

Arithmetic mirror symmetry for the 2-torus

YANKI LEKILI AND TIMOTHY PERUTZ

This paper explores a refinement of homological mirror symmetry which relates exact symplectic topology to arithmetic algebraic geometry. We establish a derived equivalence of the Fukaya category of the 2-torus, relative to a basepoint, with the category of perfect complexes of coherent sheaves on the Tate curve over the formal disc $\mathrm{Spec} \mathbb{Z} \llbracket q \rrbracket$. It specializes to a derived equivalence, over \mathbb{Z} , of the Fukaya category of the punctured torus with perfect complexes on the curve $y^2 + xy = x^3$ over $\mathrm{Spec} \mathbb{Z}$, the central fibre of the Tate curve; and, over the ‘punctured disc’ $\mathrm{Spec} \mathbb{Z} \llbracket q \rrbracket$, to an integral refinement of the known statement of homological mirror symmetry for the 2-torus. We also prove that the wrapped Fukaya category of the punctured torus is derived-equivalent over \mathbb{Z} to coherent sheaves on the central fiber of the Tate curve.

1 Introduction

This paper explores a basic case of what we believe is a general connection between exact Lagrangian submanifolds in the complement to an ample divisor D in a complex Calabi–Yau manifold X —we view $X \setminus D$ as an exact symplectic manifold—and coherent sheaves on a scheme *defined over* $\mathrm{Spec} \mathbb{Z}$, the ‘mirror’ to $X \setminus D$. We take X to be an elliptic curve; its complex structure is irrelevant, so it is really a 2-torus T . We take D to be a point z . The mirror is the Weierstrass cubic $Y^2Z + XYZ = X^3$, the restriction to $q = 0$ of the Tate curve $\mathcal{T} \rightarrow \mathrm{Spec} \mathbb{Z} \llbracket q \rrbracket$.

Kontsevich’s 1994 homological mirror symmetry (HMS) conjecture [31] claims that the Fukaya A_∞ -category $\mathcal{F}(X)$ of a polarized Calabi–Yau manifold should have a formal enlargement—precisely formulated a little later as the closure $\mathrm{tw}^\pi \mathcal{F}(X)$ under taking mapping cones and passing to idempotent summands—which is A_∞ -quasi-equivalent to a dg enhancement for the derived category of coherent sheaves on the ‘mirror’ \check{X} , a Calabi–Yau variety over the field of complex Novikov series.¹ The HMS conjecture has inspired a great deal of work in symplectic geometry, algebraic geometry and mathematical physics; the HMS paradigm has been adapted so as to apply not only to varieties whose canonical bundle \mathcal{K} is trivial, but also to those where either \mathcal{K}^{-1} or

¹Beware: the circumstances under which one expects to find such an \check{X} are more subtle than those claimed by our one-sentence précis of Kontsevich’s conjecture.

\mathcal{K} is ample, with such varieties playing either symplectic or algebro-geometric roles. Meanwhile, progress on the original case of Calabi–Yau manifolds has been slow. There are currently complete mirror-symmetric descriptions of the Fukaya category only for the 2-torus $\mathbb{R}^2/\mathbb{Z}^2$ and of the square 4-torus $\mathbb{R}^4/\mathbb{Z}^4$ [6]. The case of Calabi–Yau hypersurfaces in projective space has been solved up to a certain ambiguity in identifying the mirror variety [47, 54]. There are significant partial results for linear symplectic tori of arbitrary dimension [33].

Our contention is that even in the solved cases, there is more to be said about HMS. The Fukaya category for the 2-torus has a natural model which is defined over $\mathbb{Z}[[q]]$, a subring of the complex Novikov field. This model has a mirror-symmetric description as the perfect complexes on the Tate curve \mathcal{T} over $\mathbb{Z}[[q]]$. The symplectic geometry of the torus is thereby connected to the arithmetic algebraic geometry of \mathcal{T} . Establishing this connection is the task of this article.

Experts have certainly been aware that, in principle, homological mirror symmetry should have an arithmetic-geometric dimension (cf. Kontsevich’s lecture [34], for instance), but we believe that this article is the first to treat this idea in detail. Whilst refining existing proofs of HMS for the 2-torus might be a viable option, our method is also new: we identify a generating subalgebra \mathcal{A} of the Fukaya category, and show that Weierstrass cubic curves precisely parametrize the possible A_∞ -structures on it (Theorem C). The mirror to (T, z) is then the unique Weierstrass curve corresponding to the A_∞ -structure belonging to the Fukaya category. Our identification of this mirror parallels an argument of Gross [23] but also has a novel aspect, relating the multiplication rules for theta-functions on the Tate curve to *counts of lattice points* in triangles (not areas of triangles). Our identification of the *wrapped* Fukaya category of the punctured torus with coherent complexes on $\mathcal{T}|_{q=0}$ appears to be a basic case of an unexplored aspect of mirror symmetry for Calabi–Yau manifolds.

1.1 Statement

Let T be a closed, orientable surface of genus 1; ω a symplectic form on T ; $z \in T$ a basepoint; $T_0 = T \setminus \{z\}$; and θ a primitive for ω on T_0 . Fix also a grading for the symplectic manifold T , that is, an unoriented line-field ℓ . These data suffice to specify the relative Fukaya category $\mathcal{F}(T, z)$ up to quasi-isomorphism. It is an A_∞ -category linear over $\mathbb{Z}[[q]]$ whose objects are embedded circles $\gamma \subset T_0$ which are exact ($\int_\gamma \theta = 0$) and are equipped with orientations, double covers $\tilde{\gamma} \rightarrow \gamma$ and gradings (a grading is a homotopy from $\ell|_\gamma$ to $T\gamma$ in $T(T_0)|_\gamma$).

Let $\mathcal{T} \rightarrow \text{Spec } \mathbb{Z}[[q]]$ denote the Tate curve, the cubic curve in $\mathbb{P}^2(\mathbb{Z}[[q]])$ with equation

$$(1) \quad Y^2Z + XYZ = X^3 + a_4(q)XZ^2 + a_6(q)Z^3,$$

where

$$(2) \quad a_4(q) = -5 \sum_{n>0} \frac{n^3 q^n}{1 - q^n}, \quad a_6(q) = -\frac{1}{12} \sum_{n>0} \frac{(5n^3 + 7n^5) q^n}{1 - q^n}$$

(note that $n^2(5 + 7n^2)$ is always divisible by 12).

Let $\text{vect}(\mathcal{T})$ denote the $\mathbb{Z}[[q]]$ -linear differential graded (dg) category whose objects are locally free sheaves of finite rank over \mathcal{T} , and whose morphism spaces are Čech complexes with respect to a fixed affine open cover: $\text{hom}_{\text{vect}(\mathcal{T})}(\mathcal{E}, \mathcal{F}) = \check{C}^\bullet(\text{Hom}(\mathcal{E}, \mathcal{F}))$.

Before we state the theorem, we caution that when working with (cohomologically unital) A_∞ -categories over a base ring \mathbb{K} , one should take care to demarcate different notions of “equivalence” for such categories, notions which over fields are indistinguishable. An A_∞ -functor $\psi: \mathcal{A} \rightarrow \mathcal{B}$ is said to be *quasi-invertible* if there is an A_∞ -functor $\phi: \mathcal{B} \rightarrow \mathcal{A}$ such that $\phi \circ \psi$ is isomorphic to $\text{id}_{\mathcal{A}}$ as objects of $H^0(\text{fun}(\mathcal{A}, \mathcal{A}))$ (here $\text{fun}(\mathcal{A}, \mathcal{A})$ is the A_∞ -category of cohomologically unital functors $\mathcal{A} \rightarrow \mathcal{A}$); and similarly for $\psi \circ \phi$. A weaker condition is that ψ is a *quasi-equivalence*, meaning that $H\psi: H\mathcal{A} \rightarrow H\mathcal{B}$ is an equivalence of \mathbb{K} -linear categories.

Theorem A *A choice of basis (α, β) for $H_1(T)$, with $\alpha \cdot \beta = 1$, determines a $\mathbb{Z}[[q]]$ -linear A_∞ -functor*

$$\psi: \mathcal{F}(T, z) \rightarrow \text{tw}(\text{vect}(\mathcal{T}))$$

from the relative Fukaya category to the dg category of twisted complexes in $\text{vect}(\mathcal{T})$. The functor ψ is canonical up to isomorphism in $H^0(\text{fun}(\mathcal{F}(T, z), \text{tw}(\text{vect}(\mathcal{T})))$. Moreover,

- (i) *the functor ψ maps an object $L_0^\#$ representing β to the structure sheaf \mathcal{O} . It maps an object $L_\infty^\#$ representing α to the complex $[\mathcal{O} \rightarrow \mathcal{O}(\sigma)]$, where $\sigma = [0 : 1 : 0]$ is the section at infinity of \mathcal{T} , and the map is the inclusion. (This complex is quasi-isomorphic to the skyscraper sheaf $\mathcal{O}_\sigma = \sigma_* \mathcal{O}_{\text{Spec } \mathbb{Z}[[q]]}$ at the section at infinity.) It is an embedding on the full subcategory \mathcal{A} on $\{L_0^\#, L_\infty^\#\}$; and is characterized, up to isomorphism in $H^0(\text{fun}(\mathcal{F}(T, z), \text{tw}(\text{vect}(\mathcal{T})))$, by its restriction to \mathcal{A} . See Figure 1.*

- (ii) *ψ extends to a quasi-invertible A_∞ -functor*

$$\psi: \text{tw}^\pi \mathcal{F}(T, z) \rightarrow \text{tw}(\text{vect}(\mathcal{T}))$$

Thus one has an equivalence of triangulated categories

$$D^\pi \mathcal{F}(T, z) \rightarrow \text{Perf}(\mathcal{T}) \simeq H^0(\text{tw } \text{vect}(\mathcal{T}))$$

from the split-closed derived Fukaya category to the triangulated category of perfect complexes on \mathcal{T} .

(iii) The specialization of ψ to $q = 0$ is a \mathbb{Z} -linear functor

$$\psi_0: \mathcal{F}(T_0) \rightarrow \text{tw vect}(\mathcal{T}|_{q=0})$$

from the exact Fukaya category of (T_0, θ) to the category of perfect complexes on the central fiber of the Tate curve, inducing an equivalence on derived categories

$$D\psi_0: D\mathcal{F}(T_0) \rightarrow \text{Perf}(\mathcal{T}|_{q=0})$$

(both of these derived categories are already split-closed).

(iv) $D\psi_0$ extends to an equivalence of triangulated categories

$$D\mathcal{W}(T_0) \rightarrow D^b \text{Coh}(\mathcal{T}|_{q=0})$$

from the derived wrapped Fukaya category to the bounded derived category of coherent sheaves on $\mathcal{T}|_{q=0}$ (these derived categories are again split-closed).

Remark. The functor ψ has an additional property, which is that it is ‘trace-preserving’, in a sense to be discussed later.

Clause (ii) has the following corollary:

Corollary 1.1 *There is an A_∞ quasi-equivalence $\text{QC}(\mathcal{T}) \rightarrow \text{Mod-}\mathcal{F}(T, z)$ from a DG enhancement of the derived category of unbounded quasi-coherent complexes on the Tate curve to the category of cohomologically unital $\mathcal{F}(T, z)$ -modules.*

Indeed, the functor $\text{QC}(\mathcal{T}) \rightarrow \text{Mod-}\text{vect}(\mathcal{T})$ (the restriction of the Yoneda embedding for $\text{QC}(\mathcal{T})$ to $\text{vect}(\mathcal{T})$) is a quasi-equivalence, as an instance of the general theory of [60] or [9].

Comparison to the standard formulation. The A_∞ -structure in the ‘relative’ Fukaya category $\mathcal{F}(T, z)$ is based on counting holomorphic polygons weighted by powers q^s , where s counts how many times the polygon passes through the basepoint z . The ‘absolute’ Fukaya category $\mathcal{F}(T)$, in the version most popular for mirror symmetry, has as objects Lagrangian branes $L^\#$ in T equipped with $U(1)$ local systems $E \rightarrow L$. In the latter version, holomorphic polygons are weighted by (holonomy) q^{area} . The coefficient-ring for $\mathcal{F}(T)$ is usually taken to be $\Lambda_{\mathbb{C}}$, the field of complex Novikov series $\sum_{k>0} a_k q^{r_k}$: here $a_k \in \mathbb{C}$, $r_k \in \mathbb{R}$, and $r_k \rightarrow \infty$.

To explain the relation between the relative and absolute versions, note first that there is an equation of currents $\omega = \delta_D + d\Theta$, where Θ is a 1-current. We take θ to be the (smooth) restriction of Θ to M .




Figure 1: The torus T and the mirror correspondence ψ , for one possible choice of the line field ℓ .

Lemma 1.2 *There is a fully faithful ‘inclusion’ functor*

$$e: \mathcal{F}(T, z) \otimes_{\mathbb{Z}[[q]]} \Lambda_{\mathbb{C}} \rightarrow \mathcal{F}(T),$$

linear over $\Lambda_{\mathbb{C}}$ and acting as the identity on objects. For each exact Lagrangian L , select a function $K_L \in C^\infty(L)$ such that $dK_L = \theta|_L$. Then define e on morphism-spaces $\text{hom}(L_0^\#, L_1^\#) = CF(\phi(L_0^\#), L_1^\#)$ by

$$e(x) = q^{A(x)}x, \quad x \in \phi(L_0) \cap L_1,$$

where $A(x) = A_{\phi(L_0), L_1}(x)$ is the symplectic action, defined via the K_L , and ϕ is the exact symplectomorphism used to obtain transversality. The higher A_∞ -terms for e are identically zero.

Proof The symplectic action is defined as follows. For a path $\gamma: ([0, 1]; 0, 1) \rightarrow (M; L_0, L_1)$ (for instance, a constant path at an intersection point) we put

$$A_{L_0, L_1}(\gamma) = - \int_0^1 \gamma^* \theta - K_{L_0}(\gamma(0)) + K_{L_1}(\gamma(1)).$$

For any disc $u: (D, \partial D) \rightarrow (X, L)$, we have

$$\int_D u^* \omega - D \cdot u = \int_D u^* (\omega - \delta_D) = \int_{\partial D} u|_{\partial D}^* \theta = \int_{\partial D} d(u|_{\partial D}^* K_L) = 0.$$

Similarly, if $u: D \rightarrow X$ is a polygon attached to a sequence of Lagrangians (L_0, L_1, \dots, L_d) (where $d \geq 1$) at corners $x_1 \in L_0 \cap L_1, \dots, x_{d+1} \in L_d \cap L_0$, then

$$\int_D u^* \omega - D \cdot u = \int_D u^* (\omega - \delta_D) = A_{L_{d+1}, L_0}(x_{d+1}) + \sum_{i=1}^{d+1} A_{L_{i-1}, L_i}(x_i).$$

From this it follows that $e \circ \mu_{\mathcal{F}(T, z)}^d(x_1, \dots, x_d) = \mu_{\mathcal{F}(T)}^d \circ (ex_1, \dots, ex_d)$, which proves that e is a functor. Note that the perturbations that are used to define hom-spaces in $\mathcal{F}(T, z)$ serve equally well in $\mathcal{F}(T)$. It is clear that e is fully faithful. \square

The ‘standard’ statement of mirror symmetry is as follows. Let $\mathcal{T}_{\Lambda_{\mathbb{C}}} = \mathcal{T} \times_{\mathbb{Z}[[q]]} \Lambda_{\mathbb{C}}$; it is an elliptic curve over the field $\Lambda_{\mathbb{C}}$. When ω is normalized so that $\int_T \omega = 1$, there is a functor

$$\Phi: \mathcal{F}(T) \rightarrow \tilde{D}^b \text{Coh}(\mathcal{T}_{\Lambda_{\mathbb{C}}}),$$

where $\tilde{D}^b \text{Coh}$ is the unique dg enhancement of the bounded derived category $D^b \text{Coh}$ [36], inducing a derived equivalence; and that this functor is again canonically determined by a choice of basis for $H_1(T)$: see [42, 40, 41, 6] for one proof; [23] for an expository account of another, occasionally missing technical details (e.g. certain signs); and [52] for yet another. Our result is an arithmetic refinement of this standard one:

Theorem 1.3 *The diagram*

$$\begin{array}{ccc} \mathcal{F}(T, z) \otimes \Lambda_{\mathbb{C}} & \xrightarrow{\psi \otimes 1} & \text{tw vect}(\mathcal{T}) \otimes \Lambda_{\mathbb{C}} \\ e \downarrow & & \downarrow i \\ \mathcal{F}(T) & \xrightarrow{\Phi} & \tilde{D}^b \text{Coh}(\mathcal{T}_{\Lambda_{\mathbb{C}}}). \end{array}$$

is homotopy-commutative under composition of A_{∞} -functors.

Since $\mathcal{T} \times_{\mathbb{Z}[[q]]} \Lambda_{\mathbb{C}}$ is a non-singular variety over the field $\Lambda_{\mathbb{C}}$, we may take $\text{tw vect}(\mathcal{T}_{\Lambda_{\mathbb{C}}})$ as our dg enhancement of $D \text{Coh}(\mathcal{T}_{\Lambda_{\mathbb{C}}})$. Then i is the base-change functor $\text{tw vect}(\mathcal{T}) \rightarrow \text{tw vect}(\mathcal{T}_{\Lambda_{\mathbb{C}}})$. For this theorem to make sense, ψ and Φ must be set up so that $i \circ (\psi \otimes 1)$ and $\Phi \circ e$ agree precisely (not just up to quasi-isomorphism) on objects.

1.2 The Tate curve

Useful references for this material include [57, 26, 13, 23]. The *Tate curve* is the plane projective curve \mathcal{T} over $\mathbb{Z}[[q]]$ whose affine equation is the Weierstrass cubic

$$(3) \quad y^2 + xy = x^3 + a_4x + a_6,$$

where a_4 and a_6 are as at (2). So \mathcal{T} is a projective curve in $\mathbb{P}^2(\mathbb{Z}[[q]])$. Like any Weierstrass curve $w(x, y) = 0$, \mathcal{T} comes with a canonical differential with poles at the singularities,

$$\Omega = dx/w_y = -dy/w_x = dx/(2y + x) = -dy/(y - 3x^2 - a_4).$$

Notation:

$$(4) \quad \hat{\mathcal{T}} = \mathcal{T} \text{ specialized to } \mathbb{Z}((q)) (= \mathbb{Z}[[q]][q^{-1}]).$$

The analytic significance of the Tate curve is the following. Define $q = \exp(2\pi i\tau)$, where $\text{Im } \tau > 0$. As τ ranges over the upper half-plane, q ranges over the punctured unit disc \mathbb{D}^* . Consider the Riemann surface $E_\tau = \mathbb{C}/\langle 1, \tau \rangle$. The exponential map $z \mapsto \exp(2\pi iz)$ identifies E_τ with $\mathbb{C}^*/q^\mathbb{Z}$. As q varies over \mathbb{D}^* , the Riemann surfaces $\mathbb{C}^*/q^\mathbb{Z}$ form a holomorphic family $\mathbf{E} \rightarrow \mathbb{D}^*$. The Weierstrass function \wp_q for the modular parameter q defines an embedding

$$\mathbf{E} \rightarrow \mathbb{CP}^2 \times \mathbb{D}^*; \quad (z, q) \mapsto [(2\pi i)^{-2}\wp_q(z) : (2\pi i)^{-3}\wp'_q(z) : 1], q).$$

This embedding is cut out by an equation $y^2 = 4x^3 - g_2(q)x - g_3(q)$, which is a Weierstrass cubic in (x, y) varying holomorphically with q . The functions g_2 and g_3 are holomorphic at $q = 0$, and moreover are defined over $\mathbb{Z}[\frac{1}{6}][[q]]$ (making this so is the purpose of the powers of $2\pi i$ in the definition of the embedding). We can change coordinates, writing $x' = x - \frac{1}{12}$ and $2y' + x' = y$, so as to put the equation in the form $y'^2 + x'y' = x'^3 + a_4(q)x' + a_6(q)$. The benefit of the coordinate-change is that the coefficients now lie in $\mathbb{Z}[[q]]$. The series a_4 and a_6 are those given above—so the algebraic curve $y'^2 + x'y' = x'^3 + a_4(q)x' + a_6(q)$ is the Tate curve \mathcal{T} .

The series $a_4(q)$ and $a_6(q)$ have radius of convergence 1 (this is implicit in the previous paragraph, but is easy to check directly). By evaluating these series, we can interpret the defining equation of \mathcal{T} not as the equation of a scheme over $\mathbb{Z}[[q]]$, but rather, that of an analytic space $\mathcal{T}^{\text{an}} \subset \mathbb{CP}^2 \times \mathbb{D}$. Projection on the second factor defines a proper holomorphic map $\mathcal{T}^{\text{an}} \rightarrow \mathbb{D}$. The previous paragraph tells us that this holomorphic family of curves, when restricted to \mathbb{D}^* , is analytically isomorphic to the family of elliptic curves $\mathbb{Z}((q))^*/q^\mathbb{Z}$ for $q \in \mathbb{D}^*$.

One can also specialize \mathcal{T} to $q = 0$. The result is the curve $\mathcal{T}_0 = \mathcal{T}|_{q=0}$ in $\mathbb{P}^2(\mathbb{Z})$ given by

$$(5) \quad y^2 + xy = x^3.$$

We can characterize this Weierstrass curve as follows:

Lemma 1.4 *The curve $\mathcal{T}_0 \rightarrow \operatorname{Spec} \mathbb{Z}$ has a section $s = [0 : 0 : 1]$ which is a node of $\mathcal{T}_0 \times_{\mathbb{Z}} \mathbb{F}_p$, the mod p reduction of \mathcal{T}_0 , for every prime p . Any Weierstrass curve $C \rightarrow \operatorname{Spec} \mathbb{Z}$ possessing a section s with this property can be transformed by integral changes of variable to \mathcal{T}_0 .*

Proof Consider a general Weierstrass curve $C = [a_1, a_2, a_3, a_4, a_6]$, given as the projective closure of

$$(6) \quad y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6, \quad a_i \in \mathbb{Z}.$$

Integral points of $C \subset \mathbb{P}_{\mathbb{Z}}^2$, other than $[0 : 1 : 0]$, can be represented as rational points on the affine curve. The point $[0 : 1 : 0]$ is regular over any field, and is the unique point of C with $Z = 0$. Suppose $[X : Y : Z]$ is an integral point that is nodal mod p for all primes p . Then Z must be non-zero mod p for every prime p , and hence Z is a unit of \mathbb{Z} . Consider the \mathbb{Z} -point $(x_0, y_0) = (X/Z, Y/Z)$ of the affine curve. The partial derivatives vanish, since they vanish mod p for all p :

$$(7) \quad 2y_0 + a_1x_0 + a_3 = 0, \quad a_1y_0 = 3x_0^2 + 2a_2x_0 + a_4.$$

The nodal condition is that the Hessian is non-singular, that is,

$$(8) \quad a_1^2 + 2(6x_0 + 2a_2) \neq 0 \pmod{p}.$$

(We note in passing that conditions (7, 8) hold for the point $[0 : 0 : 1]$ of $\mathcal{T}|_0$ at all primes p .) Since (8) holds for all p , we have

$$(9) \quad a_1^2 + 12x_0 + 4a_2 = \pm 1.$$

We shall use the criterion (9) to make three changes of variable, successively making a_1 , a_2 and a_3 equal to their counterparts for \mathcal{T}_0 .

First, (9) tells us that a_1 is odd. Hence by a change of variable $x = x'$, $y = y' + c$, we may assume that $a_1 = 1$, whereupon $6x_0 + 2a_2$ is either 0 or -1 . The latter possibility is absurd, so $3x_0 + a_2 = 0$. Being divisible by 3, a_2 can be removed altogether by a change of variable $x = x' + dy$, $y = y'$ without interfering with a_1 . Thus we can assume additionally that $a_2 = 0$. We now find from (9) that $x_0 = 0$. Hence $2y_0 + a_3 = 0$, so a_3 is even. It follows that a_3 can be set to zero by a change of variable $x = x'$, $y = y' + e$, leaving a_1 and a_2 untouched. Equations (7) now tell us that $y_0 = 0 = a_4$, while the equation (6) for C tells us that $a_6 = 0$. \square

More abstractly, if we define a curve $\pi: C \rightarrow \operatorname{Spec} \mathbb{Z}$ by taking $\mathbb{P}_{\mathbb{Z}}^1$ and identifying the sections $[0 : 1]$ and $[1 : 1]$, so as to make every geometric fiber nodal, then the parametrization $\mathbb{P}_{\mathbb{Z}}^1 \rightarrow \mathbb{P}_{\mathbb{Z}}^2$ given by $[s : t] \mapsto [st(s - t) : s(s - t)^2 : t^3]$ identifies C with \mathcal{T}_0 .

Outline of method and algebraic results. This article is long partly because it contains rather more than a single proof of Theorem A, and partly because working over \mathbb{Z} presents significant technicalities beyond those that would be present if one worked over fields (or in some cases, of fields in which 6 is invertible). Part I—a large chunk—is purely algebraic; it refines and elaborates the method of [35]. The basic point is that for any Weierstrass curve C , one has a 2-object subcategory \mathcal{B}_C of $\text{Perf } C$ —the dg category of perfect complexes of coherent sheaves—with objects \mathcal{O} (the structure sheaf) and \mathcal{O}_p (the skyscraper sheaf at the point at infinity), and this subcategory split-generates $\text{Perf } C$. The cohomology category $A = H^*\mathcal{B}_C$ is independent of C , but the dg structure of \mathcal{B}_C knows C . One can transfer the dg structure to a minimal A_∞ -structure on A . This procedure defines a functor from the category of Weierstrass curves to the category of minimal A_∞ -structures on A . We prove in Theorem C that this functor is an equivalence. A slightly coarsened statement of Theorem C is as follows:

Theorem 1.5 *Let R be an integral domain which is either noetherian and normal of characteristic zero, or an arbitrary field. Let $(\mathcal{B}, \mu_{\mathcal{B}}^*)$ be an R -linear A_∞ -category together with a Calabi–Yau structure of dimension 1. Assume that \mathcal{B} is minimal, has just two objects a and b , both spherical of dimension 1 and forming an A_2 -chain (i.e. $\text{hom}(a, a) \cong \Lambda^*(R[-1]) \cong \text{hom}(b, b)$ as graded R -algebras; and $\text{hom}(a, b) \cong R$, $\text{hom}(b, a) \cong R[-1]$ as graded R -modules; and $\mu_{\mathcal{B}}^1 = 0$). Then \mathcal{B} is trace-preservingly quasi-equivalent to \mathcal{B}_C for a unique Weierstrass curve $C \rightarrow \text{Spec } R$, where \mathcal{B}_C has the Calabi–Yau structure arising from its Weierstrass differential $\Omega \in \Omega_{C/\text{Spec } R}^1$.*

The proof of Theorem C invokes the Hochschild cohomology $\text{HH}^*(A, A)$. We computed this cohomology additively in [35], but here we give a complete calculation, as a Gerstenhaber algebra, by interpreting $\text{HH}^*(A, A)$ as the Hochschild cohomology $\text{HH}^*(C_{\text{cusp}})$ of a cuspidal Weierstrass curve C_{cusp} (Theorem B).

In Part II, we identify the unique curve C_{mirror} for which $\mathcal{A}_{C_{\text{mirror}}}$ is quasi-isomorphic to the 2-object subcategory $\mathcal{A}_{\text{symp}}$ of the Fukaya category $\mathcal{F}(T, z)$ on objects of slopes 0 and $-\infty$, equipped with non-trivial double coverings. In [35], we used Abouzaid’s plumbing model [3] to prove that $\mathcal{A}_{\text{symp}}|_{q=0}$ is not formal, which implies that C_{mirror} is not cuspidal. Here we identify C_{mirror} precisely. In fact, we identify the specialization $C_{\text{mirror}}|_{q=0}$ in three independent ways: (i) by eliminating the possibility that C_{mirror} is smooth or cuspidal after reduction to an arbitrary prime p , by means of the ‘closed open string map’ from symplectic cohomology to Hochschild cohomology of the Fukaya category; (ii) by calculating “Seidel’s mirror map” [65], or more precisely, by determining the affine coordinate ring of $C_{\text{mirror}}|_{q=0}$ via a calculation in the exact Fukaya category; and (iii) via theta-functions. The third proof extends to a proof of mirror symmetry for $\mathcal{F}(T, z)$, not just its restriction to $q = 0$. We use an intrinsic model

for the Tate curve, and the integral theta-functions for this curve which played a major role in Gross’s proof [23]. The nub is the multiplication rule for these theta-functions and its relation to counts of lattice-points in triangles. The proof of mirror symmetry for the wrapped category is a rather formal extension of that for the exact category.

We should perhaps make one more remark about exposition. The authors’ background is in symplectic topology. We imagine that typical readers will have an interest in mirror symmetry, perhaps with a bias towards the symplectic, algebro-geometric or physical aspects, but, like us, will not be expert in arithmetic geometry. We would be delighted to have readers who *do* come from an arithmetic geometry background, but ask for their patience in an exposition which we fear belabors what is obvious to them and rushes through what is not.

Higher dimensions? We believe that there should be an arithmetic refinement to homological mirror symmetry for Calabi–Yau manifolds in higher dimensions, but will leave the formulation of such conjectures for elsewhere; the 2-torus is, we think, far from being an isolated case. The case of 2-tori with several basepoints can be treated rather straightforwardly starting from the one-pointed case, but we shall also leave that for another article.

Acknowledgements. YL was partly supported by the Herchel Smith Fund and by Marie Curie grant EU-FP7-268389; TP by NSF grant DMS 1049313. Paul Seidel provided notes outlining the algebro-geometric approach to computing Hochschild cohomology for A . Conversations with several mathematicians proved useful as this work evolved, and we would particularly like to thank Mohammed Abouzaid, David Ben-Zvi, Mirela Cîrîpăni, Brian Conrad, Kevin Costello, and Paul Seidel. We thank the Simons Center for Geometry and Physics for its generous hospitality.

Part I

Algebraic aspects

2 Background material

2.1 A_∞ -categories over commutative rings

Our basic reference for A_∞ -categories \mathcal{C} is Chapter 1 of Seidel’s book [50]. However, Seidel’s base ring \mathbb{K} is always a field, while we want to take for \mathbb{K} a commutative ring

with unit.

There is a sense in which to work with A_∞ -categories over non-fields is to go against the grain: part of the philosophy of A_∞ -categories is that in the A_∞ world ‘everything is derived’: for instance, if \mathbb{K} is a field, and M_1 and M_2 modules over an associative \mathbb{K} -algebra A , viewed as an A_∞ -algebra, the A_∞ -morphisms $M_1 \rightarrow M_2$ —more precisely, the space $\mathrm{Hom}_{H(\mathcal{C} \text{ mod})}(M_1, M_2)$ —coincides with $\mathrm{Ext}_A^*(M_1, M_2)$. Principles such as this one fail when one works over non-fields. The paper [?] provides a notion of a ‘derived A_∞ -category’ over a general base \mathbb{K} .

Nonetheless, Fukaya categories, over fields or not, are A_∞ -categories, and it is natural to ask to what extent one can make the theory work. A good deal of the basic algebra of A_∞ -categories does still work, enough for us to get a satisfactory result about the 2-torus. In this section, we run through the basics, noting what carries through and what does not, and in the latter case explaining partial results that will help us.

2.1.1 Basic notions

An A_∞ -category \mathcal{C} has a set $\mathrm{ob} \mathcal{C}$ of objects, and for any pair (X_0, X_1) of objects a morphism-space $\mathrm{hom}_{\mathcal{C}}(X_0, X_1)$ which is a graded \mathbb{K} -module. The multilinear operations

$$\mu_{\mathcal{C}}^d: \mathrm{hom}_{\mathcal{C}}(X_{d-1}, X_d) \otimes \cdots \otimes \mathrm{hom}_{\mathcal{C}}(X_1, X_2) \otimes \mathrm{hom}_{\mathcal{C}}(X_0, X_1) \rightarrow \mathrm{hom}_{\mathcal{C}}(X_0, X_d)[2-d]$$

satisfy the A_∞ -equations [50, (1.2)]. In particular, $\mu_{\mathcal{C}}^1$ is a degree 1 endomorphism of $\mathrm{hom}_{\mathcal{C}}(X_0, X_1)$ and $\mu_{\mathcal{C}}^1 \circ \mu_{\mathcal{C}}^1 = 0$. The cohomology spaces

$$\mathrm{Hom}_{\mathcal{C}}(X_0, X_1) := H(\mathrm{hom}_{\mathcal{C}}(X_0, X_1), \mu_{\mathcal{C}}^1)$$

form the morphism spaces in the *cohomology category* $H\mathcal{C}$, a graded \mathbb{K} -linear category. The composition in this category is given by $[a_2] \cdot [a_1] = (-1)^{|a_1|} [\mu_{\mathcal{C}}^2(a_2, a_1)]$, the sign-change designed to force associativity [50, (1.3)]. The existence of identity morphisms $\mathrm{id}_X \in H\mathrm{hom}_{\mathcal{C}}(X, X)$, for each object X , is not automatic, but is instead taken as an assumption—that is, our A_∞ -categories are assumed *cohomologically unital*, or *c-unital* for short.

There is a notion of A_∞ -functor [50, (1b)] $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$, and a strictly associative composition operation for functors [50, (1.8)]. There is a notion A_∞ -(pre)-natural transformation [50, (1d)], and an A_∞ -category $\mathrm{nu} \mathrm{fun}(\mathcal{A}, \mathcal{B})$ whose objects are A_∞ -functors $\mathcal{A} \rightarrow \mathcal{B}$ and whose morphism-spaces are of pre-natural transformations. Here ‘nu fun’ is for ‘non-unital functor’; there is a full subcategory $\mathrm{fun}(\mathcal{A}, \mathcal{B})$ of c-unital functors. There are notions of *formal diffeomorphism* [50, (1c)], *composition functors* [50, (1e)], as well as of A_∞ -module. One has *Yoneda modules* $\mathcal{Y}(X) = \mathrm{hom}_{\mathcal{A}}(X, \cdot)$

and the resulting Yoneda embedding $\mathcal{A} \rightarrow \text{mod-}\mathcal{A}$ [50, (1k)]. All these notions, being based on algebraic operations and polynomial relations, make equal sense over commutative rings.

2.1.2 Comparison arguments

Comparison arguments for spectral sequences of bounded filtrations are a basic tool of the trade. We wish to work with non-convergent spectral sequences, and shall use the following statement:

Lemma 2.1 *Suppose that A^\bullet and B^\bullet are cochain complexes equipped with decreasing filtrations $A^\bullet = F^0 A^\bullet \supset F^1 A^\bullet \supset \dots$, and $B^\bullet = F^0 B^\bullet \supset F^1 B^\bullet \supset \dots$ that are complete and bounded above by F^0 . Suppose that $f: A^\bullet \rightarrow B^\bullet$ is a map of filtered complexes which induces an isomorphism on the r th page of the spectral sequences of the filtrations:*

$$f_r: E_r^{\bullet\bullet}(A) \xrightarrow{\cong} E_r^{\bullet\bullet}(B).$$

Then f is a quasi-isomorphism.

A filtration $C = \dots \supset F^{-1}C \supset F^0C \supset F^1C \supset \dots$ is *complete* if the natural map $C \rightarrow \varprojlim C/F^p C$ is an isomorphism. It is *bounded above by F^k* if $C = F^k C$.

Proof This is [?, Theorem 7.4] (in fact Eilenberg–Moore show that one can replace boundedness from above by exhaustiveness of the filtration). Alternatively, one can infer the result from material in Weibel’s book [62], as follows: The mapping cone $C = \text{cone } f$ acquires a filtration from those of A and B , and $E_r^{\bullet\bullet}(C) = 0$ [62, exercise 5.4.4]. The filtration on C , like those of A and B , is complete and bounded above; unlike the spectral sequences for the filtrations on A and B , the spectral sequence for the filtration of C is *regular*, i.e., for each (p, q) one has $d_s^{pq} = 0$ for $s \gg 0$. By [62, Theorem 5.5.10, clauses 1 and 2], the spectral sequence converges to $H^\bullet(C)$, whence $H^\bullet(C) = 0$, which shows that f is a quasi-isomorphism. \square

If now \mathcal{G}_0 and \mathcal{G}_1 are A_∞ -functors $\mathcal{A} \rightarrow \mathcal{B}$, one has the *length filtration* on the space of pre-natural transformations $T: \mathcal{G}_0 \rightarrow \mathcal{G}_1$ [50, (1f)]. The E_1 -term in the length spectral sequence $E_1^{rs} = CC^{r+s}(A, B)^s$ is a classical graded Hochschild bar complex for $A = H^* \mathcal{A}$, where $B = H^* \mathcal{B}$ becomes an A -bimodule via $H^* \mathcal{G}_0$ (for the right action) and $H^* \mathcal{G}_1$ (for the left action). E_2 is the Hochschild cohomology $E_2^{rs} = HH^{r+s}(A, B)^s$.

By construction,

- the length filtration is complete and bounded above by zero.

Thus, in suitable circumstances, one can use the spectral sequences of two length filtrations for comparison arguments based on Lemma [?]. This applies to the comparison arguments of [50, Lemmas 1.6–7], and to the construction of functors [50, Lemmas 1.9–10].

2.1.3 Units

One can still transform a cohomologically unital A_∞ -category into a strictly unital one by a formal diffeomorphism, by the argument of [50, Lemma 2.1].

2.1.4 Homological perturbation theory

Over a field, any cochain complex C^\bullet is related to its cohomology by homotopy-inverse cochain maps $C^\bullet \rightarrow H^*C$ and $HC^\bullet \rightarrow C^\bullet$, both inducing the identity map on cohomology. This observation is the starting point for the *transfer principle*, a case of the *homological perturbation lemma* (HPL), via which an A_∞ -structure on \mathcal{C} gives rise to an A_∞ -structure on $H^*\mathcal{C}$. The observation about cochain complexes fails over rings (e.g. over \mathbb{Z}), and the transfer principle is therefore not available in full generality.

The form of the HPL proved in [38] and quoted in [50, (1i)] is the following: suppose we have an A_∞ -category \mathcal{B} over a ring \mathbb{K} , and for each pair of objects a diagram

$$\begin{array}{ccc} \text{hom}_{\mathcal{A}}(X_0, X_1) & \xrightarrow{\mathcal{F}^1} & \text{hom}_{\mathcal{B}}(X_0, X_1) \\ & \xleftarrow{\mathcal{G}^1} & \uparrow T^1 \\ & & \text{hom}_{\mathcal{B}}(X_0, X_1) \end{array}$$

Here \mathcal{A} is not yet an A_∞ -category; rather, $(\text{hom}_{\mathcal{A}}(X_0, X_1), \mu_{\mathcal{A}}^1)$ is a cochain complex. The maps \mathcal{F}^1 and \mathcal{G}^1 are degree 0 cochain-maps; and T^1 (which has degree -1) a nullhomotopy of $\mathcal{F}^1 \circ \mathcal{G}^1 - \text{id}_{\mathcal{B}}$. The *homological perturbation lemma* (HPL) then supplies explicit formulae defining

- An A_∞ -structure $\{\mu_{\mathcal{A}}^d\}$ on \mathcal{A} extending $\mu_{\mathcal{A}}^1$ (not necessarily cohomologically unital);
- A_∞ -functors $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{G}: \mathcal{B} \rightarrow \mathcal{A}$ with \mathcal{F}^1 and \mathcal{G}^1 as leading terms; and
- a nullhomotopy T of $\mathcal{F} \circ \mathcal{G} - \text{id}_{\mathcal{B}}$ with T^1 as its leading term.

When \mathcal{B} carries a strict action of a group Γ , meaning that the spaces $\text{hom}_{\mathcal{B}}(X_0, X_1)$ carry Γ -actions which exactly intertwine the composition maps, and when the spaces $\text{hom}_{\mathcal{B}}(X_0, X_1)$ also carry Γ -actions intertwining $\mu_{\mathcal{B}}^1$, \mathcal{F}^1 , \mathcal{G}^1 and T^1 , the outcome

is that the A_∞ -structure \mathcal{A} again carries a strict action of Γ , making \mathcal{F} , \mathcal{G} and T equivariant. One way to see this is to apply HPL over the group algebra $\mathbb{K}[\Gamma]$ (the explicit formulae are the same whether one thinks of the base ring as \mathbb{K} or as $\mathbb{K}[\Gamma]$).

We shall usually apply the HPL in the following case:

$$\begin{array}{ccc} & \xrightarrow{\mathcal{F}^1} & \\ \mathrm{Hom}_{H\mathcal{B}}(X_0, X_1) & & \mathrm{hom}_{\mathcal{B}}(X_0, X_1) \\ & \xleftarrow{\mathcal{G}^1} & \end{array} \quad \begin{array}{c} \text{---} T^1 \text{---} \\ \text{---} \end{array}$$

On the left we have a Hom-space in the cohomology category, regarded as a cochain complex with vanishing differential. We suppose given a *splitting of the cochains*, that is, a direct sum decomposition

$$(10) \quad \mathrm{hom}_{\mathcal{B}}(X_0, X_1) = \mathrm{im} \mu_{\mathcal{B}}^1 \oplus H \oplus I,$$

where $\mathrm{im} \mu_{\mathcal{B}}^1 \oplus H = \ker \mu_{\mathcal{B}}^1$. The map \mathcal{F}^1 , a ‘cycle-choosing map’, sends a cohomology class $[c] \in \mathrm{Hom}(X_0, X_1)$ to its unique representative $c \in H$, while \mathcal{G}^1 is projection $\mathrm{hom}_{\mathcal{B}}(X_0, X_1) \rightarrow \ker \mu_{\mathcal{B}}^1 \rightarrow \ker \mu_{\mathcal{B}}^1 / \mathrm{im} \mu_{\mathcal{B}}^1$. The homotopy T^1 is zero on $\ker \mu_{\mathcal{B}}^1$, while $T^1: \mathrm{im} \mu_{\mathcal{B}}^1 = \mu_{\mathcal{B}}^1(I) \rightarrow I$ is characterized by $T(\mu^1 a) = a$.

Proposition 2.2 (transfer principle) *When a splitting of the cochains is given, there are explicit formulae defining*

- An A_∞ -structure $\{\mu_{H\mathcal{B}}^d\}$ on $H\mathcal{B}$ with $\mu_{H\mathcal{B}}^1 = 0$.
- A_∞ -functors $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{G}: \mathcal{B} \rightarrow \mathcal{A}$ with \mathcal{F}^1 and \mathcal{G}^1 as leading terms; and
- a nullhomotopy T of $\mathcal{F} \circ \mathcal{G} - \mathrm{id}_{\mathcal{B}}$ with T^1 as its leading term.

Notice that $H^*\mathcal{F}: H^*\mathcal{B} \rightarrow H^*\mathcal{B}$ is the identity functor. Three addenda:

- We have $\mu_{H\mathcal{B}}^2([a_2], [a_1]) = [\mu^2(a_2, a_1)]$ (follows from the explicit formula for $\mu_{H\mathcal{B}}^2$, see [50, equation (1.18)]); and
- A nullhomotopy of $\mathcal{G} \circ \mathcal{F} - \mathrm{id}_{H\mathcal{B}}$ exists (by the spectral sequence argument of [50, Remark 1.13], which applies equally over a commutative ring).
- When the input data are Γ -equivariant for some group Γ , the same holds for the output data.

We shall find it useful to have the following variant available:

Lemma 2.3 (Kadeishvili’s transfer principle [28]) *Let \mathcal{C} be an A_∞ -category whose Hom-space $\mathrm{Hom}_{H\mathcal{C}}(X_0, X_1)$ are projective \mathbb{K} -modules. Assume given a cycle-choosing map, i.e., for each pair of objects (X_1, X_2) , a map*

$$s: \mathrm{Hom}_{H\mathcal{C}}(X_1, X_2) \rightarrow \ker \mu_{\mathcal{C}}^1 \subset \mathrm{hom}_{\mathcal{C}}(X_1, X_2)$$

such that $[s(c)] = c$. (Projectivity insures that a cycle-choosing map exists.) Then there exist

- (1) A_∞ -operations $\{\nu^d\}$ on $H^*\mathcal{C}$, with $\nu^1 = 0$ and $\nu^2([a_2], [a_1]) = (-1)^{|a_1|}[\mu^2(a_2, a_1)]$;
- (2) and an A_∞ -functor $\mathcal{F}: H^*\mathcal{C} \rightarrow \mathcal{C}$ acting as the identity on objects with leading term $\mathcal{F}^1 = s$. The induced functor $F = H^*\mathcal{F}$ is the identity map.

The functor \mathcal{F} depends on additional choices of maps t^d , as explained in the proof.

Proof We run through Kadeishvili's argument from [28]. The only differences in our account are that Kadeishvili focuses on the case where \mathcal{C} is a DG category (though he says that the result holds for A_∞ -categories) and that he assumes the Hom-spaces to be free (not projective) modules. However, for transparency we write out the calculations that Kadeishvili leaves to the reader.

Introduce the shorthand

$$\hat{a} := (-1)^{|a|-1}a$$

for homogeneous elements in a hom-space $\text{hom}_{\mathcal{C}}(X_0, X_1)$ (we will write $\hat{\mu}^d$ for the \wedge -operation following μ^d). With a sequence of objects X_0, X_1, \dots, X_d fixed, and $0 \leq p \leq q \leq d$, write \mathbb{I}_p^q for the identity map on the tensor product

$$\text{hom}_{\mathcal{C}}(X_{q-1}, X_q) \otimes \cdots \otimes \text{hom}_{\mathcal{C}}(X_p, X_{p+1}) \otimes \text{hom}_{\mathcal{C}}(X_{p-1}, X_p),$$

and $\hat{\mathbb{I}}_p^q$ for the signed variant of \mathbb{I}_p^q which acts on homogeneous tensors as $a_q \otimes \cdots \otimes a_p \mapsto \hat{a}_q \otimes \cdots \otimes \hat{a}_p$. Note that $\mu_{\mathcal{C}}^r(\hat{\mathbb{I}}_p^{p+r}) = -\hat{\mu}_{\mathcal{C}}^r(\mathbb{I}_p^{p+r})$.

We need to define ν and \mathcal{F} , starting from $\mathcal{F}^1 = s$ and $\nu^1 = 0$, so as to satisfy the A_∞ -structure equations

$$(11) \quad \sum_{m,n} \nu^{d-m+1}(\mathbb{I}_{n+m+1}^d, \nu^m(\mathbb{I}_{n+1}^{n+m}), \hat{\mathbb{I}}_1^n) = 0,$$

and the functor equations

$$(12) \quad \sum_r \sum_{s_1, \dots, s_r} \mu^r(\mathcal{F}^{s_r}(\mathbb{I}_{d-s_r+1}^d), \dots, \mathcal{F}^{s_1}(\mathbb{I}_1^{s_1})) = \sum_{m,n} \mathcal{F}^{d-m+1}(\mathbb{I}_{n+1}^d, \nu^m(\mathbb{I}_{n+1}^{n+m}), \hat{\mathbb{I}}_1^n).$$

Take $r > 1$, and assume that \mathcal{F}^r and ν^r have been defined for all $r < d$, and that they satisfy the structure equations (11) for $\leq d$ inputs (these only use ν^i for $i < d$, because $\nu^1 = 0$) and the functor equations (12) for $< d$ inputs. Following [28], define $U^d = S^d - R^d$ acting on composable morphisms a_1, \dots, a_d in $H^*\mathcal{C}$ as follows:

$$S^d = \sum_{r \geq 2} \sum_{s_1, \dots, s_r} \mu^r(\mathcal{F}^{s_r}(\mathbb{I}_{d-s_r+1}^d), \dots, \mathcal{F}^{s_1}(\mathbb{I}_1^{s_1})),$$

$$R^d = \sum_n \sum_{m < d} \mathcal{F}^{d-m+1} (\mathbb{I}_{n+m+1}^d, \nu^m(\mathbb{I}_{n+1}^{n+m}), \hat{\mathbb{I}}_1^n).$$

Since $\nu^1 = 0$, the summation defining R^d is effectively over $2 \leq m \leq d-1$.

In what follows, for easier legibility we sometimes omit sub- and superscripts from \mathbb{I} . We also sometimes do not label the indexing sets of our sums; the default is to sum over all variables except d , which is fixed.

