

EXTENDING THE DOUBLE RAMIFICATION CYCLE BY RESOLVING THE ABEL-JACOBI MAP

DAVID HOLMES

*Universiteit Leiden Mathematisch Instituut, Mathematics, Leiden, 2300 RA,
The Netherlands*

(Received 21 June 2018; revised 8 April 2019; accepted 12 April 2019)

Abstract Over the moduli space of smooth curves, the double ramification cycle can be defined by pulling back the unit section of the universal jacobian along the Abel–Jacobi map. This breaks down over the boundary since the Abel–Jacobi map fails to extend. We construct a ‘universal’ resolution of the Abel–Jacobi map, and thereby extend the double ramification cycle to the whole of the moduli of stable curves. In the non-twisted case, we show that our extension coincides with the cycle constructed by Li, Graber, Vakil via a virtual fundamental class on a space of rubber maps.

Keywords: curves; jacobians; cycles

2010 *Mathematics subject classification:* Primary 14H10; 14H40

Contents

1	Introduction	2
2	Notation and setup	7
3	Construction of \mathcal{M}^\diamond	9
4	Universal property of \mathcal{M}^\diamond	15
5	Properness of DRL_\diamond	18
	5.1 In characteristic zero	18
	5.2 The universal line bundle	19
	5.3 The proof of properness in general	20
6	Proof of Theorems 1.1 and 1.2	22
7	Proof of Theorem 1.3	25
	References	28

1. Introduction

Fix integers $g, n \geq 0$ satisfying $2g - 2 + n > 0$, and integers $a_1 \dots a_n$ and k such that $\sum_i a_i = k(2g - 2)$. Over the moduli stack $\mathcal{M}_{g,n}$ we have the universal curve $C_{g,n}$ with tautological sections x_1, \dots, x_n . Write J for the universal jacobian (an abelian scheme over $\mathcal{M}_{g,n}$) and σ for the section of J given by the line bundle $\omega^{\otimes k}(-\sum_i a_i x_i)$ on $C_{g,n}$. The pullback of the unit section along σ defines a codimension- g cycle class on $\mathcal{M}_{g,n}$, the *double ramification cycle* (DRC). The problem of producing a ‘reasonable’ extension of the DRC to the Deligne–Mumford–Knudsen compactification $\overline{\mathcal{M}}_{g,n}$, and computing the class of the resulting cycle in the tautological ring, was proposed by Eliashberg.

In the case $k = 0$, an extension of the class to the whole of $\overline{\mathcal{M}}_{g,n}$ was constructed by Li [25], [26], Graber and Vakil [12]. This class was computed in the compact-type case by Hain [15], and this was extended to tree-like curves with one loop by Grushevsky and Zakharov [13]. More recently, Janda, Pandharipande, Pixton and Zvonkine [20] computed this class on the whole of $\overline{\mathcal{M}}_{g,n}$, proving a conjecture of Pixton.

A construction for arbitrary k was proposed by Guéré [14], but the situation here is more complicated. Pixton’s conjecture makes sense for all k , but is purely combinatorial in origin. A more geometric conjecture is given by Janda, Pandharipande, Pixton, and Zvonkine in the appendix of [6] for $k = 1$, generalised by Schmitt [32] to all $k \geq 1$; their formulae are moreover conjectured to coincide with Pixton’s. However, it is at present unclear whether Guéré’s construction is compatible with these conjectures.

In this paper we give a simple construction of an extension of the DRC for arbitrary k , and for $k = 0$ we verify that it agrees with the construction of Li, Graber and Vakil. Given the situation described above, it seems very interesting to know whether it coincides with the construction of Guéré when k is arbitrary; this is the subject of current work in progress. Our construction is valid in arbitrary characteristic, and produces a class in $\text{CH}(\overline{\mathcal{M}}_{g,n})$; we do not need to work with rational coefficients.

Kass and Pagani [21] have recently constructed large numbers of extensions of the DRC (for every $k \in \mathbb{Z}$). In [18] we show that certain of their classes coincide with the one constructed in this article.

Statement of main results

We write $J/\overline{\mathcal{M}}_{g,n}$ for the unique semiabelian extension of the universal jacobian (also denoted above by J), and ω for the relative dualising sheaf of the universal stable curve over $\overline{\mathcal{M}}_{g,n}$. The section $\sigma = \omega^{\otimes k}(-\sum_i a_i x_i)$ does not in general extend over the whole of $\overline{\mathcal{M}}_{g,n}$. This can be partially resolved by blowing up $\overline{\mathcal{M}}_{g,n}$. Let $f: X \rightarrow \overline{\mathcal{M}}_{g,n}$ be a proper birational map from a normal stack (a ‘normal modification’). The section σ is then defined on some dense open of X . We write \mathring{X} for the largest open of X on which this rational map can be extended to a morphism, and $\sigma_X: \mathring{X} \rightarrow J$ for the extension.

We define the *double ramification locus* $\text{DRL}_X \rightarrow \mathring{X}$ to be the schematic pullback of the unit section of J along σ_X , and the *double ramification cycle* DRC_X to be the cycle-theoretic pullback, *as a cycle supported on* DRL_X (actually we pull back the unit section in $J \times_{\overline{\mathcal{M}}_{g,n}} \mathring{X}$ along the induced map $\mathring{X} \rightarrow J \times_{\overline{\mathcal{M}}_{g,n}} \mathring{X}$, since the latter is a regular

closed immersion). Now the map $\mathring{X} \rightarrow \overline{\mathcal{M}}_{g,n}$ is rarely proper, but the map $\text{DRL}_X \rightarrow \overline{\mathcal{M}}_{g,n}$ is quite often proper. More precisely, we have:

Theorem 1.1 (Theorem 6.5). In the directed system of all normal modifications of $\overline{\mathcal{M}}_{g,n}$, those X such that $\text{DRL}_X \rightarrow \overline{\mathcal{M}}_{g,n}$ is proper form a cofinal system.

Now DRC_X is supported on DRL_X , so when the map $\text{DRL}_X \rightarrow \overline{\mathcal{M}}_{g,n}$ is proper, we can take the pushforward of DRC_X to $\overline{\mathcal{M}}_{g,n}$. Writing $\pi_{X*} \text{DRC}_X$ for the resulting cycle on $\overline{\mathcal{M}}_{g,n}$, we might hope that these cycles ‘converge’ in some way as we move up in the tower of modifications. In fact, this is true in a very strong sense:

Theorem 1.2 (Theorem 6.7). The net $\pi_{X*} \text{DRC}_X$ is eventually constant in the Chow ring $\text{CH}(\overline{\mathcal{M}}_{g,n})$. We denote the limit by $\overline{\text{DRC}}$.

In other words, there exists a normal modification $X \rightarrow \overline{\mathcal{M}}_{g,n}$ such that, for every normal modification X' dominating X , we have $\pi_{X*} \text{DRC}_X = \pi_{X'*} \text{DRC}_{X'}$. Since every two normal modifications can be dominated by a third, this implies the existence of a uniquely determined ‘limit’ of this collection of $\pi_{X'*} \text{DRC}_{X'}$, and we denote this limit by $\overline{\text{DRC}}$.

Theorem 1.3 (Theorem 7.4). Suppose that $k = 0$. The class $\overline{\text{DRC}}$ coincides with the extension of the double ramification cycle defined by Li, Graber and Vakil [12] (and computed by [20]).

Strategy of proof

The proofs of Theorems 1.1 and 1.2 proceed by constructing a ‘universal’ stack over $\overline{\mathcal{M}}_{g,n}$ making the map σ extend. More formally, we call a morphism $t: T \rightarrow \overline{\mathcal{M}}_{g,n}$ σ -*extending* if the section σ can be extended to J after pulling back to T . For this to make sense we need the pullback of $\mathcal{M}_{g,n}$ to be dense in T , and we also require T to be normal and to admit a resolution of singularities (the latter being automatic in characteristic zero by [16]).

Our *universal resolution of the Abel–Jacobi map* will then be a terminal object in the category of σ -extending morphisms. In Corollary 4.6 we prove existence; we denote this object by $\mathcal{M}_{g,n}^\diamond$. The section σ extends over $\mathcal{M}_{g,n}^\diamond$ by definition, and moreover $\mathcal{M}_{g,n}^\diamond$ is the ‘universal’ stack over $\overline{\mathcal{M}}_{g,n}$ over which σ extends. Note that $\mathcal{M}_{g,n}^\diamond$ depends non-trivially on the ramification data a_1, \dots, a_n .

The proof of existence of $\mathcal{M}_{g,n}^\diamond$ is constructive, and equips it with a logarithmic structure making it logarithmically étale over $\overline{\mathcal{M}}_{g,n}$. The construction is given locally using toric geometry; we write down explicit combinatorial recipes for the fans of toric varieties, and glue them to build $\mathcal{M}_{g,n}^\diamond$. The combinatorial recipe is unavoidably rather complicated, but is amenable to implementation on a computer; in this way we produced Figure 1 showing one of these fans.




Figure 1. Slice of F_Γ .

In §5.2 we also give an explicit description of the universal line bundle on the universal curve over $\mathcal{M}_{g,n}^\diamond$ in terms of this toric data, and in §5.3 we use this to prove the key properness result needed for Theorem 1.1 (though in characteristic zero there is a simpler proof, see §5.1).

The morphism $\mathcal{M}_{g,n}^\diamond \rightarrow \overline{\mathcal{M}}_{g,n}$ may be viewed as a universal resolution of the indeterminacies of the Abel–Jacobi map. This solves a problem proposed by Grushevsky and Zakharov [13, Remark 6.3]. Together with our other results, it seems also to solve a problem of Cavalieri, Marcus and Wise proposed in [2, §1.4].

To conclude the proofs of the main results, we will choose a suitable compactification $\mathcal{M}_{g,n}^\blacklozenge$ of $\mathcal{M}_{g,n}^\diamond$. Now if $X \rightarrow \overline{\mathcal{M}}_{g,n}$ is any normal modification which factors via a map $f : X \rightarrow \mathcal{M}_{g,n}^\blacklozenge$, we will verify that

$$\mathring{X} = f^{-1}\mathcal{M}_{g,n}^\diamond, \quad \text{DRL}_X = f^* \text{DRL}_{\mathcal{M}_{g,n}^\blacklozenge}, \quad \text{and} \quad f_* \text{DRC}_X = \text{DRC}_{\mathcal{M}_{g,n}^\blacklozenge}.$$

(Here \mathring{X} still denotes the largest open of X on which the rational map σ can be extended to a morphism.) We will establish that DRL_\diamond is proper over $\overline{\mathcal{M}}_{g,n}$, whereupon the analogous properness result will hold for any normal modification which factors via $\mathcal{M}_{g,n}^\blacklozenge$, establishing Theorem 1.1. Theorem 1.2 will then follow fairly formally. Finally, when $k = 0$ we use the deformation-theoretic tools of Marcus, Wise and Cavalieri to establish Theorem 1.3.

A simple formula for the double ramification cycle. Suppose we work over a field of characteristic zero, or perhaps more generally over a ring R such that resolution of singularities is known for all schemes of finite type over R . Then there is a simpler way to prove the *existence* of the universal σ -extending morphism: the normalisation of the closure of the schematic image of σ in J satisfies the universal property (details are given in §5.1). This does not tell us a huge amount since we have no explicit description of this closure; for example, it is *a priori* far from clear that it admits a log structure making it

log étale over $\overline{\mathcal{M}}$. However, it does allow us to simplify the proof of Theorem 1.1; details are given in § 5.1.

This allows us to give a very simple formula for the extension of the double ramification cycle. We define DR_{naive} to be the cycle on $\overline{\mathcal{M}}_{g,n}$ obtained by pulling back the schematic image of σ along the unit section. Still in characteristic zero, a small argument with the projection formula (of which details can be found in [18]) shows that the resulting cycle is equal to the one constructed in this paper, and hence also to that of Li, Graber and Vakil.

Note that we do need resolution of singularities since, in order to verify that the normalisation of the closure of σ is σ -extending, we need in particular that it admits a (local) resolution of singularities, which is in general far from clear. Of course, one could drop from the definition of σ -extending that T admit a local resolution of singularities, but then our methods break down (in particular we are no longer able to prove the critical Lemma 4.3), so we can no longer give an explicit description of the universal object, or equip it with a natural logarithmic structure etc.

Conjectural relationship to a cycle of Pixton

Given the data g, n, k and \underline{a} , Pixton introduced a cycle $P_g^{g,k}(A)$ in the tautological ring of $\overline{\mathcal{M}}_{g,n}$, given in terms of decorated graphs – here $A = (a_1 + k, \dots, a_n + k)$; the details of the construction can be found in [20, § 1.1]. The main result of [20] shows that, when $k = 0$, there is an equality of cycles

$$\overline{\text{DRC}}_{\underline{a}} = 2^{-g} P_g^{g,0}(A)$$

(here we write $\overline{\text{DRC}}_{\underline{a}}$ to make the dependence on $\underline{a} = (a_1, \dots, a_n)$ explicit). Now that we have a construction of $\overline{\text{DRC}}_{\underline{a}}$ valid for all k , it seems natural to propose the following.

