Gromov–Witten theory via roots and logarithms

LUCA BATTISTELLA

NAVID NABIJOU

DHRUV RANGANATHAN

msp

Gromov–Witten theory via roots and logarithms

Luca Battistella

Navid Nabijou

Dhruv Ranganathan

Orbifold and logarithmic structures provide independent routes to the virtual enumeration of curves with tangency orders for a simple normal crossings pair $(X|D)$. The theories do not coincide and their relationship has remained mysterious. We prove that the genus-zero orbifold theories of multiroot stacks of strata blowups of $(X|D)$ converge to the corresponding logarithmic theory of $(X|D)$. With fixed numerical data, there is an explicit combinatorial criterion that guarantees when a blowup is sufficiently refined for the theories to coincide. The result identifies birational invariance as the crucial property distinguishing the logarithmic and orbifold theories. There are two key ideas in the proof. The first is the construction of a naive Gromov–Witten theory, which serves as an intermediary between roots and logarithms. The second is a smoothing theorem for tropical stable maps; the geometric theorem then follows via virtual intersection theory relative to the universal target. The results import a new set of computational tools into logarithmic Gromov–Witten theory. As an application, we show that the genus-zero logarithmic Gromov–Witten theory of a pair is determined by the absolute Gromov–Witten theories of its strata.

Introduction

0.1 Overview

Gromov–Witten theory for a pair $(X|D)$ of a projective manifold and a simple normal crossings divisor models the enumerative geometry of curves in X with prescribed tangency along D. Two approaches to this problem of significant recent interest proceed via logarithmic structures and via root constructions

along the divisor; see [Abramovich and Chen 2014; Chen 2014; Gross and Siebert 2013] and [Abramovich and Vistoli 2002; Cadman 2007; Cadman and Chen 2008; Abramovich et al. 2017; Tseng and You 2023a]. In both, the idealised object is a map of pairs from a smooth curve

$$(C \mid p_1 + \cdots + p_n) \to (X \mid D = D_1 + \cdots + D_k),$$

where the tangency of p_i along D_j is a fixed nonnegative integer c_{ij}. The approaches above yield two paths to a proper moduli space with a virtual fundamental class. In both, the degenerate objects in the moduli space parametrise stable maps with additional structure.

In the orbifold approach, the map is required to lift to a representable map of stacks, obtained from root constructions along the divisors D_j and the markings p_i. The tangency is recorded by the morphisms on isotropy groups, but the orbifold structure along D is not canonical and provides an auxiliary choice — the rooting parameter(s).

In the logarithmic approach, one demands that the map can be enhanced to a logarithmic one, with tangency orders encoded in the logarithmic structure sheaves.

The two theories are very different in nature, provide complementary computational tools, and offer different insights into enumerative geometry. We compare them. There are two orthogonal sources of complexity:

  (i) *source complexity*, the genus of the source curve, where the problem breaks naturally into the cases of genus zero and positive genus, and

  (ii) *target complexity*, the logarithmic rank of the target, where the problem breaks naturally into the cases of logarithmic rank 1, where the divisor is smooth, and arbitrary rank.

When the logarithmic rank is 1 there is a complete picture. Once the rooting parameter r is sufficiently high, the theories coincide in genus zero by work of Abramovich, Cadman, Marcus, and Wise [Abramovich et al. 2014; 2017], following calculations of Cadman and Chen [2008]. In higher genus, Maulik showed that the orbifold theory depends on the rooting parameter r. Work of Janda, Pandharipande, Pixton, and Zvonkine [Janda et al. 2017; 2020] asserts that the contribution to the Gromov–Witten theory that is local to the divisor D is polynomial in r, and the r–free term is the logarithmic theory. Finally, Tseng and You [2020] used the compatibility of these statements with degeneration to deduce that the r–free term of the orbifold theory of $(X \mid D)$ is the logarithmic theory.

We study the complementary situation of the genus-zero theory in arbitrary logarithmic rank. Examples from [Nabijou and Ranganathan 2022, Section 1] show that even in genus zero the theories do not coincide. We offer an explanation: the logarithmic theory is invariant under logarithmic modifications of the boundary divisor [Abramovich and Wise 2018], whereas the orbifold theory is not. We show that the invariant part of the latter is the logarithmic theory.

0.2 Main results

Let $(X|D)$ be a simple normal crossings pair with components $D_1, \dots, D_k$. We fix numerical data Λ for the enumerative problem, which is a tuple $(g, \beta, n, [c_{ij}])$, where g is the genus of the source curve, β is the curve class, n is the number of marked points, and c_{ij} is the tangency order at the i^{th} marking with respect to D_j. The collection of iterated blowups

$$(X^\dagger|D^\dagger) \to (X|D)$$

along the strata of D forms an inverse system. In Section 4 we introduce a property of such blowups called Λ–*sensitivity*, and observe that blowups with this property always exist. Moreover, this condition is stable under taking further blowups.

For the first result, we fix the genus to be zero. Given any iterated blowup $(X^\dagger|D^\dagger) \to (X|D)$ along logarithmic strata, there is a canonical lifting of the numerical data Λ to $(X^\dagger|D^\dagger)$, obtained by formally computing the strict transform of a stable map; see Section 3.5.

Theorem X   *Let $(X^\dagger|D^\dagger) \to (X|D)$ be a Λ–sensitive blowup. The genus-zero logarithmic Gromov–Witten invariants of $(X|D)$ with numerical data Λ agree with the orbifold Gromov–Witten invariants of $(X^\dagger|D^\dagger)$ with the same numerical data.*

The result is proved on the level of virtual cycles; see Section 0.3 for details. In particular, the result implies that the multiroot orbifold theories of strata blowups eventually stabilise, which appears to be a nonobvious consequence.

Combining this with results on the Gromov–Witten theory of root stacks [Tseng and You 2016; Chen et al. 2020], gerbes [Andreini et al. 2015], projective bundles, and blowups [Fan 2021], we obtain the following *reconstruction principle for logarithmic Gromov–Witten invariants*:

Corollary Y   *The genus-zero logarithmic Gromov–Witten theory of $(X|D)$ can be uniquely reconstructed from the absolute Gromov–Witten theories of all the logarithmic strata of $(X|D)$.*

In genus zero, this result generalises a well-known theorem from relative Gromov–Witten theory due to Maulik and Pandharipande [2006, Theorem 2]. The authors are not aware of a purely logarithmic argument that yields this result. Although it may be possible to use the degeneration formula to deduce it [Ranganathan 2022], the result illustrates the type of consequences that can be extracted from our theorem.

0.3 Naive theory

The route to Theorem X passes through a new version of logarithmic Gromov–Witten theory, referred to as the naive theory. This appeared in an early form in [Nabijou 2018, Chapter 3; Nabijou and Ranganathan 2022]. Here we construct the theory in full generality: in arbitrary genus and without positivity assumptions on the divisor.
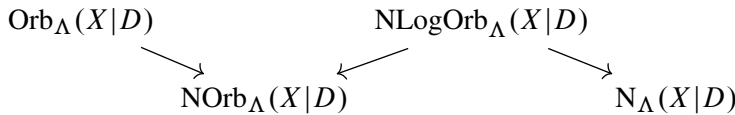
Let $\mathrm{Log}_\Lambda(X|D)$ denote the moduli space of logarithmic stable maps to a pair $(X|D)$, and $\mathrm{M}_\Lambda(X)$ the space of absolute stable maps; the symbol Λ records the numerical data. Given $(X|D)$, each divisor component D_i determines a smooth pair $(X|D_i)$, and a logarithmic stable map to $(X|D)$ induces one to each $(X|D_i)$. The geometry of the interior of the space $\mathrm{Log}_\Lambda(X|D)$ suggests an alternative compactification that we call the moduli space of *naive stable maps*:

$$\mathrm{N}_\Lambda(X|D) := \mathrm{Log}_\Lambda(X|D_1) \times_{\mathrm{M}_\Lambda(X)} \cdots \times_{\mathrm{M}_\Lambda(X)} \mathrm{Log}_\Lambda(X|D_k).$$

The fibre product is taken in the category of algebraic stacks, disregarding the logarithmic structures. The following theorem holds for arbitrary genus:

Theorem Z *The moduli space $\mathrm{N}_\Lambda(X|D)$ of naive stable maps is equipped with natural evaluation morphisms to the strata of $(X|D)$, and a virtual fundamental class whose virtual dimension coincides with the virtual dimension of the corresponding space of logarithmic stable maps.*

In genus zero the resulting system of naive Gromov–Witten invariants agrees with the orbifold theory, as follows: there is a diagram of moduli spaces equipped with virtual fundamental classes

$$
\mathrm{Orb}_\Lambda(X|D) \qquad\qquad \mathrm{NLogOrb}_\Lambda(X|D)
$$
$$
\searrow \qquad\qquad \swarrow \qquad\qquad \searrow
$$
$$
\qquad \mathrm{NOrb}_\Lambda(X|D) \qquad\qquad\qquad \mathrm{N}_\Lambda(X|D)
$$

with each arrow identifying virtual fundamental classes via proper pushforward.

The spaces appearing in the diagram above will be defined in Section 2. In simple terms, there is no direct morphism from $\mathrm{Orb}_\Lambda(X|D)$ to $\mathrm{N}_\Lambda(X|D)$, but a map exists after introducing appropriate root constructions; see for instance [Abramovich et al. 2014; 2017]. The intermediate spaces give an explicit moduli-theoretic factorisation of these root constructions. Product formulas in orbifold Gromov–Witten theory play an important role in Theorem Z; see [Andreini et al. 2016; Battistella et al. 2023]. The fact that logarithmic Gromov–Witten theory does not satisfy product-type formulas can be seen as an instance of the phenomena that we encounter here; see [Ranganathan 2019; Nabijou and Ranganathan 2022].

Theorem X can now be phrased as stating that the natural map

$$\mathrm{Log}_\Lambda(X^\dagger|D^\dagger) \to \mathrm{N}_\Lambda(X^\dagger|D^\dagger)$$

is virtually birational for any Λ–sensitive blowup $(X^\dagger|D^\dagger) \to (X|D)$. In Section 5 this is proved by reducing to a tropical smoothing problem, which we solve by exploiting the combinatorics of Λ–sensitive blowups. The solution identifies the main component of the naive moduli space.

0.4 Context

The enumerative geometry of orbifolds obtained by roots along a simple normal crossings divisor was considered in Cadman's foundational paper [2007, Section 2], and its systematic study was initiated by Tseng and You [2023a]. They observed that the theory satisfies a number of remarkable properties, owing

to the structural similarities between ordinary and orbifold Gromov–Witten theory. They ask about the relationship with logarithmic Gromov–Witten theory, and our main theorem answers this question in genus zero.

In recent years, a number of results have appeared in the logarithmic geometry literature asserting that natural geometric claims are true *after sufficiently blowing up*; see for instance [Holmes et al. 2019; Nabijou and Ranganathan 2022; Ranganathan 2019; 2022]. However our results *do not* fit into this general principle. In our work the moduli space of stable maps is never directly blown up, but rather only the target is, leading to a much stronger conclusion. In the cited works, once the moduli spaces of curves or of maps is blown up, the canon of intersection theory methods, including localisation, do not apply. These methods remain accessible for us, leading to conclusions such as Corollary Y. In short, we only blow up the target, while earlier results blow up the moduli space.

The comparison between logarithmic and orbifold Gromov–Witten theory should also be contrasted with those relating different models *within* logarithmic Gromov–Witten theory, including [Abramovich et al. 2014; Abramovich and Wise 2018; Ranganathan 2022]. While the latter statements are essentially uniform, the comparison we obtain is more subtle and only true after appropriate invariance properties are forced.

The main result has consequences for the types of computational techniques which are available. Computations in logarithmic enumerative geometry have been achieved via tropical correspondences [Mandel and Ruddat 2020; Nishinou and Siebert 2006], degeneration formulas [Abramovich et al. 2020a; 2020b; Ranganathan 2022; Wu 2024], and scattering diagrams [Argüz and Gross 2022; Gross et al. 2010]. These methods rely on tropical geometry in an essential fashion.

The orbifold theory is equipped with a different set of computational techniques, including torus localisation, Givental formalism, and universal relations such as WDVV and Virasoro constraints; see [Tseng and You 2023a]. None of these are available in logarithmic geometry. The difference in computational techniques illustrates the different nature of the two sides, and is a notable feature of Theorem X. As a coarse consequence, the main result, together with virtual localisation, determines the genus-zero logarithmic Gromov–Witten theory of a pair $(X \mid D)$ for X toric and D a subset of the toric boundary in terms of known calculations [Liu 2013, Section 9]. In a different way, once the problem has been moved out of logarithmic geometry, quasimap wall-crossing and enumerative mirror symmetry techniques become available [Battistella and Nabijou 2021; Cheong et al. 2015]. It seems plausible that mirror transformations may help to further clarify the relationship between the logarithmic and orbifold theories; see for instance [Shafi 2024; You 2022].

The process of taking the r–free term of the higher-genus orbifold theory can also be seen as a way of forcing an invariance property in the orbifold theory [Janda et al. 2017; 2020; Tseng and You 2023a]. Together with the main result here, it suggests that invariance under roots and blowups are precisely what distinguishes the orbifold theory from the logarithmic theory.

0.5 Future work

Two directions for theoretical generalisation are available. The first is the higher-genus comparison problem. As discussed earlier, if the logarithmic rank is 1, the invariants are known to become polynomial in the rooting parameter, with the r–free term being the logarithmic theory. The polynomiality result is known to hold for simple normal crossings divisors [Tseng and You 2023a]: the invariants are polynomial in variables $r_1, \ldots, r_k$ corresponding to the rooting parameters for divisors $D_1, \ldots, D_k$. It is natural to conjecture:

Conjecture W *For each choice of numerical data Λ in arbitrary genus, there is a sufficiently refined blowup $(X^\dagger | D^\dagger) \to (X | D)$ such that the constant-term orbifold Gromov–Witten invariants of $(X^\dagger | D^\dagger)$ agree with the logarithmic Gromov–Witten invariants of $(X | D)$.*

The result appears plausible, but it seems likely that new input coming from higher double ramification cycles will be necessary, going beyond [Holmes et al. 2019; Holmes and Schwarz 2022; Molcho and Ranganathan 2024].

The second direction is to incorporate negative tangency. In the logarithmic theory this is the subject of [Abramovich et al. 2020b], while in the orbifold theory, it can be extracted by studying invariants of fixed coage [Fan et al. 2020; 2021]. The latter papers investigate the comparison between logarithmic and orbifold geometries in logarithmic rank 1 with negative contact orders and establish a comparison, albeit using a different model, for the logarithmic negative contact theory via localisation and graph summation on the normal cone of D. Negative contact orders are the key input in the construction of intrinsic mirrors, via degree-zero relative quantum cohomology for pairs [Gross and Siebert 2019]. The associativity of this ring, which is the main result of [loc. cit.], is shown via a remarkable calculation. Our results offer the possibility of deducing the associativity from the simpler associativity of orbifold quantum cohomology, which also offers a route to a logarithmic quantum cohomology for pairs [Tseng and You 2023a].

0.6 Strategy

A simplified overview is as follows: Given a pair $(X | D)$, each smooth pair $(X | D_i)$ determines a cycle in the space $\mathrm{M}(X)$ of absolute maps. The intersection of these cycles is the image of the naive moduli space of maps in $\mathrm{M}(X)$. We equip it with a virtual class that is compatible with its presentation as an intersection. The orbifold geometry attached to $(X | D)$ is the fibre product of those attached to $(X | D_i)$, over X. The Gromov–Witten cycles for the orbifold theory and, by definition, the naive theory both satisfy a product rule in $\mathrm{M}(X)$. Once they are equated, we dispense with the orbifold theory. In the main text, this is refined to work on the naive spaces themselves rather than their images in $\mathrm{M}(X)$.

A map to X lifts to a logarithmic map to $(X | D_i)$ if and only if it satisfies an intersection-theoretic criterion due to Gathmann [2002]. However, the existence of lifts to all $(X | D_i)$ does not guarantee a lift to $(X | D)$. The obstruction is tropical: a logarithmic map determines a combinatorial type of a tropical map, and

when X is replaced with the universal simple normal crossings pair $[\mathbb{A}^k/\mathbb{G}_m^k]$, the tropical data essentially determines the logarithmic data. However, these tropical data are not compatible with the natural product decomposition of $[\mathbb{A}^k/\mathbb{G}_m^k]$. Concretely, assembling two rank-1 tropical types to a single type of rank 2 is a subtle problem.

However, explicit polyhedral geometry arguments show that the obstructions vanish after a sufficiently fine subdivision of the tropical target. This yields a universal version of the main theorem. Combined with techniques from virtual intersection theory, we deduce the main result.

Conventions and notation We work over an algebraically closed field of characteristic zero. We assume familiarity with toroidal, logarithmic, and tropical geometry, and the role they play in logarithmic Gromov–Witten theory. An overview is provided in Section 1.

We will work with various flavours of stable maps throughout the paper. To reduce clutter, we adopt the simplified notation

$$\mathrm{M}_\Lambda(X), \quad \mathrm{Log}_\Lambda(X|D), \quad \text{and} \quad \mathrm{Orb}_\Lambda(X|D)$$

for the absolute, logarithmic, and orbifold moduli spaces. A number of auxiliary moduli spaces will be introduced at various points and denoted similarly, as will the corresponding spaces of prestable maps to the universal targets.

1 Tropical essentials

We provide a review of the crucial aspects of logarithmic and tropical geometry that we require. The treatment strives for brevity, but we provide detailed references.

1.1 Cone complexes and Artin fans

Let $(X|D)$ be a simple normal crossings pair with components $D_1, \ldots, D_k$. A stratum of $(X|D)$ is a connected component of an intersection of some of the D_j (this includes the empty intersection X). All strata are smooth. We momentarily make the simplifying assumption that every intersection of D_j is connected.

Let $\mathcal{A}$ denote the stack $[\mathbb{A}^1/\mathbb{G}_m]$. A morphism from X to $\mathcal{A}$ is precisely the data of a line bundle–section pair, referred to here as a *generalised Cartier divisor*. Each D_j furnishes a morphism $X \to \mathcal{A}$. Combining these, the divisor D induces a smooth morphism

$$X \to \mathcal{A}^k.$$

Let $\mathcal{A}_X$ be the minimal open substack of $\mathcal{A}^k$ through which the map from X factors. It has one point for each stratum. Let $\mathcal{D}_X$ be the complement of the dense open point. Consider the sequence of maps

$$X \to \mathcal{A}_X \hookrightarrow \mathcal{A}^k.$$

Each space is a logarithmic algebraic stack, with logarithmic structure induced by the natural simple normal crossings divisor. Each arrow is strict. The stack $\mathcal{A}_X$ is called the *Artin fan* of X; see [Abramovich and Wise 2018, Section 2] for the original construction and [Abramovich et al. 2016, Section 5] for further exposition.