One checks that $\mu^1 \circ U^d = 0$, as follows. By the functor equations applied to $(\mathbb{I}, \nu^m(\mathbb{I}_{n+1}^{n+m}), \hat{\mathbb{I}}_1^n)$, summed over $m < d$ and over n , we have

$$\begin{aligned} \mu^1 \circ R^d &= - \sum_{r \geq 2} \sum_{s_1, \dots, s_r} \sum_{m < d} \mu^r (\mathcal{F}^{s_r}(\mathbb{I}_{r-s_r+1}^d), \dots, \mathcal{F}^{s_1-m+1}(\mathbb{I}, \nu^m(\mathbb{I}), \hat{\mathbb{I}}), \dots, \mathcal{F}^{s_1}(\hat{\mathbb{I}}_1^{s_1})) \\ &\quad + \sum \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^{m+p+1} (\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}_{n+1}^{n+p}), \mathbb{I}_1^n) \\ &\quad + \sum \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}, \hat{\nu}^m(\mathbb{I}_{n+1}^{n+m}), \mathbb{I}_1^n) \\ &\quad + \sum \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^m(\mathbb{I}), \hat{\mathbb{I}}, \nu^q(\hat{\mathbb{I}}_{n+1}^{n+q}), \mathbb{I}_1^n). \end{aligned}$$

Of the four sums after the $=$ sign, the second is zero by (11), while the third cancels the fourth. The A_∞ -equations for μ give

$$-\mu^1 S^d = \sum_{m \leq r-1} \mu^{r-m+1} (\mathcal{F}^{s_r}(\mathbb{I}), \dots, \mu^n (\mathcal{F}^{s_{n+m}}(\mathbb{I}), \dots, \mathcal{F}^{s_{n+1}}(\mathbb{I})), \dots, \hat{\mathcal{F}}^{s_1}(\hat{\mathbb{I}}_1^{s_1})).$$

By the functor equations (12) we conclude that $\mu^1 \circ U^d = 0$. We put

$$\nu^d(a_d, \dots, a_1) = [U^d(a_d, \dots, a_1)].$$

The functor equation for d inputs now reads $\mu^1 \circ \mathcal{F}^d = U^d$. Choose a map

$$t^d: \text{Hom}_{H\mathcal{C}}(X_{d-1}, X_d) \otimes \dots \otimes \text{Hom}_{H\mathcal{C}}(X_0, X_1) \rightarrow \text{hom}_{\mathcal{C}}(X_0, X_d)$$

such that $\mu^1 \circ t^d = U^d$. Existence of t^d follows from projectivity of the Hom-spaces, which is preserved by tensor products. Truth of the functor equation for d inputs is then assured.

One next has to check the structure equations with $d+1$ inputs. We compute

$$\begin{aligned} \sum_m S^d (\mathbb{I}_{m+3}^{d+1}, \nu^2(\mathbb{I}_{m+1}^{m+2}), \hat{\mathbb{I}}_1^m) \\ &= \sum_{r \geq 2} \mu^r (\mathcal{F}^{s_r}(\mathbb{I}_{d-s_r+1}^d), \dots, \mathcal{F}^{s_1}(\mathbb{I}, \nu^2(\mathbb{I}), \hat{\mathbb{I}}), \dots, \mathcal{F}^{s_1}(\hat{\mathbb{I}}_1^{s_1})), \\ &= - \sum_{r \geq 2; q \geq 3} \mu^r (\mathcal{F}^{s_r}(\mathbb{I}_{d-s_r+1}^d), \dots, \mathcal{F}^{s_i}(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}), \dots, \mathcal{F}^{s_1}(\hat{\mathbb{I}}_1^{s_1})) \\ &\quad + \sum_{r \geq 2} \mu^r (\mathcal{F}^{s_r}(\mathbb{I}_{d-s_r+1}^d), \dots, \mu^q (\mathcal{F}^{t_q}(\mathbb{I}), \dots, \mathcal{F}^{t_1}(\mathbb{I})), \dots, \mathcal{F}^{s_1}(\hat{\mathbb{I}}_1^{s_1})) \end{aligned}$$

and observe that the last sum vanishes modulo $\text{im } \mu^1$, by the A_∞ structure equations. Modulo $\text{im } \mu^1$, we have

$$\begin{aligned}
 & \sum_m S^d (\mathbb{I}_{m+3}^{d+1}, \nu^2(\mathbb{I}_{m+1}^{m+2}), \hat{\mathbb{I}}_1^m) \\
 &= - \sum_{q \geq 3} \mathcal{F}^{d-m-q+3} (\mathbb{I}, \nu^m(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}), \mathbb{I}) \\
 & \quad - \sum_{q \geq 3} \mathcal{F}^{d-m-q+3} (\mathbb{I}, \nu^m(\mathbb{I}), \hat{\mathbb{I}}, \hat{\nu}^q(\mathbb{I}), \mathbb{I}) - \sum_{q \geq 3} \mathcal{F}^{d-m-q+3} (\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}, \nu^m(\hat{\mathbb{I}}), \mathbb{I}) \\
 &= - \sum_{q \geq 3} \mathcal{F}^{d-m-q+3} (\mathbb{I}, \nu^m(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}), \mathbb{I}) \\
 & \quad - \sum_{q \geq 3} \mathcal{F}^{d-q+1} (\mathbb{I}, \nu^2(\mathbb{I}), \hat{\mathbb{I}}, \hat{\nu}^q(\mathbb{I}), \mathbb{I}) - \sum_{q \geq 3} \mathcal{F}^{d-q+1} (\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}, \nu^2(\hat{\mathbb{I}}), \mathbb{I}).
 \end{aligned}$$

And

$$\begin{aligned}
 & \sum_m R^d (\mathbb{I}_{m+3}^{d+1}, \nu^2(\mathbb{I}_{m+1}^{m+2}), \hat{\mathbb{I}}_1^m) \\
 &= \sum_{m < d} \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^m(\mathbb{I}, \nu^2(\mathbb{I}), \hat{\mathbb{I}}), \mathbb{I}) \\
 & \quad + \sum \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^2(\mathbb{I}), \hat{\mathbb{I}}, \nu^m(\hat{\mathbb{I}}), \mathbb{I}) + \sum \mathcal{F}^{d-m+1} (\mathbb{I}, \nu^m(\mathbb{I}), \hat{\mathbb{I}}, \hat{\nu}^2(\mathbb{I}), \mathbb{I}).
 \end{aligned}$$

We find (again modulo $\text{im } \mu^1$) that

$$\begin{aligned}
 \sum_m U^d (\mathbb{I}_{m+3}^{d+1}, \nu^2(\mathbb{I}_{m+1}^{m+2}), \hat{\mathbb{I}}_1^m) &= - \sum_{m < d} \mathcal{F}^{d-m-q+2} (\mathbb{I}, \nu^m(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}}), \mathbb{I}) \\
 & \quad - \sum_{q \geq 3} \mathcal{F}^1 (\nu^{d-q+2}(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}})) \\
 &= - \sum_{q \geq 3} \mathcal{F}^1 (\nu^{d-q+2}(\mathbb{I}, \nu^q(\mathbb{I}), \hat{\mathbb{I}})).
 \end{aligned}$$

Passing to cohomology, we get $\sum_{m,n} \nu^{d-m+2}(\mathbb{I}, \nu^m(\mathbb{I}), \hat{\mathbb{I}}_1^n) = 0$. \square

Corollary 2.4 Suppose \mathcal{A} and \mathcal{B} are A_∞ -categories such that (i) the graded \mathbb{K} -modules $\text{Hom}_{H\mathcal{A}}(X_0, X_1)$ are all projective, and (ii) each hom-space in \mathcal{B} admits a splitting of the cochains. Then any A_∞ -functor $\mathcal{K}: \mathcal{A} \rightarrow \mathcal{B}$ which induces an isomorphism $K: H^*\mathcal{A} \rightarrow H^*\mathcal{B}$ has a two-sided homotopy inverse.

Proof The hypotheses allow us to construct minimal A_∞ -structures on $H^*\mathcal{A}$ and $H^*\mathcal{B}$; A_∞ -functors $\alpha: H^*\mathcal{A} \rightarrow \mathcal{A}$ and $\beta: H^*\mathcal{B} \rightarrow \mathcal{B}$, both inducing an identity map

on cohomology; and a homotopy-inverse γ to β . In the diagram

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{\mathcal{K}} & \mathcal{B} \\ \alpha \uparrow & & \beta \uparrow \downarrow \gamma \\ H^*\mathcal{A} & \xrightleftharpoons[\phi^{-1}]{\phi} & H^*\mathcal{B} \end{array}$$

the composite $\phi := \gamma \circ \mathcal{K} \circ \alpha$ is an A_∞ -functor whose leading term is the isomorphism $\phi^1 = K$. The inverse K^{-1} is the leading term in a strict inverse ϕ^{-1} to ϕ . Put $\mathcal{L} = \alpha \circ \phi^{-1} \circ \gamma$. We have $\text{id}_{H\mathcal{B}} = \gamma \circ \mathcal{K} \circ \alpha \circ \phi^{-1}$. Hence $\beta \simeq \mathcal{K} \circ \alpha \circ \phi^{-1}$, and so $\mathcal{K} \circ \mathcal{L} = \mathcal{K} \circ \alpha \circ \phi^{-1} \circ \gamma \simeq \text{id}_{\mathcal{B}}$. We have $\text{id}_{H\mathcal{A}} = \phi^{-1} \circ \gamma \circ \mathcal{K} \circ \alpha$, so $\alpha = \mathcal{L} \circ \mathcal{K} \circ \alpha$. Right composition

$$\mathcal{R}_\alpha : \text{hom}_{\text{nu fun}(\mathcal{A}, \mathcal{A})}(\text{id}_{\mathcal{A}}, \mathcal{L} \circ \mathcal{K}) \rightarrow \text{hom}_{\text{nu fun}(H\mathcal{A}, \mathcal{A})}(\alpha, \mathcal{L} \circ \mathcal{K} \circ \alpha)$$

is a quasi-isomorphism of cochain complexes [50, Lemma 1.7], preserving the length filtration, and therefore acting as a quasi-isomorphism on the first level F^1 of the length filtration (i.e., on pre-natural transformations T with $T^0 = 0$). Define $T \in \text{hom}_{\text{nu fun}(\mathcal{A}, \mathcal{A})}^0(\text{id}_{\mathcal{A}}, \mathcal{L} \circ \mathcal{K})$ by $T^0 = 0$ and $T^d = (\text{id}_{\mathcal{A}} - \mathcal{L} \circ \mathcal{K})^d$ for $d > 0$. Then $\mathcal{R}_\alpha(T) = 0$, and hence $T = \mu^1(S)$ for some S with $S^0 = 0$; that is, $\text{id}_{\mathcal{A}}$ is homotopic to $\mathcal{L} \circ \mathcal{K}$. \square

This last result is woefully inadequate for wholesale application over non-fields. Nonetheless, we shall find it useful in specific situations.

Quasi-equivalences. Let $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{B}$ be a c-unital A_∞ -functor. We shall adopt the following terminology:

- \mathcal{F} is called a *quasi-isomorphism* if $H\mathcal{F} : H\mathcal{A} \rightarrow H\mathcal{B}$ is an isomorphism of categories.
- \mathcal{F} is *homotopy-invertible* if there is a $\mathcal{G} : \mathcal{B} \rightarrow \mathcal{A}$ such that $\mathcal{G} \circ \mathcal{F}$ acts trivially on objects and is homotopic to $\text{id}_{\mathcal{A}}$, and similarly for $\mathcal{F} \circ \mathcal{G}$.
- \mathcal{F} is called a *quasi-equivalence* if $H\mathcal{F} : H\mathcal{A} \rightarrow H\mathcal{B}$ is an equivalence of categories.
- \mathcal{F} is *quasi-invertible* if there is a $\mathcal{G} : \mathcal{B} \rightarrow \mathcal{A}$ such that $\mathcal{G} \circ \mathcal{F}$ and $\text{id}_{\mathcal{A}}$ are isomorphic objects in $H^0(\text{fun}(\mathcal{A}, \mathcal{A}))$, and $\mathcal{F} \circ \mathcal{G} \cong \text{id}_{\mathcal{B}}$ in a similar sense.

A homotopy-invertible functor is a quasi-isomorphism, and a quasi-invertible functor is a quasi-equivalence. Over a field, a quasi-isomorphism is always homotopy-invertible, and a quasi-equivalence always quasi-invertible, but these principles fail over more general rings.

Lemma 2.5 ([50], Lemma 2.8) *Suppose that $\tilde{\mathcal{A}} \subset \mathcal{A}$ is a full subcategory such that the inclusion is a quasi-equivalence. Then, for any c -unital A_∞ -category \mathcal{B} , restriction of c -unital A_∞ -functors and their pre-natural transformations yields a quasi-equivalence $\mathcal{Q} := \text{fun}(\mathcal{A}, \mathcal{B}) \rightarrow \tilde{\mathcal{Q}} := \text{fun}(\tilde{\mathcal{A}}, \mathcal{B})$.*

Proof A reading of Seidel’s argument reveals that it uses only tools that are available over a ring \mathbb{K} , so it carries through unchanged. Indeed, the fact that the restriction map is cohomologically full and faithful is proved by observing that it induces a map of filtered complexes $\text{hom}_{\mathcal{Q}}(\mathcal{G}_0, \mathcal{G}_1) \rightarrow \text{hom}_{\tilde{\mathcal{Q}}}(\tilde{\mathcal{G}}_0, \tilde{\mathcal{G}}_1)$ (filtered by length) and hence a map of spectral sequences. The latter map is an isomorphism on E_2 by the Morita invariance of Hochschild cohomology, proved via an explicit homotopy equivalence that Seidel writes down. It is therefore an isomorphism on E_∞ , i.e., $\text{gr Hom}_{H\mathcal{Q}}(\mathcal{G}_0, \mathcal{G}_1) \rightarrow \text{gr Hom}_{H\tilde{\mathcal{Q}}}(\tilde{\mathcal{G}}_0, \tilde{\mathcal{G}}_1)$ is an isomorphism of graded modules; therefore $\text{Hom}_{H\mathcal{Q}}(\mathcal{G}_0, \mathcal{G}_1) \rightarrow \text{Hom}_{H\tilde{\mathcal{Q}}}(\tilde{\mathcal{G}}_0, \tilde{\mathcal{G}}_1)$ is an isomorphism. That any functor $\tilde{\mathcal{G}}: \tilde{\mathcal{A}} \rightarrow \tilde{\mathcal{B}}$ extends to $\mathcal{G}: \mathcal{A} \rightarrow \mathcal{B}$ is proved via Seidel’s Lemma 1.10, which as we have already observed, still applies over \mathbb{K} . \square

Proposition 2.6 *Suppose that $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$ restricts to a quasi-isomorphism $\tilde{\mathcal{F}}: \tilde{\mathcal{A}} \rightarrow \tilde{\mathcal{B}}$, where the inclusions $\tilde{\mathcal{A}} \subset \mathcal{A}$ are quasi-equivalences. If $\tilde{\mathcal{F}}$ is homotopy-invertible then \mathcal{F} is quasi-invertible.*

Proof As already observed in the proof of the last lemma, there exists (by [50, Lemma 1.10]) a functor $\mathcal{J}: \mathcal{A} \rightarrow \tilde{\mathcal{A}}$ restricting to the identity functor on $\tilde{\mathcal{A}}$, and quasi-inverse to the inclusion functor incl . Similarly, $\mathcal{J}: \mathcal{B} \rightarrow \tilde{\mathcal{B}}$ restricts to the identity functor on $\tilde{\mathcal{B}}$ and is quasi-inverse to the inclusion. If $\tilde{\mathcal{G}}$ is a homotopy-inverse to $\tilde{\mathcal{F}}$, so that $\tilde{\mathcal{F}} \circ \tilde{\mathcal{G}}$ is homotopic to $\text{id}_{\tilde{\mathcal{B}}}$, then it is also true that $\tilde{\mathcal{G}} \circ \tilde{\mathcal{F}} \cong \text{id}_{\tilde{\mathcal{A}}}$ in $H^0\text{fun}(\tilde{\mathcal{B}}, \tilde{\mathcal{A}})$; similarly for $\tilde{\mathcal{G}} \circ \tilde{\mathcal{F}}$. Thus $\mathcal{G} := \text{incl} \circ \tilde{\mathcal{G}} \circ \mathcal{J}: \mathcal{B} \rightarrow \mathcal{A}$, as a composite of quasi-invertible functors, is itself quasi-invertible. \square

2.1.5 Exact triangles

The notion of a *quasi-representative* for an A_∞ -modules [50, (3a)] goes through unchanged. So too do the notions of the *abstract mapping cone* of a morphism $c \in \text{hom}_{\mathcal{C}}^0(X_0, X_1)$ such that $\mu_{\mathcal{C}}^1 = 0$ [50, (3e)]; an *exact triangle* in an A_∞ -category [50, (3f)]; and of a *triangulated A_∞ -category* [50, (3h)]. One can detect exact triangles by the criterion of [50, Lemma 3.7].

The *additive enlargement* $\Sigma\mathcal{C}$ of the A_∞ -category \mathcal{C} has objects $\bigoplus_{i \in I} V_i \otimes X_i$, where the index set I is finite, where $X_i \in \text{ob } \mathcal{C}$, and where the ‘multiplicity space’ V_i is a graded module which is free in each degree and has finite total rank. The morphism

spaces and composition are defined in a fairly obvious way (the details can be found in [50, (3k)]). A *connection* on $X \in \text{ob } \Sigma\mathcal{C}$ is a morphism $\delta_X \in \text{hom}_{\Sigma\mathcal{C}}^1(X, X)$; the pair (X, δ_X) is called a *pre-twisted complex*.

If the multiplicity spaces V_i in a pre-twisted complex have submodules U_i such that δ_X preserves $\bigoplus_{i \in I} U_i \otimes X_i$, we call $\bigoplus_{i \in I} U_i \otimes X_i$ a *subcomplex*; for the purpose at hand, it doesn't matter whether the U_i are free modules. A pre-twisted complex (X, δ_X) is called *strictly lower-triangular* if there exists a finite filtration by subcomplexes $X = F^0 X \supset F^1 X \supset \dots \supset F^n X = 0$ such that $\delta_X(F^k X) \subset F^{k+1} X$ for all k . We then define a *twisted complex* to be a strictly lower-triangular pre-twisted complex (X, δ_X) satisfying the Maurer–Cartan equation

$$\sum_{r \geq 1} \mu_{\Sigma\mathcal{C}}^r(\delta_X, \dots, \delta_X) = 0 \quad \text{in } \text{hom}_{\Sigma\mathcal{C}}^2(X, X).$$

The reason to require lower triangularity is to make the sum on the LHS finite. There is an A_∞ -category of twisted complexes $\text{tw } \mathcal{C}$, together with a fully faithful inclusion functor $\mathcal{C} \rightarrow \text{tw } \mathcal{C}$.

Passing to twisted complexes is a natural operation: there is an A_∞ -functor

$$\text{tw}: \text{nu fun}(\mathcal{A}, \mathcal{B}) \rightarrow \text{nu fun}(\text{tw } \mathcal{A}, \text{tw } \mathcal{B}).$$

The assumption of c-unitality for \mathcal{A} implies c-unitality of $\text{tw } \mathcal{A}$, and in fact tw is a c-unital A_∞ -functor [50, Lemma 3.24].

Proposition 2.7 *Take a c-unital functor $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$, inducing $\text{tw } \mathcal{F}: \text{tw } \mathcal{A} \rightarrow \text{tw } \mathcal{B}$.*

- (1) *If \mathcal{F} is homotopy-invertible then so is $\text{tw } \mathcal{F}$.*
- (2) *If \mathcal{F} is quasi-invertible then so is $\text{tw } \mathcal{F}$.*
- (3) *If \mathcal{F} is cohomologically full and faithful then so is $\text{tw } \mathcal{F}$.*
- (4) *If \mathcal{F} is a quasi-isomorphism then so is $\text{tw } \mathcal{F}$.*
- (5) *If \mathcal{F} is a quasi-equivalence then so is $\text{tw } \mathcal{F}$.*

incomplete

Proof (1) Let \mathcal{G} be a homotopy-inverse to \mathcal{F} . Then $\text{tw } \mathcal{G} \circ \text{tw } \mathcal{F}$ acts as the identity on objects. Let $\mathcal{Q} = \text{nu fun}(\mathcal{A}, \mathcal{A})$. A homotopy from $\mathcal{G} \circ \mathcal{F}$ to $\text{id}_{\mathcal{A}}$ is an element $T \in \text{hom}_{\mathcal{Q}}^0(\mathcal{G} \circ \mathcal{F}, \text{id}_{\mathcal{A}})$ such that $T^0 = 0$ and $(\mu^1 T)^d = (\text{id}_{\mathcal{A}} - \mathcal{G} \circ \mathcal{F})^d$ for all d . One has the functor $\text{tw}: \mathcal{Q} \rightarrow \mathcal{T} := \text{nu fun}(\text{tw } \mathcal{A}, \text{tw } \mathcal{A})$, and $\text{tw } T \in \text{hom}_{\mathcal{T}}^0(\text{tw } \mathcal{G} \circ \text{tw } \mathcal{F} - \text{id}_{\text{tw } \mathcal{A}})$ is the required homotopy (we note in passing that $\text{tw } \mathcal{G} \circ \text{tw } \mathcal{F} = \text{tw}(\mathcal{G} \circ \mathcal{F})$). Similarly for $\text{tw } \mathcal{F} \circ \text{tw } \mathcal{G}$.

(2)

(3) As in [50, Lemma 3.23].

(4) Every object in \mathcal{B} is isomorphic in $H^0\mathcal{B}$ to an object in the image of \mathcal{F} . Take $(Y = \bigoplus_i V_i \otimes Y_i, \delta_Y) \in \text{ob tw } \mathcal{B}$. Find objects $X_i \in \mathcal{A}$ and cocycles $c_i \in \text{hom}^0(\mathcal{F}(X_i), Y_i)$ for which $[c_i]$ are isomorphisms. Then $\text{tw } \mathcal{F}(\bigoplus V_i \otimes X_i) \dots$

(5)

□

In the A_∞ -category $\text{tw } \mathcal{A}$, there is a canonical construction of the mapping cone of a closed, degree-zero morphism. There is also a shift operation. Thus $\text{tw } \mathcal{A}$ is a triangulated A_∞ -category.

Lemma 2.8 ([50], Lemma 3.31) *A triangle in $H\mathcal{A}$ is exact if and only if its image in $H(\text{tw } \mathcal{A})$ is exact.*

check this carefully!

Several properties of exact triangles follow [50, Corollaries 3.8-9 and 3.15-19].

2.1.6 Idempotents

We now discuss idempotent completion. Suppose that $X \in \text{ob } \mathcal{A}$, and that $p \in \text{Hom}_{H\mathcal{A}}^0(X, X)$ is an idempotent endomorphism ($p^2 = p$).

Consider the Yoneda \mathcal{A} -module $\mathcal{Y}(X) = \text{hom}_{\mathcal{A}}(\cdot, X)$. Under the Yoneda isomorphism $\text{Hom}_{H\mathcal{A}}(X, X) \rightarrow \text{Hom}_{H(\text{mod-}\mathcal{A})}(\mathcal{Y}(X), \mathcal{Y}(X))$, an idempotent $p \in \text{Hom}_{H\mathcal{A}}^0(X, X)$ maps to an endomorphism π of $\mathcal{Y}(X)$ such that $\pi^2 = \pi$ in $H(\text{mod-}\mathcal{A})$. An *abstract image* for an idempotent $p \in \text{Hom}_{H\mathcal{A}}^0(X, X)$ is an \mathcal{A} -module \mathcal{Z} such that \mathcal{Z} is an *image* for π in $H(\text{mod-}\mathcal{A})$. That is, there are morphisms r and k in $H(\text{mod-}\mathcal{A})$,

$$\mathcal{Y}(X) \begin{matrix} \xleftarrow{r} \\ \xrightarrow{k} \end{matrix} \mathcal{Z}$$

such that $r \circ k = \text{id}_{\mathcal{Z}}$, $k \circ r = \pi$.

Lemma 2.9 *If \mathcal{Z} is an abstract image for $p \in \text{hom}^0(X, X)$ then one has natural isomorphisms $\text{Hom}_{H(\text{mod-}\mathcal{A})}(\mathcal{M}, \mathcal{Z}) \cong (H\mathcal{M})(X)p$ for $\mathcal{M} \in \text{ob mod-}\mathcal{A}$.*

Proof One has a quasi-isomorphism of cochain complexes

$$\lambda_{\mathcal{M}}: \text{hom}_{\text{mod-}\mathcal{A}}(\mathcal{M}, \mathcal{Y}(X)) \rightarrow \mathcal{M}(X)$$

[]. Thus one has maps

$$[\lambda_{\mathcal{M}}] \circ k^*: \text{Hom}_{H(\text{mod-}\mathcal{A})}(\mathcal{M}, \mathcal{Z}) \rightarrow H(\mathcal{M}(X)) = (H\mathcal{M})(X).$$

and

$$r^* \circ [\lambda_{\mathcal{M}}]^{-1}: (H\mathcal{M})(X) \rightarrow \text{Hom}_{H(\text{mod-}\mathcal{A})}(\mathcal{M}, \mathcal{Z})$$

whose composites are id (on $\text{Hom}_{H(\text{mod-}\mathcal{A})}(\mathcal{M}, \mathcal{Z})$) and composition on the right with p (on $(H\mathcal{M})(X)$). Since $[\lambda_{\mathcal{M}}]$ is natural in \mathcal{M} , the same is true of these isomorphisms.

□

Proposition 2.10 *If \mathcal{A} is strictly unital and has a zero-object 0 (so $\text{hom}_{\mathcal{A}}(0, \cdot) = 0$ and $\text{hom}_{\mathcal{A}}(\cdot, 0) = 0$) then every idempotent $p \in \text{Hom}_{H\mathcal{A}}^0(X, X)$ has an abstract image.*

Proof This is proved in two stages. In the first stage, we lift p to an *idempotent up to homotopy*. Introduce the A_∞ -category \mathbb{K} with one object $*$ whose endomorphism algebra is \mathbb{K} ; of the operations, only μ^2 is non-zero. An idempotent up to homotopy for an object $X \in \text{ob } \mathcal{A}$ and an idempotent endomorphism $p \in \text{Hom}_{H\mathcal{A}}^0(X, X)$ is an A_∞ functor $\mathcal{P}: \mathbb{K} \rightarrow \mathcal{A}$ such that $\mathcal{P}(*) = X$ and $[\mathcal{P}^1(\text{id}_{\mathbb{K}})] = p$. According to [50, Lemma 4.2], which is valid over rings, an idempotent endomorphism $p \in \text{Hom}_{H\mathcal{A}}(X, X)$ always arises from some idempotent up to homotopy.

The second step is to construct an abstract image \mathcal{Z} in terms of \mathcal{P} . Seidel offers two such constructions, but the first uses a ‘power series trick’, and is potentially problematic over rings, because [50, Lemmas 4.5-6] rely on projectivity of Hom-spaces in $H\mathcal{A}$. For this reason, we shall not use it. We shall instead adopt Seidel’s alternative construction, which involves $\text{tw}^- \mathcal{A}$, an A_∞ -category of ‘infinite twisted complexes’. There is a fully faithful inclusion functor $\mathcal{A} \rightarrow \text{tw}^- \mathcal{A}$, and a cohomologically full and faithful functor $\text{tw}^- \mathcal{A} \rightarrow \text{mod-}\mathcal{A}$ extending the Yoneda embedding $\mathcal{Y}: \mathcal{A} \rightarrow \text{mod-}\mathcal{A}$. Thus one can view p as an idempotent endomorphism in $\text{tw}^- \mathcal{A}$. Seidel assumes that \mathcal{A} is strictly unital and has a zero-object. Under those assumptions, \mathcal{P} determines a certain infinite twisted complex $\mathcal{Z}^- \in \text{ob } \text{tw}^- \mathcal{A}$. There are morphisms $r^-: X \rightarrow \mathcal{Z}^-$ and $k^-: \mathcal{Z}^- \rightarrow X$ in $H(\text{tw}^- \mathcal{A})$ such that $k^- \circ r^- = p$ and $r^- \circ k^- = \text{id}$. Under the functor $\text{tw}^- \mathcal{A} \rightarrow \text{mod-}\mathcal{A}$, \mathcal{Z}^- maps to an abstract image \mathcal{Z} for p . \square

- An A_∞ -category \mathcal{B} with a zero-object to be *split-closed* if for every idempotent up to homotopy \mathcal{P} has an abstract image \mathcal{Z} which is quasi-represented by some $Z \in \text{ob } \mathcal{B}$. That is, $\mathcal{Z} \cong \mathcal{Y}(Z)$ in $H^0(\text{mod-}\mathcal{B})$.
- A *split-closure* for an A_∞ -category \mathcal{A} is a split-closed A_∞ -category \mathcal{B} and a cohomologically full and faithful functor $\mathcal{A} \rightarrow \mathcal{B}$ such that every object in $H^0(\mathcal{B})$ is the image of an idempotent endomorphism of $H^0(\mathcal{A})$.

As in Seidel’s account, \mathcal{B} is split-closed if and only if $H\mathcal{B}$ is split-closed (in the sense that every idempotent $p \in \text{Hom}_{H\mathcal{B}}^0(X, X)$ has an image $Z \in \text{ob } H\mathcal{B}$).

If \mathcal{A} is strictly unital with a zero-object, declare $\Pi\mathcal{A}$ to be the full A_∞ -subcategory of $\text{tw}^- \mathcal{A}$ whose objects are objects of \mathcal{A} (viewed as twisted complexes) and the abstract images of idempotents up to homotopy \mathcal{P} . Then $\Pi\mathcal{A}$ is a split closure for \mathcal{A} .

If \mathcal{C} is strictly unital and non-empty, then $\text{tw } \mathcal{C}$ is strictly unital with a zero-object; thus we can form the split-closure $\Pi \text{tw } \mathcal{C}$. (If \mathcal{C} is merely c-unital then one can construct a formal diffeomorphism ϕ of \mathcal{C} such that $\phi^*\mathcal{C}$ is strictly unital. One can extend ϕ to a formal diffeomorphism Φ of $\text{tw } \mathcal{C}$ such that $\Phi^*\mathcal{C}$ is strictly unital. One has $\Phi^*(\text{tw } \mathcal{C}) \cong \text{tw}(\phi^*\mathcal{C})$. Thus $\Pi \text{tw}(\phi^*\mathcal{C})$ is a split closure of $\text{tw } \mathcal{C}$.) Thus we have

Proposition 2.11 *For every A_∞ -category \mathcal{A} , $\text{tw } \mathcal{A}$ has a split-closure.*

2.1.7 Spherical twists

2.1.8 Weak Calabi–Yau structures

If \mathcal{C} is a \mathbb{K} -linear A_∞ -structure, one has the A_∞ -category $\mathcal{M} = \text{bimod}(\mathcal{C}, \mathcal{C})$ of $(\mathcal{C}, \mathcal{C})$ -bimodules.

Let \mathcal{C}_Δ denote the identity (or ‘diagonal’) $(\mathcal{C}, \mathcal{C})$ -bimodule, and \mathcal{C}_Δ^\vee its \mathbb{K} -linear dual. A *weak CY structure of dimension n* is a bimodule homomorphism $\beta: \mathcal{C}_\Delta \rightarrow \mathcal{C}_\Delta^\vee[-n]$ such that

- the class $[\beta] \in \text{Hom}_{H\mathcal{M}}^0(\mathcal{C}_\Delta, \mathcal{C}_\Delta^\vee[-n])$ is an isomorphism in $H\mathcal{M}$; and
- the bimodule homomorphism $\tilde{\beta}: \mathcal{C}_\Delta \otimes_{\mathbb{K}} \mathcal{C}_\Delta \rightarrow \mathbb{K}[-n]$ adjoint to β is *symmetric*. That is, if $\sigma: \mathcal{C} \otimes_{\mathbb{K}} \mathcal{C} \rightarrow \mathcal{C} \otimes_{\mathbb{K}} \mathcal{C}$ is the transposition map, then

$$[\tilde{\beta} \circ \sigma] = [\tilde{\beta}] \in \text{Hom}_{H\mathcal{M}}^0(\mathcal{C}_\Delta \otimes_{\mathbb{K}} \mathcal{C}_\Delta, \mathbb{K}[-n]).$$

Suppose now that $\mathcal{F}: \mathcal{A} \rightarrow \mathcal{B}$ is an A_∞ -functor, where \mathcal{A} and \mathcal{B} come with respective weak CY structures α and β , both of dimension n . One has arrows in $\text{bimod}(\mathcal{A}, \mathcal{A})$ as follows:

$$\begin{array}{ccc} \mathcal{A}_\Delta & \xrightarrow{\mathcal{F}_\Delta} & \mathcal{F}^* \mathcal{B}_\Delta \\ \alpha \downarrow & & \downarrow \beta \\ \mathcal{A}_\Delta^\vee[-n] & \xleftarrow{\mathcal{F}_\Delta^\vee} & \mathcal{F}^* \mathcal{B}_\Delta^\vee[-n]. \end{array}$$

Here \mathcal{F}_Δ is the functor \mathcal{F} recast as a bimodule homomorphism, and \mathcal{F}_Δ^\vee its dual. We call \mathcal{F} a *weak CY functor* if the diagram commutes in $H^0 \text{bimod}(\mathcal{A}, \mathcal{A})$.

2.2 Derived categories and A_∞ -categories

Thomason’s theorem. Let \mathcal{T} be a triangulated category. A full triangulated subcategory $\mathcal{A} \subset \mathcal{T}$ is called *dense* if every object in \mathcal{T} is isomorphic to a direct summand in an object of \mathcal{A} .

Theorem 2.12 (Thomason [58]) *Suppose that \mathcal{T} is triangulated and essentially small, and that the full triangulated subcategory $\mathcal{A} \subset \mathcal{T}$ is strictly full (i.e. closed under isomorphism in \mathcal{T} and dense, and that the natural map on Grothendieck groups $K_0(\mathcal{A}) \rightarrow K_0(\mathcal{T})$ is an isomorphism. Then $\mathcal{A} = \mathcal{T}$.*

Corollary 2.13 *Suppose \mathcal{B} is a split-closed, triangulated A_∞ -category over \mathbb{K} , and $\mathcal{A} \subset \mathcal{B}$ a full triangulated A_∞ -subcategory which split-generates \mathcal{B} . If the natural map $K_0(H^0\mathcal{A}) \rightarrow K_0(H^0\mathcal{B})$ is an isomorphism then the inclusion functor $\mathcal{A} \rightarrow \mathcal{B}$ is quasi-invertible.*

Proof The full triangulated subcategory $H^0\mathcal{A} \subset H^0\mathcal{B}$ is dense, by definition of split-generation. Let \mathcal{A} be the smallest strictly full triangulated subcategory of $H^0\mathcal{B}$ that contains $H^0\mathcal{A}$. By Thomason's theorem, $\mathcal{A} = H^0\mathcal{B}$. Hence the inclusion $H^0\mathcal{A} \rightarrow H^0\mathcal{B}$ is essentially surjective on objects, and $\mathcal{A} \rightarrow \mathcal{B}$ is a quasi-equivalence. By Lemma 2.5, $\mathcal{A} \rightarrow \mathcal{B}$ has a quasi-inverse. \square

Perfect complexes. Let X be a scheme. A *strictly perfect complex* is a bounded complex of locally free, finite rank \mathcal{O}_X -modules. A *perfect complex* is a cohomologically bounded complex \mathcal{P}^\bullet of coherent sheaves of \mathcal{O}_X -modules which is locally quasi-isomorphic to a strictly perfect complex. Inside the bounded derived category of coherent sheaves $D^b \text{Coh}(X)$, one has a full triangulated subcategory $\text{Perf}(X)$ of perfect complexes.

We will need to consider dg enhancements of $\text{Perf}(X)$; that is, we want a pre-triangulated dg category \mathcal{C} and an equivalence of triangulated categories $\varepsilon: H^0(\mathcal{C}) \rightarrow \text{Perf}(X)$. When X is a projective scheme over a field \mathbb{K} , $\text{Perf}(X)$ has a dg enhancement $(\mathcal{C}, \varepsilon)$ which is unique: if $(\mathcal{C}', \varepsilon')$ is another then there is a *quasi-functor* $\phi: \mathcal{C} \rightarrow \mathcal{C}'$ such that $\varepsilon' \circ H^0(\phi) = \varepsilon$ [36]. Since we wish to work over more general base rings, and for computational purposes, we specify a dg enhancement of $\text{Perf}(X)$, valid for X a projective noetherian scheme, as follows.

Assume X is separated and noetherian. Fix an affine open covering \mathcal{U} of X . Define a dg category $\text{vect}(X)$ whose objects are locally free sheaves (=vector bundles) of finite rank, and whose hom-spaces, denoted $\text{R hom}^\bullet(\mathcal{E}, \mathcal{F})$, are Čech complexes:

$$\text{R hom}^\bullet(\mathcal{E}, \mathcal{F}) = (\check{C}^\bullet(\mathcal{U}; \underline{\text{Hom}}(\mathcal{E}, \mathcal{F})), \delta),$$

with δ the Čech differential. The cohomology of the Čech complex is

$$\text{R Hom}^\bullet(\mathcal{E}, \mathcal{F}) = \check{H}^\bullet(\mathcal{U}; \underline{\text{Hom}}(\mathcal{E}, \mathcal{F})) \cong \text{Ext}^\bullet(\mathcal{E}, \mathcal{F})$$

by [24, Theorem III.4.5] and the fact that $\text{Ext}^\bullet(\mathcal{E}, \mathcal{F}) \cong H^\bullet(\mathcal{E}^\vee \otimes \mathcal{F})$. Composition combines the cup product of Čech cochains with the composition of sheaf-morphisms. Whilst $\text{vect}(X)$ depends on the open covering, different choices lead to quasi-isomorphic dg categories (take the union of the two coverings). We now pass to the pre-triangulated dg category $\text{tw vect}(X)$ of twisted complexes. There is an embedding $H^0(\text{tw vect}(X)) \rightarrow \text{Perf}(X)$, mapping a twisted complex to its total

complex. This embedding is a quasi-equivalence, because every perfect complex is quasi-isomorphic to a strictly perfect complex [59, Prop. 2.3.1(d)].

Another approach to dg enhancement is to use injective resolutions; the equivalence of the injective and Čech approaches is shown in [47, Lemma 5.1], over fields; the proof remains valid over rings.

Grothendieck–Serre duality defines a CY structure for $\mathrm{tw} \, \mathrm{vect}(X)$ when X is equipped with a trivialization of the relative dualizing sheaf $\omega_{X/\mathbb{L}}$. In Theorem A, the functor ψ is trace-preserving.

2.3 Geometry of Weierstrass curves

2.3.1 Genus-one curves

We shall need to work with curves over the rings \mathbb{Z} and $\mathbb{Z}[[q]]$, and to this end we note some terminology for curves over schemes (cf. for example [17, 14]). A *curve* over a noetherian scheme S is a morphism of schemes $f: C \rightarrow S$ that is separated, flat and finitely presented, such that for every closed point $s \in S$ the fiber C_s is non-empty of pure dimension 1. The Euler characteristic $\chi(C_s, \mathcal{O}_{C_s})$ is then locally constant; when it is constant and equal to $1 - g$, and the geometric fibers are connected, we say that C has arithmetic genus g .

We shall always apply the restrictions that curves are to be proper, and that the fibres C_s are Cohen–Macaulay. This implies that one has a dualizing sheaf $\omega_{C/S}$, and where $C \rightarrow S$ is regular it coincides with the sheaf of differentials $\Omega_{C/S}^1$. A reminder on duality [13]: there is an intrinsic residue isomorphism of sheaves on S

$$\mathrm{res}: R^1 f_*(\omega_{C/S}) \rightarrow \mathcal{O}_S.$$

With the Yoneda (composition) product \smile , this defines the Serre duality pairing,

$$R\mathrm{Hom}_S^{1-i}(\mathcal{F}, \omega_{C/S}) \otimes R^i f_*(\mathcal{F}) \xrightarrow{\smile} R^1 f_*(\omega_{C/S}) \xrightarrow{\mathrm{res}} \mathcal{O}_S,$$

for any coherent sheaf \mathcal{F} .

A curve has arithmetic genus one if and only if $\mathcal{O}_C \cong \omega_{C/S}$, i.e., if and only if \mathcal{O}_C is a dualizing sheaf. If $\omega: \mathcal{O}_C \rightarrow \omega_{C/S}$ is an isomorphism then it composes with the residue map to give an isomorphism

$$\mathrm{tr}_\omega: R^1 f_*(\mathcal{O}_C) \rightarrow \mathcal{O}_S,$$

and a Serre duality pairing

$$R\mathrm{Hom}_S^{1-i}(\mathcal{F}, \mathcal{O}_C) \otimes R^i f_*(\mathcal{F}) \xrightarrow{\smile} R^1 f_*(\mathcal{O}_C) \xrightarrow{\mathrm{tr}_\omega} \mathcal{O}_S$$

which induces a perfect pairing on stalks at any closed point $s \in S$.

2.3.2 Weierstrass curves: definitions

Definition 2.14 An *abstract Weierstrass curve* (C, σ, ω) over S is a curve $f: C \rightarrow S$ of arithmetic genus one, such that each geometric fiber C_s is irreducible, equipped with a section $\sigma: S \rightarrow C$ of f and a specific isomorphism $\omega: \mathcal{O}_C \rightarrow \omega_{C/S}$. An *isomorphism* of abstract Weierstrass curves $(C_1, \sigma_1, \omega_1)$ and $(C_2, \sigma_2, \omega_2)$ over S is an isomorphism $g: C_1 \rightarrow C_2$ of S -schemes such that $g \circ \sigma_1 = \sigma_2$, and such that the map $g^*\omega_2: g^*\mathcal{O}_{C_2} \rightarrow g^*\omega_{C_2/S}$ coincides with ω_1 under the identifications $g^*\mathcal{O}_{C_2} \cong \mathcal{O}_{C_1}$ and $g^*\omega_{C_2/S} \cong \omega_{C_1/S}$ induced by g .

Definition 2.15 An *embedded Weierstrass curve* over $S = \text{Spec } R$ is a curve $C \subset \mathbb{P}_S^2$ embedded as a cubic

$$(13) \quad Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3 \quad (a_i \in R).$$

Such a curve comes with its point at infinity $p = [0 : 1 : 0]$, which defines a section σ of $C \rightarrow S$. It also comes with a standard differential ω , possibly with poles at the singular points: Writing $x = X/Z$ and $y = Y/Z$ away from σ , and writing the cubic equation as $w(x, y) := y^2 - x^3 + \dots = 0$, one has $\omega = dx/w_y$ at points where $w_y \neq 0$, and $\omega = -dy/w_x$ at points where $w_x \neq 0$.

Lemma 2.16 Assume that R is a normal ring (i.e., R is reduced and integrally closed in its total quotient ring). Then ω defines a section of the dualizing sheaf $\omega_{C/R}$.

Proof Let $R[\underline{a}] = R[a_1, a_2, a_3, a_4, a_6]$ —another normal ring. It will suffice to prove the assertion for the ‘universal Weierstrass curve’ $p: \mathcal{C} \rightarrow \text{Spec } R[\underline{a}]$ defined by (13), since the formation of the dualizing sheaf is compatible with the specialization to particular values of the a_i .

The scheme \mathcal{C} is normal: in the open set U where (13) is valid, a_6 is a function of the other variables, so projection $U \rightarrow \text{Spec } R[x, y, a_1, a_2, a_3, a_4]$ is an isomorphism, and $R[x, y, a_1, a_2, a_3, a_4]$ is normal. Along the section at infinity $\sigma = [0 : 1 : 0]$, the fibers of p are regular, and the base normal, so the total space is normal. The relative dualizing sheaf $\omega_{\mathcal{C}/R[\underline{a}]}$ is an invertible sheaf, since all its fibers are Gorenstein (being local complete intersections). The locus where the fibers of p are singular is defined by $w_x = w_y = 0$. This locus has codimension 2 in \mathcal{C} : it maps to the codimension 1 locus $\{\Delta = 0\} \subset \text{Spec } R[\underline{a}]$ defined by the vanishing of the discriminant, and it has codimension 1 in each fiber. Since ω is a section of $\omega_{\mathcal{C}/R[\underline{a}]}$ defined outside a codimension 2 subset of a normal scheme, it extends to a global section, by the algebraic counterpart to Hartogs’s theorem.² \square

²One can take this to be the statement that an integrally closed subring A of a field K is the intersection of the valuation rings in K which contain A [8, 5.22].

Thus an embedded Weierstrass curve functorially defines an abstract Weierstrass curve (C, σ, ω) . By Riemann–Roch, every abstract Weierstrass curve is isomorphic to an embedded one. To specify the embedding into \mathbb{P}^2 , one must give a basis of $H^0(\mathcal{O}_C(3\sigma))$ of the form $(y, x, 1)$, where 1 is the regular function with value 1, and $x \in H^0(\mathcal{O}_C(2\sigma))$. The denominator-free form of the argument is given at [29, p. 68], for instance.

2.3.3 Reparametrization group

The algebraic group $G \subset PGL(3)$ of elements which preserve Weierstrass form consists of matrices (up to scale) of the shape

$$(14) \quad \begin{bmatrix} u^3 & s & t \\ 0 & u^2 & r \\ 0 & 0 & 1 \end{bmatrix}, \quad u \in \mathbb{G}_m.$$

acting on homogeneous coordinates $[Y : X : Z]$ (the funny order for the coordinates makes the matrices triangular).

We shall call G the *reparametrization group* for embedded Weierstrass curves. It acts on embedded Weierstrass curves via the substitutions

$$x = u^2 x' + r, \quad y = u^3 y' + s x' + t.$$

The effects of a substitution on the Weierstrass coefficients are listed in [16] or [55]:

$$(15) \quad u a'_1 = a_1 + 2s$$

$$(16) \quad u^2 a'_2 = a_2 - s a_1 + 3r - s^2$$

$$(17) \quad u^3 a'_3 = a_3 + r a_1 + 2t$$

$$(18) \quad u^4 a'_4 = a_4 - s a_3 + 2r a_2 - (t + rs) a_1 + 3r^2 - 2st$$

$$(19) \quad u^6 a'_6 = a_6 + r a_4 + r^2 a_2 + r^3 - t a_3 - t^2 - r t a_1.$$

The unipotent subgroup $U \leq G$ of elements where $u = 1$ is the subgroup which preserves the differential ω . Thus if $g \in U$ then $g: C \rightarrow g(C)$ is an isomorphism of abstract Weierstrass curves.

The Lie algebra. The Lie algebra \mathfrak{g} of G is spanned by four vectors:

$$\partial_s := \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \partial_r := \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad \partial_t := \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

(these three span the Lie algebra \mathfrak{u} of U) and

$$\partial_u := \begin{bmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

The derivative of the G -action on

$$(20) \quad W := \operatorname{Spec} \mathbb{K}[a_1, a_2, a_3, a_4, a_6]$$

is an action of \mathfrak{g} on W by a Lie algebra homomorphism

$$(21) \quad \rho: \mathfrak{g} \rightarrow \operatorname{vec}(W),$$

which we think of as a map $\rho: \mathfrak{g} \times W \rightarrow W$. The partial derivative $(\partial\rho/\partial w)|_{w=0}: \mathfrak{g} \times W \rightarrow W$ makes W a \mathfrak{g} -module. We can form a differential graded Lie algebra (DGLA) concentrated in degrees 0 and 1,

$$(22) \quad \mathcal{L} = \{\mathfrak{g} \xrightarrow{d} W\}, \quad d(\xi) = \rho(\xi, 0),$$

whose bracket combines the Lie bracket of \mathfrak{g} with the module structure of W . Thus \mathcal{L} captures the truncation of ρ where we only work in a first-order neighborhood of $0 \in W$.

There are \mathbb{K}^\times -actions on \mathfrak{g} and on W , intertwined by d . The action on W is given by $\tau \cdot a_j = \tau^{-j} a_j$; that on \mathfrak{g} by $\tau \cdot \partial_s = \tau^{-1} \partial_s$, $\tau \cdot \partial_r = \tau^{-2} \partial_r$, $\tau \cdot \partial_t = \tau^{-3} \partial_t$, $\tau \cdot \partial_u = \partial_u$. We make W into a graded \mathbb{K} -module by declaring a_j to have degree $-j$. We make \mathfrak{g} into a graded \mathbb{K} -module by declaring ∂_s , ∂_t , ∂_r and ∂_u to have respective degrees -1 , -2 , -3 and 0 .