Conjecture 1.4. For all k , there is an equality of cycles

$$\overline{\text{DRC}}_{\underline{a}} = 2^{-g} P_g^{g,k}(A)$$

in $\text{CH}^g(\overline{\mathcal{M}}_{g,n})$.

Some evidence for this conjecture is given in the following section.

Multiplicativity of the double ramification cycle

In [19] we use the results of this paper to construct an extension of the double ramification cycle to the small b -Chow ring of $\overline{\mathcal{M}}_{g,n}$ – this is the colimit of the Chow rings of the smooth blowups of $\overline{\mathcal{M}}_{g,n}$, with transition maps given by pulling back cycles. Note that the ‘asymptotic’ approach we adopt here is essential for this construction. Given vectors $\underline{a}, \underline{b}$ of ramification data, we will show that the basic multiplicativity relation

$$\text{DRC}_{\underline{a}} \cdot \text{DRC}_{\underline{b}} = \text{DRC}_{\underline{a}} \cdot \text{DRC}_{\underline{a+b}} \tag{1}$$

holds in the small b -Chow ring of $\overline{\mathcal{M}}_{g,n}$, but fails in the Chow ring of $\overline{\mathcal{M}}_{g,n}$.

One consequence of this is that, if Conjecture 1.4 is true, the relation (1) should also hold for Pixton’s cycles on the locus of compact-type curves. This relation can

be independently checked using known relations in the tautological ring (see [19]), which may be seen as evidence for Conjecture 1.4.

Conjectural relationship to a cycle of Janda, Pandharipande, Pixton, and Zvonkine

In the appendix of [7], Janda, Pandharipande, Pixton, and Zvonkine define a cycle $H_{g,a}$ in the tautological ring of $\overline{\mathcal{M}}_{g,n}$ in the case where $k = 1$ and at least one $a_i < 0$, and conjecture that $H_{g,a}$ coincides with Pixton’s cycle $P_g^{g,k}(A)$. We are currently engaged (jointly with Johannes Schmitt) in verifying the equality $H_{g,a} = \overline{\text{DRC}}_a$, which may be viewed as a step towards Conjecture 1.4, or towards the conjecture of Janda, Pandharipande, Pixton, and Zvonkine, or both.

Comparison to the approach of Li, Graber and Vakil

The approach of Li [25], [26], Graber, and Vakil [12] when $k = 0$ is based on thinking of the DRC as the locus of curves admitting a map to \mathbb{P}^1 with specified ramification over 0 and ∞ . They define a stack of stable maps to ‘rubber \mathbb{P}^1 ’, i.e. to $[\mathbb{P}^1/\mathbb{G}_m]$. They then define the DRC as the pushforward of a virtual fundamental class from this stack of stable maps. This enables them to apply the well-developed machinery of virtual classes and spaces of stable maps. In contrast, our approach is in a sense very naive; using blowups to resolve the indeterminacies of rational maps goes back to classical algebraic geometry (and in the non-proper case to Raynaud and Gruson [30]). The more elementary nature of our approach makes it very easy to extend to the case $k \neq 0$, and we hope will allow further extensions; we are particularly interested in developing further the Gromov–Witten theory of $B\mathbb{G}_m$, extending the results of [9] beyond the ‘admissible’ case.

Comparison to the cycle of Kass and Paganì

The preprint [21] of Kass and Paganì (posted on the same day as the first version of this preprint) gives a very different approach to resolving the Abel–Jacobi map, by studying families of stability conditions on the space of rank 1 torsion-free sheaves. Moving through a suitable family produces a series of flips of a certain compactified jacobian, after which the Abel–Jacobi map extends over $\overline{\mathcal{M}}_{g,n}$ (this series of flips depends on the ramification data). In essence, we modify the source of the Abel–Jacobi map, whereas Kass and Paganì modify the target. In this way, they produce a number of different extensions of the DRC. In [18] we show that, for certain choices of stability conditions, the resulting double ramification cycle coincides with that given in this article.

Comparison to some other recent results

More recently, Marcus and Wise [28] have given another approach to resolving the Abel–Jacobi map when $k = 0$, rather closer in spirit to the present preprint. They also use logarithmic geometry to modify $\overline{\mathcal{M}}_{g,n}$, but their construction is based on stacks of stable maps rather than a universal property as in the present preprint. We hope to understand the relation between these approaches more fully in future.

An extension of the Abel–Jacobi map over a large locus in $\overline{\mathcal{M}}_{g,n}$ (when $k = 0$) was produced some time ago by Dudin [5]. His locus depended on a choice of ramification data, as in our construction. But he did not make modifications of $\overline{\mathcal{M}}_{g,n}$, and so was not able to extend over the whole of the boundary.

2. Notation and setup

Base ring

We work over the fixed base scheme $\Lambda := \text{Spec } \mathbb{Z}$ equipped with the trivial log structure. The reader who prefers to take $\Lambda = \text{Spec } \mathbb{C}$ can freely do so with no modifications to what follows, and a substantial simplification to the proofs of Proposition 5.3 and Lemma 6.1. All our constructions commute with arbitrary base change over Λ .

Remark 2.1. It seems that the definition of the double ramification cycle given in [12] can readily be extended over \mathbb{Z} , though we have not verified this. The computation of [20] is carried out over \mathbb{C} , and gives the class of the double ramification cycle as the value at zero of a polynomial P in a variable r . The value of P at a given value of r is computed using r th roots of line bundles (cf. Chiodo’s formulae [3]), which may give problems in characteristic dividing r . But in fixed characteristic ℓ , the polynomial P is completely determined by its values on integers coprime to ℓ , so this problem can be circumvented (the author is grateful to Felix Janda for pointing this out). So it remains likely that the results of [20] can be extended to arbitrary characteristic.

Stack of weighted stable curves

For us, ‘curve’ means proper, flat, finitely presented, with reduced connected nodal geometric fibres. Rather than treating each $\overline{\mathcal{M}}_{g,n}$ and weighting a_1, \dots, a_n separately, we denote by $\overline{\mathcal{M}}$ the stack of stable pointed curves together with a k -twisted integer weighting – in other words, points of $\overline{\mathcal{M}}$ consist of tuples

$$(C, x_1, \dots, x_n, a_1, \dots, a_n, k)$$

where the x_i are the marked sections of our stable curve, and the a_i and k are integers satisfying

$$\sum_i a_i = k(2g - 2).$$

This stack is smooth over Λ , but is far from being connected – it is a countably infinite disjoint union of substacks, each proper over Λ .

We write $\mathcal{C}/\overline{\mathcal{M}}$ for the universal curve, and $\mathcal{J} = \text{Pic}_{\mathcal{C}/\overline{\mathcal{M}}}^0$ for the universal jacobian (a semiabelian scheme, the fibrewise connected component of the identity in $\text{Pic}_{\mathcal{C}/\overline{\mathcal{M}}}$). Let \mathcal{M} denote the open substack of $\overline{\mathcal{M}}$ parametrising smooth curves. We write x_i for the tautological sections, and Σ for the Cartier divisor on \mathcal{C} given by $\sum_i a_i x_i$. Then $\sigma \in \mathcal{J}_{\mathcal{M}}(\mathcal{M})$ denotes the tautological section given by $\omega_{\mathcal{C}}^{\otimes k}(-\Sigma)$.

Log structures

We work with log structures in the sense of Fontaine–Illusie, using Olsson’s generalisation to stacks [29]. We put log structures on $\mathcal{C}/\overline{\mathcal{M}}$ following Kato [22], and the log structure on $\overline{\mathcal{M}}$ will be denoted by $\alpha_{\overline{\mathcal{M}}}: P_{\overline{\mathcal{M}}} \rightarrow \mathcal{O}_{\overline{\mathcal{M}}}$, etc.

If P is a (sheaf of) monoid(s), we write $\tilde{P} := P/P^\times$; this notation does not sit well with the notation $\overline{\mathcal{M}}$ for the moduli stack of stable curves, but both are very standard, and there is no actual ambiguity.

Weightings on a graph

A graph consists of a finite set V of vertices, a finite set H of half-edges, a map ‘end’ from the half-edges to the vertices, an involution i on the half-edges, a *genus* $g: V \rightarrow \mathbb{Z}_{\geq 0}$, and a *twist* $k \in \mathbb{Z}$. Graphs are assumed connected unless stated otherwise, and the genus of a graph is its first Betti number plus the sum of the genera of the vertices.

Self-loops are when two distinct half-edges have the same associated vertex and are swapped by i . We define *edges* as sets $\{h, h'\}$ (of cardinality 2) with $i(h) = h'$. *Legs* are fixed points of i . A *directed edge* h is a half-edge that is not a leg; we call $\text{end}(h)$ its *source* and $\text{end}(i(h))$ its *target*, and sometimes write it as $h: \text{end}(h) \rightarrow \text{end}(i(h))$. We write $E = E(\Gamma)$ for the set of edges, and \vec{E} for the set of directed edges.

The *valence* $\text{val}(v)$ of a vertex is the number of non-leg half-edges incident to it, and we define the *canonical degree* $\kappa(v) = 2g(v) - 2 + \text{val}(v)$, so that

$$2g(\Gamma) - 2 = \sum_v \kappa(v).$$

A *closed walk* in Γ is a sequence of directed edges so that the target of one is the source of the next, and which begins and ends at the same vertex. We call it a *cycle* if it does not repeat any vertices or (undirected) edges.

Definition 2.2. A *G-weighting* is a function w from the half-edges to a group G such that:

- (1) If $i(h) = h'$ and $h \neq h'$ then $w(h) + w(h') = 0$;
- (2) For all vertices v , $\sum_{\text{end}(h)=v} w(h) + k\kappa(v) = 0$.

When the twist $k = 0$, a G -weighting can be thought of as a flow of an incompressible fluid around the graph.

A G -leg-weighted graph is a graph together with a function from the legs to G , such that the sum over all the legs is $-k(2g(\Gamma) - 2)$. If Γ is a G -leg-weighted graph we write $W(\Gamma)$ for the set of weights on Γ which restrict to the given values on the legs. It is easy to see that $W(\Gamma)$ is never empty (this uses the running assumption of connectedness). After choosing an oriented basis of $H_1(\Gamma, \mathbb{Z})$, the set $W(\Gamma)$ becomes a torsor under $H_1(\Gamma, G)$. In this article we will use only the cases $G = \mathbb{Z}$ and $G = \mathbb{Q}$. We will refer to weightings taking values in \mathbb{Z} just as *weightings*.

Combinatorial charts

If $p: \text{Spec } k \rightarrow \overline{\mathcal{M}}$ is a geometric point, we have an associated graph Γ_p . If a node of the curve over p has local equation $xy - r$ for some $r \in \mathcal{O}_{\overline{\mathcal{M}}, p}^{et}$, then the image of r in the monoid $\overline{P}_{\overline{\mathcal{M}}, p}$ is independent of the choice of local equation. In this way we define a map $\ell: E(\Gamma_p) \rightarrow \overline{P}_{\overline{\mathcal{M}}, p}$, recalling that edges of the graph correspond to nodes of the curve. This is a logarithmic version of the labelling defined in [17].

Given a leg-weighted graph Γ with edge set E , set $\overline{\mathcal{M}}_\Gamma = \text{Spec } \Lambda[\mathbb{N}^E]$, equipped with the toric divisorial log structure. As usual we write $\alpha: P_{\overline{\mathcal{M}}_\Gamma} \rightarrow \mathcal{O}_{\overline{\mathcal{M}}_\Gamma}$ for the map from the sheaf of monoids to the structure sheaf.

To any point p in $\overline{\mathcal{M}}_\Gamma$ we associate the graph Γ_p obtained from Γ by contracting exactly those edges e such that the corresponding basis elements of \mathbb{N}^E specialise to units at p . We define a map $\ell: E(\Gamma_p) \rightarrow \overline{P}_{\overline{\mathcal{M}}_\Gamma, p}$ by sending an edge e to the image of the associated basis element of \mathbb{N}^E . The map naturally lifts to $P_{\overline{\mathcal{M}}_\Gamma, p}$, and does not send any edge to a unit, by definition of Γ_p .

