if $\langle V \rangle \subsetneq X$.

Definition 4.4. Let X be a mixed Shimura variety and $V \subset X$ a subvariety. A subvariety $A \subset V$ is an optimal subvariety (of V in X) if A is maximal for its defect, i.e. there is no subvariety $A' \subset V$ such that $A \subsetneq A'$ and $\delta(A') \leq \delta(A)$.

Remark 4.5. Observe that V is always an optimal subvariety of itself.

Proposition 4.6. *Let X be a mixed Shimura variety and $V \subset X$ a subvariety.*

- (1) *If $A \subset V$ is a proper optimal subvariety of V in X , then A is an atypical component of V in X .*

(2) *If $A \subset V$ is a maximal atypical component of V in X , then A is an optimal subvariety of V in X .*

PROOF. Let X be a mixed Shimura variety and $V \subset X$ a subvariety. Suppose first that $A \subset V$ is a proper optimal subvariety. Then $\delta(A) < \delta(V)$. Hence,

$$\dim\langle A \rangle - \dim A < \dim\langle V \rangle - \dim V,$$

and so

$$\begin{aligned} \dim A &> \dim V + \dim\langle A \rangle - \dim\langle V \rangle \\ &\geq \dim V + \dim\langle A \rangle - \dim X. \end{aligned}$$

Since certainly $A \subset V \cap \langle A \rangle$ and $\langle A \rangle$ is a special subvariety, it follows immediately that A is atypical.

Now let $A \subset V$ be a maximal atypical component. Since A is atypical, we have that

$$\delta(A) < \dim X - \dim V.$$

Suppose that $B \supset A$ is such that $\delta(B) \leq \delta(A)$. Then

$$\delta(B) \leq \delta(A) < \dim X - \dim V,$$

and so B is atypical. Hence, we must have that $A = B$, because A is a maximal atypical component. But then A is maximal for its defect, and therefore optimal. \square

The Zilber–Pink conjecture may be reformulated in terms of optimal subvarieties.

Conjecture 4.7 (Zilber–Pink conjecture, version 2). *Let X be a mixed Shimura variety and $V \subset X$ an algebraic subvariety. Then there are only finitely many optimal subvarieties of V in X .*

These two versions of the conjecture are in fact equivalent.

Proposition 4.8 ([39, Lemma 2.7]). *Conjectures 4.2 and 4.7 are equivalent.*

Let X be a mixed Shimura variety. It is clear from Proposition 4.6 (2) that Conjecture 4.7 implies Conjecture 4.2. For the converse, one needs to argue by induction on $\dim V$. In particular, for a fixed subvariety $V \subset X$, the finiteness of the optimal subvarieties of V implies the finiteness of the maximal atypical components of V , but the converse does not hold in general.

4.2. Multiplicatively dependent images of special points

Let Y be a modular or Shimura curve and $V \subset Y \times \mathbb{G}_m$ a correspondence (not necessarily defined over $\overline{\mathbb{Q}}$). For $n \geq 1$, set $X_n = Y^n \times \mathbb{G}_m^n$. Then X_n is a mixed Shimura variety and $V^n \subset X_n$ a subvariety. The special subvarieties of X_n were described in Chapter 2. In the remainder of this chapter, we show that the Zilber–Pink conjecture for V^n is intimately related to the results on multiplicatively dependent V -images of special points contained in the previous two chapters. Some of the results in this chapter appeared previously in the author’s paper [32].

Suppose that $(s, x) \in V^n$ is such that $s_1, \dots, s_n \in Y$ are special points and $x_1, \dots, x_n \in \mathbb{G}_m$ are multiplicatively dependent. So $x_1^{a_1} \cdots x_n^{a_n} = 1$ for

some $a_i \in \mathbb{Z}$, not all 0. Then

$$(s, x) \in V^n \cap T,$$

where $T \subset X_n$ is the special subvariety of X

$$T = \{(u_1, \dots, u_n, t_1, \dots, t_n) : u_i = s_i \text{ for } i = 1, \dots, n \text{ and } t_1^{a_1} \cdots t_n^{a_n} = 1\}.$$

Observe that $\dim T = n - 1$. Hence,

$$\dim V^n + \dim T - \dim X_n = n + (n - 1) - 2n = -1,$$

and thus $\{(s, x)\}$ is an atypical component of V^n in X .

Multiplicatively dependent images of special points thus give rise to atypical components of V^n . The resulting components are not necessarily maximal atypical or optimal though, and so the finiteness of multiplicatively dependent n -tuples of images of special points for a given $n \geq 1$ would not follow from the Zilber–Pink conjecture alone.

Now recall from Chapter 2 the notion of an exemplary component of V^n . In particular, we saw there that every multiplicative dependence among V -images of special points was contained in an exemplary component of V^n . We have the following result.

Proposition 4.9. *An exemplary component of V^n is an optimal subvariety.*

PROOF. The proof is identical to [63, Proposition 3.2], which treats the elliptic curve case. We denote by π_1, π_2 the projections of X_n to Y^n and \mathbb{G}_m^n respectively. Suppose that W is an exemplary component of V^n . Then there are special subvarieties $S \subset Y^n$ and $B \subset \mathbb{G}_m^n$ such that $W \subset$

$V^n \cap (S \times B)$ and $\pi_1(W) = S$ and $\langle \pi_2(W) \rangle = B$. Observe that

$$\langle W \rangle = S \times B.$$

Since V is a correspondence, we have that $\dim W = \dim S$. Therefore,

$$\delta(W) = \dim \langle W \rangle - \dim W = \dim B.$$

Suppose now that $W' \subset V^n$ is a subvariety such that $W \subset W'$ and $\delta(W') \leq \delta(W)$. We may write

$$\langle W' \rangle = S' \times B',$$

for some special subvarieties $S' \subset Y^n$ and $B' \subset \mathbb{G}_m^n$. Note that $\dim W' \leq \dim S'$ since V is a correspondence. Hence,

$$\begin{aligned} \dim B' &\leq \dim S' + \dim B' - \dim W' \\ &= \dim \langle W' \rangle - \dim W' \\ &= \delta(W') \\ &\leq \delta(W) = \dim B \leq \dim B', \end{aligned}$$

where the last inequality follows since $B \subset B'$. Hence, equality holds throughout. In particular, $B = B'$ and $\dim S' = \dim W'$. Therefore, W' is a distinguished component of V^n which contains W and projects into B . Since W is exemplary, we must then have that $W = W'$. Therefore, W is maximal for its defect, and thus optimal. \square

Consequently, for each $n \geq 1$, the finiteness of exemplary components of V^n would follow from the finiteness of optimal components of V^n . In particular, Theorem 2.3 would follow from the Zilber–Pink conjecture. (Indeed, by assuming the full conjecture, one could also remove the assumption that V is defined over $\overline{\mathbb{Q}}$ from the statement of Theorem 2.3.)

In general though, the Zilber–Pink conjecture remains open. Indeed, we can prove the finiteness of optimal components of V^n only when $n = 1$. In the remainder of this section, we describe the situation when $n \leq 2$.

Proposition 4.10. *There are only finitely many optimal components of V .*

PROOF. V itself is always optimal. Since V dominates both Y and \mathbb{G}_m , we have that $\langle V \rangle = X_1$. Hence, $\delta(V) = 1$ and V is not atypical. Any proper optimal component of V must therefore have defect 0, and thus be a special subvariety contained in V . So the only proper optimal components of V come from the special points $(s, x) \in V$. The finiteness of these was established by Proposition 2.17. \square

Definition 4.11. A point $(x_1, x_2, \zeta_1, \zeta_2) \in V^2$ is called a modular–torsion tuple of V^2 if ζ_1, ζ_2 are roots of unity and there exists a proper strongly special subvariety $M \subset Y^2$ such that $(x_1, x_2) \in M$.

Proposition 4.12. *Suppose that:*

- (1) *there are only finitely many exemplary components of V^2 ;*
- (2) *there are only finitely many modular–torsion tuples of V^2 which are not contained in a dependent exemplary component of V^2 .*

Then there are only finitely many optimal components of V^2 .

The restriction in (2) to dependent exemplary components is needed to avoid triviality, since V^2 itself is a (non-dependent) exemplary component which contains every modular–torsion tuple of V^2 .

PROOF. Note that V^2 is an optimal component with $\delta(V^2) = 2$. Recall that any proper optimal component of V^2 is atypical. An optimal component of V^2 of dimension 1 must therefore be contained in the intersection of V^2 with a special subvariety of codimension 2.

Suppose $T \subset X_2$ is a special subvariety of codimension 2. If both conditions defining T apply to the Y^2 (respectively, \mathbb{G}_m^2) factor, then V^2 cannot intersect T atypically. So we may suppose that $T = T_1 \times T_2$, where $T_1 \subset Y^2$ and $T_2 \subset \mathbb{G}_m^2$ are each special subvarieties of codimension 1.

If T_1, T_2 are both defined by constant conditions, then $V^2 \cap T$ is atypical in dimension only if either $T = \{s\} \times Y \times \{x\} \times \mathbb{G}_m$ or $T = Y \times \{s\} \times \mathbb{G}_m \times \{x\}$ for some special point $(s, x) \in V$. By the proof of Proposition 4.10, there are only finitely many such special points. Hence, there are only finitely many optimal components of this form, and these all have defect 1.

If one of T_1, T_2 is defined by a constant coordinate and the other by a relation, then $V^2 \cap T$ will have dimension 0, and hence not be atypical. So we are left with the case that T_1, T_2 are both defined by relations. In this case, a component of $V \cap T$ is atypical if it has dimension 1, in which case it is an exemplary component of V^2 . So all the optimal components of this form are exemplary components and they all have defect 1.

So we are left with the optimal components of dimension 0, i.e. the optimal points. These must have defect either 0 or 1. If they have defect

0, then they are special points of V^2 , and so finiteness follows from the proof of Proposition 4.10 again.

So we just need to identify the optimal points of V^2 which have defect 1 (and so are not special). Let (x_1, x_2, t_1, t_2) be such a point. Then (x_1, x_2, t_1, t_2) must be contained in the intersection of V^2 with a special subvariety of codimension 3. Suppose then that $T \subset X_2$ is a special subvariety of codimension 3 such that $(x_1, x_2, t_1, t_2) \in V^2 \cap T$.

Either T is defined by two Y^2 conditions and one \mathbb{G}_m^2 condition, or T is defined by one Y^2 condition and two \mathbb{G}_m^2 conditions. We may assume that the two conditions of the same type are both fixed coordinates. If the third condition is also a fixed coordinate, then (x_1, x_2, t_1, t_2) cannot be optimal, since it is contained in one of the above-identified positive-dimensional optimal components of defect 1. So we may suppose that the third condition is a relation.

If the relation applies to the \mathbb{G}_m^2 coordinates, then neither of t_1, t_2 is a root of unity (since the relation implies that one is a root of unity if and only if the other is). So either (x_1, x_2, t_1, t_2) is an exemplary component of V^2 , or there is a strongly special subvariety $M \subset Y^2$ which contains (x_1, x_2) and whose V -image satisfies the same multiplicative relation as t_1, t_2 . In the latter case, there is thus an exemplary component of defect 1 containing the point (x_1, x_2, t_1, t_2) , and so (x_1, x_2, t_1, t_2) is not optimal. So we may assume that (x_1, x_2, t_1, t_2) is an exemplary component of V^2 . So every optimal point of this form gives rise to an exemplary component of V^2 .

So we reduce to the case that the relation applies to the Y^2 coordinates (and so t_1, t_2 are roots of unity). In this case, (x_1, x_2, t_1, t_2) is a modular-torsion tuple of V^2 . If (x_1, x_2, t_1, t_2) is contained in a dependent exemplary component, then this must be one of the exemplary components coming from a relation on each side. Such a component has defect 1, and so the point (x_1, x_2, t_1, t_2) is not optimal in this case. So any optimal point of this form is a modular-torsion tuple of V^2 which is not contained in a dependent exemplary component of V^2 .

We have thus seen that the optimal components V^2 are the exemplary components of V^2 (which include V^2 itself) and the modular-torsion tuples of V^2 which are not contained in a dependent exemplary component of V^2 . The proposition thus follows. \square

Observe that (1) of Proposition 4.12 holds if V is defined over $\overline{\mathbb{Q}}$, thanks to Theorem 2.3. In this case, establishing the finiteness of optimal components of V^2 therefore reduces to showing that (2) holds. In general, we cannot prove (2) though. For some special cases, we can prove (2). Thus, we are able to establish the Zilber–Pink conjecture in these cases. This we do in the next section.

4.3. Multiplicatively dependent f -special points

We now restrict to the setting of Chapter 3. So, throughout this section $V \subset Y(1) \times \mathbb{G}_m$ is the graph of a non-constant rational function R and, for $n \geq 1$, we set $X_n = Y(1)^n \times \mathbb{G}_m^n$. Define the modular function $f: \mathbb{H} \rightarrow \mathbb{C}$ by $f(z) = R(j(z))$.

Lemma 4.13. *Suppose that f satisfies Condition 3.15. Let $n \geq 1$. An n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that $\{\sigma_1, \dots, \sigma_n\}$ is a minimally multiplicatively dependent set gives rise to an atypical point $\hat{\sigma} \in V^n$ which is not contained in any atypical component of V^n of positive dimension.*

PROOF. As we saw in the previous section, the tuple σ gives rise to an atypical point $\hat{\sigma}$ of V^n . We now show that $\hat{\sigma}$ cannot be contained in an atypical component of V^n of positive dimension. The proof is similar to [65, Lemma 6.1].

If $\hat{\sigma}$ were contained in a special subvariety of X_n defined by two independent multiplicative conditions on the σ_i coordinates, then we could eliminate one of these coordinates, contradicting the fact that the set $\{\sigma_1, \dots, \sigma_n\}$ is minimally multiplicatively dependent.

A special subvariety $M \times \mathbb{G}_m^n$, where M is a special subvariety of $Y(1)^n$, never intersects V^n atypically. Similarly, no special subvariety of the form $Y(1)^n \times T$, where T is a special subvariety of \mathbb{G}_m^n , intersects V^n atypically.

It thus remains to consider special subvarieties of the form $M \times T$, where M is a proper special subvariety of $Y(1)^n$ and T is a special subvariety of \mathbb{G}_m^n defined by one multiplicative condition. Then $V^n \cap (M \times T)$ is equal to the set

$$\{(u_1, \dots, u_n, R(u_1), \dots, R(u_n)) : (u_1, \dots, u_n) \in M, (R(u_1), \dots, R(u_n)) \in T\}.$$

This would typically have dimension $\dim M - 1$. To be atypical, we would thus require that $R(M) \cap \mathbb{G}_m^n \subset T$. Hence, if $V^n \cap (M \times T)$ were also positive-dimensional, then Condition 3.15 and the fact that no proper subset of

$\{\sigma_1, \dots, \sigma_n\}$ is multiplicatively dependent implies that $R(M)$ must have two identically equal coordinates. But then the σ_i cannot be pairwise distinct, a contradiction. \square

Therefore, provided f satisfies Condition 3.15, the finiteness of n -tuples of distinct f -special points that are multiplicatively dependent and minimal for this property would follow from the Zilber–Pink conjecture.

We now consider the Zilber–Pink conjecture for V^n when $n \leq 2$.

Proposition 4.14. *There are only finitely many optimal components of V .*

PROOF. This is just Proposition 4.10 in this setting. \square

Now for the $n = 2$ case. Recall that Proposition 4.12 showed that the optimal components of V^2 comprised the exemplary components of V^2 and the modular–torsion tuples of V^2 not contained in dependent exemplary components of V^2 .

Definition 4.15. We say that f satisfies the finiteness condition for pairs if there are only finitely many 2-tuples (σ_1, σ_2) of distinct f -special points σ_i such that the set $\{\sigma_1, \sigma_2\}$ is multiplicatively dependent and minimal for this property.

Note that the finiteness condition for pairs holds by Theorem 3.7 if f satisfies the divisor condition and by Proposition 3.16 if $f \in \overline{\mathbb{Q}}(j)$ satisfies Condition 3.15.

Proposition 4.16. *Suppose that f satisfies Condition 3.15 and the finiteness condition for pairs. Then there are only finitely many exemplary components of V^2 .*

Remark 4.17. If V is defined over $\overline{\mathbb{Q}}$, then the conclusion of Proposition 4.16 holds (unconditionally) by Theorem 2.3.

PROOF. This follows from Proposition 3.31, thanks to the finiteness condition for pairs and Proposition 3.27. \square

Proposition 4.12 and Theorem 3.29 together imply the following.

Proposition 4.18. *Suppose that f satisfies Condition 3.15 and the finiteness condition for pairs. Suppose also that there are only finitely many modular-torsion tuples $(x_1, x_2, \zeta_1, \zeta_2)$ of V^2 such that $x_1 \neq x_2$ and neither of x_1, x_2 is j -special. Then there are only finitely many optimal components of V^2 .*

In the remainder of this section, we prove the necessary finiteness statement for modular-torsion tuples in the case that f satisfies Condition 3.15 and also $f \in \overline{\mathbb{Q}}(j)$.

We say that $(x_1, x_2) \in Y(1)^2$ satisfies a modular relation if $\Phi_N(x_1, x_2) = 0$ for some $N \geq 1$. Recall from Section 1.3 that Φ_N denotes the N th classical modular polynomial. We note that (x_1, x_2) satisfies a modular relation if and only if $(x_1, x_2) = (j(z), j(gz))$ for some $z \in \mathbb{H}$, $g \in \mathrm{GL}_2^+(\mathbb{Q})$. In particular, $(x_1, x_2) \in M$ for a proper strongly special subvariety $M \subset Y(1)^2$ if and only if (x_1, x_2) satisfies a modular relation. See [9, §2].

Proposition 4.19. *Suppose that f satisfies Condition 3.15 and $f \in \overline{\mathbb{Q}}(j)$. Then there are only finitely many modular-torsion tuples $(x_1, x_2, \zeta_1, \zeta_2) \in V^2$ such that $x_1 \neq x_2$ and neither x_1 nor x_2 is j -special.*

We need the following lemma.

Lemma 4.20. *For every $\alpha > 1$, there exists a constant $c = c(\alpha) > 0$ such that, for each modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2) \in V^2$, one has that $\deg \zeta_1 \leq c(\deg \zeta_2)^\alpha$ and $\deg \zeta_2 \leq c(\deg \zeta_1)^\alpha$.*

PROOF. Let K be a number field over which the rational function R is defined. Since R is a rational function, there is some integer n such that R is at most n -to-1. In this proof, all constants are positive and depend on f, R, K, n ; any other dependencies will be explicitly indicated.

Let

$$Y = \{(\nu, g, u) \in F_j \times \mathrm{GL}_2^+(\mathbb{R}) \times F_e : g\nu \in F_j, R(j(g\nu)) = e(u)\}.$$

The set Y is definable. We view Y as a definable family of fibres

$$Y_\nu = \{(g, u) \in \mathrm{GL}_2^+(\mathbb{R}) \times F_e : g\nu \in F_j, R(j(g\nu)) = e(u)\},$$

where $\nu \in F_j$. In particular, since the Counting Theorem of [66] is uniform in definable families (Theorem A.16), for every $\epsilon > 0$, there exists a constant $c(\epsilon) > 0$ such that, for every $\nu \in F_j$, either $N_1(Y_\nu, T) \leq c(\epsilon)T^\epsilon$ for all $T \geq 1$, or else Y_ν contains a positive-dimensional, connected semialgebraic set. As usual, $N_1(Y_\nu, T)$ denotes the number of rational points in the set Y_ν which have height $\leq T$.

Suppose then that the lemma is false. Fix some $\alpha > 1$ for which it fails. Then, for every $M \geq 1$, there is some modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ with either $\deg \zeta_2 > M(\deg \zeta_1)^\alpha$ or $\deg \zeta_1 > M(\deg \zeta_2)^\alpha$. Fix some suitably large M and such a modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$. Write $d_i = \deg \zeta_i$. Without loss of generality, we may assume that $d_1 \leq d_2$, and so we must have that $d_2 > Md_1^\alpha$. Write m_i for the order of ζ_i . So $d_i = \phi(m_i)$. Let

$\nu \in F_j$ be such that $j(\nu) = x_1$. We show that $(x_1, x_2, \zeta_1, \zeta_2)$ gives rise to a rational point of the set Y_ν in the following way.