Explicitly, taking $(\partial_s, \partial_r, \partial_t, \partial_u)$ as basis for \mathfrak{g} , and $(a_1, a_2, a_3, a_4, a_6)$ as coordinates for W , one has

$$d = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

and

$$(23) \quad W^{\mathfrak{g}} := \operatorname{coker} d = \frac{\mathbb{K}}{(2)}^{a_1} [1] \oplus \frac{\mathbb{K}}{(3)}^{a_2} [2] \oplus \frac{\mathbb{K}}{(2)}^{a_3} [3] \oplus \mathbb{K}^{a_4} [4] \oplus \mathbb{K}^{a_6} [6].$$

$\ker d$ always contains the line $\mathbb{K} \cdot \partial_u$. When \mathbb{K} is a field, one can write $\ker d$ compactly as

$$(24) \quad \ker d = \mathbb{K} \oplus \frac{\mathbb{K}}{(2)}^{\partial_s} [1] \oplus \frac{\mathbb{K}}{(3)}^{\partial_r} [2] \oplus \frac{\mathbb{K}}{(2)}^{\partial_t} [3].$$

2.3.4 The cuspidal cubic

The cuspidal Weierstrass curve

$$(25) \quad C_{\text{cusp}} = \{y^2 - x^3 = 0\}$$

will play a special role in our story, stemming from the fact that the full subcategory of its derived category whose objects are the structure sheaf and the skyscraper at infinity is *formal*. Let $X = \mathbb{P}^1$, and let $p: \text{Spec } \mathbb{K} \rightarrow X$ be the \mathbb{K} -point $[0 : 1]$. Let z denote the standard affine coordinate $\mathbb{A}^1 \rightarrow \mathbb{P}^1$, $z \mapsto [z : 1]$. One has the structure sheaf \mathcal{O}_X , and inside it the sheaf $\mathcal{O}^p \subset \mathcal{O}_X$ of functions f such that $Df(p) = 0$ (that is, in terms of the local coordinate z , functions $f(z) = f(0) + O(z^2)$). Let X_{cusp} denote the scheme (X, \mathcal{O}^p) .

Lemma 2.17 *The abstract Weierstrass curve underlying C_{cusp} is isomorphic to*

$$(X_{\text{cusp}}, p, z^{-2}dz),$$

Proof The normalization of C_{cusp} is a non-singular rational curve \tilde{C} with a distinguished point c which maps to the cusp under the normalization map $\nu: \tilde{C} \rightarrow C_{\text{cusp}}$. We fix an isomorphism $X \rightarrow \tilde{C}$ mapping p to c . The map ν is a homeomorphism in the Zariski topology, and so defines a scheme-theoretic isomorphism $(\tilde{C}, \nu^*\mathcal{O}_{C_{\text{cusp}}}) \rightarrow (C_{\text{cusp}}, \mathcal{O}_{C_{\text{cusp}}})$. One has $\mathcal{O}_{C_{\text{cusp}}} \cong \mathcal{O}^p$: the local model near the cusp is the map of \mathbb{K} -algebras $\mathbb{K}[x, y]/(y^2 - x^3) \rightarrow \mathbb{K}[z]$ given by $x \mapsto z^2$ and $y \mapsto z^3$, whose image is $\mathbb{K}.1 \oplus z^2\mathbb{K}[z]$.

The \mathcal{O}^p -module of differentials $\Omega_{X_{\text{cusp}}}^1$ is given by the submodule of $\Omega_X^1(2p)$ (meromorphic differentials on X with a double pole at p) formed by the differentials with vanishing residue at p . In terms of the affine coordinate z near p , the differential of a function $g(z) = a + bz^2 + \dots$ is $dg = g'(z)dz$. The Weierstrass differential ω is given by $\omega = dx/(2y) = dy/(3x^2)$ (in characteristics 2 and 3 only one of these expressions makes sense). In terms of z , one has $\omega = z^{-2}dz$; this makes global sense because $\omega = -d(z^{-1})$. \square

3 Perfect complexes on Weierstrass curves

3.1 The two-object dg category associated with a Weierstrass curve

In this subsection we explain how to pass from a Weierstrass curve C to a two-object dg category \mathcal{B}_C with standard cohomology. Consider a genus-one curve C over

a noetherian affine scheme S . It has a dg category $\text{vect}(C)$, defined via an affine open covering, linear over the ring \mathcal{O}_S . The dg category $\text{tw vect}(C)$ for an abstract Weierstrass curve (C, σ, ω) over S has extra structure in the form of a trace pairing tr , as described in the introduction. It also has distinguished split-generators, namely, the structure sheaf and the skyscraper $\mathcal{O}_{C,\sigma} = \sigma_* \mathcal{O}_S$ at σ (more properly, its locally-free resolution $\mathcal{O} \rightarrow \mathcal{O}(\sigma)$):

Lemma 3.1 *For a Weierstrass curve $C \rightarrow \mathcal{O}_S$ over a noetherian affine scheme, one has*

$$\text{tw vect}(C) = \langle \mathcal{O}_C, \mathcal{O}_{C,\sigma} \rangle.$$

Here $\langle \cdot \rangle$ denotes the smallest dg subcategory of $\text{tw vect}(C)$ containing the objects \mathcal{O}_C and $\mathcal{O}_{C,\sigma}$ that is closed under quasi-isomorphisms, shifts, mapping cones and passing to idempotents.

Proof We claim first that, if $\mathcal{O}(1)$ is a very ample line bundle on C , then $\text{tw vect}(C)$ is split-generated (i.e., generated under quasi-isomorphisms, shifts, mapping cones and passing to idempotents) by the twists $\{\mathcal{O}(n)\}_{n < 0}$. The argument is as in [47, Lemma 5.4], which Seidel attributes to Kontsevich. Take a locally free sheaf V on C . By Serre's theorem that very ample implies ample [24, Theorem II.5.17], which is valid for noetherian projective schemes, one can find an epimorphism $\mathcal{O}(m)^{\oplus r} \rightarrow V$ for some $m \ll 0$. Iteratively, one can find for each k a left resolution

$$0 \rightarrow V' \rightarrow \mathcal{O}(m_k)^{\oplus r_k} \rightarrow \cdots \rightarrow \mathcal{O}(m_1)^{\oplus r_1} \rightarrow V \rightarrow 0.$$

There results an exact triangle in $D(C)$

$$\{\mathcal{O}(m_k)^{\oplus r_k} \rightarrow \cdots \rightarrow \mathcal{O}(m_1)^{\oplus r_1}\} \rightarrow V \rightarrow V'[k] \xrightarrow{+}.$$

Now, $\text{Ext}_R^k(V, V') = H^k(C, V^\vee \otimes V')$, and if we take $k > \dim C = 1 + \dim R$, this Ext-module must vanish. Consequently, the exact triangle splits and defines a quasi-isomorphism

$$\{\mathcal{O}(m_k)^{\oplus r_k} \rightarrow \cdots \rightarrow \mathcal{O}(m_1)^{\oplus r_1}\} \rightarrow V \oplus V'[k].$$

Thus V is a direct summand in the object on the LHS. Note also that V' is a perfect complex, because it is the mapping cone of a map of perfect complexes; therefore, V' is quasi-isomorphic to a strictly perfect complex. This proves that every locally free sheaf lies in the split-closure of the collection $\{\mathcal{O}(n)\}_{n < 0}$. It follows that the same is true of every object of $\text{tw vect}(C)$.

Note next that $\langle \mathcal{O}_C, \mathcal{O}_{C,\sigma} \rangle$ includes $\mathcal{O}_C(n\sigma)$ for each $n \leq 0$, by a straightforward induction. But $\mathcal{O}(3\sigma)$ is a very ample line bundle, so now the claim completes the proof. \square

A two-object subcategory. Let \mathcal{B}_C denote the full dg subcategory of $\text{tw vect}(C)$ with the two objects \mathcal{O}_C and $\mathcal{O}_{C,\sigma}$ and with the trace map tr_ω . It is defined up to quasi-isomorphisms acting trivially on cohomology. To be precise, we shall define \mathcal{B}_C using the Čech complexes associated with an affine open covering \mathcal{U} . If we pick two coverings \mathcal{U}_1 and \mathcal{U}_2 , we get dg categories $\mathcal{B}_{\mathcal{U}_1}$ and $\mathcal{B}_{\mathcal{U}_2}$, and a zigzag of quasi-isomorphisms

$$\mathcal{B}_{\mathcal{U}_1} \leftarrow \mathcal{B}_{\mathcal{U}_1 \cup \mathcal{U}_2} \rightarrow \mathcal{B}_{\mathcal{U}_2}.$$

The cohomology category $H^*\mathcal{B}_C$ is truly canonical—defined up to canonical isomorphism.

An isomorphism $g: C_1 \rightarrow C_2$ of abstract Weierstrass curves is a homeomorphism g , together with a local isomorphism of sheaves of S -modules $g^\#: \mathcal{O}_{C_2} \rightarrow g_*\mathcal{O}_{C_1}$, respecting the sections and differentials. That means, first, that $\sigma_2: S \rightarrow C_2$ is the composite $g \circ \sigma_1$; this implies a canonical isomorphism $\mathcal{O}_{C_2, \sigma_2} \rightarrow g_*\mathcal{O}_{C_1, \sigma_1}$. The isomorphism g induces isomorphisms between the abelian categories of coherent sheaves on C_1 and C_2 , preserving the objects \mathcal{O} and \mathcal{O}_σ . This naturally extends to an isomorphism of dg categories $g_*: \mathcal{B}_{C_1} \rightarrow \mathcal{B}_{C_2}$, provided that we use an open covering \mathcal{U} for C_1 and $g(\mathcal{U})$ for C_2 . Thus, if we have $g_{12}: C_1 \rightarrow C_2$ and $g_{23}: C_2 \rightarrow C_3$ with composite g_{13} , then the composite isomorphism $g_{23*} \circ g_{12*}: \mathcal{B}_{C_1} \rightarrow \mathcal{B}_{C_3}$ coincides with g_{13*} , provided again that we use the coverings \mathcal{U} , $g_{12}(\mathcal{U})$ and $g_{13}(\mathcal{U})$. If we do not, then we get instead the formal composite of chains of quasi-isomorphisms:

$$\begin{array}{ccccccc} & & \mathcal{B}_{g_{12}(\mathcal{U}_1) \cup \mathcal{U}_2} & & \mathcal{B}_{g_{23}(\mathcal{U}_2) \cup \mathcal{U}_3} & & \\ & & \downarrow & \searrow & \downarrow & \searrow & \\ \mathcal{B}_{\mathcal{U}_1} & \longrightarrow & \mathcal{B}_{g_{12}(\mathcal{U}_1)} & \longrightarrow & \mathcal{B}_{\mathcal{U}_2} & \longrightarrow & \mathcal{B}_{g_{23}(\mathcal{U}_2)} & \longrightarrow & \mathcal{B}_{\mathcal{U}_3} \end{array}$$

The cohomology category. We shall be interested in the map which assigns to each Weierstrass curve (C, ω, σ) a graded-linear cohomology category A_C and a dg category-with-trace,

$$(C, \omega, \sigma) \mapsto \mathcal{B}_C,$$

with an isomorphism $H^*\mathcal{B}_C \cong A_C$, defined up to quasi-isomorphisms acting trivially on A_C .

Proposition 3.2 *The category A_C is independent of the abstract Weierstrass curve. Precisely: There is an \mathcal{O}_S -linear graded category A with two objects \mathcal{O} and \mathcal{O}_σ , equipped with a trace map tr , such that the following holds: For any abstract Weierstrass curve $(f: C \rightarrow S, \omega, \sigma)$ the cohomology category $A_C = H^*(\mathcal{B}_C)$ is trace-preservingly isomorphic to A in such a way that if $C_1 \rightarrow C_2$ is any isomorphism of Weierstrass curves then the resulting map $A \rightarrow A$ is the identity.*

In other words, the category of Ext-modules between \mathcal{O}_C and $\mathcal{O}_{C,\sigma}$ is independent of (C, σ, ω) as a graded \mathbb{K} -linear category with trace.

To prove the proposition, we examine the structure of $H^*\mathcal{B}_C$. Writing $\mathcal{O} = \mathcal{O}_C$ and $\mathcal{O}_\sigma = \mathcal{O}_{C,\sigma} := \sigma_*\mathcal{O}_S$, one has canonical isomorphisms

$$\begin{aligned} \mathrm{RHom}_S(\mathcal{O}, \mathcal{O}) &\cong \mathrm{R}f_*(\mathcal{O}) \cong \mathcal{O}_S \oplus \mathrm{R}^1f_*(\mathcal{O}), \\ \mathrm{RHom}_S(\mathcal{O}_\sigma, \mathcal{O}_\sigma) &\cong \Lambda^*(\sigma^*TC) \cong \mathcal{O}_S \oplus \sigma^*\mathcal{T}_C \end{aligned}$$

Thus both endomorphism spaces are 2-dimensional and sit in degrees 0 and 1. The trace isomorphisms

$$\begin{aligned} \mathrm{tr}_\omega: \mathrm{R}^1f_*(\mathcal{O}) &\rightarrow \mathcal{O}_S \\ \mathrm{tr}_\omega: \sigma^*\mathcal{T}_C &\rightarrow \mathcal{O}_S \end{aligned}$$

are, in the first case, the one that we have discussed (the composite of ω and the residue pairing) and in the second case the pullback by σ of the composite

$$\mathcal{T}_C \rightarrow \mathcal{T}_C \otimes_{\mathcal{O}} \mathcal{O} \xrightarrow{\mathrm{id} \otimes \omega} \mathcal{T}_C \otimes \omega_{C/S} \xrightarrow{\mathrm{ev}_\sigma} \mathcal{O},$$

where ev_σ is the map defined by evaluating ω —viewed as a differential—on tangent vectors at σ . One has

$$\mathrm{RHom}_S(\mathcal{O}, \mathcal{O}_\sigma) = \mathrm{R}^0\mathrm{Hom}_S(\mathcal{O}, \mathcal{O}_\sigma) \cong f_*(\mathcal{O}_\sigma) \cong \mathcal{O}_S.$$

Finally, one has isomorphisms

$$\begin{aligned} \mathrm{RHom}_S(\mathcal{O}_\sigma, \mathcal{O}) &= \mathrm{R}^1\mathrm{Hom}_S(\mathcal{O}_\sigma, \mathcal{O}) \\ &\cong \mathrm{R}^1\mathrm{Hom}_S(\mathcal{O}_\sigma, \mathcal{O}) \otimes_{\mathcal{O}_S} \mathcal{O}_S \\ &\cong \mathrm{R}^1\mathrm{Hom}_S(\mathcal{O}_\sigma, \mathcal{O}) \otimes_{\mathcal{O}_S} \mathrm{R}^0\mathrm{Hom}_S(\mathcal{O}, \mathcal{O}_\sigma) \\ &\xrightarrow{\sim} \mathrm{R}^1\mathrm{Hom}_S(\mathcal{O}_\sigma, \mathcal{O}_\sigma) \\ &\xrightarrow{\mathrm{tr}_\omega} \mathcal{O}_S. \end{aligned}$$

We now describe the category A demanded by the proposition above. Let $X = \mathcal{O}$ and $Y = \mathcal{O}_\sigma$. We have seen how to use ω to obtain algebra isomorphisms

$$\mathrm{End}(Y) \cong \Lambda^*(\mathcal{O}_S[-1]) \cong \mathrm{End}(Y)$$

such that the trace maps correspond to the identity map of \mathcal{O}_S . We also have exhibited isomorphisms $\mathrm{Hom}(X, Y) = \mathcal{O}_S$ and $\mathrm{Hom}(Y, X) = \mathcal{O}_S[-1]$. The composition maps are mostly dictated by the requirements of grading and unitality. The interesting ones are

$$\mathrm{Hom}^1(Y, X) \otimes \mathrm{Hom}^0(X, Y) \rightarrow \mathrm{Hom}^1(Y, Y), \quad \mathrm{Hom}^0(X, Y) \otimes \mathrm{Hom}^1(Y, X) \rightarrow \mathrm{Hom}^1(X, Y).$$

These are both given by the multiplication of functions

$$\mathcal{O}_S \otimes \mathcal{O}_S \rightarrow \mathcal{O}_S.$$

The objects X and Y and their morphisms form a graded-linear CY category (A, tr) , independent of C .

Differential graded structure. While the cohomology category $A = A_C$ is independent of C , the dg structure of \mathcal{B}_C is fully sensitive to the curve C :

Theorem 3.3 (dg comparison theorem) *Work over a field \mathbb{K} .*

- (1) *Let \mathcal{B} be a dg category with trace such that $H^*(\mathcal{B}) \cong A$. Then there exist an abstract Weierstrass curve (C, ω, σ) and a trace-preserving A_∞ -quasi-isomorphism $\mathcal{B} \rightarrow \mathcal{B}_C$.*
- (2) *If (C, ω, σ) and (C', ω', σ') give rise to quasi-isomorphic dg categories with trace, i.e., \mathcal{B}_C is related to $\mathcal{B}_{C'}$ by a zig-zag of trace-preserving isomorphisms, then $(C, \sigma, \omega) \cong (C', \sigma', \omega')$.*

We shall prove clause (2) presently. The existence clause (1) is plausible because one has $H^0(\Lambda^2 T_C^*) = 0$ and $H^2(\mathcal{O}_C) = 0$. As a result, $\text{Perf}(C)$ has no Poisson deformations and no non-commutative deformations, and it is reasonable to expect all deformations of $\text{Perf}(C)$ to be geometric. However, clause (1) will be obtained as a corollary of a more detailed statement, Theorem C, and it will in fact be the latter result which we eventually use, not Theorem 3.3.

Lemma 3.4 *If (C, σ, ω) and (C', σ', ω') are abstract Weierstrass curves over a ring \mathbb{K} , and $\theta: \mathcal{B}_C \rightarrow \mathcal{B}_{C'}$ is an A_∞ quasi-isomorphism, then there is an isomorphism of \mathbb{K} -schemes $C \rightarrow C'$ carrying σ to σ' .*

Proof Let $\mathcal{S}: \text{tw vec } C \rightarrow \text{tw vec } C$ denote the twist functor along the spherical object \mathcal{O}_σ . One also has the functor $\mathcal{O}(\sigma) \otimes \cdot$, and \mathcal{S} and $\mathcal{O}(\sigma)$ induce the same functor on $H^0(\text{tw vec } C)$, up to natural isomorphism, by [51, (3.11)]. The argument does not require the base ring to be a field, nor does it require the curve to be an elliptic curve; it just requires C to be smooth in an open neighborhood of the image of σ .

Define twisted complexes $T_n \in \text{ob tw}(\text{vec } C)$ recursively by setting $T_0 = \mathcal{O}$ and $T_{n+1} = \mathcal{S}(T_n)$. Notice that T_n lies in $\text{ob tw}(\mathcal{B}_C)$, and since spherical twist functors are determined by the twisting object, T_n is determined by \mathcal{B}_C .

More precisely, introduce the corresponding twisted complexes T'_n over C' . Notice that θ necessarily carries \mathcal{O}_C to $\mathcal{O}_{C'}$ and \mathcal{O}_σ to $\mathcal{O}_{\sigma'}$, because of the asymmetry in

the dimensions of the the graded Hom-spaces. The extension $\mathrm{tw} \theta: \mathrm{tw} \mathcal{B}_C \rightarrow \mathcal{B}_{C'}$ satisfies $(\mathrm{tw} \theta)(T_n) \cong T'_n$ in $H^0(\mathrm{tw} \mathcal{B}_{C'})$.

In particular, θ induces isomorphisms $\mathrm{Hom}_{H(\mathrm{tw} \mathcal{B}_C)}(T_m, T_n) \cong \mathrm{Hom}_{H(\mathrm{tw} \mathcal{B}_{C'})}(T'_m, T'_n)$, compatible with composition. But by induction, T_n is isomorphic in $H^0(\mathrm{tw} \mathrm{vec} C)$ to the line bundle $\mathcal{L}_n := \mathcal{O}(n\sigma)$. Hence one has a categorical isomorphism from the full subcategory of $H^0(\mathrm{tw} \mathrm{vec} C)$ on objects $\{\mathcal{L}_n\}_{n \geq 0}$ to the corresponding category $\{\mathcal{L}'_n\}_{n \geq 0}$ inside $H^0(\mathrm{tw} \mathrm{vec} C')$.

There are isomorphisms $H^0(\mathcal{L}_m) = \mathrm{Hom}^0(\mathcal{O}, \mathcal{L}_m) \rightarrow \mathrm{Hom}^0(\mathcal{L}_n, \mathcal{L}_{m+n})$ given by tensoring with the identity map of \mathcal{L}_n . Using these isomorphisms one identifies the composition map

$$\mathrm{Hom}(\mathcal{L}_m, \mathcal{L}_{m+n}) \otimes \mathrm{Hom}^0(\mathcal{L}_0, \mathcal{L}_m) \rightarrow \mathrm{Hom}^0(\mathcal{L}_0, \mathcal{L}_{m+n})$$

with the tensor product of sections $H^0(\mathcal{L}_n) \otimes H^0(\mathcal{L}_m) \rightarrow H^0(\mathcal{L}_m \otimes \mathcal{L}_n) = H^0(\mathcal{L}_{m+n})$. Consequently, the graded ring $R = \bigoplus_{m \geq 0} H^0(\mathcal{L}_m)$ is determined up to isomorphism by \mathcal{B}_C . Since \mathcal{L}_1 is ample, this graded ring determines the pair consisting of the \mathbb{K} -scheme C together with the line bundle \mathcal{L}_1 : we have $C \cong \mathrm{Proj} R$, with \mathcal{L}_1 the twisting sheaf. Thus we obtain isomorphisms $(C, \mathcal{L}_1) \rightarrow (C', \mathcal{L}'_1)$. Sections of $\mathcal{L}_1^{\otimes 2}$ either have no zeros, or they vanish (to order 2) at σ . Hence the isomorphism $C \rightarrow C'$ must carry σ to σ' . \square

Proof of Theorem 3.3 (2). Over a field, dg quasi-isomorphisms have A_∞ quasi-inverses by the HPL (see subsection (2.1.4)). Thus if a zig-zag of dg quasi-isomorphisms relates \mathcal{B}_C to $\mathcal{B}_{C'}$ then there is an A_∞ quasi-isomorphism $\mathcal{B}_C \rightarrow \mathcal{B}'_C$. By the lemma, there is then an isomorphism of pointed curves $(C, \sigma) \rightarrow (C', \sigma')$. If the dg quasi-isomorphisms are trace-preserving then one can take the A_∞ quasi-isomorphism $\mathcal{B}_C \rightarrow \mathcal{B}'_C$ to preserve traces also. In particular, the isomorphism $H^1(\mathcal{O}_C) \rightarrow H^1(\mathcal{O}_{C'})$ respects the generators supplied by ω and ω' . That is, one has a Weierstrass isomorphism $(C, \sigma, \omega) \rightarrow (C', \sigma', \omega')$. \square

The cuspidal cubic $C_{\mathrm{cusp}} = \{y^2 - x^3 = 0\}$ has the following special property, which already appeared in [35]:

Lemma 3.5 *The dga $\mathcal{B}_{\mathrm{cusp}} := \mathcal{B}_{C_{\mathrm{cusp}}}$ is formal.*

move
proof
later!
the
to

Proof We may transfer the dg structure of $\mathcal{B}_{\mathrm{cusp}}$ to a minimal A_∞ -structure on $A = H^* \mathcal{B}_{\mathrm{cusp}}$. The transfer of dg structure will be described in detail in the proof of Lemma 5.2. The goal, then, is to prove that the A_∞ structure maps μ^d vanish for $d > 2$.

C_{cusp} is the curve $\{Y^2Z = X^3\} \subset \mathbb{P}^2$. The multiplicative group \mathbb{G}_m acts on C_{cusp} by $t \cdot (X, Y, Z) = (t^{-2}X, t^{-3}Y, Z)$, preserving the point $\sigma = [0 : 1 : 0]$, and therefore acts on $\mathcal{B}_{\text{cusp}}$. The action of \mathbb{G}_m on $\mathcal{B}_{\text{cusp}}$ induces an action on the cohomology A , and the transfer will be set up equivariantly so that the resulting A_∞ -structure has the property that $\mu^d(t \cdot a_d, \dots, t \cdot a_1) = t \cdot \mu^d(a_d, \dots, a_1)$.

A short computation leads to the following conclusion: the weight of the \mathbb{G}_m -action on a hom-space $\text{Hom}_A^k(X, X')$ (where X is \mathcal{O} or \mathcal{O}_σ , ditto X') is equal to the degree k . Now take (X_0, \dots, X_d) a sequence of objects (\mathcal{O} or \mathcal{O}_σ), and take $a_j \in \text{hom}_{\mathcal{B}_{\text{cusp}}}^{k_j}(X_{j-1}, X_j)$. For the equation $\mu^d(t \cdot a_d, \dots, t \cdot a_1) = t \cdot \mu^d(a_d, \dots, a_1)$ to hold, one must have $k_1 + \dots + k_d + 2 - d = k_1 + \dots + k_d$, i.e., $d = 2$. \square

3.2 Stable vector bundles on \mathcal{T}_0

When we come to prove Theorem A, clause (iii), we will need to apply Thomason's theorem about Grothendieck groups [58], and for that we shall need to know $K_0(\mathcal{T}_0)$. We think of \mathcal{T}_0 , the central fiber of the Tate curve, as the curve over $\text{Spec } \mathbb{Z}$ obtained from \mathbb{P}^1 by identifying $p = [1 : 0]$ and $q = [0 : 1]$. By definition, $K_0(\mathcal{T}_0)$ is the Grothendieck group of the abelian category of vector bundles (locally free sheaves of finite rank) on \mathcal{T}_0 . It can also be thought of as $K_0(\text{Perf } \mathcal{T}_0)$, the Grothendieck group of the triangulated category of perfect complexes.³ The proof of the following lemma is more substantial than one might expect. As partial justification, we point out that $K_0(\mathcal{T}_0)$ is an *absolute* invariant of the scheme \mathcal{T}_0 —it is not defined ‘relative to $\text{Spec } \mathbb{Z}$ ’—and that \mathcal{T}_0 is *2-dimensional* as a scheme.

Lemma 3.6 *The map $(\text{rank}, \det): K_0(\mathcal{T}_0) \rightarrow \mathbb{Z} \oplus \text{Pic}(\mathcal{T}_0)$ is an isomorphism. Thus a vector bundle on \mathcal{T}_0 with trivial determinant is stably trivial.*

Proof Let $K_0(R)$ denote the Grothendieck group of finitely-generated projective modules over the commutative ring R . We also have the reduced group $\widehat{K}_0(R) = \ker(\text{rank}: K_0(R) \rightarrow \mathbb{Z})$ and the group of ‘stable endomorphisms’ $K_1(R)$. For the following standard results in K-theory we refer to the text [63] (see in particular the ‘Fundamental theorem for K_1 ’ (3.6)). We have $\widehat{K}_0(\mathbb{Z}) = 0$ since \mathbb{Z} is a PID. The units R^\times are always a subgroup of $K_1(R)$, and one has $K_1(\mathbb{Z}) = \mathbb{Z}^\times$. Since \mathbb{Z} is a regular ring, the inclusion-induced maps $\widehat{K}_0(\mathbb{Z}) \rightarrow \widehat{K}_0(\mathbb{Z}[t])$ and $K_1(\mathbb{Z}) \rightarrow K_1(\mathbb{Z}[t])$ are isomorphisms. One has a split injection $K_1(\mathbb{Z}[t]) \rightarrow K_1(\mathbb{Z}[t, t^{-1}])$, induced by the natural map $\mathbb{Z}[t] \rightarrow \mathbb{Z}[t, t^{-1}]$, whose cokernel is $K_0(\mathbb{Z})$. Hence $K_1(\mathbb{Z}[t, t^{-1}]) \cong \mathbb{Z} \oplus \mathbb{Z}^\times$.

³That is, the abelian group generated by the objects, with a relation $[B] = [A] + [C]$ for each distinguished triangle $A \rightarrow B \rightarrow C \rightarrow A[1]$.

There is a group $K_0(\mathbb{Z}[t] \text{ on } (t))$ of complexes of f.g. projective $\mathbb{Z}[t]$ -modules whose cohomology is bounded and supported on the ideal (t) ; and an exact sequence

$$K_1(\mathbb{Z}[t]) \rightarrow K_1(\mathbb{Z}[t, t^{-1}]) \xrightarrow{\partial} K_0(\mathbb{Z}[t] \text{ on } (t)) \rightarrow \widehat{K}_0(\mathbb{Z}[t])$$

[59, (5.1)]. From our discussion, we see that the first map has cokernel \mathbb{Z} and that $\widehat{K}_0(\mathbb{Z}[t]) = 0$, whence $K_0(\mathbb{Z}[t] \text{ on } (t)) \cong \mathbb{Z}$.

Now consider the ‘node’ $R = \mathbb{Z}[x, y]/(xy)$, and the normalization map $\nu: R \rightarrow \mathbb{Z}[t] \oplus \mathbb{Z}[t]$, namely, $\nu(f) = (f(x, 0), f(0, y))$. One has $\widehat{K}_0(\mathbb{K}[t] \oplus \mathbb{K}[t]) = 0$. The kernel of $\nu_*: \widehat{K}_0(R) \rightarrow \widehat{K}_0(\mathbb{Z}[t] \oplus \mathbb{Z}[t])$ is also zero, because a f.g. projective R -module M is determined by $\nu_* M$ —which is given by a pair of f.g. projective $\mathbb{Z}[t]$ -modules M_1 and M_2 —and ‘descent data’, an isomorphism $\theta: M_1/tM_1 \rightarrow M_2/tM_2$. By stabilizing, we may assume that M_1 and M_2 are free, and choosing appropriate bases we can make θ the identity matrix. Hence $\widehat{K}_0(R) = 0$.

We have (from [59, (5.1)] again) a commutative diagram with exact rows
(26)

$$\begin{array}{ccccccc} K_1(R) & \longrightarrow & K_1(R[x^{-1}, y^{-1}]) & \longrightarrow & K_0(R \text{ on } (x, y)) & \longrightarrow & 0 \\ \downarrow \nu_* & & \cong \downarrow \nu_* & & \downarrow \nu_* & & \\ 0 \longrightarrow & \bigoplus^2 K_1(\mathbb{Z}[t]) & \longrightarrow & \bigoplus^2 K_1(\mathbb{Z}[t, t^{-1}]) & \longrightarrow & \bigoplus^2 K_0(\mathbb{Z}[t] \text{ on } (t)) & \longrightarrow 0 \end{array}$$

The middle vertical arrow is an isomorphism, since it is induced by a ring isomorphism. Note that $K_1(R)$ contains the units \mathbb{Z}^\times . Chasing the diagram, we see that $K_0(R \text{ on } (x, y)) \cong \mathbb{Z}^2 \oplus S$, where the summand \mathbb{Z}^2 mapped isomorphically by ν_* to $\bigoplus^2 K_0(\mathbb{Z}[t] \text{ on } (t))$, and S is either \mathbb{Z}^\times or 0.

We now switch from rings to schemes, referring to [59] for foundational matters. If Z is a closed subscheme of X , $K_0(X \text{ on } Z)$ is the Grothendieck group of the triangulated category of perfect complexes on X whose cohomology sheaves are coherent and supported on Z .

The map $\det: \widehat{K}_0(\mathcal{T}_0) \rightarrow \text{Pic}(\mathcal{T}_0)$ is surjective—apply it to a line bundle. We must prove its injectivity. Let $\nu: \mathbb{P}^1 \rightarrow \mathcal{T}_0$ be the normalization map. Let Z be the closure of the image of the nodal section of \mathcal{T}_0 , and $j: U \rightarrow \mathcal{T}_0$ the inclusion of the open complement of Z . There is a commutative diagram with exact rows

$$(27) \quad \begin{array}{ccccccc} K_1(U) & \xrightarrow{\partial} & K_0(\mathcal{T}_0 \text{ on } Z) & \xrightarrow{\mu} & \widehat{K}_0(\mathcal{T}_0) & \xrightarrow{j^*} & \widehat{K}_0(U) \\ \nu^* \downarrow = & & \nu^* \downarrow & & \nu^* \downarrow & & \nu^* \downarrow = \\ K_1(U) & \xrightarrow{\partial'} & K_0(\mathbb{P}^1 \text{ on } \overline{\{p, q\}}) & \xrightarrow{\mu'} & \widehat{K}_0(\mathbb{P}^1) & \longrightarrow & \widehat{K}_0(U) \end{array}$$

The groups $K_0(X \text{ on } Z)$ have an excision property [59, (3.19)], which tells us for instance that restriction induces an isomorphism $K_0(\mathcal{T}_0 \text{ on } Z) = K_0(V \text{ on } Z)$, where V is any open neighborhood of Z . Further, it tells us that $K_0(V \text{ on } Z) = K_0(\widehat{V}_Z \text{ on } \widehat{Z})$, where \widehat{V}_Z is the completion of V along Z , and \widehat{Z} the image of Z . Similarly,

$$K_0(\mathbb{P}^1 \text{ on } \overline{\{p, q\}}) = K_0(\widehat{\mathbb{P}^1_{\{p, q\}}} \text{ on } \{\widehat{p}, \widehat{q}\}) = \bigoplus^2 K_0(\mathbb{Z} \llbracket t \rrbracket \text{ on } (t)) = \bigoplus^2 K_0(\mathbb{Z}[t] \text{ on } (t)).$$

Now, \widehat{V}_Z is isomorphic to $\text{Spec } \mathbb{Z} \llbracket x, y \rrbracket / (xy)$, and by naturality of normalization, the map $\nu: \widehat{V}_Z \rightarrow \widehat{\mathbb{P}^1_{\{p, q\}}}$ corresponds to the normalization map $\nu: \text{Spec } \mathbb{Z} \llbracket x, y \rrbracket / (xy) \rightarrow \text{Spec } \bigoplus^2 \mathbb{Z} \llbracket t \rrbracket$, which is itself the completion of $\nu: \text{Spec } \mathbb{Z}[x, y] / (xy) \rightarrow \text{Spec } \bigoplus^2 \mathbb{Z}[t]$. The upshot is that the map $\nu^*: K_0(\mathcal{T}_0 \text{ on } Z) \rightarrow K_0(\mathbb{P}^1 \text{ on } \overline{\{p, q\}})$ in (27) can be identified with the arrow $\nu_*: K_0(R \text{ on } (x, y)) \rightarrow \bigoplus^2 K_0(\mathbb{Z}[t] \text{ on } (t))$ in (26). Hence $\ker [K_0(\mathcal{T}_0 \text{ on } Z) \rightarrow K_0(\mathbb{P}^1 \text{ on } \overline{\{p, q\}})] = S$.

Now take $e \in \widehat{K}_0(\mathcal{T}_0)$ with $\det e = 1 \in \text{Pic}(\mathcal{T}_0)$. Then $\det \nu^* e = 1 \in \text{Pic}(\mathbb{P}^1)$, and $\nu^* e = 0 \in \widehat{K}_0(\mathbb{P}^1)$ [59, (4.1)]. Hence $j^* e = 0 \in \widehat{K}_0(U)$. Thus e is the image $\mu(f)$ of some $f \in K_0(\mathcal{T}_0 \text{ on } Z)$, and $\nu^* f$ maps to zero in $\widehat{K}_0(\mathbb{P}^1)$. So $\nu^* f = \partial' g$ for some $g \in K_1(U)$. Then $f - \partial g \in S \subset K_0(\mathcal{T}_0 \text{ on } Z)$, and so $e \in \mu(S)$.

Let $\ell \in \text{Pic}(\mathcal{T}_0)$ be the line-bundle obtained from the trivial line bundle on \mathbb{P}^1 by gluing the fibers over the $[0 : 1]$ and $[1 : 0]$ by the map $(-1) \in \mathbb{Z}^\times$. Let $x = [\ell] - 1 \in \widehat{K}_0(\mathcal{T}_0)$. Then $\det x \neq 1 \in \text{Pic}(\mathcal{T}_0)$, so $x \neq 0$. We have $j^* x = 0$, so $x = \mu(y)$, say. Also $\nu^* x = 0$; hence $\mu'(\nu^* y) = 0$, and so $\nu^* y = \partial' z$, say, and $y - \partial z \in S$. If S were zero, we would have $x = \mu \circ \partial z = 0$. Hence S is not zero, which we have seen implies that $S = \mathbb{Z}^\times$, and therefore $\mu(S) = \{0, x\}$. So our class e from the previous paragraph must be zero. \square

4 Hochschild cohomology via algebraic geometry

In the previous section we set up a two-object, graded linear category A , the cohomology $A = A_C = H^* \mathcal{B}_C$ of a dg category associated with an arbitrary Weierstrass curve. We can view (the direct sum of the hom-spaces in) A as a \mathbb{K} -algebra. We stress that A does not depend on the choice of C , except via its base ring \mathbb{K} . In this section, assuming \mathbb{K} to be a field, we compute the Hochschild cohomology $\text{HH}^*(A, A)$ as a bigraded algebra; or more precisely, the truncated version $\text{HH}^*(A, A)^{\leq 0}$ relevant to non-curved A_∞ deformations [47]. Later (Theorem 5.6) we shall complete the picture by determining the Gerstenhaber bracket. Our main result is as follows:

Theorem B *Let \mathbb{K} be any field, and let A be the graded \mathbb{K} -algebra described in the previous section. Introduce the commutative graded \mathbb{K} -algebra*

$$T = \mathbb{K}[x, y]/(y^2 - x^3, 2y, -3x^2),$$

concentrated in degree zero. Make it a bigraded algebra by assigning the following internal degrees s to the generators:

$$s(x) = 2; \quad s(y) = 3.$$

Introduce also the free graded-commutative graded algebra

$$S^\bullet = \mathbb{K}[\beta, \gamma] = \mathbb{K}[\beta] \otimes \Lambda[\gamma], \quad \deg \beta = 2, \quad \deg \gamma = 1,$$

made a bigraded algebra by assigning internal degrees

$$s(\beta) = -6; \quad s(\gamma) = \begin{cases} 0, & 6 \neq 0 \in \mathbb{K} \\ -2, & 3 = 0; \\ -3, & 2 = 0. \end{cases}$$

Define the bigraded algebra $Q^\bullet = T \otimes S^\bullet$, where $\deg(\tau \otimes \sigma) = \deg \sigma + \deg \tau$ and $s(\tau \otimes \sigma) = s(\sigma) + s(\tau)$ for bihomogeneous elements $\tau \in T$ and $\sigma \in S^\bullet$. Let $Q^{\bullet, \leq 0}$ denote the subalgebra of Q^\bullet spanned by those classes with $s \leq 0$. Then there is a canonical isomorphism of bigraded algebras

$$\mathrm{HH}^\bullet(A, A)^{\leq 0} \rightarrow Q^{\bullet, \leq 0}.$$

(The bihomogeneous classes of Q^\bullet with $s > 0$ —i.e., the bihomogeneous classes omitted in $Q^{\bullet, \leq 0}$ —are spanned by x and $x\gamma$ when $6 \neq 0$; by x , x^2 and $x^2\gamma$ when $3 = 0$; and by x , y , xy and $xy\gamma$ when $2 = 0$.)

Remark. By definition, $\mathrm{HH}^{r+s}(A, A)^s = \mathrm{Ext}_{(A, A)}^r(A, A[s])$, the r th derived homomorphism $A \rightarrow A[s]$ in the category of graded (A, A) -bimodules. Composition of bimodule-homomorphisms yields, on the derived level, an associative product making $\mathrm{HH}^\bullet(A, A)$ a graded \mathbb{K} -algebra with an additional internal grading s . The product is graded-commutative with respect to the cohomological grading \bullet [22]. The truncation $\mathrm{HH}^\bullet(A, A)^{\leq 0}$, in which s is required to be non-positive, is a subalgebra. Theorem B does *not* make any claims about the untruncated Hochschild algebra.

The rank of $\mathrm{HH}^r(A, A)^s$ was computed by the authors in [35] by a different method. Another approach has been found, in characteristic 0, by Fisette [19], who also makes the link between A_∞ -structures and elliptic curves. Although the additive result is the only part that is essential, we choose to present here this more complete calculation, proved via algebraic geometry on a cuspidal cubic curve C_{cusp} , by a method explained to us by Paul Seidel.

Lemma 4.1 *Let C_{cusp} be a cuspidal Weierstrass curve over the field \mathbb{K} . Then $\text{HH}^\bullet(A, A) \cong \text{HH}^\bullet(C_{\text{cusp}})$.*

This lemma is restated as part of Prop. 4.4 and proved there. One can compute $\text{HH}^\bullet(C_{\text{cusp}})$ using sheaf theory; the heart of the calculation is that of $\text{HH}^\bullet(R, R)$, where $R = \mathbb{K}[x, y]/(y^2 - x^3)$. The latter calculation can be done—as for any complete intersection singularity—by using a Koszul resolution to replace R by a smooth affine dg manifold, for which a version of the Hochschild–Kostant–Rosenberg theorem is available.

One virtue of the method is that it adapts easily to yield a computation of the Hochschild cohomology of A with a non-trivial A_∞ -structure arising from a Weierstrass curve. We shall carry out the computation for the case of a nodal curve, which we will be able to identify with the symplectic cohomology SH^* for the punctured torus.

4.1 Hochschild cohomology for varieties

Note: The ideas that we review in (4.1) have undergone a lengthy evolution, which we have not attempted to trace in detail. Our citations do not necessarily reflect priority.

4.1.1 The global HKR isomorphism

Let X be a quasi-projective scheme of dimension d over a field \mathbb{K} . Its Hochschild cohomology $\text{HH}^\bullet(X)$ is defined as

$$\text{HH}^\bullet(X) := \text{Ext}_{X \times X}^\bullet(\delta_* \mathcal{O}_X, \delta_* \mathcal{O}_X),$$

where $\delta: X \rightarrow X \times X$ is the diagonal map [56]. When X is smooth, there is a natural morphism of complexes of sheaves

$$(\mathbb{L}\delta^*)(\delta_* \mathcal{O}_X) \rightarrow \bigoplus_r \Omega_X^r[r]$$

(trivial differential in the complex on the right), the Hochschild–Kostant–Rosenberg (HKR) map; see [64] or the précis in [12]. When $d!$ is invertible in \mathbb{K} —which means that $\text{char}(\mathbb{K})$ is either 0 or $> d$ —the HKR map is a quasi-isomorphism.⁴ From the HKR map and the adjunction $\text{R Hom}_{X \times X}^\bullet(\delta_* \mathcal{O}_X, \delta_* \mathcal{O}_X) \cong \text{R Hom}_X^\bullet((\mathbb{L}\delta^*)(\delta_* \mathcal{O}_X), \mathcal{O}_X)$, one obtains, under this assumption on \mathbb{K} , an isomorphism of graded \mathbb{K} -vector spaces

$$\text{HKR}^n: \text{HH}^n(X) \rightarrow \bigoplus_{p+q=n} H^p(\Lambda^q T_X).$$

⁴For a general field \mathbb{K} , this map may not be a quasi-isomorphism but it is nevertheless true ([62, Ex. 9.1.3],) that $\text{HH}^\bullet(X, X) \cong \Lambda^\bullet T_X$ when $X = \mathbb{A}^n$.

The Hodge spectral sequence. For X quasi-projective over \mathbb{K} , there is a sheaf of graded algebras \mathcal{HH}^* on X , the sheafification of a natural presheaf whose sections are $\Gamma(U) = \mathrm{HH}^*(U)$. There is by [56] a local-to-global (or ‘Hodge’) spectral sequence E_*^{**} converging to $\mathrm{HH}^*(X)$, with

$$(28) \quad E_2^{pq} \cong H^p(X, \mathcal{HH}^q).$$

When X is smooth of dimension d , and $d!$ is invertible in \mathbb{K} , one has $\mathcal{HH}^q \cong \Lambda^q T_X$ by the HKR isomorphism. Comparing dimensions of $\mathrm{HH}^n(X)$ and E_2^n , one sees that the spectral sequence degenerates at E_2 . One therefore has

$$E_\infty^{pq} = E_2^{pq} \cong H^p(X, \Lambda^q T_X).$$

As a sheaf of graded algebras, \mathcal{HH} has cohomology spaces which form a bigraded algebra $H^*(X, \mathcal{HH}^*)$. Like any local-to-global spectral sequence computing self-Exts, the Hodge spectral sequence is multiplicative. From this one sees:

Lemma 4.2 *For X a quasi-projective \mathbb{K} -scheme, the edge-map $\mathrm{HH}^*(X) \rightarrow E_2^{0,*} = H^0(X; \mathcal{HH}^*)$ is an algebra homomorphism.*

4.1.2 Derived categories

Let $\mathrm{QC}(X)$ denote a dg enhancement for the unbounded derived category of quasi-coherent sheaves on the \mathbb{K} -scheme X (e.g. [36, 60]). Thus one has an equivalence of $H^0(\mathrm{QC}(X))$ with the unbounded derived category. There is a natural isomorphism of rings [60]

$$\mathrm{HH}^*(X) \rightarrow \mathrm{HH}^*(\mathrm{QC}(X)),$$

where by the latter we mean the Hochschild cohomology of the dg category QC (i.e., self-Ext of the identity functor). Let $\mathrm{Perf}(X)$ be the full dg subcategory of $\mathrm{QC}(X)$ of objects which map to perfect complexes in the unbounded derived category. Since $\mathrm{Perf}(X)$ is a full subcategory of $\mathrm{QC}(X)$, there is a restriction (ring) map

$$\mathrm{HH}^*(\mathrm{QC}(X)) \rightarrow \mathrm{HH}^*(\widetilde{\mathrm{Perf}(X)}),$$

and this too is an isomorphism when X is quasi-projective, because $\mathrm{QC}(X)$ is the *ind-completion* of $\mathrm{Perf}(X)$ [9]. When X is quasi-projective, we also have $\mathrm{Perf}(X) \simeq \mathrm{tw vect}(X)$. Putting these facts together, we obtain the following:

Lemma 4.3 *For any quasi-projective \mathbb{K} -scheme X , one has a canonical algebra-isomorphism*

$$\mathrm{HH}^*(\mathrm{tw vect}(X)) \cong \mathrm{HH}^*(X).$$

If T is a split-generator for $\text{tw vect}(X)$, the Morita invariance of Hochschild cohomology (e.g. [60]) implies that the restriction map

$$\text{HH}^\bullet(\text{tw vect}(X)) \rightarrow \text{HH}^\bullet(\text{End}(T))$$

is an isomorphism. Hence we have $\text{HH}^\bullet(\text{End}(T)) \cong \text{HH}^\bullet(X)$, and consequently

Proposition 4.4 *Take an abstract Weierstrass curve (C, σ, ω) and let $T = \mathcal{O} \oplus \mathcal{O}_\sigma$. Then there is a canonical algebra-isomorphism*

$$\text{HH}^\bullet(C) \rightarrow \text{HH}^\bullet(\text{End}(T)).$$

In particular, by Lemma 3.5,

$$\text{HH}^\bullet(C_{\text{cusp}}) \cong \text{HH}^\bullet(A, A).$$

Note that the latter isomorphism respects bigradings. The internal grading of $\text{HH}^\bullet(A, A)$ comes from the grading of A . The internal grading of $\text{HH}^\bullet(C_{\text{cusp}})$ arises because of the \mathbb{K}^\times -action on C_{cusp} discussed in the proof of Lemma 3.5, which gives rise to an action on $\text{tw vect}(C_{\text{cusp}})$, hence on Hochschild cohomology. They agree because the action on C_{cusp} induces the grading of $\text{End}(T)$ (cf. the proof of Lemma 3.5).

4.1.3 Curves

We now specialize to quasi-projective curves C over \mathbb{K} . HKR reads as follows:

Lemma 4.5 *For a non-singular curve C over an arbitrary field \mathbb{K} , one has an HKR isomorphism of \mathbb{K} -modules*

$$\text{HKR}^n: \text{HH}^n(C) \rightarrow H^n(\mathcal{O}) \oplus H^{n-1}(TC).$$

For more general curves, we have the following

Lemma 4.6 *For any quasi-projective curve C , one has short exact sequences*

$$0 \rightarrow H^1(C, \mathcal{H}\mathcal{H}^{n-1}) \rightarrow \text{HH}^n(C) \rightarrow H^0(C, \mathcal{H}\mathcal{H}^n) \rightarrow 0.$$

Proof Since curves have cohomological dimension 1, the Hodge spectral sequence (28) is supported at E_2 in two adjacent columns $p \in \{0, 1\}$, and therefore degenerates. The edge-maps then give rise to these short exact sequences. \square

One has $\mathcal{H}\mathcal{H}^0 = \mathcal{O}_C$, while $\mathcal{H}\mathcal{H}^1 \cong \text{Der}(\mathcal{O}_C)$, the sheaf of derivations of \mathcal{O}_C , also known as the tangent sheaf \mathcal{T}_C .