A *combinatorial chart of $\overline{\mathcal{M}}$* consists of a leg-weighted graph Γ and a diagram of log stacks

$$\overline{\mathcal{M}} \xleftarrow{f} U \xrightarrow{g} \overline{\mathcal{M}}_\Gamma$$

satisfying the following five conditions:

- (1) U is a connected log scheme;
- (2) $g: U \rightarrow \overline{\mathcal{M}}_\Gamma$ is strict and log smooth;
- (3) $f: U \rightarrow \overline{\mathcal{M}}$ is strict and log étale;
- (4) the image of g meets the minimal stratum of $\overline{\mathcal{M}}_\Gamma$.

Let $p: \text{Spec } k \rightarrow U$ be any geometric point, yielding natural maps

$$\overline{P}_{\overline{\mathcal{M}}, f \circ p} \xrightarrow{f^b} \overline{P}_{U, p} \xleftarrow{g^b} \overline{P}_{\overline{\mathcal{M}}_\Gamma, g \circ p}.$$

- (5) We require the existence of an isomorphism

$$\varphi_p: \Gamma_{f \circ p} \rightarrow \Gamma_{g \circ p}$$

such that $f^b(\ell(e)) = g^b(\ell(\varphi_p(e)))$ for every edge e (which necessarily makes this φ_p unique if it exists). Moreover, the map φ_p sends the leg-weighting on $\Gamma_{f \circ p}$ coming from the $-a_i$ to the leg-weighting on $\Gamma_{g \circ p}$ coming from that on Γ .

It is clear from [23] that $\overline{\mathcal{M}}$ can be covered by combinatorial charts. We will first construct $\mathcal{M}_\Gamma^\diamond$ over the $\overline{\mathcal{M}}_\Gamma$, and then descend it to $\overline{\mathcal{M}}$.

3. Construction of \mathcal{M}^\diamond

The cone c_w associated to a weighting

For the remainder of this subsection we fix a combinatorial chart with leg-weighted graph Γ , writing E for the edge set. To a weighting $w \in W(\Gamma)$ we will associate a rational polyhedral cone c_w inside the positive orthant of $\mathbb{Z}^E \otimes_{\mathbb{Z}} \mathbb{Q}$. Recalling that

$\overline{\mathcal{M}}_\Gamma = \text{Spec } \Lambda[\mathbb{N}^E]$, such a cone will induce an affine toric scheme over $\overline{\mathcal{M}}_\Gamma$ in the usual way, cf. [11]. Such an object has a natural log structure. We will build $\mathcal{M}_\Gamma^\diamond$ by glueing together affine patches of this form.

Fix a weighting $w \in W(\Gamma)$ and let γ be an oriented cycle in Γ . If e is a directed edge appearing in γ , we define $w_\gamma(e)$ to be the value of w on the first half-edge of e – we might think of this as the flow along e in the direction given by γ .

Definition 3.1. Let $t \in \mathbb{Q}_{\geq 0}^E$; we refer to such an element as a *thickness*. We say t is *compatible with w* if for every cycle γ we have

$$\sum_{e \in \gamma} w_\gamma(e)t(e) = 0. \tag{2}$$

One checks easily that the set of all thicknesses t which are compatible with a given weighting w form a rational polyhedral cone in $\mathbb{Q}_{\geq 0}^E$, which we denote by c_w .

Lemma 3.2. Suppose $w, w' \in W(\Gamma)$ and that $c_w \cap c_{w'}$ contains a thickness t which does not vanish on any edge. Then $w = w'$.

Proof. Writing $\tilde{w} = w - w'$, we see that \tilde{w} is a ‘weighting’ for the graph Γ with all the leg decorations removed; we can think of this \tilde{w} as a flow of an incompressible fluid around Γ , with no sources and sinks. Suppose that \tilde{w} is not everywhere zero. We then claim that there exists a directed cycle γ in Γ such that, for every directed edge $e \in \gamma$, we have $\tilde{w}_\gamma(e) > 0$. To build such a cycle, we begin on any directed edge with positive flow \tilde{w} . The incompressibility condition then implies that we can continue the path along another edge, still having $\tilde{w} > 0$. Continuing in this way, the finiteness of the graph forces this path to intersect itself at some point. Possibly discarding the beginning of this path, we have the desired cycle γ . Then

$$\sum_{e \in \gamma} \tilde{w}_\gamma(e)t(e) = 0,$$

but all the $t(e) > 0$, and all the $\tilde{w}_\gamma(e) > 0$, a contradiction. □

Definition 3.3. We write F_Γ for the set of faces of the cones c_w as w runs over weightings in $W(\Gamma)$.

Remark 3.4 (Example: 2-gon, $k = 0$). Suppose the graph Γ has two edges and two (non-loop) vertices u and v . Suppose the leg-weighting is $+n$ at u and $-n$ at v . Weightings consist of a flow of a along one edge from u to v , and $n - a$ along the other (again from u to v), for $a \in \mathbb{Z}$. The cone c_w is non-zero if and only if both a and $n - a$ are non-negative, and for such a the cone c_w is the ray in $\mathbb{Q}_{\geq 0}^2$ through the point $(n - a, a)$. Thus we get exactly $n + 1$ rays in the positive quadrant. This is the fan F_Γ .

Remark 3.5 (Example: 3-edge banana, $k = 0$). Suppose again that the graph has two vertices u and v , but now three edges between them. Suppose that the weighting is $+10$

on u and -10 on v . Figure 1 shows the slice through the (incomplete) fan $F_\Gamma \subseteq \mathbb{Q}_{\geq 0}^3$ where the sum of the values of the thickness on the edges is 1.

In the next subsection we will verify that F_Γ is a *finite* fan (in the sense of toric geometry). The reader might prefer to skip the details, as they play little role in what follows.

F_Γ is a finite fan

Lemma 3.6. Let $w_1, w_2 \in W(\Gamma)$ be two weightings. Then the intersection of the cones c_{w_1} and c_{w_2} is a face of c_{w_1} .

Proof. Let $t \in c_{w_1} \cap c_{w_2}$. For an edge e with $t(e) \neq 0$ we claim $w_1(e) = w_2(e)$. To see this, let Γ_t be the graph obtained from Γ by contracting those edges on which t vanishes. Then each $w_i|_{\Gamma_t}$ is a weighting compatible with $t|_{\Gamma_t}$, and the latter does not vanish on any edge, so $w_1|_{\Gamma_t} = w_2|_{\Gamma_t}$ by Lemma 3.2.

Define

$$E_{=} = \{e \in E(\Gamma) : w_1(e) = w_2(e)\}, \quad E_{\neq} = E(\Gamma) \setminus E_{=}$$

Writing $\Gamma_{=}$ for the graph obtained from Γ by contracting E_{\neq} , and set $w_{=} := w_1|_{\Gamma_{=}} = w_2|_{\Gamma_{=}}$. Define $c_{=}$ to be the cone in $\mathbb{Q}_{\geq 0}^{E_{=}}$ corresponding to the weighting $w_{=}$. By the claim above we see that every $t \in c_w \cap c_{w'}$ vanishes on every edge of E_{\neq} .

I now claim that $c_{w_1} \cap c_{w_2} = c_{=} \times \{\mathbf{0}\}$, where $\mathbf{0}$ is the zero vector in $\mathbb{Q}_{\geq 0}^{E_{\neq}}$. Well, it is clear that $c_{w_1} \cap c_{w_2} \subseteq c_{=} \times \{\mathbf{0}\}$ by definition of $E_{=}$ and E_{\neq} . For the other inclusion, let $t \in c_{=} \times \{\mathbf{0}\}$, and let γ be a cycle in Γ . Write $\gamma_{=}$ for the closed walk in $\Gamma_{=}$ resulting from contracting γ . By definition of $c_{=}$ we have that

$$\sum_{e \in \gamma'} w_{=}(e)t(e) = 0,$$

and t vanishes on E_{\neq} , so we see

$$\sum_{e \in \gamma} w_1(e)t(e) = 0 = \sum_{e \in \gamma} w_2(e)t(e) = 0$$

as required. □

Now we know that any two c_w intersect in a face. The following well-known lemma shows that F_Γ is a fan.

Lemma 3.7. Let Φ_0 be a set of cones in \mathbb{Q}^n , and assume that for all $C, C' \in \Phi_0$, the intersection $C \cap C'$ is a face of C . Let Φ be the set of all faces of cones in Φ_0 . Then Φ is a fan.

Lemma 3.8. The set of cones $\{c_w : w \in W(\Gamma)\}$ is finite.

Proof. The proof is by induction on $h^1(\Gamma)$. Recalling that $W(\Gamma)$ is a torsor under $H_1(\Gamma, \mathbb{Z})$, we see that W is finite whenever $h^1(\Gamma) = 0$, so in this case there is nothing to prove.

In general, we say that a weighting w on Γ admits a positive cycle if there exists a cycle γ in Γ such that $w(e) > 0$ for every $e \in \gamma$. In the next two lemmas, we will show the following.

- (1) For fixed Γ , all but finitely many weightings admit a positive cycle (Lemma 3.9).
- (2) If γ is a positive cycle for w and Γ/γ is the graph obtained from Γ by contracting every edge in γ , then

$$c_w = c_{w|_{\Gamma/\gamma}} \times \{\mathbf{0}\},$$

where $c_{w|_{\Gamma/\gamma}} \subseteq \mathbb{Q}_{\geq 0}^{E(\Gamma/\gamma)}$ is the cone associated to the restricted weighting $w|_{\Gamma/\gamma}$, and $\mathbf{0}$ is the zero vector in $\mathbb{Q}_{\geq 0}^{E(\gamma)}$ (Lemma 3.10).

Now there are only finitely many cycles in Γ , and for every cycle γ we have that $h^1(\Gamma/\gamma) < h^1(\Gamma)$; hence by our induction hypothesis we have that there are only finitely many cones for Γ/γ . Putting these ingredients together concludes the proof. \square

Lemma 3.9. For fixed Γ , all but finitely many w admit a positive cycle.

Proof. Step 1: setup.

Fix a weighting w , and fix a basis B of $H_1(\Gamma, \mathbb{Z})$ consisting of cycles. We can think of $b \in B$ as a function from the set \vec{E} of directed edges of Γ , sending a directed edge e to 0 if $e \notin b$, and ± 1 otherwise (depending on whether the orientation of e agrees with b). Given an element $\mathbf{v} \in \mathbb{Z}^B$, we define $\text{cycle}(\mathbf{v})$ to be the function $\vec{E} \rightarrow \mathbb{Z}$ given by $\sum_b \mathbf{v}_b b$. This is somewhat clumsy notation, as it would be nicer just to think of \mathbf{v} as an element in the cycle space, but distinguishing carefully between \mathbf{v} and $\text{cycle}(\mathbf{v})$ seems important to avoid confusion in this proof. In this way we see that every weighting on Γ is of the form $w + \text{cycle}(\mathbf{v})$ for a unique $\mathbf{v} \in \mathbb{Z}^B$.

Define recursively a function $\varphi: \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}$ by setting $\varphi(0) = 1$ and $\varphi(n) = \sum_{0 \leq j < n} \varphi(j)$. Write $m := \max_{e \in E(\Gamma)} |w(e)|$, and $h := h^1(\Gamma)$. Define $N = m\varphi(h)$ (this is rigged exactly to make step 3 of this argument work). Define

$$B(N) = \left\{ \mathbf{v} \in \mathbb{Z}^B : \max_{b \in B} |\mathbf{v}_b| \leq N \right\},$$

a finite set. In the remainder of this argument, we will show that for every $\mathbf{v} \in \mathbb{Z}^B \setminus B(N)$, the weighting $w + \text{cycle}(\mathbf{v})$ admits a positive cycle. To this end, we fix for the remainder of the argument a $\mathbf{v} \in \mathbb{Z}^B \setminus B(N)$.

Step 2: ordering the \mathbf{v}_b .

Up to changing the orientations of the elements of B , we may assume that all the integers \mathbf{v}_b are non-negative. Put an ordering on B so that the \mathbf{v}_{b_i} are in increasing order:

$$0 \leq \mathbf{v}_{b_1} \leq \mathbf{v}_{b_2} \leq \dots \leq \mathbf{v}_{b_h},$$

with $\mathbf{v}_{b_h} > N$.

Step 3: choosing a critical b_r .

We now show by a small computation that there exists $1 \leq r \leq h$ such that

$$m + \sum_{1 \leq i \leq r-1} \mathbf{v}_{b_i} < \mathbf{v}_{b_r}.$$

Indeed, suppose no such r exists. Then for each $1 \leq j \leq h$ we have

$$m + \sum_{1 \leq i \leq j-1} \mathbf{v}_{b_i} \geq \mathbf{v}_{b_j},$$

and induction on j yields $\mathbf{v}_{b_j} \leq m\varphi(j)$ for all $1 \leq j \leq h$, contradicting our assumption that $\mathbf{v}_{b_h} > N = m\varphi(h)$.

From now on, we fix such an r .

Step 4: finding the positive cycle γ .