We have that $\zeta_2 = e(q)$, where $q = a/m_2$ for some $a \in \mathbb{Z}$ with $0 < a < m_2$ and $\gcd(a, m_2) = 1$. In particular, $H(q) = m_2$. Let E_i be an elliptic curve with j -invariant x_i . The modular relation satisfied by x_1, x_2 implies that the elliptic curves E_1, E_2 are isogenous. We will bound the degree of this isogeny.

First, we bound the degrees of x_1, x_2 . Write

$$R(t) = \frac{p(t)}{q(t)},$$

where $p(t), q(t) \in K[t]$. Let $d = [K : \mathbb{Q}]$ and $l = \max\{\deg p, \deg q\}$. Note that x_i is a root of the non-zero polynomial $f_i(t) = p(t)^{m_i} - q(t)^{m_i}$. The polynomial f_i has degree $\leq lm_i$ and coefficients in K . Hence

$$[\mathbb{Q}(x_i) : \mathbb{Q}] \leq [K(x_i) : \mathbb{Q}] = [K(x_i) : K][K : \mathbb{Q}] \leq lm_i d.$$

So x_i has degree bounded by $c_1 m_i$.

Next, we bound the logarithmic heights of x_1, x_2 . Viewing the rational function R as a morphism $\mathbb{P}^1 \rightarrow \mathbb{P}^1$, one may use [40, Theorem B.2.5] and the fact that $R(x_i) = \zeta_i$ is a root of unity to obtain that $h(x_1), h(x_2) \leq c_2$.

We then use [65, (5.8)] to bound the semistable Faltings height $h_F(E_1)$ of the elliptic curve E_1 . We obtain that

$$h_F(E_1) \leq c_3 \max\{1, h(x_1)\}.$$

In particular, the above bound on $h(x_1)$ thus implies that $h_F(E_1) \leq c_4$. Combining this bound on the Faltings height with the above degree bounds

on x_1, x_2 , we may use [65, (5.10)] to deduce that there is an isogeny between E_1 and E_2 of degree $N \leq c_5 \max\{m_1, m_2\}^5$. Consequently, by [38, Lemma 5.2], there exists $g \in \mathrm{GL}_2^+(\mathbb{Q})$ such that $g\nu \in F_j$, $j(g\nu) = x_2$, and the height of g (regarded as the vector of its entries) is bounded by $c_6 \max\{m_1, m_2\}^{50}$.

The modular-torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ thus gives rise to the rational point $(g, q) \in Y_\nu$, and this point has height $\leq c_7 \max\{m_1, m_2\}^{50}$. Since $\phi(m_i) = d_i$, we may use the elementary lower bound $\phi(m_i) \geq \sqrt{m_i}/2$ to obtain that the height of this rational point is $\leq c_8 d_2^{100}$.

Now consider the conjugates of $(x_1, x_2, \zeta_1, \zeta_2)$ over the field $K(\zeta_1)$. Observe that

$$[K(\zeta_1, \zeta_2) : K(\zeta_1)][K(\zeta_1) : \mathbb{Q}] = [K(\zeta_1, \zeta_2) : \mathbb{Q}] \geq [\mathbb{Q}(\zeta_2) : \mathbb{Q}] = d_2,$$

so

$$\begin{aligned} [K(\zeta_1, \zeta_2) : K(\zeta_1)] &\geq \frac{d_2}{[K(\zeta_1) : \mathbb{Q}]} \\ &= \frac{d_2}{[K(\zeta_1) : \mathbb{Q}(\zeta_1)][\mathbb{Q}(\zeta_1) : \mathbb{Q}]} \\ &\geq \frac{d_2}{[K : \mathbb{Q}] \cdot d_1}. \end{aligned}$$

Recall that $d_2 > Md_1^\alpha$. Hence, there are at least

$$[K(\zeta_1, \zeta_2) : K(\zeta_1)] \geq c_9 M^{1/\alpha} d_2^{(\alpha-1)/\alpha}$$

conjugates of $(x_1, x_2, \zeta_1, \zeta_2)$ over the field $K(\zeta_1)$.

We may enumerate the conjugates as $(x_1^{(i)}, x_2^{(i)}, \zeta_1, \zeta_2^{(i)})$. Note that $R(x_1^{(i)}) = \zeta_1$ for all of these conjugates. Hence, there are at most n distinct

coordinates $x_1^{(i)}$ among these conjugates since R is at most n -to-1. Let $\nu_1, \dots, \nu_k \in F_j$ be such that $j(\nu_1), \dots, j(\nu_k)$ are all distinct and give all the possible $x_1^{(i)}$ coordinates. Note that $k \leq n$.

Now each distinct conjugate $(x_1^{(i)}, x_2^{(i)}, \zeta_1, \zeta_2^{(i)})$ gives rise, in the same way as above, to a distinct rational point which lies on one of the definable sets $Y_{\nu_1}, \dots, Y_{\nu_k}$. Further, these rational points all have height bounded by $c_8 d_2^{100}$ since every $\zeta_2^{(i)}$ is again a root of unity of order m_2 . Thus, there must be at least one $r \in \{1, \dots, k\}$ such that Y_{ν_r} has at least $c_{10} M^{1/\alpha} d_2^{(\alpha-1)/\alpha}$ rational points of height $\leq c_8 d_2^{100}$.

We now apply the above-stated form of the Counting Theorem with $\epsilon = (\alpha - 1)/(100\alpha)$ and $T = c_8 d_2^{100}$. Provided M is suitably large (which we may always assume), there must then exist some $r \in \{1, \dots, k\}$ such that Y_{ν_r} contains a positive-dimensional, connected, semialgebraic set. Fix such an r and let S be the corresponding semialgebraic set.

Define the map $\Theta: \mathrm{GL}_2^+(\mathbb{R}) \times \mathbb{C} \rightarrow \mathbb{H} \times \mathbb{C}$ by $(g, u) \mapsto (g\nu_r, u)$. In particular, the map Θ is semialgebraic. The set $\Theta(S) \subset \mathbb{H} \times \mathbb{C}$ is positive-dimensional and semialgebraic. Observe that $\Theta(S) \subset \pi^{-1}(V)$ since $S \subset Y_{\nu_r}$. Arguing as in the proof of Proposition 3.16, we may find a complex algebraic subvariety $W \subset \mathbb{C}^2$ and a positive-dimensional complex-analytically irreducible component $A \subset (\mathbb{H} \times \mathbb{C}) \cap W$ such that $A \subset \pi^{-1}(V_1)$. Ax–Schanuel then implies that V_1 must contain a positive-dimensional weakly special subvariety. The only weakly special subvarieties of V_1 are points though, because $V_1 = \{(x, t) \in X_1 : R(x) = t\}$ and the map R is non-constant. We thus obtain a contradiction, and the result is proved. \square

We may now prove our finiteness result.

PROOF OF PROPOSITION 4.19. Once again, let K be a number field over which R is defined. In this proof, all constants are positive and may depend possibly on f, R, K ; any other dependencies will be explicitly indicated.

Observe that the proofs of Propositions 4.12 and 4.18 imply that every modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ with $x_1 \neq x_2$ both not j -special gives rise to a maximal atypical component $\{(x_1, x_2, \zeta_1, \zeta_2)\}$ of V^2 . In particular, no such modular–torsion tuple is contained in a positive-dimensional atypical component of V^2 .

We define the complexity Δ of a modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ by $\Delta = \min\{\deg \zeta_1, \deg \zeta_2\}$. The degree bound in Lemma 4.20 and the fact that R is finite-to-one together imply that, for any $C > 0$, there are only finitely many modular–torsion tuples with complexity $\leq C$. Suppose then, for contradiction, that there are infinitely many modular–torsion tuples. In particular, there are modular–torsion tuples of arbitrarily large complexity.

Now let

$$Y = \{(g, z_1, z_2, u_1, u_2) \in \mathrm{GL}_2^+(\mathbb{R}) \times F_j^2 \times F_e^2 : \\ gz_1 = z_2, R(j(z_1)) = e(u_1), R(j(z_2)) = e(u_2)\},$$

and let

$$Z = \{(g, u_1, u_2) : \exists z_1, z_2 (g, z_1, z_2, u_1, u_2) \in Y\}$$

be its projection. Both sets are definable.

Suppose $(x_1, x_2, \zeta_1, \zeta_2)$ is a modular–torsion tuple of complexity Δ . We show that this tuple leads to a rational point on Z of bounded height. Say ζ_1 is a primitive m_1 th root of unity and ζ_2 is a primitive m_2 th root of unity.

So $\deg \zeta_1 = \phi(m_1)$ and $\deg \zeta_2 = \phi(m_2)$. Write $d_i = \deg \zeta_i$. Without loss of generality, we assume that $d_2 \geq d_1$.

We have that $\zeta_1 = e(k/m_1)$ for some $0 \leq k < m_1$ with $\gcd(k, m_1) = 1$ and $\zeta_2 = e(l/m_2)$ for some $0 \leq l < m_2$ with $\gcd(l, m_2) = 1$. In particular, $H(k/m_1) = m_1$ and $H(l/m_2) = m_2$.

Let E_1, E_2 be elliptic curves with j -invariants x_1, x_2 respectively. As in the proof of Lemma 4.20, we may then bound the degrees and logarithmic heights of x_1, x_2 . Once again, we obtain that the degree of an isogeny between E_1, E_2 is bounded by $c_1(\max\{m_1, m_2\})^{c_2}$. Let $\tau_1, \tau_2 \in F_j$ such that $x_i = j(\tau_i)$. By [38, Lemma 5.2], there then exists $g \in \mathrm{GL}_2^+(\mathbb{Q})$ such that $g\tau_1 = \tau_2$ and the height of g (viewed as a vector of its entries) is bounded by $c_3(\max\{m_1, m_2\})^{c_4}$.

The modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ thus gives rise to the rational point $\sigma = (g, k/m_1, l/m_2) \in Z$. The height bound on g , Lemma 4.20, and the elementary estimate that $\phi(m_i) \geq \sqrt{m_i}/2$ together imply that this point σ has height bounded by $c_5\Delta^{c_6}$.

Observe that

$$[K(\zeta_i) : K][K : \mathbb{Q}] = [K(\zeta_i) : \mathbb{Q}] \geq d_i \geq \Delta$$

since $\Delta = \min\{d_1, d_2\}$, and so

$$[K(\zeta_i) : K] \geq c_7\Delta.$$

The modular–torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ thus has $\geq c_7\Delta$ conjugates over K . Each of these conjugates $(x_1^{(r)}, x_2^{(r)}, \zeta_1^{(r)}, \zeta_2^{(r)})$ is again a modular–torsion tuple. In particular, note that each $\zeta_i^{(r)}$ is again a root of unity of order

m_i . Each distinct conjugate $(x_1^{(r)}, x_2^{(r)}, \zeta_1^{(r)}, \zeta_2^{(r)})$ thus gives rise, in the same way as above, to a rational point $(g_r, k_r/m_1, l_r/m_2) \in Z$ which has height bounded by $c_5\Delta^{c_6}$. Moreover, there must also be $\geq c_7\Delta$ distinct coordinates k_r/m_1 and $\geq c_7\Delta$ distinct coordinates l_r/m_2 among the resulting rational points.

We now apply the Counting Theorem in the form of [59, Theorem 3.5] (Theorem A.18), with $\epsilon = 1/2c_6$ say. We may write the resulting basic block families as W_i , where $i = 1, \dots, J$, for some constant J . In particular, the rational points of Z with height at most $T = c_5\Delta^{c_6}$ are contained in $\leq c_8T^{1/2c_6} = c_9\Delta^{1/2}$ basic blocks, each of which is a fibre of one of W_1, \dots, W_J .

The structure $\mathbb{R}_{\text{an,exp}}$ has analytic cell decomposition [30] (see Section A.2). We may thus decompose each of the arising basic block families W_1, \dots, W_J into finitely many analytic cells P_i , in such a way that this induces a decomposition of each fibre of the basic block family over its base.

Recall that a modular–torsion tuple of complexity Δ gives rise, via its K -conjugates, to a collection of rational points (g, u_1, u_2) on Z which are all of height $\leq c_5\Delta^{c_6}$ and among which there $\geq c_7\Delta$ distinct u_1 coordinates and $\geq c_7\Delta$ distinct u_2 coordinates l_i/m_2 . Provided that Δ is suitably large, we thus see that at least one of the analytic cells P_i must contain one of these rational points and have non-constant projection to both its F_e coordinates. Fix such a cell P and such a rational point $p = (g_p, u_{1,p}, u_{2,p}) \in P$.

The cell P is contained in a basic block B that is itself contained in Z . As in the proof of [59, Proposition 3.4(1)], one may thus find an open

neighbourhood $\Omega_{\mathbb{R}} \subset \mathrm{GL}_2^+(\mathbb{R}) \times \mathbb{C}^2$ of p such that the intersection $\Omega_{\mathbb{R}} \cap B$ is a real semialgebraic set, which we denote S . Note that $p \in S$.

Provided that Δ is sufficiently large, we must have that

$$e(u_{1,p}), e(u_{2,p}) \notin \{f(w) : w \in \mathbb{H} \text{ and } f'(w) = 0\}.$$

Therefore, the function f is locally invertible at $e(u_{1,p}), e(u_{2,p})$. Given a sufficiently small open neighbourhood $\Omega \subset \mathrm{GL}_2(\mathbb{C}) \times \mathbb{C}^2$ of p , we may then define an analytic subvariety $A \subset \Omega$ by

$$A = \{(g, e(u_1), e(u_2)) \in \Omega : g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \text{ and} \\ \prod_{h_1, h_2} ((ch_1(e(u_1)) + d)h_2(e(u_2)) - (ah_1(e(u_1)) + b)) = 0\},$$

where the product runs over all the local inverses h_1, h_2 of f at $e(u_{1,p}), e(u_{2,p})$ respectively which hit the fundamental domain F_e . Without loss of generality, we may assume that $\Omega_{\mathbb{R}} = \Omega \cap (\mathrm{GL}_2^+(\mathbb{R}) \times \mathbb{C}^2)$, and thus $S \subset A$.

As in the proof of Proposition 3.16, we may apply the results of [1]. There exists an open neighbourhood $\Omega' \subset \Omega$ of p and a set $Q \subset \Omega'$, where Q may be written as a finite union of irreducible Nash subsets of Ω' which each contain p , such that $p \in \Omega' \cap S \subset Q \subset A$.

If every complex-analytically irreducible component of Q projects constantly to one or more of its F_e coordinates, then the cell P would have constant projection to (at least) one of its F_e coordinates by real analytic continuation. This cannot happen. Hence, there must exist a complex

algebraic subvariety $W \subset \mathrm{GL}_2(\mathbb{C}) \times \mathbb{C}^2$ such that there is a complex-analytically irreducible component $D \subset \Omega' \cap W$ such that $p \in D \subset A$ and D projects non-constantly to both its F_e coordinates.

Denote by $\nu: W^\nu \rightarrow W$ the normalisation of the variety W . There exists an open neighbourhood in $\nu^{-1}(W \cap A)$ of some preimage p^ν of p . We may choose a point q in this neighbourhood such that, writing $\nu(q) = (g, u_1, u_2)$, we have that $u_1 \neq u_{1,p}$ and $u_2 \neq u_{2,p}$. By [21, Corollary 1.9], we may then find an irreducible complex algebraic curve $T \subset W^\nu$ which passes through the points p^ν, q .

The Zariski-closure of the set

$$\{(z, gz, u_1, u_2) : z \in \mathbb{C}, (g, u_1, u_2) \in \nu(T)\}$$

is then a complex algebraic surface, which we denote H . There exists a positive-dimensional complex-analytically irreducible component H_0 of $H \cap \pi^{-1}(V^2)$ which has non-constant projection to both its u_i coordinates and also contains a point $(z, g_p z, u_{1,p}, u_{2,p})$ whose image under π is a modular-torsion tuple.

The Ax–Schanuel results of Section 2.3 then imply that there exists a weakly special subvariety W_0 of U such that $H_0 \subset W_0 \subset \pi^{-1}(V_2)$. Since this weakly special subvariety W_0 has no constant coordinates, it is in fact a special subvariety. It must also have codimension at least two, in order to be contained in $\pi^{-1}(V^2)$. Taking the image under π , we thus obtain a positive-dimensional atypical component of V^2 which contains $\pi(z, g_p z, u_{1,p}, u_{2,p})$. This is a positive-dimensional atypical component of V^2 which contains a modular-torsion tuple $(x_1, x_2, \zeta_1, \zeta_2)$ with $x_1 \neq x_2$

both not j -special. However, we know already that there are no such components. We thus obtain the desired contradiction, and so the proof is complete. \square

From Propositions 3.16, 4.18, and 4.19, one may deduce the following case of the Zilber–Pink conjecture.

Proposition 4.21. *Suppose that f satisfies Condition 3.15 and $f \in \overline{\mathbb{Q}}(j)$. Then there are only finitely many optimal components of V^2 .*

We now return briefly to the setting of Section 4.2. So $V \subset Y \times \mathbb{G}_m$ is an irreducible correspondence. Suppose that V is defined over $\overline{\mathbb{Q}}$. Then Proposition 4.12 and Theorem 2.3 combined show that the finiteness of optimal components of V^2 would follow from the finiteness of the modular–torsion tuples of V^2 which are not contained in an exemplary component.

The proof of Proposition 4.19 suggests a strategy for proving a finiteness statement for modular–torsion tuples. Suppose that $(x_1, x_2, \zeta_1, \zeta_2) \in V^2$ is a modular–torsion tuple. Then $(x_1, x_2) \in M$ for some strongly special subvariety $M \subset Y^2$. The strongly special subvariety M is the image under the uniformisation $q: \mathbb{H} \rightarrow Y$ of a set

$$\{(z, gz) : z \in \mathbb{H}\} \subset \mathbb{H}^2,$$

where g is a rational point (under a suitable embedding) of $\mathrm{GL}_2^+(\mathbb{R})$.

To proceed as in the proof of Proposition 4.19, we would need to be able to bound the height of this rational point g in terms of the degrees of ζ_1, ζ_2 . In the $Y(1)$ case, we did this using isogeny estimates for elliptic curves and bounds for the difference between the Faltings height of an elliptic curve

and the height of its j -invariant. A similar approach may well be possible in the general case.

4.4. The finite rank case

4.4.1. The Zilber–Pink setting. We now put Theorem 3.8 in the context of the Zilber–Pink conjecture.

For $n, k \geq 1$, we set $X = X_{n,n+k} = Y(1)^n \times \mathbb{G}_m^{n+k}$, $U = U_{n,n+k} = \mathbb{H}^n \times \mathbb{C}^{n+k}$, and let $\pi: U \rightarrow X$ be given by

$$\pi(z_1, \dots, z_n, u_1, \dots, u_{n+k}) = (j(z_1), \dots, j(z_n), e(u_1), \dots, e(u_{n+k})).$$

We define (weakly) special subvarieties of U, X as in Section 2.1.

For the remainder of this chapter, we let $f \in \mathbb{C}(j)$ be a non-constant modular function and write $f(z) = R(j(z))$ for some rational function R . We consider the Zilber–Pink conjecture for the family of subvarieties

$$V_{n,k,\bar{b}} = \{(x_1, \dots, x_n, t_1, \dots, t_{n+k}) \in X_{n,n+k} : \\ t_i = R(x_i) \text{ for } i = 1, \dots, n \text{ and } t_{n+i} = b_i \text{ for } i = 1, \dots, k\},$$

where $n, k \geq 1$ and $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$. The relevant statement of the Zilber–Pink conjecture in this context is as follows.