The datum $(X|D)$ determines a cone complex with integral structure, also called the *tropicalisation of $(X|D)$*. It is denoted by Σ_X, with D suppressed from notation, and is constructed as follows. It has k distinct rays $\rho_1, \ldots, \rho_k$. For each nonempty subset $I \subset \{1, \ldots, k\}$ we write D_I for the intersection of the corresponding D_i. If nonempty, the intersection is closed of codimension equal to the size of I. For each nonempty D_I, the complex Σ_X contains a copy of the cone σ_I, generated by the rays ρ_i for i in I. There is a natural identification of σ_I as a face of σ_J whenever J contains I. Define

$$\mathrm{Trop}(X|D) := \Sigma_X := \varinjlim_I \sigma_I$$

as the colimit under these identifications. It comes with a canonical set of integral points, denoted by $\Sigma_X(\mathbb{N})$. The construction first appeared in a more general form in [Kempf et al. 1973, Chapter II]. An exposition may be found in [Abramovich et al. 2016, Section 2] or [Cavalieri et al. 2020, Section 2].

Remark 1.1 The partially ordered set of faces of Σ_X can be equipped with its Alexandrov topology, and this coincides with the underlying topological space of $\mathcal{A}_X$.

Remark 1.2 If the intersections D_I are not connected, both the Artin fan and cone complex constructions generalise in a natural way. Following the construction above, the map from X to $\mathcal{A}^k$ has disconnected fibres. By replacing the $\mathcal{A}_X$ above with a nonseparated cover, we can guarantee connectedness of the fibres; see [Abramovich and Wise 2018]. The cone complex construction can also be generalised by taking one cone for each connected component in the intersection.

1.2 Tropical curves and tropical maps

The tropical curves we consider will have genus zero. The *dual graph* Γ of a stable genus-zero curve C has vertices in bijection with the irreducible components of C, and the set of edges connecting two vertices is in bijection with the set of nodes connecting the corresponding components. If C is marked by smooth points $p_1, \ldots, p_n$, this determines a map to the vertex set

$$m \colon \{1, \ldots, n\} \to V(\Gamma).$$

A *tropical curve* is an enhancement of Γ obtained by assigning a positive real length to each edge

$$\ell \colon E(\Gamma) \to \mathbb{R}_{>0}.$$

The triple (Γ, m, ℓ) determines a metric graph enhancing Γ in two steps. First, identify an edge e of Γ with an interval of length $\ell(e)$. Second, for each marking i, attach a copy of $\mathbb{R}_{\geq 0}$ to this metric graph by identifying $0 \in \mathbb{R}_{\geq 0}$ with the vertex $m(i)$.

The notion of a *family of tropical curves* is defined as follows. Fix a sharp toric monoid Q and let τ be the dual cone. Given a marked dual graph, a family of tropical curves over τ is determined by a map

(1) $$\ell \colon E(\Gamma) \to Q \setminus \{0\}.$$

Each point of τ is a homomorphism from Q to $\mathbb{R}_{\geq 0}$, and by applying this homomorphism to the edge assignment above, each edge e acquires a length. On the interior of τ, this is a tropical curve in the previous sense. Fixing the data (1), there is a cone complex $\sqsubset$ together with a map

$$\sqsubset \to \tau$$

whose fibre over a point in the interior is a tropical curve as previously discussed. The boundary parametrises tropical curves where a subset of edges have been contracted. The total space $\sqsubset$ is a cone complex of relative dimension 1 over τ and its cones correspond to the faces of the underlying dual graph Γ. We typically work with this total space $\sqsubset$, with the labelling of the graph Γ implicit; see [Cavalieri et al. 2020, Section 3] for details and generalisations.

Let $\sqsubset \to \tau$ be a family of tropical curves and let Σ be a cone complex. A *tropical map from $\sqsubset$ to Σ* is a map of cone complexes

$$\sqsubset \to \Sigma.$$

The fibres over τ can be viewed as tropical curves equipped with a piecewise linear map to Σ. When $\Sigma = \mathbb{R}_{\geq 0}$ we will refer to this as a *piecewise linear function* on $\sqsubset$.

Remark 1.3 (Slice picture) In order to speak of a single tropical curve, as opposed to a family, formally one takes the monoid Q above to be $\mathbb{N}$. In reality, this produces a family of tropical curves $\sqsubset \to \tau = \mathbb{R}_{\geq 0}$. In combinatorial arguments it is traditional to slice this map of cone complexes and take the fibre over the point 1 in τ. The family can be uniquely reconstructed from this slice by taking the cone over the tropical curve.

1.3 Tropicalisation of curves

Let $(\operatorname{Spec}\Bbbk, Q)$ denote a closed point equipped with the logarithmic structure associated to a toric monoid Q. Let τ be the dual cone of Q. Given a logarithmic curve

$$\mathcal{C} \to (\operatorname{Spec}\Bbbk, Q),$$

there is an associated tropical curve $\sqsubset \to \tau$. Let Γ be the dual graph of $\mathcal{C}$. The étale local structure of logarithmic curves was determined by Kato [2000], and associates an element of Q to each edge of Γ; see [Cavalieri et al. 2020, Section 7.2] for a tropical perspective. The *tropicalisation of* $\mathcal{C} \to (\operatorname{Spec}\Bbbk, Q)$ is the associated family of tropical curves $\sqsubset \to \tau$. We denote the tropicalisation by $\operatorname{Trop}(\mathcal{C})$ or $\sqsubset$.

Remark 1.4 The construction globalises to relate the moduli space of logarithmic curves and the moduli space of tropical curves; see [Cavalieri et al. 2020]. We will not require this generality.

1.4 Line bundles and blowups

We will use two combinatorial dictionaries. First, if $\mathcal{C}$ is a logarithmic curve over a geometric point $(\operatorname{Spec}\Bbbk, Q)$, then a piecewise linear function on the tropicalisation $\sqsubset$ determines a generalised Cartier divisor on $\mathcal{C}$. The construction generalises the toric dictionary between toric Cartier divisors and piecewise linear functions on the fan.

If $\alpha \colon \sqsubset \to \mathbb{R}_{\geq 0}$ is a piecewise linear function, it can be identified with an element of $H^0(\mathcal{C}, \overline{M}_{\mathcal{C}}^{\mathrm{gp}})$. The line bundle $\mathcal{O}_{\mathcal{C}}(\alpha)$ is the image of $-\alpha$ in $H^1(\mathcal{C}, \mathcal{O}_{\mathcal{C}}^{\star})$ under the coboundary map induced by

$$0 \to \mathcal{O}_{\mathcal{C}}^{\star} \to M_{\mathcal{C}}^{\mathrm{gp}} \to \overline{M}_{\mathcal{C}}^{\mathrm{gp}} \to 0.$$

The logarithmic structure map $M_{\mathcal{C}} \to \mathcal{O}_{\mathcal{C}}$ furnishes a section of this bundle; see eg [Ranganathan et al. 2019, Section 2]. The resulting generalised Cartier divisor is denoted by $(\mathcal{O}_{\mathcal{C}}(\alpha), s_\alpha)$. A piecewise linear function on Σ_X similarly determines a divisor on X supported along D.

The second construction is related to birational modifications. A morphism of cone complexes $\Sigma' \to \Sigma_X$ (see Section 1.1 and [Kempf et al. 1973, Chapter II]) is a *subdivision* if it is a bijection on the underlying sets. The most basic type of subdivision is the stellar subdivision of a simplicial cone. Given a simplicial cone σ with primitive generators $u_1, \ldots, u_k$, the *stellar subdivision at* σ is the unique cone complex that refines Σ_X and has precisely one more ray, namely the one spanned by $\sum_{i=1}^{k} u_i$. A subdivision $\Sigma' \to \Sigma_X$ induces a birational modification $X' \to X$ which modifies X along the boundary D. The stellar subdivision at a cone σ is the blowup of X at the smooth stratum dual to σ. Further combinatorial details can be found in the text on toric geometry [Cox et al. 2011, Section 3.3]. The generalisation to the context of simple normal crossings pairs is detailed in [Kempf et al. 1973].

A more general version will also be used. If $(\operatorname{Spec}\Bbbk, Q)$ is a geometric point, then a subdivision of τ gives rise to a logarithmic modification of $\operatorname{Spec}\Bbbk$. We can identify this logarithmic point with a point in the closed orbit of the toric variety with coordinate ring $\Bbbk[Q]$, equipped with the pullback logarithmic

structure. The subdivision of τ birationally modifies this toric variety. The corresponding modification of Spec $\Bbbk$ is obtained by pullback along the inclusion; see [Kato 1994, Section 9] for a detailed study.

1.5 Moduli of curves with fixed tropicalisation

Let C be genus-zero prestable curve over Spec $\Bbbk$. Choose an identification of the dual graph of C with a labelled graph Γ, and let σ_Γ denote the smooth cone $\operatorname{Hom}(\mathbb{N}^{E(\Gamma)}, \mathbb{R}_{\geq 0})$. Points of σ_Γ correspond to choices of edge lengths, ie metric enhancements of Γ. This gives rise to a universal family of tropical curves $\sqsubset_\Gamma \to \sigma_\Gamma$. Choose a rational ray

$$\rho \subset \sigma_\Gamma$$

whose interior is contained in the interior of σ_Γ, and let v be the primitive integral generator of ρ. We describe the set of all logarithmic curves $\mathcal{C}$ over the standard logarithmic point (Spec $\Bbbk$, $\mathbb{N}$) such that the underlying curve is C, the dual graph is identified with Γ, and the tropicalisation is given by the restriction of $\sqsubset_\rho \to \rho$ of the universal tropical family to ρ.

The moduli space of logarithmic curves with such fixed tropicalisation is described by [Cavalieri et al. 2020, Theorem 4] and the formalism surrounding it. As we have chosen an identification of the dual graph with a fixed copy of Γ and only require the case where the underlying base is a geometric point, the stack-theoretic subtleties in [loc. cit.] do not arise, and the formalism simplifies.

This moduli space is an algebraic torus, which can be identified with the parameter space of logarithmic morphisms

$$(\operatorname{Spec} \Bbbk, \mathbb{N}) \to (\operatorname{Spec} \Bbbk, \mathbb{N}^{E(\Gamma)})$$

whose tropicalisation is $\rho \hookrightarrow \sigma_\Gamma$. Indeed, a logarithmic lift is given by specifying the image of every generator

$$(1, e_i) \in \Bbbk^\times \oplus \mathbb{N}^{E(\Gamma)}$$

in $\Bbbk^\times \oplus \mathbb{N}$. The discrete part is determined by the dual of $\rho \hookrightarrow \sigma_\Gamma$, but the continuous part is free.

A geometric view is provided below. A closely related concept in tropical geometry is the initial degeneration; see [Maclagan and Sturmfels 2015, Chapters 2–3].

Construction 1.5 Given $\rho \subset \sigma_\Gamma$, choose a subdivision of σ_Γ that includes ρ as a face. Let $\widetilde{U}$ be the associated logarithmic modification of (Spec $\Bbbk$, $\mathbb{N}^{E(\Gamma)}$). Define

$$\mathrm{M}(\rho) \subseteq \widetilde{U}$$

to be the locally closed stratum of codimension 1 which is dual to ρ.

Each point of $\mathrm{M}(\rho)$ gives rise to a logarithmic curve $\mathcal{C}$ over a logarithmic point (Spec $\Bbbk$, $\mathbb{N}$) enhancing C whose tropicalisation is $\sqsubset_\rho \to \rho$, and which is compatible with the fixed identification of the dual graph with Γ. In fact, $\mathrm{M}(\rho)$ is precisely the moduli space of such data.

More generally, if we choose any rational cone $\tau \subset \sigma_\Gamma$ with Q the dual monoid of τ, we can ask for logarithmic enhancements of C over $(\mathrm{Spec}\,\Bbbk, Q)$, again with a fixed identification of the dual graph with Γ. The moduli space of such enhancements is constructed similarly.

1.6 Logarithmic maps to Artin fans

We will use the following method for constructing logarithmic maps to Artin fans. Given an Artin fan $(\mathcal{A}_X | \mathcal{D}_X)$ and a logarithmic curve $\mathcal{C} \to (\mathrm{Spec}\,\Bbbk, Q)$, the data of a logarithmic map

$$\mathcal{C} \to (\mathcal{A}_X | \mathcal{D}_X)$$

is equivalent to the data of a map of cone complexes

(2) $$\mathrm{Trop}\,\mathcal{C} = \square \to \Sigma = \mathrm{Trop}(\mathcal{A}_X | \mathcal{D}_X).$$

See [Abramovich et al. 2020a, Proposition 2.10]. Here

$$\mathrm{Trop}\,\mathcal{C} = \square \to \tau = \mathrm{Hom}(Q, \mathbb{R}_{\geq 0})$$

is the tropicalisation of $\mathcal{C} \to (\mathrm{Spec}\,\Bbbk, Q)$.

Typically we start with a scheme-theoretic curve C, and proceed by first constructing an appropriate tropical curve $\square \to \tau$ and accompanying map $\square \to \Sigma$. Once this has been done, it remains to construct any logarithmic enhancement of C which tropicalises to the given tropical curve, which is achieved via the process outlined in Section 1.5.

2 Naive–orbifold comparisons

A naive stable map to a simple normal crossings pair $(X|D)$ is a stable map to X together with a logarithmic lift to each pair $(X|D_i)$. The theory was studied under positivity assumptions in genus zero in [Nabijou 2018; Nabijou and Ranganathan 2022]. The spaces were introduced as part of a framework to reduce questions in logarithmic Gromov–Witten theory to questions in the simpler relative theory. It was later realised that both the geometry of these spaces and their enumerative invariants closely mirror the orbifold theory [Battistella et al. 2023, Theorem B].

We now construct naive Gromov–Witten theory for simple normal crossings pairs without positivity or genus restrictions. In genus zero, we equate the naive invariants with the orbifold invariants, removing the spurious assumptions in [Battistella et al. 2023, Theorem B]. In higher genus, we conjecture that the naive invariants coincide with the constant-term orbifold invariants described in [Tseng and You 2023a].

A reader less familiar with orbifold geometry may wish to examine the definition of naive Gromov–Witten theory, take the equality with the orbifold theory as a given, and skip to the later sections where the orbifold theory is dispensed with.

2.1 Setup

Consider a simple normal crossings pair $(X|D)$. We equip X with the divisorial logarithmic structure corresponding to D. Fix numerical data Λ for a moduli space of logarithmic stable maps to $(X|D)$. This consists of a genus g, a labelling set $\{1, \ldots, n\}$ for the marked points, a curve class $\beta \in A_1(X)$, and tangency orders to D at each marked point. The tangency orders are encoded by a list of integral points

$$\alpha_1, \ldots, \alpha_n \in \Sigma_X(\mathbb{N}).$$

Let $D_1, \ldots, D_k$ be the smooth components of the divisor. The tangency conditions are equivalently given by nonnegative integers c_{ij} for $1 \leq i \leq k$ and $1 \leq j \leq n$. If w_i is the primitive lattice vector generating the ray corresponding to D_i, then for each j we have

$$\alpha_j = \sum_{1 \leq i \leq k} c_{ij} w_i.$$

Remark 2.1 If every intersection of divisor components $\bigcap_{i \in I} D_i$ is connected, then the data c_{ij} ranging over all i are equivalent to α_j. Connectivity of strata intersections can always be achieved after additional blowups. If the intersections are disconnected, then α_j contains more information than the c_{ij}: it specifies which connected component of the intersection contains p_j.

Remark 2.2 The numerical data satisfies the *global balancing condition* with respect to each divisor component if

$$(3) \qquad \sum_{1 \leq j \leq n} c_{ij} = D_i \cdot \beta.$$

If the global balancing condition is not satisfied, the relevant moduli spaces of maps will be empty. We therefore impose it from the outset.

The data Λ induces numerical data for a moduli space of logarithmic stable maps to any pair $(X|E)$ with $E \subseteq D$ (including $E = \varnothing$). The inclusion $E \subseteq D$ determines a map

$$\mathrm{Trop}(X|D) \to \mathrm{Trop}(X|E)$$

obtained by projecting away from the coordinates corresponding to divisors not contained in E. Given an integral vector α_j in $\mathrm{Trop}(X|D)$, its image in $\mathrm{Trop}(X|E)$ determines the contact data for the new pair. We abuse notation and denote the compatible numerical data by the same symbol Λ. Spaces of logarithmic stable maps are denoted by

$$\mathrm{Log}_\Lambda(X|D).$$

The numerical data Λ induces numerical data for a moduli of twisted[1] stable maps [Abramovich and Vistoli 2002; Cadman 2007]. We consider the multiroot stack [Cadman 2007; Tseng and You 2023a]

$$\mathcal{X} = X_{D,\vec{r}},$$

---

[1]For readers less familiar with orbifold Gromov–Witten theory, the lectures of Abramovich [2008] may prove helpful.

where $\vec{r} = (r_1, \ldots, r_k)$ is a vector of rooting parameters. The moduli problem considers representable maps from twisted curves, namely 1–dimensional Deligne–Mumford stacks with stack structure imposed at nodes and markings. Let $\mathcal{C}$ be a twisted curve, and consider

$$f : \mathcal{C} \to \mathcal{X}.$$

Let D_i / r_i in $\mathcal{X}$ denote the substack lying over D_i in X. The pullback $f^\star \mathcal{O}_{\mathcal{X}}(D_i / r_i)|_{p_j}$ is a character of the isotropy group at p_j. If the order of the stabiliser at p_j is s_j, then this character is an integer k_j lying between 0 and $s_j - 1$. The age of $f^\star \mathcal{O}_{\mathcal{X}}(D_i / r_i)|_{p_j}$ is

$$k_j / s_j \in [0, 1) \cap \mathbb{Q}.$$

The age is related to the tangency order c_{ij} via the formula

$$c_{ij} = r_i \cdot (k_j / s_j).$$

See [Cadman and Chen 2008, (2.1.4)]. The same applies for any $E \subseteq D$. Spaces of twisted stable maps to $\mathcal{X}$ will be denoted by

$$\mathrm{Orb}_\Lambda (X | D),$$

with the auxiliary rooting parameters understood to be arbitrarily large.