Define a function $f: \vec{E} \rightarrow \mathbb{Z}$ sending a directed edge e to

$$\sum_{r \leq j \leq h} \left\{ \begin{array}{l} 1 \text{ if } e \in b_j \text{ and has same direction as } b_j \\ -1 \text{ if } e \in b_j \text{ and has opposite direction to } b_j \\ 0 \text{ otherwise (i.e. } e \notin b_j) \end{array} \right\}$$

In the remainder of this step we will show that there exists a cycle γ with $f(e) > 0$ for every directed edge $e \in \gamma$. In step 5 we will see that any such γ is necessarily a positive cycle.

First, because the b_j are part of a basis, we see that f is not identical to zero. Hence there is a directed edge e with $f(e) > 0$. We build a path in Γ starting with e by the following procedure: whenever we hit a vertex v , choose an edge e_v out of v such that $f(e_v) > 0$. Why is this always possible? Note that the sum over all edges e'' into v of $f(e'')$ is necessarily zero, and since we arrived at v along an edge with $f > 0$, there must also be an edge leaving v with $f > 0$.

Since Γ is finite, this path must eventually meet itself, say at a vertex v_0 . Deleting the start of the path up to v_0 yields the cycle γ that we sought.

Step 5: showing that γ is indeed a positive cycle for the weighting $w + \text{cycle}(\mathbf{v})$.

Choose any γ as in step 4. Define functions $F, G: \vec{E} \rightarrow \mathbb{Z}$ by the following formulae:

$$F(e) = \sum_{r \leq j \leq h} \left\{ \begin{array}{l} \mathbf{v}_{b_j} \text{ if } e \in b_j \text{ and has same direction as } b_j \\ -\mathbf{v}_{b_j} \text{ if } e \in b_j \text{ and has opposite direction to } b_j \\ 0 \text{ otherwise (i.e. } e \notin b_j) \end{array} \right\}$$

$$G(e) = w(e) + \sum_{1 \leq j \leq r-1} \left\{ \begin{array}{l} \mathbf{v}_{b_j} \text{ if } e \in b_j \text{ and has same direction as } b_j \\ -\mathbf{v}_{b_j} \text{ if } e \in b_j \text{ and has opposite direction to } b_j \\ 0 \text{ otherwise (i.e. } e \notin b_j) \end{array} \right\}.$$

Observe that $F + G = w + \text{cycle}(\mathbf{v})$ as functions $\vec{E} \rightarrow \mathbb{Z}$. Let $e \in \gamma$ be a directed edge; we will show that $F(e) + G(e) > 0$ as required.

Define $\epsilon = m + \sum_{1 \leq i \leq r-1} \mathbf{v}_{b_i}$ (cf. step 3). Note that $f(e) \geq 1$, and $F(e) > \epsilon f(e) \geq \epsilon$ since $\mathbf{v}_{b_j} > \epsilon$ for every $j \geq r$. Now

$$|G(e)| \leq m + \sum_{1 \leq j \leq r-1} \mathbf{v}_{b_j} = \epsilon,$$

hence

$$(w + \text{cycle}(\mathbf{v}))(e) = F(e) + G(e) > \epsilon - \epsilon = 0. \quad \square$$

Lemma 3.10. Fix a weighting w . If γ is a positive cycle for w and Γ/γ is the graph obtained from Γ by contracting every edge in γ , then

$$c_w = c_{w|_{\Gamma/\gamma}} \times \{\mathbf{0}\},$$

where $c_{w|_{\Gamma/\gamma}} \subseteq \mathbb{Q}_{\geq 0}^{E(\Gamma/\gamma)}$ is the cone associated to the restricted weighting $w|_{\Gamma/\gamma}$, and $\mathbf{0}$ is the zero vector in $\mathbb{Q}_{\geq 0}^{E(\gamma)}$.

Proof. The inclusion

$$c_w \supseteq c_{w|_{\Gamma/\gamma}} \times \{\mathbf{0}\}$$

follows easily from the definition of the cone of a weighting, since every cycle in Γ/γ arises by restricting some cycle in Γ .

We need to show the other inclusion, so let $t \in c_w$ be a thickness. The t satisfies the equation

$$\sum_{e \in \gamma} t(e)w(e) = \mathbf{0},$$

and since all the $w(e)$ are positive this forces all the $t(e)$ to vanish. □

Putting together Lemmas 3.6–3.8 we immediately deduce:

Corollary 3.11. The set of cones F_Γ is a finite fan inside $\mathbb{Q}_{\geq 0}^E$.

The construction of \mathcal{M}^\diamond

Definition 3.12. We define $\mathcal{M}_\Gamma^\diamond$ to be the toric scheme over $\overline{\mathcal{M}}_\Gamma$ defined by the fan F_Γ . We equip it with the toric log structure.

If $w \in W(\Gamma)$ we define \mathcal{M}_w^\diamond to be the affine toric variety associated to c_w , an affine patch of $\mathcal{M}_\Gamma^\diamond$.

Remark 3.13. It follows from [11] that $\mathcal{M}_\Gamma^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$ is separated, of finite presentation, and normal. It is moreover logarithmically étale, since it is given by patches of toric blowups.

Given a combinatorial chart $\overline{\mathcal{M}} \leftarrow U \rightarrow \overline{\mathcal{M}}_\Gamma$ we define \mathcal{M}_U^\diamond by pulling back $\mathcal{M}_\Gamma^\diamond$ from $\overline{\mathcal{M}}_\Gamma$. Such U form an étale cover of $\overline{\mathcal{M}}$, and the collection of \mathcal{M}_U^\diamond is easily upgraded to a descent datum.

Definition 3.14. We define $\pi^\diamond: \mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ to be the algebraic space obtained by descending the \mathcal{M}_U^\diamond .

Theorem 3.15. The stack \mathcal{M}^\diamond is normal, and the map $\pi^\diamond: \mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is separated, of finite presentation, relatively representable by algebraic spaces, birational, and logarithmically étale.

Note that π^\diamond is almost never proper.

Proof. Toric varieties in this sense are always normal, see [11]. The properties of π^\diamond are all local on the target, so it is enough to check them for the $\mathcal{M}_\Gamma^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$. Separatedness and logarithmic-étaleness are automatic for toric varieties (in the sense of Fulton), finite presentation follows from the finiteness of the fans, see Corollary 3.11. The maps are clearly isomorphisms over the locus \mathcal{M} of smooth curves, hence are birational. □

4. Universal property of \mathcal{M}^\diamond

Definition 4.1. We say a stack T is *locally desingularisable* if étale-locally it admits a proper surjective finitely presented map $T' \rightarrow T$ with T' regular and $T' \rightarrow T$ inducing an isomorphism between some dense open substacks of T' and T .

For example, this is true in characteristic zero, and for arithmetic surfaces and threefolds, and for log regular schemes. We can now give a more precise variant of the notion of a σ -*extending* morphism from the introduction. Recall that $\sigma: \mathcal{M} \rightarrow \mathcal{J}$ is the morphism given by $\omega^{\otimes k}(-\sum_i a_i x_i)$.

Given any map $t: T \rightarrow \overline{\mathcal{M}}$, we can consider the open subset $t^{-1}\mathcal{M} = T \times_{\overline{\mathcal{M}}} \mathcal{M} \hookrightarrow T$, and the semiabelian scheme $t^*\mathcal{J} = T \times_{\overline{\mathcal{M}}} \mathcal{J}$ over T . We denote by $t^*\sigma$ the canonical map $id_T \times \sigma: t^{-1}\mathcal{M} \rightarrow t^*\mathcal{J}$.

Definition 4.2. We say a map $t: T \rightarrow \overline{\mathcal{M}}$ is *non-degenerate* if T is normal and locally desingularisable, and $t^{-1}\mathcal{M}$ is dense in T . We say t is σ -*extending* if in addition the section $t^*\sigma \in (t^*\mathcal{J})(t^{-1}\mathcal{M})$ admits a (necessarily unique) extension to $t^*\mathcal{J}(T)$.

For example, the open immersion $\mathcal{M} \rightarrow \overline{\mathcal{M}}$ is clearly σ -extending, and in general the identity on $\overline{\mathcal{M}}$ is not σ -extending. In this section we will show that $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is σ -extending, and moreover that \mathcal{M}^\diamond is universal with respect to this property: that it is terminal in the 2-category of σ -extending morphisms (Corollary 4.6). Our main technical result is the following:

Lemma 4.3. Fix a combinatorial chart $\overline{\mathcal{M}} \leftarrow U \rightarrow \overline{\mathcal{M}}_\Gamma$, and let $t: T \rightarrow U$ be such that the composite $T \rightarrow \overline{\mathcal{M}}$ is non-degenerate. The following are equivalent:

- (1) Locally on T there exists a weighting w on Γ such that $T \rightarrow \overline{\mathcal{M}}_\Gamma$ factors via $\mathcal{M}_w^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$.
- (2) $T \rightarrow \overline{\mathcal{M}}$ is σ -extending.

Before giving the proof we set up a little notation, which will also be useful in § 5. The map $T \rightarrow \overline{\mathcal{M}}$ defines a stable curve over T , which we shall denote by C . For each edge $e \in E = E(\Gamma)$ we define $\ell(e) \in \mathcal{O}_{\overline{\mathcal{M}}_\Gamma}(\overline{\mathcal{M}}_\Gamma)$ to be the image of the standard basis vector $\delta_e \in \mathbb{N}^E$ under the log structure map. If $w \in W(\Gamma)$ is a weighting and $\gamma \subseteq \Gamma$ a cycle, we let

$$\delta_\gamma = \prod_{e \in \gamma} \delta_e^{w_\gamma(e)} \in \mathbb{N}^E$$

where $w_\gamma(e) \in \mathbb{Z}$ is the value of w on e in the direction dictated by γ . One easily verifies the following lemma.

Lemma 4.4. The dual cone $c_w^\vee \subseteq \mathbb{Z}^E$ is the intersection of \mathbb{Z}^E with the rational span of the positive orthant in \mathbb{Z}^E together with the δ_γ for γ running over cycles in Γ .

Proof of Lemma 4.3. (1) \implies (2): we may assume T is local, since an extension is unique if it exists. Perhaps shrinking the combinatorial chart we may assume that Γ is the graph over the closed point of T .

The map $T \rightarrow \overline{\mathcal{M}}_\Gamma$ corresponds to a map $t^\#: \Lambda[\mathbb{N}^E] \rightarrow \mathcal{O}_T(T)$, and each $t^\#\delta_e$ is a unit on $t^{-1}\mathcal{M}$. For each cycle γ we obtain an element $t^\#\delta_\gamma \in \text{Frac } \mathcal{O}_T(T)$, and the factorisation of t via \mathcal{M}_w^\diamond says that $t^\#\delta_\gamma \in \mathcal{O}_T(T) \subseteq \text{Frac } \mathcal{O}_T(T)$. If we write $i(\gamma)$ for the cycle with the same edges as γ but in the reverse direction, we see that $t^\#\delta_{i(\gamma)} = t^\#\delta_\gamma^{-1}$ in $\text{Frac } \mathcal{O}_T(T)$, and hence that actually each $t^\#\delta_\gamma$ is a unit, i.e. lies in $\mathcal{O}_T(T)^\times \subseteq \mathcal{O}_T(T)$.

Since the product around each cycle in Γ of the $t^\#\delta_e^{w_\gamma(e)}$ lies in $\mathcal{O}_T(T)^\times$, we can choose elements $r_v \in \text{Frac } \mathcal{O}_T(T)^\times$ for each vertex v of Γ such that for each directed edge $e: u \rightarrow v$ we have

$$\frac{r_u}{r_v} = t^\#\delta_e^{w_\gamma(e)} \cdot (\text{unit in } \mathcal{O}_T(T)). \tag{3}$$

For a vertex v , write η_v for the generic point of the component of the special fibre of C corresponding to v . Now we define a Weil divisor Y on C by specifying that Y is trivial over $t^{-1}\mathcal{M}$, and that locally around η_v it is cut out by r_v . Then (3) and a small computation implies that Y is actually a Cartier divisor.

Now we claim that the line bundle $\omega_C^{\otimes k}(\Sigma + Y)$ defines an extension of σ in $t^*J = \text{Pic}_{C/T}^0$. Clearly it coincides with σ over $t^{-1}\mathcal{M}$, so all we need to check is that $\Sigma + Y$ has degree zero on every irreducible component of the special fibre of C . Fix a vertex v . We need to check that the degree of $\mathcal{O}_C(Y)$ on the component C_v of the special fibre corresponding to v is exactly the sum of the weights of the non-leg half-edges out of v . After adjusting Y by the pullback of a divisor on T we may assume that $r_v = 1$. Let $e: u \rightarrow v$ be an edge out of v . Then the completed étale local ring at the singular point corresponding to e is isomorphic to

$$\widehat{\mathcal{O}_T(T)^{et}}[[x, y]]/(xy - t^\#\delta_e),$$

where we take x to be the coordinate vanishing on C_v and y to be vanishing on C_u . We may assume that r_u is given by $t^\#\delta_e^{w(e)}$ (since we can ignore T -units), so Y is locally defined by $y^{w(e)}$, and the order of vanishing on C_v is exactly $w(e)$ as required.