Conjecture 4.22 (Zilber–Pink conjecture). *Let $n, k \geq 1$ and $\bar{b} \in \mathbb{G}_m^k$. Then there are only finitely many maximal atypical components of $V_{n,k,\bar{b}}$.*

Suppose $\Gamma \leq \mathbb{G}_m$ is a subgroup of finite, positive rank. If f satisfies Condition 3.15, then Conjecture 4.22 implies the finiteness of n -tuples of

distinct f -special points that are Γ -dependent and minimal for this property, i.e. Proposition 3.32. To show this, we first prove the following lemma via an easy modification of the proof of Lemma 4.13.

Lemma 4.23. *Suppose that f satisfies Condition 3.15 and $\Gamma \leq \mathbb{G}_m$ is a subgroup of finite, positive rank. Then there exists $k \geq 1$ and $\bar{b} \in \mathbb{G}_m^k$ with the following property: an n -tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is minimally Γ -dependent gives rise to a point $\hat{\sigma} \in V_{n,k,\bar{b}}$ such that $\{\hat{\sigma}\}$ is a maximal atypical component of $V_{n,k,\bar{b}}$.*

PROOF. There is a finitely generated $\Gamma_0 \subset \Gamma$ such that, for every $\gamma \in \Gamma$, there exists $m \geq 1$ with $\gamma^m \in \Gamma_0$. Write $\Gamma_0 = \langle b_1, \dots, b_k \rangle$ for some $b_1, \dots, b_k \in \mathbb{G}_m$. We may assume that the set $\{b_1, \dots, b_k\}$ is multiplicatively independent. Let $n \geq 1$. Set $X = X_{n,n+k}$ and $V = V_{n,k,\bar{b}}$, where $\bar{b} = (b_1, \dots, b_k)$.

Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be an n -tuple of distinct f -special points such that $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent. Then

$$\sigma_1^{a_1} \cdots \sigma_n^{a_n} b_1^{\alpha_1} \cdots b_k^{\alpha_k} = 1$$

for some integers a_i, α_i with the a_i not all zero. There exist j -special points x_1, \dots, x_n such that each $\sigma_i = R(x_i)$. The point

$$\hat{\sigma} = (x_1, \dots, x_n, \sigma_1, \dots, \sigma_n, b_1, \dots, b_k) \in V$$

then lies in the intersection of V with the special subvariety T of X defined by

$$T = \{(\bar{z}, \bar{t}, \bar{u}) : z_i = x_i \text{ for } i = 1, \dots, n \text{ and } t_1^{a_1} \cdots t_n^{a_n} u_1^{\alpha_1} \cdots u_k^{\alpha_k} = 1\}.$$

Here $\text{codim } T = n + 1$ and so $\dim T = n + k - 1$. Thus,

$$\dim V + \dim T - \dim X = n + (n + k - 1) - (2n + k) = -1.$$

So σ gives rise to an atypical component $\{\hat{\sigma}\}$ of V .

Since Condition 3.15 holds for f and the b_i are multiplicatively independent, one may then show, as in the proof of Lemma 4.13, that this atypical point $\hat{\sigma}$ cannot be contained in any atypical component of V of positive dimension. \square

The desired result then follows straightforwardly.

Proposition 4.24. *Suppose Condition 3.15 and Conjecture 4.22 hold for f . Let $\Gamma \leq \mathbb{G}_m$ be a subgroup of finite, positive rank. Then, for each $n \geq 1$, there are only finitely many n -tuples $(\sigma_1, \dots, \sigma_n)$ of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent.*

PROOF. There is a finitely generated $\Gamma_0 \subset \Gamma$ such that, for every $\gamma \in \Gamma$, there exists $m \geq 1$ with $\gamma^m \in \Gamma_0$. Let $\{b_1, \dots, b_k\}$ be a multiplicatively independent set of generators for Γ_0 . Set $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$.

Suppose then that f satisfies Condition 3.15 and Conjecture 4.22. Let $n \geq 1$. Suppose that $\sigma = (\sigma_1, \dots, \sigma_n)$ is an n -tuple of distinct f -special points such that the set $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent, but no proper subset of $\{\sigma_1, \dots, \sigma_n\}$ is Γ -dependent. Then Lemma 4.23 shows that σ gives rise to a maximal atypical component $\{\hat{\sigma}\}$ of the subvariety $V_{n,k,\bar{b}}$. Further, distinct n -tuples σ of this kind give rise, in this way, to distinct points $\hat{\sigma} \in V_{n,k,\bar{b}}$. By Conjecture 4.22, there are only finitely many maximal

atypical components of $V_{n,k,\bar{b}}$. Hence, there are only finitely many such n -tuples. \square

In particular, Theorem 3.8 would follow from Conjecture 4.22. However, we are not able to prove Conjecture 4.22 in general. We can though show that Conjecture 4.22 holds for $V_{1,k,\bar{b}}$ if either $f \in \overline{\mathbb{Q}}(j)$ or $\bar{b} \in \mathbb{G}_m(\overline{\mathbb{Q}})^k$. When $n = 2$, we show that Conjecture 4.22 holds for those $f \in \overline{\mathbb{Q}}(j)$ which satisfy both Condition 3.15 and also a suitable finiteness assumption for points satisfying certain special conditions.

4.4.2. The Zilber–Pink conjecture for $n = 1$.

Proposition 4.25. *Let $k \geq 1$ and $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$. Define $\Gamma = \langle b_1, \dots, b_k \rangle \leq \mathbb{G}_m$. Suppose that there are only finitely many f -special points σ such that the set $\{\sigma\}$ is Γ -dependent. Then there are only finitely many maximal atypical components of $V_{1,k,\bar{b}}$.*

PROOF. Fix $k \geq 1$ and $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$. We look for the atypical components of the subvariety

$$V = \{(x, t, \bar{b}) : t = R(x)\} \subset X = Y(1) \times \mathbb{G}_m^{1+k}.$$

So $\dim V = 1$ and $\dim X = k + 2$. If the set $\{b_1, \dots, b_k\}$ is multiplicatively dependent, then V is itself atypical and hence the only maximal atypical component. So we may assume that the set $\{b_1, \dots, b_k\}$ is multiplicatively independent; in particular, no b_i is a root of unity.

Consider the possible special subvarieties T of X . Clearly X cannot itself intersect V atypically, so we look only at proper special subvarieties.

We may write $T = T_1 \times T_2$, where T_1 is a special subvariety of $Y(1)$ and T_2 is a special subvariety of \mathbb{G}_m^{1+k} .

Look first at those T where T_1 is a proper special subvariety. The condition on T_1 must be a fixed coordinate. So $T = \{x\} \times T_2$, where x is a j -special point and T_2 is a special subvariety of \mathbb{G}_m^{k+1} . If $T = \{x\} \times \mathbb{G}_m^{1+k}$, then $V \cap T = \{(x, R(x), b_1, \dots, b_k)\}$, which is not atypical since

$$\dim V + \dim T - \dim X = 1 + (k + 1) - (k + 2) = 0.$$

So we must have that $T = \{x\} \times T_2$ with T_2 a proper special subvariety of \mathbb{G}_m^{1+k} . Then $\dim T \leq k$, and so the intersection $V \cap T$ is atypical if it is non-empty.

Suppose for now that T_2 has at least one fixed coordinate (i.e. a root of unity). If $V \cap T$ is non-empty, then this fixed coordinate ζ of T_2 must be in the first of the \mathbb{G}_m^{1+k} coordinates, since no b_i is a root of unity. We then require for $V \cap T \neq \emptyset$ also that $\zeta = R(x)$, and so ζ is both f -special and a root of unity. By Proposition 3.27, there are at most finitely many such points. Thus there are only finitely many atypical components of this form.

So we may suppose that T_2 has no fixed coordinate. Then some multiplicative relation must hold on T_2 since T_2 is proper. Then $V \cap T$ is non-empty (and hence atypical) only if $R(x), b_1, \dots, b_k$ satisfy this multiplicative relation, which must involve $R(x)$ since the b_i are multiplicatively independent. Note that $R(x)$ is an f -special point. Hence, $R(x)$ is a Γ -dependent f -special point, and so there are only finitely many atypical components of this form by assumption.

It thus remains to consider those cases where $T_1 = Y(1)$. So $T = Y(1) \times T_2$, where T_2 must be a proper special subvariety of \mathbb{G}_m^{k+1} since $T \subsetneq X$. Then

$$V \cap T = \{(t, R(t), b_1, \dots, b_k) : (R(t), b_1, \dots, b_k) \in T_2\}.$$

The subvariety T_2 of \mathbb{G}_m^{1+k} is proper, so either some coordinate is fixed in T_2 or a multiplicative relation holds. If $V \cap T \neq \emptyset$, then any fixed coordinate in T_2 must be the first coordinate since no b_i is a root of unity. Hence in this case $V \cap T$ must be finite. Now suppose some multiplicative relation holds on T_2 . Since the b_i are multiplicatively independent, if $V \cap T \neq \emptyset$, then this multiplicative relation must involve the first coordinate of T_2 . Write $t_1^{a_1} \cdots t_{k+1}^{a_{k+1}} = 1$ for this relation (so $a_1 \neq 0$). The equation

$$t^{a_1} b_1^{a_2} \cdots b_k^{a_{k+1}} = 1$$

has only finitely many solutions t . So $V \cap T$ is finite.

In either of these cases, $\dim(V \cap T) = 0$. So the components of the intersection $V \cap T$ are then atypical only if

$$\begin{aligned} 0 &> \dim V + \dim T - \dim X \\ &= 1 + (1 + \dim T_2) - (k + 2) \\ &= \dim T_2 - k, \end{aligned}$$

i.e. if $\dim T_2 < k$. So at least two independent conditions must hold on T_2 . Any fixed coordinate condition must apply to the first coordinate, as otherwise $V \cap T$ would be empty since no b_i is a root of unity. Since the two conditions are independent, at least one of them must therefore be

a multiplicative relation. Further, if $V \cap T \neq \emptyset$, then this multiplicative relation must involve the first coordinate since the b_i are multiplicatively independent. This then rules out the possibility of the first coordinate of T_2 being fixed because that would then imply the existence of a multiplicative relation among the b_i if $V \cap T \neq \emptyset$. So we must have that the second condition is also a multiplicative relation, and this multiplicative relation must again involve the first coordinate of T_2 for the same reason as before. So we have T_2 defined by conditions

$$t^a t_1^{a_1} \cdots t_k^{a_k} = 1,$$

$$t^{a'} t_1^{a'_1} \cdots t_k^{a'_k} = 1$$

with $a, a' \neq 0$. If $V \cap T \neq \emptyset$, then there is some t such that

$$t^a b_1^{a_1} \cdots b_k^{a_k} = 1,$$

$$t^{a'} b_1^{a'_1} \cdots b_k^{a'_k} = 1.$$

Since the two multiplicative conditions are independent we may then eliminate t to get a multiplicative dependence among b_1, \dots, b_k , a contradiction. Hence there are no atypical components of this form, and the proof is complete. \square

Corollary 4.26. *Let $k \geq 1$ and $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$. Suppose that either $f \in \overline{\mathbb{Q}}(j)$ or $b_1, \dots, b_k \in \overline{\mathbb{Q}}$. Then there are only finitely many maximal atypical components of $V_{1,k,\bar{b}}$.*

PROOF. If b_1, \dots, b_k are multiplicatively dependent, then $V_{1,k,\bar{b}}$ is itself the only maximal atypical component. If b_1, \dots, b_k are multiplicatively

independent, then $\Gamma = \langle b_1, \dots, b_k \rangle \leq \mathbb{G}_m$ is a subgroup of finite, positive rank. The desired result then follows immediately from Proposition 4.25 combined with Propositions 3.33 and 3.35 applied to Γ . \square

4.4.3. The Zilber–Pink conjecture for $n = 2$. In this subsection, we show that Conjecture 4.22 for $V_{2,k,\bar{b}}$ follows from Condition 3.15 together with a finiteness statement for pairs $(x_1, x_2) \in Y(1)^2$ satisfying certain special conditions.

Definition 4.27. Let $\Gamma \leq \mathbb{G}_m$. We define an (f, Γ) -pair to be a pair $(x_1, x_2) \in Y(1)^2$ such that (x_1, x_2) satisfies a modular relation, $R(x_1)$ is Γ -dependent, and $R(x_2)$ is Γ -dependent.

We now prove the following conditional version of Conjecture 4.22 for $n = 2$.

Proposition 4.28. *Suppose that $f \in \overline{\mathbb{Q}}(j)$ satisfies Condition 3.15. Let $k \geq 1$ and $\bar{b} = (b_1, \dots, b_k) \in \mathbb{G}_m^k$. Set $\Gamma = \langle b_1, \dots, b_k \rangle \leq \mathbb{G}_m$. Suppose that there are only finitely many (f, Γ) -pairs (x_1, x_2) with $x_1 \neq x_2$ and x_1, x_2 not j -special. Then there are only finitely many maximal atypical components of $V_{2,k,\bar{b}}$.*

PROOF. We look for the atypical components of

$$V = \{(x_1, x_2, t_1, t_2, b_1, \dots, b_k) : t_1 = R(x_1), t_2 = R(x_2)\} \subset X = Y(1)^2 \times \mathbb{G}_m^{2+k}.$$

So $\dim V = 2$ and $\dim X = k + 4$. A special subvariety T of X gives rise to an atypical component of V if

$$\dim(V \cap T) > \dim V + \dim T - \dim X = 2 - \text{codim } T.$$

If the set $\{b_1, \dots, b_k\}$ is multiplicatively dependent, then V is itself atypical and hence the only maximal atypical component. Thus we may assume that the set $\{b_1, \dots, b_k\}$ is multiplicatively independent. The group Γ is of finite rank and so, by Proposition 3.32, for every $n \geq 1$ there are only finitely many n -tuples of distinct f -special points which are Γ -dependent and minimal for this property.

Consider the possible special subvarieties T of X . We find those T that give rise to atypical components of V . We split into cases according to $\text{codim } T$.

(1) Clearly X cannot itself intersect V atypically, so we only need to look at proper special subvarieties.

(2) Suppose $\text{codim } T = 1$. Then some component of $V \cap T$ is atypical if and only if $\dim(V \cap T) \geq 2$. If T is defined by one fixed multiplicative coordinate, then $\dim(V \cap T)$ is either 0 or 1 depending on whether this fixed coordinate is one of the first two \mathbb{G}_m coordinates or not. In either case, the intersection is not atypical. If T is defined by specifying one fixed modular coordinate, then $\dim(V \cap T) = 1$ and the intersection is not atypical. Similarly if T is defined by a single modular relation. So the only case left is where T is defined by a multiplicative relation. Then, since the b_i are multiplicatively independent, this multiplicative relation must involve at least one of the first two \mathbb{G}_m coordinates if $V \cap T \neq \emptyset$. But then $\dim(V \cap T) = 1$ and the intersection is not atypical.

(3) Next we look at $\text{codim } T = 2$. Then a component of $V \cap T$ is atypical if and only if it is positive-dimensional. Clearly then T cannot be defined by two independent modular conditions. Suppose T is defined by two independent multiplicative conditions (either fixed coordinates or multiplicative relations). If $V \cap T \neq \emptyset$, then each of these conditions must involve at least one of the first two \mathbb{G}_m coordinates since the b_i are multiplicatively independent. In all such cases, one then sees that $V \cap T$ must be finite and hence its components cannot be atypical.

So T must be defined by one modular condition and one multiplicative condition. If both conditions are fixed coordinates, then $V \cap T$ is positive-dimensional only if the two conditions either both apply to the respective first coordinate or both apply to the respective second coordinate, and they also satisfy $\zeta = R(x)$, where ζ is the multiplicative fixed coordinate and x is the fixed modular coordinate. In such cases, ζ is both an f -special point and a root of unity, and so by Proposition 3.27 there are only finitely many atypical components of this kind.

If T is defined by a modular relation and a multiplicative relation, then, by Condition 3.15 and the multiplicative independence of the b_i , the intersection $V \cap T$ cannot be positive-dimensional unless the multiplicative relation has the form $x = y$. Since no non-constant modular function is invariant under a larger subgroup of $\text{GL}_2^+(\mathbb{Q})$ than $\mathbb{Q}^\times \cdot \text{SL}_2(\mathbb{Z})$, the modular relation must also be of the form $x = y$. So there is just one such atypical component.

Suppose T is defined by a modular relation and a fixed multiplicative coordinate. Then this fixed coordinate must be one of the first two \mathbb{G}_m^{2+k} coordinates. But then the modular relation on the $Y(1)^2$ coordinates implies that $V \cap T$ is finite and so not atypical.

If T is defined by a multiplicative relation and a fixed modular coordinate, then $V \cap T$ is positive-dimensional only if the multiplicative relation involves the \mathbb{G}_m^{2+k} coordinate corresponding to the fixed $Y(1)^2$ coordinate, but not the other of the first two \mathbb{G}_m^{2+k} coordinates. If the fixed modular coordinate is given by the j -special point x , then the f -special point $R(x)$ is Γ -dependent since it satisfies a multiplicative relation over the b_i . Hence, by Proposition 3.32, there are only finitely many atypical components of this form to consider.

(4) When $\text{codim } T = 3$, the components of the intersection $V \cap T$ are atypical if and only if $V \cap T \neq \emptyset$. If all three conditions defining T are fixed coordinates, then $V \cap T \neq \emptyset$ implies that $V \cap T$ must already be contained in one of the positive-dimensional atypical components arising from T defined by a fixed modular coordinate and a fixed multiplicative coordinate. So we may assume that at least one condition defining T is a relation.

If there were three independent multiplicative conditions defining T , then $V \cap T \neq \emptyset$ would imply a multiplicative relation among the b_i , which is impossible. Clearly one cannot have three independent conditions on the $Y(1)^2$ coordinates. So there must be at least one modular condition and at least one multiplicative condition defining T .

If there are two independent modular conditions, then we may assume these are fixed coordinates. The other condition must then be a multiplicative relation, and this relation must involve at least one of the first two \mathbb{G}_m^{2+k} coordinates if $V \cap T \neq \emptyset$ since the b_i are multiplicatively independent. If the multiplicative relation involves only one of the first two \mathbb{G}_m^{2+k} coordinates, then the corresponding fixed modular coordinate x gives rise to a Γ -dependent f -special point $R(x)$. We are thus in one of the positive-dimensional atypical components arising from a special subvariety of codimension 2.

If the multiplicative relation involves both the first two \mathbb{G}_m^{2+k} coordinates, then the fixed modular coordinates x_1, x_2 give rise to a Γ -dependent pair of f -special points $(R(x_1), R(x_2))$. If some $R(x_i)$ is Γ -dependent, then we are in one of the previously identified atypical components. We may thus assume that no subtuple of $(R(x_1), R(x_2))$ is Γ -dependent. Since there are only finitely many pairs of j -special points (x_1, x_2) with $x_1 \neq x_2$ and $R(x_1) = R(x_2)$, as shown in the proof of Proposition 4.16, we may also assume that $R(x_1) \neq R(x_2)$. It thus follows from Proposition 3.32 that there are at most finitely many maximal atypical components of this form.

So we may suppose that T is defined by two multiplicative conditions and one modular condition. If the multiplicative conditions are both fixed coordinates, then the modular condition is a relation. We thus have that $V \cap T = \{(x_1, x_2, \zeta_1, \zeta_2, b_1, \dots, b_k)\}$, where (x_1, x_2) satisfies a modular relation, ζ_1, ζ_2 are roots of unity, and $\zeta_i = R(x_i)$ for $i = 1, 2$. By Proposition 4.19, there are only finitely many such components satisfying the additional restrictions $x_1 \neq x_2$ and neither x_1 nor x_2 is j -special. If these

additional restrictions are not met, then we are in one of the previously identified atypical components.