Finally, for our comparison arguments we will make use of the following hybrid version of the previous two moduli spaces:

Definition 2.3 We denote by $\mathrm{LogOrb}_\Lambda (X | D)$ the moduli space parametrising

  (i) a logarithmic twisted curve $\mathcal{C}$ in the sense of Olsson [2007], and

  (ii) a commuting diagram of representable logarithmic morphisms

$$
\begin{array}{ccc}
(\mathcal{C}, M_{\mathcal{C}}) & \longrightarrow & (\mathcal{X}, \mathcal{D}) \\
\downarrow & & \downarrow \\
(C, M_C) & \longrightarrow & (X, D)
\end{array}
$$

Denoting the coarse moduli map by $\rho \colon \mathcal{C} \to C$, the twisting is encoded in a Kummer extension of logarithmic structures

$$\rho^* M_C \to M_{\mathcal{C}}.$$

It is by now standard to verify that this problem is represented by a proper Deligne–Mumford stack with logarithmic structure. Its virtual fundamental class is discussed in the proof of Theorem 2.9. This moduli space admits forgetful morphisms to both $\mathrm{Log}_\Lambda (X | D)$ and $\mathrm{Orb}_\Lambda (X | D)$.

2.2 Naive invariants

The *naive space* $N_\Lambda(X|D)$ is defined as the fibre product of logarithmic mapping spaces with respect to the smooth divisor components

(4)
$$
\begin{array}{ccc}
N_\Lambda(X|D) & \longrightarrow & \prod_{i=1}^k \mathrm{Log}_\Lambda(X|D_i) \\
\downarrow & \square & \downarrow \\
M_\Lambda(X) & \longrightarrow & M_\Lambda(X)^k
\end{array}
$$

Crucially, the above fibre product is taken in the category of algebraic stacks, and not in the category of fine and saturated logarithmic algebraic stacks.

We now manipulate this cartesian diagram in order to highlight the existence and role of the *universal naive space*. This will allow us to endow every (geometric) naive space with a virtual fundamental class, and to define naive Gromov–Witten invariants in all genera.

Let $\mathcal{A} = [\mathbb{A}^1/\mathbb{G}_m]$ and $\mathcal{D} = [0/\mathbb{G}_m] \subseteq \mathcal{A}$ so that $(\mathcal{A}|\mathcal{D})$ is the universal smooth pair. We shall denote by $\mathfrak{M}(\mathcal{A})$ the moduli space of prestable curves endowed with a line bundle–section pair (or generalised Cartier divisor). We shall denote by $\mathfrak{Log}(\mathcal{A}|\mathcal{D})$ the moduli space of logarithmic prestable maps to $(\mathcal{A}|\mathcal{D})$, ie logarithmically smooth curves together with a line bundle–section pair such that the section is a logarithmic morphism, where the line bundle is endowed with the sum of the logarithmic structure pulled back from the curve *and* the logarithmic structure induced by the zero section. This space is known to admit a minimal logarithmic structure. Since

$$
\mathrm{Trop}(\mathcal{A}|\mathcal{D}) = \mathbb{R}_{\geq 0},
$$

tropicalising a logarithmic prestable map to $(\mathcal{A}|\mathcal{D})$ produces a tropical curve equipped with a piecewise linear function.

For each smooth divisor component D_i there is a cartesian square [Abramovich et al. 2014, Lemma A(iv)]

(5)
$$
\begin{array}{ccc}
\mathrm{Log}_\Lambda(X|D_i) & \longrightarrow & \mathfrak{Log}(\mathcal{A}|\mathcal{D}) \\
\downarrow & \square & \downarrow \\
M_\Lambda(X) & \longrightarrow & \mathfrak{M}(\mathcal{A})
\end{array}
$$

where the morphism $M_\Lambda(X) \to \mathfrak{M}(\mathcal{A})$ performs postcomposition with the map $X \to \mathcal{A}$ induced by $D_i \subseteq X$. We obtain

$$
\begin{array}{ccc}
N_\Lambda(X|D) & \longrightarrow & \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \\
\downarrow & \square & \downarrow \\
M_\Lambda(X) & \longrightarrow & \mathfrak{M}(\mathcal{A})^k
\end{array}
$$

and this square fits into the commutative diagram

$$
(6)
$$

$$
\begin{array}{ccccc}
\mathrm{N}_\Lambda(X|D) & \longrightarrow & \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}^k} \mathfrak{M} & \longrightarrow & \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \\
\downarrow & A & \downarrow & B & \downarrow \\
\mathrm{M}_\Lambda(X) & \xrightarrow{\ p\ } & \mathfrak{M}(\mathcal{A})^k \times_{\mathfrak{M}^k} \mathfrak{M} & \longrightarrow & \mathfrak{M}(\mathcal{A})^k \\
 & & \downarrow & C & \downarrow \\
 & & \mathfrak{M} & \xrightarrow{\ \Delta\ } & \mathfrak{M}^k
\end{array}
$$

The squares C and $B + C$ are cartesian, and therefore B is cartesian. Since $A + B$ is cartesian, it follows that A is cartesian. Finally, $\mathfrak{M}(\mathcal{A})^k \times_{\mathfrak{M}^k} \mathfrak{M} = \mathfrak{M}(\mathcal{A}^k)$, and so we obtain the desired square

$$
(7)
$$

$$
\begin{array}{ccc}
\mathrm{N}_\Lambda(X|D) & \longrightarrow & \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}^k} \mathfrak{M} \\
\downarrow & \square & \downarrow \\
\mathrm{M}_\Lambda(X) & \xrightarrow{\ p\ } & \mathfrak{M}(\mathcal{A}^k)
\end{array}
$$

Definition 2.4 The *naive space for the universal target* is the stack

$$
\mathfrak{NLog}(\mathcal{A}^k|\partial\mathcal{A}^k) := \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}^k} \mathfrak{M}.
$$

Remark 2.5 A diagram chase shows that the naive space for the universal target is equivalently given by any of the three fibre products

$$
\mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}^k} \mathfrak{M} = \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}(\mathcal{A})^k} \mathfrak{M}(\mathcal{A}^k) = \prod_{i=1}^{k} \mathfrak{Log}(\mathcal{A}^k|\mathcal{D}_i) \times_{\mathfrak{M}(\mathcal{A}^k)^k} \mathfrak{M}(\mathcal{A}^k),
$$

where $\mathcal{D}_i \subseteq \mathcal{A}^k$ is the i^{th} coordinate hyperplane. The third description is closest to the definition of the naive space for geometric targets. The first description is the most useful for proofs.

We claim that $p\colon \mathrm{M}_\Lambda(X) \to \mathfrak{M}(\mathcal{A}^k)$ has a relative perfect obstruction theory. If f denotes the universal stable map and π is the universal curve, the obstruction theory is the dual of

$$
(8) \qquad\qquad\qquad\qquad\qquad \boldsymbol{R}^\bullet \pi_\star f^\star T^{\log}_{X|D}.
$$

When $k = 1$ the result is [Abramovich et al. 2014, Proposition 3.1.1]. The general case has the same proof, noting that when X is equipped with the divisorial logarithmic structure from D, we have

$$
\Omega_{X/\mathrm{Log}} = \Omega_{X/\mathcal{A}^k} = \Omega^{\log}_{X|D}.
$$

The relative obstruction theory (8) for p induces a virtual pullback $p_v^!$ on Chow groups [Manolache 2012].

The naive space for the universal target is typically not pure dimensional when $k > 1$ (see Remark 2.8). Nevertheless, its presentation as an intersection of Artin stacks as above endows it with a class of the expected dimension, by Fulton–Kresch intersection theory [Kresch 1999]. Fixing the numerical

data, each factor $\mathfrak{Log}(\mathcal{A}|\mathcal{D})$ is irreducible and carries a fundamental class [Abramovich and Wise 2018, Proposition 1.6.1]. The *virtual class for the naive space* is therefore defined as

$$[\mathrm{N}_\Lambda(X|D)]^{\mathrm{vir}} := p_{\mathrm{v}}^! \Delta^! \big([\mathfrak{Log}(\mathcal{A}|\mathcal{D})] \times \cdots \times [\mathfrak{Log}(\mathcal{A}|\mathcal{D})]\big).$$

Unusually, this virtual class arises via virtual pullback $p_{\mathrm{v}}^!$ from a reducible base. On the other hand, the composition $p_{\mathrm{v}}^! \Delta^!$ of pullbacks can be viewed as endowing $\mathrm{N}_\Lambda(X|D)$ with a perfect obstruction theory over the irreducible $\mathfrak{Log}(\mathcal{A}|\mathcal{D})^k$.

Remark 2.6 We verify that Kresch's hypotheses are satisfied in our case. The stabilisers in $\mathfrak{M}$ are affine, and hence by [Kresch 1999, Proposition 3.5.9] the stack admits a stratification by global quotients. Since it is also smooth, by [Kresch 1999, page 530] it admits a refined intersection product; see also [Bae and Schmitt 2022, Section 3.1] for a discussion of these issues in a similar context.

We now define numerical naive Gromov–Witten invariants. The space $\mathrm{N}_\Lambda(X|D)$ admits a map

$$\mathrm{N}_\Lambda(X|D) \to \mathrm{M}_\Lambda(X),$$

and by composing with the evaluation maps $\mathrm{ev}_i : \mathrm{M}_\Lambda(X) \to X$ we obtain evaluation maps to X. Fix a marked point p_i and let W_i be the intersection of divisor components D_j with respect to which p_i has positive contact order. By definition of the naive space, the evaluation morphism at p_i on $\mathrm{N}_\Lambda(X|D)$ factors through this intersection. We obtain

$$\mathrm{ev}_i : \mathrm{N}_\Lambda(X|D) \to W_i.$$

Definition 2.7 *Naive Gromov–Witten invariants* of $(X|D)$ with numerical data Λ are defined as tautological integrals over the naive virtual class,

$$\langle \tau_{a_1}(\gamma_1), \ldots, \tau_{a_n}(\gamma_n) \rangle_\Lambda^{\mathrm{N}(X|D)} = \int_{[\mathrm{N}_\Lambda(X|D)]^{\mathrm{vir}}} \prod_{i=1}^n \psi_i^{a_i} \, \mathrm{ev}_i^\star(\gamma_i),$$

where $W_i \subset X$ is the stratum of X determined by the marked point p_i, the class γ_i is in the Chow cohomology of W_i, and the ψ–classes are pulled back from $\mathrm{M}_\Lambda(X)$.

Remark 2.8 The stack $\mathfrak{NLog}(\mathcal{A}^k|\partial\mathcal{A}^k)$ is a fibre product of toroidal, ie logarithmically smooth, stacks. This fibre product is not transverse along certain boundary strata, and may consist of several irreducible components. The failure of transversality can be detected from the tropicalisations [Abramovich and Karu 2000; Molcho 2021]. The fact that there are many irreducible components can be seen from the examples in [Nabijou and Ranganathan 2022, Section 1]. These examples consider the space of naive and logarithmic maps to $\mathbb{P}^2$ relative to two lines; for brevity we call it the *geometric* naive space. The examples imply the lack of irreducibility in the universal setting. Indeed, this geometric target is logarithmically convex, and therefore the space of logarithmic maps to this target is smooth over the space of logarithmic maps to the Artin fan. Similarly, by the description of the obstruction theory above, the map from this geometric naive space to the universal one is smooth. The examples in [Nabijou and Ranganathan 2022] identify points

in the geometric naive space of $\mathbb{P}^2$ that do not lie in the closure of the locus of nondegenerate maps. It follows that there exist naive maps to the Artin fan that cannot be deformed to nondegenerate maps.

The strata of the geometric and universal naive spaces can be understood as fibre products of strata of the geometric and universal spaces of maps to smooth pairs. The relationship between the strata of these latter spaces can be understood, for example, using [Kennedy-Hunt et al. 2022, Section 1.1].

2.3 Naive–orbifold correspondence

We now restrict our attention to *genus zero*. We show that naive invariants and orbifold invariants coincide in this case. This removes the convexity hypothesis of [Battistella et al. 2023, Theorem B].

Theorem 2.9 *There is a diagram of spaces connected by virtually birational morphisms*

$$\mathrm{Orb}_\Lambda(X|D) \to \mathrm{NOrb}_\Lambda(X|D) \leftarrow \mathrm{NLogOrb}_\Lambda(X|D) \to \mathrm{N}_\Lambda(X|D).$$

This means that every space in the above diagram is endowed with a virtual fundamental class, and these classes coincide under pushforward along the proper morphisms in the diagram. The auxiliary spaces $\mathrm{NOrb}_\Lambda(X|D)$ and $\mathrm{NLogOrb}_\Lambda(X|D)$ will be defined below.

Corollary 2.10 *The naive and orbifold invariants coincide*:

$$\langle \tau_{a_1}(\gamma_1), \ldots, \tau_{a_n}(\gamma_n) \rangle_\Lambda^{\mathrm{N}(X|D)} = \langle \tau_{a_1}(\gamma_1), \ldots, \tau_{a_n}(\gamma_n) \rangle_\Lambda^{\mathrm{Orb}(X|D)}.$$

Proof This follows immediately from the previous theorem and the projection formula □

Proof of Theorem 2.9 We start with the case of a smooth pair. This is a minor modification of the main results of [Abramovich et al. 2017]. For a fixed rooting index r, the morphism

$$\mathcal{A}_r \to \mathcal{A}$$

will denote the morphism $[\mathbb{A}^1/\mathbb{G}_\mathrm{m}] \to [\mathbb{A}^1/\mathbb{G}_\mathrm{m}]$ given by $t \mapsto t^r$. This is the *universal rth root construction*. While $\mathcal{A}_r$ is abstractly isomorphic to $\mathcal{A}$, we retain the subscript to indicate that the morphism $\mathcal{A}_r \to \mathcal{A}$ is not the identity.

Consider the moduli stack $\mathfrak{Orb}(\mathcal{A}_r)$ parametrising *representable* maps $\mathcal{C} \to \mathcal{A}_r$ from a prestable twisted curve, ie the data of a generalised Cartier divisor $(\mathcal{L}, s)$ on $\mathcal{C}$ such that, for every point p in $\mathcal{C}$ at which $s(p) = 0$, the induced action $\mu_{r(p)} \curvearrowright \mathcal{L}_p$ is faithful, where the integer $r(p)$ is the twisting index. Furthermore, the twisting index is required to divide r, so that the rth tensor power of $\mathcal{L}$ is pulled back from the coarse curve of $\mathcal{C}$.

On the other hand, we define a moduli stack $\mathfrak{LogOrb}(\mathcal{A}_r | \mathcal{D}_r)$ whose objects consist of the following data:

  (i)   a logarithmic twisted curve $\mathcal{C}$,

  (ii) a representable generalised Cartier divisor $(\mathcal{L}, s)$ on $\mathcal{C}$, and

  (iii) an enhancement of the section to a logarithmic morphism $s: \mathcal{C} \to \mathcal{L}$, where $\mathcal{L}$ is equipped with the sum of the logarithmic structure induced by the pullback of the logarithmic structure on $\mathcal{C}$ and the divisorial logarithmic structure corresponding to the zero section.

It is not hard to see that this moduli space admits a minimal logarithmic structure too. The last condition above ensures that taking the r^{th} power of $(\mathcal{L}, s)$ and forgetting the Kummer extension of the logarithmic structure defines a morphism $\mathfrak{Log}\mathfrak{Orb}(\mathcal{A}_r | \mathcal{D}_r) \to \mathfrak{Log}(\mathcal{A} | \mathcal{D})$. On the other hand, forgetting the logarithmic structure in toto, we obtain a morphism $\mathfrak{Log}\mathfrak{Orb}(\mathcal{A}_r | \mathcal{D}_r) \to \mathfrak{Orb}(\mathcal{A}_r)$.

We have thus produced a roof

$$\text{(9)} \qquad \mathfrak{Orb}(\mathcal{A}_r) \leftarrow \mathfrak{Log}\mathfrak{Orb}(\mathcal{A}_r | \mathcal{D}_r) \to \mathfrak{Log}(\mathcal{A} | \mathcal{D}).$$

As is typical when dealing with an Artin stack which is not of finite type, we must pass to an appropriate open substack which contains the image of the geometric moduli space under consideration. In our context this has been discussed carefully in [Abramovich et al. 2017], where primed superscripts are used to indicate the chosen open substacks. From now on we replace each of the entries in (9) by the appropriate open substack.

With this replacement, minor modifications of the arguments of Abramovich, Cadman, and Wise then show that the morphisms in (9) are birational. The passage from universal to geometric targets is then achieved using the following cartesian squares:

$$
\begin{array}{ccc}
\text{Log}_\Lambda(X|D) \longrightarrow \text{M}_\Lambda(X) & \text{Orb}_\Lambda(X|D) \longrightarrow \text{M}_\Lambda(X) & \text{LogOrb}_\Lambda(X|D) \longrightarrow \text{M}_\Lambda(X) \\
\downarrow \quad \square \quad \downarrow & \downarrow \quad \square \quad \downarrow & \downarrow \quad \square \quad \downarrow \\
\mathfrak{Log}(\mathcal{A}|\mathcal{D}) \longrightarrow \mathfrak{M}(\mathcal{A}) & \mathfrak{Orb}(\mathcal{A}_r) \longrightarrow \mathfrak{M}(\mathcal{A}) & \mathfrak{Log}\mathfrak{Orb}(\mathcal{A}_r|\mathcal{D}_r) \longrightarrow \mathfrak{M}(\mathcal{A})
\end{array}
$$

The first square is cartesian by [Abramovich et al. 2014, Lemma A(iv)]. Using the definition of the root stack as $X \times_{\mathcal{A}} \mathcal{A}_r$, it is easy to see that the second square is cartesian. The third square is taken as the definition of $\text{LogOrb}_\Lambda(X|D)$.

In each case, the virtual fundamental class is the pullback of the fundamental class of the universal moduli space, using the virtual pullback morphism induced by the perfect obstruction theory for the map $\text{M}_\Lambda(X) \to \mathfrak{M}(\mathcal{A})$ constructed in the previous section. For the first square, this follows by a direct comparison of obstruction theories. For the second square, it follows by the identity

$$\text{(10)} \qquad p^\star T_{X|D}^{\log} = T_{\mathcal{X}|\mathcal{D}}^{\log},$$

where $p \colon \mathcal{X} \to X$ is the morphism from the root stack to its coarse moduli space. This holds because p is logarithmically étale. Finally, for the third square this is taken as the definition of the virtual class on $\text{LogOrb}_\Lambda(X|D)$.