(2) \implies (1): Again, we may assume T is local. We consider first the case where T is regular. The argument above is almost reversible; we are assuming this extension of σ exists, and it is necessarily given by a line bundle \mathcal{L} of degree 0 on every irreducible component of every fibre. Then $\mathcal{L}(-\Sigma) = \mathcal{O}(Y)$ for some vertical Cartier divisor Y . For each vertex v of Γ , let $r_v \in \text{Frac } \mathcal{O}_T(T)^\times$ be a local equation for Y at η_v .

By the regularity of T we can apply [17, Theorem 4.1] to see that for each directed edge $e: u \rightarrow v$, an equation of the form

$$r_u/r_v = t^\# \delta_e^a \cdot (\text{unit in } \mathcal{O}_T(T)) \tag{4}$$

holds for some $a \in \mathbb{Z}$. Assigning to the directed edge e the integer a is easily verified to give a weighting w on Γ , and for each cycle γ we see that $\prod_{e \in \gamma} t^\# \delta_e^{w_\gamma(e)} \in \text{Frac } \mathcal{O}_T(T)^\times$ actually lies in the subgroup $\mathcal{O}_T(T)^\times \subseteq \text{Frac } \mathcal{O}_T(T)^\times$. Hence the map

$$\begin{aligned} \Lambda[c_w^\vee] &\rightarrow \text{Frac } \mathcal{O}_T(T) \\ \delta_e &\mapsto t^\# \delta_e \end{aligned}$$

factors via $\mathcal{O}_T(T) \subseteq \text{Frac } \mathcal{O}_T(T)$, and we are done.

It remains to reduce the general case to the case when T is regular. So assume T is local, normal and (locally) desingularisable, and let $T' \rightarrow T$ be a desingularisation, so T' is regular and $T' \rightarrow T$ is proper, surjective, and birational. By Zariski's Main Theorem, the fibres of $T' \rightarrow T$ are connected. Write \mathfrak{t} for the closed point of T .

Since we know the result in the regular case, we can apply this to T' to find an open cover $\{V_i\}_{i \in I}$ of T' , and weightings w_i on the V_i , such that each $V_i \rightarrow \overline{\mathcal{M}}_\Gamma$ factors via $\mathcal{M}_{w_i}^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$. Adjusting the cover, we may assume that each V_i is connected and meets the fibre $T'_\mathfrak{t}$ of T' over the closed point of T . If Γ' is the graph over \mathfrak{t} , the weightings w_i on Γ need not be unique, but their restrictions to the contracted graph Γ' are unique. To simplify the notation we will assume $\Gamma' = \Gamma$, since the value taken by the weightings on the contracted edges never plays any role.

Now $T'_\mathfrak{t}$ is connected and is covered by the $V_i \cap T'_\mathfrak{t}$, and the weightings w_i must agree on overlaps of the $V_i \cap T'_\mathfrak{t}$, so we see that actually all the w_i are equal. Write w for this weighting; we will show that $T \rightarrow \overline{\mathcal{M}}_\Gamma$ factors via $\mathcal{M}_w^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$.

Note that each $t^\# \delta_e$ is a regular element in $\mathcal{O}_T(T)$ since its restriction to $t^{-1}\mathcal{M}$ is invertible. We write $\text{div}(t^\# \delta_e)$ for the associated Cartier divisor on T .

Fix a directed loop γ in Γ . By a similar argument as in the regular case, to construct the map $T \rightarrow \mathcal{M}_w^\diamond$ it is enough to show that, for each cycle γ we have

$$\sum_{e \in \gamma} w_\gamma(e) \text{div}(t^\# \delta_e) = 0. \tag{5}$$

But we know that (5) holds on each V_i after pulling back (since we have maps $V_i \rightarrow \mathcal{M}_w^\diamond$) so the result follows from [17, Lemma 2.23].

Corollary 4.5. Let $t: T \rightarrow \overline{\mathcal{M}}$ be any σ -extending morphism. Then t factors uniquely via $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$.

Proof. Uniqueness is clear since the map is determined on \mathcal{M} , whose pullback is dense in T . Existence then follows immediately from Lemma 4.3, since $\mathcal{M}_\Gamma^\diamond$ is formed by

glueing together the \mathcal{M}_w^\diamond , and the uniqueness ensures that the maps we obtain glue on overlaps. □

Corollary 4.6. $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is the terminal object in the 2-category of σ -extending morphisms to $\overline{\mathcal{M}}$.

Proof. We know by Theorem 3.15 that \mathcal{M}^\diamond is normal, and it is locally desingularisable by [24] since it is locally toric (log regular). Applying Lemma 4.3 we see that $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is σ -extending, since $\mathcal{M}_w^\diamond \rightarrow \overline{\mathcal{M}}_\Gamma$ is clearly w -aligned. Corollary 4.5 then shows that it is terminal. □

5. Properness of DRL_\diamond

Write $\sigma'_\diamond : \mathcal{M}^\diamond \rightarrow J$ for the extension of σ (we reserve the notation σ_\diamond for the induced map $\mathcal{M}^\diamond \rightarrow \mathcal{M}^\diamond \times_{\overline{\mathcal{M}}} J$). Write DRL_\diamond for the schematic pullback of the unit section of the universal jacobian along σ'_\diamond , so DRL_\diamond is a closed substack of \mathcal{M}^\diamond . We will show that DRL_\diamond is proper over $\overline{\mathcal{M}}$ (recalling that the map $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is in general far from proper, cf. Figure 1).

5.1. In characteristic zero

Over a field of characteristic zero (or more generally, over a ring over which all schemes of finite type admit a resolution of singularities) this result can be proven more directly. Namely, we will show that \mathcal{M}^\diamond coincides with the normalisation of the closure of the image of σ in J , from which properness follows immediately. More precisely, it is the underlying scheme of \mathcal{M}^\diamond which is given by this construction, since the normalisation of the closure does not come with a natural logarithmic structure.

Write S for the schematic image of the section σ in J (in other words, for the closure of the image, with suitable reduced structure), and S' for its normalisation. We will show that S' is a universal σ -extending morphism. Indeed, the pullback of \mathcal{M} is evidently dense in S' , and S' is automatically normal, and admits a resolution of singularities by [16] (here we use characteristic zero). Moreover, S' comes with a map to J making it σ -extending.

Suppose $T \rightarrow \mathcal{M}$ is also σ -extending, then the map $T \rightarrow J$ factors via the inclusion $S \hookrightarrow J$ by definition of the schematic image, and the resulting map $T \rightarrow S$ lifts to S' since T is assumed normal. We thus see that S' is the universal σ -extending morphism, and so is canonically isomorphic to \mathcal{M}^\diamond .

To prove properness of DRL_\diamond over $\overline{\mathcal{M}}$ we simply observe that $S' \rightarrow S$ is proper, so DRL_\diamond is proper over the intersection of S with the unit section in J , and that intersection is automatically proper over $\overline{\mathcal{M}}$ since it is closed in the image of the unit section.

In positive characteristic, we do not know that S' is σ -extending (Definition 4.2), since we do not know if it admits a resolution of singularities. In the remainder of this section we will give a more hands-on proof of properness in this case. Some of the details (in particular the description of the universal line bundle) may also be of independent interest.

5.2. The universal line bundle

Before giving the proof in the general case, we describe a certain line bundle on the universal curve over \mathcal{M}^\diamond , which will play a crucial role in the proof – this may be of some independent interest. The proof itself is then mainly a matter of keeping careful track of isomorphisms and valuations. To simplify the notation, we write ω for the relative dualising sheaf of the universal curve over $\overline{\mathcal{M}}$.

Let $p: \text{Spec } k \rightarrow \mathcal{M}^\diamond$ be a geometric point, and write C_p for the stable curve over k . The map $\sigma'_\diamond: \mathcal{M}^\diamond \rightarrow \mathcal{J}$ determines an isomorphism class of line bundles on C_p (necessarily with degree zero on every component). In this section we will give representatives of this isomorphism class.

Remark 5.1. In fact, we do not need to assume that the field k is algebraically closed; it is enough to assume that all irreducible components of C_p are geometrically irreducible, and that all preimages of all nodes in the normalisation of C_p are k -rational points.

Choose a combinatorial chart $\overline{\mathcal{M}}_\Gamma \leftarrow U \rightarrow \overline{\mathcal{M}}$ containing p , and such that Γ is the graph of C_p . Let w be the weighting on Γ such that p lies in \mathcal{M}_w^\diamond . If v is a vertex of Γ , write C_v for the corresponding irreducible component of C_p , and define a line bundle \mathcal{F}_v on C_v by the formula

$$\mathcal{F}_v = \omega^{\otimes k}|_{C_v} \left(\sum_e w(e) \cdot [e_v] \right). \tag{6}$$

Here the sum runs over directed edges e out of v , and e_v is the point on C_v corresponding to the node e on C_p . If e is a self-loop then the point e_v is not a Cartier divisor on C_v , but (by the definition of a weighting) it appears with coefficient 0 in the above formula, so we do not worry about this. Writing \tilde{C}_v for the normalisation of C_v , we have $\omega|_{\tilde{C}_v} = \Omega_{\tilde{C}_v}(\sum_e [e_v])$ (where the sum runs over all directed edges with an end at v), and the restriction of the latter to any node e_v is canonically trivialised by the residue map.

Now we need to glue the \mathcal{F}_v together along the non-self-loop edges e to give a line bundle on C_p . If $e: u \rightarrow v$ is a non-self-loop edge then we see

$$\begin{aligned} \text{Hom}_k(\mathcal{F}_v|_{e_v}, \mathcal{F}_u|_{e_u}) &= \left(\omega^{\otimes k}|_{C_v}(w(e)[e_v]) \right)|_{e_v}^{\otimes -1} \otimes \left(\omega^{\otimes k}|_{C_u}(-w(e)[e_u]) \right)|_{e_u} \\ &= \mathcal{O}_{C_v}(w(e)[e_v])|_{e_v}^{\otimes -1} \otimes \mathcal{O}_{C_u}(-w(e)[e_u])|_{e_u} \\ &= \left(\mathcal{O}_{C_v}([e_v])|_{e_v} \otimes \mathcal{O}_{C_u}([e_u])|_{e_u} \right)^{\otimes -w(e)}. \end{aligned} \tag{7}$$

By the deformation theory of stable curves the vector space $\mathcal{O}_{C_v}([e_v])|_{e_v} \otimes \mathcal{O}_{C_u}([e_u])|_{e_u}$ is naturally a sub-space of the tangent space to p in $\overline{\mathcal{M}}$. The choice of combinatorial chart then yields a canonical generator of this summand of the tangent space, giving us a canonical isomorphism

$$\mathcal{O}_{C_v}([e_v])|_{e_v} \otimes \mathcal{O}_{C_u}([e_u])|_{e_u} \xrightarrow{\sim} k. \tag{8}$$

We can use this to explicitly describe how to glue the \mathcal{F}_v together to a line bundle on C_p . The map $p: \text{Spec } k \rightarrow \mathcal{M}_w^\diamond$ corresponds to a map $p^\#: \Lambda[c_w^\vee] \rightarrow k$, and for each cycle

γ we see $p^\# \delta_\gamma \in k^\times$. Choose a function

$$\lambda: \vec{E} \rightarrow k^\times \tag{9}$$

such that $\lambda(i(e)) = \lambda(e)^{-1}$ and such that for every cycle γ we have

$$\prod_{e \in \gamma} \lambda(e)^{w(e)} = p^\# \delta_\gamma. \tag{10}$$

For an edge $e: u \rightarrow v$, the $-w(e)$ th power of the element $\lambda(e) \in k^\times$ gives (via (8) and (7)) an isomorphism $\mathcal{F}_v|_{e_v} \xrightarrow{\sim} \mathcal{F}_u|_{e_u}$. We use these isomorphisms to glue the \mathcal{F}_v to a line bundle on the whole of C_p (cf. [8]), which we denote by \mathcal{F}_λ . Clearly \mathcal{F}_λ depends on the choice of λ , but a different choice of λ will yield an isomorphic \mathcal{F}_λ .

Proposition 5.2. The isomorphism class $[\mathcal{F}_\lambda(-\Sigma)]$ of line bundles on C_p corresponds to the image of the section σ'_\diamond in $\text{Pic}_{C_p/k}^0$.