So we may assume that at least one of the multiplicative conditions is a relation; clearly, this relation must involve at least one of the first two \mathbb{G}_m^{2+k} coordinates. Suppose the modular condition is a fixed coordinate x . If the second multiplicative condition is a fixed coordinate ζ , then clearly this must be one of the first two \mathbb{G}_m^{2+k} coordinates. Further since ζ is a root of unity, we can eliminate this coordinate from the multiplicative relation. If $V \cap T \neq \emptyset$, then, according to whether the two fixed coordinates are in the same respective position or not, either $\zeta = R(x)$ is both a root of unity and f -special or $(R(x), b_1, \dots, b_k)$ satisfy the multiplicative relation and so $R(x)$ is a Γ -dependent 1-tuple. Thus such components are contained in positive-dimensional atypical components arising from special subvarieties of codimension 2.

So now suppose the modular condition is a fixed coordinate x and both the multiplicative conditions are relations. If $V \cap T \neq \emptyset$, then both relations must involve at least one of the first two \mathbb{G}_m^{2+k} coordinates. They cannot both involve only the first (respectively the second) of the first two \mathbb{G}_m^{2+k} coordinates, since by their independence we would then be able to obtain a relation among the b_i . If both relations involve both the first two \mathbb{G}_m^{k+2} coordinates, then we may eliminate either of these two coordinates. Thus we may assume that the first relation involves the first but not the second \mathbb{G}_m^{2+k} coordinate and that the second relation involves the second but not the first \mathbb{G}_m^{2+k} coordinate. The components of $V \cap T$ are therefore contained in some of the already identified positive-dimensional atypical components.

We thus reduce to considering when the modular condition is a relation and at least one of the multiplicative conditions is a relation. Since the b_i are multiplicatively independent, if $V \cap T \neq \emptyset$, then any component of $V \cap T$ must have the form $\{(x_1, x_2, R(x_1), R(x_2), b_1, \dots, b_k)\}$ where (x_1, x_2) satisfies a modular relation and $R(x_1), R(x_2)$ are both (individually) Γ -dependent. Any such (x_1, x_2) is an (f, Γ) -pair. The finiteness of the resulting maximal atypical components then follows from the assumption on (f, Γ) -pairs in the hypotheses of the proposition. This is because if $x_1 = x_2$ or some x_i is j -special, then the corresponding component of $V \cap T$ is already contained in one of the atypical components above.

(5) We now consider the case when $\text{codim } T \geq 4$. If T is defined by ≥ 3 independent multiplicative conditions and $V \cap T \neq \emptyset$, then we can eliminate the first two \mathbb{G}_m^{2+k} coordinates from these relations and obtain a multiplicative dependency among b_1, \dots, b_k . This obviously cannot happen. Clearly, there can also be no more than two independent modular conditions defining T . Thus the only case to consider is when T is defined by two modular conditions and two multiplicative conditions.

The components of the intersection $V \cap T$ are then atypical if and only if $V \cap T \neq \emptyset$. We may assume that the two modular conditions are both fixed coordinates. If one of the multiplicative conditions is a fixed coordinate, then this must be one of the first two \mathbb{G}_m^{2+k} coordinates. If $V \cap T \neq \emptyset$, then the root of unity corresponding to this fixed coordinate must also be an f -special point because the respective modular coordinate is j -special. In this case, the intersection is contained in one of the already identified positive-dimensional atypical components. Therefore we may assume that

both the multiplicative conditions are relations. The components that can arise here are thus all contained in larger atypical components identified in (3). \square

We note here the similarity between the finiteness assumption on (f, Γ) -pairs contained in the hypotheses of Proposition 4.28 and the finiteness statement on modular-torsion tuples in Proposition 4.19. In both cases, one is dealing with the atypical components which arise when one has a modular relation and two independent multiplicative relations. However, in the case of Proposition 4.19 we are able to prove finiteness, whereas in Proposition 4.28 we must assume it.

The difference is that in Proposition 4.19 the points $R(x_1), R(x_2)$ are roots of unity, since if $R(x_1), R(x_2)$ satisfy two independent multiplicative relations then they must both be roots of unity. One thus obtains that x_1, x_2 are algebraic and their heights are bounded, which is crucial for the finiteness proof. In contrast, for an (f, Γ) -pair (x_1, x_2) the multiplicative relations satisfied by $R(x_1), R(x_2)$ involve also the generators b_1, \dots, b_k of Γ . One thus obtains only the weaker condition of each $R(x_i)$ being Γ -dependent, rather than a root of unity. In particular, we thus do not seem able to obtain suitable bounds on the heights of x_1, x_2 (they might not even be algebraic) in order to prove the finiteness of the pairs (x_1, x_2) by a modified version of the argument for Proposition 4.19.

CHAPTER 5

Effective results on products of singular moduli

In this thesis, we have so far proved numerous finiteness results (e.g. Theorems 2.3, 3.7, 3.8, and Proposition 4.19). One thing these results have in common is that they are all ineffective; in other words, there is no way of knowing just how many is “finite many” or how “complicated” the finitely many examples may be.

There are several sources of ineffectivity in the proofs of these results, which have all made use of o-minimal methods. In particular, the constants in the Pila–Wilkie Counting Theorem and in the lower bound for the size of the Galois orbit of a special point are ineffective. There is therefore no straightforward way to modify our proofs in order to obtain effective versions of our results.

On the other hand, we have a very explicit understanding of the arithmetic properties of singular moduli, as was indicated in Section 1.7. Therefore, if we restrict just to the case of multiplicative relations among singular moduli, then proving some effective results is feasible, at least in low dimensions. In particular, fully effective results are known already in one and two dimensions for the question of multiplicatively dependent singular moduli. In this chapter, we will prove a first effective result in three dimensions. All these results are naturally formulated in terms of the André–Oort conjecture, which we now discuss briefly.

5.1. The André–Oort conjecture

Definition 5.1. Let X be a Shimura variety and $V \subset X$ a subvariety. A special subvariety $T \subset X$ is called a maximal special subvariety (of V) if $T \subset V$ and there is no special subvariety $T' \subset X$ such that $T \subsetneq T' \subset V$.

Conjecture 5.2. *Let X be a Shimura variety and $V \subset X$ a subvariety. Then V contains only finitely many maximal special subvarieties.*

The André–Oort conjecture, formulated in [3] and [55], is an analogue, in the setting of Shimura varieties, of the Manin–Mumford conjecture (which was proved by Raynaud [73]). Pink’s formulation of the Zilber–Pink conjecture arose (in part) as a generalisation of the André–Oort conjecture.

Proposition 5.3. *The Zilber–Pink conjecture implies the André–Oort conjecture.*

PROOF. Let X be a Shimura variety. Suppose that $A \subsetneq X$ is a proper subvariety which contains a special subvariety $T \subset X$. Then $A \cap T = T$ and

$$\dim T > \dim A + \dim T - \dim X.$$

Hence, T is atypical for A in X .

Now fix $V \subset X$ a subvariety. If $V = X$, then V itself is a special subvariety and so the only maximal special subvariety. Hence, we may and do assume that $V \subsetneq X$. By the Zilber–Pink conjecture, V contains only finitely many maximal atypical components. Enumerate these as A_i , where

$A_i \subset V \cap T_i$ for a special subvariety $T_i \subset X$ such that

$$\dim A_i > \dim V + \dim T_i - \dim X.$$

Suppose then that S is a maximal special subvariety of V . Then S is a special subvariety of X contained in V . Hence, S is atypical for V in X . So S is contained in some A_i . Now either $A_i = T_i$ and hence $S = A_i = T_i$, or else S is atypical for A_i in T_i . In the first case, we are done. In the second case, by the Zilber–Pink conjecture again, there are only finitely many maximal atypical components of A_i in T_i , and S must be contained in one of them. We may continue this process, which must terminate after a finite number of steps. We thus see that S must belong to a finite list of special subvarieties of X . \square

We will be interested in the case that $X = Y(1)^n$, where $n \geq 1$. In this case (and many others), the André–Oort conjecture is known.

Theorem 5.4. *Let $n \geq 1$ and set $X = Y(1)^n$. Let $V \subset X$ be a subvariety. Then V contains only finitely many maximal special subvarieties.*

For $n = 1$, this is trivial. The case of $n = 2$ was proved by André [4]. In general, it was proved by Pila [60] using o-minimal methods. These results were both ineffective: they do not give effective bounds on the degrees of the maximal special subvarieties. An effective version of André’s result for $n = 2$ was proved by Kühne [43] and also independently in [12]. Some further special cases of Theorem 5.4 for restricted classes of subvarieties have been proved effectively [8, 13].

Let $F \in \mathbb{C}[x_1, x_2]$. The effective André–Oort result for subvarieties of $Y(1)^2$ allows one to bound by an effective constant (depending on F)

the discriminants of all singular moduli x, y such that $F(x, y) = 0$. In principle, one may then find by a computation all such pairs (x, y) . The relevant effective constant here though depends on the polynomial F . It is therefore of interest to try and find all the solutions in singular moduli for families of polynomials F which are of interest. Various such results have been obtained, e.g. [2, 10, 11, 46, 74].

5.2. The multiplicative case

Fix $n \geq 1$ and let $X = Y(1)^n$. Given $a_i \in \mathbb{Z} \setminus \{0\}$, define the subvariety

$$V_a = \{(x_1, \dots, x_n) \in X : x_1^{a_1} \dots x_n^{a_n} = 1\}.$$

Theorem 3.2 implies that the only positive-dimensional special subvarieties of V_a come from equations of the form $x_i = x_j$ for $i \neq j$.

Suppose then that $\sigma_1, \dots, \sigma_n$ are pairwise distinct singular moduli such that $\sigma_1^{a_1} \dots \sigma_n^{a_n} = 1$ for some $a_i \in \mathbb{Z} \setminus \{0\}$. Let $a = (a_1, \dots, a_n)$. Then $\{(\sigma_1, \dots, \sigma_n)\}$ is a zero-dimensional maximal special subvariety of V_a . Conversely, every zero-dimensional maximal special subvariety is of this form.

Therefore, proving an effective André–Oort statement for the family of subvarieties V_a reduces to the problem of finding effective bounds on solutions in pairwise distinct singular moduli to equations of the form $x_1^{a_1} \dots x_n^{a_n} = 1$, where $a_i \in \mathbb{Z} \setminus \{0\}$. It is therefore a closely related problem to the question of effectivising Theorem 3.1.

Note that the André–Oort statement Theorem 5.4 does not imply Theorem 3.1. André–Oort implies that, given $a_i \in \mathbb{Z} \setminus \{0\}$, there are only finitely many n -tuples σ of pairwise distinct singular moduli satisfying the equation $\sigma_1^{a_1} \dots \sigma_n^{a_n} = 1$. To deduce Theorem 3.1, one would need, for

each $n \geq 1$, an André–Oort statement which was uniform across all the $V_a \subset Y(1)^n$. If this statement was not only uniform but also effective, then one could (at least in principle) find all the possible n -tuples for a given $n \geq 1$. It is this task which we will consider in the remainder of this chapter. In fact, we will consider the more general family of subvarieties $V_{a,q}$, where $a_i \in \mathbb{Z} \setminus \{0\}$ and $q \in \mathbb{Q}^\times$, which are defined

$$V_{a,q} = \{(x_1, \dots, x_n) \in X : x_1^{a_1} \dots x_n^{a_n} = q\}.$$

When $n \leq 2$, the singular moduli lying on subvarieties of this form have been classified completely by Riffaut [74], generalising earlier work of [10].

Proposition 5.5 ([74]). *Suppose that x is a singular modulus such that $x^n \in \mathbb{Q}$ for some $n \in \mathbb{Z} \setminus \{0\}$. Then $x \in \mathbb{Q}$.*

Theorem 5.6 ([74, Theorem 1.6]). *Let x_1, x_2 be non-zero singular moduli and $m, n \in \mathbb{Z} \setminus \{0\}$. Suppose that $x_1^m x_2^n \in \mathbb{Q}^\times$. Then one of the following holds:*

- (1) $x_1, x_2 \in \mathbb{Q}^\times$;
- (2) $x_1 = x_2$ and $m + n = 0$;
- (3) $m = n$ and $x_1 \neq x_2$ are of degree 2 and conjugate over \mathbb{Q} .

Since 0 is a singular modulus, we must exclude the case of product 0 to obtain any kind of finiteness statement in the second result. Both results are clearly best possible. There are 13 rational singular moduli (including 0) and 29 conjugate pairs of singular moduli of degree 2. The list of these is known explicitly, and so one may easily find the set of $q \in \mathbb{Q}$ such that there are singular moduli x_1, x_2 with $x_1^m x_2^n = q$ for some $m, n \in \mathbb{Z}$.

The following result, which we prove in this chapter, is a first step towards proving the corresponding result for $n = 3$. It is the analogue of the result for $n = 2$ of [10], which was subsequently generalised by Riffaut [74] to prove Theorem 5.6.

Theorem 5.7. *Suppose that x_1, x_2, x_3 are singular moduli such that $x_1x_2x_3 \in \mathbb{Q}^\times$. Then one of the following holds:*

- (1) $x_1, x_2, x_3 \in \mathbb{Q}^\times$;
- (2) some $x_i \in \mathbb{Q}^\times$ and the remaining x_j, x_k are distinct, of degree 2, and conjugate over \mathbb{Q} ;
- (3) x_1, x_2, x_3 are pairwise distinct, of degree three, and conjugate over \mathbb{Q} .

Conversely, if one of (1)–(3) holds, then it is clear that $x_1x_2x_3 \in \mathbb{Q}^\times$. Further, each of these cases is achieved. Theorem 5.7 is therefore best possible. Theorem 5.7 provides a completely explicit André–Oort statement for the family of subvarieties $V_{a,q} \subset Y(1)^3$, where $a_1 = a_2 = a_3 = \pm 1$ and $q \in \mathbb{Q}$.

There are 13 rational singular moduli, one of which is 0; 29 pairs of conjugate singular moduli of degree 2; and 25 triples of conjugate singular moduli of degree 3. The list of these may be computed in PARI. There are thus 364 unordered triples (x_1, x_2, x_3) as in (1); 348 unordered triples (x_1, x_2, x_3) as in (2); and 25 unordered triples (x_1, x_2, x_3) as in (3). One may straightforwardly compute the corresponding products $x_1x_2x_3$. There are 13 rational numbers which are the product of two distinct triples in (1) and 16 rational numbers which are the products of triples in (1) and (2). No other rational number is the product of more than one such triple. There

are thus 708 distinct non-zero rational numbers which arise as the product of three singular moduli, and the list of both these rational numbers and the corresponding triples of singular moduli is known. Since singular moduli are algebraic integers, we note that if $x_1x_2x_3 \in \mathbb{Q}$, then in fact $x_1x_2x_3 \in \mathbb{Z}$ and so these 708 distinct rational numbers are all rational integers.

The plan for our proof of Theorem 5.7 is as follows. Section 5.3 contains the facts about singular moduli that we need for the proof of Theorem 5.7. The proof itself is split over Sections 5.4 and 5.5. In Section 5.4, we reduce to an effective finite list the possible triples (x_1, x_2, x_3) of pairwise distinct singular moduli with non-zero rational product which do not belong to one of the trivial cases (1)–(3) of Theorem 5.7. Then in Section 5.5 we explain how to use a PARI script [57] to eliminate all the triples on this list. The PARI scripts used are available from <https://github.com/guyfowler/rationaltriples>. The results of this chapter were previously published in the author’s paper [33].

5.3. Background on singular moduli

We begin with some upper and lower bounds for singular moduli. These bounds will be used without special reference in Section 5.4. For every non-zero singular modulus x of discriminant Δ , we have ([10, (12)]) the lower bound

$$|x| \geq \min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}.$$

Recall from Section 1.7 that the singular moduli of a given discriminant Δ may be explicitly described in terms of the set T_Δ . For a singular

modulus x corresponding to a triple $(a, b, c) \in T_\Delta$, we have that ([**31**, §2]):

$$e^{\pi|\Delta|^{1/2}/a} - 2079 \leq |x| \leq e^{\pi|\Delta|^{1/2}/a} + 2079.$$

This follows from the q -expansion of the j -function given in Proposition 1.15. We will apply this bound variously with $a = 2, 3, 4, 5$, making use of Lemma 5.8 below. For a singular modulus x corresponding to a triple $(a, b, c) \in T_\Delta$ with $a = 1$ and $|\Delta| \geq 23$, one obtains also that $|x| \geq 0.9994e^{\pi|\Delta|^{1/2}}$, see [**10**, (11)].

Lemma 5.8. *For a given discriminant Δ , there exists:*

- (1) *a unique singular modulus, corresponding to a triple with $a = 1$;*
- (2) *at most two singular moduli, corresponding to triples with $a = 2$, and if $\Delta \equiv 4 \pmod{16}$, then there are no such singular moduli;*
- (3) *at most two singular moduli corresponding to triples with $a = 3$;*
- (4) *at most two singular moduli corresponding to triples with $a = 4$;*
- (5) *at most two singular moduli corresponding to triples with $a = 5$.*

We call the unique singular modulus corresponding to a triple $(a, b, c) \in T_\Delta$ with $a = 1$ the dominant singular modulus of discriminant Δ .

PROOF. The first two claims are Proposition 2.6 of [**10**]. We show the remaining claims.

Suppose $a = 3$. Let $(3, b_1, c_1), (3, b_2, c_2)$ be two such tuples. Then $b_1, b_2 \in \{-2, -1, 0, 1, 2, 3\}$. Since $\Delta = b^2 - 4ac$ for all $(a, b, c) \in T_\Delta$, one has that $b_1^2 - 12c_1 = b_2^2 - 12c_2$. Thus $b_1^2 - b_2^2 \equiv 0 \pmod{12}$. Therefore, it must be that $b_1 = \pm b_2$. Since a_i, b_i together uniquely determine c_i , there are at most two tuples in T_Δ with $a = 3$.

Now let $a = 4$. Suppose $(4, b_1, c_1), (4, b_2, c_2)$ are two such tuples. Then $b_1, b_2 \in \{-3, -2, -1, 0, 1, 2, 3, 4\}$. Since $\Delta = b^2 - 4ac$ for all $(a, b, c) \in T_\Delta$, one has that $b_1^2 - 16c_1 = b_2^2 - 16c_2$. Thus $b_1^2 - b_2^2 \equiv 0 \pmod{16}$. Therefore, it must be that either $b_1 = \pm b_2$ or $\{b_1, b_2\} = \{0, 4\}$. Since a_i, b_i together uniquely determine c_i , there are at most two tuples in T_Δ with $a = 4$.

Let $a = 5$. Suppose $(5, b_1, c_1), (5, b_2, c_2)$ are two such tuples. Then $b_1, b_2 \in \{-4, -3, -2, -1, 0, 1, 2, 3, 4, 5\}$. Since $\Delta = b^2 - 4ac$ for all $(a, b, c) \in T_\Delta$, one has that $b_1^2 - 20c_1 = b_2^2 - 20c_2$. Thus $b_1^2 - b_2^2 \equiv 0 \pmod{20}$. Therefore, it must be that $b_1 = \pm b_2$. Since a_i, b_i together uniquely determine c_i , there are at most two tuples in T_Δ with $a = 5$. \square

Our proof of Theorem 5.7 will also rely on some results about the fields generated by singular moduli. The first of these is a result on when two singular moduli generate the same field. It was proved mostly in [2], as Corollary 4.2 and Proposition 4.3. For the “further” claim in (2), see [10, §3.2.2].

Lemma 5.9. *Let x_1, x_2 be singular moduli with discriminants Δ_1, Δ_2 respectively. Suppose that $\mathbb{Q}(x_1) = \mathbb{Q}(x_2)$, and denote this field L . Then $h(\Delta_1) = h(\Delta_2)$, and we have that:*

- (1) *If $\mathbb{Q}(\sqrt{\Delta_1}) \neq \mathbb{Q}(\sqrt{\Delta_2})$, then the possible fields L are listed in [2, Table 4.1]. Further, the field L is Galois and the discriminant of any singular modulus x with $\mathbb{Q}(x) = L$ is also listed in this table.*
- (2) *If $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2})$, then either: $\Delta_1, \Delta_2 \in \{-3, -12, -27\}$ and $L = \mathbb{Q}$; or: $\Delta_1/\Delta_2 \in \{1, 4, 1/4\}$ and, further, if $\Delta_1 = 4\Delta_2$, then $\Delta_2 \equiv 1 \pmod{8}$.*

We now establish a similar result on when one singular modulus generates a degree 2 subfield of the field generated by another singular modulus. We split the proof into the next two lemmas.