From this, it follows that the birational diagram (9) pulls back to a virtually birational diagram

$$\text{Orb}_\Lambda(X|D) \leftarrow \text{LogOrb}_\Lambda(X|D) \to \text{Log}_\Lambda(X|D).$$

This completes the discussion of smooth pairs. Now let $D = D_1 + \cdots + D_k$ be simple normal crossings. Each of the three enumerative theories for smooth pairs considered above defines a corresponding naive

theory for simple normal crossings pairs. To accompany Definition 2.4, we set forth some naturally defined universal naive spaces

$$\mathfrak{NOrb}(\mathcal{A}^k | \partial \mathcal{A}^k) := \prod_{i=1}^{k} \mathfrak{Orb}(\mathcal{A}_{r_i}) \times_{\mathfrak{M}^k} \mathfrak{M}, \quad \mathfrak{NLogOrb}(\mathcal{A}^k | \partial \mathcal{A}^k) := \prod_{i=1}^{k} \mathfrak{LogOrb}(\mathcal{A}_{r_i} | \mathcal{D}_{r_i}) \times_{\mathfrak{M}^k} \mathfrak{M}.$$

While even the *universal* naive spaces may be badly behaved beyond the base case of a smooth pair, each carries a virtual fundamental class obtained by capping the product of fundamental classes of the factors with the diagonal in the space of curves. Since for smooth pairs the morphisms in (9) preserve the fundamental classes, a diagram chase gives a virtually birational diagram

$$\mathfrak{NOrb}(\mathcal{A}^k | \partial \mathcal{A}^k) \leftarrow \mathfrak{NLogOrb}(\mathcal{A}^k | \partial \mathcal{A}^k) \rightarrow \mathfrak{NLog}(\mathcal{A}^k | \partial \mathcal{A}^k).$$

The geometric naive spaces are defined by fibring the universal naive spaces over the map

$$\mathrm{M}_\Lambda(X) \rightarrow \mathfrak{M}(\mathcal{A})^k \times_{\mathfrak{M}^k} \mathfrak{M} = \mathfrak{M}(\mathcal{A}^k).$$

Their virtual classes are pulled back from the virtual classes of the universal spaces, and the standard comparison of obstruction theories gives a virtually birational diagram

$$\mathrm{NOrb}_\Lambda(X|D) \leftarrow \mathrm{NLogOrb}_\Lambda(X|D) \rightarrow \mathrm{NLog}_\Lambda(X|D) = \mathrm{N}_\Lambda(X|D).$$

It remains to show that the naive orbifold theory coincides with the orbifold theory. We will show that there is a virtually birational morphism

$$\mathrm{Orb}_\Lambda(X|D) \rightarrow \mathrm{NOrb}_\Lambda(X|D).$$

This is a generalised form of the relative product formula for orbifold invariants established in [Battistella et al. 2023, Theorem 3.1]. Consider the diagram

$$
\begin{array}{ccccc}
\mathrm{Orb}_\Lambda(X|D) & \longrightarrow & \mathrm{NOrb}_\Lambda(X|D) & \longrightarrow & \mathrm{M}_\Lambda(X) \\
\downarrow & A & \downarrow & B & \downarrow \\
\mathfrak{Orb}\big(\prod_{i=1}^{k} \mathcal{A}_{r_i}\big) & \xrightarrow{\ \varphi\ } & \mathfrak{NOrb}(\mathcal{A}^k | \partial \mathcal{A}^k) & \longrightarrow & \mathfrak{M}(\mathcal{A}^k)
\end{array}
$$

The morphism φ is obtained by projecting $\mathcal{C} \rightarrow \prod_{i=1}^{k} \mathcal{A}_{r_i}$ onto each factor of the target and taking the minimal partial coarsening of $\mathcal{C}$ which ensures that the composed map is representable. This partial coarsening is discussed in [Abramovich and Vistoli 2002, Section 9; Abramovich et al. 2011, Theorem 3.1], where it is referred to as the "relative coarse moduli space".

The square B is cartesian by definition, and $A + B$ is clearly cartesian. It follows that A is cartesian. The vertical arrows in A are equipped with perfect obstruction theories, which are compatible by the identity (10). The bases of both vertical arrows in A are equipped with virtual classes, recalled below. These virtual classes virtually pull back to the virtual classes on the top row of A. By Manolache's formalism [2012, Theorem 4.1], in order to show that $\mathrm{Orb}_\Lambda(X|D) \rightarrow \mathrm{NOrb}_\Lambda(X|D)$ is virtually birational it suffices to show that φ identifies virtual classes.

We first recall the virtual classes on the bottom row of A. The space $\mathfrak{Orb}\left(\prod_{i=1}^{k} \mathcal{A}_{r_i}\right)$ is equipped with a virtual class by a construction of Abramovich, Cadman, and Wise. The forgetful morphism to the moduli stack $\mathfrak{M}^{\mathrm{tw}}$ of twisted curves factors as

$$\mathfrak{Orb}\left(\prod_{i=1}^{k} \mathcal{A}_{r_i}\right) \to \mathfrak{Orb}(\mathcal{B}\mathbb{G}_{\mathrm{m}}^{k}) \to \mathfrak{M}^{\mathrm{tw}}.$$

The first arrow takes a tuple of line bundles and sections and forgets the sections. The second arrow then forgets the line bundles. By [Abramovich et al. 2017, Section 5.2] the morphism

$$\mathfrak{Orb}\left(\prod_{i=1}^{k} \mathcal{A}_{r_i}\right) \to \mathfrak{Orb}(\mathcal{B}\mathbb{G}_{\mathrm{m}}^{k})$$

is virtually smooth and equips the universal orbifold space $\mathfrak{Orb}\left(\prod_{i=1}^{k} \mathcal{A}_{r_i}\right)$ with a virtual class. It coincides with the standard virtual class construction in orbifold Gromov–Witten theory.

On the other hand, the virtual class on $\mathfrak{NOrb}(\mathcal{A}^k | \partial \mathcal{A}^k)$ is defined via diagonal pullback, analogously to the virtual class of $\mathfrak{NLog}(\mathcal{A}^k | \partial \mathcal{A}^k)$ defined in Section 2.2. The orbifold product formula for the target $\prod_{i=1}^{k} \mathcal{A}_{r_i}$ then shows that the morphism

$$\varphi \colon \mathfrak{Orb}\left(\prod_{i=1}^{k} \mathcal{A}_{r_i}\right) \to \mathfrak{NOrb}(\mathcal{A}^k | \partial \mathcal{A}^k)$$

identifies virtual classes [Andreini et al. 2016; Behrend 1999]. Note that the product formula is typically stated for schematic or Deligne–Mumford targets and over stable curves, but the arguments apply without change in this setting. □

Remark 2.11 In higher genus, we conjecture (Conjecture W) that the naive invariants coincide with the constant-term orbifold invariants of [Tseng and You 2023a]. The missing step is a universal formulation of the higher-genus smooth divisor logarithmic-orbifold correspondence [Janda et al. 2020; Tseng and You 2020]. We will return to this in future work.

3 Geometry of the naive space

We henceforth restrict to genus zero. In Theorem 2.9 we proved that orbifold and naive invariants coincide. We now begin a study of the more delicate relationship between naive and logarithmic invariants.

3.1 Numerical criteria

Consider the finite morphism from the naive space to the absolute mapping space

$$\mathrm{N}_{\Lambda}(X | D) \to \mathrm{M}_{\Lambda}(X).$$

The geometric points of the image have an elementary characterisation, extending an insight originally due to Gathmann. In practice this allows one to explore the geography of the naive space with relative ease.

Gathmann's condition [Gathmann 2002, Remark 1.4] For each divisor component D_i, consider a connected component Z of the subscheme $f^{-1}(D_i) \subseteq C$. Then:

(i) If Z is zero-dimensional, it must be supported at a marked point p_j and have length c_{ij}.

(ii) Otherwise, $Z^{\mathrm{red}} = C_0 \subseteq C$ is a subcurve. In this case, let $C_1, \ldots, C_r$ denote the irreducible components of C adjacent to C_0 and let $q_1, \ldots, q_r$ be the connecting nodes. For $1 \leq j \leq r$ let $m_j \in \mathbb{N}$ be the contact order of $f|_{C_j}$ to D_i at the point q_j in C_j. Then we must have

$$\sum_{p_j \in C_0} c_{ij} - \sum_{1 \leq j \leq r} m_j = D_i \cdot f_\star[C_0] = \deg f^\star \mathcal{O}_X(D_i)|_{C_0}.$$

Gathmann's condition is closely related, though not identical, to the tropical balancing condition, which we now recall:

Balancing condition Consider the logarithmic scheme $S = (\mathrm{Spec}\,\Bbbk, Q)$ with monoid Q and let $f : \mathcal{C} \to X$ be a logarithmic stable map over S. The tropicalisation of $\mathcal{C} \to S$ is a family of tropical curves

$$\Box \to \sigma,$$

and tropicalising f gives a morphism of cone complexes

$$\mathrm{f} \colon \Box \to \Sigma.$$

A piecewise linear function on $\Box$ (or on Σ) gives rise to a generalised Cartier divisor on C (or on X); see Section 1.4. We can now state the balancing condition [Gross and Siebert 2013, Section 1.4].

Let $f : \mathcal{C} \to X$ be as above and fix a piecewise linear function $\alpha : \Sigma \to \mathbb{R}$. There are two ways to produce a line bundle on $\mathcal{C}$: take the line bundle $f^\star \mathcal{O}_X(\alpha)$ where $\mathcal{O}_X(\alpha)$ is the line bundle on X associated to α, or take the line bundle $\mathcal{O}_{\mathcal{C}}(\mathrm{f}^\star \alpha)$ on C associated to the piecewise linear function

$$\mathrm{f}^\star \alpha \colon \Box \xrightarrow{\mathrm{f}} \Sigma \xrightarrow{\alpha} \mathbb{R}.$$

The logarithmic morphism f encodes an isomorphism $\mathcal{O}_{\mathcal{C}}(\mathrm{f}^\star \alpha) \cong f^\star \mathcal{O}_X(\alpha)$. The *balancing condition* is the coarser statement that for every irreducible component $C' \subseteq C$ we have

$$(11) \qquad \qquad \deg \mathcal{O}_{\mathcal{C}}(\mathrm{f}^\star \alpha)|_{C'} = \deg f^\star \mathcal{O}_X(\alpha)|_{C'}.$$

Lemma 3.1 *A stable map in $\mathrm{M}_\Lambda(X)$ belongs to the image of $\mathrm{N}_\Lambda(X|D)$ if and only if it satisfies Gathmann's condition with respect to every divisor component D_i.*

Proof Assume that D is smooth and fix a stable map $f : C \to X$ which admits a logarithmic enhancement. Its tropicalisation satisfies the balancing condition, and we will use this to deduce Gathmann's condition.

Consider the identity piecewise linear function on $\Sigma = \mathbb{R}_{\geq 0}$. The associated line bundle is $\mathcal{O}_X(D)$. On the other hand, this piecewise linear function pulls back to the piecewise linear function on $\Box$ with slopes along the edges determined by the tropical map $\mathrm{f} \colon \Box \to \Sigma$. By the balancing condition these two line bundles have the same multidegree. Summing the degrees over all irreducible components of a given $C_0 \subseteq f^{-1}(D)$, we obtain Gathmann's condition.

For the converse, consider the commuting diagram of moduli spaces

$$\text{Kim}_\Lambda(X|D)$$

$$\text{Log}_\Lambda(X|D) \qquad\qquad \text{Li}_\Lambda(X|D)$$

$$\text{M}_\Lambda(X)$$
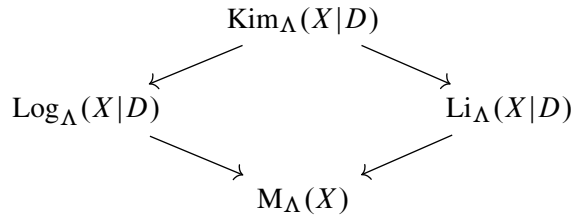
where $\text{Li}_\Lambda(X|D)$ is Li's space [2001; 2002] of relative stable maps to expansions, and $\text{Kim}_\Lambda(X|D)$ is Kim's space [2010] of logarithmic stable maps to expansions.

Fix a stable map $f: C \to X$ satisfying Gathmann's condition. This guarantees the existence of a meromorphic section of $f^\star \mathcal{O}_X(D)$ on each irreducible component of $f^{-1}(D)$ with specified zero and pole locus. These can be assembled to produce a relative stable map to the expanded target; see [Gathmann 2003, Lemma 5.1.12]. It follows that our given stable map admits a lift to $\text{Li}_\Lambda(X|D)$.

It is a direct consequence of Li's original construction that the morphism $\text{Kim}_\Lambda(X|D) \to \text{Li}_\Lambda(X|D)$ is surjective, the difference being given by saturation; see [Abramovich et al. 2014, Lemma 4.2.2; Gross and Siebert 2013, Section 6]. It follows that the given stable map admits a lift to $\text{Kim}_\Lambda(X|D)$, and hence, by commutativity of the diagram, lies in the image of $\text{Log}_\Lambda(X|D) \to \text{M}_\Lambda(X)$.

This completes the proof in the smooth divisor case. The general case follows from the definition (4) of the naive space, which implies that the image of $\text{N}_\Lambda(X|D) \to \text{M}_\Lambda(X)$ is the intersection of the images of $\text{Log}_\Lambda(X|D_i) \to \text{M}_\Lambda(X)$ for $1 \le i \le k$. □

3.2 The data structure of combinatorial types

A logarithmic map $\mathcal{C} \to (X|D)$ over the logarithmic point $(\text{Spec } \Bbbk, Q)$ produces a family of tropical maps $\square \to \Sigma$ over the dual cone τ. The combinatorics of this tropical map enhances the dual graph of the source curve with data recording the vanishing orders of the divisor equations at generic points of components, and at marked and singular points. This data is referred to as the *combinatorial type* [Abramovich et al. 2020a, Section 2.5].

As we now explain, a naive map also has a combinatorial type, which can be packaged in the same data structure. Recall that the cones of a cone complex form a poset, where $\sigma_1 \preceq \sigma_2$ if σ_1 is a face of σ_2.

Definition 3.2 A *combinatorial type* of tropical stable map to Σ consists of the following data:

(i) a finite graph Γ consisting of vertices, finite edges, and semi-infinite legs, where each leg ℓ in $L(\Gamma)$ carries a marking label $j_\ell \in \{1, \ldots, n\}$, and each vertex v in $V(\Gamma)$ carries genus and degree labels

$$g_v \in \mathbb{N} \quad \text{and} \quad \beta_v \in A_1(X),$$

(ii) cones of Σ associated to every vertex, edge, and leg of Γ

$$v \rightsquigarrow \sigma_v, \quad e \rightsquigarrow \sigma_e, \quad \text{and} \quad \ell \rightsquigarrow \sigma_\ell,$$

such that if $v \preceq e$ then $\sigma_v \preceq \sigma_e$ for every vertex v and adjacent edge or leg e,

(iii) for every oriented edge or leg $\vec{e}$ an associated slope $m_{\vec{e}} \in N_{\sigma_e}$ satisfying $m_{\vec{e}} = -m_{\overleftarrow{e}}$.

3.3 Combinatorial type of a naive map

We will associate a combinatorial type to each naive map. It will record the dual graph of the curve, its interaction with the target stratification, and the tangency orders at the special points.

Remark 3.3 The combinatorial type of a naive map fits into the same data structure as the combinatorial type of a logarithmic map. However, there will exist naive types that do not arise from logarithmic maps, or equivalently, naive types that are not the underlying combinatorial type of a tropical map. The latter are said to be *smoothable*. The detection of smoothable types among naive types is the main issue at play in Section 5.

Fix a genus-zero naive stable map $f : C \to (X|D)$.

Step I (Labelled graph) The graph Γ is the dual graph of C, with vertices corresponding to irreducible components, finite edges to nodes, and semi-infinite legs to markings. The leg and vertex labels are determined in the obvious way. Since we work in genus zero, all $g_v = 0$.

Step II (Associated cones) Let $C_v \subseteq C$ be an irreducible component. Let $X_v \subseteq X$ be the unique minimal stratum of $(X|D)$ through which the map $C_v \to X$ factors. This corresponds to a cone $\sigma_v \preceq \Sigma$ which is defined to be the cone associated to v. The cones σ_e and σ_ℓ are defined similarly.

Step III (Slopes along legs) The slopes along the legs are part of the input data, corresponding to the specified tangencies at the markings. In the notation of Section 2.1, we have

$$m_{\vec{\ell}} = \alpha_{j(\ell)} \in \Sigma(\mathbb{N}),$$

where $\vec{\ell}$ is oriented to point away from its supporting vertex. Gathmann's condition implies that

$$m_{\vec{\ell}} \in \sigma_\ell(\mathbb{N}) \subseteq N_{\sigma_\ell}.$$

Step IV (Slopes along edges) The slopes $m_{\vec{e}}$ along the finite edges are determined by formally imposing the balancing condition at every vertex. Fix a divisor component D_i. For each vertex v of Γ, the balancing equation gives

$$(12) \qquad \sum_{p_j \in C_v} c_{ij} + \sum_{v \preceq e} (m_{\vec{e}})_i = D_i \cdot \beta_v,$$

where every edge is oriented to point away from v. The $(m_{\vec{e}})_i$ are unknown integers to be solved for. The first term is a sum over markings p_j in C_v, or equivalently legs ℓ_j containing v. The tangency order c_{ij} is equal to the component of α_j in the direction corresponding to D_i.

The system of equations (12) will appear repeatedly in this section, and will be referred to as the *balancing equations*.

Lemma 3.4 *The system of balancing equations* (12) *has a unique solution in the* $(m_{\vec{e}})_i$.

Proof There are $\#V(\Gamma)$ balancing equations (12), and a single linear dependency; summing all the equations, the $(m_{\vec{e}})_i$ cancel and we are left with the global balancing constraint

$$\sum_{1 \leq j \leq n} c_{ij} = D_i \cdot \beta, \tag{13}$$

which holds by assumption; see Remark 2.2. We obtain $\#V(\Gamma) - 1$ linearly independent equations. Since Γ has genus zero, $\#V(\Gamma) - 1 = \#E(\Gamma)$, which coincides with the number of unknowns. Hence there is a unique solution for the $(m_{\vec{e}})_i$.