Proof. This follows by a rather messy unravelling of the constructions in §3. □

5.3. The proof of properness in general

Proposition 5.3. The map $\text{DRL}_\diamond \rightarrow \overline{\mathcal{M}}$ is proper.

Proof. Step 1: Setup.

The map is clearly separated and of finite presentation since the same holds for $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$. We need to show that the dashed arrow in the following diagram can be filled in:

$$\begin{array}{ccc}
 \eta & \longrightarrow & \text{DRL}_\diamond \\
 \downarrow & \nearrow & \downarrow \\
 T & \xrightarrow{t} & \overline{\mathcal{M}}
 \end{array} \tag{11}$$

where T is a strictly henselian trait with generic point η and closed point p . Choose a combinatorial chart $\overline{\mathcal{M}}_\Gamma \leftarrow U \rightarrow \overline{\mathcal{M}}$ containing p , and such that $\Gamma = \Gamma_p$ is the graph of C_p .

We write Γ_η for the graph over η , with edge set E_η and vertex set V_η , and similarly over p , so we have a contraction map $\Gamma_p \rightarrow \Gamma_\eta$, and $E_\eta \subseteq E_p$. Given $v \in V_\eta$ we define C_v to be the corresponding irreducible component of C_η .

Let w_η be a weighting on Γ_η such that η lands in $\mathcal{M}_{w_\eta}^\diamond$.

Step 2: Extending the weighting to Γ_p .

Because the trait T is strictly henselian, we know that all irreducible components of C_η are geometrically irreducible, and all nodes and their tangent directions are rational over η . Because of this, by Remark 5.1, we can apply Proposition 5.2 over η . Accordingly, we choose

- a line bundle \mathcal{L} on C_η ;
- for each $v \in V_\eta$ an isomorphism

$$\mathcal{L}|_{C_v} \xrightarrow{\sim} \omega^{\otimes k}|_{C_v} \left(\sum_v w(e)[e_v] \right) = \mathcal{F}_v. \tag{12}$$

- an isomorphism

$$\mathcal{L}(-\Sigma) \xrightarrow{\sim} \mathcal{O}_{C_\eta}; \tag{13}$$

(the last is possible exactly because η lands in DRL_\diamond). Putting together (13) and (12) we obtain for each $e: u \rightarrow v$ an isomorphism

$$\mathcal{F}_v|_{e_v} \xrightarrow{\sim} \mathcal{F}_u|_{e_u}. \tag{14}$$

Perhaps after replacing η by a finite extension, we can choose elements $\lambda(e) \in \mathcal{O}_T(\eta)$ as in (9), satisfying (10) and inducing via (7) and (8) the isomorphisms (14).

We write \bar{C}_v for the closure of C_v in the curve $C = C_T$ (so the \bar{C}_v are the irreducible components of C), and we write $\bar{\mathcal{F}}_v$ for the line bundle on \bar{C}_v given by

$$\bar{\mathcal{F}}_v = \omega^{\otimes k}|_{\bar{C}_v} \left(\sum_s w(e)[e_v] \right), \tag{15}$$

where now we view e_v as an element of $\bar{C}_v(T)$ (cf. (6)). Putting together (13) and (12) again we obtain trivialisations $\bar{\mathcal{F}}_v(-\Sigma) \otimes \omega^{\otimes -k} \xrightarrow{\sim} \mathcal{O}_{C_v}$, which we can think of as being trivialisations of $\bar{\mathcal{F}}_v(-\Sigma) \otimes \omega^{\otimes -k}$ over η , which yield Cartier divisors Y_v on \bar{C}_v supported on the special fibre.

If $v' \in V_p$ is a vertex mapping to $v \in V_\eta$, we define $\mathcal{Y}(v')$ to be the multiplicity along the generic point of $C_{v'}$ of the divisor Y_v ; this gives a function $\mathcal{Y}: V_p \rightarrow \mathbb{Z}$. Given $e: u \rightarrow v \in E_p \setminus E_\eta$ we define

$$w(e) = \frac{\mathcal{Y}(v) - \mathcal{Y}(u)}{\text{thickness}(e)} \in \mathbb{Z}. \tag{16}$$

One now checks easily that this w extends the weighting w_η to a weighting w onto the whole of Γ_p .

Step 3: Computing $\text{ord}_T \lambda(e)$.

Let $e: u \rightarrow v \in E_\eta$, so we have $\lambda(e) \in \mathcal{O}_T(\eta)$ giving the glueing $\mathcal{F}_v|_{e_v} \xrightarrow{\sim} \mathcal{F}_u|_{e_u}$. Over T we have isomorphisms

$$\bar{\mathcal{F}}_v(Y_v)|_{e_v} = \bar{\mathcal{F}}_v(-\Sigma + Y_v)|_{e_v} \xrightarrow{a} \omega^{\otimes k}|_{e_v} = \omega^{\otimes k}|_{e_u} \xrightarrow{b} \bar{\mathcal{F}}_u(-\Sigma + Y_u)|_{e_u} = \bar{\mathcal{F}}_u(Y_u)|_{e_u},$$

where the isomorphisms a and b come from the definition of Y_v . Now $e \in E_\eta$ lifts to a unique edge $e' \in E_p$, and we write $v' \in V_p$ for the endpoint of e' which maps to v , and similarly define u' . A small computation then shows that

$$\text{ord}_T \lambda(e) = \mathcal{Y}(v') - \mathcal{Y}(u'). \tag{17}$$

Step 4: Constructing a lift $T \rightarrow \mathcal{M}_w^\diamond$.

We want to construct a map $\Lambda[c_w^\vee] \rightarrow \mathcal{O}_T(T)$ (recall the description of c_w^\vee from §4). The map $\Lambda[\mathbb{N}^{E_p}] \rightarrow \mathcal{O}_T(T)$ vanishes exactly on δ_e for $e \in E_\eta \subseteq E_p$, so we have a map $\Lambda[\mathbb{Z}^{E_p \setminus E_\eta}] \rightarrow \mathcal{O}_T(\eta)$, which we extend to a map $\varphi: \Lambda[\mathbb{Z}^{E_p}] \rightarrow \mathcal{O}_T(\eta)$ by sending δ_e to $\lambda(e)$ for each $e \in E_\eta$.

Let γ be a cycle in Γ_p . Then

$$\sum_{(e: u \rightarrow v) \in \gamma} \text{ord}_T \varphi(\delta_e)^{w_\gamma(e)} = \sum_{(e: u \rightarrow v) \in \gamma} \mathcal{Y}(v) - \mathcal{Y}(u) = 0,$$

so $\varphi: \Lambda[\mathbb{Z}^{E_p}] \rightarrow \mathcal{O}_T(\eta)$ restricts to a map $\Lambda[c_w^\vee] \rightarrow \mathcal{O}_T(T)$ as required.

Step 5: Verifying the valuative criterion.

We have constructed a map $T \rightarrow \mathcal{M}^\diamond$ over t whose restriction to η is as in (11). Since the inclusion $\text{DRL}_\diamond \rightarrow \mathcal{M}^\diamond$ is proper it follows that $T \rightarrow \mathcal{M}^\diamond$ factors via DRL_\diamond as required. \square

6. Proof of Theorems 1.1 and 1.2

We now have all the tools to easily prove Theorems 1.1 and 1.2. We begin by giving slightly more precise statements.

An *admissible modification* of $\overline{\mathcal{M}}$ is a morphism $x: X \rightarrow \overline{\mathcal{M}}$ satisfying:

- x is proper and surjective;
- x is birational (i.e. there exists a dense open $U \subseteq \overline{\mathcal{M}}$ such that $x^{-1}U$ is dense in X and $x^{-1}U \rightarrow U$ is an isomorphism);
- X is normal and locally desingularisable (Definition 4.1).

Admissible modifications together with maps over $\overline{\mathcal{M}}$ form a directed system. We check that \mathcal{M}^\diamond can be compactified to an admissible modification (perhaps after modification of \mathcal{M}^\diamond itself):

Lemma 6.1. There exist a proper birational map $\widetilde{\mathcal{M}}^\diamond \rightarrow \mathcal{M}^\diamond$, an admissible modification $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$, and an open immersion $\widetilde{\mathcal{M}}^\diamond \rightarrow \mathcal{M}^\diamond$ over $\overline{\mathcal{M}}$. In characteristic zero we may take $\widetilde{\mathcal{M}}^\diamond = \mathcal{M}^\diamond$.

Note that this $\widetilde{\mathcal{M}}^\diamond$ is then the largest open of \mathcal{M}^\diamond which admits a map to \mathcal{M}^\diamond over $\overline{\mathcal{M}}$.

Proof. We begin by giving a simple argument in characteristic zero. First construct some compactification \mathcal{M}^\diamond of $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ following [31, §6], then apply Hironaka [16] to see that \mathcal{M}^\diamond is desingularisable hence $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$ is an admissible modification.

In arbitrary characteristic the proof is slightly more involved. It is enough to treat the connected components of $\overline{\mathcal{M}}$ separately, so we fix one such component; abusing notation, we will still write it as $\overline{\mathcal{M}}$, and similarly write $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$.

Choose a finite cover $\mathfrak{U} = \{U_i \rightarrow \overline{\mathcal{M}}\}$ by combinatorial charts (each U_i having graph Γ_i), then choose a finite cover $\mathfrak{V} = \{V_j \rightarrow \mathfrak{U} \times_{\overline{\mathcal{M}}} \mathfrak{U}\}$ of $\mathfrak{U} \times_{\overline{\mathcal{M}}} \mathfrak{U}$ by combinatorial charts, each V_j having graph Γ'_j . For each U_i we have a fan F_i in $\mathbb{Q}^{E(\Gamma_i)}$ from Definition 3.3, and similarly for each V_j a fan F'_j in $\mathbb{Q}^{E(\Gamma'_j)}$.

Fix a V_j . Recalling that combinatorial charts are by definition connected, composing the map $V_j \rightarrow \mathfrak{U} \times_{\overline{\mathcal{M}}} \mathfrak{U}$ with one of the projections yields a map $V_j \rightarrow U_i$ for some i . This yields maps $\Gamma_i \rightarrow \Gamma'_j$, $E(\Gamma'_j) \rightarrow E(\Gamma_i)$, and $\mathbb{Q}^{E(\Gamma'_j)} \rightarrow \mathbb{Q}^{E(\Gamma_i)}$, and F_i necessarily pulls back to F'_j (we say the F_i are ‘compatible on overlaps’).

For each i we can choose a finite refinement \bar{F}_i of F_i which is complete in the sense that it fills the positive orthant. Hence if $\mathcal{M}_i^\diamond/U_i$ is the toric variety associated to \bar{F}_i then the map $\mathcal{M}_i^\diamond \rightarrow U_i$ is proper. If we write \tilde{F}_i for the fan obtained by restricting \bar{F}_i to the support of F_i then the associated toric variety $\widetilde{\mathcal{M}}^\diamond_i$ over U_i is a blowup of $\mathcal{M}^\diamond_{U_i}$ and has a natural open immersion to \mathcal{M}_i^\diamond . Since toric varieties are desingularisable we have proven the lemma ‘locally on $\overline{\mathcal{M}}$ ’.

Since the transition maps $\mathbb{Q}^{E(\Gamma'_j)} \rightarrow \mathbb{Q}^{E(\Gamma_i)}$ are just inclusions of coordinate sub-spaces, it is not hard to check that the \bar{F}_i can be chosen to be compatible on overlaps, whence they will glue to give a global construction. □

Definition 6.2. Given an admissible modification $x: X \rightarrow \overline{\mathcal{M}}$, the map σ induces a map $\sigma': X \times_{\overline{\mathcal{M}}} \mathcal{M} \rightarrow J$. We write \mathring{X} for the largest open of X over which the closure of the graph of σ' inside $X \times_{\overline{\mathcal{M}}} J$ is flat. We write $\sigma'_X: \mathring{X} \rightarrow J$ for the induced map (informally, we call \mathring{X} the ‘largest open of X over which σ extends’). Clearly \mathring{X} contains $x^{-1}\mathcal{M}$. We write e for the unit section in the jacobian, and $\text{DRL}_X := (\sigma'_X)^*e$ for the pullback as a closed subscheme of \mathring{X} .