To go further, we suppose that C is given with a point $s \in C(\mathbb{K})$, and that $C \setminus \{s\}$ is non-singular. We take an affine open cover $\mathcal{U} = \{U, V\}$ of C such that $s \notin V$. We compute sheaf cohomology $H^*(C, \mathcal{H}\mathcal{H}^\bullet)$ using the Čech complex $\check{C}^* = \check{C}^*(\mathcal{U}, \mathcal{H}\mathcal{H}^q)$. To validate such a computation, we must check that $H^i(\mathcal{H}\mathcal{H}^q|_Y) = 0$ when $Y = U, V$ or $U \cap V$ and $i > 0$. As coherent sheaves, $\mathcal{H}\mathcal{H}^0$ and $\mathcal{H}\mathcal{H}^1$ have no higher cohomology on affine open subsets. When $q > 1$, $\mathcal{H}\mathcal{H}^q$ is supported at s , by (4.5). It therefore has no higher cohomology. We now proceed with the computation.

Since V is non-singular, (4.5) implies that $\text{HH}^0(V) \cong \Gamma(V, \mathcal{O}_V)$, $\text{HH}^1(V) \cong \Gamma(V, \mathcal{T}_V)$ and $\text{HH}^q(V) = 0$ for $q > 1$. The same applies over $U \cap V$. So the Čech complex is

$$0 \rightarrow \text{HH}^q(U) \oplus \text{HH}^q(V) \rightarrow \text{HH}^q(U \cap V) \rightarrow 0$$

supported in degrees 0 and 1. Hence:

Lemma 4.7 *For $q > 1$, one has*

$$H^0(C, \mathcal{H}\mathcal{H}^q) \cong \text{HH}^q(U), \quad H^1(C, \mathcal{H}\mathcal{H}^q) = 0.$$

Therefore, for $n > 2$,

$$\text{HH}^n(C) \xrightarrow{\cong} H^0(C, \mathcal{H}\mathcal{H}^n) \cong \text{HH}^n(U) = \text{HH}^n(\Gamma(\mathcal{O}_U), \Gamma(\mathcal{O}_U)).$$

In low degrees, we have an isomorphism $\text{HH}^0(C) \cong H^0(\mathcal{O}_C)$ and two short exact sequences

$$(29) \quad 0 \rightarrow H^1(C, \mathcal{O}_C) \rightarrow \text{HH}^1(C) \rightarrow H^0(C, \mathcal{T}_C) \rightarrow 0$$

$$(30) \quad 0 \rightarrow H^1(C, \mathcal{T}_C) \rightarrow \text{HH}^2(C) \rightarrow \text{HH}^2(U) \rightarrow 0.$$

In summary,

■ $\text{HH}^\bullet(C)$ coincides with $\text{HH}^\bullet(\Gamma(\mathcal{O}_U), \Gamma(\mathcal{O}_U))$ with corrections in degrees ≤ 2 from the cohomology of the functions and of the vector fields on C .

4.1.4 Plane curve singularities

If $R = \mathbb{K}[x, y]/(f)$ is a plane curve singularity, over a field \mathbb{K} , one can compute $\text{HH}^\bullet(R, R)$ via Koszul resolutions and a version of HKR. We have been informed that such calculations go back to Quillen [43], but an explicit recipe, valid for complete intersections, is explained by Kontsevich in [32]. The result is an algebra isomorphism

$$(31) \quad \text{HH}^\bullet(R, R) \cong H^\bullet(D, d_D),$$

where (D, d_D) is a certain dga, namely, the supercommutative \mathbb{K} -algebra

$$R \otimes \mathbb{K}[\beta, x^*, y^*], \quad \deg \beta = 2, \quad \deg x^* = \deg y^* = 1$$

and $d_D(R) = 0$, $d_D(\beta) = 0$, $d_D x^* = f_x \beta$, $d_D y^* = f_y \beta$. Thus

$$D_{2n} = R\beta^n \oplus R\beta^{n-1}x^*y^*, \quad D_{2n+1} = R\beta^n x^* \oplus R\beta^n y^*.$$

With the isomorphism (31) understood, one immediately reads off the even Hochschild cohomology of R . Let $T = R/(f_x, f_y)$. In the literature T is often called the *Tjurina algebra* of this isolated hypersurface singularity; it parameterizes a miniversal deformation of the singularity; see for instance [25, Theorem 14.1] for the case of curve singularities.

Lemma 4.8 *Assume that f_x and f_y are not both zero. Then the map*

$$T \rightarrow \mathrm{HH}^{2n}(R, R), \quad [\alpha] \mapsto [\alpha\beta^n]$$

is an isomorphism for each $n > 0$. Hence $\bigoplus_{n>0} \mathrm{HH}^{2n}(R, R) \cong T \otimes \beta\mathbb{K}[\beta]$.

Let M be the R -submodule of $R \oplus R$ of pairs (α_1, α_2) with $\alpha_1 f_x + \alpha_2 f_y = 0 \in R$. It has a submodule N generated by $(f_y, -f_x)$. There is a skew-symmetric pairing

$$\omega: M \otimes M \rightarrow T, \quad (\alpha_1, \alpha_2; \gamma_1, \gamma_2) \mapsto [\alpha_1 \gamma_2 - \alpha_2 \gamma_1]$$

such that $\omega(N \otimes M) = 0$. When f is irreducible and $df \neq 0$, one has $\omega = 0$. Indeed, R is then an integral domain, and the matrix

$$\begin{bmatrix} \alpha_1 & \alpha_2 \\ \gamma_1 & \gamma_2 \end{bmatrix},$$

over R , has a non-trivial kernel and therefore vanishing determinant.

Observe also that the R -module M/N is actually a T -module.

Lemma 4.9 *One has a surjective map*

$$M \rightarrow \mathrm{HH}^{2n+1}(R, R), \quad (\alpha_1, \alpha_2) \mapsto \beta^n [\alpha_1 x^* + \alpha_2 y^*]$$

The kernel is 0 when $n = 0$ and is N when $n > 0$.

The proof is an easy check, in light of (31). Note that $\mathrm{HH}^1(R, R) = \mathrm{Der} R$, the derivations of R ; the isomorphism $M = H^1(D, d_D) \rightarrow \mathrm{Der}(R)$ induced by (31) is the map $(\alpha_1, \alpha_2) \mapsto \alpha_1 \partial_x + \alpha_2 \partial_y$. We deduce:

Proposition 4.10 *There is a canonical map of algebras*

$$\mathrm{HH}^\bullet(R, R) \rightarrow Q^\bullet,$$

which is an isomorphism in degrees ≥ 2 . Here

$$Q^\bullet = \mathbb{K}[\beta] \otimes \left(T \oplus \frac{M}{N}[-1] \right), \quad \deg T = 0, \quad \deg \beta = 2.$$

The product in Q^\bullet combines the algebra structure of T , the left and right T -module structures of M/N , and the skew pairing ω :

$$(\beta^n \otimes (t + m)) \cdot (\beta^{n'} \otimes (t' + m')) = \beta^{n+n'} \otimes (t \cdot t' + t \cdot m' + m \cdot t') + \beta^{n+n'+1} \otimes \omega(m, m').$$

The following observation, whose proof is immediate, is helpful in computing M/N :

Lemma 4.11 *M/N is the middle homology $H_1(f_x, f_y; R)$ of the Koszul complex $K(f_x, f_y)$, i.e., the chain complex*

$$R \xrightarrow{\begin{bmatrix} f_y \\ -f_x \end{bmatrix}} R^2 \xrightarrow{\begin{bmatrix} f_x & f_y \end{bmatrix}} R.$$

$K(f_x, f_y)$ (notation from Serre [53, ch. IV]) is the tensor product $K(f_x) \otimes_R K(f_y)$ of Koszul complexes for the elements f_x and f_y . Here $K(a)$ denotes the complex $0 \rightarrow R \xrightarrow{a} R \rightarrow 0$ where the map—multiplication by a —maps degree 1 to degree 0.

4.2 Hochschild cohomology for the cuspidal cubic

4.2.1 Hochschild cohomology for the affine curve

We take $f(x, y) = y^2 - x^3$. The Tjurina algebra T (4.8) is as follows:

$$T \cong \begin{cases} \mathbb{K}[x]/(x^2) & \text{if } 6 \neq 0, \\ \mathbb{K}[x]/(x^3) & \text{if } 3 = 0, \\ \mathbb{K}[x, y]/(x^2, y^2) & \text{if } 2 = 0. \end{cases}$$

(The isomorphism takes $x \in T$ to x , and takes $y \in T$ to 0 in the first two cases and to y in the third.)

Lemma 4.12 *There are isomorphisms of R -modules*

$$M/N \cong \begin{cases} \mathrm{Ann}_{R/f_y}(f_x) = (x) & \text{if } 6 \neq 0, \\ \mathrm{Ann}_{R/f_y}(f_x) = R/f_y & \text{if } 3 = 0, \\ \mathrm{Ann}_{R/f_x}(f_y) = R/f_x & \text{if } 2 = 0 \end{cases}$$

and

$$T \rightarrow M/N, \quad t \mapsto \begin{cases} [2xt, 3yt] & \text{if } 6 \neq 0, \\ [t, 0] & \text{if } 3 = 0, \\ [0, t] & \text{if } 2 = 0. \end{cases}$$

Proof By Lemma 4.11, M/N is equal as an R -module to $H_1(f_x, f_y; R)$, and hence isomorphic to $H_1(K(f_x; R) \otimes K(f_y; R))$.

If $2 \neq 0 \in \mathbb{K}$ then the element $f_y = 2y$ is not a zero-divisor in R . Therefore the complex $K(f_y)$ has homology only in degree zero, and $H_0(K(f_y)) = R/(f_y)$. Projection $R \rightarrow R/(f_y)$ defines a quasi-isomorphism $K(f_y) \rightarrow R/(f_y)$. Hence $K(f_x, f_y; R) \simeq K(f_x; R) \otimes R/f_y$, and the latter complex is

$$0 \rightarrow R/(f_y) \xrightarrow{f_x} R/(f_y) \rightarrow 0,$$

where the differential maps degree 1 to degree 0. So $H_1(f_x, f_y; R)$ is isomorphic as an R -module (hence also as a T -module) to the annihilator of f_x in R/f_y . Explicitly, the annihilator is $(x) \subset \mathbb{K}[x, y]/(x^3, y)$ if $3 \neq 0$ and $\mathbb{K}[x, y]/(x^3, y)$ if $3 = 0$. If $2 = 0$ then the element $f_x = -3x^2$ is not a zero-divisor in R . We can then run the same argument with the roles of f_x and f_y interchanged, with a similar outcome. Here R/f_x is $\mathbb{K}[x, y]/(x^2, y^2)$, and $f_y = 0$.

At this stage, we can easily see that $\dim_{\mathbb{K}} T = \dim_{\mathbb{K}}(M/N)$. The given map $T \rightarrow M/N$ is injective, by another easy check, hence an isomorphism. \square

Since f is irreducible and $df \neq 0$, we have the

Lemma 4.13 *The skew pairing $\omega: (M/N) \times (M/N) \rightarrow T$ is zero.*

Collating results (Prop. 4.10 and Lemmas 4.12, 4.13) we obtain

Lemma 4.14 *Consider the free graded-commutative algebra $Q^\bullet = T \otimes \mathbb{K}[\beta, \gamma]$, where $\deg \beta = 2$ and $\deg \gamma = 1$. One then has a map of graded \mathbb{K} -algebras*

$$HH^\bullet(R, R) \rightarrow Q^\bullet$$

which is an isomorphism except in degrees 0 and 1.

Proof In light of Prop. 4.10, we need only see that $T \otimes \mathbb{K}[\beta, \gamma]$ agrees with the algebra named Q^\bullet there. In view of the isomorphism $T \cong M/N$, and the vanishing of ω , the latter algebra is

$$\mathbb{K}[\beta] \otimes (T \oplus T[-1]) \cong \mathbb{K}[\beta, \gamma] \otimes T, \quad \deg \gamma = 1.$$

\square

We can be more explicit. We know that $\mathrm{HH}^\bullet(R, R) \cong H^\bullet(D)$. The map $H^\bullet(D) \rightarrow Q^\bullet$ is R -linear. It maps β to β . It maps a certain class $[ax^* + by^*] \in H^1(D)$ to γ . The coefficients are given by $(a, b) = (2x, 3y)$ when $6 \neq 0$; $(a, b) = (1, 0)$ when $3 = 0$; and $(a, b) = (0, 1)$ when $2 = 0$.

4.2.2 Global calculation

Lemma 4.15 *The Lie algebra $H^0(C_{\mathrm{cusp}}, \mathcal{T}_{C_{\mathrm{cusp}}})$ contains linearly independent vector fields v_0, v_1 which restrict to U as the vector fields*

$$v_0|_U = 2x\partial_x + 3y\partial_y, \quad v_1|_U = 2y\partial_x + 3x^2\partial_y.$$

When $6 \neq 0$, v_0 and v_1 span. When $3 = 0$, $H^0(C_{\mathrm{cusp}}, \mathcal{T}_{C_{\mathrm{cusp}}})$ is 3-dimensional, spanned by v_0 and v_1 together with

$$v_{-2} = -\partial_x.$$

When $2 = 0$, it is 4-dimensional, spanned by v_0, v_1 and

$$v_{-1} = x\partial_y, \quad v_{-3} = \partial_y.$$

(We will see presently that the subscript s of v_s represents the internal degree.)

Proof As shown to us by Seidel, the calculation of H^0 is straightforward when one takes the ‘abstract’ view of C_{cusp} (Lemma 2.17), as \mathbb{P}^1 with the marked point c , and non-standard structure sheaf \mathcal{O}^c . Thus $\mathcal{T}_{C_{\mathrm{cusp}}} = \mathrm{Der}(\mathcal{O}^c)$ is the sheaf of meromorphic vector fields θ on \mathbb{P}^1 , with a pole only at c , which preserve \mathcal{O}^c . The latter condition forces θ to be regular at c except when $6 = 0$. When θ is regular at c , it must vanish there, except in characteristic 2. One has $H^0(T_{\mathbb{P}^1}) = \mathfrak{pgl}_2$ (the Lie algebra of the automorphism group of \mathbb{P}^1), spanned by $v_{-1} = \partial_z$, $v_0 = z\partial_z$ and $v_1 = z^2\partial_z$. In characteristic 2, all three lie in $H^0(\mathrm{Der}(\mathcal{O}^c))$, while when $2 \neq 0$, only v_0 and v_1 do (they span the Borel subalgebra $\mathfrak{b} \subset \mathfrak{pgl}_2$ of upper triangular 2×2 matrices modulo scalars).

We now account for the additional vector fields not regular at c : when $3 = 0$, one must include $v_{-2} = z^{-1}\partial_z$; when $2 = 0$, one must include $v_{-3} = z^{-2}\partial_z$ (in addition to $v_{-1} = \partial_z$).

To interpret these vector fields as derivations of R , we observe that the normalization map $\nu: \mathbb{P}^1 \rightarrow C_{\mathrm{cusp}}$ corresponds to the map of rings $R = \mathbb{K}[x, y]/(y^2 - x^3) \rightarrow \mathbb{K}[z]$, $x \mapsto z^2$, $y \mapsto z^3$. Using this we compute that the v_i restrict to U in the way stated above. \square

Lemma 4.16 *One has $H^1(C_{\text{cusp}}, \mathcal{T}_{C_{\text{cusp}}}) = 0$.*

Proof Consider \mathbb{P}^1 with the cusp modifying its structure sheaf at $c = [0 : 1]$. Let $b = [1 : 0]$. We use the affine coordinate $[z : 1] \mapsto z$. Take a derivation θ of $\mathcal{O}_{\mathbb{P}^1}$ (i.e., a vector field) over $\mathbb{P}^1 \setminus \{b, c\}$. We must show that it is the difference $u - v$ of vector fields u on $U = \mathbb{P}^1 \setminus \{b\}$ and v on $V = \mathbb{P}^1 \setminus \{c\}$. We can extend θ to a meromorphic vector field on \mathbb{P}^1 . We proceed by induction on the order d of the pole of θ at b . If θ is regular at b (i.e., $d \leq 0$) then we take $u = 0$ and $v = -\theta$. For the inductive step, say $\theta \sim az^d$ at b , where $a \neq 0$ and $d > 0$. Let $u = az^d$. Then u defines a derivation of \mathcal{O}^c near c , because $d > 0$. Moreover, $\theta' := u - \theta$ has a pole of order $< d$ at b , so by induction we can write $\theta' = u' - v'$ for u' on U and v' on V ; then $\theta = (u - u') + v'$ and we are done. \square

At this stage it becomes useful to bring internal gradings into play. The Hochschild cohomology $\text{HH}^\bullet(C_{\text{cusp}})$ carries an internal grading s , arising from the \mathbb{K}^\times -action on C_{cusp} . Equivalently, this is the internal grading of $\text{HH}^\bullet(A, A)$ arising from the grading of A . There is also an action on \mathcal{HH}^\bullet , hence an internal grading on $H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)$.

Under the \mathbb{K}^\times -action on C_{cusp} , the functions X, Y and Z on C_{cusp} have respective weights $w(X) = 2$, $w(Y) = 3$ and $w(Z) = 0$. Hence on the affine part, the functions $x = X/Z$ and $y = Y/Z$ have weights $w(x) = 2$ and $w(y) = 3$. The dga (D^*, d_D) inherits a \mathbb{K}^\times -action, i.e., a grading s , in which the weights of x and y are $s(x) = 2$ and $s(y) = 3$. The variables x^* and y^* have weights $s(x^*) = -2$ and $s(y^*) = -3$. It then follows that $s(\beta) = -6$.

In Lemma 4.15, one should understand ∂_x to have weight $s = -2$ and ∂_y weight $s = -3$. Then the vector field v_k has weight k .

Introduce the internal grading on Q^\bullet as in the statement of Theorem B. It is set up so that the homomorphism $\text{HH}^\bullet(R, R) \rightarrow Q^\bullet$ respects it. Let $Q^{\bullet, \leq 0}$ be the subalgebra of Q^\bullet where the internal grading is non-positive. Similarly, define $H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)^{\leq 0}$.

Lemma 4.17 *The map $\text{HH}^\bullet(C_{\text{cusp}})^{\leq 0} \rightarrow H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)^{\leq 0}$ is an isomorphism.*

Proof The kernel of the surjective map $\text{HH}^\bullet(C_{\text{cusp}}) \rightarrow H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)$ is spanned by $H^1(\mathcal{T}_{C_{\text{cusp}}})$ in degree 2—this vanishes by Lemma 4.16—and $H^1(\mathcal{O})$ in degree 1. The action of \mathbb{K}^\times has weight 1 on $H^1(\mathcal{O})$. Hence the restricted map $\text{HH}^\bullet(C_{\text{cusp}})^{\leq 0} \rightarrow H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)^{\leq 0}$ is injective, and so an isomorphism. \square

Lemma 4.18 *One has an isomorphism of graded algebras $H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)^{\leq 0} \rightarrow Q^{\bullet, \leq 0}$.*

Proof There is a restriction map $r: H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet) \rightarrow HH^\bullet(U) = HH^\bullet(R, R)$. By Lemmas 4.7 and 4.16, r is an isomorphism in degrees ≥ 2 . In degree 1, it is the restriction map $H^0(C_{\text{cusp}}, \mathcal{T}_{C_{\text{cusp}}}) \rightarrow H^0(U, \mathcal{T}_U)$, and therefore again injective. The composite

$$H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet) \xrightarrow{r} HH^\bullet(U) = HH^\bullet(R, R) \rightarrow Q^\bullet$$

of r with the map from Lemma 4.14, restricted so as to map between the non-positively graded subalgebras, is the sought-for map. It is certainly an isomorphism in degrees $\bullet > 1$. It is also an isomorphism in degree 0, where both sides reduce to \mathbb{K} .

The non-positively graded part of $H^0(C_{\text{cusp}}, \text{Der}(\mathcal{O}_{C_{\text{cusp}}})^{\leq 0})$ has dimension $\dim T - 1$, by Lemma 4.15: it is spanned by v_0 , together with v_{-2} in characteristic 3, and v_{-1} , v_{-3} in characteristic 2. One also has $Q^1 = T$, and $\dim(Q^1)^{\leq 0} = \dim T - 1$.

We have identified v_j as derivations of R , hence as elements of $H^1(D, d_D) = M$. One has $Q^1 \cong M/N$, in such a way that the map $H^1(D, d_D) \rightarrow Q^1$ corresponds to the quotient map $H^1(D, d_D) = M \rightarrow M/N$. To show that $H^0(C_{\text{cusp}}, \text{Der}(\mathcal{O}_{C_{\text{cusp}}})^{\leq 0}) \rightarrow (Q^1)^{\leq 0}$ is injective, one need only show that the relevant elements v_j are linearly independent in M/N . But as elements of M , one has $v_0 = [2x, 3y]$; $v_{-2} = [-1, 0]$; $v_{-1} = [0, x]$; and $v_{-3} = [0, 1]$. Linear independence in M/N is easily seen. \square

Proof of Theorem B The isomorphism we want is the composite of the isomorphisms described in the last two lemmas:

$$HH^\bullet(C_{\text{cusp}})^{\leq 0} \rightarrow H^0(C_{\text{cusp}}, \mathcal{HH}^\bullet)^{\leq 0} \rightarrow Q^{\bullet, \leq 0}.$$

\square

Before leaving C_{cusp} , it will be worthwhile to spell out our findings about $HH^1(C_{\text{cusp}})^{\leq 0}$ in a clean form. They are as follows:

Proposition 4.19 *Over any field \mathbb{K} ,*

- (1) *The canonical map*

$$HH^1(C_{\text{cusp}})^{\leq 0} \rightarrow H^0(C_{\text{cusp}}, \mathcal{T}_{C_{\text{cusp}}})^{\leq 0}$$

is an isomorphism.

- (2) *Let $\ker d \subset W$ be as at (24); it is the Lie subalgebra of $\mathfrak{pgl}_3(\mathbb{K})$ of those vector fields on $\mathbb{P}^2(\mathbb{K})$ which preserve C_{cusp} . Thus there is a canonical map*

$$\ker d \rightarrow H^0(C_{\text{cusp}}, \mathcal{T}_{C_{\text{cusp}}})^{\leq 0},$$

natural in \mathbb{K} . The latter map is an isomorphism.

Proof Only the second clause has not already been proved. As a vector field on \mathbb{P}^2 , we have $\partial_u = (\partial x/\partial u)\partial_x + (\partial y/\partial u)\partial_y = -2x\partial_x - 3y\partial_y$. Similarly, $\partial_s = -x\partial_y$; $\partial_r = -\partial_x$; and $\partial_t = -\partial_y$. Now, ∂_u is tangent to C_{cusp} , and the restriction of ∂_u to C_{cusp} is the vector field v_0 . When $2 = 0$, ∂_s and ∂_t are also tangent to C_{cusp} , and they restrict to C_{cusp} as the respective vector fields v_{-1} and v_{-3} . When $3 = 0$, ∂_r is tangent to C_{cusp} and restricts to C_{cusp} as v_{-2} . Thus, by Lemma 4.15, the map $\ker d \rightarrow H^0(C_{\text{cusp}}, \mathcal{T}_{C_{\text{cusp}}})^{\leq 0}$ is an isomorphism in every case. \square

4.3 Hochschild cohomology of the nodal cubic curve

Note: this subsection is not used elsewhere in the paper.

Let $\mathcal{T}_0 \rightarrow \text{Spec } \mathbb{Z}$ be the central fiber of the Tate curve, defined by $w(x, y) = 0$ where $w(x, y) = y^2 + xy - x^3$: the proposed mirror to the punctured 2-torus T_0 . Consider the \mathbb{K} -points $\mathcal{T}_0(\mathbb{K})$ for an arbitrary field \mathbb{K} . To simplify notation, we shall write \mathcal{T}_0 for $\mathcal{T}_0(\mathbb{K})$. The Hochschild cohomology $\text{HH}^\bullet(\mathcal{T}_0)$ is also the Hochschild cohomology of $\text{tw vect}(\mathcal{T}_0)$. Since Hochschild cohomology is invariant under A_∞ quasi-equivalences, Theorem A says that $\text{HH}^\bullet(\mathcal{T}_0) \cong \text{HH}^\bullet(\mathcal{F}(T_0)^{\text{ex}})$. So, taking that theorem for granted for now, one can regard this subsection as a computation of $\text{HH}^\bullet(\mathcal{F}(T_0)^{\text{ex}})$. In general, there is a natural map to the Hochschild cohomology of the exact Fukaya category from the *symplectic cohomology* [48] of the manifolds, and in certain cases [46] this is expected to be an isomorphism. The graded ring we compute here is indeed isomorphic to $SH^\bullet(T_0)$, though we do not check that the map is an isomorphism.

Theorem 4.20 *Over any field \mathbb{K} , there is an isomorphism of graded algebras*

$$\text{HH}^\bullet(\mathcal{T}_0) \rightarrow U^\bullet/I.$$

Here $U^\bullet := \mathbb{K}[\beta, \gamma_1, \gamma_2, \epsilon]$ is a free supercommutative algebra on variables of the following degrees:

$$\deg \gamma_1 = \deg \gamma_2 = 1, \quad \deg \beta = 2, \quad \deg \epsilon = 3,$$

and

$$I = \gamma_1 \cdot U^{\geq 1} + \gamma_2 \cdot U^{\geq 1}$$

The computation follows similar lines to that for the cuspidal cubic, and we shall be terse. Taking $R = \mathbb{K}[x, y]/(w)$, one readily checks the following:

Lemma 4.21 *The Tjurina algebra T is reduced to \mathbb{K} . Its module M/N is 1-dimensional, and therefore the skew pairing $\omega: M/N \otimes M/N \rightarrow \mathbb{K}$ is zero. Hence there is an isomorphism of non-unital algebras*

$$\text{HH}^{\bullet \geq 2}(R, R) \rightarrow \mathbb{K}[\beta, \epsilon]^{\bullet \geq 2}, \quad \deg \beta = 2, \quad \deg \epsilon = 3,$$

with kernel concentrated in degree 1. (The notation $\bullet \geq 2$ means the part spanned by homogeneous elements of degree at least 2.)

The normalization of \mathcal{T}_0 is isomorphic to \mathbb{P}^1 ; the normalization map $\nu: \mathcal{T}_0 \rightarrow \mathbb{P}^1$ carries two points, p and q , say, to the node. The pullback $\nu^*\mathcal{O}_{\mathcal{T}_0}$ is contained in $\mathcal{O}_{\mathbb{P}^1}$ as the sheaf of functions f such that $f(p) = f(q)$. Similarly, $\nu^*\mathrm{Der}(\mathcal{O}_{\mathcal{T}_0})$ is contained in $\mathrm{Der}(\mathcal{O}_{\mathbb{P}^1})$ as the sheaf of vector fields v with $v(p) = v(q) = 0$. Hence

$$h^0(\mathcal{T}_0, \mathrm{Der}(\mathcal{O}_{\mathcal{T}_0})) = h^0(\mathbb{P}^1, \mathrm{Der}(\nu^*\mathcal{O}_{\mathcal{T}_0})) = 1.$$

By Riemann–Roch for \mathbb{P}^1 ,

$$h^1(\mathbb{P}^1, \mathrm{Der}(\nu^*\mathcal{O}_{\mathcal{T}_0})) = 0.$$

In Čech terms, this means that if we cover \mathbb{P}^1 by $U = \mathbb{P}^1 \setminus \{p, q\}$ and $V = \mathbb{P}^1 \setminus \{r\}$, then every vector field on $U \cap V$ can be expressed as the difference $u - v$ of vector fields on U and V which descend to $\nu(U)$ and $\nu(V)$ respectively. Hence $H^1(\mathcal{T}_0, \mathrm{Der}(\mathcal{O}_{\mathcal{T}_0})) = 0$.

We have

$$H^0(\mathcal{T}_0, \mathcal{HH}^1) = H^0(\mathcal{T}_0, \mathrm{Der}(\mathcal{O})) \cong \mathbb{K}.$$

Lemma 4.22 *The restriction map $H^0(\mathcal{T}_0, \mathcal{HH}^1) \rightarrow HH^1(R, R)$ is zero.*

Proof The vector field $v = w_y \partial_x - w_x \partial_y$ lies in $N \subset M$, and so is zero in $HH^1(R, R)$, yet it extends to a global vector field; in fact, in the coordinates $s = x/y$ and $t = 1/y$ at infinity, one has $v = (1 + s)\partial_s + (t - 3s^2)\partial_t$. Since $H^0(\mathcal{T}_0, \mathcal{HH}^1)$ is 1-dimensional, it is spanned by v . \square

Proof of Theorem 4.20 One has isomorphisms of non-unital algebras

$$HH^{\bullet \geq 2}(\mathcal{T}_0) \rightarrow H^0(\mathcal{T}_0, \mathcal{HH}^{\bullet \geq 2}) \rightarrow HH^{\bullet \geq 2}(R, R) \rightarrow \mathbb{K}[\beta, \epsilon]^{\bullet \geq 2}.$$

That the first map is an isomorphism even in degree 2 follows from the vanishing of $H^1(\mathcal{T}_0, \mathrm{Der}(\mathcal{O}_{\mathcal{T}_0}))$. It remains, then, to define the desired isomorphism in degree 1.

Now, $HH^1(\mathcal{T}_0)$ is 2-dimensional, spanned by the class η_1 of a generator for $H^1(\mathcal{O}) \subset HH^1(\mathcal{T}_0)$, and by a lift η_2 to $HH^1(\mathcal{T}_0)$ of a global vector field v on \mathcal{T}_0 . The E_∞ -page of the Hodge spectral sequence is an algebra isomorphic to $\mathrm{gr} HH^\bullet(C)$, the associated graded algebra for the filtration giving rise to the spectral sequence. The E_∞ -page here is supported along the 0th row and the 0th column (in non-negative total degrees). Hence $\mathrm{gr}^i HH^\bullet(\mathcal{T}_0) = HH^i(\mathcal{T}_0)$ for $i \neq 1$. To show that $\beta\eta_1 = 0$ it therefore suffices to show that it vanishes in $H^1(\mathcal{T}_0, \mathcal{HH}^2)$, which is obvious since this module vanishes. Likewise, $\epsilon\eta_1 = 0$. It follows from the last lemma that $[v] = 0 \in HH^1(R, R)$. Now, $\beta\eta_2 \in HH^3(\mathcal{T}_0)$ maps to $\beta \cdot [v] \in H^0(\mathcal{T}_0, \mathcal{HH}^3)$, and $\beta \cdot v = 0$ because this is so in U . Similarly $\epsilon\eta_2 = 0$ and $\eta_1\eta_2 = 0$. Thus we can obtain the desired isomorphism by mapping η_1 to γ_1 and η_2 to γ_2 . \square

5 Weierstrass curves versus A_∞ -structures

In this section we shall prove our ‘dg comparison theorem’ 3.3, and refinements of it.

5.1 Cochain models and splittings

Our plan is to reformulate Theorem 3.3 in terms of minimal A_∞ -structures on the fixed algebra A , and prove it in sharper form in that language. To do so, we need the homological perturbation theory reviewed in subsection 2.1.4.

Definition 5.1 Let (C^\bullet, δ) be a cochain complex over a commutative ring R , with an action of a group Γ by automorphisms. A *splitting* for C^\bullet is an internal direct sum decomposition

$$C^k = \underbrace{\mathcal{H}^k \oplus \operatorname{im} \delta^{k-1}}_{\ker \delta^k} \oplus \mathcal{J}^k$$

for each k , with Γ -invariant summands. Equivalently, it is a Γ -equivariant linear map $s \in \operatorname{hom}^{-1}(C, C)$ such that $\delta s \delta = \delta$ (for given the direct sum decomposition we put $s|_{\mathcal{H} \oplus \mathcal{J}} = 0$ and $s|_{\operatorname{im} \delta} = \delta^{-1}: \operatorname{im} \delta \rightarrow \mathcal{J}$, while given such an s we have $C^\bullet = \underbrace{[\ker \delta \cap \ker(\delta \circ s)]}_{\ker \delta} \oplus \operatorname{im} \delta \oplus \operatorname{im}(s \circ \delta)$).

The set of splittings will be denoted by $\operatorname{spl}_\Gamma(\mathcal{A})$. We noted in subsection 2.1.4 that when \mathcal{A} is a Γ -equivariant A_∞ -algebra, with cohomology A , a splitting for \mathcal{A} as a cochain complex gives rise, via the homological perturbation lemma to a canonical, Γ -equivariant A_∞ -structure $\mu_{\mathcal{A}}^\bullet$ on A , together with equivariant, homotopy-inverse A_∞ -homomorphisms $i: A \rightarrow \mathcal{A}$ and $p: \mathcal{A} \rightarrow A$ whose leading terms i^1 and p^1 induce the identity map id_A on cohomology.

If A has projective Hom-spaces, and if one chooses a *cycle choosing map*, that is, a map $s: A \rightarrow \ker \delta$ such that $[s(c)] = c$, but does not complement the cocycles, then by Lemma 2.3, one gets an A_∞ -structure $\mu_{\mathcal{A}}^\bullet$ on A and an A_∞ -homomorphisms $i: A \rightarrow \mathcal{A}$ whose leading term is $i^1 = s$ (whence i^1 induces the identity on cohomology). If one wants $\mu_{\mathcal{A}}^\bullet$ and i to be Γ -equivariant, one should assume that the Hom-spaces are projective over $\mathbb{K}[\Gamma]$, with the cycle-choosing map Γ -equivariant.

5.1.1 A Čech model for Weierstrass curves

Let $\mathcal{O}_W = \mathbb{K}[a_1, \dots, a_6]$, and let \mathcal{C} be the universal embedded Weierstrass curve $y^2z + a_1xyz + a_3yz^2 = x^3 + a_2x^2z + a_4xz^2 + a_6z^3$ over \mathcal{O}_W . We are interested in

cochain models $\mathcal{B}_{\mathcal{C}}$ for the endomorphism algebra $\text{Ext}^\bullet(T, T)$, where $T = \mathcal{O} \oplus \mathcal{O}_\sigma$ over the curve \mathcal{C} . A technical irritation is that \mathcal{O}_σ , though a perfect complex, is not locally free. However, there is an autoequivalence τ of $\text{tw vect}(\mathcal{C})$ —a twist along the spherical object \mathcal{O} [51]—such that $\tau(\mathcal{O}) = \mathcal{O}$ and $\tau(\mathcal{O}_\sigma) = \mathcal{O}(-\sigma)$. There is an autoequivalence τ' which is inverse to τ in that $\tau \circ \tau'$ and $\tau' \circ \tau$ are both naturally isomorphic to the identity functor. There is an induced isomorphism $\text{Ext}^\bullet(T, T) \ni [h] \mapsto [\tau' \circ h \circ \tau] \in \text{Ext}^\bullet(T', T')$, where $T' = \mathcal{O} \oplus \mathcal{O}(-\sigma)$. We shall set up a cochain complex $\mathcal{B}'_{\mathcal{C}}$ which computes $\text{Ext}^\bullet(T', T')$, and use τ and τ' to obtain from it a complex $\mathcal{B}_{\mathcal{C}}$ which computes $\text{Ext}^\bullet(T, T)$. We observe that a splitting for the former complex will transfer to one for the latter.

Our dg model for $\mathcal{B}'_{\mathcal{C}}$ will be the Čech complex

$$(32) \quad \mathcal{B}'_{\mathcal{C}} := (\check{C}^*(\mathcal{U}; \underline{\text{End}}(T')), \delta), \quad \mathcal{U} = \{U, V\},$$

where $U := \{z \neq 0\}$ is the complement of the section at infinity $\sigma = [0 : 1 : 0]$, and $V = \{y \neq 0\}$. The multiplicative group \mathbb{K}^\times acts on \mathcal{C} , covering its action on W . The action preserves each of the two sets in the covering \mathcal{U} ; hence \mathbb{K}^\times acts on $\mathcal{B}'_{\mathcal{C}}$ by automorphisms—a *strict* action [50, (10b)].

Lemma 5.2 *The set $\text{spl}_{\mathbb{K}^\times}(\mathcal{B}'_{\mathcal{C}})$ of \mathbb{K}^\times -equivariant retractions is non-empty; there exists a distinguished splitting r .*

Proof The sheaf $\underline{\text{End}}(T')$ has four line-bundle summands (‘matrix entries’): \mathcal{O} (twice), $\mathcal{O}(\sigma)$, and $\mathcal{O}(-\sigma)$. The complex $\mathcal{B}'_{\mathcal{C}}$ is the direct sum of the Čech complexes for these line bundles, so it suffices to handle them separately. Evidently, we must put $\mathcal{H}^1 = 0$ and $\mathcal{H}^0 = \ker \delta^0$. Our tasks are to identify \mathcal{H}^1 and \mathcal{H}^0 for the line bundles \mathcal{O} , $\mathcal{O}(\sigma)$ and $\mathcal{O}(-\sigma)$.

1. We consider endomorphisms of \mathcal{O} or of $\mathcal{O}(-\sigma)$. We have $\underline{\text{End}}(\mathcal{O}) = \underline{\text{End}}(\mathcal{O}(-\sigma)) = \mathcal{O}$. Let $\mathbb{L} = \mathbb{K}[W]$, and let $F = Y^2Z + \cdots \in \mathbb{L}[X, Y, Z]$ be the universal Weierstrass polynomial. Then $\mathcal{O}_U = \mathbb{L}[X, Y, Z^{\pm 1}]/F$, $\mathcal{O}_V = \mathbb{L}[X, Y^{\pm 1}, Z]/F$, and $\mathcal{O}_{U \cap V} = \mathbb{L}[X, Y^{\pm 1}, Z^{\pm 1}]/F$. We have $H^1(\mathcal{O}) = \text{coker } \delta^0$, where δ^0 is the map $\mathcal{O}_U \oplus \mathcal{O}_V \rightarrow \mathcal{O}_{U \cap V}$ given by $\delta^0(u, v) = u - v$. We claim that $\text{coker } \delta^0$ is a free \mathbb{L} -module of rank 1 generated by $g := X^2/(YZ)$.

The function $s = X/Z$ is a local uniformizer at σ . Indeed, if R denotes the localization of $\mathbb{L}[X, Y, Z]/F$ away from $Z = 0$, and \hat{R} the completion along $Z = 0$, then one has $\hat{R} = \mathbb{L}[[s]]$. There exists a homogeneous element $f \in \mathbb{L}[X, Y^{\pm 1}, Z]/F$ whose meromorphic extension across σ has leading term $s^{-2} \in \mathbb{L}((s))$. Using this, one sees that $\text{coker } \delta$ is spanned by any homogeneous element of $\mathbb{L}[X, Y^{\pm 1}, Z]/F$ which has a simple pole of residue 1 at σ . The element $X^2/(YZ)$ has this property.

Turning to the degree 0 part, $\ker \delta^0$ consists of pairs (c, c) with c constant. A complement \mathcal{I}^0 is given by

$$\mathcal{I}^0 = \mathcal{O}_U \oplus I \subset \mathcal{O}_U \oplus \mathcal{O}_V, \quad I = \{v \in \mathcal{O}_V : v(\sigma) = 0\}.$$

2. We consider $\mathcal{H}om(\mathcal{O}, \mathcal{O}(\sigma)) = \mathcal{O}(\sigma)$. One has $\check{\mathcal{C}}^1(\mathcal{U}, \mathcal{O}(\sigma)) = \text{im } \delta^0$, so $\mathcal{H}^1 = 0$; and

$$\ker \delta^0 = \{(c, c) \in \mathcal{O}(\sigma)_U \oplus \mathcal{O}(\sigma)_V : c \text{ constant}\}.$$

Thus

$$\check{\mathcal{C}}^0(\mathcal{U}, \mathcal{O}(\sigma)) = \ker \delta \oplus \mathcal{I}^0, \quad \mathcal{I}^0 = \mathcal{O}(\sigma)_U \oplus I,$$

where I consists of sections over V which vanish at σ .

3. We consider $\mathcal{H}om(\mathcal{O}(\sigma), \mathcal{O}) = \mathcal{O}(-\sigma)$. There are no global sections, so we put $\mathcal{I}^0 = \check{\mathcal{C}}^0(\mathcal{U}, \mathcal{O}(-\sigma))$. We have

$$\check{\mathcal{C}}^1(\mathcal{U}, \mathcal{O}(-\sigma)) = \text{im } \delta^0 \oplus \mathcal{H}^1,$$

where $\mathcal{H}^1 = \mathbb{L} \cdot g$, where g is as in case 1, for similar reasons to those explained in case 1. \square

We now come to a key point in the construction:

Take r from Lemma 5.2. By applying the homological perturbation lemma to $(\mathcal{B}'_{\mathcal{C}}, r)$, we obtain a minimal A_∞ structure $\mathcal{A}_W = (A \otimes W, \mu_W^\bullet)$, linear over $\mathbb{K}[W]$.

Said another way, we obtain a family of minimal A_∞ -structures $\mathcal{A}_w = (A, \mu_w^\bullet)$ parametrized by $w \in W$, whose structure coefficients depend polynomially on w .

The \mathbb{K}^\times -equivariance of r implies equivariance of μ_W^\bullet ; precisely, for $w \in W_k$ (the k th graded part), we have

$$(33) \quad \mu_{\epsilon^k w}^d = \epsilon^{d-2} \mu_w^d.$$

The dg comparison theorem was stated in terms of abstract Weierstrass curves, but the A_∞ -version will be formulated using embedded Weierstrass curves. First we set up the relevant categories:

- (1) The groupoid \mathcal{W} whose objects are embedded Weierstrass curves, thought of as elements $w \in W$. The group of reparametrizations G , from (2.3.3), acts on $\text{ob } \mathcal{W}$; we set

$$\text{mor}_{\mathcal{W}}(w_1, w_2) = \{g \in G : g(w_1) = w_2\}.$$

- (2) The groupoid \mathcal{M} of minimal A_∞ -structures \mathcal{A} on A . Let $C^{r+s}(A, A)^s$ denote the part of the r th Hochschild cochain group in cohomological degree $r + s$ and internal degree s ,

$$C^{r+s}(A, A)^s = \text{Hom}_{\mathbb{K}[W]}^s(A^{\otimes r}, A).$$

and write $C^k(A, A)^{\leq 0} = \prod_{s \leq 0} C^k(A, A)^s$: these are the truncated Hochschild cochains, and they govern deformations of A_∞ -structures [47]. Let \mathcal{G} denote the group of ‘rescalings and gauge transformations’, namely, the group of $u \in C^1(A, A)^{\leq 0}$ whose leading term $u^1 \in \text{Hom}(A, A)$ is $r \cdot \text{id}$ for some $r \in \mathbb{K}^\times$. Then \mathcal{G} acts on the objects of \mathcal{M} , $(u, \mathcal{A}) \mapsto u_* \mathcal{A}$. We think of u as an A_∞ -functor $\mathcal{A} \rightarrow u_* \mathcal{A}$ acting trivially on objects. The morphisms in \mathcal{M} are given by $\text{mor}_{\mathcal{M}}(\mathcal{A}, \mathcal{A}') = \{u \in \mathcal{G} : u_* \mathcal{A} = \mathcal{A}'\} / \sim$, where \sim is the equivalence relation which identifies homotopic functors. ‘Homotopy’ has the following meaning: say u_1 and u_2 are A_∞ -functors $\mathcal{A} \rightarrow \mathcal{A}'$ with the same action on objects. Their difference $D = u_1 - u_2$ then defines a natural transformation, i.e., a morphism $D = (D^0, D^1, D^2, \dots) \in \text{hom}_{\mathcal{Q}}^1(u_1, u_2)$, where \mathcal{Q} is the A_∞ -category of non-unital functors $\text{funct}(\mathcal{A}, \mathcal{A}')$, satisfying $\mu_{\mathcal{Q}}^1 D = 0$. One puts $D^0 = 0$ and $D^k = u_1^k - u_2^k$ for $k > 0$. A *homotopy* from u_1 to u_2 is a $T \in \text{hom}_{\mathcal{Q}}^0(u_1, u_2)$, with $T^0 = 0$, such that $D = \delta_{\mathcal{Q}} T$. Functors are called *homotopic* if a homotopy exists.

Theorem C (A_∞ comparison theorem) *Let \mathbb{K} be a normal commutative ring. The passage from Weierstrass curves over \mathbb{K} to \mathbb{K} -linear minimal A_∞ -structures on A defines a functor $\mathcal{F}: \mathcal{W} \rightarrow \mathcal{M}$. Precisely:*

- (1) *Each embedded Weierstrass curve $C = C_w$ gives rise to a minimal A_∞ -structure $\mathcal{F}(C) = \mathcal{A}_w = (A, \mu_w^\bullet)$, by applying the homological perturbation lemma to \mathcal{B}_C , defined via the open cover \mathcal{U} and the splitting r . These A_∞ -structures depend polynomially on $w \in W$, and are \mathbb{K}^\times -equivariant, meaning that $\mu_{t \cdot w}^d = t^{d-2} \mu_w^d$. In particular, the cuspidal cubic C_0 gives rise to the minimal A_∞ -structure \mathcal{A}_0 with no higher products: $\mu_0^d = 0$ for $d \neq 2$.*
- (2) *Each pair $(g, w) \in G \times W$, where G is the group of projective transformations acting on Weierstrass curves, gives rise to an element $\mathcal{F}(g, w) = u^g(w) \in \mathcal{G}$ such that $u^g(w)_* \mathcal{A}_w = \mathcal{A}_{g(w)}$. Moreover, $u^g(t \cdot w) = t^{d-1} u^g(w)$ for $t \in \mathbb{K}^\times$.*
- (3) *One has $u^1(w) = \text{id}$. For any $(g_1, g_2) \in G \times G$, there exists a homotopy*

$$u^{g_2 g_1}(w) \simeq u^{g_2}(g_1 w) \circ u^{g_1}(w)$$

depending algebraically on w .

The functor \mathcal{F} commutes with base-change (i.e., the operation $\cdot \otimes_{\mathbb{K}} \mathbb{K}'$ when $\mathbb{K} \rightarrow \mathbb{K}'$ is a ring homomorphism). If \mathbb{K} is either an integrally closed noetherian domain of characteristic zero, or a field then \mathcal{F} is an equivalence of categories.

The equivalence clause says that the following three properties hold:

- *Essential surjectivity*: every minimal A_∞ -structure on A is isomorphic in \mathcal{M} to one of the form \mathcal{A}_w .
- *Faithfulness*: If $g_1(w) = g_2(w)$, and if $u^{g_1}(w), u^{g_2}(w) \in \text{ob funct}(\mathcal{A}_w, \mathcal{A}_{g_1(w)})$ are homotopic functors, then $g_1 = g_2$.
- *Fullness*: if $u \in \mathcal{G}$ and $u_*\mathcal{A}_{w_1} = \mathcal{A}_{w_2}$ then $u \in \text{ob funct}(\mathcal{A}_{w_1}, \mathcal{A}_{w_2})$ is homotopic to $u^g(w_1)$ for some $g \in G$ such that $g(w_1) = w_2$.

The normality condition is there because that is the condition under which we know that the Weierstrass differential ω on the universal Weierstrass curve C defines a section of the dualizing sheaf.

Proof of Theorem 3.3 assuming Theorem C. Over a field \mathbb{K} , every dg structure \mathcal{B} can be transferred to a quasi-isomorphic A_∞ -structure \mathcal{A} on the cohomology $A \cong H^*(\mathcal{B}_C)$. By Theorem C, \mathcal{A} is gauge-equivalent to \mathcal{A}_w for some w , which in turn is quasi-isomorphic to the dg category \mathcal{B}_w . Hence $\mathcal{B} \simeq \mathcal{B}_w$ in the A_∞ -sense. Furthermore, if $\mathcal{B}_{C_1} \simeq \mathcal{B}_{C_2}$, realize C_i as an embedded Weierstrass curve C_{w_i} ; so $\mathcal{B}_{w_1} \simeq \mathcal{B}_{w_2}$, and hence $\mathcal{A}_{w_1} \simeq \mathcal{A}_{w_2}$. Pick a gauge-equivalence $u: \mathcal{A}_{w_1} \rightarrow \mathcal{A}_{w_2}$, and then use the theorem to replace u by a homotopic gauge-transformation $u(g)$, where $g(w_1) = w_2$. Hence $C_1 \cong C_2$. \square

Coherence. One can ask whether the homotopies $H(g_1, g_2)$ from $u^{g_2 g_1}$ to $u^{g_2}(g_1 \cdot) \circ u^{g_1}$ can be chosen coherently in the sense of [50, (10b)]. Let $\tilde{\mathcal{M}}$ be the category of minimal A_∞ -structures on A , over \mathbb{K} , in which morphisms $\mathcal{A} \rightarrow \mathcal{A}'$ are rescaled gauge transformations $u \in \mathcal{G}$ such that $u_*\mathcal{A} = \mathcal{A}'$. Let $M = \text{Aut}(\text{id}_{\tilde{\mathcal{M}}})$; it is an abelian group under composition (by an Eckmann–Hilton argument), and a $G(\mathbb{K})$ -module. The obstruction to coherence is a group cohomology class $o \in H^2(G(\mathbb{K}); M)$. If one works over a field \mathbb{K} and asks that the homotopies be continuous in g_1 and g_2 the obstruction lies in the continuous group cohomology $H_{\text{cont}}^2(G(\mathbb{K}); M)$; if one wants them algebraic, it lies in an algebraic version of group cohomology.