Lemma 5.10. *Let x_1, x_2 be singular moduli with discriminants Δ_1, Δ_2 respectively. Suppose that $[\mathbb{Q}(x_1) : \mathbb{Q}(x_2)] = 2$. If $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2})$, then either $\Delta_1 \in \{9\Delta_2/4, 4\Delta_2, 9\Delta_2, 16\Delta_2\}$ or $x_2 \in \mathbb{Q}$.*

PROOF. We modify the proof in [2] of Lemma 5.12. Given an imaginary quadratic field K of (fundamental) discriminant D and an integer $f \geq 1$, we write $\text{RiCF}(K, f)$ for the ring class field of the imaginary quadratic order of discriminant $\Delta = f^2D$.

Let x_1, x_2 be singular moduli with discriminants Δ_1, Δ_2 respectively. Suppose that $[\mathbb{Q}(x_1) : \mathbb{Q}(x_2)] = 2$ and $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2})$. Denote by K the field $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2})$. The singular moduli x_1, x_2 have the same fundamental discriminant (the discriminant of the field K), which we denote D . We may then write $\Delta_1 = f_1^2D$, $\Delta_2 = f_2^2D$. Note that $\Delta_1 \neq \Delta_2$. Let $f = \text{lcm}(f_1, f_2)$.

Suppose in addition that $D \neq -3, -4$. Then [2, Proposition 3.1]

$$\text{RiCF}(K, f_1)\text{RiCF}(K, f_2) = \text{RiCF}(K, f),$$

where the left hand side denotes the compositum of $\text{RiCF}(K, f_1)$ and $\text{RiCF}(K, f_2)$. By the theory of complex multiplication $\text{RiCF}(K, f_i) = K(x_i)$ for $i = 1, 2$. Thus $\text{RiCF}(K, f) = K(x_1)$ since $x_2 \in \mathbb{Q}(x_1)$. Therefore

$$h(f^2D) = [\text{RiCF}(K, f) : K] = [\text{RiCF}(K, f_1) : K] = h(f_1^2D).$$

Also,

$$[\mathbb{Q}(x_1) : \mathbb{Q}] = 2[\mathbb{Q}(x_2) : \mathbb{Q}],$$

and thus $h(f^2D) = h(f_1^2D) = 2h(f_2^2D)$. As in the proof of [2, Proposition 4.3], one may then use the class number formula [2, (6)] to obtain that

$$\frac{f}{f_1} \prod_{\substack{p|f \\ p \nmid f_1}} \left(1 - \left(\frac{D}{p}\right) \frac{1}{p}\right) = 1$$

and

$$\frac{f}{f_2} \prod_{\substack{p|f \\ p \nmid f_2}} \left(1 - \left(\frac{D}{p}\right) \frac{1}{p}\right) = 2,$$

where $\left(\frac{D}{p}\right)$ denotes the Kronecker symbol. This implies that $f/f_1 \in \{1, 2\}$ and $f/f_2 \in \{2, 3, 4\}$. One thus has that $f_1/f_2 \in \{3/2, 2, 3, 4\}$ since $f_1 \neq f_2$. Hence $\Delta_1 \in \{9\Delta_2/4, 4\Delta_2, 9\Delta_2, 16\Delta_2\}$.

Now let $D \in \{-3, -4\}$. If $\gcd(f_1, f_2) > 1$, then, by [2, Proposition 3.1] again,

$$\text{RiCF}(K, f_1)\text{RiCF}(K, f_2) = \text{RiCF}(K, f),$$

and the above proof works. If $f_1 = 1$, then $\mathbb{Q}(x_1) = \mathbb{Q}$, a contradiction. If $f_2 = 1$, then $x_2 \in \mathbb{Q}$ and we are done.

So we may now assume that $f_1, f_2 > 1$ and $\gcd(f_1, f_2) = 1$. So $f = f_1f_2$. In this case, by [2, Proposition 3.1],

$$2h(f_2^2D) = h(f_1^2D) = l^{-1}h(f^2D),$$

where $l = 2$ for $D = -4$ and $l = 3$ for $D = -3$. We now apply again the class number formula to obtain that

$$f_1 \prod_{p|f_1} \left(1 - \left(\frac{D}{p}\right) \frac{1}{p}\right) = 2l$$

and

$$f_2 \prod_{p|f_2} \left(1 - \left(\frac{D}{p}\right) \frac{1}{p}\right) = l.$$

Inspecting the possibilities for f_2 , we see that

$$\Delta_2 \in \{-12, -16, -27\}.$$

But then $x_2 \in \mathbb{Q}$. □

Lemma 5.11. *Let x_1, x_2 be singular moduli with discriminants Δ_1, Δ_2 respectively. Suppose that $[\mathbb{Q}(x_1) : \mathbb{Q}(x_2)] = 2$. If $\mathbb{Q}(\sqrt{\Delta_1}) \neq \mathbb{Q}(\sqrt{\Delta_2})$, then one of the following holds:*

- (1) *at least one of Δ_1 or Δ_2 is listed in [2, Table 2.1] and the corresponding field $\mathbb{Q}(x_i)$ is Galois;*
- (2) *$h(\Delta_1) \geq 128$.*

PROOF. This proof is a modified version of [2, Theorem 4.1]. If $\mathbb{Q}(x_1)$ is Galois over \mathbb{Q} , then by Corollaries 3.3 and 2.2 and Remark 2.3 of [2], either Δ_1 is listed in [2, Table 2.1] or $h(\Delta_1) \geq 128$. If $\mathbb{Q}(x_2)$ is Galois over \mathbb{Q} , then similarly either Δ_2 is listed in [2, Table 2.1] or $h(\Delta_2) \geq 128$ (and so certainly $h(\Delta_1) \geq 128$).

So we may now suppose that neither $\mathbb{Q}(x_1)$ nor $\mathbb{Q}(x_2)$ is Galois over \mathbb{Q} . We will show this leads to a contradiction. Let M_1 be the Galois closure of $\mathbb{Q}(x_1)$ over \mathbb{Q} . Then $M_1 = \mathbb{Q}(\sqrt{\Delta_1}, x_1) \supset \mathbb{Q}(x_2)$. Let M_2 be the Galois

closure of $\mathbb{Q}(x_2)$ over \mathbb{Q} . Then $M_2 = \mathbb{Q}(\sqrt{\Delta_2}, x_2)$. Also $M_2 \subset M_1$ since M_1 is Galois and contains $\mathbb{Q}(x_2)$. Since $\mathbb{Q}(x_1), \mathbb{Q}(x_2)$ are not Galois, one has that $\sqrt{\Delta_i} \notin \mathbb{Q}(x_i)$. Hence $[M_1 : \mathbb{Q}] = 2h(\Delta_1)$ and $[M_2 : \mathbb{Q}] = 2h(\Delta_2)$. In particular, $[M_1 : M_2] = 2$ since $h(\Delta_1) = 2h(\Delta_2)$.

Now let $G = \text{Gal}(M_1/\mathbb{Q})$, $H = \text{Gal}(M_1/\mathbb{Q}(\sqrt{\Delta_1}, \sqrt{\Delta_2}))$, and $H_i = \text{Gal}(M_1/\mathbb{Q}(\sqrt{\Delta_i}))$ for $i = 1, 2$. So $H = H_1 \cap H_2$, $[H_1 : H] = 2$, and $[H_2 : H] = 2$. As in the proof of [2, Theorem 4.1], one has that H is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^n$ for some n . Each of H_1, H_2 contains H as an index 2 subgroup. So H_1, H_2 must each be isomorphic to either $(\mathbb{Z}/2\mathbb{Z})^{n+1}$ or $(\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$. If $H_1 \cong (\mathbb{Z}/2\mathbb{Z})^{n+1}$, then by [2, Corollary 3.3], the field $\mathbb{Q}(x_1)$ is Galois, a contradiction. So we may assume that $H_1 \cong (\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$.

Suppose next that $H_2 \cong (\mathbb{Z}/2\mathbb{Z})^{n+1}$. Observe that H_2 is abelian, and hence all its subgroups are normal. Therefore, considering the extension $M_1/M_2/\mathbb{Q}(\sqrt{\Delta_2})$, we obtain that $\text{Gal}(M_2/\mathbb{Q}(\sqrt{\Delta_2}))$ is isomorphic to the quotient $H_2/\text{Gal}(M_1/M_2)$. Thus, since $H_2 \cong (\mathbb{Z}/2\mathbb{Z})^{n+1}$, we must have that $\text{Gal}(M_2/\mathbb{Q}(\sqrt{\Delta_2})) \cong (\mathbb{Z}/2\mathbb{Z})^n$. This implies, by [2, Corollary 3.3] again, that $\mathbb{Q}(x_2)$ is Galois, a contradiction.

Hence we must have that $H_1 \cong H_2 \cong (\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$. Exactly as in [2, Theorem 4.1], this implies that $G \cong D_8 \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$, where D_8 denotes the dihedral group with 8 elements. The group $D_8 \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$ has only one subgroup isomorphic to $(\mathbb{Z}/4\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})^{n-1}$, and hence one must have that $H_1 = H_2$. This though implies that $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2})$, a contradiction. \square

The other result we use is on the fields generated by products of pairs of non-zero singular moduli, and establishes that such a field is “close to” the field generated by the pair of singular moduli. This result is due to Faye and Riffaut [31].

Lemma 5.12 ([31, Theorem 1.3]). *Let x_1, x_2 be distinct non-zero singular moduli of discriminants Δ_1, Δ_2 respectively. Then $\mathbb{Q}(x_1x_2) = \mathbb{Q}(x_1, x_2)$ unless $\Delta_1 = \Delta_2$, in which case $[\mathbb{Q}(x_1, x_2) : \mathbb{Q}(x_1x_2)] \leq 2$.*

5.4. An effective bound

We now turn to the proof of Theorem 5.7 itself. Suppose x_1, x_2, x_3 are singular moduli such that $x_1x_2x_3 = \alpha \in \mathbb{Q}^\times$. Write Δ_i for their respective discriminants and h_i for the corresponding class numbers $h(\Delta_i)$. In this section, we will reduce the possible $(\Delta_1, \Delta_2, \Delta_3)$ to an (effectively) finite list. Without loss of generality, assume that $h_1 \geq h_2 \geq h_3$.

If the x_i are not pairwise distinct, then (1) of Theorem 5.7 must hold by Theorem 5.6, which classifies all pairs (x, y) of singular moduli satisfying $x^m y^n \in \mathbb{Q}^\times$ for some $m, n \in \mathbb{Z} \setminus \{0\}$. So we may and do assume that x_1, x_2, x_3 are pairwise distinct. If $h_3 = 1$, then $x_3 \in \mathbb{Q}^\times$ and hence $x_1x_2 \in \mathbb{Q}^\times$. Thus, by Theorem 5.6 (in fact, the earlier, weaker result of [10] suffices), either $x_1, x_2 \in \mathbb{Q}$ or x_1, x_2 are of degree 2 and conjugate over \mathbb{Q} . Thus either (1) or (2) in Theorem 5.7 holds and we are done. We therefore assume subsequently that $h_1 \geq h_2 \geq h_3 \geq 2$.

Clearly we have that $\mathbb{Q}(x_1) = \mathbb{Q}(x_2x_3)$. Thus, by Lemma 5.12, we have that

$$[\mathbb{Q}(x_1) : \mathbb{Q}] = [\mathbb{Q}(x_2x_3) : \mathbb{Q}] = \begin{cases} \text{either } [\mathbb{Q}(x_2, x_3) : \mathbb{Q}], \\ \text{or } \frac{1}{2}[\mathbb{Q}(x_2, x_3) : \mathbb{Q}]. \end{cases}$$

Note that $h_2 = [\mathbb{Q}(x_2) : \mathbb{Q}]$ and $h_3 = [\mathbb{Q}(x_3) : \mathbb{Q}]$ each divide $[\mathbb{Q}(x_2, x_3) : \mathbb{Q}]$, so we must have that $h_2, h_3 \mid 2[\mathbb{Q}(x_1) : \mathbb{Q}]$. Hence, $h_2, h_3 \mid 2h_1$. Symmetrically, we have also that $h_1, h_2 \mid 2h_3$ and $h_1, h_3 \mid 2h_2$. Then, since $h_1 \geq h_2 \geq h_3$, one of the following must hold: either $h_1 = h_2 = h_3$, or $h_1 = h_2 = 2h_3$, or $h_1 = 2h_2 = 2h_3$. We consider each of these cases in turn.

We will write (x'_1, x'_2, x'_3) for a conjugate of (x_1, x_2, x_3) , where x'_i is the conjugate of x_i associated to an element $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$. Computations in this section were carried out in PARI [57].

5.4.1. The case $h_1 = h_2 = h_3$. Write $h = h_1 = h_2 = h_3$. We split this situation into subcases, depending as to whether the Δ_i are equal.

5.4.1.1. *The subcase $\Delta_1 = \Delta_2 = \Delta_3$.* Write Δ for this shared discriminant. The x_i are thus all singular moduli of discriminant Δ and hence are all conjugate. Since the x_i are pairwise distinct, they must then be of degree at least 3, so $h \geq 3$. If $h = 3$, then we are in case (3) of Theorem 5.7. So we may assume that $h \geq 4$.

Taking conjugates as necessary, we may assume that x_1 is dominant. Since $h \geq 4$, one certainly has $|\Delta| \geq 23$. Thus, by the bounds in Section 5.3, one has the lower bound for $|\alpha|$ given by

$$|\alpha| \geq (0.9994e^{\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\})^2.$$

We now establish upper bounds for $|\alpha|$, incompatible with this lower bound for suitably large $|\Delta|$. The larger the class number h , the better these bounds will be. Let $\sigma_1, \dots, \sigma_h$ be the automorphisms of $\mathbb{Q}(x_1)$. Then $\sigma_i(x_1, x_2, x_3) = (x'_1, x'_2, x'_3)$, where x'_1, x'_2, x'_3 are themselves singular moduli of discriminant Δ , since these singular moduli form a complete Galois orbit over \mathbb{Q} . Further, if $\sigma_i(x_k) = \sigma_j(x_k)$, then $i = j$ since the action is sharply transitive. Thus each singular moduli of discriminant Δ occurs at most once among the $\sigma_i(x_k)$ for $i = 1, \dots, h$.

Let $k, m_1, m_2, m_3 \in \mathbb{Z}$ be as given in the following table. When $h \geq k$, we can by Lemma 5.8 find a conjugate (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_\Delta$ with $a'_i \geq m_i$.

m_1	m_2	m_3	k
3	3	4	12
3	4	4	14
4	4	4	16
4	4	5	18
4	5	5	20

Since $x'_1 x'_2 x'_3 = \alpha$, each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$|\alpha| \leq (e^{\pi|\Delta|^{1/2}/m_1} + 2079)(e^{\pi|\Delta|^{1/2}/m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079).$$

For (m_1, m_2, m_3) as in the above table, these bounds are incompatible with the earlier lower bound for $|\alpha|$ when $|\Delta|$ is suitably large. Explicitly, we obtain that one of the following holds:

- (1) $4 \leq h \leq 11$;
- (2) $12 \leq h \leq 13$ and $|\Delta| \leq 30339$;¹
- (3) $14 \leq h \leq 15$ and $|\Delta| \leq 4124$;
- (4) $16 \leq h \leq 17$ and $|\Delta| \leq 1045$;
- (5) $18 \leq h \leq 19$ and $|\Delta| \leq 488$;
- (6) $20 \leq h$ and $|\Delta| \leq 334$.

5.4.1.2. *The subcase where the Δ_i are not all equal.* Without loss of generality assume that $|\Delta_1| > |\Delta_2|$. Then $\mathbb{Q}(x_3) = \mathbb{Q}(x_1x_2) = \mathbb{Q}(x_1, x_2)$, where the last equality holds by Lemma 5.12 since $\Delta_1 \neq \Delta_2$. Thus $\mathbb{Q}(x_1), \mathbb{Q}(x_2) \subset \mathbb{Q}(x_3)$. Since $h_1 = h_2 = h_3$, these inclusions are in fact equalities. Thus $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) = \mathbb{Q}(x_3)$. Denote this field L .

Suppose first that $\mathbb{Q}(\sqrt{\Delta_i}) \neq \mathbb{Q}(\sqrt{\Delta_j})$ for some i, j . Then by (1) of Lemma 5.9, we have that the field L is listed in [2, Table 4.1], as are the possible discriminants $\Delta_1, \Delta_2, \Delta_3$.

So we reduce to the situation where $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2}) = \mathbb{Q}(\sqrt{\Delta_3})$. Then by (2) of Lemma 5.9 either $h = 1$ or, for every i, j , we have that $\Delta_i/\Delta_j \in \{1, 4, 1/4\}$. Since $h \geq 2$, we must be in the second case. Write $\Delta = \Delta_2$. Then we have that $\Delta \equiv 1 \pmod{8}$, $\Delta_1 = 4\Delta_2 = 4\Delta$, and either $\Delta_3 = \Delta$ or $\Delta_3 = 4\Delta = \Delta_1$. Also, $|\Delta_1| \geq 23$ since $\Delta_1 = 4\Delta_2$ and $h \geq 2$ implies $|\Delta_2| \geq 15$.

¹We observe that the bound on $|\Delta|$ obtained here is superfluous, since in fact $h \leq 13$ already implies $|\Delta| \leq 20563$, as may be demonstrated in Sage [76] using the function `cm_orders()`.

Suppose first that $\Delta_3 = \Delta$. Taking conjugates, assume that x_1 is dominant. Then by the bounds in Section 5.3 we have the lower bound

$$\begin{aligned} |\alpha| &\geq (0.9994e^{\pi|\Delta_1|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500|\Delta_2|^{-3}\}) \\ &\quad (\min\{4.4 \times 10^{-5}, 3500|\Delta_3|^{-3}\}), \\ &\geq (0.9994e^{2\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\})^2. \end{aligned}$$

Since $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) = \mathbb{Q}(x_3)$, the Galois orbit of (x_1, x_2, x_3) has exactly h elements. Each conjugate of x_i occurs exactly once as the i th coordinate of a conjugate (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) . Since $\Delta \equiv 1 \pmod{8}$ and $\Delta_1 = 4\Delta$, there are no tuples $(a, b, c) \in T_{\Delta_1}$ with $a = 2$ because $\Delta_1 \equiv 4 \pmod{16}$. Let $k, m_1, m_2, m_3 \in \mathbb{Z}$ be as given in the following table. When $h \geq k$, we can by Lemma 5.8 find a conjugate (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$.

m_1	m_2	m_3	k
3	2	2	4
3	3	2	6
3	3	3	8

Since $x'_1 x'_2 x'_3 = \alpha$, each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$\begin{aligned} |\alpha| &\leq (e^{\pi|\Delta_1|^{1/2}/m_1} + 2079)(e^{\pi|\Delta_2|^{1/2}/m_2} + 2079)(e^{\pi|\Delta_3|^{1/2}/m_3} + 2079) \\ &= (e^{2\pi|\Delta|^{1/2}/m_1} + 2079)(e^{\pi|\Delta|^{1/2}/m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079). \end{aligned}$$

For (m_1, m_2, m_3) as in the above table, these bounds are incompatible with the earlier lower bound for $|\alpha|$ when $|\Delta|$ is suitably large. We obtain therefore that one of the following must hold:

- (1) $2 \leq h \leq 3$;
- (2) $4 \leq h \leq 5$ and $|\Delta| \leq 367$;
- (3) $6 \leq h \leq 7$ and $|\Delta| \leq 163$;
- (4) $8 \leq h$ and $|\Delta| \leq 93$.