The induced map $\sigma_X: \mathring{X} \rightarrow J \times_{\overline{\mathcal{M}}} \mathring{X}$ is then a regular closed immersion, since $J \times_{\overline{\mathcal{M}}} \mathring{X}$ is smooth over \mathring{X} . We denote by e_X the unit section $\mathring{X} \rightarrow J \times_{\overline{\mathcal{M}}} \mathring{X}$. Then we can pull back the fundamental class of e_X to a cycle class $\text{DRC}_X := \sigma'_X[e_X]$ on

$$\mathring{X} \times_{\sigma_X, J \times_{\overline{\mathcal{M}}} \mathring{X}, e} \mathring{X} = \text{DRL}_X .$$

Remark 6.3. In the above construction, we could alternatively have defined DRC_X by pulling back the class of σ_X along e_X . The resulting cycle would be the same. To see this, note first that, since σ_X and e_X are regular closed immersions, the fundamental class of \mathring{X} is exactly the refined gysin pullback (along either σ_X or e_X) of the fundamental class of $J \times_{\overline{\mathcal{M}}} \mathring{X}$. We then apply the commutativity of the intersection pairing, see [33, Theorem 3.13], cf. [10, Theorem 6.4].

Remark 6.4. In general \mathring{X} is strictly smaller than X . For example, if $X = \bar{\mathcal{M}}_{g,n}$ itself then $\mathring{\mathcal{M}}_{g,n}$ contains the compact-type (even tree-like) locus, but need not be the whole of $\bar{\mathcal{M}}_{g,n}$. If $g = 1, n = 2$, and $a_1 = -a_2 = 2$ then $\mathring{\mathcal{M}}_{g,n}$ is the complement of the boundary point corresponding to a pair of \mathbb{P}^1 s glued at 0 and ∞ . If we take X to be the blowup of $\bar{\mathcal{M}}_{g,n}$ at that point then \mathring{X} will be the complement of the two singular points in the boundary that are contained within the exceptional curve of the blowup.

Given any \mathcal{M}^\diamond as in Lemma 6.1, note that $\overset{\circ}{\mathcal{M}}^\diamond = \widetilde{\mathcal{M}}^\diamond$. The inclusion $\overset{\circ}{\mathcal{M}}^\diamond \supset \widetilde{\mathcal{M}}^\diamond$ holds because \mathcal{M}^\diamond and hence $\widetilde{\mathcal{M}}^\diamond$ are σ -extending. For the other implication, let $x \in \overset{\circ}{\mathcal{M}}^\diamond$, and let T be a trait in $\overset{\circ}{\mathcal{M}}^\diamond$ through x with generic point lying over \mathcal{M} . Since T is σ -extending it lifts to \mathcal{M}^\diamond , and because $\mathcal{M}^\diamond \rightarrow \overset{\circ}{\mathcal{M}}^\diamond$ is proper it lifts further to $\widetilde{\mathcal{M}}^\diamond$, so $x \in \widetilde{\mathcal{M}}^\diamond$.

Theorem 6.5. Choose any \mathcal{M}^\diamond as in Lemma 6.1. Let $X \rightarrow \overline{\mathcal{M}}$ be an admissible modification which factors via $\mathcal{M}^\diamond \rightarrow \overline{\mathcal{M}}$. Then $\text{DRL}_X \rightarrow \overline{\mathcal{M}}$ is proper.

Proof. Write $f: X \rightarrow \mathcal{M}^\diamond$ for the factorisation. Since \mathcal{M}^\diamond is σ -extending we see that $f^{-1}\widetilde{\mathcal{M}}^\diamond \subseteq \overset{\circ}{X}$. In fact they are equal; let $x \in \overset{\circ}{X}$, and choose a trait T in $\overset{\circ}{X}$ through x with generic point lying over \mathcal{M} . Then T is σ -extending hence lifts to \mathcal{M}^\diamond , and it lifts further to $\widetilde{\mathcal{M}}^\diamond$ since $\mathcal{M}^\diamond \rightarrow \overset{\circ}{\mathcal{M}}^\diamond$ is proper.

Write $\overset{\circ}{f}: \overset{\circ}{X} \rightarrow \mathcal{M}^\diamond$ for the restricted map (which is proper since f is proper). Then $\text{DRL}_X = \overset{\circ}{f}^{-1} \text{DRL}_\diamond$, so we have a proper map $\overset{\circ}{f}: \text{DRL}_X \rightarrow \text{DRL}_\diamond$. By Proposition 5.3 we know $\text{DRL}_\diamond \rightarrow \overline{\mathcal{M}}$ is proper, so we are done. \square

Observing that the admissible modifications which factor via \mathcal{M}^\diamond form a cofinal system among all admissible modifications, we have established Theorem 1.1.

Whenever $\text{DRL}_X \rightarrow \overline{\mathcal{M}}$ is proper, we can form the pushforward of DRC_X to $\overline{\mathcal{M}}$; write $\pi_* \text{DRC}_X$ for this cycle on $\overline{\mathcal{M}}$. We defined DRL_\diamond at the beginning of §5 as the schematic pullback of the unit section e of J along the section $\sigma'_\diamond: \mathcal{M}^\diamond \rightarrow J$, just as in Definition 6.2.

Definition 6.6. Write $\sigma_\diamond: \mathcal{M}^\diamond \rightarrow \mathcal{M}^\diamond \times_{\overline{\mathcal{M}}} J$ for the section induced by σ'_\diamond , a regular closed immersion, and write $[e_\diamond]$ for the fundamental class of \mathcal{M}^\diamond , viewed as a cycle on $\mathcal{M}^\diamond \times_{\overline{\mathcal{M}}} J$ via the map e_\diamond . Then define $\text{DRC}_\diamond := \sigma_\diamond^! [e_\diamond]$, as a class on DRL_\diamond , the refined gysin pullback in the pullback diagram

$$\begin{array}{ccc} \text{DRL}_\diamond & \xrightarrow{\sigma'} & \mathcal{M}^\diamond \\ \downarrow e' & & \downarrow e_\diamond \\ \mathcal{M}^\diamond & \xrightarrow{\sigma_\diamond} & \mathcal{M}^\diamond \times_{\overline{\mathcal{M}}} J. \end{array}$$

This is exactly the cycle as we would obtain applying Definition 6.2 with $X = \mathcal{M}^\diamond$. Since DRC_\diamond is a cycle class on DRL_\diamond , we can define $\overline{\text{DRC}}$ as the pushforward of this cycle to $\overline{\mathcal{M}}$.

Theorem 6.7. Choose any \mathcal{M}^\diamond as in Lemma 6.1. If an admissible modification $x: X \rightarrow \overline{\mathcal{M}}$ factors via \mathcal{M}^\diamond , then $\pi_* \text{DRC}_X = \overline{\text{DRC}}$.

This would be a formality if we could pull back the cycle DRC_\diamond to X , but since \mathcal{M}^\diamond and X may be singular we must take some care with pulling back cycles.

Proof. For simplicity we give the proof only in the case where $\widetilde{\mathcal{M}}^\diamond \rightarrow \mathcal{M}^\diamond$ is an isomorphism, i.e. we have an open immersion $\mathcal{M}^\diamond \rightarrow \overset{\circ}{\mathcal{M}}^\diamond$ (this can always be arranged

in characteristic zero). To extend to the general case one uses a very similar argument to show that $\pi_* \text{DRC}_\diamond = \pi_* \text{DRC}_{\mathcal{M}^\diamond}$.

As in the proof of Theorem 6.5, the factorisation $f: X \rightarrow \mathcal{M}^\diamond$ restricts to a proper map $f: \mathring{X} \rightarrow \mathcal{M}^\diamond$. Write $J_{\mathcal{M}^\diamond} := J \times_{\mathcal{M}} \mathcal{M}^\diamond$, and $J_{\mathring{X}} := J \times_{\mathcal{M}} \mathring{X}$. Then we can view $\sigma_{\mathcal{M}^\diamond}$ as a section of the projection $J_{\mathcal{M}^\diamond} \rightarrow \mathcal{M}^\diamond$, and similarly for \mathring{X} . And we write $e_{\mathcal{M}^\diamond}$ for the unit section of $J_{\mathcal{M}^\diamond}$, and similarly for \mathring{X} . Writing f_J for the induced map $J_X \rightarrow J_{\mathcal{M}^\diamond}$, we obtain a commutative diagram

$$\begin{array}{ccc}
 J_{\mathring{X}} & \xrightarrow{f_J} & J_{\mathcal{M}^\diamond} \\
 \left(\begin{array}{c} \uparrow \\ \downarrow \\ \uparrow \\ \downarrow \end{array} \right) & & \left(\begin{array}{c} \uparrow \\ \downarrow \\ \uparrow \\ \downarrow \end{array} \right) \\
 \sigma_X & & \sigma_{\mathcal{M}^\diamond} \\
 & e_X & & e_{\mathcal{M}^\diamond} \\
 X & \xrightarrow{f} & \mathcal{M}^\diamond
 \end{array}$$

Note that the upward-pointing arrows are closed immersions, so we can see them as algebraic cycles. Since f is proper and birational, we see immediately that $f_{J*}[e_X] = [e_{\mathcal{M}^\diamond}]$. By the commutativity of proper pushforward and the refined Gysin homomorphism ([33, Theorem 3.12], cf. [10, Theorem 6.2(a)]), we see that

$$\sigma_{\mathcal{M}^\diamond}^! f_{J*}[e_X] = f_* \sigma_X^![e_X]$$

hence

$$\text{DRC}_\diamond = \sigma_{\mathcal{M}^\diamond}^![e_{\mathcal{M}^\diamond}] = \sigma_{\mathcal{M}^\diamond}^! f_{J*}[e_X] = f_* \sigma_X^![e_X] = f_* \text{DRC}_X. \quad \square$$

This immediately implies Theorem 1.2, and shows moreover that the limit is given by the pushforward of DRL_\diamond .

7. Proof of Theorem 1.3

In this section, we will use results of Cavalieri, Marcus, and Wise [2, 27] to check that our double ramification cycle $\overline{\text{DRC}}$ coincides with the cycle constructed in [12] and computed in [20]. We will temporarily denote the latter cycle by DRC_{LGV} . Their construction only applies in the ‘non-twisted’ case $k = 0$, so for the remainder of this section we restrict to that case without further comment.

We begin by briefly recalling the setup and notation from [27]. We denote by J the universal (semiabelian) jacobian over $\overline{\mathcal{M}}$. Over $\overline{\mathcal{M}}$ they construct a stack $\overline{M}(\mathcal{P})$ of stable maps to (expansions of) $\mathcal{P} := [\mathbb{P}^1/\mathbb{G}_m]$. We will see in the next proposition that $\overline{M}(\mathcal{P}) \rightarrow \overline{\mathcal{M}}$ is birational, and their constructions yield a map $\overline{M}(\mathcal{P}) \rightarrow J$, extending the map σ on \mathcal{M} (if $\overline{M}(\mathcal{P})$ were normal we would call it σ -extending). Writing $e: \overline{\mathcal{M}} \rightarrow J$ for the closed immersion from the unit section ([27] denote this by Z), we define $\overline{M}(\mathcal{P}/B\mathbb{G}_m)$ to be the fibre product of e and $\overline{M}(\mathcal{P})$ over J .

Since $e: \overline{\mathcal{M}} \rightarrow J$ is a regular closed immersion (say with ideal sheaf \mathcal{I}), it comes with a natural perfect relative obstruction theory, namely the cotangent complex $L_{e/J} = \iota^* \mathcal{I}/\mathcal{I}^2[1]$ itself. The associated relative virtual class is just the fundamental class $[e]$ as a cycle on e (cf. [1, Example 7.6]).

In [2] they construct directly a perfect relative obstruction theory for $\bar{M}(\mathcal{P}/B\mathbb{G}_m) \rightarrow \bar{M}(\mathcal{P})$, and show (see [2, Remark after Proposition 3.5]) that it coincides with the pullback of $L_{e/J}$ in the fibre product diagram

$$\begin{array}{ccc} \bar{M}(\mathcal{P}/B\mathbb{G}_m) & \longrightarrow & \bar{\mathcal{M}} \\ \downarrow & & \downarrow e \\ \bar{M}(\mathcal{P}) & \longrightarrow & J. \end{array}$$

The remaining key properties we need are summarised below.

Proposition 7.1.

- (1) $\bar{M}(\mathcal{P}) \rightarrow \bar{\mathcal{M}}$ is birational.
- (2) $\bar{M}(\mathcal{P}/B\mathbb{G}_m)$ is proper over $\bar{\mathcal{M}}$.
- (3) The pushforward along $\bar{M}(\mathcal{P}/B\mathbb{G}_m) \rightarrow \bar{M}(\mathcal{P})$ of the class associated to the pullback of $L_{e/J}$ to $\bar{M}(\mathcal{P}/B\mathbb{G}_m) \rightarrow \bar{M}(\mathcal{P})$ is the double ramification cycle DRC_{LGV} of [12].

Proof.