We do not pursue these obstructions in detail, but content ourselves with an easy case. Work over \mathbb{C} , and let G_{an} denote $G(\mathbb{C})$ with its analytic (not Zariski) topology. We ask that the homotopies be continuous on $G_{\text{an}} \times G_{\text{an}} \times \mathbb{C}^5$. If we restrict g_1 and g_2 to the uni-triangular normal subgroup $U_{\text{an}} \subset G_{\text{an}}$ then the obstruction lies in $H_{\text{cont}}^2(U_{\text{an}}; M)$. But U_{an} is contractible, hence $H_{\text{cont}}^2(U_{\text{an}}; M) = H^2(BU_{\text{an}}; M) = H^2(\{\text{pt.}\}, M) = 0$. Thus we *can* make our homotopies coherent for U_{an} .

The Lyndon–Hochschild–Serre spectral sequence for $U_{\text{an}} \subset G_{\text{an}}$ is concentrated in a single row, which tells us that $H_{\text{cont}}^2(G_{\text{an}}; M) \cong H_{\text{cont}}^2(\mathbb{C}^\times, M^{U_{\text{an}}}) \cong M^{U_{\text{an}}}$. So $M^{U_{\text{an}}}$ is where to find the obstruction to extending the coherent homotopies from U_{an} to G_{an} .

5.2 The functor \mathcal{F} on hom-spaces

The proof of the equivalence clause of Theorem C will be given in the next section; for now, we shall set up the functor.

We have already set up the functor on objects—namely, we have constructed \mathcal{A}_W . Now take $g \in G$ and $w \in W$, and notice that they define an isomorphism $g_*: C_w \rightarrow C_{g(w)}$. We have a diagram of A_∞ quasi-isomorphisms

$$\begin{array}{ccccc}
 & & \mathcal{B}_{g(w), g(\mathcal{U}) \cup \mathcal{U}} & & \\
 & \swarrow f & & \searrow f' & \\
 \mathcal{B}_{w, \mathcal{U}} & \xrightarrow{g_*} & \mathcal{B}_{g(w), g(\mathcal{U})} & & \mathcal{B}_{g(w), \mathcal{U}} \\
 \downarrow p_w \uparrow i_w & & \swarrow \text{dotted } s & & \downarrow p_{g(w)} \uparrow i_{g(w)} \\
 \mathcal{A}_w & & & & \mathcal{A}_{g(w)}
 \end{array}$$

Here f and f' are maps of dga; they forget one of the three open sets in the covering $g(\mathcal{U}) \cup \mathcal{U} = \{U, V, g(V)\}$. The dotted arrow marked s is an A_∞ -homomorphism which is inverse to f , up to homotopy; s is still to be constructed. The A_∞ -maps i_w and p_w are mutual inverses up to homotopy; they are associated with the splitting r , via the homological perturbation lemma [50, Remark 1.13].

Once s has been constructed, we shall define $\mathcal{F}(g, w)$ as the A_∞ -composite $p_{g(w)} \circ f' \circ s \circ g_* \circ i_w$. For composition of A_∞ -functors, see [50, (1e)]. To obtain s , we consider the following picture (cf. [50, Cor. 1.14]):

$$\begin{array}{ccc}
 \mathcal{B}_{g(w), g(\mathcal{U})} & \xleftarrow{f} & \mathcal{B}_{g(w), g(\mathcal{U}) \cup \mathcal{U}} \\
 g^* p_w \downarrow & & \uparrow \tilde{i}_{g(w)} \\
 g^* \mathcal{A}_w & \xrightleftharpoons[v]{v^{-1}} & \tilde{\mathcal{A}}_{g(w)}
 \end{array}$$

The left vertical arrow is the pullback by g of the A_∞ -morphism $p_w: \mathcal{B}_{w, \mathcal{U}} \rightarrow \mathcal{A}_w$, produced by means of the splitting r and the homological perturbation lemma.


Claim: the $\mathbb{K}[W]$ -cochain complex $\mathcal{B}_{g(w), g(\mathcal{U}) \cup \mathcal{U}}$ admits a \mathbb{K}^\times -invariant splitting for its cocycles.

The proof of the claim will be given below. The splitting gives rise to a minimal A_∞ -structure $\tilde{\mathcal{A}}_{g(w)}$ on its cohomology A , and the A_∞ -quasi-isomorphism $\tilde{i}_{g(w)}$. The composite $(g^* p_w) \circ f \circ \tilde{i}_{g(w)}$ is an A_∞ -morphism inducing the identity map $A \rightarrow A$ on cohomology. Thus it is a gauge transformation $v \in C^1(A, A)^{\leq 0}$. As such, it has a strict inverse v^{-1} . We put $s = \tilde{i}_{g(w)} \circ v^{-1} \circ (g^* p_w)$. Then $s \circ f$ and $f \circ s$ are homotopic

to identity maps. (To prove this, use the fact that composition with a fixed functor preserves homotopy and observe that therefore, since $v = (g^*p_w) \circ f \circ \tilde{i}_{g(w)}$, one has $(g^*i_w) \circ v \circ \tilde{p}_w^* \simeq f$.)

We now consider the existence of homotopies $\mathcal{F}(h, g(w)) \circ \mathcal{F}(g, w) \simeq \mathcal{F}(hg, w)$. For this, contemplate the diagram


(34)



In the lower part of the diagram we see the juxtaposition of the maps that go into the definitions of $\mathcal{F}(g, w)$ (red arrows) and $\mathcal{F}(h, g(w))$ (magenta arrows). In the top row is the Čech complex associated with a 4-set open cover. The long, curved arrow pointing to it is a homotopy-inverse to the four-step composite formed by the blue arrows; it must be constructed.

The arrows marked g_* or h_* have strict inverses g_*^{-1} and h_*^{-1} . Now consider the arrow pointing down and right from the top of the diagram, shown in cyan. It is homotopic to the composite of five arrows going the other way round the hexagonal region. Moreover, the composite of the four blue arrows in (34) is equal to the forgetful map $\mathcal{B}_{hg(w), hg(U) \cup h(U) \cup U} \rightarrow \mathcal{B}_{hg(w), hg(U)}$ followed by $(hg)_*^{-1}$. As a result, we see that $\mathcal{F}(h, g(w)) \circ \mathcal{F}(g, w)$ is homotopic to the map $\mathcal{A}_w \rightarrow \mathcal{A}_{hg(w)}$ which factors through the curving and cyan arrows:

(35)



Therefore it will suffice to construct a homotopy-inverse to the latter forgetful map (this homotopy-inverse is indicated by the straight dotted arrow in (35)). For this, it is sufficient to prove the

Claim: the $\mathbb{K}[W]$ -cochain complex $\mathcal{B}_{hg(w), hg(\mathcal{U}) \cup h(\mathcal{U}) \cup \mathcal{U}}$ admits a \mathbb{K}^\times -invariant splitting for its cocycles.

The A_∞ -homomorphism $\mathcal{A}_w \rightarrow \mathcal{A}_{hg(w)}$ indicated by (35) is homotopic to $\mathcal{F}(hg, w)$. That is because the solid arrows make the following diagram commutative; hence the dotted homotopy-inverses form a homotopy-commutative diagram:

$$\begin{array}{ccc}
 & \mathcal{B}_{hg(w), hg(\mathcal{U}) \cup h(\mathcal{U}) \cup \mathcal{U}} & \\
 \nearrow & \downarrow & \searrow \\
 \mathcal{B}_{hg(w), hg(\mathcal{U})} & \mathcal{B}_{hg(w), hg(\mathcal{U}) \cup \mathcal{U}} & \mathcal{B}_{hg(w), \mathcal{U}}
 \end{array}$$

Hence $\mathcal{F}(hg, w) \simeq \mathcal{F}(h, g(w)) \circ \mathcal{F}(g, w)$.

It remains to prove the two claims highlighted above. Recall that we constructed a splitting for $\mathcal{B}_{w, \mathcal{U}}$ by constructing one for $\mathcal{B}'_{w, \mathcal{U}}$, meaning that we considered $\underline{\text{End}}(\mathcal{O} \oplus \mathcal{O}(-\sigma))$ instead of $\underline{\text{End}}(\mathcal{O} \oplus \mathcal{O}_\sigma)$. We shall do the same here. And as before, it suffices to consider the line bundles \mathcal{O} , $\mathcal{O}(\sigma)$ and $\mathcal{O}(-\sigma)$ which form the matrix entries for $\underline{\text{End}}(\mathcal{O} \oplus \mathcal{O}(-\sigma))$. These sheaves have cohomology in any given degree which is either zero (in which case the splitting of the cocycles is trivial) or is a free module of rank 1, which means that one can complement $\text{im } \delta$ in $\ker \delta$ by choosing any representative for the generator of cohomology—and the resulting splitting will be \mathbb{K}^\times -invariant. This establishes the two claims, and thereby completes the construction of the functor \mathcal{F} .

5.3 Comparison of deformation theories

Introduce the shifted truncated Hochschild cochain complex \mathcal{D}^\bullet , given by

$$\mathcal{D}^\bullet = (C^\bullet(A, A)^{\leq 0})[1].$$

A minimal A_∞ -structure on the algebra A is a sequence of maps $\mu^i \in \text{Hom}_{\mathbb{K}}^{2-i}(A^{\otimes i}, A)$ for $i \geq 2$ such that μ^2 is the multiplication for A . We can view the structure as a truncated Hochschild cochain $\mu^\bullet \in \mathcal{D}^1$. Thus the functor \mathcal{F} is defined, on objects, by a map $\mu^\bullet: W \rightarrow \mathcal{D}^1$. Pick a \mathbb{K}^\times -equivariant splitting of the cocycles for $\mathcal{B}_{\mathcal{C}}$, where \mathcal{C} is the universal Weierstrass curve over $\mathbb{K}[W]$; then we obtain, for each $w \in W$, a minimal A_∞ -structure $\alpha(w) := \mu_w^\bullet$ on A .

Lemma 5.3 *Pick $i > 2$ and homogeneous elements $x_1, \dots, x_i \in A$. The map $p: W_d \rightarrow A_{2-i+\sum |x_i|}$ given by $w \mapsto \mu_w^i(x_1, \dots, x_i)$ is zero except when $i - 2$ is a*

multiple of d . It is linear if $i = d + 2$, and in general is homogeneous of degree $(i - 2)/d$.

Proof The map p is a polynomial function, equivariant under \mathbb{K}^\times , meaning that $p(\epsilon^d w) = \epsilon^{i-2} p(w)$, and is therefore homogeneous, of degree $(i - 2)/d$. \square

We now consider the derivative at 0 of the map $w \mapsto \mathcal{A}_w$,

$$\lambda := D\mu^\bullet|_{w=0}: W \rightarrow \mathcal{D}^1.$$

Take $w \in W_d$. By definition, $\lambda_d(w)$ is a sequence (x^2, x^3, x^4, \dots) where $x^i \in \text{Hom}_{\mathbb{K}}^{2-i}(A^{\otimes i}, A)$. By Lemma 5.3, $x^i = 0$ except when $i = d + 2$. Thus, we consider $\lambda|_{W_d}$ as a map

$$\lambda_d = \lambda|_{W_d}: W_d \rightarrow (\mathcal{D}^1)_{-d}.$$

It is the map which assigns to w its *obstruction cocycle* (see e.g. [47, (3a)]). Because $\mu_w^i = 0$ for $i < d + 2$, we have $\delta \circ \lambda_d = 0$.

The effect of \mathcal{F} on morphisms is encoded in a map

$$u: G \times W \rightarrow \mathcal{D}^0,$$

which has a partial derivative

$$\left. \frac{\partial u}{\partial g} \right|_{g=1}: \mathfrak{g} \times W \rightarrow \mathcal{D}^0.$$

We define

$$\kappa^0: \mathfrak{g} \rightarrow \mathcal{D}^0, \quad \xi \mapsto \left. \frac{\partial u}{\partial g} \right|_{g=1}(\xi, 0).$$

When $\xi \in \mathfrak{g}_d$, the only non-vanishing component of $\kappa^0(\xi)$ lies in \mathcal{D}^0 .

Lemma 5.4 *The vertical maps κ^j in the diagram*

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathfrak{g} & \xrightarrow{d} & W & \longrightarrow & 0 \\ & & \downarrow \kappa^0 & & \downarrow \kappa^1 = \lambda & & \downarrow \kappa^2 \\ 0 & \longrightarrow & \mathcal{D}^{-1} & \xrightarrow{\delta} & \mathcal{D}^0 & \xrightarrow{\delta} & \mathcal{D}^1 & \xrightarrow{\delta} & \mathcal{D}^2 \end{array}$$

define a map κ^\bullet of cochain complexes.

Proof We have seen that $\delta \circ \kappa^1 = 0$. To prove that $\delta \circ \kappa^0 = \kappa^1 \circ d$, observe that the action of gauge transformations on A_∞ -structures is defined through a map $a: \mathcal{D}^0 \times \mathcal{D}^1 \rightarrow \mathcal{D}^1$. We have $a(u^g(0), \mu_0^\bullet) = \mu_{g(0)}^\bullet$; differentiating this relation with respect to g , and setting $g = 1$, we obtain the sought equation. \square

Theorem 5.5 *The maps κ^0 and κ^1 induce isomorphisms*

$$(36) \quad [\kappa^0]: \ker d \rightarrow \mathrm{HH}^1(A, A)^{\leq 0}$$

$$(37) \quad [\kappa^1]: \mathrm{coker} d \rightarrow \mathrm{HH}^2(A, A)^{\leq 0}$$

when the base ring R is

- (i) a field \mathbb{K} ; or
- (ii) the ring of integers \mathbb{Z} ; or more generally
- (iii) an integral domain of characteristic zero.

Proof (i) We begin with $[\kappa^1]$. We claim that $\ker[\kappa^1] = 0$. Indeed, if w lies in this kernel, let C_{wt} be the Weierstrass curve C over $\mathbb{K}[t]/t^2$ with parameters $a_i = tw_i$. Thus C_{wt} specializes to C_{cusp} at $t = 0$. The resulting minimal A_∞ -algebra \mathcal{A}_{wt} is then formal over $\mathbb{K}[t]/t^2$; indeed, its class $[\kappa^1(wt)] \in t\mathrm{HH}^2(A, A)$ exactly measures non-triviality of \mathcal{A}_{wt} as an A_∞ -deformation of A . Since the quasi-isomorphism class of $\mathcal{A}_{C_{wt}}$ determines the curve and its basepoint (Lemma 3.4), one finds that $C_{wt} \cong C_{\mathrm{cusp}} \times_{\mathrm{Spec} \mathbb{K}} \mathrm{Spec} \mathbb{K}[t]/t^2$. It follows that wt defines a trivial first-order deformation of C_{cusp} . Thus there is a Weierstrass automorphism $I + t\xi$ of $\mathbb{P}^2(\mathbb{K}[t]/t^2)$, carrying C_{cusp} to C_{wt} . Here $\xi \in \mathfrak{g}$ and one has $w = d\xi$.

We quote from [35] or Theorem B the result that, as graded \mathbb{K} -modules, we have

$$\mathrm{HH}^2(A, A)^{\leq 0} = \mathbb{K}/(2)^{s=-1} \oplus \mathbb{K}/(3)^{-2} \oplus \mathbb{K}/(2)^{-3} \oplus \mathbb{K}^{-4} \oplus \mathbb{K}^{-6}.$$

Thus $\mathrm{coker} \delta$ is abstractly isomorphic to $\mathrm{HH}^2(A, A)^{\leq 0}$ as a graded vector space, and hence the injection $[\kappa^1]$ is an isomorphism.

Now consider $[\kappa^0]: \ker d \rightarrow \mathrm{HH}^1(A, A)^{\leq 0}$. Recall that W is a \mathfrak{g} -module. On the level of cohomology, $W_{\mathfrak{g}} = \mathrm{coker} d$ is a $\ker d$ -module. This module is easily checked to be faithful: that is, the action homomorphism $\ker d \rightarrow \mathrm{End}(W_{\mathfrak{g}})$ is injective. Moreover, HH^2 is a HH^1 -module, and the map $[\kappa^\bullet]$ respects the actions on cohomology:

$$[\kappa^0\xi] \cdot [\kappa^1w] = [\kappa^1(\xi \cdot w)].$$

Given $\xi \in \ker d$, pick a $w \in W_{\mathfrak{g}}$ such that $\xi \cdot w \neq 0 \in W_{\mathfrak{g}}$. We then have $[\kappa^0(\xi)] \neq 0$. Hence $[\kappa^0]$ is injective. Both domain and codomain of $[\kappa^0]$ are isomorphic to

$$\mathbb{K}^{s=0} \oplus \mathbb{K}/(2)^{-1} \oplus \mathbb{K}/(3)^{-2} \oplus \mathbb{K}/(2)^{-3};$$

therefore $[\kappa^1]$ is an isomorphism.

(ii) It will be helpful to note at the outset that for any ring R one has $CC_R^\bullet(A \otimes R, A \otimes R) = CC^\bullet(A, A) \otimes R$, that $\mathfrak{g}_R = \mathfrak{g}_{\mathbb{Z}} \otimes R$, and that these canonical isomorphisms are compatible

with the construction of the map κ^\bullet . Let $\mathbb{K} = \mathbb{Z}/(p)$ be a residue field of \mathbb{Z} . Form the \mathbb{Z} -cochain complex $K^\bullet = \text{cone } \kappa^\bullet$, a complex of free abelian groups, and note, using (i), that $H^1(K^\bullet \otimes \mathbb{K}) = H^2(K^\bullet \otimes \mathbb{K}) = 0$ while $\text{HH}_{\mathbb{K}}^0(A \otimes \mathbb{K}, A \otimes \mathbb{K})^{\leq 0} \rightarrow H^0(K^\bullet \otimes \mathbb{K})$ is an isomorphism (so the latter vector space is 1-dimensional). By universal coefficients [62, 3.6.2], $H^j(K^\bullet \otimes \mathbb{K})$ has a direct summand $H^j(K^\bullet) \otimes \mathbb{K}$. Consequently $H^1(K^\bullet) \otimes \mathbb{K} = H^2(K^\bullet) \otimes \mathbb{K} = 0$. Since $H^j(K^\bullet)$ is a finitely-generated abelian group, one deduces $H^1(K^\bullet) = H^2(K^\bullet) = 0$. One then has $\text{Tor}_{\mathbb{Z}}^1(H^1(K^\bullet), \mathbb{K}) = 0$, and so by universal coefficients again, $H^0(K^\bullet) \otimes \mathbb{K} \cong H^0(K^\bullet \otimes \mathbb{K}) \cong \mathbb{K}$. Hence $H^0(K^\bullet)$ is free of rank 1. Over \mathbb{Z} , one has $\ker \delta \cong \mathbb{Z}$, and a look at the exact sequence of the mapping cone then tells us that $\text{HH}^0(A, A)^{\leq 0} \rightarrow H^0(K^\bullet)$ is an isomorphism. The same exact sequence then tells us that $[\kappa^0]$ and $[\kappa^1]$ are isomorphisms.

(iii) Take the \mathbb{Z} -cochain complex K^\bullet as before. Universal coefficients now tells us that

$$\begin{aligned} H^0(K^\bullet \otimes_{\mathbb{Z}} R) &\cong H^0(K^\bullet) \otimes_{\mathbb{Z}} R \oplus \text{Tor}_{\mathbb{Z}}^1(H^1(K^\bullet), R) \cong R, \\ H^1(K^\bullet \otimes_{\mathbb{Z}} R) &\cong H^1(K^\bullet) \otimes_{\mathbb{Z}} R \oplus \text{Tor}_{\mathbb{Z}}^1(H^2(K^\bullet), R) = 0, \\ H^2(K^\bullet \otimes_{\mathbb{Z}} R) &\cong H^2(K^\bullet) \otimes_{\mathbb{Z}} R \oplus \text{Tor}_{\mathbb{Z}}^1(H^3(K^\bullet), R) = \text{Tor}_{\mathbb{Z}}^1(H^3(K^\bullet), R). \end{aligned}$$

The ring R is torsion-free as an abelian group, so $\text{Tor}_{\mathbb{Z}}^1(H^3(K^\bullet), R) = \text{Tor}_{\mathbb{Z}}^1(R, H^3(K^\bullet)) = 0$. The exact sequence of the mapping cone then tells us that $[\kappa^1]$ is an isomorphism. Over R , one has $\ker[\delta: \mathfrak{g} \rightarrow W] \cong R$, again because R is torsion-free. Part of the exact sequence of the mapping cone reads

$$0 \rightarrow \text{HH}^0(A_R, A_R)^{\leq 0} \rightarrow R \rightarrow R \xrightarrow{[\kappa^0]} \text{HH}^1(A_R, A_R)^{\leq 0} \rightarrow 0,$$

and since $\text{HH}^0(A_R, A_R)^{\leq 0}$ is non-zero, and R an integral domain, the map $R \rightarrow R$ must be zero. Hence $[\kappa^0]: R \rightarrow \text{HH}^1(A_R, A_R)^{\leq 0}$ is also an isomorphism. \square

Remark. Out of caution, work over a field \mathbb{K} in this remark. The map $[\kappa^1]$ has a straightforward deformation-theoretic meaning (a first-order deformation of Weierstrass curves gives a first-order deformation of A_∞ -structures on A). The map $[\kappa^0]$ may then be characterized as the unique map that makes the following diagram commute:

$$\begin{array}{ccc} \ker d & \xrightarrow{\quad} & \text{End } W_{\mathfrak{g}} \\ \downarrow [\kappa^0] & & \downarrow [\kappa^1] \\ \text{HH}^1(A, A)^{\leq 0} & \xrightarrow{\quad} & \text{End } \text{HH}^2(A, A)^{\leq 0}. \end{array}$$

The horizontal arrows are the module actions. This leads to a derived-categorical construction of $[\kappa^0]$ as the composite of the canonical isomorphisms of graded \mathbb{K} -

modules

$$\begin{aligned}
 \ker d &\xrightarrow{\cong} H^0(C_{\text{cusp}}, \mathcal{T})^{\leq 0} && \text{Prop. 4.19 (2)} \\
 &\xrightarrow{\cong} HH^1(C_{\text{cusp}})^{\leq 0} && \text{Lemma 4.17} \\
 &\xrightarrow{\cong} HH^1(\text{tw vect}(C_{\text{cusp}}))^{\leq 0} && \text{Prop. 4.3} \\
 &\xrightarrow{\cong} HH^1(\mathcal{B}'_{\text{cusp}}, \mathcal{B}'_{\text{cusp}})^{\leq 0} && \text{Prop. 4.4} \\
 &\xrightarrow{\cong} HH^1(A, A)^{\leq 0} && \text{Prop. 3.5.}
 \end{aligned}$$

Remark. Before embarking on the proof of Theorem C, we say a word about the methodology. We are very close here to the framework for deformation theory which uses differential graded Lie algebras (DGLA). For instance, the DGLA $\mathcal{K}^\bullet = C^{\bullet-1}(A, A)^{\leq 0}$ determines, in characteristic zero, a deformation functor which assigns to a local artinian \mathbb{K} -algebra A the solutions in $\mathcal{K}^1 \otimes m_A$ to the Maurer–Cartan equation $\delta\mu^\bullet + \frac{1}{2}[\mu^\bullet, \mu^\bullet] = 0$, modulo the group $\exp(\mathcal{K}^0 \otimes m_A)$ (see e.g. [37]). A standard approach would be to show that κ^\bullet is a map of DGLA, and conclude, given its effect on cohomology, that the deformation theories controlled by $\mathfrak{g} \oplus W$ and \mathcal{K}^\bullet coincide [37].

The route we have actually taken is a variant of this standard approach. Everything here works over arbitrary fields; there is no need for characteristic zero. A minimal A_∞ -structure on A is a Hochschild cochain $\mu^\bullet \in C^2(A, A)^{\leq 0}$ which satisfies $\delta\mu^\bullet + \mu^\bullet \circ \mu^\bullet = 0$. The Gerstenhaber square $\mu^\bullet \circ \mu^\bullet$ agrees with $\frac{1}{2}[\mu^\bullet, \mu^\bullet]$ when 2 is invertible. There are no artinian rings in the picture, but the length filtration of Hochschild cochains serves as a substitute. The map κ^\bullet is not quite a Lie algebra homomorphism; it might be possible to promote it to an L_∞ -homomorphism, but we have chosen to use the functor \mathcal{F} more directly.

Proof of Theorem C We have set up the functor $w \mapsto \mathcal{A}_w$ over regular rings R , and must prove that it is an equivalence in the stated sense.

Essential surjectivity: If two minimal A_∞ -structures on A , with composition maps \mathbf{m}^k and \mathbf{n}^k , agree for $k \leq 8$, then the two structures are gauge-equivalent. Indeed, one proves inductively that if $\mathbf{m}^k = \mathbf{n}^k$ for $k < d$ then one can find a gauge transformation u such that $(u_*\mathbf{m})^k = \mathbf{m}^k$ for $k < d$ and $(u_*\mathbf{m})^d = \mathbf{n}^d$. To do so, one notes that $\mathbf{m}^d - \mathbf{n}^d$ defines a class in $HH^2(A, A)^{2-d}$. This Hochschild module is zero for $d > 8$: when R is a field, this holds by Theorem B; when $R = \mathbb{Z}$, it then follows by the universal coefficients; when R is an integral domain of characteristic zero, it then follows by universal coefficients from the \mathbb{Z} case (cf. the proof of Theorem 5.5). One uses a trivialization of $\mathbf{m}^d - \mathbf{n}^d$ to define the gauge transformation (cf. [47]).

Now let $\mathcal{A} = (A, \mu^\bullet)$ be a minimal A_∞ -structure on A , over \mathbb{Z} . Our goal is to show that \mathcal{A} is gauge-equivalent to \mathcal{A}_w for some $w \in W$. We shall repeatedly apply gauge transformations to \mathcal{A} , without notating them. If \mathcal{A} is formal then $\mathcal{A} \simeq \mathcal{A}_0$. If it is not formal then we apply a gauge transformation so as to arrange that $\mu^k = 0$ for $2 < k < d$ but $[\mu^d] \neq 0 \in \mathrm{HH}^2(A, A)^{2-d}$. By Theorem 5.5, one has $[\mu^d] = [\kappa^1(w)]$ for a unique $w \in W_{d-2}$; we may assume, by applying another gauge transformation, that in fact $\mu^d = \kappa^1(w)$. By Lemma 5.3, $\mu_w^k = 0$ for $k < d$. Thus $\mu^k = \mu_w^k$ for $k \leq d$. The difference $\mu^{d+1} - \mu_w^{d+1}$ is a cocycle, as one checks using the A_∞ -relations. If it is exact, one can adjust μ^d by a gauge transformation which leaves μ^k untouched for $k \leq d$ such that μ^{d+1} equals μ_w^{d+1} , whereupon $\mu^{d+2} - \mu_w^{d+2}$ is a cocycle, and we can repeat the process. What we find is that either μ^\bullet is gauge-equivalent to μ_w^\bullet , or else there is a $d' \leq 8 - d$ such that, after applying a gauge transformation to μ^\bullet , one has $\mu^{d+k} = \mu_w^{d+k}$ for $k = 0, \dots, d' - 1$ and $[\mu^{d+d'} - \mu_w^{d+d'}] \neq 0$. In the latter case, write $\mu^{d+d'} - \mu_w^{d+d'} = \kappa^1(w')$, and consider $\mu_{w+w'}^\bullet$. The differences $\mu^k - \mu_{w+w'}^k$ can be killed by gauge transformations for $k \leq d + d'$. We continue in the same fashion; the process stops once has made μ^k agree with $\mu_{w+w'+\dots}^k$ for $k \leq 8$.

Faithfulness: We shall consider the case of automorphisms: say $g(w) = w$, and that $u^g(w) \simeq \mathrm{id}$ (\simeq means ‘is homotopic to’); we shall show that then $g = 1$. From this, the general case will follow: say $g_1(w) = g_2(w)$ and $u^{g_1}(w) \simeq u^{g_2}(w)$. Let $g = g_1^{-1}g_2$. Then $g(w) = w$ and $u^g(w) \simeq (u^{g_1})^{-1} \circ u^{g_2} \simeq \mathrm{id}$.

Recall that the Lie algebra \mathfrak{g} is a graded R -module: $\mathfrak{g} = \mathfrak{g}_0 \oplus \mathfrak{g}_{-1} \oplus \mathfrak{g}_{-2} \oplus \mathfrak{g}_{-3}$. There is an induced filtration by Lie subalgebras

$$\mathfrak{g} = \mathfrak{g}_{\leq 0} \supset \mathfrak{g}_{\leq -1} \supset \mathfrak{g}_{\leq -2} \supset \mathfrak{g}_{\leq -3} \supset \mathfrak{g}_{\leq -4} = 0, \quad \mathfrak{g}_{\leq j} = \bigoplus_{k \leq j} \mathfrak{g}_k.$$

One easily finds algebraic subgroups $G \supset G_{-1} \supset G_{-2} \supset G_{-3} \supset G_{-4} = \{1\}$, with $\mathrm{Lie} G_j = \mathfrak{g}_{\leq j}$; for instance, G_{-1} is the unipotent subgroup U .

We also filter the group \mathcal{G} , so that $\mathcal{G} = \mathcal{G}_0 \supset \mathcal{G}_{-1} \supset \dots$, as follows: take $v = (v^1, v^2, v^3, \dots) \in \mathcal{G}$ (so $v^j \in C^1(A, A)^{1-j}$, and $v^1 = c \mathrm{id}$ with c a unit). If $v^1 = \mathrm{id}$, say $v \in \mathcal{G}_{-1}$; if in addition, $v^k = 0$ for $1 < k \leq d - 1$, say $v \in \mathcal{G}_{1-d}$. If $g \in G_{1-d}$ then $u^g(w) \in \mathcal{G}_{1-d}$. If in addition $d > 2$ then we have $[u_d] \neq 0 \in \mathrm{HH}^1(A, A)^{1-d}$. Hence we have the following observation:

■ if $g(w) = w$, if $u = u^g(w) \in \mathcal{G}_{1-d}$ for some $d > 2$, and if $[u^d] \neq 0$, then $g \in G_{-d}$.

Now let $u^g(w) = (u^1, u^2, u^3, \dots)$, where $u^j \in C^1(A, A)^{1-j}$, with $g(w) = w$ and $\mathrm{id} \simeq u^g(w)$. Since $u^g(w) \simeq \mathrm{id}$, u^1 induces the identity on $A = H^* \mathcal{A}_w$. Hence g lies in the unipotent subgroup $U = G_{-1} \subset G$. Hence $u^1 = \mathrm{id}$, i.e., $u^g \in \mathcal{G}_{-1}$. Thus $\delta u^2 = 0$, but the fact that $u^g(w) \simeq \mathrm{id}$ implies that u^2 is a coboundary. By

the highlighted observation, we deduce that $g \in G_{-2}$. We argue similarly that u^2 is a coboundary, hence that $g \in G_{-3}$, and finally that u^4 is a coboundary, hence that $g \in G_{-4} = \{1\}$.

Fullness. Next consider the assertion that given $v \in \mathcal{G}$ and given w_1, w_2 such that $v_*\mathcal{A}_{w_1} = \mathcal{A}_{w_2}$, there is some $g \in G$ such that $g(w_1) = w_2$ and $v \simeq u^g(w_1)$.

We first note a non-emptiness statement: if $v_*\mathcal{A}_{w_1} = \mathcal{A}_{w_2}$ where $v \in \mathcal{G}$, then there is some $h \in G$ such that $h(w_1) = w_2$. This is true because the formal diffeomorphism-type of the A_∞ -structure \mathcal{A}_w determines the affine coordinate ring, and hence the curve (Lemma 3.4). That is:

$$\text{mor}_{\mathcal{M}}(\mathcal{A}_{w_1}, \mathcal{A}_{w_2}) \neq \emptyset \Rightarrow \text{mor}_{\mathcal{W}}(w_1, w_2) \neq \emptyset.$$

Likewise, if $v_*\mathcal{A}_{w_1} = \mathcal{A}_{w_2}$ where $v \in \mathcal{G}_{-1}$, then there is some $h \in U = G_{-1}$ such that $h(w_1) = w_2$. Indeed, if we know \mathcal{A}_w up to gauge-equivalence then we can reconstruct not only the affine coordinate ring of the curve, but also the Weierstrass differential on the curve.

Hence, since \mathcal{W} and \mathcal{M} are groupoids, and \mathcal{F} functorial, it will be enough to prove fullness when $w_1 = w_2$. We wish to show that the map

$$\mathcal{F}_{w,w}: \text{mor}_{\mathcal{W}}(w, w) \rightarrow \text{mor}_{\mathcal{M}}(\mathcal{A}_w, \mathcal{A}_w),$$

which we already know to be injective, is also surjective. Moreover, by a rescaling argument, we see that it suffices to prove this under the assumption that $v \in \mathcal{G}_{-1}$.

Take some $v \in \mathcal{G}_{-1}$ with $v_*\mathcal{A}_w = \mathcal{A}_w$. If v is homotopic to the identity, we are done. If not then there is a $d > 1$ such that v is homotopic to some $x \in \mathcal{G}_{-d}$, $x_*\mathcal{A}_w = \mathcal{A}_w$, where $[x^d] \neq 0 \in \text{HH}^1(A, A)^{1-d}$. By Theorem 5.5, $[x^d] = [\kappa^0(\xi_1)]$ for some $\xi_1 \in \mathfrak{g}_d$. We then have $g_1 := 1 + \xi_1\epsilon \in U(\mathbb{K}[\epsilon]/\epsilon^2)$, and $g_1([w]) = [w] \in W \otimes_{\mathbb{K}} \mathbb{K}[\epsilon]/\epsilon^2$. It is possible to lift g_1 to $g_2 = 1 + \xi_1\epsilon + \xi_2\epsilon^2 \in U(\mathbb{K}[\epsilon]/\epsilon^3)$, with $\xi_2 \in \mathfrak{g}_{-d}$. We wish to do so in such a way that $g_2([w]) = [w] \in W \otimes_{\mathbb{K}} \mathbb{K}[\epsilon]/\epsilon^2$. The obstruction is the class of $g_2(w) - w$ in $T_w W / \text{im } \rho_w = \text{coker } \rho_w$. Now, $[\kappa^1]$ maps W to $\text{HH}^2(A, A)^{\leq 0}$, and carries $\text{im } \rho_w$ to $[\mathcal{A}_w, \text{HH}^1(A, A)^{\leq 0}]$. Because x^d extends to a gauge transformation which preserves \mathcal{A}_w , $[\kappa^1(g_2(w) - w)] \in [w, \text{HH}^1(A, A)^{\leq 0}]$. Hence $g_2(w) - w \in \text{im } \rho_w$. Inductively, we extend g_1 to $1 + \xi_1\epsilon + \dots + \xi_3\epsilon^3 \pmod{\epsilon^4}$ with $\xi_j \in \mathfrak{g}_{\leq j}$. We can then find a homomorphism $\theta: \mathbb{K} \rightarrow \mathfrak{g}$ which has this series as its 3-jet, and we put $g = \theta(1)$. One has $g(w) = w$.

We next ask whether $(u^g)^{-1} \circ v$ is homotopic to $\text{id}: \mathcal{A}_w \rightarrow \mathcal{A}_w$. The obstruction we encounter now lies in $\text{HH}^1(A, A)^{1-d'}$ with $d' > d$. Repeating the argument, we eventually obtain an element $g \in U$ with $u^g \simeq v$. \square

5.4 The Gerstenhaber bracket on $\mathrm{HH}^\bullet(A, A)^{\leq 0}$

Note: This subsection will not be used elsewhere, except in the variant method (7.1.1) of our proof-by-elimination that \mathcal{T}_0 is mirror to T_0 .

The truncated Hochschild cohomology $H^\bullet = \mathrm{HH}^\bullet(A, A)^{\leq 0}$, in addition to its graded algebra structure, carries the Gerstenhaber bracket $[\cdot, \cdot]$. These two operations make H^\bullet a Gerstenhaber algebra over \mathbb{K} : the product is graded commutative, the bracket makes $H^\bullet[1]$ a Lie superalgebra, and the bracket is a bi-derivation for the product. Moreover, the internal grading s is additive under the bracket. Since the brackets can readily be computed from our results, we record it here. We use the notation of Theorem B.

Theorem 5.6 $Q^{\bullet, \leq 0}$ carries a Gerstenhaber bracket $[\cdot, \cdot]$, respecting the internal grading s and making the (canonical) isomorphism $\mathrm{HH}^\bullet(A, A)^{\leq 0} \rightarrow Q^{\bullet, \leq 0}$ from Theorem B a map of Gerstenhaber algebras. The Gerstenhaber bracket on $Q^{\bullet, \leq 0}$ is as follows. The Lie algebra $L = (Q^1)^{\leq 0} = (T \otimes \gamma)^{\leq 0}$ is given by

$$\begin{aligned} & \mathbb{K}\gamma, & & \text{if } 6 \neq 0 \\ & \mathbb{K}\gamma \oplus \mathbb{K}x\gamma, & [\gamma, x\gamma] = -\gamma, & \text{if } 3 = 0 \\ & \mathbb{K}\gamma \oplus \mathbb{K}x\gamma \oplus \mathbb{K}y\gamma, & [\gamma, x\gamma] = 0, [\gamma, y\gamma] = \gamma, [x\gamma, y\gamma] = x\gamma, & \text{if } 2 = 0. \end{aligned}$$

One has $[Q^2, Q^2] = 0$. The brackets $[L, Q^2]$ are given by

$$\begin{aligned} & [\gamma, x\beta] = -4x\beta, \quad [\gamma, \beta] = -6\beta, & & \text{if } 6 \neq 0; \\ & [x\gamma, x^2\beta] = x^2\beta, \quad [x\gamma, x\beta] = -x\beta, \quad [x\gamma, \beta] = 0, \\ & [\gamma, x^2\beta] = x\beta, \quad [\gamma, x\beta] = -\beta, \quad [\gamma, \beta] = 0 & & \text{if } 3 = 0; \\ & [y\gamma, xy\beta] = xy\beta, \quad [y\gamma, y\beta] = y\beta, \quad [y\gamma, x\beta] = 0, \quad [y\gamma, \beta] = 0, \\ & [x\gamma, xy\beta] = 0, \quad [x\gamma, y\beta] = x\beta, \quad [x\gamma, x\beta] = 0, \quad [x\gamma, \beta] = 0, \\ & [\gamma, xy\beta] = x\beta, \quad [\gamma, y\beta] = \beta, \quad [\gamma, x\beta] = 0, \quad [\gamma, \beta] = 0, & & \text{if } 2 = 0. \end{aligned}$$

$Q^{\bullet, \leq 0}$ is generated as a unital \mathbb{K} -algebra by L and Q^2 . The remaining brackets are therefore determined by the Leibniz rule.

Proof The structure of L was already computed in Prop. 4.19; the first assertion here is essentially a restatement. A Hochschild 2-cocycle c for A determines a first-order deformation of A as an A_∞ -algebra, and that the Gerstenhaber square $[c \circ c] \in \mathrm{HH}^3(A, A)$ is the obstruction to lifting it to a second-order deformation. By Theorem 5.5, all first-order deformations of A come (via κ^1) from deformations of C_{cusp} as a Weierstrass curve. These Weierstrass deformations lift to second order;

hence the algebraic deformation of A also lifts. This shows that $c \circ c = 0$. The bracket on HH^2 is given by $[a, b] = a \circ b + b \circ a = (a + b) \circ (a + b) - a \circ a - b \circ b$, so $[\mathrm{HH}^2, \mathrm{HH}^2] = 0$. The adjoint action of the Lie algebra $\mathrm{HH}^1(A, A)^{\leq 0}$ on HH^2 is the natural action of infinitesimal (A_∞) automorphisms of A on first order deformations. But the Lie algebra $\mathrm{H}^0(\mathcal{T}_{C_{\mathrm{cusp}}})^{\leq 0}$ acts via Lie derivatives on Weierstrass deformations of C_{cusp} . Namely, take $\xi \in \ker d = \mathrm{H}^0(\mathcal{T}_{C_{\mathrm{cusp}}})^{\leq 0}$. We have a vector field $\rho(\xi)$ on W . Take $w \in W$, and regard it as a translation-invariant vector field on W ; then the adjoint action of ξ on w is given by the Lie bracket of vector fields: $w \mapsto [\rho(\xi), w](0)$. One computes these brackets using the formulae (15). The maps $[\kappa^*]$ from (5.5) intertwine this Lie derivative with the adjoint action of HH^1 on HH^2 . The statement about generation is clear from the definition of $Q^{\bullet, \leq 0}$. \square

Remark. Once the Lie algebra L has been computed, another approach to obtaining the rest of the brackets is to compute the BV operator $\Delta: Q^{\bullet, s} \rightarrow Q^{\bullet-1, s}$. Along with the product, that determines the brackets, and it is sharply constrained by the bidegrees (which, for instance, force $\Delta\beta = 0$). We have used this method as an independent check of the above formulae for $[L, L]$ and $[L, Q^2]$ —in particular, as a check on the signs.

Part II

Symplectic geometry and the mirror map

6 Fukaya categories

6.1 The exact Fukaya category

T_0 as a Liouville manifold.

Lemma 6.1 *Let (S, ω) be a symplectic surface, and γ a simple closed curve in S . Then there exists a vector field λ defined near γ which has the Liouville property that $\mathcal{L}_\lambda \omega = \omega$, which is nowhere tangent to γ , and which determines a chosen co-orientation for γ .*

Proof By the Lagrangian neighborhood theorem, we may assume that $S = T^*S^1$ with the canonical 2-form $-d(p dq)$ (here p is the fiber coordinate, q the base coordinate) with γ the zero-section. The vector fields $\lambda_c = (p \pm 1)\partial_p$ are Liouville and point in opposite directions along the zero-section. \square

If now (S, ω) is a closed, connected symplectic surface, and $c = \partial D$ the loop bounding a closed disc $D \subset S$, let $S^\circ = S \setminus \text{int}(D)$. Choose a Liouville vector field λ near c in S , pointing strictly out of S° , which is *exact* in the sense that $\int_{S^\circ} \omega = \int_{\partial S} \iota(\lambda)\omega$. There is then no obstruction to extending the 1-form $\iota(\lambda)\omega$ (defined in an annular neighborhood of c) to a primitive θ for ω on S° . The 1-form θ makes S° a Liouville domain [48], i.e. $\omega = d\theta$ is symplectic, and the Liouville vector field λ , defined by $\iota(\lambda)\omega = \theta$, points strictly outwards along the boundary.

Let $\hat{S} = S^\circ \cup_c (c \times [0, \infty))$ be the cylindrical completion of S° ; it comes with a 1-form $\hat{\theta}$ that agrees with θ on S° , given by $\hat{\theta} = e^s \theta|_c$ on $c \times [0, \infty)$. We write $\hat{\omega} = d\hat{\theta}$. The data $(\hat{S}, \hat{\theta}, c)$ define a *Liouville manifold* [48]: an exact symplectic manifold (M, θ) whose Liouville vector field λ is complete, with a given compact hypersurface H enclosing a compact domain N in M , such that λ is nowhere vanishing on $M \setminus \text{int}(N)$. Notice that if $z \in D$, one can identify \hat{S} smoothly with $S \setminus \{z\}$, but not symplectically, because $\hat{\omega}$ has infinite area.

Proposition 6.2 *If (S, ω) is a closed and connected surface, the Liouville-isomorphism class of $(\hat{S}, \hat{\theta}, c)$ depends only on (S, ω) .*

A Liouville isomorphism $f: (M_0, \theta_0) \rightarrow (M_1, \theta_1)$ is a diffeomorphism such that $f^*\theta_1 - \theta_0 = dK$ for a compactly supported function $K \in C_c^\infty(M)$.

Proof The choices involved are of the disc D , the exact Liouville vector field λ near $c = \partial D$, and the primitive θ for ω . Let $(D', c', \lambda', \theta')$ be another set of choices, resulting in a Liouville manifold \hat{S}' . Notice first that if the primed data are obtained by pulling back the unprimed data via some $f \in \text{Aut}(S, \omega)$ then there is a Liouville isomorphism $\hat{S} \cong \hat{S}'$. We may therefore assume that $\text{int}(D) \cap \text{int}(D') \neq \emptyset$. There is then a disc $D'' \subset \text{int}(D) \cap \text{int}(D')$. From this we infer that it suffices to consider the case where $D' \subset \text{int}(D)$. In that case, we have a symplectic cylinder $C \subset S$ cobounding c and c' which comes with a Liouville vector field $\lambda|_C$ that points inwards along c and outwards along c' . Let E (resp. E') be the complement of $\text{int}(S^\circ)$ in \hat{S} (resp. in \hat{S}'). By matching up the Liouville flows, one can find a Liouville isomorphism $E \rightarrow E'$ which is the identity on the boundary curve c .

The remaining task is then to consider the following situation: one has two 1-forms $\hat{\theta}$ and $\hat{\theta}'$ on \hat{S} , both primitives for $\hat{\omega}$, and the closed 1-form $\eta := \hat{\theta} - \hat{\theta}'$ has compact support. We want to find a compactly supported symplectic automorphism ϕ_1 such that $\phi_1^*\hat{\theta} - \hat{\theta}' = dK$, where K has compact support. The required ϕ_1 is the time 1 map of the flow ϕ_t of the symplectic vector field v dual to $-\eta$. Indeed, one computes that $\frac{d}{dt}\phi_t^*\hat{\theta} = -\phi_t^*\eta + d(\phi_t^*(\iota(v)\theta))$, whence

$$\phi_1^*\hat{\theta} - \hat{\theta}' = -\int_0^1 \phi_t^*\eta + dH,$$

where H has compact support in \hat{S} . For any closed loop $L \subset S^\circ$, one therefore has

$$\int_\gamma \phi_1^* \hat{\theta} - \hat{\theta} = - \int_0^1 \left(\int_{\phi_t(\gamma)} \eta \right) dt = - \int_\gamma \eta,$$

i.e., $[\phi_1^* \hat{\theta}] = [\hat{\theta}'] \in H^1(S^\circ; \mathbb{R})$. Since $H^1(S^\circ, \partial S^\circ) \subset H^1(S^\circ)$, it follows that $\phi_1^* \hat{\theta} - \hat{\theta}' = dK$ for some compactly supported K . \square

Now suppose that $T = \mathbb{R}^2/\mathbb{Z}^2$ is the 2-torus, and ω a symplectic form of total area 1. Let $z \in T$, and write $T_0 = T \setminus \{z\}$. Let c_z be a loop encircling z . By identifying T_0 with the cylindrical completion \hat{T} of the complement of a disc $D \supset z$, we endow \hat{T} with the structure of a Liouville manifold with symplectic form $\hat{\omega} = d\hat{\theta}$ and contact cross-section c_z . According to the proposition, the Liouville isomorphism type of T_0 is determined by (T, ω) .

Grading. We specify a *grading* of T_0 as a symplectic manifold, that is, a trivialization of the square of the canonical line bundle; in two dimensions, that amounts to an unoriented line field $\ell \subset T(T_0)$. Anticipating our discussion of $\mathcal{F}(T, z)$, we shall always choose ℓ extends over z to a line field on T . Such line fields form a torsor for $C^\infty(T, \mathbb{R}P^1)$; one has $\pi_0 C^\infty(T, \mathbb{R}P^1) = H^1(T; \mathbb{Z})$.

Branes. The objects of $\mathcal{F}(T_0)^{\text{ex}}$, the exact Fukaya category of the Liouville manifold T_0 , as defined in [50, chapter 2], are ‘exact Lagrangian branes’ $L^\#$:

embedded closed curves $L \subset T_0$ such that $\int_L \theta = 0$, equipped with spin structures and gradings.