Now suppose that $\Delta_1 = \Delta_3 = 4\Delta_2$, where $\Delta_2 = \Delta \equiv 1 \pmod{8}$. Taking conjugates, assume that x_1 is dominant. Then by the bounds in Section 5.3 we have the lower bound

$$\begin{aligned} |\alpha| &\geq (0.9994e^{\pi|\Delta_1|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500|\Delta_2|^{-3}\}) \\ &\quad (\min\{4.4 \times 10^{-5}, 3500|\Delta_3|^{-3}\}), \\ &\geq (0.9994e^{2\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}) \\ &\quad (\min\{4.4 \times 10^{-5}, 3500 \times 4^{-3}|\Delta|^{-3}\}). \end{aligned}$$

Since $\Delta \equiv 1 \pmod{8}$ and $\Delta_1 = \Delta_3 = 4\Delta$, as before there are no tuples $(a, b, c) \in T_{\Delta_1} = T_{\Delta_3}$ with $a = 2$. Therefore, when $h \geq k$, we can by Lemma 5.8 find conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$, and $k, m_1, m_2, m_3 \in \mathbb{Z}$ given in the following table.

m_1	m_2	m_3	k
3	2	3	4
3	3	3	6
4	3	3	8
4	3	4	10

Each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$\begin{aligned} |\alpha| &\leq (e^{\pi|\Delta_1|^{1/2}/m_1} + 2079)(e^{\pi|\Delta_2|^{1/2}/m_2} + 2079)(e^{\pi|\Delta_3|^{1/2}/m_3} + 2079) \\ &= (e^{2\pi|\Delta|^{1/2}/m_1} + 2079)(e^{\pi|\Delta|^{1/2}/m_2} + 2079)(e^{2\pi|\Delta|^{1/2}/m_3} + 2079). \end{aligned}$$

We thus obtain that one of the following holds:

- (1) $2 \leq h \leq 3$;
- (2) $4 \leq h \leq 5$ and $|\Delta| \leq 5781$;
- (3) $6 \leq h \leq 7$ and $|\Delta| \leq 650$;
- (4) $8 \leq h \leq 9$ and $|\Delta| \leq 192$;
- (5) $10 \leq h$ and $|\Delta| \leq 92$.

5.4.2. The case $h_1 = h_2 = 2h_3$. Since $h_3 \neq h_1, h_2$, we have that $\Delta_3 \neq \Delta_1, \Delta_2$. Then $\mathbb{Q}(x_2) = \mathbb{Q}(x_1x_3) = \mathbb{Q}(x_1, x_3)$, where the last equality holds by Lemma 5.12 since $\Delta_1 \neq \Delta_3$. Hence, $\mathbb{Q}(x_1) \subset \mathbb{Q}(x_2)$. Since $h_1 = h_2$, this is in fact an equality $\mathbb{Q}(x_1) = \mathbb{Q}(x_2)$.

Suppose $\Delta_1 \neq \Delta_2$. Then $\mathbb{Q}(x_3) = \mathbb{Q}(x_1x_2) = \mathbb{Q}(x_1, x_2)$ by Lemma 5.12. But then $\mathbb{Q}(x_1) \subset \mathbb{Q}(x_3)$, and so $h_1 = [\mathbb{Q}(x_1) : \mathbb{Q}] \leq [\mathbb{Q}(x_3) : \mathbb{Q}] = h_3$. This though is a contradiction as $h_3 < h_1$ by assumption.

So we must have that $\Delta_1 = \Delta_2$. Note also that $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) \supset \mathbb{Q}(x_3)$. Since $h_1 = 2h_3$, one therefore has that $[\mathbb{Q}(x_1) : \mathbb{Q}(x_3)] = 2$.

If $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_3})$, then by Lemma 5.10 one has that either $x_3 \in \mathbb{Q}$ or $\Delta_1 \in \{9\Delta_3/4, 4\Delta_3, 9\Delta_3, 16\Delta_3\}$. The former cannot happen since $h_3 \geq 2$, so we must have that $\Delta_1 \in \{9\Delta_3/4, 4\Delta_3, 9\Delta_3, 16\Delta_3\}$. Note also that $h_1 \geq 4$ and so certainly $|\Delta_1| \geq 23$.

Suppose first that $\Delta_1 = \Delta_2 = 9\Delta_3/4$ and write $\Delta = \Delta_3$. We may assume that x_1 is dominant, and so obtain the lower bound

$$|\alpha| \geq (0.9994e^{3\pi|\Delta|^{1/2}/2})(\min\{4.4 \times 10^{-5}, 3500 \times \left(\frac{9}{4}\right)^{-3} |\Delta|^{-3}\}) \\ (\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}).$$

Since $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) \supset \mathbb{Q}(x_3)$, there are exactly h_1 conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) . Each conjugate of x_1, x_2 occurs exactly once as the coordinate x'_1, x'_2 respectively of a conjugate (x'_1, x'_2, x'_3) . Further, each conjugate x'_3 of x_3 must appear at least once among the conjugates (x'_1, x'_2, x'_3) .

Let $k, m_1, m_2, m_3 \in \mathbb{Z}$ be as given in the following table. When $h_3 \geq k$, we can, by Lemma 5.8 as usual, find conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$.

m_1	m_2	m_3	k
3	3	3	10
4	4	2	12
4	4	3	14
4	4	4	16

Each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$\begin{aligned} |\alpha| &\leq (e^{\pi|\Delta_1|^{1/2}/m_1} + 2079)(e^{\pi|\Delta_2|^{1/2}/m_2} + 2079)(e^{\pi|\Delta_3|^{1/2}/m_3} + 2079) \\ &= (e^{3\pi|\Delta|^{1/2}/2m_1} + 2079)(e^{3\pi|\Delta|^{1/2}/2m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079). \end{aligned}$$

For (m_1, m_2, m_3) as in the above table, these bounds are incompatible with the earlier lower bound for $|\alpha|$ when $|\Delta|$ is suitably large. Explicitly, we obtain that one of the following holds:

- (1) $2 \leq h_3 \leq 9$;
- (2) $10 \leq h_3 \leq 11$ and $|\Delta| \leq 5076$;
- (3) $12 \leq h_3 \leq 13$ and $|\Delta| \leq 1430$;
- (4) $14 \leq h_3 \leq 15$ and $|\Delta| \leq 255$;
- (5) $16 \leq h_3$ and $|\Delta| \leq 164$.

Suppose next that $\Delta_1 = \Delta_2 = 4\Delta_3$ and write $\Delta = \Delta_3$. We may assume that x_1 is dominant, and so obtain the lower bound

$$\begin{aligned} |\alpha| &\geq (0.9994e^{2\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500 \times 4^{-3}|\Delta|^{-3}\}) \\ &\quad (\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}). \end{aligned}$$

As before, since $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) \supset \mathbb{Q}(x_3)$, there are exactly h_1 conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) . Each conjugate of x_1, x_2 occurs exactly once as the coordinate x'_1, x'_2 respectively of a conjugate (x'_1, x'_2, x'_3) . Further, each conjugate x'_3 of x_3 appears at least once among the conjugates (x'_1, x'_2, x'_3) .

Let $k, m_1, m_2, m_3 \in \mathbb{Z}$ be as given in the following table. When $h_3 \geq k$, we can by Lemma 5.8 as usual, find conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) ,

where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$.

m_1	m_2	m_3	k
3	3	3	10
3	3	4	12
3	3	5	14
3	4	4	16

Each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$\begin{aligned} |\alpha| &\leq (e^{\pi|\Delta_1|^{1/2}/m_1} + 2079)(e^{\pi|\Delta_2|^{1/2}/m_2} + 2079)(e^{\pi|\Delta_3|^{1/2}/m_3} + 2079) \\ &= (e^{2\pi|\Delta|^{1/2}/m_1} + 2079)(e^{2\pi|\Delta|^{1/2}/m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079). \end{aligned}$$

These bounds are incompatible with the above lower bound for $|\alpha|$ when $|\Delta|$ is large. Hence, we must have one of:

- (1) $2 \leq h_3 \leq 9$;
- (2) $10 \leq h_3 \leq 11$ and $|\Delta| \leq 650$;
- (3) $12 \leq h_3 \leq 13$ and $|\Delta| \leq 317$;
- (4) $14 \leq h_3 \leq 15$ and $|\Delta| \leq 236$;
- (5) $16 \leq h_3$ and $|\Delta| \leq 129$.

Now suppose that $\Delta_1 = \Delta_2 = 9\Delta_3$ and write $\Delta = \Delta_3$. We may assume that x_1 is dominant, and so obtain the lower bound

$$\begin{aligned} |\alpha| &\geq (0.9994e^{3\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500 \times 9^{-3}|\Delta|^{-3}\}) \\ &\quad (\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}). \end{aligned}$$

As before, there are exactly h_1 conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) and each conjugate of x_1, x_2 occurs exactly once as the coordinate x'_1, x'_2 respectively of a conjugate (x'_1, x'_2, x'_3) . Further, each conjugate x'_3 of x_3 appears at least once among the conjugates (x'_1, x'_2, x'_3) .

As previously, when $h_3 \geq k$, we can find conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$, and $k, m_1, m_2, m_3 \in \mathbb{Z}$ are as given in the following table.

m_1	m_2	m_3	k
3	3	2	8
3	4	2	10

Each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$|\alpha| \leq (e^{3\pi|\Delta|^{1/2}/m_1} + 2079)(e^{3\pi|\Delta|^{1/2}/m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079).$$

For (m_1, m_2, m_3) as in the above table, these bounds are incompatible with the lower bound for $|\alpha|$ when $|\Delta|$ is large. One thus has that:

- (1) $2 \leq h_3 \leq 7$;
- (2) $8 \leq h_3 \leq 9$ and $|\Delta| \leq 255$;
- (3) $10 \leq h_3$ and $|\Delta| \leq 85$.

Finally suppose that $\Delta_1 = \Delta_2 = 16\Delta_3$ and write $\Delta = \Delta_3$. We may assume that x_1 is dominant, and so obtain the lower bound

$$|\alpha| \geq (0.9994e^{4\pi|\Delta|^{1/2}})(\min\{4.4 \times 10^{-5}, 3500 \times 16^{-3}|\Delta|^{-3}\})$$

$$(\min\{4.4 \times 10^{-5}, 3500|\Delta|^{-3}\}).$$

As before, there are exactly h_1 conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) and each conjugate of x_1, x_2 occurs exactly once as the coordinate x'_1, x'_2 respectively of a conjugate (x'_1, x'_2, x'_3) . Further, each conjugate x'_3 of x_3 appears at least once among the conjugates (x'_1, x'_2, x'_3) .

When $h_3 \geq k$, we can then by Lemma 5.8 as usual, find conjugates (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) , where each x'_i is a singular modulus corresponding to a tuple $(a'_i, b'_i, c'_i) \in T_{\Delta_i}$ with $a'_i \geq m_i$, and $k, m_1, m_2, m_3 \in \mathbb{Z}$ are as given in the following table.

m_1	m_2	m_3	k
3	3	2	8
3	3	3	10

Each such conjugate gives rise to an upper bound for $|\alpha|$ of the form

$$|\alpha| \leq (e^{4\pi|\Delta|^{1/2}/m_1} + 2079)(e^{4\pi|\Delta|^{1/2}/m_2} + 2079)(e^{\pi|\Delta|^{1/2}/m_3} + 2079).$$

For (m_1, m_2, m_3) as in the above table, these bounds are incompatible with the earlier lower bound for $|\alpha|$ when $|\Delta|$ is suitably large. Explicitly, we obtain that one of the following holds:

- (1) $2 \leq h_3 \leq 7$;
- (2) $8 \leq h_3 \leq 9$ and $|\Delta| \leq 79$;
- (3) $10 \leq h_3$ and $|\Delta| \leq 52$.

Now consider the case when $\mathbb{Q}(\sqrt{\Delta_1}) \neq \mathbb{Q}(\sqrt{\Delta_3})$. Then by Lemma 5.11 one of the following holds:

- (1) at least one of Δ_1 or Δ_3 is listed in [2, Table 2.1];
- (2) $h_1 \geq 128$.

If Δ_i is listed in [2, Table 2.1], then we can find all possible $(\Delta_1, \Delta_2, \Delta_3)$ satisfying the condition $[\mathbb{Q}(x_1) : \mathbb{Q}(x_3)] = 2$.

So suppose $h_1 \geq 128$. Write $\Delta = \max\{|\Delta_1|, |\Delta_2|, |\Delta_3|\}$. Taking conjugates as necessary, we may assume that x_i is dominant, where $\Delta = |\Delta_i|$. Since $h_i \geq 64$, certainly $|\Delta_i| \geq 23$. Then by the bounds in Section 5.3

$$|\alpha| \geq (0.9994e^{\pi\Delta^{1/2}})(\min\{4.4 \times 10^{-5}, 3500\Delta^{-3}\})^2.$$

Since $h_1 \geq 128$, we can always find a conjugate (x'_1, x'_2, x'_3) of (x_1, x_2, x_3) with the associated a'_1, a'_2, a'_3 satisfying $a'_1, a'_2 \geq 4$ and $a'_3 \geq 5$. This gives rise to the upper bound

$$|\alpha| \leq (e^{\pi\Delta^{1/2}/4} + 2079)(e^{\pi\Delta^{1/2}/4} + 2079)(e^{\pi\Delta^{1/2}/5} + 2079).$$

Together these bounds imply that $\Delta \leq 488$. Hence, we must have that $h_1 \geq 128$ and $|\Delta_1|, |\Delta_3| \leq 488$.

5.4.3. The case $h_1 = 2h_2 = 2h_3$. Since $h_1 \neq h_2, h_3$, one has that $\Delta_1 \neq \Delta_2, \Delta_3$. Therefore, $\mathbb{Q}(x_3) = \mathbb{Q}(x_1x_2) = \mathbb{Q}(x_1, x_2)$. The last equality holds by Lemma 5.12 since $\Delta_1 \neq \Delta_2$. Thus $\mathbb{Q}(x_1) \subset \mathbb{Q}(x_3)$ and so $h_1 = [\mathbb{Q}(x_1) : \mathbb{Q}] \leq [\mathbb{Q}(x_3) : \mathbb{Q}] = h_3$. This though is a contradiction as $h_3 < h_1$ by assumption, and so we may eliminate this case.

5.5. Eliminating non-trivial cases

Recall that we assumed x_1, x_2, x_3 are singular moduli with $x_1x_2x_3 = \alpha \in \mathbb{Q}^\times$. We write Δ_i for their respective discriminants and h_i for the corresponding class numbers $h(\Delta_i)$. Without loss of generality $h_1 \geq h_2 \geq h_3$. Assuming that we are not in one of the trivial cases (1)–(3) of Theorem 5.7,

then we have shown in Section 5.4 that x_1, x_2, x_3 are pairwise distinct and that we must be in one of the following cases.

(1) $h_1 = h_2 = h_3$.

Write $h = h_1 = h_2 = h_3$.

(a) $\Delta_1 = \Delta_2 = \Delta_3$.

Write $\Delta = \Delta_1 = \Delta_2 = \Delta_3$.

(i) $4 \leq h \leq 11$.

(ii) $12 \leq h \leq 13$ and $|\Delta| \leq 30339$.

(iii) $14 \leq h \leq 15$ and $|\Delta| \leq 4124$.

(iv) $16 \leq h \leq 17$ and $|\Delta| \leq 1045$.

(v) $18 \leq h \leq 19$ and $|\Delta| \leq 488$.

(vi) $20 \leq h$ and $|\Delta| \leq 334$.

(b) $|\Delta_1| > |\Delta_2|$.

In this case, $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) = \mathbb{Q}(x_3)$.

(i) $\mathbb{Q}(\sqrt{\Delta_i}) \neq \mathbb{Q}(\sqrt{\Delta_j})$ for some i, j .

The list of all possible $\Delta_1, \Delta_2, \Delta_3$ is given in [2, Table 4.1].

(ii) $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_2}) = \mathbb{Q}(\sqrt{\Delta_3})$.

In this case, $\Delta_1 = 4\Delta_2$, $\Delta_3 \in \{\Delta_1, \Delta_2\}$, and $\Delta_2 \equiv 1 \pmod{8}$.

(A) $\Delta_3 = \Delta_2$.

Write $\Delta = \Delta_2 = \Delta_3$.

(I) $2 \leq h \leq 3$.

(II) $4 \leq h \leq 5$ and $|\Delta| \leq 367$.

(III) $6 \leq h \leq 7$ and $|\Delta| \leq 163$.

(IV) $8 \leq h$ and $|\Delta| \leq 93$.

(B) $\Delta_3 = \Delta_1$.

Write $\Delta = \Delta_2$.

(I) $2 \leq h \leq 3$.

(II) $4 \leq h \leq 5$ and $|\Delta| \leq 5781$.

(III) $6 \leq h \leq 7$ and $|\Delta| \leq 650$.

(IV) $8 \leq h \leq 9$ and $|\Delta| \leq 192$.

(V) $10 \leq h$ and $|\Delta| \leq 92$.

(2) $h_1 = h_2 = 2h_3$.

In this case, $\mathbb{Q}(x_1) = \mathbb{Q}(x_2) \supset \mathbb{Q}(x_3)$ and $[\mathbb{Q}(x_1) : \mathbb{Q}(x_3)] = 2$.

(a) $\Delta_1 \neq \Delta_2$.

This case cannot arise.

(b) $\Delta_1 = \Delta_2$.

(i) $\mathbb{Q}(\sqrt{\Delta_1}) = \mathbb{Q}(\sqrt{\Delta_3})$.

(A) $\Delta_1 = \Delta_2 = 9\Delta_3/4$.

Write $\Delta = \Delta_3$.

(I) $2 \leq h_3 \leq 9$.

(II) $10 \leq h_3 \leq 11$ and $|\Delta| \leq 5076$.

(III) $12 \leq h_3 \leq 13$ and $|\Delta| \leq 1430$.

(IV) $14 \leq h_3 \leq 15$ and $|\Delta| \leq 255$.

(V) $16 \leq h_3$ and $|\Delta| \leq 164$.

(B) $\Delta_1 = \Delta_2 = 4\Delta_3$.

Write $\Delta = \Delta_3$.

(I) $2 \leq h_3 \leq 9$;

(II) $10 \leq h_3 \leq 11$ and $|\Delta| \leq 650$.

(III) $12 \leq h_3 \leq 13$ and $|\Delta| \leq 317$.

(IV) $14 \leq h_3 \leq 15$ and $|\Delta| \leq 236$.

(V) $16 \leq h_3$ and $|\Delta| \leq 129$.

(C) $\Delta_1 = \Delta_2 = 9\Delta_3$.

Write $\Delta = \Delta_3$.

(I) $2 \leq h_3 \leq 7$.

(II) $8 \leq h_3 \leq 9$ and $|\Delta| \leq 255$.

(III) $10 \leq h_3$ and $|\Delta| \leq 85$.

(D) $\Delta_1 = 16\Delta_3$.

Write $\Delta = \Delta_3$.

(I) $2 \leq h_3 \leq 7$.

(II) $8 \leq h_3 \leq 9$ and $|\Delta| \leq 79$.

(III) $10 \leq h_3$ and $|\Delta| \leq 52$.

(ii) $\mathbb{Q}(\sqrt{\Delta_1}) \neq \mathbb{Q}(\sqrt{\Delta_3})$.

(A) Δ_1 is listed in [2, Table 2.1].

(B) Δ_3 is listed in [2, Table 2.1].

(C) $h_1 \geq 128$ and $|\Delta_1|, |\Delta_3| \leq 488$.

(3) $h_1 = 2h_2 = 2h_3$.

This case cannot arise.

The finite list of discriminants $(\Delta_1, \Delta_2, \Delta_3)$ satisfying one of these conditions may be computed in PARI. In fact, there are 2888 triples $(\Delta_1, \Delta_2, \Delta_3)$ satisfying one of the above conditions. We now show how each such choice of $(\Delta_1, \Delta_2, \Delta_3)$ may be eliminated by another computation in PARI.