The uniqueness may be deduced constructively. Select a root vertex and orient the edges of Γ to point towards the root. Place a level structure on Γ which respects the orientation, ie an assignments of heights $h(v) \in \mathbb{R}$ such that $h(v_1) < h(v_2)$ if there is an oriented edge $\vec{e}$ connecting v_1 to v_2. Iterate along Γ towards the root, from the lowest level to the highest. At each step, the vertex encountered has several incoming edges whose $(m_{\vec{e}})_i$ have previously been determined, and a single outgoing edge whose $(m_{\vec{e}})_i$ is currently unknown. The latter is determined uniquely by solving the balancing equation at the current vertex. This process terminates at the root vertex, where the global balancing constraint (13) ensures that the balancing equation at the root vertex is satisfied. $\qquad\square$

By the above lemma, (12) has a solution for each D_i, and this produces integers $(m_{\vec{e}})_i$ for all $\vec{e}$ and i. We now use the assumption that $f : C \to (X|D)$ is a naive map to assemble these integers into slope vectors $m_{\vec{e}} \in N_{\sigma_e}$.

Fix an oriented edge $\vec{e}$ and let $w_{i_1}, \ldots, w_{i_c}$ be the primitive generators of the cone σ_e, corresponding to divisor components $D_{i_1}, \ldots, D_{i_c}$.

Lemma 3.5 *We have* $(m_{\vec{e}})_i = 0$ *for* $i \notin \{i_1, \ldots, i_c\}$.

Once this is proved, we can define the slope along $\vec{e}$ as

$$m_{\vec{e}} := (m_{\vec{e}})_{i_1} w_{i_1} + \cdots + (m_{\vec{e}})_{i_c} w_{i_c} \in N_{\sigma_e}, \tag{14}$$

which completes the construction.

It remains to establish Lemma 3.5. This follows from the more general Lemma 3.8, which asserts that the formal tangency orders $(m_{\vec{e}})_i$ constructed above coincide with the geometric tangency orders, provided the latter can be defined. This makes crucial use Gathmann's condition.

Definition 3.6 A node q in C has *well-defined geometric tangency* with respect to D_i if there exists a neighbourhood U of q such that $f(U) \not\subseteq D_i$.

Equivalently, a point q in C has well-defined geometric tangency if $f(C_1) \nsubseteq D_i$ or $f(C_2) \nsubseteq D_i$, where $C_1, C_2 \subseteq C$ are the irreducible components adjacent to q.

Definition 3.7 For every node q with well-defined geometric tangency, let e in $E(\Gamma)$ be the corresponding edge and denote the *geometric tangency* by

$$(m_{\vec{e}})'_i \in \mathbb{N}.$$

Here if $f(C_1) \nsubseteq D_i$ and $f(C_2) \subseteq D_i$, we orient e to point from v_1 to v_2. If neither component is mapped into D_i then Gathmann's condition, combined with the global balancing constraint (3), implies that $f(q_e)$ does not lie in D_i and so the geometric tangency is zero.

Lemma 3.8 *If q in C is a node with well-defined geometric tangency and e in $E(\Gamma)$ is the corresponding edge, then*

$$(m_{\vec{e}})_i = (m_{\vec{e}})'_i.$$

Proof Consider the collection of geometric tangencies $(m_{\vec{e}})'_i$ and formal tangencies $(m_{\vec{\ell}})_i$ for each leg ℓ. Note that the latter satisfy the global balancing constraint (13). For all remaining edges, introduce unknowns $(m_{\vec{e}})'_i$ and consider the resulting system of balancing equations, as in the preceding discussion.

The system coincides with the one in (12), with a minor difference: the $(m_{\vec{e}})'_i$ corresponding to nodes with well-defined tangency are specified at the start. As such, there are fewer unknowns.

There are also more linear dependencies. In fact, the system of equations decomposes into a sum of systems of equations with independent sets of variables, one for each maximal subcurve $C_0 \subseteq f^{-1}(D_i)$. The same argument we have used previously, applied to each of these independent systems, guarantees a unique solution for the $(m_{\vec{e}})'_i$. Thus we have two solutions for the balancing equations: the $(m_{\vec{e}})_i$ and the $(m_{\vec{e}})'_i$. By uniqueness they coincide, so the formal tangency orders $(m_{\vec{e}})_i$ coincide with the geometric tangency orders $(m_{\vec{e}})'_i$ whenever the latter are defined. $\square$

Proof of Lemma 3.5 Fix an index $i \notin \{i_1, \ldots, i_c\}$. By the definition of σ_e we have $f(q_e) \notin D_i$. Let C_1 and C_2 be the irreducible components of C adjacent to q_e. Certainly either $f(C_1)$ or $f(C_2)$ is not contained in D_i, so q_e has well-defined tangency. Since $f(q_e)$ does not lie in D_i this geometric tangency is zero, so $(m_{\vec{e}})_i = 0$ as claimed. $\square$

This completes the construction of the combinatorial type associated to a naive stable map.

Remark 3.9 Given a naive stable map $f : C \to (X|D)$ the preceding construction assigns formal tangency vectors $m_{\vec{e}}$ along the finite edges of the dual graph. Lemma 3.8 can be restated as follows:

Consider a vertex v of Γ and an edge or leg e of Γ containing v, and orient e to point away from v. Suppose that D_i is a divisor component which contains the stratum $X_e = X_{\sigma_e}$ but does not contain the larger stratum $X_v = X_{\sigma_v}$. Then the component $(m_{\vec{e}})_i \in \mathbb{Z}$ of the vector $m_{\vec{e}} \in N_{\sigma_e}$ must agree with the geometric tangency order of $f|_{C_v}$ to D_i at the special point q_e in C_v. The geometric tangency order is well defined because $f(C_v) \nsubseteq D_i$. Note in particular that we must have $(m_{\vec{e}})_i > 0$.

Example 3.10 Let $X = \mathbb{P}^2$ with $D = L_1 + L_2$ the sum of two lines. Consider a degree 4 stable map with the following dual graph:
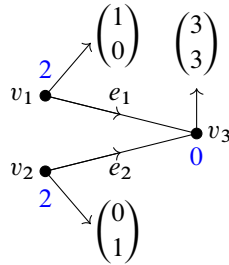
The numbers give the degree of the stable map on each component, and the column vectors give the tangency orders at the marked points. Orienting the edges as above, the system of balancing equations at the vertices is

$$v_1: \quad m_{\vec{e}_1} + \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 2 \\ 2 \end{pmatrix}, \qquad v_2: \quad m_{\vec{e}_2} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 2 \end{pmatrix}, \qquad \text{and} \qquad v_3: \quad -m_{\vec{e}_1} - m_{\vec{e}_2} + \begin{pmatrix} 3 \\ 3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

to which there is a unique solution, namely

$$m_{\vec{e}_1} = \begin{pmatrix} 1 \\ 2 \end{pmatrix} \quad \text{and} \quad m_{\vec{e}_2} = \begin{pmatrix} 2 \\ 1 \end{pmatrix}.$$

3.4 Naive types vs logarithmic types

We briefly address the combinatorial question of how the naive and logarithmic types differ. Although this discussion is not logically necessary, it provides context for the arguments of Section 5.

Definition 3.11 A combinatorial type of tropical stable map is *smoothable* if there exist assignments of edge lengths metrising the graph Γ

$$\ell_e \in \mathbb{R}_{>0},$$

and a continuous map $f: \Gamma \to \Sigma$ such that

(i) every vertex, edge, and leg p of Γ is mapped into the interior of the cone σ_p, and

(ii) on every edge or leg e, the map $f|_e: e \to \sigma_e$ is affine with slope equal to $m_{\vec{e}}$.

The positivity of the ℓ_e and the fact that faces of Γ must map to the interiors of their cones is crucial. Without it, every combinatorial type is smoothable for trivial reasons. The following rephrases [Gross and Siebert 2013, Remark 1.21] and the surrounding discussion. The proof is omitted.

Proposition 3.12 *Let $f: C \to (X|D)$ be a logarithmic map over a geometric point. The combinatorial type of f is smoothable.*

The next example shows that naive combinatorial types are not always smoothable:

Example 3.13 Consider Example 3.10 above, and suppose that the associated cones are

$$\sigma_{v_1} = \sigma_{v_2} = 0 \quad \text{and} \quad \sigma_{v_3} = \mathbb{R}^2_{\geq 0}.$$

There is then no smoothing, because a distance-preserving map f would impose the simultaneous equations on the edge lengths

$$\ell(e_1) = 2\ell(e_2) \quad \text{and} \quad 2\ell(e_1) = \ell(e_2),$$

which are only satisfied when $\ell(e_1) = \ell(e_2) = 0$.

The following purely combinatorial question arises:

Question 3.14 Which combinatorial types of naive stable maps are smoothable?

In fact, the answer to this combinatorial question dictates the entire comparison between logarithmic and orbifold Gromov–Witten theory. In Section 5 we prove that after performing suitably many blowups of the target, every combinatorial type of naive stable map is smoothable.

3.5 Naive maps and target blowups

We examine the behaviour of naive maps under target blowups and describe the associated transformation on combinatorial types.

A *simple blowup* of a simple normal crossings pair is the blowup along a stratum. The reduced preimage of the boundary divisor is again simple normal crossings. We examine an iteration of such simple blowups

$$\pi : (X^\dagger | D^\dagger) \to (X | D).$$

It is often useful to note that the strata of a simple normal crossings pair are regularly embedded, with canonically split normal bundles, and that π is a local complete intersection morphism. Since π is a logarithmic modification, it corresponds to a subdivision of cone complexes

$$\Sigma^\dagger \to \Sigma.$$

Consider numerical data Λ for a moduli space of stable maps to $(X | D)$. The data Λ lifts uniquely to numerical data for maps to $(X^\dagger | D^\dagger)$, as follows. The genus and marking set are unchanged. For $1 \leq i \leq n$ we use the identification of integral lattice points of the tropicalisation to lift the tangency data at the marked point p_i:

$$\alpha_i \in \Sigma(\mathbb{N}) = \Sigma^\dagger(\mathbb{N}).$$

It remains to lift the curve class. We show how to do this in the case where π is a single blowup; the general case follows by induction. Suppose then that we have

$$\pi : X^\dagger = \mathrm{Bl}_Z X \to X,$$

a blowup along $Z \subseteq X$ with exceptional divisor $E \subseteq X^\dagger$. Consider the projection

$$\mathrm{p}_E : \Sigma^\dagger \to \mathbb{R}_{\geq 0}$$

corresponding to the forgetful logarithmic morphism $(X^\dagger|D^\dagger) \to (X^\dagger|E)$. Alternatively, this is the piecewise linear function corresponding to the Cartier divisor E. Let

$$d_E := \sum_{i=1}^n \mathrm{p}_E(\alpha_i) \in \mathbb{N}.$$

Let $L \in A_1(X^\dagger)$ be the class of a line in a fibre of the projective bundle $E \to Z$, and note that $E \cdot L = -1$. We define

$$\beta^\dagger := \pi^\star \beta - d_E L.$$

The global balancing constraint (3) is clearly satisfied, and $\pi_\star \beta^\dagger = \beta$. These define the lifted numerical data Λ for $(X^\dagger|D^\dagger)$.

Remark 3.15 Suppose each α_i in $\Sigma(\mathbb{N})$ lies on a ray of Σ. Geometrically, each marking carries tangency to at most one irreducible component of the divisor, and we refer to this as the *disjoint case*. Then $\beta^\dagger = \pi^\star \beta$. We may always reduce to the disjoint case by performing finitely many initial blowups of the target.

The next task is to explain how combinatorial types behave under blowups. A geometric point s of $\mathrm{N}_\Lambda(X|D)$ has a combinatorial type. By the discussion of Section 3.3 this type is determined by the numerical data Λ, and the image of s in $\mathrm{M}_\Lambda(X)$.

Proposition 3.16 *Consider the diagram of mapping spaces with compatible numerical data*

$$\begin{array}{ccc} \mathrm{N}_\Lambda(X^\dagger|D^\dagger) & & \mathrm{N}_\Lambda(X|D) \\ {\scriptstyle \theta^\dagger}\downarrow & & \downarrow{\scriptstyle \theta} \\ \mathrm{M}_\Lambda(X^\dagger) & \xrightarrow{\ \varphi\ } & \mathrm{M}_\Lambda(X) \end{array}$$

Let s be a geometric point of $\mathrm{N}_\Lambda(X^\dagger|D^\dagger)$. Then the set

$$\theta^{-1}(\varphi \circ \theta^\dagger(s))$$

is nonempty and finite, and every element has the same combinatorial type.

The following lemma will be key to the proof. Let C be an n–pointed curve, and fix tangency orders $c_1, \ldots, c_n$ in $\mathbb{N}$ for a map to a smooth pair. A morphism $f : C \to \mathcal{A}$ *admits a logarithmic lift* if f is the underlying map of a logarithmic map with these tangency orders.

Lemma 3.17 *Let C be an n–marked nodal curve over $\mathrm{Spec}\,\mathbb{k}$, and let*

$$f_1, f_2 : C \to \mathcal{A}$$

be two morphisms with (L_1, s_1) and (L_2, s_2) the associated generalised Cartier divisors. Let

$$h : C \to \mathcal{A}$$

be the morphism determined by $(L_1 \otimes L_2, s_1 \otimes s_2)$. Suppose f_1 and f_2 admit logarithmic lifts, where f_i has contact order c_{ij} at the marked point p_j. Then h also admits a logarithmic lift with contact order $c_{1j} + c_{2j}$ at p_j.

Proof Since f_1 and f_2 admit logarithmic enhancements, we may consider the combinatorial types of tropical maps to $\mathbb{R}_{\geq 0}$ of these enhancements. We first predict the combinatorial type of h. The graph Γ is the dual graph of C. If p is a vertex, edge, or leg of Γ define the target cone by the following rule: $\sigma_p = 0$ if and only if p is assigned the cone 0 in the types of both f_1 and f_2. If $\vec{e}$ is a directed edge of Γ, we let $m_{\vec{e}}^i$ denote the slope along this directed edge in the combinatorial type of f_i, viewed as an element in N_{σ_e}. We then define

$$m_{\vec{e}} := m_{\vec{e}}^1 + m_{\vec{e}}^2 \in N_{\sigma_e}.$$

The parallel definition is taken at the legs, as demanded by the lemma. This gives a well-defined combinatorial type for a map to $\mathbb{R}_{\geq 0}$.

Since all types of tropical maps to $\mathbb{R}_{\geq 0}$ are smoothable, we can find a logarithmic structure on $\mathrm{Spec}\,\Bbbk$ with monoid Q and dual cone τ, an enhancement of C to a logarithmic curve $\mathcal{C}$, and a map from the associated family of tropical curves

$$\begin{array}{ccc} \sqsubset & \longrightarrow & \mathbb{R}_{\geq 0} \\ \downarrow & & \\ \tau & & \end{array}$$

As discussed in Section 1.6 we deduce that there is a logarithmic morphism

$$\tilde{h} \colon \mathcal{C} \to \mathcal{A}.$$

Now, the map $\tilde{h}$ satisfies all the demands placed on h in the statement of the lemma, so we need only show that h is isomorphic to the underlying morphism of $\tilde{h}$. Both h and $\tilde{h}$ produce a generalised Cartier divisor, and we must show that these are isomorphic.

At each vertex v of Γ, the line bundle associated to f_i is given by a divisor, specifically the sum of distinguished points of C_v weighted by the slopes of the associated piecewise linear function; see eg [Ranganathan et al. 2019, Proposition 2.4.1]. The line bundle associated to their tensor product is the line bundle associated to the sum of these divisors. It follows immediately from the definition that the line bundles associated to h and $\tilde{h}$ are isomorphic. It is also easy to see that the isomorphism preserves the sections. $\square$

Proof of Proposition 3.16 It follows from Section 3.3 that the fibres of both vertical morphisms in the proposition parametrise maps with the same combinatorial type. The finiteness follows from [Wise 2019, Corollary 1.2]. The main work is to demonstrate nonemptiness.

Consider a single blowup along a smooth stratum

$$(X^{\dagger}|D^{\dagger}) \to (X|D).$$

Let $D_1, \ldots, D_k$ be the divisor components of D. Let $\tilde{D}_1, \ldots, \tilde{D}_k$ denote their strict transforms and E the exceptional divisor. The pullback of D_i is either $\tilde{D}_i$ or $\tilde{D}_i + E$.

Let $C \to (X^\dagger | D^\dagger)$ be a naive stable map and consider the map

$$f \colon C \to X^\dagger \to X$$

obtained by composition, without stabilising. We verify that for each D_i, f admits a logarithmic lift $f_i \colon \mathcal{C}_i \to (X|D_i)$. By (7) we may equivalently check that the generalised Cartier divisor

$$(f^\star \mathcal{O}_X(D_i), s_{D_i})$$

admits a logarithmic lift, when viewed as a map to $\mathcal{A}$. If D_i pulls back to $\widetilde{D}_i$ then there is nothing to prove. If it pulls back to $\widetilde{D}_i + E$ then the result follows from Lemma 3.17 by taking $L_1 = f^\star \mathcal{O}_X(\widetilde{D}_i)$ and $L_2 = f^\star \mathcal{O}_X(E)$ so that $L_1 \otimes L_2 = f^\star \mathcal{O}_X(D_i)$. Since the stabilisation of a logarithmic stable map carries a natural logarithmic structure [Abramovich et al. 2014, Appendix B], the claim follows. $\square$

Remark 3.18 Proposition 3.16 gives a multivalued map between naive spaces, which is a weak form of functoriality. In case it may be useful in future work, we record here that the proof shows that the arrow φ restricts to a morphism between the loci of stable maps that admit a naive lift. In fact, if we use spaces of fine but not necessarily saturated logarithmic maps, as introduced by Chen [2014, Lemma 3.7.4] and advocated by Wise [2019], then the multivalued map above becomes a morphism.