A grading for L is a homotopy-class of paths from $\ell|_L$ to TL inside $T(T_0)|_L$. If L and L' are graded curves then a transverse intersection point $y \in L \cap L'$ has a degree $i(y) = \lfloor \alpha/\pi \rfloor + 1$, where α is the net rotation of the path from $T_y L$ to ℓ_y to $T_y L'$. When ℓ is *oriented*—pointing along β , say—the grading for a curve L induces an orientation for L . The sign $(-1)^{i(y)}$ is then the intersection sign $[L'] \cdot [L]$ (note the order!), regardless of which orientation for ℓ was selected; see Figure 1 for relevant examples. There are four inequivalent spin structures on each curve L , defined by the two orientations and the two double coverings.

The morphism-spaces in $\mathcal{F}(T_0)^{\text{ex}}$ are \mathbb{Z} -linear cochain complexes; $\text{hom}_{\mathcal{F}(T_0)^{\text{ex}}}(L^\#, L'^\#)$ is a cochain complex $CF(\phi(L^\#), L'^\#) = \mathbb{Z}^{\phi(L) \cap L'}$ computing the Floer cohomology $HF(L^\#, L'^\#)$; here $\phi(L)$ is the image of L under the time-1 map ϕ of a specified Liouville isotopy. The degrees $i(y)$ define the grading of the complex; the spin structures determine the sign-contributions of the holomorphic bigons defining the differential.

Floer cohomology is invariant under Liouville isotopies ϕ_t (i.e., 1-parameter families of Liouville automorphisms) in that $HF(\phi_t(L^\#), L^\#) \cong HF(L^\#, L^\#)$. As a result, Liouville-isotopic objects are quasi-isomorphic in $\mathcal{F}(T_0)^{\text{ex}}$.

One has $\text{hom}_{\mathcal{F}(T_0)^{\text{ex}}}(L^\#, L^\#) \simeq \mathbb{Z}[-1] \oplus \mathbb{Z}$: this is just a formula for the Morse cochain complex for L . If L^* denotes the brane sharing the same oriented, graded curve as $L^\#$ but with the other double covering, then one has $\text{hom}_{\mathcal{F}(T_0)^{\text{ex}}}(L^\#, L^*) \simeq \{\mathbb{Z} \xrightarrow{2} \mathbb{Z}[-1]\}$, which is the Morse complex for L with the local system with fiber \mathbb{Z} and holonomy -1 .

The composition map $\mu^2: \text{hom}(L_1, L_2) \otimes \text{hom}(L_0, L_1) \rightarrow \text{hom}(L_0, L_2)$ can be understood purely combinatorially [50, (13b)], provided that L_0, L_1 and L_2 are in general position. The coefficient for y_0 in $\mu^2(y_2, y_1)$ is a signed count of immersed triangles bounding L_0, L_1 and L_2 in cyclic order, with convex corners at y_1, y_2 and y_0 . We shall discuss the sign when we discuss the relative Fukaya category. When (L_0, L_1, L_2) are not in general position, one moves L_1 and L_2 by exact isotopies so as to make them so.

Beyond the differential μ^1 and the composition μ^2 , the higher A_∞ -structure maps μ^d are defined through inhomogeneous pseudo-holomorphic polygons for a complex structure j on T_0 compatible with the orientation and the Liouville structure at infinity. We refer to [50] for the foundations; since we shall not make any direct calculations with the higher structure maps, we need not say more here.

Describing the objects. Oriented simple closed curves in T_0 do not realize all free homotopy classes of loops, but just one free homotopy class per non-zero homology class. Indeed, suppose that γ_0 and γ_1 are two such curves representing the same class in $H_1(T_0)$. It follows from the ‘bigon criterion’ [18] that they can be disjointed by isotopies in T_0 . Assuming that they are disjoint, they divide T into two annuli. Only one of the annuli contains z , and hence an isotopy from γ_0 to γ_1 can be realized in T_0 . If, moreover, γ_0 and γ_1 are isotopic simple closed curves which are both exact then they cobound an immersed annulus of area zero, and hence are Liouville-isotopic. A non-exact simple closed curve representing a non-trivial homology class can be isotoped to an exact one (just take an isotopy of an appropriate flux).

Generation. Pick a basis (α, β) for $H_1(T_0; \mathbb{Z})$ with $\alpha \cdot \beta = 1$. Let $L_0^\#$ and $L_\infty^\#$ be objects of $\mathcal{F}(T_0)^{\text{ex}}$ of respective slopes α and β . In Figure 1, the line field ℓ is chosen to be parallel to β . With that choice, we grade L_0 by the trivial homotopy from $\ell|_{L_0}$ to TL_0 , and L_∞ by the homotopy depicted, so that $i(x) = 0$ where x is the generator for $CF(L_0, L_\infty)$. We also orient these curves in the respective directions of α and β (so, if ℓ is oriented in the direction of β , the orientations are those obtained from the

orientation of ℓ and the grading of the curve). It is permissible to choose a different line field ℓ' , so long as it is defined on all of T , twisting say, a times along L_∞ and b times along L_0 ; but in that case we put the same number of extra twists into the gradings of L_0 and L_∞ , so that one still has $i(x) = 0$. We take the spin-structures on L_0 and L_∞ to be the *non-trivial* double covers. A convenient way to keep track of double coverings of curves L is to mark a point $\star_L \in L$, and declare the double cover to be trivial on $L \setminus \{\star_L\}$ and to exchange the sheets over \star_L . Such stars appear in Figure 1.

Let \mathcal{A} be the full A_∞ sub-category of $\mathcal{F}(T_0)^{\text{ex}}$ with objects $(L_0^\#, L_\infty^\#)$. Every object of $\mathcal{F}(T_0)^{\text{ex}}$ whose double covering is non-trivial is quasi-isomorphic, in $\text{tw } \mathcal{F}(T_0)^{\text{ex}}$, to a twisted complex in \mathcal{A} . This is because such objects are given, up to possible reversal of orientation and shift in degree, by iterated Dehn twists along $L_0^\#$ and $L_\infty^\#$ and the effect of the Dehn twists $\tau_{L_0^\#}^{\pm 1}$ and $\tau_{L_\infty^\#}^{\pm 1}$ on the Fukaya category are given by the twist functors along the spherical objects $L_0^\#$ and $L_\infty^\#$, provided that the double covering of $L_0^\#$ and $L_\infty^\#$ are non-trivial. For a systematic account, including the double-cover condition, see [50, Theorem 17.16]; a low-tech account of the case at hand is given in [35].

Curves with trivial spin structures are not quasi-isomorphic in $\text{tw } \mathcal{F}(T_0)^{\text{ex}}$ to objects of $\text{tw } \mathcal{A}$; in fact, we shall see later that these objects do not represent classes in the subgroup $K_0(\text{tw } \mathcal{A})$ of $K_0(\text{tw } \mathcal{F}(T_0)^{\text{ex}})$. What is true, however, is that for any object X of $\mathcal{F}(T_0)^{\text{ex}}$, the direct sum $X \oplus X[2]$ is quasi-isomorphic, in $\text{tw } \mathcal{F}(T_0)^{\text{ex}}$, to a twisted complex in \mathcal{A} . One can write down an explicit twisted complex representing $X \oplus X[2]$ using [50, Cor. 5.8, Theorem 17.16, and formula 19.4], bearing in mind that there is an exact isotopy $(\tau_{L_0} \tau_{L_\infty})^6 \simeq \tau_{c_z}$ where τ_γ is a symplectic Dehn twist along γ . The isotopy takes place in a compact domain containing the surface bounded by c_z , and is trivial near the boundary. The relation arises from the monodromy of an anticanonical Lefschetz pencil on $\mathbb{C}P^2$ with 8 of the 9 base-points blown up. Consequently, \mathcal{A} *split-generates* $\mathcal{F}(T_0)^{\text{ex}}$, i.e., $\text{tw}^\pi \mathcal{A} \rightarrow \text{tw}^\pi \mathcal{F}(T_0)^{\text{ex}}$ is a quasi-equivalence.

6.2 The wrapped category

The *wrapped* Fukaya A_∞ -category $\mathcal{W}(M)$ of a Liouville manifold (M, θ) is set up in [5]. It is a \mathbb{Z} -linear A_∞ -category containing $\mathcal{F}(M)^{\text{ex}}$ as a full subcategory. Its objects are again certain Lagrangian branes $L^\#$. Precisely, L must be a properly embedded, eventually conical Lagrangian submanifold, exact in the strong sense that $\theta|_L = dK$ for some $K \in C_c^\infty(L)$. The brane structure consists of a spin structure and a grading on L , and both can be taken to be eventually translation-invariant on the conical end. The morphism spaces $\text{hom}_{\mathcal{W}(M)}(L^\#, L'^\#)$ are cochain complexes $CW^*(L^\#, L'^\#)$ computing wrapped Floer cohomology $HW^*(L^\#, L'^\#)$, which is Lagrangian Floer cohomology $HF^*(\phi(L^\#), L'^\#)$ for a Hamiltonian diffeomorphism ϕ which ‘accelerates’ on the conical

end. That is, ϕ acts on the conical end $N \times \mathbb{R}_+$ as $\phi(x, r) = (\phi_{f(r)}^{\text{Reeb}}(x), r)$, where $\{\phi_t^{\text{Reeb}}\}$ is the time t Reeb flow on the contact cross-section N and $f(r)$ is a function which increases rapidly in a precise sense. As a consequence ϕ ‘wraps’ L around the end many times, typically producing an infinity of intersections with L' .

An object in $\mathcal{W}(T_0)$ is either an object of $\mathcal{F}(T_0)$, or an eventually-straight oriented arc A , with a spin-structure, of which there is only one isomorphism class per orientation, and a grading. A grading is, again, a way to rotate from $\ell|_\alpha$ to TA inside $T(T_0)|_A$.

The oriented arc A has an initial segment $\{a_{\text{in}}\} \times [r, \infty)$ and a final segment $\{a_{\text{out}}\} \times [r, \infty)$, where a_{in} and a_{out} are points on the circle c_z . Exact arcs with fixed a_{in} and a_{out} are Liouville-isotopic if and only if they have the same slope, i.e., represent the same primitive class in $H_1(T, \{z\})$. Exact arcs with the same slope but different endpoints are not Liouville-isotopic, but are nonetheless quasi-isomorphic in $\mathcal{W}(T_0)$.

If $L^\#$ is an object of $\mathcal{W}(T_0)$ such that L is closed, and $\Lambda^\#$ any object of $\mathcal{W}(T_0)$, one can apply a Liouville isotopy to L so as to reduce the number of intersections to the minimal intersection number $i(L, \Lambda)$. There are then no bigons, so $HW^*(L^\#, \Lambda^\#) \cong \mathbb{Z}^{i(L, \Lambda)}$. In particular, if A is an arc of the same slope as the closed curve L , one has $HW^*(A^\#, L^\#) = 0$.

One can show that $HW^*(A^\#, A^\#)$ is the non-commutative (tensor) algebra $T(u, v)$ on generators of degree 1 modulo the two-sided ideal (u^2, v^2) , but we shall not need to use this assertion.

Generation. Let A be any non-compact exact arc, equipped with a brane structure so as to make it an object of $\mathcal{W}(T_0)$. Let L_0 be a simple closed curve of the same slope as A , and L_∞ a curve which intersects L_0 transversely at a single point. Equip these three curves with brane structures, with the double-coverings of the two closed curves both non-trivial. Then $\{A^\#, L_0^\#, L_\infty^\#\}$ split-generates $\mathcal{W}(T_0)$. Indeed, we have already seen that $\{L_0^\#, L_\infty^\#\}$ split-generates $\mathcal{F}(T_0)^{\text{ex}}$, so we need only consider the arcs. Any arc can be obtained from A by a sequence of Dehn twists along L_0 and L_∞ , and so by an easy adaptation of [50, Theorem 17.16], or by a much more elementary argument which applies to the surface case, can be represented as a twisted complex in $A^\#, L_0^\#$ and $L_\infty^\#$.

6.3 The relative Fukaya category

The relative Fukaya category $\mathcal{F}(T, z)$ (cf. [46, 47, 54]) has the same objects as $\mathcal{F}(T_0)^{\text{ex}}$; and

$$\text{hom}_{\mathcal{F}(T, z)}(L^\#, L'^\#) = \text{hom}_{\mathcal{F}(T_0)^{\text{ex}}}(L^\#, L'^\#) \otimes_{\mathbb{Z}} \mathbb{Z} \llbracket q \rrbracket,$$

i.e., $\text{hom}_{\mathcal{F}(T,z)}(L^\#, L'^\#)$ is a free $\mathbb{Z}[[q]]$ -module on the intersections $L \cap L'$. The line field ℓ on T defines a grading, and hence makes the hom-spaces graded modules.

The A_∞ -structure $\{\mu^d\}$ is defined through inhomogeneous pseudo-holomorphic polygons, now in T . Such polygons u count with a weight $\varepsilon(u)q^{u \cdot z}$, where $\varepsilon(u)$ is a sign, defined just as in $\mathcal{F}(T_0)^{\text{ex}}$, and $u \cdot z$ is the intersection number with z . A formula for $\varepsilon(u)$ is given in [49] (see also [35]). We shall give here only a special case in which all corners of the polygon (i.e., $y_{k+1} \in \text{hom}(L_k, L_{k+1})$ for $k = 0, \dots, d-1$ and $y_0 \in \text{hom}(L_d, L_0)$) have *even index* $i(y_k)$. In that case, $\varepsilon = (-1)^s$, where s is the number of stars on the boundary (recall that the stars designate non-trivial monodromies for the double covers).

The same perturbations that define the A_∞ -structure in $\mathcal{F}(T_0)^{\text{ex}}$ succeed in defining an A_∞ -structure here too. The proof uses automatic regularity for holomorphic maps to surfaces [50, (13a)], and is otherwise unchanged from the proof in the exact case. The resulting A_∞ -structure is an invariant of $(T, \omega, z; \theta)$. Moreover, up to quasi-equivalence it is independent of θ ; the proofs of these assertions are straightforward adaptations of their analogues for T_0 given in [50].

Generation. The objects $L_0^\#$ and $L_\infty^\#$ split-generate $\mathcal{F}(T, z)$. The proof is essentially the same as for $\mathcal{F}(T_0)$. One can see curves with non-trivial double coverings as explicit twisted complexes in $L_0^\#$ and $L_\infty^\#$, cf. [35]. For arbitrary objects X , the sum $X \oplus X[2]$ is again a twisted complex in $L_0^\#$ and $L_\infty^\#$; the proof uses the main results of [50] just as before, via the relation $(\tau_{L_0} \circ \tau_{L_\infty})^6 \simeq \text{id}[2]$ in the graded symplectic mapping class group. It is significant here that $c_1(T) = 0$, and more particularly that ℓ extends over T , since Seidel's argument depends on the presence of absolute gradings; with that point noted, the argument applies to (T, z) as it does to T_0 .

6.4 The closed-open string map

Besides the ‘open string’ invariants $\mathcal{F}(M)^{\text{ex}}$ and $\mathcal{W}(M)$, Liouville manifolds M have a ‘closed string’ invariant, the symplectic cohomology algebra $SH^\bullet(M)$ (see [48] for an exposition and foundational references). In the first place, $SH^\bullet(M)$ is a graded-commutative graded ring; the grading depends on a choice of grading for M as a symplectic manifold. It also comes with a ring map $v: H^\bullet(M) \rightarrow SH^\bullet(M)$, which in our grading convention is a map of graded rings; this pins down our normalization for the grading of $SH^\bullet(M)$, which in some other accounts (such as [48]) differs from ours by $\dim_{\mathbb{C}} M$. As a simple algebraic variant, we can work with $SH^\bullet(M; \mathbb{K})$, an algebra over the commutative ring \mathbb{K} .

Lemma 6.3 *For any commutative ring \mathbb{K} , there is an isomorphism of graded \mathbb{K} -modules*

$$\theta: \mathbb{K}[\beta, \gamma_1, \gamma_2]/(\gamma_1\gamma_2, \beta(\gamma_1 - \gamma_2)) \rightarrow SH^\bullet(T_0; \mathbb{K}),$$

canonical after a choice of basis of $H_1(T_0)$, where $\deg \gamma_1 = \deg \gamma_2 = 1$ and $\deg \beta = 2$. Here the variables γ_1, γ_2 and β are super-commutative. Moreover, $\theta(\gamma_1) \cdot \theta(\gamma_2) = 0$.

When \mathbb{K} is a field, the algebra on the right is isomorphic to $HH^\bullet(\mathcal{F}_\mathbb{K}(\mathcal{T}_0)^{\text{ex}})$ by Theorem 4.20. We will see presently that θ respects products.

Proof We begin with a comment about grading. The grading of $SH(T_0; \mathbb{K})$ depends, *a priori*, on the choice of line field ℓ . If $\ell' = h(\ell)$ is another choice, obtained from ℓ by a map $g: T_0 \rightarrow \mathbb{R}P^1$ representing a class $c = g^*(o) \in H^1(T_0; \mathbb{Z})$, where $H^1(\mathbb{R}P^1; \mathbb{Z}) = \mathbb{Z}o$, the degrees of generators, which are 1-periodic Hamiltonian orbits y , change according to the formula $|y|_{\ell'} = |y|_\ell \pm 2\langle c, [y] \rangle$ (we do not bother with the sign). In the cochain complex for $SH(T_0)$ we shall describe, all orbits are null-homologous, and hence the grading of $SH(T_0)$ is independent of ℓ .

The graded \mathbb{K} -module $SH(T_0; \mathbb{K})$ is described in [48, ex. 3.3]. One uses an autonomous Hamiltonian which is a perfect Morse function h , accelerating appropriately on the cylindrical end. One has a natural map of algebras $v: H^\bullet(T_0; \mathbb{K}) \rightarrow SH^\bullet(T_0; \mathbb{K})$ which is an isomorphism onto $SH^{\leq 1}(T_0; \mathbb{K})$; this is the contribution of the minimum m and two saddle-points s_1 and s_2 of h . We choose h so that $([s_1], [s_2])$ is dual to the chosen basis for the Morse cohomology $H_1(T_0; \mathbb{K})$. Define $\theta(1) = 1$, $\theta(\gamma_1) = s_1$ and $\theta(\gamma_2) = s_2$.

For each $q \geq 1$, there is a Reeb orbit o_q which winds q times around the puncture; $H^*(o_q; \mathbb{K})$ contributes classes $c_{2q} \in SH^{2q}(T_0; \mathbb{K})$ and $c_{2q+1} \in SH^{2q+1}(T_0; \mathbb{K})$ which span those respective modules. We define $\theta(\beta^q) = c_{2q}$ and $\theta(\gamma_1\beta^q) = c_{2q+1}$. We have $[s_1] \cdot [s_2] = 0$. Indeed, $[s_1] \cdot [s_2]$, having degree 2, is a multiple of $[o_1]$, but since o_1 is not contractible, there is no pair-of-pants that bounds o_1 and the constant loops s_1 and s_2 . \square

The exact Fukaya category of a Liouville manifold M is tied to its symplectic cohomology via the ‘closed-open string map’ [46] to its Hochschild cohomology

$$(38) \quad \mathcal{CO}: SH^\bullet(M; \mathbb{K}) \rightarrow HH^\bullet(\mathcal{F}(M)_\mathbb{K}^{\text{ex}}).$$

For the details of the construction of \mathcal{CO} we refer to S. Ganatra’s doctoral thesis [21].

Theorem 6.4 (see [21]) *\mathcal{CO} is a homomorphism of \mathbb{K} -algebras: it intertwines the pair-of-pants product on SH^\bullet with the cup product on HH^\bullet .*

7 The punctured torus

This part of the paper pinpoints the Weierstrass curve $C_{\text{mirror}} \rightarrow \text{Spec } \mathbb{Z} \llbracket q \rrbracket$ such that the minimal A_∞ -structure $\mathcal{A}_{\text{mirror}}$ it induces on the algebra A is gauge-equivalent to the A_∞ -structure $\mathcal{A}_{\text{symp}}$ obtained from the Fukaya category $\mathcal{F}(T, z)$. Theorem C assures us that C_{mirror} exists, and is unique as an abstract Weierstrass curve. Our aim is to show that $C_{\text{mirror}} \cong \mathcal{T}$. This will be accomplished in the final section of the paper by an argument involving θ -functions. In this section we offer two alternative proofs that $C_{\text{mirror}}|_{q=0} \cong \mathcal{T}_0$. One is by eliminating all possibilities other than \mathcal{T}_0 ; the other is by a calculation of ‘Seidel’s mirror map’. We also prove our mirror-symmetry theorem for the wrapped category.

7.1 First proof that $C_{\text{mirror}}|_{q=0} = \mathcal{T}_0$: by elimination

Lemma 1.4 characterized the Weierstrass curve $\mathcal{T}_0 \rightarrow \text{Spec } \mathbb{Z}$ as having a section which is a node at p , for any $p \in \text{Spec } \mathbb{Z}$. Symplectic topology now enters the picture:

Proposition 7.1 *Take $p \in \text{Spec } \mathbb{Z}$, and $C \rightarrow \text{Spec } \mathbb{F}_p$ a Weierstrass curve (here $\mathbb{F}_0 = \mathbb{Q}$). Suppose that the exact category $\text{tw}^\pi \mathcal{F}(T_0)^{\text{ex}}$, taken with coefficients in \mathbb{F}_p , is quasi-equivalent to $\text{tw vect } C$. Then C is nodal.*

Proof Let $\mathbb{K} = \mathbb{F}_p$. One has Viterbo’s map $v: H^\bullet(T_0; \mathbb{K}) \rightarrow SH^\bullet(T_0; \mathbb{K})$, a graded algebra homomorphism [61, 48], and restriction homomorphisms

$$HH^\bullet(\mathcal{F}(T_0)^{\text{ex}}) \rightarrow HH^\bullet(\text{hom}_{\mathcal{F}(T_0)^{\text{ex}}}(L^\#, L^\#)),$$

(everything over \mathbb{K}), one for each Lagrangian brane $L^\# \in \text{ob } \mathcal{F}(T_0)^{\text{ex}}$. The Hochschild cochain complex C^\bullet for an A_∞ -algebra [50] has a filtration $F^r C^\bullet$ by the length of cochains (so $C^\bullet = F^0 C^\bullet \supset F^1 C^\bullet \supset \dots$), and so for each L one has a quotient map

$$HH^\bullet(\text{hom}(L^\#, L^\#)) \rightarrow H(C^\bullet / F^1 C^\bullet) = H^\bullet(\text{hom}(L^\#, L^\#)) = HF(L^\#, L^\#).$$

Now, $HF(L^\#, L^\#) \cong H^\bullet(L; \mathbb{K})$, ordinary cohomology, by a canonical isomorphism [50, (8c)]. It follows easily from the definitions, plus the gluing theorem for Hamiltonian Floer theory—of which [45] has a meticulous account—that the composite of the maps

$$H^\bullet(T_0; \mathbb{K}) \xrightarrow{v} SH^\bullet(T_0; \mathbb{K}) \xrightarrow{\text{co}} HH^\bullet(\mathcal{F}(T_0)^{\text{ex}}) \rightarrow HH^\bullet(\text{hom}(L^\#, L^\#)) \rightarrow H^\bullet(L; \mathbb{K})$$

is the classical restriction map $H^\bullet(T_0; \mathbb{K}) \rightarrow H^\bullet(L; \mathbb{K})$ —a *surjective* map. Taking the sum of these composite maps for the objects $L_0^\#$ and $L_\infty^\#$ produces an *isomorphism*

$$(39) \quad H^1(T_0; \mathbb{K}) \rightarrow H^1(L_0; \mathbb{K}) \oplus H^1(L_\infty; \mathbb{K}).$$

On the algebro-geometric side, suppose that $C = C_{\text{mirror}} \rightarrow \text{Spec } \mathbb{K}$ is a Weierstrass curve mirror to T_0 over \mathbb{K} . By Lemma 4.3, one has $\text{HH}^\bullet(\text{tw vect } C) \cong \text{HH}^\bullet(C)$. Hochschild cohomology is invariant under passing to tw^π , by a form of Morita invariance [30, 60]. Hence, under the hypotheses of the proposition, one has $\text{HH}^\bullet(\mathcal{F}(T_0)^{\text{ex}}) \cong \text{HH}^\bullet(C)$. Consequently, using (39) we obtain a map of \mathbb{K} -algebras $SH^\bullet(T_0; \mathbb{K}) \rightarrow \text{HH}^\bullet(C)$ such that the composite

$$H^1(T_0; \mathbb{K}) \rightarrow SH^1(T_0; \mathbb{K}) \rightarrow \text{HH}^1(C) \rightarrow \text{Ext}^1(\mathcal{O}_C, \mathcal{O}_C) \oplus \text{Ext}^1(\mathcal{O}_{C,\sigma}, \mathcal{O}_{C,\sigma})$$

is an isomorphism of \mathbb{K} -modules. We assert that such a homomorphism exists only if C is nodal. To prove this, we must eliminate the cuspidal and smooth cases.

We claim that if C were smooth, one would have a \mathbb{K} -algebra isomorphism $\text{HH}^\bullet(C) \cong \Lambda^\bullet[\alpha_1, \alpha_2]$, where $\deg \alpha_1 = 1 = \deg \alpha_2$. Additively, this follows from the degeneration of the Hodge spectral sequence (28). The spectral sequence is multiplicative, and so the \mathbb{K} -algebra $\bigoplus_{p,q} H^p(\Lambda^q \mathcal{T}_C)$ is isomorphic to the associated graded algebra $E^\infty = \text{gr } \text{HH}^\bullet(C)$. We have $\alpha_1 \cdot \alpha_2 \neq 0 \in \text{HH}^2(C)$, since this is even true in the associated graded algebra. This is sufficient to establish the claim.

Any homomorphism of graded unital \mathbb{K} -algebras

$$\theta: SH^\bullet(T_0) \rightarrow \mathbb{K}[\alpha_1, \alpha_2]$$

obeys, in the notation of Lemma 6.3, $\theta(\gamma_1)\theta(\gamma_2) = 0$, and therefore θ fails to surject onto the 1-dimensional part $\{\alpha_1, \alpha_2\}$. The composite of θ with the maps to $\text{Ext}^\bullet(\mathcal{O}_C, \mathcal{O}_C)$ and $\text{Ext}^\bullet(\mathcal{O}_{C,\sigma}, \mathcal{O}_{C,\sigma})$ cannot then both be surjective.

We can rule out the possibility that C is cuspidal, i.e., that $C \cong C_{\text{cusp}}$, by noting that $\text{Ext}^1(\mathcal{O}_{C_{\text{cusp}}}, \mathcal{O}_{C_{\text{cusp}}}) \cong H^1(\mathcal{O}_{C_{\text{cusp}}}) = \mathbb{K} \cdot \omega$. This module transforms under the \mathbb{K}^\times -action on C_{cusp} , and it has weight $+1$. The restriction map $\text{HH}^1(C_{\text{cusp}}) \rightarrow \text{Ext}^1(\mathcal{O}_{C_{\text{cusp}}}, \mathcal{O}_{C_{\text{cusp}}})$ respects the weight of the \mathbb{K}^\times -action—that is, it is a map of graded vector spaces—and hence its restriction to $\text{HH}^1(C_{\text{cusp}})^{\leq 0}$ is zero. \square

7.1.1 Variants

There are other methods for ruling out C_{cusp} :

- When $6 \neq 0$, the map that $\text{HH}^1(C_{\text{cusp}})^{\leq 0}$ is 1-dimensional, and therefore the composite map $H^\bullet(T_0; \mathbb{K}) \rightarrow \text{Ext}^1(\mathcal{O}, \mathcal{O}) \oplus \text{Ext}^1(\mathcal{O}_\sigma, \mathcal{O}_\sigma)$ unavoidably has a kernel.
- In [35], we used Abouzaid’s model [3] for the Fukaya category of a plumbing to describe the structure maps of \mathcal{A} and thereby prove non-formality, assuming $6 \neq 0$.

- When $6 = 0$, the Gerstenhaber bracket on $\mathrm{HH}^1(C_{\mathrm{cusp}})^{\leq 0}$ is non-zero by Theorem 5.6. For any Liouville domain M , the bracket is zero on the image of Viterbo's map $v: H^\bullet(M) \rightarrow SH^\bullet(M)$. Indeed, v is a ring homomorphism, and it is easy to see that $\Delta \circ v = 0$, where Δ is the BV operator on symplectic cohomology [48], and the bracket is, in accordance with the rules of BV algebras,

$$[x, y] = (-1)^{|x|} \Delta(x \cdot y) - x \cdot \Delta y - (-1)^{|x|} (\Delta x) \cdot y,$$

which implies that $[v(a), v(b)] = 0$. Moreover, \mathcal{CO} preserves Gerstenhaber brackets: this was first stated by Seidel [46], but there is no published proof. In characteristic 2 we have verified it for ourselves using standard gluing methods. Over other fields the signs are tricky, so we regard it as conjectural. However, taking this assertion for granted, the composite map $H^1(T_0; \mathbb{K}) \rightarrow \mathrm{HH}^\bullet(C_{\mathrm{cusp}})^{\leq 0}$ preserves brackets and so cannot be injective.

Proof of Theorem A clause (iii). We want to construct an A_∞ -functor

$$\psi: \mathcal{F}(T_0) \rightarrow \mathrm{tw} \, \mathrm{vect}(\mathcal{T}|_{q=0}).$$

We already have an isomorphism $H\psi: A = H^* \mathcal{A} \rightarrow H^* \mathcal{B}_C$, valid for any Weierstrass curve C defined over \mathbb{Z} . Theorem C implies that there is a unique C for which $H\psi$ lifts to a quasi-isomorphism $\psi: \mathcal{A} \rightarrow \mathcal{B}_C$. For that particular C , ψ extends naturally to a quasi-isomorphism

$$\mathrm{tw}^\pi \mathcal{A} \rightarrow \mathrm{tw}^\pi \mathcal{B}_C.$$

Since the inclusion maps $\mathrm{tw}^\pi \mathcal{A} \rightarrow \mathrm{tw}^\pi \mathcal{F}(T_0)^{\mathrm{ex}}$ and $\mathrm{tw}^\pi \mathcal{B}_C \rightarrow \mathrm{tw}^\pi \mathrm{vect}(C)$ are quasi-equivalences, one obtains a quasi-isomorphism

$$\mathrm{tw}^\pi \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathrm{tw}^\pi \mathrm{vect}(C).$$

Composing this with a quasi-inverse to the inclusion $\mathrm{tw} \, \mathrm{vect}(C) \rightarrow \mathrm{tw}^\pi \mathrm{vect}(C)$, one obtains a quasi-isomorphism

$$\mathrm{tw}^\pi \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathrm{tw} \, \mathrm{vect}(C)$$

whose restriction to $\mathcal{F}(T_0)$ is the functor we want.

Our task, then, is to identify the mirror Weierstrass curve C . In light of Prop. 7.1, it must be a curve which is nodal over \mathbb{Q} . *A priori*, the node is only a $\overline{\mathbb{Q}}$ -point. However, the normalization of C , defined over \mathbb{Q} , has two points (over $\overline{\mathbb{Q}}$) which map to the node, and since the normalization is a rational curve, these points are actually defined over \mathbb{Q} . Hence the same is true of the node. By clearing denominators, we obtain integer coordinates for the node. It then defines a section of $C \rightarrow \mathrm{Spec} \, \mathbb{Z}$, which by the proposition must map to a node over \mathbb{F}_p for every prime p . By Lemma 1.4, C is therefore equivalent to $\mathcal{T}|_{q=0}$.

We have proved a weakened form of Theorem A clause (iii): we have shown that $\mathrm{tw}^\pi \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathrm{tw} \mathrm{vect}(C)$ is a quasi-equivalence, while the theorem claims that $\mathrm{tw} \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathrm{tw} \mathrm{vect}(C)$ is already an equivalence. We formulate this step as a separate statement, Prop. 7.2, whose proof completes that of the theorem. \square

Proposition 7.2 *$\mathrm{tw} \psi: \mathrm{tw} \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathrm{Perf}(C)$ is a quasi-equivalence.*

The triangulated A_∞ -category $\mathrm{tw} \mathrm{vect}(\mathcal{T}_0)$ is split-closed. Let $\mathcal{J} \subset \mathrm{tw} \mathrm{vect}(\mathcal{T}_0)$ denote the image of $\mathrm{tw} \mathcal{F}(T_0)^{\mathrm{ex}}$ under $\mathrm{tw} \psi$. We must show that the inclusion $\mathcal{J} \rightarrow \mathrm{tw} \mathrm{vect}(\mathcal{T}_0)$ is a quasi-equivalence. By Thomason’s theorem [58], it is sufficient to prove equality of Grothendieck groups: $K_0(\mathcal{J}) = K_0(\mathrm{tw} \mathrm{vect}(\mathcal{T}_0))$.⁵

Lemma 7.3 *Consider the map $s: \mathrm{ob} \mathcal{F}(T_0)^{\mathrm{ex}} \rightarrow \mathbb{Z}/2$ which maps an exact Lagrangian brane L to 0 if and only if the double covering $\tilde{L} \rightarrow L$ (part of the brane structure) is trivial. This map descends to a homomorphism $s: K_0(\mathrm{tw} \mathcal{F}(T_0)^{\mathrm{ex}}) \rightarrow \mathbb{Z}/2$.*

Proof This is an adaptation of a result of Abouzaid [1, Prop. 6.1]. One views a spin-structure, lifting a given orientation, as a local system for the group $\{\pm 1\} \subset U(1)$. The switch from closed, higher-genus surfaces to T_0 is irrelevant. \square

Completion of the proof of 7.2. We must show that

$$K_0(\psi_0): K_0(\mathrm{tw} \mathcal{F}(T)^{\mathrm{ex}}) \rightarrow K_0(\mathrm{tw} \mathrm{vect} \mathcal{T}_0)$$

is onto. We have $K_0(\mathrm{tw} \mathrm{vect} \mathcal{T}_0) = K_0(\mathcal{T}_0)$, and by Lemma 3.6, $(\mathrm{rank}, \det): K_0(\mathcal{T}_0) \rightarrow \mathbb{Z} \oplus \mathrm{Pic}(\mathcal{T}_0)$ is an isomorphism. One also has an isomorphism $(\mathrm{deg}, \rho): \mathrm{Pic}(\mathcal{T}_0) \rightarrow \mathbb{Z} \oplus \mathbb{Z}^\times$, where deg is the degree and ρ describes the ‘descent data’ under normalization, as at the end of the proof of that lemma. The image of ψ_0 contains \mathcal{O} (rank 1, degree 0) and $\mathcal{O}(\sigma)$ (rank 1, degree 1). Hence $(\mathrm{rank}, \mathrm{deg}) \circ K_0(\psi_0): K_0(\mathrm{tw} \mathcal{F}(T)^{\mathrm{ex}}) \rightarrow \mathbb{Z}^2$ is onto. Now take some $L^\# \in \mathrm{ob} \mathcal{F}(T_0)^{\mathrm{ex}}$, and let L^* be the same object with the other double covering. Then the class $[L^\#] - [L^*]$ is 2-torsion in $K_0(\mathrm{tw} \mathcal{F}(T)^{\mathrm{ex}})$, by the last lemma and the fact that the change of covering can be accomplished by an involution of $\mathrm{tw} \mathcal{F}(T_0)^{\mathrm{ex}}$ (namely, tensoring spin-structures on Lagrangians by the restrictions of some real line-bundle $\ell \rightarrow T_0$). Thus $[\psi_0(L)] - [\psi_0(L^*)]$ is again 2-torsion in $K_0(\mathrm{tw} \mathrm{vect} \mathcal{T}_0)$. There is a unique 2-torsion class in $K_0(\mathrm{tw} \mathrm{vect} \mathcal{T}_0)$, detected by ρ . Therefore $\rho \circ K_0(\psi_0)$ is surjective, and hence $K_0(\psi_0)$ is surjective. \square

⁵By K_0 of a triangulated A_∞ -category \mathcal{C} , we mean $K_0(H^* \mathcal{C})$, the Grothendieck group of the classical triangulated category $H^* \mathcal{C}$ —in this case, $K_0(\mathrm{Perf} \mathcal{T}_0)$.

Remark. There is an alternative to the argument just given which does not appeal to Abouzaid’s analysis, but instead observes that the objects $L_0^\#$ and L_0^* (which differ only in their double coverings) map under ψ to perfect complexes whose K_0 -classes differ by the generator of $\mathbb{Z}/2 \in K_0(\mathcal{T})$. For this, we regard L_0^* as $L_0^\#$ with a local system with fiber \mathbb{Z} and holonomy -1 . We have $\psi(L_0^\#) = \mathcal{O}$, and it follows from an easy adaptation of Lemma 8.1 below that $\psi(L_0^*) = \mathcal{O}(\sigma - \sigma')$, where σ' is the 2-torsion section of $\mathcal{T}_0^{\text{sm}} = \mathbb{G}_m(\mathbb{Z}) \rightarrow \text{Spec } \mathbb{Z}$ and, as usual, σ is the identity section. From this point, the argument is straightforward.

7.2 A second identification of the central fiber of the mirror curve: Seidel’s mirror map

The affine coordinate ring. Suppose given an abstract Weierstrass curve (C, σ, Ω) over $\text{Spec } S$. There is then a Weierstrass cubic embedding carrying σ to $[0 : 1 : 0]$; the affine complement to the closure of $\text{im } \sigma$ is $\text{Spec } R_C$, where R_C , the affine coordinate ring, is the ring of functions on C with poles only at σ :

$$R_C = \varinjlim_n H^0(C, \mathcal{O}(n\sigma)).$$

7.2.1 The Dehn twist ring.

We want to compute $R_{C_{\text{mirror}}}$ by determining an equation that holds in $H^0(C_{\text{mirror}}, \mathcal{O}(6\sigma))$. (This is almost the method of [65], but we are concerned here with the affine rather than homogeneous coordinate ring.) We shall in fact carry out this computation only for $C_{\text{mirror}}|_{q=0}$.

Since \mathcal{A} split-generates $\mathcal{F}(T, z)$, the quasi-isomorphism $\mathcal{A} \rightarrow \mathcal{B}_{C_{\text{mirror}}}$ extends to an A_∞ -functor $\psi: \mathcal{F}(T, z) \rightarrow \text{tw vect}(C_{\text{mirror}})$. Consider the object $L_\infty^\#$. The Dehn twist $\tau = \tau_{L_\infty^\#}$, acting as an autoequivalence of $\mathcal{F}(T_0)$, is homotopic to the twist functor along the spherical object $L_\infty^\#$: this is elementary in the present case [35], but is an instance of a general result of Seidel’s [50]. Now, $\psi(L_\infty^\#) = \mathcal{O}_\sigma$. The twist along the spherical object $\mathcal{O}_\sigma \in \text{ob tw vect}(C_{\text{mirror}})$ is homotopic to the functor $\mathcal{O}(\sigma) \otimes \cdot$ of tensoring with $\mathcal{O}(\sigma)$ (see [51, (3.11)]; the argument is carried out over fields, but over $\mathbb{Z}[[q]]$, Seidel–Thomas’s map f must be a unit times the restriction map by base-changing to reduce to the case of fields). Thus ψ induces an isomorphism $R_{C_{\text{mirror}}} \cong R_\tau$, where

$$R_\tau := \varinjlim_n HF^*(L_0^\#, \tau^n(L_0^\#)).$$

The ‘Dehn twist ring’ R_τ , needs explanation—neither the direct system, nor the ring structure, is obvious. The maps

$$\sigma_{n,m+n}: HF^*(L_0^\#, \tau^n(L_0^\#)) \rightarrow HF^*(L_0^\#, \tau^{m+n}(L_0^\#))$$

which form the direct system are defined via holomorphic sections of a Lefschetz fibration over a strip; this interpretation is part of Seidel’s analysis of Dehn twists [50]. The ring structure is easier: the m th power of the Dehn twist defines a map

$$(\tau_*)^m: HF^*(L_0^\#, \tau^n(L_0^\#)) \rightarrow HF^*(\tau^m(L_0^\#), \tau^{m+n}(L_0^\#))$$

which applies the Dehn twist to the intersection points between Lagrangians. The product in the ring is given by composing this with the triangle product \cdot ,

$$\begin{aligned} HF^*(L_0^\#, \tau^n(L_0^\#)) \otimes HF^*(L_0^\#, \tau^m(L_0^\#)) &\xrightarrow{\tau_*^m \otimes \text{id}} HF^*(\tau^m(L_0^\#), \tau^{m+n}(L_0^\#)) \otimes HF^*(L_0^\#, \tau^m(L_0^\#)) \\ &\xrightarrow{\cdot} HF^*(L_0^\#, \tau^{m+n}(L_0^\#)). \end{aligned}$$

Associativity of this product is easily seen, as is the fact that the unit element $e \in HF^*(L_0^\#, L_0^\#)$ is a 2-sided unit for the multiplication in the Dehn twist ring. However, commutativity is something that we learn from the isomorphism $R_\tau \cong R_{C_{\text{mirror}}}$.

7.2.2 Avoiding the direct system

If one knew the maps $\sigma_{m,m+n}$ explicitly, one would be able to proceed by perfect analogy with the algebro-geometric side of the mirror, as follows. Take the unit $e \in HF^*(L_0^\#, L_0^\#)$ and its images $e_n = \sigma_{0,n}(e) \in HF^*(L_0^\#, \tau^n(L_0^\#))$. Take a basis $\{e_2, x\}$ for $HF^*(L_0^\#, \tau^2(L_0^\#))$ and a basis $\{e_2, \sigma_{2,3}x, y\}$ for $HF^*(L_0^\#, \tau^3(L_0^\#))$. Identify the Weierstrass equation as the unique relation satisfied by the seven monomials $\{y^2, x^3, xy e_1, x^2 e_2, y e_3, x e_4, e_6\}$.

We can extract nearly complete information about the maps $\sigma_{0,n}$ as follows. We choose L_0 and L_∞ to have just one, transverse intersection point. Then $L_0 \cap \tau(L_0)$ consists of a single point z' , and $HF^*(L_0^\#, \tau(L_0^\#)) = \mathbb{Z}[[q]]z'$. Hence $e_1 = f(q)z'$ for some $f(q) \in \mathbb{Z}[[q]]^\times$. From the isomorphism $R_\tau \cong R_{C_{\text{mirror}}}$, we see that $e_n \cdot e_m = e_{m+n}$, and hence that $e_n = (e_1)^n = f(q)^n z'^n$. Moreover, $\sigma_{m,m+1}(u) = e_1 u = f(q)z' \cdot u$. In practice, then, one obtains a cubic equation by picking x' so as to make $\{z'^2, x'\}$ a basis for $HF^*(L_0^\#, \tau^2(L_0^\#))$ and y' so as to make $\{z'^3, z'x', y'\}$ a basis for $HF^*(L_0^\#, \tau^3(L_0^\#))$. One computes the products

$$\{y'^2, x'^3, x'y'z', x'^2 z'^2, y'z'^3, x'z'^4, z'^6\}$$

and identifies the unique (up to scale) relation that they satisfy. This relation is necessarily of form

$$y'^2 - f(q)x'^3 = \dots,$$

hence it determines f . Now let $x = fx'$ and $y = fy'$. Then one has $y^2 - x^3 = \dots$, i.e., these coordinates satisfy the Weierstrass equation.

Remark. We learned something interesting *en route* here, though we shall not pursue it: the series f , which encodes information about sections of a Lefschetz fibration, can be computed.

7.2.3 A model for the Dehn twist

To compute with the Dehn twist ring, note that one can take for τ any compactly supported exact automorphism of T_0 that is isotopic to a Dehn twist along L_∞ . To obtain a convenient model, start with the linear symplectomorphism of $T = \mathbb{R}^2/\mathbb{Z}^2$ given by

$$\delta[x_1, x_2] = [x_1, x_2 - x_1].$$

The fixed-point set of δ is the line $L_\infty = \{x_1 = 0\}$. We will take our basepoint to be $z = (\epsilon, \epsilon)$ where $\epsilon \in (0, 1/4)$. One has $\delta(z) = (\epsilon, 0)$; let D be the ϵ^2 -neighborhood of the line segment $[\delta(z), z]$. Take $\rho \in \text{Aut}_c(D, \omega|_D)$ to be a symplectomorphism such that $\rho(\delta(z)) = z$; extend ρ to a symplectomorphism of T , still called ρ , which is trivial outside D . Let $\tau = \rho \circ \delta$. Then $\tau(z) = z$; by adjusting ρ , we may assume that τ acts as the identity in some neighborhood of z . Thus τ restricts to give $\tau_0 \in \text{Aut}_c(T_0, \omega|_{T_0})$. Let $L_{(1, -n)} = \{[x_1, x_2] \in T : nx_1 + x_2 = 0\}$. When $0 \leq n \leq (2\epsilon)^{-1}$, we have that $\tau_0^n(L_0) = L_{(1, -n)}$.

Lemma 7.4 *There is a primitive θ for $\omega|_{T_0}$ making L_0 an exact curve and τ_0 an exact symplectomorphism.*

Proof We must exhibit a primitive θ for ω such that $[\tau_0^*\theta - \theta] = 0 \in H_c^1(T_0; \mathbb{R})$ and $\int_{L_0} \theta = 0$. For the first requirement, it suffices to show that $\int_\gamma (\tau_0^*\theta - \theta) = 0$ for curves γ forming a basis for $H_1(T, \{z\}; \mathbb{R})$. Such a basis is given by $\{L_\infty, L_0\}$, with chosen orientations for these two curves. For any primitive θ for ω , one has $\int_{L_\infty} (\tau_0^*\theta - \theta) = \int_{\tau_0(L_\infty)} \theta - \int_{L_\infty} \theta = 0$. It suffices to choose θ , therefore, in such a way that L_0 and $\tau_0(L_0) = L_{(1, -1)}$ are both exact Lagrangians. It is easy to find a 1-form ι on T_0 such that $d\iota = \omega$ on a small regular neighborhood N of $L_0 \cup L_{(1, -1)}$ and such that $\int_{L_0} \iota = 0 = \int_{L_{(1, -1)}} \iota$. Now, ∂N is a circle, isotopic to a loop encircling z . There is therefore no obstruction to finding a 1-form $\kappa \in \Omega^1(T_0)$, supported outside $L_0 \cup L_{(1, -1)}$, such that $\omega = d(\iota + \kappa)$; then $\theta = \iota + \kappa$ is the required primitive. \square

Proposition 7.5 *For any natural number N , one can choose a primitive θ for $\omega|_{T_0}$ such that the curves $L_{(1, -n)}$ are exact for $n = 0, \dots, N$.*

Proof Choose $\epsilon < (2N)^{-1}$, and take θ as in the lemma. Since L_0 is exact, so too is $\tau_0^n(L_0)$ for any $n \in \mathbb{Z}$. But for $0 \leq n \leq N$ we have $\tau_0^n(L_0) = \delta^n(L_0) = L_{(1,-n)}$, \square

7.2.4 Computation in the exact case

Use a θ as in the proposition, taking N at least 6. We have

$$\begin{aligned} L_0 \cap L_{(1,-1)} &= \{z'\}, & z' &= [0, 0]; \\ L_0 \cap L_{(1,-2)} &= \{\zeta_0, \zeta_1\}, & \zeta_k &= [k/2, 0]; \\ L_0 \cap L_{(1,-3)} &= \{\eta_0, \eta_1, \eta_2\}, & \eta_k &= [k/3, 0]. \end{aligned}$$

In calculating products in the Dehn twist ring R_τ , immersed triangles in T_0 count with sign $+1$. To avoid repetition, we do not give the argument here but defer it to Section 10.4. With this understood, it is straightforward to calculate that in the Dehn twist ring of $\mathcal{F}(T_0)^{\text{ex}}$, one has

$$z'^2 = \zeta_0 + 2\zeta_1.$$

We put $x' = \zeta_1$; then $\{z'^2, x'\}$ is a \mathbb{Z} -basis for $HF^*(L_0^\#, L_{(1,-2)}^\#)$. Next, we compute

$$\begin{aligned} z'\zeta_0 &= \eta_0 + \eta_1 + \eta_2, \\ z'\zeta_1 &= \eta_1 + \eta_2. \end{aligned}$$

We deduce that

$$\begin{aligned} z'^3 &= z'(\zeta_0 + 2\zeta_1) = \eta_0 + 3\eta_1 + 3\eta_2, \\ z'x' &= \eta_1 + \eta_2. \end{aligned}$$

We put $y' = \eta_2$, and note that $\{z'^3, z'x', y'\}$ is a basis for $HF^*(L_0^\#, L_{(1,-3)}^\#)$. Further computations yield

$$\eta_2^2 = \theta_4, \quad \eta_1\eta_2 = \theta_3, \quad \zeta_1^3 = \theta_3$$

where $(\theta_0, \dots, \theta_5)$ are the intersection points $\theta_k = [k/6, 0] \in L_0 \cap L_{-6}$. These relations imply that

$$y'^2 + x'^3 = x'y'z'.$$

Putting $x = -x'$, $y = y'$ and $z = z'$, we obtain the Weierstrass relation in the desired form

$$y^2 - x^3 = -xyz.$$

8 Local systems

In this section (which lies outside the main development) we make the quasi-equivalence ψ_0 more explicitly by describing its effect on Lagrangians with finite-rank local systems. These may be regarded as twisted complexes in $\mathcal{F}(T_0)^{\text{ex}}$.