For each tuple $(\Delta_1, \Delta_2, \Delta_3)$ satisfying one of the above conditions, we show that $x_1x_2x_3 \notin \mathbb{Q}$, for any choice of x_1, x_2, x_3 pairwise distinct singular moduli of respective discriminant Δ_i . (In fact, by taking conjugates as

necessary, it is enough to eliminate all possible choices of x_2, x_3 for some fixed x_1 .) To do this, we use the following algorithm.

For each possible choice of (x_1, x_2, x_3) , let L be a number field containing all conjugates of x_1, x_2, x_3 . If $x_1x_2x_3 \in \mathbb{Q}^\times$, then

$$\frac{x_1x_2x_3}{\sigma(x_1)\sigma(x_2)\sigma(x_3)} = 1$$

for every automorphism $\sigma \in \text{Gal}(L/\mathbb{Q})$. It therefore suffices to find an automorphism of L without this property in order to eliminate the tuple (x_1, x_2, x_3) . Once all such tuples (x_1, x_2, x_3) have been eliminated, we can eliminate the tuple $(\Delta_1, \Delta_2, \Delta_3)$.

It remains to find a suitable field L . In 1(a), the x_i are all conjugate, so $L = \mathbb{Q}(\sqrt{\Delta_1}, x_1)$ suffices. In 1(b)(i), the field $\mathbb{Q}(x_1)$ is Galois, and so we may take $L = \mathbb{Q}(x_1)$. In 1(b)(ii), $L = \mathbb{Q}(\sqrt{\Delta_1}, x_1)$ works. In 2(b)(i), let $L = \mathbb{Q}(\sqrt{\Delta_1}, x_1)$. In 2(b)(ii)(A), the field $\mathbb{Q}(x_1)$ is Galois and so we may take $L = \mathbb{Q}(x_1)$. In 2(b)(ii)(B), the field $\mathbb{Q}(x_3)$ is Galois and so $L = \mathbb{Q}(\sqrt{\Delta_1}, x_1)$ suffices. There are no $(\Delta_1, \Delta_2, \Delta_3)$ satisfying the conditions in 2(b)(ii)(C), so this case may be excluded.

We implement the resulting algorithm using a PARI script. Running it, we are able to eliminate each of the above possibilities. The proof of Theorem 5.7 is thus complete. Total running time of our program is about 12 hours on a standard laptop computer.² The time taken to find and eliminate a triple of discriminants $(\Delta_1, \Delta_2, \Delta_3)$ satisfying one of the above conditions increases with the respective class numbers h_1, h_2, h_3 . Consequently, approximately 75% of the overall run time of the program is spent

²With a 2.5GHz Intel i5 processor and 8GB RAM.

dealing with case (2)(b)(ii)(B), which includes discriminants Δ_1 of class number 32, greater than in any other case.

APPENDIX A

O-minimality

In several places in this thesis, our proofs make use of results from o-minimality. We collect in this appendix all those results which we use, as an aid to the reader unfamiliar with o-minimality. We assume basic material from model theory, in particular the definitions of a structure and a definable set. Such background may be found in any introduction to model theory, for example [47]. By definable, we always mean definable with parameters.

A.1. O-minimal structures

Definition A.1. Let \mathcal{L} be a (first-order) language containing a binary relation $<$, plus possibly some other symbols. Let $\mathcal{A} = (A; <, \dots)$ be an \mathcal{L} -structure such that $<$ is a dense linear order without endpoints on the domain A . We say that \mathcal{A} is an o-minimal structure if every definable subset of A is a finite union of points and intervals.

Intervals are always open. We allow the case of intervals with endpoints at $\pm\infty$, where we adopt the obvious convention that $-\infty < a < \infty$ for every $a \in A$. One can consider o-minimal structures where the underlying order is not dense, but, as we will not need any such examples, we restrict to the dense case for convenience.

From now on, all languages we consider will extend $\{<\}$ and all structures we consider will interpret $<$ as a dense linear order without endpoints

on our domain A . Since $<$ is an order on A , we may endow the set A with a topology by taking the open intervals in A as a basis for the open sets. We then take the corresponding product topology for each A^n . Topological notions (e.g. continuity, connectedness) will always be with respect to this topology.

O-minimal structures have the following remarkable property.

Theorem A.2. *Let \mathcal{A}, \mathcal{B} be elementarily equivalent \mathcal{L} -structures. Then \mathcal{A} is o-minimal if and only if \mathcal{B} is o-minimal.*

The motivating example of an o-minimal structure is the real closed field. Fix a language \mathcal{L} with a single binary relation symbol, two binary function symbols, and two constant symbols. Let $\overline{\mathbb{R}} = (\mathbb{R}; <, +, \cdot, 0, 1)$ be the \mathcal{L} -structure given by the ordered field of the real numbers, so that $<$ is the standard order on \mathbb{R} and $+, \cdot, 0, 1$ have their usual interpretations.

A semialgebraic subset of \mathbb{R}^n is defined to be a finite union of subsets of \mathbb{R}^n defined by formulas of the form $p(x_1, \dots, x_n) = 0$ and $q(x_1, \dots, x_n) > 0$, where $p, q \in \mathbb{R}[X_1, \dots, X_n]$. Observe that the collection of semialgebraic sets in \mathbb{R}^n is obviously closed under taking complements, finite unions, and finite intersections. The semialgebraic subsets of \mathbb{R} are precisely the finite Boolean combinations of sets of the form

$$\{x : p(x) = 0\}$$

and

$$\{x : q(x) > 0\},$$

where $p, q \in \mathbb{R}[X]$. The Fundamental Theorem of Algebra shows that sets of the first form are either finite or all of \mathbb{R} , while those of the second form are finite unions of open intervals.

Theorem A.3 (Tarski–Seidenberg theorem). *Let $\pi: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ be the projection map onto the first n coordinates. Let $X \subset \mathbb{R}^{n+1}$ be a semialgebraic subset of \mathbb{R}^{n+1} . Then the set $\pi(X)$ is a semialgebraic subset of \mathbb{R}^n .*

This implies that $A \subset \mathbb{R}^n$ is definable in the structure $\overline{\mathbb{R}}$ if and only if A is semialgebraic. The Tarski–Seidenberg theorem may thus be taken (anachronistically) to say that the structure $\overline{\mathbb{R}}$ is o-minimal.

Van den Dries had noted [26] that many of the useful model theoretic properties of semialgebraic sets followed from a few basic axioms and so suggested that considering other structures on \mathbb{R} satisfying these axioms was a promising avenue of investigation. This led to the definition of an o-minimal structure. The basic concepts of o-minimality were worked out in a series of papers by Pillay and Steinhorn [68, 69], and also together with Knight [42].

A.2. The tameness of o-minimal structures

Fix an o-minimal structure $\mathcal{A} = (A; <, \dots)$. By “definable”, we mean definable in \mathcal{A} . In this section, we show that the definable subsets of A^n have some very nice tameness properties.

Theorem A.4 (Uniform Finiteness Theorem, [28, p. 53]). *Let $X \subset A^{n+1}$ be a definable set. Suppose that for every $a \in A^n$ the fibre*

$$X_a = \{b \in A : (a, b) \in X\}$$

is a finite set. Then there exists $N \geq 1$ such that $|X_a| \leq N$ for every $a \in A^n$.

Since \mathcal{A} is o-minimal, we know that the definable subsets of A are very simple; precisely, they are the finite unions of points and intervals. O-minimal structures have the particularly nice property that their definable sets in higher dimensions are also well-behaved. To make this precise, we introduce the notion of cells.

Definition A.5. For the structure \mathcal{A} , cells are defined [28, pp. 3–4] inductively:

- (1) The cells in A are precisely the points $\{a\}$ and the open intervals (a, b) .
- (2) Let $C \subset A^n$ be a cell and $f : C \rightarrow A$ a definable continuous function, then

$$\Gamma(f) = \{(x, r) \in C \times A : f(x) = r\},$$

$$(-\infty, f) = \{(x, r) \in C \times A : r < f(x)\},$$

$$(f, \infty) = \{(x, r) \in C \times A : f(x) < r\}$$

are cells in A^{n+1} .

- (3) Let $C \subset A^n$ be a cell and $f, g : C \rightarrow A$ definable continuous functions such that $f < g$ on C , then

$$(f, g) = \{(x, r) \in C \times \mathbb{R} : f(x) < r < g(x)\}$$

is a cell in A^{n+1} .

The cells in A are the points and open intervals. To say the structure \mathcal{A} is o-minimal is thus equivalent to saying that every definable set $X \subset A$ is a finite union of cells.

Definition A.6. A decomposition of A^n is defined [28, Definition 2.10] inductively.

- (1) A decomposition of A is a partition of A into finitely many cells.
- (2) A decomposition of A^{n+1} is a partition of A^{n+1} into finitely many cells C_i such that the projections $\pi(C_i)$ are a decomposition of A^n . (Here $\pi: A^{n+1} \rightarrow A^n$ is the projection map to the first n coordinates.)

We then have the following theorem.

Theorem A.7 (Cell Decomposition Theorem, [28, Theorem 2.11]). *Let $X_1, \dots, X_k \subset A^n$ be definable. Then there exists a decomposition of A^n into cells C_i such that each X_j is a union of some of the cells C_i . Further, if $f: X \rightarrow A$ is a definable function for some $X \subset A^n$, then there exists a decomposition of A^n into cells C_i such that X is a union of some of the cells C_i and, for each $C_i \subset X$, the restriction $f|_{C_i}$ is continuous.*

We now restrict to the case of o-minimal expansions of the real field $\overline{\mathbb{R}}$. Some o-minimal expansions of $\overline{\mathbb{R}}$ satisfy a stronger tameness property for

definable sets than that given by the Cell Decomposition Theorem. This stronger property is called analytic cell decomposition. For the remainder of this section, fix an o-minimal expansion \mathcal{R} of $\overline{\mathbb{R}}$. By “definable”, we mean definable in the structure \mathcal{R} .

Definition A.8 ([30, 8.4]). A function $f: A \rightarrow \mathbb{R}^m$, where $A \subset \mathbb{R}^n$, is called \mathcal{R} -analytic if A is definable and there exists a definable open neighbourhood $U \subset \mathbb{R}^n$ of A and a definable real analytic function $F: U \rightarrow \mathbb{R}^m$ such that $f = F|_A$.

Definition A.9 ([30, 8.5]). The \mathcal{R} -analytic cells are defined inductively, exactly like the cells of \mathcal{R} are defined in Definition A.5, except that, in parts (2) and (3) of Definition A.5, the cell C is required to be an \mathcal{R} -analytic cell and the functions $f, g: C \rightarrow \mathbb{R}$ are required to be \mathcal{R} -analytic.

Definition A.10 ([30, 8.6]). An \mathcal{R} -analytic decomposition of \mathbb{R}^n is defined inductively.

- (1) An \mathcal{R} -analytic decomposition of \mathbb{R} is a partition of \mathbb{R} into finitely many \mathcal{R} -analytic cells of \mathbb{R} .
- (2) An \mathcal{R} -analytic decomposition of \mathbb{R}^{n+1} is a partition of \mathbb{R}^{n+1} into finitely many \mathcal{R} -analytic cells C_i of \mathbb{R}^{n+1} such that the sets $\pi(C_i)$ form an \mathcal{R} -analytic decomposition of \mathbb{R}^n , where $\pi: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ is the projection map to the first n coordinates.

Definition A.11 ([30, 8.8]). We say that the structure \mathcal{R} admits analytic cell decomposition if the following properties hold.

- (1) Let $X_1, \dots, X_k \subset \mathbb{R}^n$ be definable. Then there exists an \mathcal{R} -analytic decomposition of \mathbb{R}^n into some \mathcal{R} -analytic cells C_i such that each X_j is the union of some of the C_i .
- (2) If $f : X \rightarrow \mathbb{R}$ is a definable function for some $X \subset \mathbb{R}^n$, then there exists an \mathcal{R} -analytic decomposition of \mathbb{R}^n into some \mathcal{R} -analytic cells C_i such that X is a union of some of the cells C_i and, for each $C_i \subset X$, the restricted function $f|_{C_i}$ is \mathcal{R} -analytic.

A.3. The Pila–Wilkie Counting Theorem

In this section, we discuss the various versions of the Pila–Wilkie Counting Theorem for o-minimal structures which we make use of in this thesis. For the remainder of this section, fix an o-minimal expansion of $\overline{\mathbb{R}}$; references to “definable” mean definable in this structure.

Definition A.12. Let $Z \subset \mathbb{R}^n$ and $k \geq 1$. For $T > 0$, we define

$$Z(k, T) = \{\bar{x} \in Z : [\mathbb{Q}(x_i) : \mathbb{Q}] \leq k, H(\bar{x}) \leq T\}$$

and

$$N_k(Z, T) = |Z(k, T)|.$$

Given a definable set Z , we would like a bound on the counting function $N_k(Z, T)$. Of course, it may be that Z contains lots of rational (or algebraic) points, for example if Z contains a segment of an algebraic curve. We therefore make the following definition.

Definition A.13. For $Z \subset \mathbb{R}^n$, the algebraic part Z^{alg} of Z is the union of all connected semialgebraic subsets of Z of positive dimension.

Pila's strengthening of the Counting Theorem (the version of [66] just treated the case $k = 1$) is the following.

Theorem A.14 ([59, Theorem 1.6]). *Let $Z \subset \mathbb{R}^n$ be definable, $k \geq 1$, and $\epsilon > 0$. There exists a constant $c(Z, k, \epsilon) > 0$ such that*

$$N_k(Z \setminus Z^{\text{alg}}, T) \leq c(Z, k, \epsilon)T^\epsilon$$

for all $T > 0$.

There is also a version of the theorem for definable families.

Definition A.15. A definable family is a definable set $X \subset \mathbb{R}^n \times \mathbb{R}^m$ considered as the family of fibres $X_{\bar{a}} = \{\bar{x} \in \mathbb{R}^n : (\bar{x}, \bar{a}) \in X\}$ for $\bar{a} \in \mathbb{R}^m$.

Note that, for a definable family $X \subset \mathbb{R}^n \times \mathbb{R}^m$, we always consider the fibres $X_{\bar{a}}$, for $\bar{a} \in \mathbb{R}^m$, as subsets of \mathbb{R}^n . When we refer to points on $X_{\bar{a}}$ of a given height and degree then, this is only with respect to the \mathbb{R}^n coordinates (and does not imply anything about the coordinates of \bar{a} itself). The Counting Theorem for definable families is the following statement.

Theorem A.16 ([59, Theorem 5.3]). *Let $X \subset \mathbb{R}^n \times \mathbb{R}^m$ be a definable family, $k \geq 1$, and $\epsilon > 0$. There exists a constant $c(Z, k, \epsilon) > 0$ such that*

$$N_k(X_{\bar{a}} \setminus X_{\bar{a}}^{\text{alg}}, T) \leq c(Z, k, \epsilon)T^\epsilon$$

for all $T > 0$ and $\bar{a} \in \mathbb{R}^m$.

There are also two refinements of the Counting Theorem which we use. To state the first of these, we introduce "blocks".

Definition A.17 ([59, Definition 3.2]).

- (1) A basic block of dimension k in \mathbb{R}^n is a connected definable set $U \subset \mathbb{R}^n$ of dimension k which is contained in a semialgebraic set $A \subset \mathbb{R}^n$ of dimension k such that, for every $x \in U$, the point x is a regular point of dimension k in both U and A .
- (2) A basic block family is a definable family Z such that each fibre of Z is a basic block.

Theorem A.18 ([59, Theorem 3.5]). *Let $Z \subset \mathbb{R}^n \times \mathbb{R}^m$ be a definable family and $\epsilon > 0$. Denote by Z_y the fibre of Z at $y \in \mathbb{R}^m$. There exists a constant $J = J(Z, \epsilon) \in \mathbb{Z}_{>0}$ and a collection of basic block families $W_j \subset \mathbb{R}^n \times (\mathbb{R}^m \times \mathbb{R}^{m_j})$ for $j = 1, \dots, J$ such that:*

- (1) *For each $j = 1, \dots, J$, there is some number w_j such that each point in each fibre of W_j is regular of dimension w_j .*
- (2) *For each $j = 1, \dots, J$ and $(y, z) \in \mathbb{R}^m \times \mathbb{R}^{m_j}$, the fibre of W_j at (y, z) is contained in the fibre of Z at y .*
- (3) *For each $y \in \mathbb{R}^m$ and all $T \geq 1$, the set $Z_y(1, T)$ is contained in the union of $O_{Z, \epsilon}(T^\epsilon)$ basic blocks, each of which is a fibre of one of the W_j at some $(y, z) \in \mathbb{R}^m \times \mathbb{R}^{m_j}$.*
- (4) *Let $W \subset \mathbb{R}^n \times \mathbb{R}^m$ be the family whose fibre at $y \in \mathbb{R}^m$ is the union over all the j with $w_j > 0$ of all the fibres of the W_j over (y, z) for some $z \in \mathbb{R}^{m_j}$. Then W is definable. Further, if we write X_ϵ for the fibre of W at y , then $X_\epsilon \subset Z_y^{\text{alg}}$ and $N_1(Z_y \setminus X_\epsilon, T) = O_{Z, \epsilon}(T^\epsilon)$.*

We also use a slightly modified version of the Counting Theorem, due to Habegger and Pila [39]. To state it, we introduce a new height function.

Definition A.19. Let $k \geq 1$. The k -height $H_k(y)$ of $y \in \mathbb{R}$ is defined to be the minimum of $\max\{|a_0|, \dots, |a_k|\}$ taken over all possible choices

of $(a_0, \dots, a_k) \in \mathbb{Z}^k \setminus \{(0, \dots, 0)\}$ such that a_0, \dots, a_k are coprime and $a_k y^k + \dots + a_0 = 0$. (We adopt the convention that $\min \emptyset = \infty$.) For a tuple $y = (y_1, \dots, y_n) \in \mathbb{R}^n$, we define the k -height of y to be $H_k(y) = \max\{H_k(y_i)\}$.

Note that $y \in \mathbb{R}$ has finite k -height if and only if $[\mathbb{Q}(k) : \mathbb{Q}] \leq k$. In this case, we may relate the k -height of y to the height of y as follows.

Proposition A.20. *Let $k \geq 1$ and $y \in \mathbb{R}$ such that $[\mathbb{Q}(y) : \mathbb{Q}] \leq k$. Then $H_k(y) \leq 2^k H(y)$.*

PROOF. See [59, §5]. □

We may now give the modified Counting Theorem.

Theorem A.21 ([39, Corollary 7.2]). *Let $F \subset \mathbb{R}^l \times \mathbb{R}^m \times \mathbb{R}^n$ be a definable family parametrised by \mathbb{R}^l . Let π_1, π_2 denote the projections $\mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $\mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ respectively. Let $\epsilon > 0$ and $k \in \mathbb{Z}_{>0}$. Then there exists a constant $c = c(F, \epsilon, k)$ with the following property. Say $x \in \mathbb{R}^l$ and denote by Z the fibre of F over x . Let $T \geq 1$ and define*

$$\tilde{Z}(k, T) = \{(y, z) \in Z : H_k(y) \leq T\}.$$

If $\Sigma \subset \tilde{Z}(k, T)$ is such that $|\pi_2(\Sigma)| > cT^\epsilon$, then there exists a continuous, definable function $\beta: [0, 1] \rightarrow Z$ such that:

- (1) $\pi_1 \circ \beta: [0, 1] \rightarrow \mathbb{R}^m$ is semialgebraic and its restriction to $(0, 1)$ is real analytic;
- (2) $\pi_2 \circ \beta: [0, 1] \rightarrow \mathbb{R}^n$ is non-constant;
- (3) $\beta(0) \in \Sigma$;

(4) and if the o-minimal structure admits analytic cell decomposition, then the restriction of β to $(0, 1)$ is real analytic.

Note that although the authors state only that $\pi_2(\beta(0)) \in \pi_2(\Sigma)$ in [39], they in fact prove the stronger fact that $\beta(0) \in \Sigma$.