Construction 3.19 (Pushforward of types) As a consequence of the above proposition, although we have not produced a morphism between naive spaces, there is a morphism between their images in the mapping spaces. Let

$$\pi \colon (X^\dagger | D^\dagger) \to (X|D)$$

be an iterated blowup along strata. Given a combinatorial type for a naive map to $(X^\dagger | D^\dagger)$, the induced type of the map to $(X|D)$ is obtained as follows. The blowup is given by a subdivision

$$\mathrm{p} \colon \Sigma^\dagger \to \Sigma.$$

The labelled source graph Γ is replaced by its stabilisation, obtained by deleting bivalent or univalent vertices whose curve classes are exceptional. For every vertex, edge, or leg p of Γ we replace the cone $\sigma_p \preceq \Sigma^\dagger$ by the minimal cone $\overline{\sigma}_p \preceq \Sigma$ containing the image of σ_p. Finally, for every oriented edge or leg $\vec{e}$, we replace $m_{\vec{e}}$ by its image

$$\mathrm{p}(m_{\vec{e}}) \in N_{\overline{\sigma}_e}.$$

4 Slope-sensitivity

The present section contains our main polyhedral arguments. We introduce the notion of a slope-sensitive subdivision. The term "slope" here refers to the edge directions in the combinatorial types of naive stable maps. Once the numerical data is fixed, these form a finite set of directions in the cone complex Σ. After subdividing Σ, we show that the combinatorial types of naive maps become tightly constrained. In Section 5 we use this to establish the main theorem.

4.1 Rank 2

The *rank* of a normal crossings pair $(X|D)$ is the dimension of the tropicalisation. The pair $(X|D)$ is of rank 1 if and only if D is smooth, and is of rank 2 if and only if it contains double but not triple intersection points. We first define slope-sensitivity in the rank-2 case.

Consider the space of naive stable maps to $(X|D)$ of class Λ. The set of combinatorial types of such maps is finite. For each combinatorial type, consider the set of slopes $m_{\vec{e}}$ in N_{σ_e} which additionally lie in the positive quadrant $\sigma_e \cap N_{\sigma_e} \subseteq N_{\sigma_e}$. Enumerate these slopes, running over all combinatorial types, as follows:

$$m_1, \ldots, m_h \in \Sigma(\mathbb{N}).$$

A smooth subdivision of Σ which includes amongst its cones the rays generated by $m_1, \ldots, m_h$ is referred to as Λ–*sensitive*. We will also use the term *slope-sensitive* when Λ is unambiguous. Such subdivisions always exist, and form a cofinal subset in the inverse system of smooth subdivisions. We refer to the induced logarithmic modification as a Λ–*sensitive modification*.

The set of slopes $m_{\vec{e}}$ appearing in combinatorial types of class Λ is unchanged by subdividing the target, as this operation does not alter the support of Σ (see also Section 3.5). This ensures that Λ–sensitive subdivisions always exist. Since $\Sigma^{\dagger}$ is required to be smooth, there is typically more than one minimal choice.

Note that while $m_{\vec{e}} \in \Sigma(\mathbb{N})$ does not change, the minimal cone $\sigma_e \preceq \Sigma$ containing the edge e can. This will alter the expression of $m_{\vec{e}}$ in terms of the ray generators of σ_e. This fact is important for motivating the results of Section 4.4.

4.2 Higher rank

Write $D = D_1 + \cdots + D_k$. For a subset of size two

$$I = \{i_1, i_2\} \subseteq \{1, \ldots, k\},$$

we let $D_I = D_{i_1} + D_{i_2} \subseteq D$ be the corresponding divisor. Note that $\Sigma_I = \mathrm{Trop}(X|D_I)$ is either 1–dimensional or 2–dimensional, depending on whether $D_{i_1} \cap D_{i_2}$ is empty or nonempty. There is a logarithmic morphism

$$(X|D) \to (X|D_I)$$

inducing a tropical projection

$$p_I : \Sigma \to \Sigma_I.$$

For each I we choose a Λ–sensitive subdivision $\Sigma_I^{\dagger}$ of the rank-2 tropical target Σ_I. This induces a subdivision

$$p_I^{-1}(\Sigma_I^{\dagger}) \to \Sigma.$$

A smooth subdivision $\Sigma^{\dagger} \to \Sigma$ is then declared to be Λ–*sensitive* if it refines all the subdivisions $p_I^{-1}(\Sigma_I^{\dagger})$ simultaneously. Again, Λ–sensitive subdivisions always exist and are cofinal in the system of smooth subdivisions.

4.3 Iterated blowups

The system of iterated stellar subdivisions of a smooth fan is cofinal in the system of all refinements of that fan; see for instance [Botero 2019, Remark 2.13]. We may therefore assume, by passing to a further refinement, that $\Sigma^\dagger \to \Sigma$ is an iterated stellar subdivision, so that $X^\dagger \to X$ is an iterated blowup in smooth logarithmic strata. Summarising:

Lemma 4.1 *Fix numerical data Λ. Then there exists an iterated strata blowup*

$$(X^\dagger | D^\dagger) \to (X|D)$$

such that $(X^\dagger | D^\dagger)$ is Λ–sensitive with respect to the lifted numerical data Λ. Moreover any smooth logarithmic modification of $(X^\dagger | D^\dagger)$ is also Λ–sensitive.

4.4 Properties: small jumping and slope negativity

Fix a Λ–sensitive modification

$$(X^\dagger | D^\dagger) \to (X|D).$$

By Section 3.5 the numerical data lifts uniquely to data for maps to $(X^\dagger | D^\dagger)$, also denoted by Λ. Fix a naive map $C \to (X^\dagger | D^\dagger)$ with numerical data Λ and consider the associated combinatorial type of tropical stable map to $\Sigma^\dagger$ as in Section 3.3.

Proposition 4.2 *Given an oriented edge $\vec{e}$ of the dual graph, label the generators of the associated cone by*

$$\sigma_e = \mathrm{Cone}(u_1, \ldots, u_c) \preceq \Sigma^\dagger$$

and express the slope along $\vec{e}$ as

$$m_{\vec{e}} = \sum_{i=1}^{c} a_i u_i \in N_{\sigma_e}.$$

Then for all $1 \le i \ne j \le c$, the pair (a_i, a_j) is not contained in $\mathbb{Z}^2_{>0}$ nor in $\mathbb{Z}^2_{<0}$.

Proof After relabelling the coordinates, we suppose for a contradiction that $(a_1, a_2) \in \mathbb{Z}^2_{>0}$. Consider the minimal cone $\overline{\sigma}_e \preceq \Sigma$ containing the image of $\sigma_e \preceq \Sigma^\dagger$ under the subdivision $\Sigma^\dagger \to \Sigma$. Label its generators as

$$\overline{\sigma}_e = \mathrm{Cone}(v_1, \ldots, v_s).$$

We claim that there exist indices $1 \le i \ne j \le s$ such that the image of the composition

$$\mathrm{Cone}(u_1, u_2) \to \sigma_e \to \overline{\sigma}_e \to \mathrm{Cone}(v_i, v_j)$$

is 2–dimensional. Otherwise, u_1 and u_2 must map to the same ray in every $\mathrm{Cone}(v_i, v_j)$ and must therefore map to the same ray in $\mathrm{Cone}(v_1, \ldots, v_s)$. Since $\sigma_e \to \overline{\sigma}_e$ is injective, it follows that u_1 and u_2 are proportional, a contradiction. By relabelling again, assume that the induced map

$$\mathrm{Cone}(u_1, u_2) \to \mathrm{Cone}(v_1, v_2)$$

has 2–dimensional image. Consider the commutative diagram of cones

$$
\begin{array}{ccc}
\sigma_e = \mathrm{Cone}(u_1, \ldots, u_c) & \xrightarrow{\ \mathrm{p}\ } & \mathrm{Cone}(u_1, u_2) \\
\Big\downarrow{\mathrm{f}} & & \Big\downarrow{\mathrm{g}} \\
\bar{\sigma}_e = \mathrm{Cone}(v_1, \ldots, v_s) & \xrightarrow{\ \mathrm{q}\ } & \mathrm{Cone}(v_1, v_2)
\end{array}
$$

The naive map to $(X^{\dagger}|D^{\dagger})$ induces a naive map to $(X|D)$; see [Construction 3.19](#). The combinatorial type of this map has slope along $\vec{e}$ given by $\mathrm{f}(m_{\vec{e}})$ in $N_{\bar{\sigma}_e}$. By commutativity of the square we have

$$
\mathrm{q} \circ \mathrm{f}(m_{\vec{e}}) = \mathrm{g} \circ \mathrm{p}(m_{\vec{e}}).
$$

Now $\mathrm{p}(m_{\vec{e}}) = (a_1, a_2) \in \mathbb{Z}^2_{>0}$ by assumption. Since g has 2–dimensional image, it maps the interior of the source cone to the interior of the target cone, so

$$
\mathrm{g} \circ \mathrm{p}(m_{\vec{e}}) \in \mathbb{Z}^2_{>0}.
$$

Every Λ–sensitive subdivision of $(X|D_{v_1} + D_{v_2})$ therefore includes the ray generated by $\mathrm{g} \circ \mathrm{p}(m_{\vec{e}})$. Consequently every Λ–sensitive subdivision of $(X|D)$ must contain the preimage of this ray in $\sigma_e = \mathrm{Cone}(u_1, \ldots, u_c)$ as a union of cones. But this preimage is a hyperplane in the interior of σ_e, so we have a contradiction. $\square$

Corollary 4.3 (Small jumping) *For every flag consisting of a vertex v incident to an edge e in the graph Γ,*

$$
0 \leq \dim \sigma_e - \dim \sigma_v \leq 1.
$$

Proof We have $\sigma_e \succeq \sigma_v$, and so $\dim \sigma_e - \dim \sigma_v \geq 0$ is immediate. For the other inequality, assume for a contradiction that $\dim \sigma_e \geq \dim \sigma_v + 2$. Order the generators of the cones so that

$$
\sigma_e = \mathrm{Cone}(w_1, \ldots, w_s, u_1, \ldots, u_c) \quad \text{and} \quad \sigma_v = \mathrm{Cone}(u_1, \ldots, u_c),
$$

with $s \geq 2$. Orient e to point away from v and express the slope $m_{\vec{e}} \in N_{\sigma_e}$ in terms of the above basis:

$$
m_{\vec{e}} = \sum_{i=1}^{s} b_i w_i + \sum_{i=1}^{c} a_i u_i.
$$

Using the equality of formal and geometric tangency orders explained in [Lemma 3.8](#) and [Remark 3.9](#), we have $b_i > 0$ for $1 \leq i \leq s$. Since $s \geq 2$ this contradicts [Proposition 4.2](#). $\square$

Corollary 4.4 (Slope negativity) *Consider a flag of a vertex v and incident edge e such that $\dim \sigma_e = \dim \sigma_v + 1$. Order the generators of the cones so that*

$$
\sigma_e = \mathrm{Cone}(u_0, u_1, \ldots, u_c) \quad \text{and} \quad \sigma_v = \mathrm{Cone}(u_1, \ldots, u_c).
$$

Orient e to point away from v and express the slope $m_{\vec{e}}$ in terms of the above basis:

$$
m_{\vec{e}} = a_0 u_0 + \sum_{i=1}^{c} a_i u_i.
$$

Then $a_0 > 0$ and $a_i \leq 0$ for all $1 \leq i \leq c$.

Proof The fact that $a_0 > 0$ follows from Lemma 3.8 and Remark 3.9. The second statement then follows from Proposition 4.2. □

5 Logarithmic–naive comparisons

We establish the main comparison result. To ease notation, we replace the initial simple normal crossings pair with a slope-sensitive modification, which we continue to denote by $(X|D)$.

We will equate the logarithmic and orbifold theories of $(X|D)$. By Theorem 2.9, it suffices to equate the logarithmic and naive theories. Theorem X is thus reduced to the following:

Theorem 5.1 *Let Λ be numerical data for a space of genus-zero logarithmic stable maps to $(X|D)$. Replace $(X|D)$ with a Λ–sensitive modification. Then the natural morphism*

$$\iota \colon \mathrm{Log}_\Lambda(X|D) \to \mathrm{N}_\Lambda(X|D)$$

is virtually birational, ie

$$\iota_\star [\mathrm{Log}_\Lambda(X|D)]^{\mathrm{vir}} = [\mathrm{N}_\Lambda(X|D)]^{\mathrm{vir}}.$$

The proof is given in two parts. In Section 5.1 we establish the theorem on the tropical level. We show that every combinatorial type of naive stable map is smoothable, thus identifying the logarithmic and naive tropical moduli spaces. In Section 5.2 we then use the geometry of the universal target to conclude the algebrogeometric Theorem 5.1 from its tropical counterpart.

5.1 Tropical lifting

Let $\mathcal{A}_X$ be the Artin fan of the pair $(X|D)$ with divisor $\mathcal{D}_X$. As in Section 2.2 let $\mathcal{A}$ be the Artin fan of a smooth pair, with universal divisor $\mathcal{D} = \mathcal{B}\mathbb{G}_{\mathrm{m}}$. For each $1 \le i \le k$ there is a well-defined projection

$$(\mathcal{A}_X|\mathcal{D}_X) \to (\mathcal{A}|\mathcal{D})$$

obtained by equipping the stack $\mathcal{A}_X$ with the Cartier divisor associated to each component of $\mathcal{D}_X$. In Section 2.2 the universal naive space was defined as

$$\mathfrak{N}\mathfrak{Log}(\mathcal{A}^k|\partial\mathcal{A}^k) := \mathfrak{Log}(\mathcal{A}|\mathcal{D})^k \times_{\mathfrak{M}^k} \mathfrak{M}.$$

Consider a naive map $C \to (\mathcal{A}_X|\mathcal{D}_X)$ over a geometric point. By definition, we have logarithmic maps $\mathcal{C}_i \to (\mathcal{A}|\mathcal{D})$ for each $1 \le i \le k$. Each gives a family of tropical maps

$$
\begin{array}{ccc}
\mathsf{C}_i & \longrightarrow & \mathbb{R}_{\ge 0} \\
\downarrow & & \\
\tau_i & &
\end{array}
$$

where τ_i is the corresponding tropical moduli space of edge lengths and target offsets, dual to the base monoid of $\mathcal{C}_i$. Setting $\sigma_\Gamma = \prod_{e \in E(\Gamma)} \mathbb{R}_{\geq 0}$ where Γ is the dual graph of C, there is a universal family of tropical curves (see Sections 1.2 and 1.5)

$$\mathsf{C}_\Gamma \to \sigma_\Gamma.$$

For each $1 \leq i \leq k$ there is a natural morphism $\tau_i \to \sigma_\Gamma$ and C_i is the pullback of C_Γ.

As discussed in Section 3.3, the naive map produces a combinatorial type of tropical stable map to Σ. The following tropical lifting theorem shows that this combinatorial type is smoothable in the sense of Definition 3.11. The theorem only holds after passing to a slope-sensitive subdivision; see the counterexample given in Section 3.4.

Theorem 5.2 *The fibre product of the τ_i over σ_Γ contains a 1–dimensional cone*

$$\rho \subseteq \prod_{i=1}^{k} \tau_i \times_{\sigma_\Gamma^k} \sigma_\Gamma$$

such that the interior of ρ is contained in the interior of the fibre product cone. This satisfies the following properties. Let $\mathsf{C} \to \rho$ denote the pullback of $\mathsf{C}_\Gamma \to \sigma_\Gamma$. Then there is a tropical map

$$
\begin{array}{ccc}
\mathsf{C} & \longrightarrow & \Sigma \\
\downarrow & & \\
\rho & &
\end{array}
$$

whose combinatorial type agrees with the combinatorial type of the given naive stable map. Moreover, for each $1 \leq i \leq k$, the composition $\mathsf{C} \to \Sigma \to \mathbb{R}_{\geq 0}$, where $\Sigma \to \mathbb{R}_{\geq 0}$ is the projection corresponding to the logarithmic morphism $(X|D) \to (X|D_i)$, coincides with the pullback of

$$
\begin{array}{ccc}
\mathsf{C}_i & \longrightarrow & \mathbb{R}_{\geq 0} \\
\downarrow & & \\
\tau_i & &
\end{array}
$$

along the morphism $\rho \to \tau_i$.

Proof It suffices to construct a single tropical map to Σ with all tropical parameters — edge lengths and target positions — nonzero; we then obtain a 1–parameter family by scaling. We will find strictly positive real edge lengths ℓ_e giving rise to a tropical curve C, and a tropical map

$$\mathsf{f} : \mathsf{C} \to \Sigma$$

such that for each vertex v the image $\mathsf{f}(v)$ is contained in $\mathrm{Int}\,\sigma_v$.

We construct f by iterating along Γ. The first step consists of choosing a vertex v of Γ and letting $\mathsf{f}(v)$ be an arbitrary interior point of σ_v. For the induction step, we traverse along an edge e connecting v_1 to v_2. We assume that $\mathsf{f}(v_1)$ in $\mathrm{Int}\,\sigma_{v_1}$ has already been constructed, and show that there exists a choice $\ell_e > 0$ of edge length such that if we define $\mathsf{f}(v_2) := \mathsf{f}(v_1) + \ell_e m_{\vec{e}}$ then $\mathsf{f}(v_2)$ is contained in $\mathrm{Int}\,\sigma_{v_2}$. Since Γ has genus zero, there are no loops to potentially cause inconsistencies.

For the induction step, there are several cases to consider. The most difficult is moving from a larger to a smaller cone, ie

$$\sigma_{v_1} = \sigma_e \succeq \sigma_{v_2}.$$

By Corollary 4.3 we know that σ_{v_2} is a facet of σ_{v_1}. Label the generators as follows:

$$\sigma_{v_1} = \mathrm{Cone}(u_0, u_1, \ldots, u_c) \quad \text{and} \quad \sigma_{v_2} = \mathrm{Cone}(u_1, \ldots, u_c).$$

Then $\mathrm{f}(v_1)$ in $\mathrm{Int}\,\sigma_{v_1}$ can be written as

$$\mathrm{f}(v_1) = \mu_0 u_0 + \sum_{i=1}^{c} \mu_i u_i,$$

with all $\mu_i > 0$. We orient the edge e from v_2 to v_1 and write the slope $m_{\vec{e}}$ as

$$m_{\vec{e}} = a_0 u_0 + \sum_{i=1}^{c} a_i u_i.$$

We have $a_0 > 0$ by Lemma 3.8 and Remark 3.9. By Corollary 4.4 it follows that $a_i \leq 0$ for all $1 \leq i \leq c$. We define

$$\ell_e := \mu_0/a_0 \in \mathbb{R}_{>0}.$$

This gives

$$\mathrm{f}(v_2) = \mathrm{f}(v_1) - \ell_e m_{\vec{e}} = \sum_{i=1}^{c} (\mu_i - \mu_0 a_i/a_0) u_i.$$

The negative sign arises from the choice of orientation for e. Since $\mu_i > 0$ and $\mu_0 a_i/a_0 \leq 0$, we conclude that $\mathrm{f}(v_2)$ lies in $\mathrm{Int}\,\sigma_{v_2}$, as required.