8.1 Rank 1 local systems

As usual, let $L_\infty^\# \in \text{ob } \mathcal{F}(T_0)$ be a Lagrangian brane of slope $(0, -1)$ with non-trivial double covering. For any $h \in \mathbb{C}^*$, let L_∞^h denote the brane L_∞ equipped with a rank 1 \mathbb{K} -local system of holonomy $h \in \mathbb{K}^\times$. This is an object in a larger Fukaya category $\mathcal{F}(T_0)_{\text{loc}}^{\text{ex}}$ whose objects are exact Lagrangian branes with local systems of finite rank free \mathbb{Z} -modules.

For a \mathbb{K} -linear A_∞ -category \mathcal{C} , let $\text{mod-}\mathcal{C}$ denote the category of finitely generated projective \mathcal{C} -modules, assigning to each object a finite cochain complex of finitely generated projective \mathbb{K} -modules. With $\mathbb{K} = \mathbb{Z}$, the object L_∞^h defines a left Yoneda-module $\mathcal{Y}_L(L_\infty^h) = \text{hom}_{\mathcal{F}(T_0)}(L_\infty^h, \cdot) \in \text{ob mod-}\mathcal{F}(T_0)$.

Let $\phi_0: \text{tw vect}(\mathcal{T}_0) \rightarrow \mathcal{F}(T_0)$ be an A_∞ functor quasi-inverse to ψ_0 . Over \mathbb{Z} , the existence of such a functor is not quite trivial. However, our earlier analysis of Čech complexes implies that one can define a quasi-inverse (or even strict inverse) $\mathcal{B}_{\mathcal{T}_0} \rightarrow \mathcal{A}$ to $\psi_0|_{\mathcal{A}}: \mathcal{A} \rightarrow \mathcal{B}_{\mathcal{T}_0}$. We then define ϕ_0 by extending the latter functor to twisted complexes. Module categories are contravariant, and so ϕ_0 induces a functor

$$\phi_0^*: \text{mod-}\mathcal{F}(T_0) \rightarrow \text{mod- vect}(\mathcal{T}_0).$$

Lemma 8.1 *Work over a base ring R which is a commutative, unital, normal, noetherian domain. Identify the normalization of \mathcal{T}_0 with $\mathbb{P}^1(R)$ by sending the preimages of the nodal section of \mathcal{T}_0 to $\{[0 : 1], [1 : 0]\}$ and σ to $[1 : 1]$. Two such identifications exist, of which one has the property that for each $h \in R^\times$, the module $P^h := \phi_0^*(L_\infty^h) \in \text{ob mod- vect}(\mathcal{T}_0)$ is represented by a locally free resolution of the skyscraper sheaf \mathcal{O}_h at the section $h = [h : 1]: \text{Spec } R \rightarrow \mathbb{P}^1(R)$.*

Proof There is, for each $h \in R^\times$, a rank 1 \mathbb{K} -local system \mathcal{E}^h over T_0 for which $\text{hol}(L_\infty) = h$ and $\text{hol}(L_0) = 1$. Moreover, $\mathcal{E}^{h_1} \otimes \mathcal{E}^{h_2} = \mathcal{E}^{h_1 h_2}$. The local system \mathcal{E}^h induces a strict autoequivalence α^h of $\mathcal{F}(T_0)_{\text{loc}}^{\text{ex}}$: on objects: leave the Lagrangian brane unchanged but tensor the local system by the restriction of \mathcal{E}^h . On morphism-spaces, for each intersection point $x \in L \cap L'$, map x to $\theta(h)x$, where $\theta: \mathbb{K} \rightarrow \text{End}_{\mathbb{K}} \mathcal{E}^h(x)$ is the isomorphism that sends 1 to id .

Define $\mathcal{A}^h = \alpha^h(\mathcal{A})$, the full subcategory of $\mathcal{F}(T_0)_{\text{loc}}^{\text{ex}}$ on the two objects L_0 with its trivial local system and $L_\infty^h = L_\infty$ with its local system of holonomy h . Then \mathcal{A}^h is a minimal A_∞ -structure. Moreover, α^h induces a trace-preserving isomorphism $H^*(\mathcal{A}) \cong H^*(\mathcal{A}^h)$. By Theorem C, the A_∞ -structure \mathcal{A}^h is gauge-equivalent to \mathcal{B}_C for a unique Weierstrass curve C . As an abstract curve, we have $C = \mathcal{T}_0$, but the Weierstrass data (basepoint, differential) might not be standard. We can think of C as \mathcal{T}_0 with standard differential, but different basepoint $\sigma(h)$. By Theorem C, the isomorphism

$\alpha^h: \mathcal{A} \rightarrow \mathcal{A}^h$ arises from a Weierstrass isomorphism $(\mathcal{T}_0, \omega, \sigma) \rightarrow (\mathcal{T}_0, \omega, \sigma(h))$. Automorphisms of (\mathcal{T}_0, ω) are the same thing as automorphisms of \mathbb{P}^1 that map $\{0, \infty\}$ to $\{0, \infty\}$; thus they form a group $(\mathbb{Z}/2) \ltimes R^\times$, where $\mathbb{Z}/2$ acts as the antipodal involution. Our construction gives rise to a homomorphism

$$\beta_R: R^\times \rightarrow \text{Aut } \mathcal{T}_0 = (\mathbb{Z}/2) \ltimes R^\times$$

mapping h to the automorphism β^h such that $(\phi_0)^*(L_\infty^h)$ is represented by $\beta_R(h)^*\mathcal{O}_\sigma = \mathcal{O}_{\beta_R(h)\circ\sigma}$. The homomorphisms β_R are by construction compatible with base change $R \rightarrow R'$. We claim that they must map R^\times to R^\times . Indeed, since R^\times is normal in $(\mathbb{Z}/2) \ltimes R^\times$, there is a quotient map $\bar{\beta}_R: R^\times \rightarrow \mathbb{Z}/2$, also natural in R . One must have $\bar{\beta}_k = 0$ when k is an algebraically closed field, since then every $z \in k^\times$ is a square. By naturality, $\bar{\beta}_k = 0$ for arbitrary fields k (embed k into an algebraic closure), and so for arbitrary domains R of the sort specified in the statement (embed R into its field of fractions).

In view of their compatibility with base change, we view the β_R collectively as a natural transformation $\beta: \mathbb{G}_m \Rightarrow \mathbb{G}_m$, where **rings** is the category of rings satisfying the conditions listed in the statement, and $\mathbb{G}_m: \mathbf{rings} \rightarrow \mathbf{groups}$ is the multiplicative group functor: $\mathbb{G}_m(R) = R^\times$.

Now, \mathbb{G}_m is co-represented by $S := \mathbb{Z}[t, t^{-1}]$, meaning that $\mathbb{G}_m \cong \text{Hom}_{\mathbf{rings}}(S, \cdot)$ —the identification makes the set $\text{Hom}_{\mathbf{rings}}(S, R)$ into a group. The isomorphism sends $r \in R^\times$ to the homomorphism $S \rightarrow R, f \mapsto f(r)$. We note that S is indeed a noetherian normal domain! By Yoneda's lemma, β must arise from some ring endomorphism $b: S \rightarrow S$. Moreover, $\text{Hom}_{\mathbf{rings}}(S, S) \cong S^\times = \{\pm t^d : d \in \mathbb{Z}\}$. These endomorphisms give rise to the set-theoretic natural transformations $R^\times \ni r \mapsto \pm r^d \in R^\times$, of which $r \mapsto r^d$ is a group homomorphism but $r \mapsto -r^d$ is not. We note also that β_R must be injective for each R , which leaves us only with the two possibilities $b(r) = r^{\pm 1}$. One of these possibilities is the right one, the other not; this is the ambiguity left in the statement of the lemma. \square

Remark. The imprecision in the previous lemma is easily resolved. To fix an identification of the normalization of \mathcal{T}_0 with \mathbb{P}^1 of the sort described, it suffices to describe its effect on the tangent space to \mathcal{T}_0 at σ . Now, $\sigma^*T\mathcal{T}_0 \cong R$ canonically, via the 1-form ω ; and $T_{[1:1]}\mathbb{P}^1 = T_1\mathbb{A}^1 = R$. The correct identification is the one given in these terms by id_R . That is because $\sigma^*T\mathcal{T}_0$ can be understood as $\text{Ext}^1(\mathcal{O}_\sigma, \mathcal{O}_\sigma)$, which is identified by ψ with $HF^1(L_\infty^\#, L_\infty^\#)$. The latter identification is the one determined by the trace-maps.

8.2 Higher rank local systems

In Lemma 8.1, we established that over a base ring R , the functor $\phi_0^*: \text{mod } \mathcal{F}(T_0)^{\text{ex}} \rightarrow \text{mod vect } \mathcal{T}_0$ maps L_∞^h , that is, $L_\infty^\#$ with a rank 1 local system of holonomy $h \in \mathbb{G}_m(R)$, to the skyscraper sheaf located at the section h of the smooth locus $V = \mathcal{T}_0^{\text{sm}} = \text{Spec } R[t, t^{-1}] = \mathbb{G}_m(R)$. Over \mathbb{Z} , rank 1 local systems are not very interesting, higher rank local systems more so. Working over \mathbb{Z} , take a local system on $L_\infty^\#$ with fiber \mathbb{Z}^n and holonomy $\phi \in GL(\mathbb{Z}^n)$; denote this object by $L_\infty^\phi \in \text{ob } \mathcal{F}(T_0)_{\text{loc}}^{\text{ex}}$. It maps under ϕ_0^* to a module for $\text{vect}(\mathcal{T}_0)$ co-represented by some quasi-coherent complex \mathcal{K}_ϕ^\bullet . That is, $\mathcal{Y}_L(\mathcal{K}_\phi^\bullet) \cong \phi_0^* \mathcal{Y}_L(L_\infty^\phi) \in H^0(\text{mod-vect}(\mathcal{T}_0))$.

Theorem 8.2 *Let V be the smooth locus in \mathcal{T}_0 ; thus $V \cong \text{Spec } \mathbb{Z}[t, t^{-1}] \cong \mathbb{G}_m(\mathbb{Z})$. For $\phi \in GL(\mathbb{Z}^n)$, let \mathbb{Z}_ϕ^n denote the $\mathbb{Z}[t, t^{-1}]$ -module \mathbb{Z}^n , on which t acts as ϕ . Let $(\mathbb{Z}_\phi^n)^\sim$ be the associated quasi-coherent sheaf on V , and let K_ϕ be the push-forward of $(\mathbb{Z}_\phi^n)^\sim$ to \mathcal{T}_0 . Then \mathcal{K}_ϕ^\bullet is quasi-isomorphic to the sheaf K_ϕ .*

Proof Notice that L_∞^ϕ is quasi-isomorphic in $\text{tw}^\pi \mathcal{F}(T_0)_{\text{loc}}^{\text{ex}}$ to an object of $\text{tw}^\pi \mathcal{A}$. As such, it is a *compact* object of the dg category of \mathcal{A} -modules (see for instance [9]). Quasi-equivalences preserve compact objects; consequently \mathcal{K}_ϕ^\bullet is compact as an object of $\text{QC}(\mathcal{T}_0)$. Therefore \mathcal{K}_ϕ^\bullet is quasi-isomorphic to a perfect complex [39, 9]. Since only its quasi-isomorphism class matters, we may assume that \mathcal{K}_ϕ^\bullet is a strictly perfect complex—a finite complex of locally free sheaves.

It will be helpful to work over base *fields* \mathbb{K} . We then take the object $L_\infty^\phi \in \mathcal{F}(T_0)_{\text{loc}}^{\text{ex}} \otimes \mathbb{K}$ associated with $\phi \in GL_n(\mathbb{K})$.

Step 1. Work over an algebraically closed field \mathbb{K} . We claim that \mathcal{K}_ϕ^\bullet is then quasi-isomorphic to a sheaf supported whose support is contained in the eigenvalue spectrum $\text{eval } \phi$ in $\mathbb{K}^* = \text{Spec } \mathbb{K}[t, t^{-1}] = V(\mathbb{K}) = V \times_{\text{Spec } \mathbb{Z}} \text{Spec } \mathbb{K}$.

If $\lambda \in \mathbb{K}^*$, we have a hyperext spectral sequence

$$E_2^{rs} = \text{Ext}_{\mathcal{T}_0}^r(\mathcal{H}^{-s}(\mathcal{K}_\phi^\bullet), \mathcal{O}_\lambda) \Rightarrow \text{RHom}_{\mathcal{T}_0}^{r+s}(\mathcal{K}_\phi^\bullet, \mathcal{O}_\lambda) \cong HF^{r+s}(L_\infty^\phi, L_\infty^\lambda).$$

If $\lambda \in \text{eval } \phi$ then (only) HF^0 and HF^1 are non-zero. If $\lambda \notin \text{eval } \phi$ then $HF^* = 0$. One has $\text{Ext}_{\mathcal{T}_0}^r(\mathcal{H}^{-s}(\mathcal{K}_\phi^\bullet), \mathcal{O}_\lambda) = \text{Ext}_V^r(\mathcal{H}^{-s}(j^* \mathcal{K}_\phi^\bullet), \mathcal{O}_\lambda)$, where $j: V \rightarrow \mathcal{T}_0$ is the inclusion. Since V is the spectrum of a regular local ring of dimension 1, the Ext^r -modules vanish except for $r \in \{0, 1\}$. Since it is supported in two adjacent columns, the spectral sequence degenerates at E_2 . Therefore $\text{Hom}_V(\mathcal{H}^{-s}(j^* \mathcal{K}_\phi^\bullet), \mathcal{O}_\lambda) = 0$ except when $-s \in \{0, 1\}$ and $\lambda \in \text{Spec } \phi$. Hence $j^* \mathcal{H}^{-s}(j^* \mathcal{K}_\phi^\bullet) = 0$ for $s \notin \{-1, 0\}$.

Next, consider the hyperext spectral sequence

$$'E_2^{rs} = \text{Ext}_{\mathcal{T}_0}^r(\mathcal{O}, \mathcal{H}^s(\mathcal{K}_\phi^\bullet)) \Rightarrow \text{RHom}_{\mathcal{T}_0}^{r+s}(\mathcal{O}, \mathcal{K}_\phi^\bullet) \cong \mathbb{K}^n.$$

The cohomology sheaf $\mathcal{H}^s(\mathcal{K}_\phi^\bullet)$ is supported on the singular section together with eval ϕ . One has $\text{Ext}_{\mathcal{T}_0}^r(\mathcal{O}, \mathcal{H}^s(\mathcal{K}_\phi^\bullet)) = H^r(\mathcal{H}^s(\mathcal{K}_\phi^\bullet))$ which is zero for $r \neq 0$ because $\mathcal{H}^s(\mathcal{K}_\phi^\bullet)$ has affine support. Thus $'E_2^{rs} = 0$ except when $r = 0$. So the spectral sequence degenerates, and we see that $\mathcal{H}^s(\mathcal{K}_\phi^\bullet) = 0$ for $s \neq 0$. We may therefore truncate the complex \mathcal{K}_ϕ^\bullet , replacing it by 0th cohomology $\mathcal{H}^0(\mathcal{K}_\phi^\bullet)$, to which it is quasi-isomorphic. Further, $\mathcal{H}^0(\mathcal{K}_\phi^\bullet)$ is a torsion sheaf since its stalks are generically zero.

Notice that if we have a short exact sequence of modules \mathbb{Z}_ϕ^n , the corresponding sheaves $\mathcal{H}^0(\mathcal{K}_\phi^\bullet)$ form a long exact triangle. From that, and the fact that over the algebraically closed field \mathbb{K} , matrices are conjugate to upper triangular matrices, we see that $\mathcal{H}^0(\mathcal{K}_\phi^\bullet)$ is supported in $V(\mathbb{K})$, and therefore in eval $\phi \subset V(\mathbb{K})$.

Step 2. Over an arbitrary field \mathbb{K} , \mathcal{K}_ϕ^\bullet is quasi-isomorphic to a torsion sheaf K'_ϕ , supported in $V(\mathbb{K})$, such that $H^0(K'_\phi) \cong \mathbb{K}^n$ canonically.

Indeed, by (flat) base change from \mathbb{K} to its algebraic closure, we see that the s th cohomology sheaf of \mathcal{K}_ϕ^\bullet vanishes for each $s \neq 0$. By truncation we may replace the complex by its zeroth cohomology $\mathcal{F} := \mathcal{H}^0(\mathcal{K}_\phi^\bullet)$. Moreover, \mathcal{F} vanishes at the generic point, and its stalk at the singular section is zero, so it is torsion and supported in $V \otimes_{\mathbb{Z}} \mathbb{K}$. Moreover, $H^0(\mathcal{F}) \cong \mathbb{K}^n$ via the spectral sequence $'E_{**}^*$ above.

Step 3. Over \mathbb{Z} , \mathcal{K}_ϕ^\bullet is quasi-isomorphic to a torsion sheaf supported in V . Moreover, $H^0(\mathcal{F}) \cong \mathbb{Z}^n$ canonically.

Since \mathcal{K}_ϕ^\bullet is a perfect complex, its cohomology sheaves are coherent. Hence $H^0(\mathcal{H}^s(\mathcal{K}_\phi^\bullet))$ is a finitely generated \mathbb{Z} -module for each s . By Step 2 and the compatibility of the construction with base change, one has $H^0(\mathcal{H}^s(\mathcal{K}_\phi^\bullet)) \otimes \mathbb{K} = 0$ for every field \mathbb{K} and every $s \neq 0$. Hence $H^0(\mathcal{H}^s(\mathcal{K}_\phi^\bullet)) = 0$ for $s \neq 0$. On the other hand, $z^*(\mathcal{H}^s(\mathcal{K}_\phi^\bullet))$ is a coherent sheaf on $\text{Spec } \mathbb{Z}$, for each s ; or in other words, it is a finitely generated abelian group G_ϕ . By the projection formula, and Step 2, $G_\phi \otimes \mathbb{K} = 0$ for any field \mathbb{K} . Therefore $G_\phi = 0$. Consequently $\mathcal{H}^s(\mathcal{K}_\phi^\bullet)$ is supported in V , and so is the module associated with its sections $H^0(\mathcal{H}^s(\mathcal{K}_\phi^\bullet))$, which is 0 if $s \neq 0$. Hence \mathcal{K}_ϕ^\bullet is quasi-isomorphic to its 0th cohomology sheaf. One has $H^0(\mathcal{K}_\phi^\bullet) = \mathbb{Z}^n$, again via the spectral sequence $'E_{**}^*$.

Step 4. Completion of the proof.

Since it is torsion and supported in V , the sheaf $\mathcal{H}^0(\mathcal{K}_\phi^\bullet)$ is the push-forward of a coherent sheaf on V . We think of this as the sheaf associated with a module M_ϕ . We stress that M_ϕ is canonically identified with \mathbb{Z}^n as a \mathbb{Z} -module, so $M_\phi = \mathbb{Z}_\phi^n$ for a well-defined matrix $\phi' \in GL_n(\mathbb{K})$. Moreover, the map $\phi \mapsto \phi'$ is compatible with conjugation of matrices (i.e. $(\chi\phi\chi^{-1})' = \chi\phi'\chi^{-1}$).

These points apply over any \mathbb{Z} -algebra \mathbb{K} , and the construction is compatible with base change. Hence the map $\phi \rightarrow \phi'$ arises from a map of \mathbb{Z} -schemes $F: GL_n(\mathbb{Z}) \rightarrow GL_n(\mathbb{Z})$. We claim that $F = \text{id}$. It will suffice to show that $F = \text{id}$ when we base-change to $\overline{\mathbb{Q}}$, an algebraic closure of the rationals. The induced map $F: GL_n(\overline{\mathbb{Q}}) \rightarrow GL_n(\overline{\mathbb{Q}})$ is the identity on the diagonal matrices, and therefore, by compatibility with conjugation, also on the Zariski-open set of diagonalizable matrices. Therefore it is the identity map. \square

9 The wrapped Fukaya category

We restate clause (iv) of Theorem A, in slightly refined form. The category $\text{tw vect}(\mathcal{T}_0)$ is a dg enhancement for $\text{Perf}(\mathcal{T}_0)$. Enlarge it to any dg enhancement $\tilde{D}^b \text{Coh}(\mathcal{T}_0)$ of the bounded derived category $D^b \text{Coh}(\mathcal{T}_0)$ (over \mathbb{Z}). That is, $\tilde{D}^b \text{Coh}(\mathcal{T}_0)$ is a dg category containing $\text{tw vect}(\mathcal{T}_0)$ as a full subcategory, with an equivalence of triangulated categories $H^0(\tilde{D}^b \text{Coh}(\mathcal{T}_0)) \rightarrow D^b \text{Coh}(\mathcal{T}_0)$ extending the canonical equivalence $H^0(\text{tw vect}(\mathcal{T}_0)) \rightarrow \text{Perf}(\mathcal{T}_0)$.

A standard method to construct such an enhancement would be to use injective resolutions for coherent sheaves, and then to exhibit equivalence of that approach to the approach via Čech complexes in the case of perfect complexes by the method of [47, Lemma 5.1]. However, there is no requirement for the enlargement to be of geometric origin.

Theorem 9.1 *There are a \mathbb{Z} -linear A_∞ -category \mathcal{D} and a zig-zag of quasi-equivalences*

$$\text{tw } \mathcal{W}(T_0) \longrightarrow \mathcal{D} \longleftarrow \tilde{D}^b \text{Coh}(\mathcal{T}_0)$$

making the diagram

$$\begin{array}{ccc} \text{tw } \mathcal{W}(T_0) & \longrightarrow & \mathcal{D} \longleftarrow \tilde{D}^b \text{Coh}(\mathcal{T}_0) \\ \uparrow \text{inclusion} & & \uparrow \text{inclusion} \\ \mathcal{F}(T_0)^{\text{ex}} & \xrightarrow{\psi_0} & \text{vect}(\mathcal{T}_0) \end{array}$$

homotopy-commutative.

9.1 Generation of the wrapped Fukaya category and the bounded derived category

Let $\Lambda^\# \in \text{ob } \mathcal{W}(T_0)$ be an arc of slope $(0, -1)$, graded so that $HW(L_0, \Lambda^\#)$ lies in degree 0, and oriented so that it runs into z . Our functor ψ_{wrap} will carry $\Lambda^\#$ to \mathcal{O}_s ,

the skyscraper sheaf along the nodal section.

Lemma 9.2 $\mathcal{W}(T_0)$ is generated by $\mathcal{F}(T_0)$ and $\Lambda^\#$.

Proof First, if orient Λ in the opposite direction, we obtain an isomorphic object of $\mathcal{W}(T_0)$. This reflects the fact that a spin-structure on a Lagrangian (which trivializes w_1 and w_2) is more data than is needed; a Pin-structure, trivializing w_2 , is sufficient [50]. Changing the orientation corresponds to an automorphism $(-1)^{\deg} \text{id}$ of the object. The brane structure also involves a double covering of Λ , but that is necessarily trivial. With these points noted, we find that any object of $\mathcal{W}(T_0)$ whose Lagrangian is non-compact is quasi-isomorphic to a shift of an iterated Dehn twist of $\Lambda^\#$ along closed, exact curves equipped with non-trivial double coverings. These Dehn twists act on $\text{tw } \mathcal{W}(T_0)$ by spherical twists [50, 35]; hence the arcs are represented by twisted complexes in $\mathcal{F}(T_0)$ and $\Lambda^\#$. \square

Lemma 9.3 $\widetilde{\text{D}} \text{Coh}(\mathcal{T}_0)$ is generated by $\text{vect}(\mathcal{T}_0)$ and the skyscraper sheaf \mathcal{O}_s along the nodal section s .

Proof It suffices to show that $G_0(\mathcal{T}_0)$ is generated, as an abelian group, by the image of $K_0(\mathcal{T}_0)$ and the class $[\mathcal{O}_s]$. Here G_0 denotes the Grothendieck group of coherent sheaves, while K_0 is the Grothendieck group of vector bundles. The quotient $G_0(\mathcal{T}_0)/\text{im } K_0(\mathcal{T}_0)$ is certainly generated by coherent sheaves supported on Z , the closure of $\text{im } z$. Thus it will suffice to show that $K_0(M_Z(\mathcal{T}_0))$, the Grothendieck group of the abelian category of coherent sheaves supported along Z , is generated by the class of \mathcal{O}_s . We now proceed as in [63, ex. II 6.3.4]: $M_Z(\mathcal{T}_0)$ is the abelian category of finitely-generated modules M for $R = \mathbb{Z}[x, y]/(y^2 + xy - x^3)$ such that $I^n M = 0$ for some n , where $I = (x, y) \subset R$. Such a module has a filtration $M \supset IM \supset I^2 M \supset \cdots \supset I^n M = 0$ and therefore $K_0(M_Z(\mathcal{T}_0))$ is generated by the factors of such filtrations, i.e., by modules N with $IN = 0$. As sheaves, those are precisely the push-forwards of sheaves on Z , or equivalently those of the form $s_* \mathcal{F}$ for some coherent sheaf \mathcal{F} on $\text{Spec } \mathbb{Z}$. Since $G_0(Z) = \mathbb{Z}$, the result follows. \square

9.2 The wrapped category embeds into \mathcal{A} -modules

Work over the base ring \mathbb{K} . Let's write \mathcal{V} for the full subcategory of $\mathcal{W}(T_0)$ consisting of \mathcal{A} together with the arc $\Lambda^\#$. There are A_∞ -functors

$$\begin{array}{ll} \mathcal{V} \xrightarrow{\mathcal{Y}} \text{mod-}\mathcal{V} & (\text{Yoneda embedding, } Y \mapsto \text{hom}_{\mathcal{V}}(\cdot, Y)) \\ \xrightarrow{r} \text{mod-}\mathcal{A} & (\text{restriction}) \end{array}$$

whose composite we denote by ρ .

Our aim in this section is to prove the

Theorem 9.4 *r is cohomologically full and faithful.*

Since the Yoneda functor is an embedding, this statement has the following

Corollary 9.5 *ρ is cohomologically full and faithful.*

The strategy is as follows. Write $A = H^*\mathcal{A}$ and $V = H^*\mathcal{V}$. In general, for graded \mathcal{V} -modules \mathcal{M} and \mathcal{N} , the length filtration on $\mathrm{hom}_{\mathrm{mod-}\mathcal{V}}^*(\mathcal{M}, \mathcal{N}[s])$ gives rise to a spectral sequence $\{^V E_q^{r,s}\}$, with

$$^V E_2^{r,s} = \mathrm{Ext}_V^r(H^*\mathcal{M}, H^*\mathcal{N}[s]).$$

Here $\mathrm{Ext}_V^r(H^*\mathcal{M}, H^*\mathcal{N}[s])$ means the r th derived functor of degree zero homomorphism between graded V -modules. It does not necessarily converge, but when it does, $^V E_\infty^{r,s}$ is the associated graded of a filtration on

$$H^* \mathrm{hom}_{\mathrm{mod-}\mathcal{V}}^{r+s}(\mathcal{M}, \mathcal{N}[s]).$$

Restricting the modules to \mathcal{A} , we get a morphism of filtered complexes

$$\mathrm{hom}_{\mathrm{mod-}\mathcal{V}}(\mathcal{M}, \mathcal{N}[s]) \rightarrow \mathrm{hom}_{\mathrm{mod-}\mathcal{A}}(\mathcal{M}, \mathcal{N}[s]),$$

and hence a map of spectral sequences $\{^V E_q^{r,s}\} \rightarrow \{^A E_q^{r,s}\}$. At the E_2 page, the latter map is the derived restriction map

$$(40) \quad \mathrm{Ext}_V^r(H^*\mathcal{M}, H^*\mathcal{N}[s]) \rightarrow \mathrm{Ext}_A^r(H^*\mathcal{M}, H^*\mathcal{N}[s]).$$

We will prove the

Lemma 9.6 *When \mathcal{M} and \mathcal{N} are Yoneda modules for objects of \mathcal{V} , derived restriction (40) is an isomorphism.*

Since the length filtrations (over \mathcal{V} and over \mathcal{A}) are complete, the spectral sequence comparison principle (Lemma 2.1) implies that the restriction map

$$r: H^* \mathrm{hom}_{\mathrm{mod-}\mathcal{V}}(\mathcal{M}, \mathcal{N}[s]) \rightarrow H^* \mathrm{hom}_{\mathrm{mod-}\mathcal{A}}(\mathcal{M}, \mathcal{N}[s])$$

is an isomorphism. That proves Theorem 9.4, modulo a proof of Lemma 9.6.

A-modules. To set the stage for the proof of Lemma 9.6, we discuss (graded, right) A -modules.

Write $a = L_0^\#$ and $b = L_\infty^\#$ as objects of A . We have $\text{Hom}_A(a, b) = \mathbb{K} \cdot u_0$, where u_0 has degree 0, and $\text{Hom}_A(b, a) = \mathbb{K} \cdot v_1$, where v_1 has degree 1. We have $u \circ v \circ u = 0$, $v \circ u \circ v = 0$. A graded right A -module, M , amounts to a diagram

$$\begin{array}{ccc} M(a) & \begin{array}{c} \xrightarrow{v_1} \\ \xleftarrow{u_0} \end{array} & M(b) \end{array}$$

where $M(a)$ are graded \mathbb{K} -modules, u_0 has degree 0, v_1 has degree 1, and $u_0 \circ v_1 \circ u_0 = 0$ and $v_1 \circ u_0 \circ v_1 = 0$.

One can view A as an algebra (by taking the direct sum of all the morphism spaces), over the ring

$$\mathbb{L} = \mathbb{K} \cdot e_a \oplus \mathbb{K} \cdot e_b$$

generated by the identity morphisms e_a and e_b . One can thereby import notions from standard algebra. The Yoneda module $\text{Hom}_A(\cdot, a)$ is a projective graded A -module P_a . Its diagram is as follows:

$$\begin{array}{ccc} \mathbb{K} \oplus \mathbb{K}[-1] & \begin{array}{c} \xrightarrow{\begin{bmatrix} 1 & 0 \end{bmatrix}} \\ \xleftarrow{\begin{bmatrix} 0 \\ 1 \end{bmatrix}} \end{array} & \mathbb{K}[-1]. \end{array}$$

Similarly, $P_b = \text{Hom}_A(\cdot, b)$ is a projective graded module with diagram

$$\begin{array}{ccc} \mathbb{K} & \begin{array}{c} \xrightarrow{\begin{bmatrix} 0 \\ 1 \end{bmatrix}} \\ \xleftarrow{\begin{bmatrix} 1 & 0 \end{bmatrix}} \end{array} & \mathbb{K} \oplus \mathbb{K}[-1]. \end{array}$$

The module $P_b[-1]$ has a simple quotient S_b , unique up to isomorphism, with diagram

$$0 \begin{array}{c} \xrightarrow{\quad} \\ \xleftarrow{\quad} \end{array} \mathbb{K}[-1]$$

Similarly, P_a has a unique simple quotient S_a :

$$\mathbb{K} \begin{array}{c} \xrightarrow{\quad} \\ \xleftarrow{\quad} \end{array} 0$$

The arc $\Lambda^\#$ defines an A -module $H^*(\rho(\Lambda^\#)) = H^*(\text{Hom}_{\mathcal{W}}(\cdot, \Lambda^\#)|_A)$ with diagram

$$\mathbb{K} \begin{array}{c} \xrightarrow{\quad} \\ \xleftarrow{\quad} \end{array} 0$$

Thus we have $M_\Lambda \cong S_a$. The projection $P_a \rightarrow M_\Lambda \rightarrow 0$ arises from the geometric generator of $HW(\Lambda^\#, L_0^\#)$.

There are exact sequences

$$\begin{aligned} 0 \longrightarrow S_b \longrightarrow P_b \longrightarrow P_a \longrightarrow S_a \longrightarrow 0 \\ 0 \longrightarrow S_a[-3] \longrightarrow P_a[-2] \longrightarrow P_b[-1] \longrightarrow S_b \longrightarrow 0, \end{aligned}$$

and these splice to form an infinite projective resolution $\mathcal{S}_a^\bullet \rightarrow S_a \rightarrow 0$, given by

$$\mathcal{S}_a^\bullet = \{\cdots \rightarrow P_b[-3] \rightarrow P_a[-3] \rightarrow P_a[-2] \rightarrow P_b[-1] \rightarrow P_b \rightarrow P_a\}.$$

The morphisms making up \mathcal{S}_a^\bullet all arise from the geometric generators (intersection points) for morphism-spaces in A .

We proceed to calculate some Ext modules.

Since P_a and P_b are projective, one has

$$\bigoplus_s \text{Ext}_A(P_a, S_a[s]) = \bigoplus_s \text{Hom}_A^s(P_a, S_a) = \mathbb{K}, \quad \text{Ext}_A(P_b, S_a[s]) = \text{Hom}_A^s(P_b, S_a) = 0.$$

Remark. One can show that A is *self-injective*, meaning A is injective as a right A -module. From this it follows that right projective modules coincide with right injective modules. From that we see that P_a and P_b are injective, whence one gets a quick verification that

$$\bigoplus_s \text{Ext}_A(S_a, P_a[s]) = \bigoplus_s \text{Hom}_A^s(S_a, P_a) = \mathbb{K}[-1], \quad \text{Ext}_A(S_a, P_b[s]) = \text{Hom}_A^s(S_a, P_b) = 0.$$

We will not make direct use of the last remark. Instead, we observe that we can calculate Exts using the projective resolution $\mathcal{S}_a^\bullet \rightarrow S_a \rightarrow 0$:

$$\text{Ext}_A^r(S_b, N) = H^* \text{Hom}_A(\mathcal{S}_a^\bullet, N).$$

We are interested in the cases $N = P_a$, P_b or S_b ; but we will not need to compute the cohomology on the right explicitly.

Ext-modules over V . Notice that P_a (an A -module) naturally extends to a projective V -module, the cohomology of the \mathcal{V} -module $\text{hom}_{\mathcal{V}}(\cdot, a)$. Similarly P_b extends to a projective V -module. We still denote these modules by P_a and P_b .

Lemma 9.7 *Restriction*

$$(41) \quad \mathrm{Hom}_{\mathrm{mod-}V}(P_i, P_j[s]) \rightarrow \mathrm{Hom}_{\mathrm{mod-}A}(P_i, P_j[s])$$

is an isomorphism for each s and all four possibilities for $i \in \{a, b\}$ and $j \in \{a, b\}$.

Proof This follows from Yoneda's lemma and the fact that the inclusion $A \rightarrow V$ is full and faithful. \square

The simple A -module S_a also extends naturally to a V -module M_Λ (it is the Yoneda module for $\Lambda^\#$). The A -module morphism $P_a \rightarrow S_a \rightarrow 0$ extends to a V -module homomorphism $P_a \rightarrow M_\Lambda \rightarrow 0$ (coming from the geometric generator of $HW(\Lambda^\#, a)$).

Lemma 9.8 *The restriction maps*

$$\mathrm{Hom}_{\mathrm{mod-}V}(P_a, M_\Lambda[s]) \rightarrow \mathrm{Hom}_{\mathrm{mod-}A}(P_a, S_a[s])$$

$$\mathrm{Hom}_{\mathrm{mod-}V}(P_b, M_\Lambda[s]) \rightarrow \mathrm{Hom}_{\mathrm{mod-}A}(P_b, S_a[s])$$

are isomorphisms for each s .

Proof For surjectivity note that $\mathrm{Hom}_{\mathrm{mod-}A}(P_a, S_a[s])$ is non-zero only when $s = 0$, in which case it is \mathbb{K} , the generator arising from a geometric morphism $\Lambda^\# \rightarrow a$; and that $\mathrm{Hom}_{\mathrm{mod-}A}(P_b, S_a[s]) = 0$. For injectivity, note that the V -module structure of M_Λ involves a non-trivial morphism $M_\Lambda(\Lambda^\#) \rightarrow M_\Lambda(a)$ to see that vanishing of a V -module map on a and b implies vanishing on $\Lambda^\#$. \square

Proof of Lemma 9.6. The restriction maps

$$r: \mathrm{Hom}_{\mathrm{mod-}V}(P_i, P_j[s]) \rightarrow \mathrm{Hom}_{\mathrm{mod-}A}(P_i, P_j[s])$$

(where $i, j \in \{a, b\}$) are isomorphisms by the Yoneda lemma. Next, restriction induces isomorphisms

$$\mathrm{Ext}_V^*(P_b, M_\Lambda) \xrightarrow{\cong} \mathrm{Ext}_A^*(P_b, S_a).$$

$$\mathrm{Ext}_V^*(P_a, M_\Lambda) \xrightarrow{\cong} \mathrm{Ext}_A^*(P_a, S_a)$$

because P_b and P_a are projective (over V and over A), so both sides reduce to module homomorphisms, and we can invoke Lemma 9.8.

Now we consider

$$r: \mathrm{Ext}_V^*(M_\Lambda, N) \rightarrow \mathrm{Ext}_A^*(S_a, N),$$

where N may be a shift of P_a , P_b or M_Λ . We compute this derived restriction map using the projective resolution $S_a^\bullet \rightarrow S_a \rightarrow 0$ for S_a as an A -module. We note that it naturally extends to a projective resolution $S_a^\bullet \rightarrow M_\Lambda \rightarrow 0$ for M_Λ as a V -module.

We have $\mathrm{Ext}_V^*(S_a, N) = H^* \mathrm{hom}_V(\mathcal{S}_a^\bullet, N)$ and $\mathrm{Ext}_A^*(S_a, N) = H^* \mathrm{hom}_A(\mathcal{S}_a^\bullet, N)$. The derived restriction map r becomes the map

$$H^* \mathrm{hom}_V(\mathcal{S}_a^\bullet, N) \rightarrow H^* \mathrm{hom}_A(\mathcal{S}_a^\bullet, N)$$

induced by restriction of module homomorphisms. For any $i \geq 0$, we have $\mathcal{S}_a^i = P_j$ (where $j \in \{a, b\}$); hence restriction

$$\mathrm{hom}_V(\mathcal{S}_a^i, N) \rightarrow \mathrm{hom}_A(\mathcal{S}_a^i, N)$$

is an isomorphism by Lemma 9.7 (when $N = P_a[s]$ or $P_b[s]$), or by Lemma 9.8 (when $N = S_a[s]$). Thus $\mathrm{hom}_V(\mathcal{S}_a^\bullet, N) \rightarrow \mathrm{hom}_A(\mathcal{S}_a^\bullet, N)$ is an isomorphism of cochain complexes, inducing an isomorphism on cohomology. \square

9.3 The coherent derived category embeds into $\mathrm{vect}(\mathcal{T}_0)$ -modules

Denote by $\mathrm{QC}(\mathcal{T}_0)$ a dg enhancement of the unbounded derived category of quasi-coherent complexes, as in [60]. One has a Yoneda functor

$$(42) \quad \mathrm{QC}(\mathcal{T}_0) \rightarrow \mathrm{mod}\text{-}\mathrm{vect}(\mathcal{T}_0).$$

Proposition 9.9 *The functor (42) induces an equivalence of cohomology categories.*

Proof This is an application of [60, Theorem 8.9] or [9]. The deduction is as follows: The dg category of modules over $\mathrm{vect}(\mathcal{T}_0)$ is, rather trivially, equivalent to the dg category of dg functors $[\mathrm{tw} \mathrm{vect}(\mathcal{T}_0), \mathrm{tw} \mathrm{vect}(\mathrm{Spec} \mathbb{Z})]$. There are further equivalences $[\mathrm{tw} \mathrm{vect}(\mathcal{T}_0), \mathrm{tw} \mathrm{vect}(\mathrm{Spec} \mathbb{Z})] \rightarrow [\mathrm{QC}(\mathcal{T}_0), \mathrm{QC}(\mathrm{Spec} \mathbb{Z})]_c$ where the c denotes that the functors respect filtered colimits (think of quasi-coherent complexes as filtered colimits of perfect complexes); and $\mathrm{QC}(\mathcal{T}_0 \times_{\mathrm{Spec} \mathbb{Z}} \mathrm{Spec} \mathbb{Z}) = \mathrm{QC}(\mathcal{T}_0) \rightarrow [\mathrm{QC}(\mathcal{T}_0), \mathrm{QC}(\mathrm{Spec} \mathbb{Z})]_c$, mapping an integral kernel to its push-pull functor. \square

Let

$$\tilde{\mathrm{D}}^b \mathrm{Coh}(\mathcal{T}_0) \subset \mathrm{QC}(\mathcal{T}_0)$$

denote the full dg subcategory whose objects are bounded complexes with coherent cohomology. The restriction of the Yoneda functor is a cohomologically full and faithful functor

$$(43) \quad \tilde{\mathrm{D}}^b \mathrm{Coh}(\mathcal{T}_0) \rightarrow \mathrm{mod}\text{-}\mathrm{vect}(\mathcal{T}_0).$$

9.4 The arc is mirror to the skyscraper sheaf at the singularity

We have a sequence of A_∞ -functors

$$(44) \quad \mathcal{W}(T_0) \xrightarrow{\mathcal{Y}_L} \text{mod-}\mathcal{A} \xrightarrow{\phi_0^*} \text{mod-}\mathcal{B}_{\mathcal{T}_0} \xleftarrow{\simeq} \text{mod-}\text{vect}(\mathcal{T}_0) \xleftarrow{\simeq} \text{QC}(\mathcal{T}_0).$$

Here \mathcal{Y}_L is the (covariant) left Yoneda functor, $X \mapsto \text{hom}_{\mathcal{A}}(X, \cdot)$. The functor $\text{mod-}\mathcal{B}_{\mathcal{T}_0} \xleftarrow{\simeq} \text{mod-}\text{vect}(\mathcal{T}_0)$ is restriction, which is a quasi-equivalence because \mathcal{B}_C split-generates $\text{tw vect}(\mathcal{T}_0)$.

We now define \mathcal{Z}^\bullet to be an object of $\text{QC}(\mathcal{T}_0)$ corresponding under mirror symmetry to the arc $\Lambda^\#$. Precisely:

Define \mathcal{Z}^\bullet to be a choice of quasi-coherent complex whose associated $\mathcal{B}_{\mathcal{T}_0}$ -module $\mathcal{Y}_L(\mathcal{Z}^\bullet) = \text{hom}_{\mathcal{B}_{\mathcal{T}_0}}(\mathcal{Z}^\bullet, \cdot)$ is quasi-isomorphic to $\phi_0^* \mathcal{Y}_L(\Lambda^\#)$.

We learn from Theorem 9.4 and Proposition 9.9 that

$$(45) \quad \text{Hom}_{\text{QC}(\mathcal{T}_0)}(\mathcal{O}, \mathcal{Z}^\bullet) \cong HW(L_0^\#, \Lambda^\#), \quad \text{Hom}_{\text{QC}(\mathcal{T}_0)}(\mathcal{O}_\sigma, \mathcal{Z}^\bullet) \cong HW(L_\infty^\#, \Lambda^\#),$$

and also that

$$(46) \quad \text{Hom}_{\text{QC}(\mathcal{T}_0)}(\mathcal{Z}^\bullet, \mathcal{Z}^\bullet) \cong HW(\Lambda^\#, \Lambda^\#).$$

A key point in the proof of mirror symmetry for the wrapped category will be the following assertion:

Proposition 9.10 *One has $\mathcal{Z}^\bullet \simeq \mathcal{O}_s$ in $\text{QC}(\mathcal{T}_0)$, where as before, \mathcal{O}_s is the push-forward of $\mathcal{O}_{\text{Spec } \mathbb{Z}}$ by the singular section $s: \text{Spec } \mathbb{Z} \rightarrow \mathcal{T}_0$.*

Lemma 9.11 *If we work over a base field \mathbb{K} , then the quasi-coherent complex $\mathcal{Z}_{\mathbb{K}}^\bullet$ mirror to the arc has the property that its cohomology sheaves $\mathcal{H}_{\mathbb{K}}^k$ are supported in the open set $U(\mathbb{K}) = \text{Spec } \mathbb{K}[x, y]/(y^2 + xy - x^3)$.*

Proof In general, if \mathcal{E}^\bullet and \mathcal{F}^\bullet are complexes of quasi-coherent sheaves over a scheme X , at least one of them bounded, one has a right half-plane spectral sequence arising from a filtration on a complex computing $\text{R Hom}_{\mathcal{O}_X}^{r+s}(\mathcal{E}^\bullet, \mathcal{F})$, with

$$(47) \quad E_2^{rs} = \bigoplus_k \text{Ext}_{\mathcal{O}_X}^r(\mathcal{H}^{k-s}(\mathcal{E}^\bullet), \mathcal{H}^k(\mathcal{F}^\bullet)),$$

abutting to $\text{R Hom}_{\mathcal{O}_X}^{r+s}(\mathcal{E}^\bullet, \mathcal{F}^\bullet)$. In particular, we have a spectral sequence

$$(48) \quad E_2^{rs} = \text{Ext}_{\mathcal{T}_0}^r(\mathcal{O}_p, \mathcal{H}_{\mathbb{K}}^s) \Rightarrow \text{R Hom}_{\mathcal{T}_0}^{r+s}(\mathcal{O}_\sigma, \mathcal{Z}_{\mathbb{K}}^\bullet).$$

where $p = [0 : 1 : 0]$. We know by (45) that $\text{R Hom}_{\mathcal{T}_0}(\mathcal{O}_\sigma, \mathcal{Z}_{\mathbb{K}}^\bullet) = 0$. This spectral sequence degenerates at E_2 . Indeed, let M^s denote the module of global sections of

$j^* \mathcal{H}_{\mathbb{K}}^s$, where j is the open inclusion of the smooth locus $V(\mathbb{K}) \cong \operatorname{Spec} \mathbb{K}[t, t^{-1}]$. Let $\mathfrak{p} = (t - 1)\mathbb{K}[t, t^{-1}]$, a maximal ideal of $\mathbb{K}[t, t^{-1}]$, and let A denote the localization of $\mathbb{K}[t, t^{-1}]$ at \mathfrak{p} . Then $\operatorname{Ext}_{\mathcal{T}_0}^r(\mathcal{O}_\sigma, \mathcal{H}_{\mathbb{K}}^s) \cong \operatorname{Ext}_A^r(\mathbb{K}, \mathcal{H}_{\mathfrak{p}}^s)$. Since A is a regular ring of dimension 1, these Ext-modules vanish except for $r = 0$ or 1. Therefore the spectral sequence is concentrated in two adjacent columns, and so degenerates. Thus $\operatorname{Hom}_A(\mathbb{K}, \mathcal{H}_{\mathfrak{p}}^s) = 0$, and so $\mathcal{H}_{\mathfrak{p}}^s = 0$.

Consequently, $\mathcal{H}_{\mathbb{K}}^k$ is for each k supported in $U(\mathbb{K})$. \square

Lemma 9.12 *Over \mathbb{Z} , the quasi-coherent complex \mathcal{Z}^\bullet has the property that its cohomology sheaves \mathcal{H}^k are supported in the open set $U = \operatorname{Spec} \mathbb{Z}[x, y]/(y^2 + xy - x^3)$*

Proof Again, let $j^* \mathcal{H}^k$ be the restriction of \mathcal{H}^k to the smooth locus $j: V \rightarrow \mathcal{T}_0$. One has $V \cong \operatorname{Spec} R$ where $R = \mathbb{Z}[t, t^{-1}]$, an identification under which the section σ is defined by the homomorphism $R \rightarrow \mathbb{Z}$, $p(t) \mapsto p(1)$. Let M^k be the R -module of sections $\Gamma(j^* \mathcal{H}^k, V)$. Since \mathcal{H}^k is quasi-coherent, $j^* \mathcal{H}^k$ is the sheaf associated with M^k . Let $\mathfrak{p} = (t - 1)R$, and let $R_{\mathfrak{p}}$ be the localization at \mathfrak{p} . Our task is to show that the $R_{\mathfrak{p}}$ -module $M_{\mathfrak{p}}^k$ is zero.