A.4. The o-minimal structure $\mathbb{R}_{\text{an,exp}}$

So far the only o-minimal structure we have seen is $\overline{\mathbb{R}}$. In this section, we describe the o-minimal structure $\mathbb{R}_{\text{an,exp}}$, which we use in all our applications in this thesis.

By the definition of o-minimality, any infinite definable set in an o-minimal structure must contain an open interval. The set

$$\{x : \sin x = 0\},$$

for example, is an infinite discrete subset of \mathbb{R} . Hence, no o-minimal structure on \mathbb{R} may have the sine function definable. Consequently, there is no o-minimal structure which has all real analytic functions definable in it. There are though o-minimal structures in which all restricted analytic functions are definable.

Let $\mathbb{R}\{X_1, \dots, X_n\}$ be the ring of formal power series which converge in some neighbourhood of $[-1, 1]^n$. For every $f \in \mathbb{R}\{X_1, \dots, X_n\}$, let the function $\tilde{f}: \mathbb{R}^n \rightarrow \mathbb{R}$ be defined by

$$\tilde{f}(\bar{x}) = \begin{cases} f(\bar{x}) & \text{if } \bar{x} \in [-1, 1]^n \\ 0 & \text{if } \bar{x} \notin [-1, 1]^n. \end{cases}$$

Such a function \tilde{f} is called a restricted analytic function. The structure \mathbb{R}_{an} is $\overline{\mathbb{R}}$ together with every restricted analytic function \tilde{f} . The o-minimality of \mathbb{R}_{an} was observed by van den Dries [27].

In fact, some global analytic functions are definable in o-minimal structures. The real exponential function $\exp: \mathbb{R} \rightarrow \mathbb{R}$ is an example of an (unrestricted) real analytic function definable in an o-minimal structure. Let $\mathbb{R}_{\text{exp}} = (\mathbb{R}; <, +, \cdot, \exp, 0, 1)$, where $<, +, \cdot, 0, 1$ have their usual interpretations and \exp is the standard real exponential function $\exp: \mathbb{R} \rightarrow \mathbb{R}$. Tarski asked the (still unanswered) question of whether the theory of the structure \mathbb{R}_{exp} is decidable. Indeed, it was work on this problem which motivated van den Dries [26] in his work which led to the definition of o-minimality. That the structure \mathbb{R}_{exp} is o-minimal was proved by Wilkie [81], as part of his own work on Tarski's question.

The o-minimal structure we use in all our applications is $\mathbb{R}_{\text{an,exp}}$. This is the expansion of $\overline{\mathbb{R}}$ by the real exponential function $\exp: \mathbb{R} \rightarrow \mathbb{R}$ and all the restricted analytic functions. Every set definable in either \mathbb{R}_{exp} or \mathbb{R}_{an} is thus definable in $\mathbb{R}_{\text{an,exp}}$ as well. The o-minimality of $\mathbb{R}_{\text{an,exp}}$ was proved in [29]. Van den Dries and Miller [30, §8] proved that $\mathbb{R}_{\text{an,exp}}$ admits analytic cell decomposition.

Every modular function, restricted to a suitable domain, is definable in $\mathbb{R}_{\text{an,exp}}$. Let $f: \mathbb{H} \rightarrow \mathbb{C}$ be a non-constant modular function, i.e. a rational function of the j -function $j: \mathbb{H} \rightarrow \mathbb{C}$. We identify (subsets of) \mathbb{C} with (subsets of) \mathbb{R}^2 by taking real and imaginary parts. The function $f: \mathbb{H} \rightarrow \mathbb{C}$ is thus identified with a map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$.

The $\text{SL}_2(\mathbb{Z})$ -invariance of the function f implies that $f(z) = f(z + 1)$ for all $z \in \mathbb{H}$. Clearly then the unrestricted function $f: \mathbb{H} \rightarrow \mathbb{C}$ cannot be

definable in any o-minimal structure. We consider instead its restriction $f|_{F_j}: F_j \rightarrow \mathbb{C}$, where F_j is the usual fundamental domain for the action of $\text{SL}_2(\mathbb{Z})$ on \mathbb{H} . On F_j , the function f is finite-to-one since $j|_{F_j}: F_j \rightarrow \mathbb{C}$ is injective. The composition of definable functions is definable (in any structure) and the field operations $+, \cdot$ are definable in $\mathbb{R}_{\text{an,exp}}$. It is thus enough to show that $j|_{F_j}: F_j \rightarrow \mathbb{C}$ is definable in $\mathbb{R}_{\text{an,exp}}$.

That $j|_{F_j}: F_j \rightarrow \mathbb{C}$ is definable in $\mathbb{R}_{\text{an,exp}}$ is an easy consequence of the q -expansion of j (see Proposition 1.15). A proof of this fact is given in [84, Example 4.14]. Consequently, we obtain that all modular functions are definable in the o-minimal structure $\mathbb{R}_{\text{an,exp}}$ when restricted to (a finite union of $\text{SL}_2(\mathbb{Z})$ -translates of) the fundamental domain F_j .

The modular function $j: \mathbb{H} \rightarrow \mathbb{C}$ is the uniformising map of the modular curve $Y(1) = \text{SL}_2(\mathbb{Z}) \backslash \mathbb{H}$. This map, restricted to a suitable semi-algebraic fundamental domain like F_j , is definable in $\mathbb{R}_{\text{an,exp}}$. Similarly, the uniformising map of any mixed Shimura variety is definable in $\mathbb{R}_{\text{an,exp}}$, when restricted to an appropriate domain. This result was proved for pure Shimura varieties in [41], and for mixed Shimura varieties by Gao [34]. The structure $\mathbb{R}_{\text{an,exp}}$ is thus the natural setting for all the proofs in this thesis, and so we always work in this structure.

References

1. J. Adamus and S. Randriambololona, *Tameness of holomorphic closure dimension in a semialgebraic set*, Math. Ann. **355** (2013), no. 3, 985–1005.
2. B. Allombert, Yu. Bilu, and A. Pizarro-Madariaga, *CM-points on straight lines*, Analytic number theory, Springer, Cham, 2015, pp. 1–18.
3. Y. André, *G-functions and geometry*, Aspects of Mathematics, E13, Friedr. Vieweg & Sohn, Braunschweig, 1989.
4. ———, *Finitude des couples d’invariants modulaires singuliers sur une courbe algébrique plane non modulaire*, J. Reine Angew. Math. **505** (1998), 203–208.
5. A. Astaneh-Asl and H. Daghigh, *Independence of Heegner points for nonmaximal orders*, Int. J. Number Theory **7** (2011), no. 3, 663–669.
6. J. Ax, *On Schanuel’s conjectures*, Ann. of Math. (2) **93** (1971), 252–268.
7. G. Baldi, *On a conjecture of Buium and Poonen*, Ann. Inst. Fourier (Grenoble) **70** (2020), no. 2, 457–477.
8. Yu. Bilu and L. Kühne, *Linear Equations in Singular Moduli*, Int. Math. Res. Not. IMRN (2020), no. 21, 7617–7643.
9. Yu. Bilu, F. Luca, and D. Masser, *Collinear CM-points*, Algebra & Number Theory **11** (2017), no. 5, 1047–1087.
10. Yu. Bilu, F. Luca, and A. Pizarro-Madariaga, *Rational products of singular moduli*, J. Number Theory **158** (2016), 397–410.
11. ———, *Trinomials, singular moduli and Riffaut’s conjecture*, preprint, arXiv:2003.06547 (2020).
12. Yu. Bilu, D. Masser, and U. Zannier, *An effective “theorem of André” for CM-points on a plane curve*, Math. Proc. Cambridge Philos. Soc. **154** (2013), no. 1, 145–152.

13. G. Binyamini, *Some effective estimates for André-Oort in $Y(1)^n$* , J. Reine Angew. Math. **767** (2020), 17–35, with an appendix by E. Kowalski.
14. E. Bombieri and W. Gubler, *Heights in Diophantine geometry*, New Mathematical Monographs, vol. 4, Cambridge University Press, Cambridge, 2006.
15. E. Bombieri, D. Masser, and U. Zannier, *Anomalous subvarieties—structure theorems and applications*, Int. Math. Res. Not. IMRN (2007), no. 19.
16. R. Borcherds, *Automorphic forms on $O_{s+2,2}(\mathbb{R})^+$ and generalized Kac–Moody algebras*, Proceedings of the International Congress of Mathematicians, Vol. 2 (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 744–752.
17. ———, *Automorphic forms on $O_{s+2,2}(\mathbb{R})$ and infinite products*, Invent. Math. **120** (1995), no. 1, 161–213.
18. ———, *Book review: Moonshine beyond the monster: The bridge connecting algebra, modular forms and physics*, Bull. Amer. Math. Soc. **45** (2008), no. 4, 675–680.
19. D. Brink, *On alternating and symmetric groups as Galois groups*, Israel J. Math. **142** (2004), 47–60.
20. A. Buïum and B. Poonen, *Independence of points on elliptic curves arising from special points on modular and Shimura curves. I. Global results*, Duke Math. J. **147** (2009), no. 1, 181–191.
21. F. Charles and B. Poonen, *Bertini irreducibility theorems over finite fields*, J. Amer. Math. Soc. **29** (2016), no. 1, 81–94.
22. H. Cohen, *A course in computational algebraic number theory*, Graduate Texts in Mathematics, vol. 138, Springer-Verlag, Berlin, 1993.
23. D. Cox, *Primes of the form $x^2 + ny^2$* , John Wiley & Sons, Inc., New York, 1989.
24. C. Daw and M. Orr, *Heights of pre-special points of Shimura varieties*, Math. Ann. **365** (2016), no. 3-4, 1305–1357.
25. C. Daw and J. Ren, *Applications of the hyperbolic Ax–Schanuel conjecture*, Compos. Math. **154** (2018), no. 9, 1843–1888.
26. L. van den Dries, *Remarks on Tarski’s problem concerning $(\mathbb{R}, +, \cdot, \exp)$* , Logic colloquium ’82 (Florence, 1982), Stud. Logic Found. Math., vol. 112, North-Holland, Amsterdam, 1984, pp. 97–121.

27. ———, *A generalization of the Tarski–Seidenberg theorem, and some nondefinability results*, Bull. Amer. Math. Soc. (N.S.) **15** (1986), no. 2, 189–193.
28. ———, *Tame topology and o-minimal structures*, London Mathematical Society Lecture Note Series, vol. 248, Cambridge University Press, Cambridge, 1998.
29. L. van den Dries, A. Macintyre, and D. Marker, *The elementary theory of restricted analytic fields with exponentiation*, Ann. of Math. (2) **140** (1994), no. 1, 183–205.
30. L. van den Dries and C. Miller, *On the real exponential field with restricted analytic functions*, Israel J. Math. **85** (1994), no. 1-3, 19–56.
31. B. Faye and A. Riffaut, *Fields generated by sums and products of singular moduli*, J. Number Theory **192** (2018), 37–46.
32. G. Fowler, *Multiplicative independence of modular functions*, preprint, arXiv:2005.13328 (2020).
33. ———, *Triples of singular moduli with rational product*, Int. J. Number Theory **16** (2020), no. 10, 2149–2166.
34. Z. Gao, *Towards the André–Oort conjecture for mixed Shimura varieties: the Ax–Lindemann theorem and lower bounds for Galois orbits of special points*, J. Reine Angew. Math. **732** (2017), 85–146.
35. G. van der Geer, *Siegel modular forms and their applications*, The 1-2-3 of modular forms, Universitext, Springer, Berlin, 2008, pp. 181–245.
36. B. Gross and D. Zagier, *On singular moduli*, J. Reine Angew. Math. **355** (1985), 191–220.
37. P. Habegger, *Singular moduli that are algebraic units*, Algebra Number Theory **9** (2015), no. 7, 1515–1524.
38. P. Habegger and J. Pila, *Some unlikely intersections beyond André–Oort*, Compos. Math. **148** (2012), no. 1, 1–27.
39. ———, *O-minimality and certain atypical intersections*, Ann. Sci. Éc. Norm. Supér. (4) **49** (2016), no. 4, 813–858.
40. M. Hindry and J. Silverman, *Diophantine geometry*, Graduate Texts in Mathematics, vol. 201, Springer-Verlag, New York, 2000.

41. B. Klingler, E. Ullmo, and A. Yafaev, *The hyperbolic Ax–Lindemann–Weierstrass conjecture*, Publ. Math. Inst. Hautes Études Sci. **123** (2016), 333–360.
42. J. Knight, A. Pillay, and C. Steinhorn, *Definable sets in ordered structures. II*, Trans. Amer. Math. Soc. **295** (1986), no. 2, 593–605.
43. L. Kühne, *An effective result of André–Oort type*, Ann. of Math. (2) **176** (2012), no. 1, 651–671.
44. L. Kühne, *Intersection of class fields*, Acta Arith. **198** (2021), no. 2, 109–127.
45. T. Loher and D. Masser, *Uniformly counting points of bounded height*, Acta Arith. **111** (2004), no. 3, 277–297.
46. F. Luca and A. Riffaut, *Linear independence of powers of singular moduli of degree three*, Bull. Aust. Math. Soc. **99** (2019), no. 1, 42–50.
47. D. Marker, *Model theory*, Graduate Texts in Mathematics, vol. 217, Springer-Verlag, New York, 2002, An introduction.
48. D. Masser, *Linear relations on algebraic groups*, New advances in transcendence theory (Durham, 1986), Cambridge Univ. Press, Cambridge, 1988, pp. 248–262.
49. J. Milne, *Canonical models of (mixed) Shimura varieties and automorphic vector bundles*, Automorphic forms, Shimura varieties, and L -functions, Vol. I (Ann Arbor, MI, 1988), Perspect. Math., vol. 10, Academic Press, Boston, MA, 1990, pp. 283–414.
50. ———, *Introduction to Shimura varieties*, Harmonic analysis, the trace formula, and Shimura varieties, Clay Math. Proc., vol. 4, Amer. Math. Soc., Providence, RI, 2005, pp. 265–378.
51. ———, *Abelian varieties (v2.00)*, 2008, available from www.jmilne.org/math/, pp. 166+vi.
52. ———, *Modular functions and modular forms (v1.31)*, 2017, available from www.jmilne.org/math/, p. 134.
53. N. Mok, J. Pila, and J. Tsimerman, *Ax–Schanuel for Shimura varieties*, Ann. of Math. (2) **189** (2019), no. 3, 945–978.
54. K. Ono, *The web of modularity: arithmetic of the coefficients of modular forms and q -series*, CBMS Regional Conference Series in Mathematics, vol. 102, Published

- for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2004.
55. F. Oort, *Canonical liftings and dense sets of CM-points*, Arithmetic geometry (Cortona, 1994), Sympos. Math., XXXVII, Cambridge Univ. Press, Cambridge, 1997, pp. 228–234.
 56. M. Orr, *Introduction to abelian varieties and the Ax-Lindemann-Weierstrass theorem*, O-minimality and diophantine geometry, London Math. Soc. Lecture Note Ser., vol. 421, Cambridge Univ. Press, Cambridge, 2015, pp. 100–128.
 57. PARI Group, Univ. Bordeaux, *PARI/GP version 2.11.4*, 2020, available from <http://pari.math.u-bordeaux.fr/>.
 58. H. Petersson, *Über die Entwicklungskoeffizienten der automorphen Formen*, Acta Math. **58** (1932), no. 1, 169–215.
 59. J. Pila, *On the algebraic points of a definable set*, Selecta Math. (N.S.) **15** (2009), no. 1, 151–170.
 60. ———, *O-minimality and the André–Oort conjecture for \mathbb{C}^n* , Ann. of Math. (2) **173** (2011), no. 3, 1779–1840.
 61. ———, *Special point problems with elliptic modular surfaces*, Mathematika **60** (2014), no. 1, 1–31.
 62. ———, *Functional transcendence via o-minimality*, O-minimality and diophantine geometry, London Math. Soc. Lecture Note Ser., vol. 421, Cambridge Univ. Press, Cambridge, 2015, pp. 66–99.
 63. J. Pila and J. Tsimerman, *Independence of CM points in elliptic curves*, to appear in J. Eur. Math. Soc. (JEMS).
 64. ———, *Ax–Schanuel for the j -function*, Duke Math. J. **165** (2016), no. 13, 2587–2605.
 65. ———, *Multiplicative relations among singular moduli*, Ann. Sc. Norm. Super. Pisa Cl. Sci. (5) **17** (2017), no. 4, 1357–1382.
 66. J. Pila and A. Wilkie, *The rational points of a definable set*, Duke Math. J. **133** (2006), no. 3, 591–616.

67. J. Pila and U. Zannier, *Rational points in periodic analytic sets and the Manin–Mumford conjecture*, Atti Accad. Naz. Lincei Rend. Lincei Mat. Appl. **19** (2008), no. 2, 149–162.
68. A. Pillay and C. Steinhorn, *Definable sets in ordered structures. I*, Trans. Amer. Math. Soc. **295** (1986), no. 2, 565–592.
69. ———, *Definable sets in ordered structures. III*, Trans. Amer. Math. Soc. **309** (1988), no. 2, 469–476.
70. R. Pink, *A combination of the conjectures of Mordell–Lang and André–Oort*, Geometric methods in algebra and number theory, Progr. Math., vol. 235, Birkhäuser Boston, Boston, MA, 2005, pp. 251–282.
71. H. Rademacher, *The Fourier Coefficients of the Modular Invariant $J(\tau)$* , Amer. J. Math. **60** (1938), no. 2, 501–512.
72. ———, *The Fourier Series and the Functional Equation of the Absolute Modular Invariant $J(\tau)$* , Amer. J. Math. **61** (1939), no. 1, 237–248.
73. M. Raynaud, *Sous-variétés d’une variété abélienne et points de torsion*, Arithmetic and geometry, Vol. I, Progr. Math., vol. 35, Birkhäuser Boston, Boston, MA, 1983, pp. 327–352.
74. A. Riffaut, *Equations with powers of singular moduli*, Int. J. Number Theory **15** (2019), no. 3, 445–468.
75. M. Rosen and J. Silverman, *On the independence of Heegner points associated to distinct quadratic imaginary fields*, J. Number Theory **127** (2007), no. 1, 10–36.
76. The Sage Developers, *Sagemath, the Sage Mathematics Software System (Version 9.0)*, 2020, available from <https://www.sagemath.org>.
77. H. Şahinoğlu, *On the independence of Heegner points on CM elliptic curves associated to distinct quadratic imaginary fields*, Proc. Amer. Math. Soc. **141** (2013), no. 3, 813–826.
78. T. Scanlon, *Counting special points: logic, Diophantine geometry, and transcendence theory*, Bull. Amer. Math. Soc. (N.S.) **49** (2012), no. 1, 51–71.
79. J. Silverman, *The arithmetic of elliptic curves*, second ed., Graduate Texts in Mathematics, vol. 106, Springer, Dordrecht, 2009.

80. H. Spence, *A note on the degree of field extensions involving classical and nonholomorphic singular moduli*, 2017, preprint, arXiv:1702.01950.
81. A. Wilkie, *Model completeness results for expansions of the ordered field of real numbers by restricted Pfaffian functions and the exponential function*, J. Amer. Math. Soc. **9** (1996), no. 4, 1051–1094.
82. D. Zagier, *Traces of singular moduli*, Motives, polylogarithms and Hodge theory, Part I (Irvine, CA, 1998), Int. Press Lect. Ser., vol. 3, Int. Press, Somerville, MA, 2002, pp. 211–244.
83. ———, *Elliptic modular forms and their applications*, The 1-2-3 of modular forms, Universitext, Springer, Berlin, 2008, pp. 1–103.
84. U. Zannier, *Some problems of unlikely intersections in arithmetic and geometry*, Annals of Mathematics Studies, vol. 181, Princeton University Press, Princeton, NJ, 2012, With appendices by D. Masser.
85. B. Zilber, *Exponential sums equations and the Schanuel conjecture*, J. London Math. Soc. (2) **65** (2002), no. 1, 27–44.