This deals with the difficult case. Now suppose that $\sigma_{v_1} \preceq \sigma_{v_2}$. In this case we are moving from a smaller to a larger cone, and any sufficiently small $\ell_e > 0$ will suffice. Note that this includes the case $\sigma_{v_1} = \sigma_{v_2}$. The final case is where σ_{v_1} and σ_{v_2} are distinct facets of σ_e:

$$\sigma_{v_1} \preceq \sigma_e \succeq \sigma_{v_2}.$$

We subdivide the edge e by introducing an additional bivalent vertex v_0, and assign $\sigma_{v_0} = \sigma_e$. We then apply the previous two cases to construct the map. The vertex v_0 may now be deleted, and we conclude the result. □

5.2 Geometric lifting

We now use the geometry of the universal target to reduce Theorem 5.1 to the tropical lifting theorem (Theorem 5.2) proved above. Consider the following diagram:

$$
\begin{array}{ccccc}
\mathrm{Log}_\Lambda(X|D) & \longrightarrow & \mathrm{N}_\Lambda(X|D) & \longrightarrow & \mathrm{M}_\Lambda(X) \\
\downarrow{\scriptstyle p} & A & \downarrow{\scriptstyle q} & B & \downarrow \\
\mathfrak{Log}(\mathcal{A}_X|\mathcal{D}_X) & \xrightarrow{\psi} & \mathfrak{NLog}(\mathcal{A}^k|\partial\mathcal{A}^k) & \longrightarrow & \mathfrak{M}(\mathcal{A}^k)
\end{array}
$$

We note that $\mathcal{A}_X$ is the open substack of $\mathcal{A}^k$ corresponding to the subcomplex inclusion[2]

$$\Sigma \subseteq \mathbb{R}_{\geq 0}^k.$$

It follows that the outer square $A + B$ is cartesian. On the other hand, we saw in (7) that the square B is cartesian. We conclude that A is cartesian.

The morphisms p and q are equipped with compatible perfect obstruction theories; see Section 2.2. By the Costello–Herr–Wise comparison theorem [Costello 2006; Herr and Wise 2023], Theorem 5.1 follows if we show that ψ is proper and birational [Herr and Wise 2023, Remark 2.6].

The arguments of [Wise 2019] apply equally when the target is a logarithmic algebraic stack. They imply that $\mathfrak{Log}(\mathcal{A}_X|\mathcal{D}_X) \to \mathfrak{M}(\mathcal{A}^k)$ is finite over its image, and hence proper. From Remark 2.5 we see that the morphism $\mathfrak{NLog}(\mathcal{A}^k|\partial\mathcal{A}^k) \to \mathfrak{M}(\mathcal{A}^k)$ is the base change of a proper morphism, and hence a fortiori separated. We conclude that ψ is proper.

To prove Theorem 5.1, it remains to show that ψ is birational. The source of ψ is logarithmically smooth, irreducible, and equal to the closure of the locus where the logarithmic structure is trivial. More explicitly, this is the locus of nondegenerate maps: the source curve is smooth and its generic point maps to the interior of $\mathcal{A}_X$. Moreover, ψ induces an isomorphism of this locus onto the same locus of nondegenerate maps in the codomain. As such, if we prove that ψ is surjective then it follows that ψ is birational. Therefore Theorem 5.1, and hence Theorem X, reduces to the following:

Theorem 5.3 *The following morphism is surjective*:

$$\psi : \mathfrak{Log}(\mathcal{A}_X|\mathcal{D}_X) \to \mathfrak{NLog}(\mathcal{A}^k|\partial\mathcal{A}^k).$$

Proof Choose a geometric point of the codomain. This consists of the following data:

- C a fixed scheme-theoretic source curve,
- $\mathcal{C}_i \to (\operatorname{Spec} \Bbbk, Q_i)$ a logarithmic enhancement of C, one for each $1 \leq i \leq k$, and
- $\Box_i \to \mathbb{R}_{\geq 0}$ a morphism of cone complexes for $1 \leq i \leq k$, where $\Box_i \to \tau_i$ is the family of tropical curves obtained by tropicalising $\mathcal{C}_i \to (\operatorname{Spec} \Bbbk, Q_i)$.

The maps $\Box_i \to \mathbb{R}_{\geq 0}$ correspond to logarithmic maps $\mathcal{C}_i \to (\mathcal{A}|\mathcal{D})$. Forgetting logarithmic structures, we obtain a map $C \to \mathcal{A}^k$ which factors through $\mathcal{A}_X$. We construct a logarithmic enhancement of this by constructing a strongly transverse map to an expansion of $\mathcal{A}_X$, following [Ranganathan 2022]. Strong transversality means that the map is disjoint from codimension-2 strata and only intersects codimension-1 strata at special points, with balanced tangency orders at nodes.

---

[2]Strictly speaking, this is only a subcomplex inclusion if all intersections $\bigcap_{i \in I} D_i$ are connected. In general $\Sigma \to \mathbb{R}_{\geq 0}^k$ is only a finite cover of a subcomplex, but in either case the map $\mathcal{A}_X \to \mathcal{A}^k$ is open, which is all we require. Alternatively, we can force the component intersections to be connected by performing further blowups.

The expansion is produced from the previously constructed tropical lift. In Theorem 5.2 we proved the existence of a 1–dimensional cone

$$\rho \subseteq \prod_{i=1}^{k} \tau_i \times_{\sigma_\Gamma^k} \sigma_\Gamma$$

and compatible tropical map $\sqsubset \to \Sigma$ metrised by ρ. Consider the graph map

$$\sqsubset \to \Sigma \times \rho.$$

Toroidal semistable reduction [Abramovich and Karu 2000] produces a logarithmic modification of $\mathcal{A}_X \times \mathbb{A}^1$ forming a flat and reduced family of expansions of $\mathcal{A}_X$. The central fibre $\mathcal{E}_X$ has irreducible components $\mathcal{E}_X(v)$ indexed by vertices of Γ which meet along divisors $\mathcal{E}_X(e)$ indexed by edges of Γ. There is a natural logarithmic collapsing morphism given as the composite $\mathcal{E}_X \to \mathcal{A}_X \times \mathbb{A}^1 \to \mathcal{A}_X$. We will construct a strongly transverse lift $g \colon C \to \mathcal{E}_X$ of $f \colon C \to \mathcal{A}_X$.

For each vertex v of Γ the restriction of $f \colon C \to \mathcal{A}_X$ to the curve component C_v factors through

$$f_v \colon C_v \to \mathcal{A}_X(v),$$

where $\mathcal{A}_X(v) \subseteq \mathcal{A}_X$ is the closed stratum corresponding to $\sigma_v \preceq \Sigma$. The collapsing morphism restricts to a map of closed strata $\mathcal{E}_X(v) \to \mathcal{A}_X(v)$. Restricting further to the locally closed strata, we obtain a principal torus bundle

$$\mathcal{E}_X^\circ(v) \to \mathcal{A}_X^\circ(v)$$

with structure group $T_{\sigma_v} = N_{\sigma_v} \otimes \mathbb{G}_{\mathrm{m}}$.

Consider the irreducible components of $\mathcal{D}_X$ containing the stratum $\mathcal{A}_X(v)$. Let $L_1, \ldots, L_k$ denote the restrictions to $\mathcal{A}_X^\circ(v)$ of their associated line bundles. Then $\mathcal{E}_X^\circ(v)$ is precisely the principal bundle associated to $\bigoplus_{i=1}^{k} L_i$ [Carocci and Nabijou 2022, Proposition 4.3]. Given a map to $\mathcal{A}_X^\circ(v)$, a lift to $\mathcal{E}_X^\circ(v)$ is given by specifying nonvanishing sections of the pullbacks of $L_1, \ldots, L_k$.

Denote by $x_1, \ldots, x_r$ the special points in C_v. For each L_i, the combinatorial type of naive map determines contact orders $c_{ij} \in \mathbb{Z}$ for $1 \leq j \leq r$. By the balancing condition, these data satisfy

$$\sum_{j=1}^{r} c_{ij} = \deg f_v^\star L_i.$$

Since we work in genus zero, C_v is a marked $\mathbb{P}^1$, and so we may choose an isomorphism

$$(15) \qquad \mathcal{O}_{C_v}\left(\sum_{j=1}^{r} c_{ij} x_j \right) \cong f_v^\star L_i.$$

The standard rational section on the left-hand side produces a rational section of $f_v^\star L_i$. This section is both regular and nonvanishing away from the special points, and hence we obtain a map

$$C_v \setminus \{x_1, \ldots, x_r\} \to \mathcal{E}_X^\circ(v).$$

Passing to the partial compactification $\mathcal{E}_X(v)$ of the codomain, the above map extends to all of C_v and produces a strongly transverse lift

$$
\begin{array}{ccc}
C_v & \dashrightarrow{\ \ g_v\ \ } & \mathcal{E}_X(v) \\
& {\scriptstyle f_v}\searrow & \downarrow \\
& & \mathcal{A}_X(v)
\end{array}
$$

with the desired tangency orders. The set of such lifts forms a torsor under the group T_{σ_v} with moduli the choice of isomorphisms (15). See [Ranganathan 2017, Proposition 3.3.3] for a parallel calculation.

We now argue that the scalar moduli above can be chosen in such a way that the lifts $g_v \colon C_v \to \mathcal{E}_X(v)$ glue to form a globally defined map $g \colon C \to \mathcal{E}_X$. The following argument uses slope-sensitivity in an essential way.

Choose a root vertex of Γ and orient the graph to flow away from the root. We will construct g by iterating along this flow. At each step, we traverse from v_1 to v_2 along an oriented edge $\vec{e}$ of the graph, and assume that we have already constructed a lift

$$
g_1 \colon C_{v_1} \to \mathcal{E}_X(v_1).
$$

Fix an arbitrary lift $g_2 \colon C_{v_2} \to \mathcal{E}_X(v_2)$ with the prescribed contact orders. We claim that there exists a scalar λ in $T_{\sigma_{v_2}}$ such that

$$
(16) \qquad\qquad\qquad\qquad \lambda(g_2(q_e)) = g_1(q_e),
$$

where $q_e \in C$ is the node connecting C_{v_1} and C_{v_2}. Note that $g_2(q_e)$ is an interior point of the join divisor $\mathcal{E}_X(e)$. This interior is a principal torus bundle over the interior of $\mathcal{A}_X(\sigma_e)$ with structure group $(N_{\sigma_e}/\mathbb{Z}m_{\vec{e}}) \otimes \mathbb{G}_{\mathrm{m}}$, where $m_{\vec{e}} \in N_{\sigma_e}$ is the slope along the oriented edge; see eg [Carocci and Nabijou 2024, Theorem 1.8]. To prove the claim it therefore suffices to show that the composition

$$
(17) \qquad\qquad\qquad\qquad N_{\sigma_{v_2}} \to N_{\sigma_e} \to N_{\sigma_e}/\mathbb{Z}m_{\vec{e}}
$$

is surjective. By Corollary 4.3 there are two cases to consider: $\dim \sigma_{v_2} = \dim \sigma_e$ and $\dim \sigma_{v_2} = \dim \sigma_e - 1$. In the first case, (17) is the composition of an isomorphism followed by a surjection and so the claim holds. In the second case, we have a hyperplane inclusion

$$
N_{\sigma_{v_2}} \hookrightarrow N_{\sigma_e}.
$$

Corollary 4.4 implies that $m_{\vec{e}}$ does not belong to the image, and hence the composition (17) is surjective as claimed. We conclude the existence of a scalar λ satisfying (16). Replacing g_2 by $\lambda \circ g_2$ we obtain a strongly transverse lift of f_2 which glues to g_1 along q_e.

By iterating, we construct the desired global lift $g \colon C \to \mathcal{E}_X$. Since this lift is strongly transverse, standard arguments (see eg [Ranganathan 2022, Section 6] or [Abramovich et al. 2020a, Section 4.2]) produce a logarithmic enhancement $\mathcal{C} \to \mathcal{E}_X$. Composing with the collapsing map produces a logarithmic morphism $\mathcal{C} \to (\mathcal{A}_X \mid \mathcal{D}_X)$, and this completes the proof. $\qquad\square$

6 Logarithmic from absolute

The main result of this paper allows for the importation of orbifold techniques into logarithmic Gromov–Witten theory. As a first application, we establish the following conceptual result:

Theorem 6.1 *The genus-zero logarithmic Gromov–Witten theory of $(X|D)$ is determined algorithmically from the genus-zero absolute Gromov–Witten theories of all the strata of $(X|D)$.*

Proof Let $(X|D)$ be a normal crossings pair and fix numerical data Λ for a specific logarithmic Gromov–Witten invariant. Choose a Λ–sensitive modification $(X^\dagger|D^\dagger) \to (X|D)$. We may additionally assume that $D^\dagger \subseteq X^\dagger$ is *simple* normal crossings.

Theorem X equates the logarithmic Gromov–Witten theory of $(X|D)$ to the orbifold Gromov–Witten theory of $(X^\dagger|D^\dagger)$. The multiroot stack associated to $(X^\dagger|D^\dagger)$ may be factored as a sequence of root stacks with respect to smooth divisors. By iteratively applying [Chen et al. 2020, Theorem 1.9], we conclude that the orbifold Gromov–Witten theory of $(X^\dagger|D^\dagger)$ is determined by the absolute Gromov–Witten theories of all the strata of $(X^\dagger|D^\dagger)$.

It remains to reduce these to the absolute Gromov–Witten theories of the strata of $(X|D)$. By induction it is sufficient to consider the case where $X^\dagger \to X$ is a single blowup along a smooth stratum $Z \subseteq X$. Every stratum $V^\dagger \subseteq X^\dagger$ is then either a blowup of, or a projective bundle over, a stratum $V \subseteq X$. In the former case, work of Fan [2021, Theorem B] reduces the theory of $V^\dagger$ to the theories of V and $V \cap Z$, both of which are strata in $(X|D)$.

In the latter case, we observe that the projective bundle $V^\dagger \to V$ is the projectivisation of a direct sum of line bundles. This is because the logarithmic structure provides a natural splitting of the normal bundle of the blowup centre Z. Therefore there is a well-defined big torus acting on the fibres of $V^\dagger \to V$. Localising with respect to this action reduces the theory of $V^\dagger$ to the theory of V (analogously to the proof of Theorem 6.2). Similar arguments have been employed by others; see eg [Maulik and Pandharipande 2006; Brown 2014]. $\square$

If we restrict to insertions coming from X, then a slightly weaker form of the statement can be proved in a direct manner. While weaker, we believe that it transparently illustrates a general theme here, namely that once the problem has been moved to the orbifold side, more techniques become available to manipulate the theory. We include this weaker statement to give a sense for this. See also [Tseng and You 2023b] where similar ideas appear.

As usual fix $(X|D = D_1 + \cdots + D_k)$ a simple normal crossings pair, $\vec{r} = (r_1, \ldots, r_k)$ a tuple of pairwise coprime positive integers, and $\mathcal{X} = X_{D,\vec{r}}$ the associated multiroot stack. For $Y = \mathcal{X}$ or Y a stratum of $(X|D)$, there is a natural morphism $Y \to X$. The *restricted Gromov–Witten theory* of Y is the system of descendant invariants with evaluation classes pulled back from the cohomology of X.

Theorem 6.2 *The genus-zero restricted Gromov–Witten theory of $\mathcal{X}$ is determined algorithmically from the genus-zero restricted absolute Gromov–Witten theories of all the strata of $(X|D)$.*

Proof We recall the projective bundle construction of [Tseng and You 2023b, Section 3.1]. For $1 \leq i \leq k$ let $P_i = \mathbb{P}_X(\mathcal{O}_X(D_i) \oplus \mathcal{O}_X)$. This contains zero and infinity sections X_{i0} and $X_{i\infty}$ satisfying

$$\mathcal{O}_{P_i}(1)|_{X_{i0}} \cong \mathcal{O}_X \quad \text{and} \quad \mathcal{O}_{P_i}(1)|_{X_{i\infty}} \cong \mathcal{O}_X(D_i).$$

There is a tautological section $s_i \in H^0(P_i, \mathcal{O}_{P_i}(1))$ and a natural isomorphism of pairs

$$(V(s_i)|V(s_i) \cap X_{i\infty}) = (X|D_i).$$

From the P_i we construct the following $(\mathbb{P}^1)^k$–bundle over X:

$$P := P_1 \times_X P_2 \times_X \cdots \times_X P_k.$$

We let $p_i \colon P \to P_i$ be the projection and consider the multiroot stack

$$\mathcal{P} := \sqrt[r_1]{(P, p_1^{-1} X_{1\infty})} \times_P \cdots \times_P \sqrt[r_k]{(P, p_k^{-1} X_{k\infty})}.$$

This is a $(\mathbb{P}(1, r_1) \times \cdots \times \mathbb{P}(1, r_k))$–bundle over X. Let $p \colon \mathcal{P} \to P$ be the map to the coarse moduli space. Then the pullbacks of the tautological sections

$$p^\star p_i^\star s_i \in H^0(\mathcal{P}, p^\star p_i^\star \mathcal{O}_{P_i}(1))$$

satisfy

$$V(p^\star p_1^\star s_1, \ldots, p^\star p_k^\star s_k) \cong \mathcal{X}.$$

Since each bundle $p^\star p_i^\star \mathcal{O}_{P_i}(1)$ is positive and pulled back from the coarse moduli space, we may apply the quantum Lefschetz theorem for orbifolds [Coates et al. 2012]. This reduces the restricted genus-zero theory of $\mathcal{X}$ to the genus-zero theory of $\mathcal{P}$ capped with the class

$$(18) \qquad \prod_{i=1}^{k} e(\boldsymbol{R}^\bullet \pi_\star f^\star p^\star p_i^\star \mathcal{O}_{P_i}(1)).$$

We compute this by localising with respect to the big torus acting on the fibres of $\mathcal{P} \to X$. Each fixed locus $\mathcal{Z}_I \subseteq \mathcal{P}$ in the target is indexed by a subset $I \subseteq \{1, \ldots, k\}$ and is isomorphic to a gerbe over X banded by $\prod_{i \in I} \mu_{r_i}$:

$$\mathcal{Z}_I = \prod_{i \in I} \sqrt[r_i]{(X, \mathcal{O}_X(D_i))} \times_{X^I} X.$$

Now consider a fixed locus in the moduli space $\mathrm{Orb}_\Lambda(\mathcal{P})$. Up to a finite cover, this is a fibre product of spaces of maps to the gerbes $\mathcal{Z}_I$. The fibring is with respect to the evaluation maps to the coarse moduli space X.