- (1) This is easy from the construction; it can be found in [2, Proposition 3.4].
- (2) This is stated in the third bullet point on page 958 of [2] (just above proposition 3.4).
- (3) This is [27, theorem, top of p. 9]. □

Remark 7.2 (On the characteristic of the base ring). Write $\bar{M}(\mathcal{P})'$ for the normalisation of $\bar{M}(\mathcal{P})$. In this section we need that $\bar{M}(\mathcal{P})' \rightarrow \bar{M}(\mathcal{P})$ is proper, and that $\bar{M}(\mathcal{P})'$ is locally desingularisable. If we work over a field of characteristic zero both conditions clearly hold. More generally, the finiteness of normalisation holds if we work over any universally Japanese base ring, for example \mathbb{Z} . In general resolution of singularities is more difficult outside characteristic zero. However, in [28] (which appeared after the first version of the present article) it is shown that $\bar{M}(\mathcal{P})$ is log regular (at least under mild assumptions on the base scheme), hence it is locally desingularisable. With this new reference available, the comparison results in this section are valid over \mathbb{Z} .

Write $\bar{M}(\mathcal{P}/B\mathbb{G}_m)'$ for the fibre product of $\bar{M}(\mathcal{P})'$ with $\bar{M}(\mathcal{P}/B\mathbb{G}_m)$ over $\bar{M}(\mathcal{P})$. Because $\bar{M}(\mathcal{P}) \rightarrow \bar{\mathcal{M}}$ admits an extension of the section σ , the same holds for its normalisation $\bar{M}(\mathcal{P})'$, so the latter is σ -extending. Hence can apply Corollary 4.6 to obtain a canonical map $\bar{M}(\mathcal{P})' \rightarrow \mathcal{M}^\diamond$ (we cannot work directly with $\bar{M}(\mathcal{P})$ since it might not be normal). Recalling that DRL_\diamond is the pullback of the unit section e from J to \mathcal{M}^\diamond , we obtain a diagram

$$\begin{array}{ccccc} \bar{M}(\mathcal{P}/B\mathbb{G}_m)' & \xrightarrow{f_\diamond} & \text{DRL}_\diamond & \longrightarrow & \bar{\mathcal{M}} \\ \downarrow & & \downarrow & & \downarrow e \\ \bar{M}(\mathcal{P})' & \longrightarrow & \mathcal{M}^\diamond & \longrightarrow & J \end{array} \tag{18}$$

where both squares are fibre products.

The perfect relative obstruction theory $L_{e/J}$ for $e: \overline{M} \rightarrow J$ pulls back to both the other vertical arrows, yielding a class κ on DRL_\diamond , and a class D' on $\overline{M}(\mathcal{P}/B\mathbb{G}_m)'$.

Lemma 7.3. We have an equality $\kappa = \text{DRC}_\diamond$ of classes on DRL_\diamond .

Proof. We begin by recalling the construction of the class κ from the relative perfect obstruction theory of e/J . The perfect obstruction theory of e/J is the cotangent complex $L_{e/J} = t^*\mathcal{I}/\mathcal{I}^2[1]$ (where \mathcal{I} is the ideal sheaf of e in J), which we pull back to DRL_\diamond . We then embed in this the relative intrinsic normal cone of DRL_\diamond over \mathcal{M}^\diamond . This is just the classical cone (as used by Fulton for schemes, and Vistoli for stacks); following [1] we are supposed to take the stack quotient by the relative tangent bundle of \mathcal{M}^\diamond over \mathcal{M}^\diamond , but this is of course trivial. The class κ is then obtained by intersecting this cone with the zero section of the cotangent complex (a vector bundle in a single degree).

Next we compare this to our class DRC_\diamond . The key pullback square is

$$\begin{array}{ccc} \text{DRL}_\diamond & \xrightarrow{\sigma'} & \mathcal{M}^\diamond \\ \downarrow e' & & \downarrow e_\diamond \\ \mathcal{M}^\diamond & \xrightarrow{\sigma_\diamond} & \mathcal{M}^\diamond \times_{\overline{M}} J, \end{array}$$

where both σ_\diamond and e_\diamond are regular closed immersions. To define DRC_\diamond , we take $[\mathcal{M}^\diamond]$ as a class in the Chow ring of \mathcal{M}^\diamond , and then apply the refined gysin pullback $\sigma_\diamond^!$ to get a class on DRL_\diamond (this is equivalent to taking $e_\diamond^!$, see Remark 6.3). How is this refined gysin pullback defined? Using that e_\diamond is a regular closed immersion, we take its normal bundle in $\mathcal{M}^\diamond \times_{\overline{M}} J$, which is just the pullback of the normal bundle of $e: \overline{M} \rightarrow J$. We then pull back further to DRL_\diamond , and embed the normal cone of $e': \text{DRL}_\diamond \rightarrow \mathcal{M}^\diamond$ inside it, and intersect with the zero section. This is precisely the same as the definition of the class κ above, and we are done. □

The above proof was just a matter of checking that Behrend–Fantechi’s relative virtual class construction degenerates in our setting to Fulton’s intersection product. We could have saved ourselves the effort of writing this proof if we had worked throughout with the language of Behrend–Fantechi, but we wished to emphasise that in our situation this construction is essentially classical.

Now apply Costello’s theorem [4, Theorem 5.0.1] to the left hand square in the diagram (18). The bottom horizontal arrow is birational, and the top arrow is proper since both the source and target are proper over \overline{M} . Hence we find that $f_{\diamond,*}D' = \text{DRC}_\diamond$ as classes on DRL_\diamond .

We can make another commutative diagram

$$\begin{array}{ccccc} \overline{M}(\mathcal{P}/B\mathbb{G}_m)' & \xrightarrow{f} & \overline{M}(\mathcal{P}/B\mathbb{G}_m) & \longrightarrow & \overline{M} \\ \downarrow & & \downarrow & & \downarrow e \\ \overline{M}(\mathcal{P})' & \longrightarrow & \overline{M}(\mathcal{P}) & \longrightarrow & J \end{array}$$

where again both squares are fibre products. Note that $\bar{M}(\mathcal{P})' \rightarrow \bar{M}(\mathcal{P})$ is birational, since it is an isomorphism over the locus of smooth curves. Write D for the class on $\bar{M}(\mathcal{P}/BG_m)$ arising by pulling back the perfect relative obstruction theory $L_{e/J}$ from $e: \mathcal{M} \rightarrow J$. Another application of Costello’s theorem yields that $f_*D' = D$.

Theorem 7.4. We have an equality of classes on $\overline{\mathcal{M}}$:

$$\text{DRC}_{LGV} = \overline{\text{DRC}}.$$

Proof. This follows immediately from the above discussion by pushing forward in the commutative diagram

$$\begin{array}{ccc} \bar{M}(\mathcal{P}/BG_m)' & \xrightarrow{f} & \bar{M}(\mathcal{P}/BG_m) \\ \downarrow & & \downarrow \\ \text{DRL}_\diamond & \longrightarrow & \overline{\mathcal{M}}. \end{array} \quad \square$$

Theorem 1.3 follows immediately.

Acknowledgements. I am very grateful to Alessandro Chiodo for many extensive discussions and much encouragement, without which this paper would not have been written. I have also benefitted greatly while writing this paper from discussions with Robin de Jong, Bas Edixhoven, Rahul Pandharipande, Paolo Rossi, David Rydh, Johannes Schmitt, Arne Smeets, Nicolo Sibilla, Mattia Talpo, Jonathan Wise, and many others. I am also grateful to Dimitri Zvonkine for encouraging me to include the extension to the twisted case, and to Nicola Pagani for pointing out that the construction works in the Chow ring with integral (rather than just rational) coefficients. Finally, I would like to thank the anonymous referee for a very helpful report.

References

1. K. BEHREND AND B. FANTECHI, The intrinsic normal cone, *Invent. Math.* **128**(1) (1997), 45–88.
2. R. CAVALIERI, S. MARCUS AND J. WISE, Polynomial families of tautological classes on $\mathcal{M}_{g,n}^{rt}$, *J. Pure Appl. Algebra* **216** (2012), 950–981.
3. A. CHIODO, Towards an enumerative geometry of the moduli space of twisted curves and r th roots, *Compos. Math.* **144**(06) (2008), 1461–1496.
4. K. COSTELLO, Higher genus Gromov-Witten invariants as genus zero invariants of symmetric products, *Ann. of Math. (2)* **164**(2) (2006), 561–601.
5. B. DUDIN, Compactified universal jacobian and the double ramification cycle, <http://arxiv.org/abs/1505.02897> (2015).
6. C. FABER AND R. PANDHARIPANDE, Relative maps and tautological classes, *J. Eur. Math. Soc. (JEMS)* **7**(1) (2005), 13–49.
7. G. FARKAS AND R. PANDHARIPANDE, The moduli space of twisted canonical divisors, *J. Inst. Math. Jussieu* **17**(3) (2018), 615–672. With an appendix by Janda, Pandharipande, Pixton, and Zvonkine.
8. D. FERRAND, Conducteur, descente et pincement, *Bull. Soc. Math. France* **131**(4) (2003), 553–585.

9. E. FRENKEL, C. TELEMAN AND A. J. TOLLAND, Gromov-Witten gauge theory, *Adv. Math.* **288** (2016), 201–239.
10. W. FULTON, *Intersection Theory* (Springer, Berlin, 1984).
11. W. FULTON, *Introduction to Toric Varieties*, Volume 131 (Princeton University Press, 1993).
12. T. GRABER AND R. VAKIL, Relative virtual localization and vanishing of tautological classes on moduli spaces of curves, *Duke Math. J.* **130**(1) (2005), 1–37.
13. S. GRUSHEVSKY AND D. ZAKHAROV, The zero section of the universal semiabelian variety, and the double ramification cycle, *Duke Math. J.* **163**(5) (2014), 889–1070.
14. J. GUÉRÉ, A generalization of the double ramification cycle via log-geometry, <http://arxiv.org/abs/1603.09213> (2016).
15. R. HAIN, *Normal Functions and the Geometry of Moduli Spaces of Curves*, *Handbook of Moduli* (ed. G. FARKAS AND I. MORRISON), Advanced Lectures in Mathematics, Volume XXIV, Volume I (International Press, Boston, 2013).
16. H. HIRONAKA, Resolution of singularities of an algebraic variety over a field of characteristic zero: II, *Ann. of Math. (2)* **79**(2) (1964), 205–326.
17. D. HOLMES, Néron models of jacobians over base schemes of dimension greater than 1, *J. Reine Angew. Math.* **2019**(747) (2017), 109–145.
18. D. HOLMES, J. KASS AND N. PAGANI, Extending the double ramification cycle using jacobians, *Eur. J. Math.* **4**(1087) (2018), 1087–1099.
19. D. HOLMES, A. PIXTON AND J. SCHMITT, Multiplicativity of the double ramification cycle, to appear in *Documenta Mathematica* (2017), <https://arxiv.org/abs/1711.10341>.
20. F. JANDA, R. PANDHARIPANDE, A. PIXTON AND D. ZVONKINE, Double ramification cycles on the moduli spaces of curves, *Publ. Math. Inst. Hautes Études Sci.* **125**(1) (2017), 221–266.
21. J. L. KASS AND N. PAGANI, The stability space of compactified universal jacobians, to appear in *Trans. Amer. Math. Soc.* (2017), <https://doi.org/10.1090/tran/7724>.
22. F. KATO, Log smooth deformation theory, *Tohoku Math. J. (2)* **48**(3) (1996), 317–354.
23. F. KATO, Log smooth deformation and moduli of log smooth curves, *Internat. J. Math.* **11**(2) (2000), 215–232.
24. K. KATO, Toric singularities, *Amer. J. Math.* **116**(5) (1994), 1073–1099.
25. J. LI, Stable morphisms to singular schemes and relative stable morphisms, *J. Differential Geom.* **57**(3) (2001), 509–578.
26. J. LI, A degeneration formula of GW-invariants, *J. Differential Geom.* **60**(2) (2002), 199–293.
27. S. MARCUS AND J. WISE, Stable maps to rational curves and the relative jacobian, <https://arxiv.org/abs/1310.5981> (2013).
28. S. MARCUS AND J. WISE, Logarithmic compactification of the Abel-Jacobi section, <https://arxiv.org/abs/1708.04471> (2017).
29. M. C. OLSSON, Log algebraic stacks and moduli of log schemes, ProQuest LLC, Ann Arbor, MI, PhD thesis, University of California, Berkeley (2001).
30. M. RAYNAUD AND L. GRUSON, Critères de platitude et de projectivité. Techniques de “platification” d’un module, *Invent. Math.* **13** (1971), 1–89.
31. D. RYDH, Compactification of tame deligne–mumford stacks, Preprint, 2011, <https://people.kth.se/~dary/tamecompactification20110517.pdf>.
32. J. SCHMITT, Dimension theory of the moduli space of twisted k -differentials, *Doc. Math.* **23** (2018), 871–894.
33. A. VISTOLI, Intersection theory on algebraic stacks and on their moduli spaces, *Invent. Math.* **97**(3) (1989), 613–670.