Since $\mathcal{O}_{P_i}(1)|_{X_{i0}} \cong \mathcal{O}_X$, it follows that for $i \notin I$ the class $e(\boldsymbol{R}^\bullet \pi_\star f^\star p^\star p_i^\star \mathcal{O}_{P_i}(1))$ is pure weight. On the other hand, for $i \in I$ we have

$$p^\star p_i^\star \mathcal{O}_{Y_i}(1) \cong s^\star \mathcal{O}_X(D_i),$$

where $s\colon \mathcal{Z}_I \to X$ is the rigidification. Let $u\colon \operatorname{Orb}_\Lambda(\mathcal{Z}_I) \to \mathrm{M}_\Lambda(X)$ denote the induced morphism on spaces of stable maps. By the standard comparison of universal curves (see eg Section 2.3) we obtain

$$e(\boldsymbol{R}^\bullet \pi_\star f^\star p^\star p_i^\star \mathcal{O}_{P_i}(1)) = u^\star e(\boldsymbol{R}^\bullet \pi_\star f^\star \mathcal{O}_X(D_i)).$$

Finally, the normal bundle term expands as a power series in equivariant weights and (coarse) psi classes [Liu 2013]. For a discussion of the comparison between coarse and gerby psi classes see [Abramovich et al. 2008, Section 8.3].

In summary, the contribution of the fixed locus is a (product of) integrals over $\operatorname{Orb}_\Lambda(\mathcal{Z}_I)$ with insertions pulled back along u. By the virtual pushforward theorem for root gerbes [Andreini et al. 2015, Theorem 4.3] this can be expressed as an integral over $\mathrm{M}_\Lambda(X)$. The integrand is a product of equivariant weights, evaluation classes, psi classes, and an Euler class term. But the Euler class term is simply

$$\prod_{i \in I} e(\boldsymbol{R}^\bullet \pi_\star f^\star \mathcal{O}_X(D_i)) = \iota_\star [\mathrm{M}_\Lambda(\cap_{i \in I} D_i)]^{\mathrm{vir}},$$

by the absolute quantum Lefschetz theorem. We have therefore successfully reduced the restricted theory of $\mathcal{X}$ to the restricted theories of the strata $\bigcap_{i \in I} D_i$ of $(X \,|\, D)$. $\square$

References

[Abramovich 2008] **D Abramovich**, *Lectures on Gromov–Witten invariants of orbifolds*, from "Enumerative invariants in algebraic geometry and string theory" (K Behrend, M Manetti, editors), Lecture Notes in Math. 1947, Springer (2008) 1–48 MR Zbl

[Abramovich and Chen 2014] **D Abramovich**, **Q Chen**, *Stable logarithmic maps to Deligne–Faltings pairs, II*, Asian J. Math. 18 (2014) 465–488 MR Zbl

[Abramovich and Karu 2000] **D Abramovich**, **K Karu**, *Weak semistable reduction in characteristic 0*, Invent. Math. 139 (2000) 241–273 MR Zbl

[Abramovich and Vistoli 2002] **D Abramovich**, **A Vistoli**, *Compactifying the space of stable maps*, J. Amer. Math. Soc. 15 (2002) 27–75 MR Zbl

[Abramovich and Wise 2018] **D Abramovich**, **J Wise**, *Birational invariance in logarithmic Gromov–Witten theory*, Compos. Math. 154 (2018) 595–620 MR Zbl

[Abramovich et al. 2008] **D Abramovich**, **T Graber**, **A Vistoli**, *Gromov–Witten theory of Deligne–Mumford stacks*, Amer. J. Math. 130 (2008) 1337–1398 MR Zbl

[Abramovich et al. 2011] **D Abramovich**, **M Olsson**, **A Vistoli**, *Twisted stable maps to tame Artin stacks*, J. Algebraic Geom. 20 (2011) 399–477 MR Zbl

[Abramovich et al. 2014] **D Abramovich**, **S Marcus**, **J Wise**, *Comparison theorems for Gromov–Witten invariants of smooth pairs and of degenerations*, Ann. Inst. Fourier (Grenoble) 64 (2014) 1611–1667 MR Zbl

[Abramovich et al. 2016] **D Abramovich**, **Q Chen**, **S Marcus**, **M Ulirsch**, **J Wise**, *Skeletons and fans of logarithmic structures*, from "Nonarchimedean and tropical geometry" (M Baker, S Payne, editors), Springer (2016) 287–336 MR Zbl

[Abramovich et al. 2017]   **D Abramovich**, **C Cadman**, **J Wise**, *Relative and orbifold Gromov–Witten invariants*, Algebr. Geom. 4 (2017) 472–500   MR Zbl

[Abramovich et al. 2020a]   **D Abramovich**, **Q Chen**, **M Gross**, **B Siebert**, *Decomposition of degenerate Gromov–Witten invariants*, Compos. Math. 156 (2020) 2020–2075   MR Zbl

[Abramovich et al. 2020b]   **D Abramovich**, **Q Chen**, **M Gross**, **B Siebert**, *Punctured logarithmic maps*, preprint (2020)   arXiv 2009.07720

[Andreini et al. 2015]   **E Andreini**, **Y Jiang**, **H-H Tseng**, *Gromov–Witten theory of root gerbes, I: Structure of genus 0 moduli spaces*, J. Differential Geom. 99 (2015) 1–45   MR Zbl

[Andreini et al. 2016]   **E Andreini**, **Y Jiang**, **H-H Tseng**, *Gromov–Witten theory of product stacks*, Comm. Anal. Geom. 24 (2016) 223–277   MR Zbl

[Argüz and Gross 2022]   **H Argüz**, **M Gross**, *The higher-dimensional tropical vertex*, Geom. Topol. 26 (2022) 2135–2235   MR Zbl

[Bae and Schmitt 2022]   **Y Bae**, **J Schmitt**, *Chow rings of stacks of prestable curves, I*, Forum Math. Sigma 10 (2022) art. id. e28   MR Zbl

[Battistella and Nabijou 2021]   **L Battistella**, **N Nabijou**, *Relative quasimaps and mirror formulae*, Int. Math. Res. Not. 2021 (2021) 7885–7931   MR Zbl

[Battistella et al. 2023]   **L Battistella**, **N Nabijou**, **H-H Tseng**, **F You**, *The local-orbifold correspondence for simple normal crossing pairs*, J. Inst. Math. Jussieu 22 (2023) 2515–2531   MR Zbl

[Behrend 1999]   **K Behrend**, *The product formula for Gromov–Witten invariants*, J. Algebraic Geom. 8 (1999) 529–541   MR Zbl

[Botero 2019]   **A M Botero**, *Intersection theory of toric b–divisors in toric varieties*, J. Algebraic Geom. 28 (2019) 291–338   MR Zbl

[Brown 2014]   **J Brown**, *Gromov–Witten invariants of toric fibrations*, Int. Math. Res. Not. 2014 (2014) 5437–5482   MR Zbl

[Cadman 2007]   **C Cadman**, *Using stacks to impose tangency conditions on curves*, Amer. J. Math. 129 (2007) 405–427   MR Zbl

[Cadman and Chen 2008]   **C Cadman**, **L Chen**, *Enumeration of rational plane curves tangent to a smooth cubic*, Adv. Math. 219 (2008) 316–343   MR Zbl

[Carocci and Nabijou 2022]   **F Carocci**, **N Nabijou**, *Tropical expansions and toric variety bundles*, preprint (2022)   arXiv 2207.12541

[Carocci and Nabijou 2024]   **F Carocci**, **N Nabijou**, *Rubber tori in the boundary of expanded stable maps*, J. Lond. Math. Soc. 109 (2024) art. id. e12874

[Cavalieri et al. 2020]   **R Cavalieri**, **M Chan**, **M Ulirsch**, **J Wise**, *A moduli stack of tropical curves*, Forum Math. Sigma 8 (2020) art. id. e23   MR Zbl

[Chen 2014]   **Q Chen**, *Stable logarithmic maps to Deligne–Faltings pairs, I*, Ann. of Math. 180 (2014) 455–521   MR Zbl

[Chen et al. 2020]   **B Chen**, **C-Y Du**, **R Wang**, *Orbifold Gromov–Witten theory of weighted blowups*, Sci. China Math. 63 (2020) 2475–2522   MR Zbl

[Cheong et al. 2015]   **D Cheong**, **I Ciocan-Fontanine**, **B Kim**, *Orbifold quasimap theory*, Math. Ann. 363 (2015) 777–816   MR Zbl

[Coates et al. 2012]   **T Coates**, **A Gholampour**, **H Iritani**, **Y Jiang**, **P Johnson**, **C Manolache**, *The quantum Lefschetz hyperplane principle can fail for positive orbifold hypersurfaces*, Math. Res. Lett. 19 (2012) 997–1005   MR Zbl

[Costello 2006]   **K Costello**, *Higher genus Gromov–Witten invariants as genus zero invariants of symmetric products*, Ann. of Math. 164 (2006) 561–601   MR Zbl

[Cox et al. 2011]   **D A Cox**, **J B Little**, **H K Schenck**, *Toric varieties*, Grad. Stud. Math. 124, Amer. Math. Soc., Providence, RI (2011)   MR Zbl

[Fan 2021]   **H Fan**, *Chern classes and Gromov–Witten theory of projective bundles*, Amer. J. Math. 143 (2021) 811–832   MR Zbl

[Fan et al. 2020]   **H Fan**, **L Wu**, **F You**, *Structures in genus-zero relative Gromov–Witten theory*, J. Topol. 13 (2020) 269–307   MR Zbl

[Fan et al. 2021]   **H Fan**, **L Wu**, **F You**, *Higher genus relative Gromov–Witten theory and double ramification cycles*, J. Lond. Math. Soc. 103 (2021) 1547–1576   MR Zbl

[Gathmann 2002]   **A Gathmann**, *Absolute and relative Gromov–Witten invariants of very ample hypersurfaces*, Duke Math. J. 115 (2002) 171–203   MR Zbl

[Gathmann 2003]   **A Gathmann**, *Gromov–Witten invariants of hypersurfaces*, Habilitationsschrift, Universität Kaiserslautern (2003)   Available at http://www.mathematik.uni-kl.de/~gathmann/pub/habil.pdf

[Gross and Siebert 2013]   **M Gross**, **B Siebert**, *Logarithmic Gromov–Witten invariants*, J. Amer. Math. Soc. 26 (2013) 451–510   MR Zbl

[Gross and Siebert 2019]   **M Gross**, **B Siebert**, *Intrinsic mirror symmetry*, preprint (2019)   arXiv 1909.07649

[Gross et al. 2010]   **M Gross**, **R Pandharipande**, **B Siebert**, *The tropical vertex*, Duke Math. J. 153 (2010) 297–362   MR Zbl

[Herr and Wise 2023]   **L Herr**, **J Wise**, *Costello's pushforward formula: errata and generalization*, Manuscripta Math. 171 (2023) 621–642   MR Zbl

[Holmes and Schwarz 2022]   **D Holmes**, **R Schwarz**, *Logarithmic intersections of double ramification cycles*, Algebr. Geom. 9 (2022) 574–605   MR Zbl

[Holmes et al. 2019]   **D Holmes**, **A Pixton**, **J Schmitt**, *Multiplicativity of the double ramification cycle*, Doc. Math. 24 (2019) 545–562   MR Zbl

[Janda et al. 2017]   **F Janda**, **R Pandharipande**, **A Pixton**, **D Zvonkine**, *Double ramification cycles on the moduli spaces of curves*, Publ. Math. Inst. Hautes Études Sci. 125 (2017) 221–266   MR Zbl

[Janda et al. 2020]   **F Janda**, **R Pandharipande**, **A Pixton**, **D Zvonkine**, *Double ramification cycles with target varieties*, J. Topol. 13 (2020) 1725–1766   MR Zbl

[Kato 1994]   **K Kato**, *Toric singularities*, Amer. J. Math. 116 (1994) 1073–1099   MR Zbl

[Kato 2000]   **F Kato**, *Log smooth deformation and moduli of log smooth curves*, Int. J. Math. 11 (2000) 215–232   MR Zbl

[Kempf et al. 1973]   **G Kempf**, **F F Knudsen**, **D Mumford**, **B Saint-Donat**, *Toroidal embeddings, I*, Lecture Notes in Math. 339, Springer (1973)   MR Zbl

[Kennedy-Hunt et al. 2022]   **P Kennedy-Hunt**, **N Nabijou**, **Q Shafi**, **W Zheng**, *Divisors and curves on logarithmic mapping spaces*, preprint (2022)   arXiv 2209.00630

[Kim 2010]   **B Kim**, *Logarithmic stable maps*, from "New developments in algebraic geometry, integrable systems and mirror symmetry" (M-H Saito, S Hosono, K Yoshioka, editors), Adv. Stud. Pure Math. 59, Math. Soc. Japan, Tokyo (2010) 167–200   MR Zbl

[Kresch 1999]   **A Kresch**, *Cycle groups for Artin stacks*, Invent. Math. 138 (1999) 495–536   MR Zbl

[Li 2001]   **J Li**, *Stable morphisms to singular schemes and relative stable morphisms*, J. Differential Geom. 57 (2001) 509–578   MR Zbl

[Li 2002]   **J Li**, *A degeneration formula of GW-invariants*, J. Differential Geom. 60 (2002) 199–293   MR Zbl

[Liu 2013]   **C-C M Liu**, *Localization in Gromov–Witten theory and orbifold Gromov–Witten theory*, from "Handbook of moduli, II" (G Farkas, I Morrison, editors), Adv. Lect. Math. 25, International, Somerville, MA (2013) 353–425   MR Zbl

[Maclagan and Sturmfels 2015]   **D Maclagan**, **B Sturmfels**, *Introduction to tropical geometry*, Grad. Stud. Math. 161, Amer. Math. Soc., Providence, RI (2015)   MR Zbl

[Mandel and Ruddat 2020]   **T Mandel**, **H Ruddat**, *Descendant log Gromov–Witten invariants for toric varieties and tropical curves*, Trans. Amer. Math. Soc. 373 (2020) 1109–1152   MR Zbl

[Manolache 2012]   **C Manolache**, *Virtual pull-backs*, J. Algebraic Geom. 21 (2012) 201–245   MR Zbl

[Maulik and Pandharipande 2006]   **D Maulik**, **R Pandharipande**, *A topological view of Gromov–Witten theory*, Topology 45 (2006) 887–918   MR Zbl

[Molcho 2021]   **S Molcho**, *Universal stacky semistable reduction*, Israel J. Math. 242 (2021) 55–82   MR Zbl

[Molcho and Ranganathan 2024]   **S Molcho**, **D Ranganathan**, *A case study of intersections on blowups of the moduli of curves*, Algebra Number Theory 18 (2024) 1767–1816   MR Zbl

[Nabijou 2018]   **N Nabijou**, *Recursion formulae in logarithmic Gromov–Witten theory and quasimap theory*, PhD thesis, Imperial College London (2018)   Available at https://core.ac.uk/download/pdf/388546480.pdf

[Nabijou and Ranganathan 2022]   **N Nabijou**, **D Ranganathan**, *Gromov–Witten theory with maximal contacts*, Forum Math. Sigma 10 (2022) art. id. e5   MR Zbl

[Nishinou and Siebert 2006]   **T Nishinou**, **B Siebert**, *Toric degenerations of toric varieties and tropical curves*, Duke Math. J. 135 (2006) 1–51   MR Zbl

[Olsson 2007]   **M C Olsson**, *(Log) twisted curves*, Compos. Math. 143 (2007) 476–494   MR Zbl

[Ranganathan 2017]   **D Ranganathan**, *Skeletons of stable maps, I: Rational curves in toric varieties*, J. Lond. Math. Soc. 95 (2017) 804–832   MR Zbl

[Ranganathan 2019]   **D Ranganathan**, *A note on the cycle of curves in a product of pairs*, preprint (2019)   arXiv 1910.00239

[Ranganathan 2022]   **D Ranganathan**, *Logarithmic Gromov–Witten theory with expansions*, Algebr. Geom. 9 (2022) 714–761   MR Zbl

[Ranganathan et al. 2019]   **D Ranganathan**, **K Santos-Parker**, **J Wise**, *Moduli of stable maps in genus one and logarithmic geometry, I*, Geom. Topol. 23 (2019) 3315–3366   MR Zbl

[Shafi 2024]   **Q Shafi**, *Logarithmic quasimaps*, Adv. Math. 438 (2024) art. id. 109469   MR Zbl

[Tseng and You 2016]   **H-H Tseng**, **F You**, *On orbifold Gromov–Witten theory in codimension one*, J. Pure Appl. Algebra 220 (2016) 3567–3571   MR Zbl

[Tseng and You 2020]   **H-H Tseng**, **F You**, *Higher genus relative and orbifold Gromov–Witten invariants*, Geom. Topol. 24 (2020) 2749–2779   MR Zbl

[Tseng and You 2023a]   **H-H Tseng**, **F You**, *A Gromov–Witten theory for simple normal-crossing pairs without log geometry*, Comm. Math. Phys. 401 (2023) 803–839   MR Zbl

[Tseng and You 2023b]   **H-H Tseng**, **F You**, *A mirror theorem for multi-root stacks and applications*, Selecta Math. 29 (2023) art. id. 6   MR Zbl

[Wise 2019]   **J Wise**, *Uniqueness of minimal morphisms of logarithmic schemes*, Algebr. Geom. 6 (2019) 50–63   MR Zbl

[Wu 2024]   **Y Wu**, *Splitting of Gromov–Witten invariants with toric gluing strata*, J. Algebraic Geom. 33 (2024) 213–263   Zbl

[You 2022]   **F You**, *Relative quantum cohomology under birational transformations*, preprint (2022)   arXiv 2204.00509

Goethe Universität Frankfurt
Frankfurt, Germany

Current address: *Università di Bologna*
Bologna, Italy

School of Mathematical Sciences, Queen Mary University of London
London, United Kingdom

Department of Pure Mathematics and Mathematical Statistics, University of Cambridge
Cambridge, United Kingdom

luca.battistella2@unibo.it,   n.nabijou@qmul.ac.uk,   dr508@cam.ac.uk

GEOMETRY & TOPOLOGY

Geometry & Topology