We observe that $M^k \otimes_{\mathbb{Z}} \mathbb{Q} = 0$. For this, consider the wrapped Fukaya category with \mathbb{Q} -coefficients; the arc now maps to $\mathcal{Z}_{\mathbb{Q}}^\bullet = \mathcal{Z}^\bullet \otimes_{\mathbb{Z}} \mathbb{Q}$. One has $\mathcal{H}^k(\mathcal{Z}^\bullet \otimes \mathbb{Q}) = \mathcal{H}^k \otimes \mathbb{Q}$. Also, $j^*(\mathcal{H}^k(\mathcal{Z}^\bullet \otimes \mathbb{Q})) = j^* \mathcal{H}^k \otimes \mathbb{Q}$. Thus, $j^*(\mathcal{H}^k(\mathcal{Z}^\bullet \otimes \mathbb{Q}))$ is the sheaf associated with its module of sections $M^k \otimes_{\mathbb{Z}} \mathbb{Q}$ over $R \otimes_{\mathbb{Z}} \mathbb{Q}$. The previous lemma then implies that $M^k \otimes_{\mathbb{Z}} \mathbb{Q} = 0$.

Now, $R_{\mathfrak{p}} = S^{-1}R$, where $S = R \setminus \mathfrak{p}$; in particular, S contains every prime of \mathbb{Z} , and hence $R_{\mathfrak{p}}$ is a \mathbb{Q} -algebra. Thus $M_{\mathfrak{p}}^k$ is a \mathbb{Q} -algebra, yet is torsion as an abelian group. Hence $M_{\mathfrak{p}}^k = 0$. \square

Proof of Prop. 9.10 Consider the spectral sequence

$$(49) \quad {}^r E_2^{rs} = \operatorname{Ext}_{\mathcal{T}_0}^r(\mathcal{O}, \mathcal{H}^s) \Rightarrow \operatorname{RHom}_{\mathcal{T}_0}^{r+s}(\mathcal{O}, \mathcal{Z}^\bullet) \cong \mathbb{Z}.$$

We used (45) to compute the target of the spectral sequence.

One has $\operatorname{Ext}_{\mathcal{T}_0}^r(\mathcal{O}, \mathcal{H}^s) = H^r(\mathcal{H}^s)$. But the foregoing lemma says that \mathcal{H}^s has affine support in U ; hence $H^r(\mathcal{H}^s) = 0$ for all $r > 0$. The spectral sequence is therefore concentrated in one column, and so degenerates. Hence $H^0(\mathcal{H}^s)$ vanishes for all $s \neq 0$. Since \mathcal{H}^s is quasi-coherent with affine support in U , it follows that $\mathcal{H}^s = 0$ for $s \neq 0$.

Because of the vanishing of the non-zero cohomology sheaves, we have a diagram of quasi-isomorphisms $\mathcal{Z}^\bullet \leftarrow \tau_{\leq 0} \mathcal{Z}^\bullet \rightarrow \mathcal{H}^0$, where $\tau_{\leq 0}$ denotes the truncation $\cdots \rightarrow$

$\mathcal{Z}^{-2} \rightarrow \mathcal{Z}^{-1} \rightarrow \ker \delta^0 \rightarrow 0$. Therefore $\mathcal{Z}_{\mathbb{K}}^{\bullet}$ is quasi-isomorphic to the sheaf \mathcal{H}^0 , with which we may replace it.

Moreover, \mathcal{H}^0 also has affine support in U , and $H^0(\mathcal{H}^0) \cong \mathbb{Z}$. Hence \mathcal{H}^0 is the push forward from U of the sheaf associated with the $\mathbb{Z}[x, y]/(y^2 + xy - x^3)$ -module \mathbb{Z} , with some actions of x and y .

The actions of x and y on the module \mathbb{Z} are multiplication by integers X and Y . We claim $(X, Y) = (0, 0)$. Take a prime $p \in \mathbb{Z}$ and work over $\mathbb{K} = \mathbb{Z}/(p)$. By (46), the derived endomorphisms of $\mathcal{Z}_{\mathbb{K}}^{\bullet}$ are infinite-dimensional. Hence (X, Y) cannot be a regular point of $U(\mathbb{K})$. Thus $(X, Y) = (0, 0) \bmod p$. Since p is arbitrary, we have $(X, Y) = (0, 0)$.

Hence $\mathcal{H}^0 = s_* \mathcal{O}_{\text{Spec } \mathbb{Z}}$. □

9.5 Proving the mirror equivalence

We have a Yoneda functor $r: \mathcal{W}(T_0) \rightarrow \text{mod-}\mathcal{F}(T_0)^{\text{ex}}$ which, by Theorem 9.4, is cohomologically full and faithful. Let \mathcal{W}' denote the smallest full subcategory of $\text{mod-}\mathcal{F}(T_0)^{\text{ex}}$ which contains every object quasi-isomorphic to one in $r(\text{ob } \mathcal{W}(T_0))$.

We also have a Yoneda functor $\tilde{D}^b \text{Coh } \mathcal{T}_0 \rightarrow \text{mod}-(\text{tw vect}(\mathcal{T}_0))$ from (43) which, by proposition 9.9 is again cohomologically full and faithful. Let \mathcal{D} denote the smallest full subcategory of $\text{mod}-(\text{tw vect}(\mathcal{T}_0))$ which contains every object quasi-isomorphic to one in the image of $\tilde{D}^b \text{Coh } \mathcal{T}_0$.

Notice that the functor $\phi_0^*: \text{mod-}\mathcal{F}(T_0)^{\text{ex}} \rightarrow \text{mod-vec } \mathcal{T}_0$ carries \mathcal{W}' to \mathcal{D} . That's because it maps the generators $\text{hom}_{\mathcal{W}(T_0)}(\cdot, \Lambda^{\#})$, $\text{hom}_{\mathcal{W}(T_0)}(\cdot, L_0^{\#})$ and $\text{hom}_{\mathcal{W}(T_0)}(\cdot, L_{\infty}^{\#})$ of \mathcal{W}' to modules quasi-isomorphic to $\text{hom}_{\text{vec } \mathcal{T}_0}(\cdot, \mathcal{O}_s)$, $\text{hom}_{\text{vec } \mathcal{T}_0}(\cdot, \mathcal{O})$ and $\text{hom}_{\text{vec } \mathcal{T}_0}(\cdot, \mathcal{O}_{\sigma})$ (we invoke Theorem 9.10). In the opposite direction ψ_0^* carries \mathcal{D} to \mathcal{W}' . These two functors are homotopy-inverses, because that is true of ϕ_0 and ψ_0 .

We sum up the situation in the homotopy-commutative diagram of A_{∞} -functors

$$\begin{array}{ccccc}
 & & \phi_0^* & & \\
 & & \curvearrowright & & \\
 \mathcal{W}(T_0) & \longrightarrow & \mathcal{W}' & \longrightarrow & \mathcal{D} \longleftarrow \tilde{D}^b \text{Coh } \mathcal{T}_0 \\
 & & \curvearrowleft & & \\
 & & \psi_0^* & & \\
 \text{inclusion} \uparrow & & & & \uparrow \text{inclusion} \\
 \mathcal{F}(T_0)^{\text{ex}} & \xrightarrow{\quad \simeq \quad} & & & \text{tw vect}(\mathcal{T}_0)
 \end{array}$$

Since the arrows on the top line are quasi-equivalences, one obtains Theorem 9.1.

Remark. In an earlier version of the paper, we erroneously claimed that one can construct a functor

$$\psi_{\text{wrap}}: \text{tw } \mathcal{W}(T_0) \rightarrow \widetilde{\mathcal{D}}^b \text{ Coh } \mathcal{T}_0$$

extending ψ_0 . We actually do not know whether that is possible. It depends on finding a quasi-inverse to the arrow

$$\mathcal{D} \leftarrow \widetilde{\mathcal{D}}^b \text{ Coh } \mathcal{T}_0.$$

A natural strategy is to consider the Yoneda functor $\mathcal{U} := \{\mathcal{O}, \mathcal{O}_\sigma, \mathcal{O}_s\} \rightarrow \text{mod-}\mathcal{B}_{\mathcal{T}_0}$. Its image is a 3-object category $\mathcal{Y}(\mathcal{U})$, and one wants to homotopy-invert $\mathcal{U} \rightarrow \mathcal{Y}(\mathcal{U})$. However, it is not clear that corollary 2.4 applies, because one would need to show that $\mathcal{Y}(\mathcal{U})$ can be taken to have splittable cochains.

10 Identifying the mirror curve over $\mathbb{Z}[[q]]$: the Tate curve and toric geometry

In this section, we offer a third proof that $C_{\text{mirror}}|_{q=0}$ has the equation $y^2 + xy = x^3$. More significantly, this proof extends to show that C_{mirror} is the Tate curve.

10.1 The Tate curve

The Tate curve was constructed by Raynaud using formal schemes; we have followed the expositions by Deligne–Rapoport [17] and Conrad [14], and Gross’s reinterpretation of the construction in toric language [23]. We review the construction.

10.1.1 Construction of a scheme with \mathbb{Z} -action $\mathcal{T}_\infty \rightarrow \text{Spec } \mathbb{Z}[t]$

Toric construction of \mathcal{T}_∞ . We begin with the toric fan picture. Fix a commutative ring R . Consider the rays $\rho_i = \mathbb{Q}_+(i, 1) \subset \mathbb{Q}^2$, where $i \in \mathbb{Z}$ (Figure 10.1.1). The convex hull of ρ_i and ρ_{i+1} is a cone $\sigma_{i+1/2} \subset \mathbb{Q}^2$; the collection of cones $\sigma_{i+1/2}$, their boundary faces ρ_j , and *their* common endpoint $\{0\}$, form a rational fan F in \mathbb{Q}^2 . Each cone c of F has a dual cone

$$c^\vee = \{x \in \mathbb{Q}^2 : \langle x, s \rangle \geq 0 \ \forall s \in c\}$$

(the use of inner products rather than dual spaces is an aid to visualization). To be explicit, ρ_i^\vee is the half-plane $\{(m, n) \in \mathbb{Q}^2 : im + n \geq 0\}$, and $\sigma_{i+1/2}^\vee = \rho_i^\vee \cap \rho_{i+1}^\vee$. The dual cones c^\vee have semigroup R -algebras $R[c^\vee \cap \mathbb{Z}^2]$ spanned by monomials e^λ


 Figure 2: Part of the fan F

where $\lambda \in c^\vee \cap \mathbb{Z}^2$. We glue together the affine schemes $U_{i+1/2} = \operatorname{Spec} R[\sigma_{i+1/2}^\vee \cap \mathbb{Z}^2]$ so as to form a toric scheme $\mathcal{T}_\infty \rightarrow \operatorname{Spec} R$ with an affine open cover $\{U_{i+1/2}\}$.

For each i , there is a map of cones $\sigma_{i+1/2} \rightarrow \mathbb{Q}_+$, $(x_1, x_2) \mapsto x_2$. The dual to this map is the map of cones $\mathbb{Q}_+ \rightarrow \sigma_{i+1/2}^\vee$ given by $1 \mapsto (0, 1)$, which induces a map of semigroup R -algebras $R[\mathbb{Q}_+^\vee \cap \mathbb{Z}] \rightarrow R[\sigma_{i+1/2}^\vee \cap \mathbb{Z}^2]$ and hence a function $U_{i+1/2} \rightarrow \mathbb{A}^1(R)$. The maps $\sigma_{i+1/2} \rightarrow \mathbb{Q}_+$ assemble to form a map of fans $F \rightarrow \mathbb{Q}_+$, whence the functions $U_{i+1/2} \rightarrow \mathbb{A}^1(R)$ consistently define a morphism $\mathcal{T}_\infty \rightarrow \mathbb{A}^1(R) = \operatorname{Spec} R[t]$. This morphism makes \mathcal{T}_∞ an $R[t]$ -scheme. We can regard t as a regular function on \mathcal{T}_∞ ; its restriction to $U_{i+1/2}$ is the monomial $t = e^{(0,1)}$.

Explicit description of the gluing maps. We now describe the gluing construction of \mathcal{T}_∞ in scheme-theoretic terms, without the toric language. The exposition follows Deligne–Rapoport’s (*op. cit.*); we recall it for convenience. The lattice points $\sigma_{i+1/2}^\vee \cap \mathbb{Z}^2$ are just the $\mathbb{Z}_{\geq 0}$ -linear combinations of $(1, -i)$ and $(-1, i+1)$. In $R[\sigma_{i+1/2}^\vee]$, let $Y_{i+1} = e^{(-1, i+1)}$ and $X_i = e^{(1, -i)}$. Then

$$U_{i+1/2} = \operatorname{Spec} \frac{R[t][X_i, Y_{i+1}]}{(X_i Y_{i+1} - t)}.$$

Now let $V_i = U_{i-1/2} \cap U_{i+1/2}$ as subsets of \mathcal{T}_∞ . In abstract terms, we have

$$\begin{aligned} V_i &\subset U_{i+1/2}, & V_i &= U_{i+1/2}[X_i^{-1}] \xrightarrow{X_i \mapsto X_i} \operatorname{Spec} R[t][X_i, X_i^{-1}] \\ V_i &\subset U_{i-1/2}, & V_i &= U_{i-1/2}[Y_i^{-1}] \xrightarrow{Y_i \mapsto Y_i} \operatorname{Spec} R[t][Y_i^{-1}, Y_i], \end{aligned}$$

and these two descriptions of V_i are matched up by putting $X_i Y_i = 1$. The scheme obtained from the union of the open sets $U_{i+1/2}$ by gluing $U_{i-1/2}$ to $U_{i+1/2}$ is \mathcal{T}_∞ . No additional gluing is required because when $j - i > 1$ one has $U_{i-1/2} \cap U_{j-1/2} \subset U_{i-1/2} \cap U_{i+1/2}$.

In the scheme $\mathcal{T}_\infty[t^{-1}] \rightarrow \operatorname{Spec} R[t, t^{-1}]$, all the open sets $U_{i+1/2}[t^{-1}]$ are identified with one another. Thus one has isomorphisms $\mathcal{T}_\infty[t^{-1}] \cong U_{1/2}[t^{-1}] \cong \mathbb{G}_m(R[t, t^{-1}])$.

The fiber $\mathcal{T}_\infty|_{t=0} = \mathcal{T}_\infty \times_{R[t]} R$ is the union of the subsets

$$U_{i+1/2}|_{t=0} = \operatorname{Spec} R[X_i, Y_{i+1}]/(X_i Y_{i+1}).$$

The sets $U_{i+1/2}|_{t=0}$ are disjoint when the indices i and j are not adjacent ($|i - j| > 1$). Moreover, $U_{i-1/2}|_{t=0}$ is glued to $U_{i+1/2}|_{t=0}$ along $V_i|_{t=0} = R[X_i, X_i^{-1}] = R[Y_i, Y_i^{-1}]$. Thus $\mathcal{T}_\infty|_{t=0}$ is an infinite chain of \mathbb{P}^1 's, say $\mathbb{P}_i^1 \subset U_{i-1/2}|_{t=0} \cap U_{i+1/2}|_{t=0}$.

The group \mathbb{Z} acts on \mathcal{T}_∞ covering the trivial action on $\operatorname{Spec} R[t]$. The action is induced by a \mathbb{Z} -action on \mathbb{Q}^2 preserving the fan F , given by $n \cdot (x_1, x_2) = (x_1 + n, x_2)$. One has $n \cdot U_{i+1/2} = U_{i+n+1/2}$; the action identifies X_i with X_{i+n} and Y_i with Y_{i+n} .

10.2 Quotients along the thickened zero-fiber

We wish to form quotients of \mathcal{T}_∞ by the group $d\mathbb{Z} \subset \mathbb{Z}$, but find ourselves unable to do so in the category of $R[t]$ -schemes. One can, however, construct quotients by $d\mathbb{Z}$ of thickened neighborhoods of the zero-fiber, i.e. of $\mathcal{T}_\infty|_{t^k=0}$, by virtue of the following fact:

■ The $U_{i+1/2}|_{t^k=0}$ form a chain: they are disjoint when i and j are not adjacent.

Indeed, when $j - i > 1$, the set $U_{i-1/2} \cap U_{j-1/2} \subset \mathcal{T}_\infty$ lies over $\operatorname{Spec} R[t, t^{-1}]$.

We write down a concrete model for the quotient $\mathcal{T}_\infty|_{t^k=0}/(d\mathbb{Z})$, when $d > 1$. We take the affine schemes $U_{i+1/2}|_{t^k=0}$ for $i = 0, \dots, d-1$, and we identify the open subset V_d of $U_{d-1/2}|_{t^k=0}$ with the open subset V_0 of $U_{1/2}|_{t^k=0}$ in just the same way as we identify V_d with an open subset of $U_{d+1/2}$ to form \mathcal{T}_∞ . The result is a proper scheme $\mathcal{T}_k^d \rightarrow \operatorname{Spec} R[t]/(t^k)$ whose specialization to $t = 0$ is a cycle of d \mathbb{P}^1 's. We find it convenient to rename t as $q^{1/d}$; so we have $\mathcal{T}_k^d \rightarrow \operatorname{Spec} R[q^{1/d}]/(q^{k/d})$. These schemes form an inverse system in k ; passing to the inverse limit as $k \rightarrow \infty$, we obtain a proper formal scheme $\hat{\mathcal{T}}^d \rightarrow \operatorname{Spf} R[[q^{1/d}]]$.

There are étale quotient maps $\hat{\mathcal{T}}^{d_1 d_2}(q^{1/d_1}) \rightarrow \hat{\mathcal{T}}^{d_1}$, where by $\hat{\mathcal{T}}^{d_1 d_2}(q^{1/d_1})$ we mean the formal base-change from $\operatorname{Spf} R[[q^{1/d_1 d_2}]]$ to $\operatorname{Spf} R[[q^{1/d_1}]]$ given by $q^{1/d_1 d_2} \mapsto q^{1/d_1}$. The construction of $\hat{\mathcal{T}}^d$ does not work in quite the same way when $d = 1$, because one is then gluing V_0 to itself, and the Zariski-open cover by the $U_{i+1/2}$ becomes merely an étale cover. Nonetheless, we can define $\hat{\mathcal{T}}^1 \rightarrow \operatorname{Spf} R[[q]]$ as the étale quotient $\hat{\mathcal{T}}^2/(\mathbb{Z}/2)$. This quotient is the locally-ringed space whose functions are the $\mathbb{Z}/2$ -invariant functions of $\hat{\mathcal{T}}^2$. To see that it is a formal scheme, we need only note that $\hat{\mathcal{T}}^1$ can be covered by two $\mathbb{Z}/2$ -invariant formal-affine open sets, which is straightforward to check.

Notice also that $\mathcal{T}_\infty \rightarrow \operatorname{Spec} R[t]$ has sections $\sigma_{i+1/2}: \operatorname{Spec} R[t] \rightarrow U_{i+1/2}$, defined by $X_i = 1$ and $Y_{i+1} = t$. The \mathbb{Z} -action intertwines these sections, and consequently, $\hat{\mathcal{T}}^1 \rightarrow \operatorname{Spf} R[[q]]$ has a distinguished section σ .

10.2.1 Polarization

One now wants to polarize the formal scheme $\hat{\mathcal{T}}^1$, that is, to identify an ample line-bundle $\hat{\mathcal{L}} \rightarrow \hat{\mathcal{T}}^1$. A very ample power of $\hat{\mathcal{L}}$ will then define a projective embedding of $\hat{\mathcal{T}}^1$, and in doing so, will ‘algebraize’ $\hat{\mathcal{T}}^1$, refining it to a true scheme over $R[[q]]$.⁶

That ample line bundles algebraize formal schemes in a general setting is a theorem of Grothendieck, but here, as noted by Gross, it is quite concrete. We will find a line bundle $\hat{\mathcal{L}} \rightarrow \hat{\mathcal{T}}^1$, by which we mean a sequence of line-bundles $\mathcal{L}_k \rightarrow \mathcal{T}_\infty|_{t^k=0}/\mathbb{Z}$ and isomorphisms $\mathcal{L}_k|_{t^{k-1}=0} \cong \mathcal{L}_{k-1}$. We will find a basis $\{\theta_{3,m/3}\}_{m=0,1,2}$ for $H^0(\hat{\mathcal{T}}^1, \hat{\mathcal{L}}^{\otimes 3})$ —by this we mean bases $\{\theta_{3,m/3}^{(k)}\}_{m=0,1,2}$ for $H^0(\mathcal{T}_\infty|_{t^k=0}/\mathbb{Z}, (\mathcal{L}_k)^{\otimes 3})$, carried one to another by the maps in the inverse system—defining a plane embedding of $\hat{\mathcal{T}}^1$. In this way we will see that $\mathcal{T}_\infty|_{t^k=0}/\mathbb{Z}$ is cut out from $\mathbb{P}^2(R[t]/(t^k))$ by an equation which reduces modulo t^{k-1} to that cutting out $\mathcal{T}_\infty|_{t^{k-1}=0}/\mathbb{Z}$. Passing to the limit, we see that $\hat{\mathcal{T}}^1$ is cut out by an equation from $\mathbb{P}^2(R[[q]])$, and hence arises from a projective scheme \mathcal{T} .

As Gross explains, one obtains a line bundle $\mathcal{L} = \mathcal{L}_\Delta \rightarrow \mathcal{T}^\infty$ by observing that F is the fan dual to an unbounded convex polygon $\Delta \subset \mathbb{Q}^2$, namely, the convex hull in \mathbb{Q}^2 of a sequence of points $w_{i+1/2} \in \mathbb{Z}^2$, $i \in \mathbb{Z}$, given by $w_{1/2} = (0, 0)$ and $w_{i-1/2} - w_{i+1/2} = (1, -i)$ (Figure 10.2.1).




Figure 3: The convex polygon Δ

The cone $\sigma_{i+1/2}^\vee$ is the tangent wedge $T_{i+1/2}$ at $w_{i+1/2}$. For each $j \in \mathbb{Z} + \frac{1}{2}$, we define $\mathcal{O}_\Delta|_{U_j}$ to be the free \mathcal{O}_{U_j} -module $\mathcal{O}_{U_j} \cdot z^{w_j}$ on a generator z^{w_j} . Here z is a formal symbol.

⁶An analogous situation more familiar to geometers accustomed to working over \mathbb{C} is that a projective embedding of a complex manifold makes this analytic object algebraic.

We assemble the $\mathcal{O}_\Delta|_{U_j}$ into an invertible sheaf $\mathcal{L} = \mathcal{O}_\Delta \rightarrow \mathcal{T}_\infty$ by declaring that the transition function from $\mathcal{L}|_{U_{i+1/2}}$ to $\mathcal{L}|_{U_{i-1/2}}$ is multiplication by $z^{w_{1/2-i}-w_{1/2+i}}$. Each lattice point $\lambda \in \Delta \cap \mathbb{Z}^2$ defines a section $z^\lambda \in H^0(\mathcal{L})$: over $U_{i+1/2}$, we view λ as a point in $(w_{i+1/2} + \sigma_{i+1/2}^\vee) \cap \mathbb{Z}^2$, and so assign to it a section $w^\lambda|_{U_{i+1/2}} = X_i^a Y_{i+1}^b z^{i+1/2}$. These local sections agree on overlaps and so define a global section.

The line bundle \mathcal{L} is *ample*: this is a general feature of line bundles associated with lattice polytopes.

The \mathbb{Z} -action on \mathcal{T}_∞ lifts to a \mathbb{Z} -action on \mathcal{L} (we continue to follow Gross). Define $\tau \in SL_2(\mathbb{Z})$ by

$$\tau(x_1, x_2) = (x_1 + 1, x_1 + x_2).$$

We have $\tau(w_{i+1/2}) = w_{i-1/2}$, from which it is easy to see that $\tau(\Delta) = \Delta$. Thus τ generates an action of \mathbb{Z} on Δ . We lift the action of $m \in \mathbb{Z}$ to \mathcal{L} as follows. Take the map $\phi_m: U_j \rightarrow U_{m+j}$ which defines the \mathbb{Z} -action (namely, $\phi_m(X_j) = X_{j+1}$ and $\phi_m(Y_{j+1}) = Y_{j+m+1}$) and lift it to a map

$$\tilde{\phi}_m: \mathcal{L} \rightarrow \mathcal{L}, \quad z^\lambda \mapsto z^{\tau^m(\lambda)}, \quad \lambda \in \Delta \cap \mathbb{Z}^2.$$

(we have specified in particular the effect of $\tilde{\phi}_d$ on z^{w_j}).

As a result, \mathcal{L} descends to a line bundle $\hat{\mathcal{L}}$ over each scheme \mathcal{T}_k^d , and indeed over the formal scheme $\hat{\mathcal{T}}^d$. The projective embeddings $\hat{\mathcal{T}}^d \rightarrow \mathbb{P}^N(\mathcal{L}^{\otimes N})^\vee$ for $N \gg 0$ cut out $\hat{\mathcal{T}}^d$ as a projective scheme $\mathcal{T}^d \rightarrow \text{Spec } \mathbb{Z}[[q]]^{1/d}$, lifting the formal scheme structure. This is even true for $d = 1$. The scheme $\mathcal{T} = \mathcal{T}^1 \rightarrow \text{Spec } \mathbb{Z}[[q]]$ is the Tate curve. It has its distinguished section σ .

The line bundle $\mathcal{L} \rightarrow \mathcal{T}$ has a global section $\theta := \sum_{k \in \mathbb{Z}} z^{w_k}$. This formula is to be interpreted initially on $U_{i+1/2}|_{t^k=0}$, where it is a finite sum. It descends to a section of \mathcal{L} over $\mathcal{T}_\infty|_{t^k=0}/(d\mathbb{Z})$, and thereby a section over $\hat{\mathcal{T}}^d$ for each d . It is instructive to write θ in the open set V_0 in the following forms:

$$\theta|_{V_0} = \sum_{k \in \mathbb{Z}} q^{k(k-1)/2} (z^{(1,0)})^k = \sum_{k \in \mathbb{Z}} (-1)^k q^{k(k-1)/2} \zeta^k,$$

where $\zeta = -z^{(1,0)}$. Up to a factor of $iq^{1/4}\zeta$, $\theta(q, \zeta)$ is exactly the Fourier expansion for the classical theta-function $\vartheta_{1,1}$, written in terms of $q = e^{2\pi i\tau}$ and $\zeta = e^{2\pi i x}$, where x is the coordinate on $\mathbb{C}/\langle 1, \tau \rangle$.

Observe (i) that $\theta|_{V_0 \cap \{q=0\}}$ vanishes precisely where $\zeta = 1$, whence \mathcal{L} has degree 1 on the geometric fibers of $\mathcal{T} \rightarrow \text{Spec } R[[q]]$; (ii) that $\theta(q, \zeta^{-1}) = -\zeta^{-1}\theta(q, \zeta)$, so θ vanishes where $\zeta = 1$; and hence (iii) that θ vanishes precisely where $\zeta = 1$, i.e. along σ . Thus $\mathcal{L} \cong \mathcal{O}(\sigma)$. Using Riemann–Roch, we see that $\mathcal{L}^{\otimes 3} = \mathcal{O}(3\sigma)$ is very ample relative to the morphism $\mathcal{T} \rightarrow \text{Spec } R[[q]]$, and hence embeds \mathcal{T} as a Weierstrass cubic

in $\mathbb{P}^2(R[[q]])$. Thus \mathcal{T} becomes a Weierstrass curve, with a canonical differential ω . Moreover, this perspective makes clear that \mathcal{T} coincides, as a Weierstrass curve, with the curve described in the introduction to this paper. (However, we shall not review the derivation of the Fourier expansions of $a_6(q)$ and $a_6(q)$).

10.3 Lattice-points and multiplication of theta-functions

We have a canonical \mathbb{Z} -basis $\{z^\lambda : \lambda \in \mathbb{Z}^2 \cap N\Delta\}$ for $H^0(\mathcal{T}_\infty; \mathcal{L}^{\otimes N})$. The action of q is given by $q \cdot z^\lambda = z^{\lambda+(0,1)}$. To get a basis over $\mathbb{Z}[[q]]$, it suffices to consider $\{z^\lambda : \lambda \in \mathbb{Z}^2 \cap \partial(N\Delta)\}$, since we then obtain the remaining lattice points by multiplying by powers of q . There is also a canonical $\mathbb{Z}[[q]]$ -basis β_N for the \mathbb{Z} -invariant part $H^0(\mathcal{T}_\infty; \mathcal{L}^{\otimes N})^\mathbb{Z} = H^0(\mathcal{T}; \mathcal{L}^{\otimes N})$,

$$\beta_N = \{\theta_{N,p} : p \in C_N\},$$

where p runs over the cyclic group $C_N := (\frac{1}{N}\mathbb{Z}) / \mathbb{Z}$. To define the ‘theta-functions’ $\theta_{N,p}$, let $\phi : \mathbb{Q} \rightarrow \mathbb{Q}$ be the piecewise-linear function whose graph is the boundary of Δ . Define automorphisms τ_N of the dilated polygon $N\Delta$, sending vertex Nw_j to the adjacent vertex Nw_{j+1} , by

$$\tau_N(x_1, x_2) = (x_1 + N, x_1 + x_2).$$

Then $(Np, N\phi(p))$ lies on the boundary of $N\Delta$, and hence so does $\tau_N^k(Np, N\phi(p))$ for each $k \in \mathbb{Z}$. Put

$$\theta_{N,p} = \sum_{k \in \mathbb{Z}} z^{\tau_N^k(Np, N\phi(p))}.$$


Again, the formula for $\theta_{N,p}$ makes sense, as a finite formal sum, on $U_{i+1/2}|_{t^k=0}$, hence on each $\hat{\mathcal{T}}^d$, and so finally on \mathcal{T} . It is clear that β_N is a basis for $H^0(\mathcal{T}, \mathcal{L})$. One has $\theta_{1,0} = \theta$.

Explicit multiplication rules for classical theta-functions are standard. Whilst these theta-functions are not precisely identical to the classical theta functions which give canonical bases for $H^0(\mathcal{O}(Np))$, they do obey a very similar multiplication rule [23]:

$$(50) \quad \theta_{n_1, p_1} \cdot \theta_{n_2, p_2} = \sum_{j \in \mathbb{Z}} q^{\lambda(p_1, p_2 + j)} \theta_{n_1 + n_2, E(p_1, p_2 + j)}.$$

Here $E(p_1, p_2)$ is the weighted average with respect to a distribution determined by n_1 and n_2 ,

$$E(a, b) = \frac{n_1 a + n_2 b}{n_1 + n_2},$$



and

$$(51) \quad \begin{aligned} \lambda(p_1, p_2) &= n_1 \phi(p_1) + n_2 \phi(p_2) - (n_1 + n_2) \phi(E(p_1, p_2)) \\ &= (n_1 + n_2) \{E(\phi(p_1), \phi(p_2)) - \phi(E(p_1, p_2))\}. \end{aligned}$$

Three points about ϕ are noteworthy:

- For any convex function f , the quantity $\Delta_f(p_1, p_2) := E(f(p_1), f(p_2)) - f(E(p_1, p_2))$ is non-negative, by Jensen's inequality. One has $\lambda = (n_1 + n_2) \Delta_\phi$.
- ϕ is a piecewise-linear approximation to the quadratic function

$$(52) \quad \psi(x) = \frac{1}{2}x(x - 1);$$

indeed, $\phi(n) = \psi(n)$ for $n \in \mathbb{Z}$, and ϕ is affine-linear on intervals $[n, n + 1]$.

- The expression

$$(53) \quad \begin{aligned} a(p_1, p_2) &:= (n_1 + n_2) \Delta_\phi(p_1, p_2) \\ &= n_1 \psi(p_1) + n_2 \psi(p_2) - (n_1 + n_2) \psi(E(p_1, p_2)) \\ &= (n_1 + n_2) \{E(\psi(p_1), \psi(p_2)) - \psi(E(p_1, p_2))\}, \end{aligned}$$

defined analogously to λ , is the area of the triangle $T(p_1, p_2)$ with vertices

$$(p_1, 0), \quad (p_2, -n_1(p_2 - p_1)), \quad (E(p_1, p_2), 0)$$

(see Figure 10.3).

The last fact is key to Gross’s derivation of ‘classical’ HMS [23, section 8.4.2].⁷ For a fixed n , the theta-functions $\theta_{n,p}$ are *almost* mirror to the intersection points of L_0 and $L_{(1,-n)}$. We say “almost” because correction factors are required, since it is $\lambda(p_1, p_2)$ and not $a(p_1, p_2)$ which appears in the product rule for theta-functions.

It turns out that λ has a similar geometric interpretation, one which will be equally central in our derivation of arithmetic HMS. We point out that whilst subsection 10.1 was a review, and the material of subsection 10.3 has so far also been standard, this interpretation is to our knowledge original:

Proposition 10.1 *Fix $\epsilon > 0$, and say a point in \mathbb{R}^2 is a perturbed lattice point if it is congruent to (ϵ, ϵ) modulo \mathbb{Z}^2 . Then, if $\epsilon \ll (n_1 + n_2)^{-1}$, the number of perturbed lattice points inside $T(p_1, p_2)$ is equal to $\lambda(p_1, p_2)$.*

Proof It is possible to do this calculation directly—a tedious and unilluminating computation that occupied almost three pages in the preliminary version of this article.⁸ We present here a more sophisticated approach based on Brion’s formula [?]. This formula, which gives a generating function for lattice points inside rational polytopes, was discovered using toric geometry. Indeed, it has an interpretation in terms of sections of an ample line bundle over a singular toric variety—in this case, a weighted projective plane—but it also has direct combinatorial proofs. The case at hand, that of a simplex with rational edge directions, has a particularly straightforward proof [?].

Write $n_3 = n_1 + n_2$ and $p_3 = E(p_1, p_2)$. Suppose, for $i \in \{1, 2, 3\}$, that $p_i = q_i + r_i/n_i$, where $q_i, r_i \in \mathbb{Z}$ and $0 \leq r_i < n_i$. Then

$$\lambda(p_1, p_2) = \frac{n_1}{2}q_1(q_1 - 1) + \frac{n_2}{2}q_2(q_2 - 1) - \frac{n_3}{2}q_3(q_3 - 1) + r_1q_1 + r_2q_2 - r_3q_3.$$

We wish to show that Λ , the number of perturbed lattice points in $T(p_1, p_2)$, equals $\lambda(p_1, p_2)$.

The generating function for the set of perturbed lattice points in a set $C \subset \mathbb{R}^2$ is by definition the 2-variable Laurent series

$$F_C(x, y) = \sum_{(a,b) \in C \cap (\mathbb{Z} + \epsilon)^2} x^a y^b.$$

Let $F = F_{T(p_1, p_2)} \in \mathbb{Z}[x^{\pm 1}, y^{\pm 1}]$. Let C_1 be the convex cone (wedge) in \mathbb{R}^2 whose two edges are the edges of $T(p_1, p_2)$ that meet at $v_1 := (p_1, 0)$ and which contains

⁷A similar multiplication rule, but for *classical* theta-functions, plays an analogous role in Polishchuk–Zaslow’s proof of cohomology-level mirror symmetry for elliptic curves [42]. For those theta-functions, $a(p_1, p_2)$ is the relevant exponent.

⁸This calculation can be found in the arXiv version 2 of this article.

$T(p_1, p_2)$. Similarly, let C_2 (resp. C_3) be the convex cone whose edges meet at $v_2 := (p_2, -n_1(p_2 - p_1))$ (resp. $v_3 := (E(p_1, p_2), 0)$) and contains $T(p_1, p_2)$. Then F_{C_i} is a rational function of x and y , i.e., an element F_i in the fraction field $\mathbb{Q}(x, y)$ of $\mathbb{Z}[x, y]$. Brion's formula says that

$$F = F_1 + F_2 + F_3.$$

Thus $\Lambda = F(1, 1) = F_1(1, 1) + F_2(1, 1) + F_3(1, 1)$.

To compute Λ , it suffices to consider the specialization $G = F|_{y=1} \in \mathbb{Z}[x^{\pm 1}]$, which is the sum of the three rational functions $G_i = F_i|_{y=1} \in \mathbb{Q}(x)$. Let $P_1 \subset C_1$ denote the parallelogram one of whose vertices is v_1 , and whose edge vectors $(1, 0)$ and $(1, -n_1)$ at v_1 are the primitive integral generators for the edges of C_1 . Similarly define P_2 and P_3 . Let g_i be the generating function for the x -coordinates of perturbed lattice points in P_i . Tiling the cones C_i by translates of P_i , one sees that

$$G_1(x) = g_1(x) \frac{1}{(1-x)^2}, \quad G_2(x) = g_2(x) \frac{x^2}{(1-x)^2}, \quad G_3(x) = g_3(x) \frac{-x}{(1-x)^2}.$$

It is straightforward to verify that, when ϵ is small compared to n_3 ,

$$\begin{aligned} g_1(x) &= (n_1 - r_1)x^{q_1+1} + r_1x^{q_1+2}, \\ g_2(x) &= (n_2 - r_2)x^{q_2-1} + r_2x^{q_1}, \\ g_3(x) &= (n_3 - r_3)x^{q_3} + r_3x^{q_3+1}. \end{aligned}$$

Let $h(x) = g_1(x) + x^2g_2(x) - xg_3(x)$. Then one checks that $h(1) = 0$ and (using the relation $r_1 + r_2 - r_3 = n_3q_3 - n_1q_1 - n_2q_2$) that $h'(1) = 0$. Hence $\Lambda = \frac{1}{2}h''(1)$, and the latter, by a short calculation, equals $\lambda(p_1, p_2)$. \square

10.4 Homogeneous coordinate rings

We earlier discussed the affine coordinate ring $\varinjlim_N H^0(\mathcal{O}(Np))$. To work with this ring, one must understand the direct system as well as the multiplication of sections of powers of $\mathcal{O}(p)$. On any Weierstrass curve $C \rightarrow \operatorname{Spec} R$, the divisor $D = 3p$ is very ample. It defines a homogeneous coordinate ring $R_C^0 = \bigoplus_{N \geq 0} H^0(\mathcal{O}(ND))$, a graded R -algebra whose isomorphism determines the curve. The homogeneous coordinate ring has the advantage over the affine one that it is directly determined by the composition maps in $\operatorname{Perf} C$. For a Weierstrass curve (C, σ, ω) over $\operatorname{Spec} S$, let R_C^* denote the extended homogeneous coordinate ring $\bigoplus_{n \geq 0} H^*(\mathcal{O}(nD))$. This bigraded ring differs from R_C^0 only by the presence of the summand $H^1(\mathcal{O})$. This summand has its trace map $\operatorname{tr}_\omega: H^1(\mathcal{O}) \rightarrow S$. The ring R_C^* determines C , and the trace map then determines the differential ω .

The truncated ring $tR_C^* = R_C^*/I$, where $I = \bigoplus_{N \geq 3} H^0(\mathcal{O}(ND))$, with its bigrading and trace $H^1(\mathcal{O}) \rightarrow S$, already determines (C, ω) , because from it a defining cubic equation (in Hesse form) can be read off in $H^0(\mathcal{O}(3D))$. However, the regular section σ might not be fully determined by tR_C^* , or even R_C^* , inasmuch as one could replace σ by a different regular section σ' such that the divisor $3\sigma'$ is linearly equivalent to 3σ .

Proof of Theorem A, clauses (i), (ii) Let C_{mirror} be an abstract Weierstrass curve over $\mathbb{Z}[[q]]$ whose category $\mathcal{B} = \mathcal{B}_{\text{mirror}}$ admits a quasi-isomorphism with \mathcal{A} such that the induced isomorphism $H\mathcal{A} \rightarrow H\mathcal{B}$ is the standard one. We must show that $C_{\text{mirror}} \cong \mathcal{T}$ as Weierstrass curves. We shall initially prove that they are isomorphic as curves with differential. For this it is sufficient to show that $tR_{C_{\text{mirror}}}^* \cong tR_{\mathcal{T}}^*$ by a trace-respecting bigraded ring isomorphism. We recall that $L_{(1,-n)}^\#$ denotes an oriented exact Lagrangian in T_0 of slope $-n$, equipped with its non-trivial double covering, and graded in such a way that $HF^*(L_0^\#, L_{(1,-n)}^\#) = HF^0(L_0^\#, L_{(1,-n)}^\#)$ for $n \neq 0$ (this is not quite a complete specification of the grading). We have

$$R_{C_{\text{mirror}}}^* \cong \bigoplus_{N \geq 0} HF^*(L_0, L_{(1,-3N)}) \cong HF^1(L_0, L_0) \oplus \bigoplus_{N \geq 0} HF^0(L_0, L_{(1,-3N)}).$$

We use Prop. 7.5 to set up the basepoint z and the 1-form θ on T_0 in such a way that the exact curve $L_{(1,-n)}$ is the image in $\mathbb{R}^2/\mathbb{Z}^2$ of a straight line through the origin in \mathbb{R}^2 —this for $0 \leq n \leq 9$. We realize the non-trivial double covering of $L_{(1,-n)}$ by selecting the point $\star = \star_{-n} = (\epsilon, -n\epsilon) \in L_{(1,-n)}$, where $0 < \epsilon \ll 1$, and declaring the double covering $\tilde{L}_{(1,-n)} \rightarrow L_{(1,-n)}$ to be trivial away from \star and to exchange the sheets at \star .

Take $1 \leq n \leq 9$. The differential on $CF^0(L_0^\#, L_{(1,-n)}^\#)$ is zero, since there are no immersed bigons bounding L_0 and $L_{(1,-n)}$. Thus we have a basis $B_n = \{x_{n,p} : p \in C_n\}$ for $HF^0(L_0^\#, L_{(1,-n)}^\#)$, where $C_n = \frac{1}{n}\mathbb{Z}/\mathbb{Z}$, and $x_{n,p} = [p, 0] \in \mathbb{R}^2/\mathbb{Z}^2 = T$. We have $L_{(1,-n)} = \tau^n(L_0)$ for $0 \leq n \leq 9$, where τ is the nearly-linear Dehn twist along L_∞ set up at (7.2.3); hence $HF(L_{(1,-n_1)}^\#, L_{(1,-n_1-n_2)}^\#)$ has basis $\tau^{n_1}(B_{n_2})$ when $0 < n_1 < n_1 + n_2 \leq 9$. We claim that the Floer product $HF(L_{(1,-n_1)}^\#, L_{(1,-n_1-n_2)}^\#) \otimes_{\mathbb{Z}[[q]]} HF(L_0^\#, L_{(1,-n_1)}^\#) \rightarrow HF(L_0, L_{(1,-n_1-n_2)})$ is given by

$$\tau^{n_1}(x_{n_2,p_2}) \cdot x_{n_1,p_1} = \sum_{j \in \mathbb{Z}} q^{\lambda(p_1,p_2+j)} x_{n_1+n_2,E(p_1,p_2+j)},$$

an expression formally identical to the multiplication rule for theta-functions (50); E and λ are as defined there. The contributions to $\tau^{n_1}(x_{n_2,p_2}) \cdot x_{n_1,p_1}$ are immersed triangles, the images of embedded triangles in \mathbb{R}^2 . The first vertex is a lift of x_{n_1,p_1} to \mathbb{R}^2 , which we assume is the point $(p_1, 0) \in \mathbb{R}^2$. The second vertex is a lift of

$\tau^{n_1}(x_{n_2, p_2}) = [p_2, n_1 p_2]$ which lies on the line of slope $-n_1$ through $(p_1, 0)$; thus it is of form $(p_2 + j, -n_1(p_2 + j - p_1))$ where $j \in \mathbb{Z}$. The third vertex is then at $(0, E(p_1, p_2 + j))$.

The contribution of the triangle just described is $\varepsilon q^{\lambda(p_1, p_2 + j)}$, where the exponent is the number of perturbed lattice points in the triangle, as computed in Proposition 10.1, and the sign $\varepsilon = \pm 1$ depends on the orientations, double coverings and the intersection signs of the corners determined by the orientations. A formula for ε is given in [49]; it is simplest when, as here, all the corners have intersection number -1 (and therefore even index for Floer cohomology). In that case, $\varepsilon = (-1)^s$, where s is the number of stars on the boundary. There are $\lceil E(p_1, p_2 + j) \rceil - \lceil p_1 \rceil$ stars on the L_0 boundary, $\lceil p_2 + j \rceil - \lceil p_1 \rceil$ on the $L_{(1, -n_1)}$ boundary, and $\lceil p_2 + j \rceil - \lceil E(p_1, p_2 + j) \rceil$ on the $L_{(1, -n_1 - n_2)}$ boundary, so $s = 2(j + \lceil p_2 \rceil - \lceil p_1 \rceil)$ and $\varepsilon = +1$. This justifies our claim.

Define a $\mathbb{Z}[[q]]$ -linear map $t\psi: tR_{C_{\text{mirror}}}^* \rightarrow tR_{\mathcal{T}}^*$ by linearly extending the assignment $t\psi(x_{3N, p}) = \theta_{3N, p}$ for $N \in \{1, 2, 3\}$ together with the canonical ring-isomorphism $t\psi: HF^*(L_0, L_0) \rightarrow H^*(\mathcal{O})$ defined by the Weierstrass differential ω . In view of (50) and Prop. 10.1, and the fact that the unit of $HF^*(L_0, L_0)$ is also the unit of $tR_{C_{\text{mirror}}}^*$, the map $t\psi$ is a bigraded ring isomorphism preserving the trace. Therefore, there is an isomorphism $\iota: (C_{\text{mirror}}, \sigma_{\text{mirror}}, \omega_{\text{mirror}}) \rightarrow (\mathcal{T}, \omega, \sigma')$, where ω is the standard Weierstrass differential on \mathcal{T} and σ' is some section of $\mathcal{T}^{\text{sm}} \rightarrow \text{Spec } \mathbb{Z}[[q]]$, the regular locus in \mathcal{T} .

Since \mathcal{T} is a *generalized elliptic curve* (see [17] or [13, Def. 2.1.4]), one has a homomorphism $\Gamma \rightarrow \text{Aut}(\mathcal{T}/\mathbb{Z}[[q]])$, from the group Γ of sections of $\mathcal{T}^{\text{sm}} \rightarrow \mathbb{Z}[[q]]$, to the automorphism group of \mathcal{T} fixing the base: sections act on \mathcal{T} by fiberwise translations. Consequently, all regular sections are related by automorphisms of \mathcal{T} . Hence $(C_{\text{mirror}}, \sigma_{\text{mirror}}, \omega_{\text{mirror}})$ is isomorphic to $(\mathcal{T}, \omega, \sigma)$.

At this point we have proved clauses (i) and (ii) of Theorem A apart from the uniqueness statement in (i). That is easily taken care of: \mathcal{A} split-generates $\text{tw}^\pi \mathcal{F}(T, z)$, and the functor ψ is determined, up to natural quasi-equivalence, by its effect on a full, split-generating subcategory. \square

The proof just given also gives another identification $C_{\text{mirror}}|_{q=0}$, which may be used as part of the argument for clauses (iii) and (iv) of Theorem A, which we have already proved.

Remark. In the argument above, we identified the mirror essentially by describing a canonical isomorphism of graded rings

$$\bigoplus_{n \geq 0} HF^0(L_0^\#, L_{(1, -3n)}^\#) \rightarrow S := \bigoplus_{n \geq 0} H^0(\mathcal{O}(3n\sigma)),$$

given, in degrees $3n \geq 3M$ where $L_{(-1,3n)}^\#$ is exact, by matching up the canonical bases: $x_{3n,p} \mapsto \theta_{3n,p}$. The argument did not depend on this assignment being the one arising from the functor ψ , but it is natural to expect that this is so. The assignment $\psi'(x_{3n,p}) = \theta_{3n,p}$, defined when $n \leq M$, respects products. Hence, taking $M \geq 3$, we see that ψ' extends to a ring isomorphism, because S is generated in degree 3 and subject only to a relation in degree 9. The functor ψ defines another such isomorphism. Thus $\psi' \circ \psi^{-1}$ extends to an automorphism of S . This must come from a Weierstrass automorphism α of \mathcal{T} . The only possibilities for α are the identity and the hyperelliptic involution, but we do not identify α here.

Proof of Theorem 1.3 We begin by clarifying how to construct Φ . We now consider $L_0^\#$ and $L_\infty^\#$ as objects of the ‘absolute’ Fukaya category $\mathcal{F}(T)$ over $\Lambda_{\mathbb{C}}$ —so their exactness plays no role. They form a full subcategory \mathcal{A}_T , which we identify with \mathcal{B}_E for some mirror Weierstrass curve $E \rightarrow \text{Spec } \Lambda_{\mathbb{C}}$. Thus we get a functor $\phi: \mathcal{F}(T) \rightarrow \text{tw vect}(E)$ inducing a quasi-equivalence $D^\pi \mathcal{F}(T_0) \rightarrow \text{Perf}(E)$.

We identify E via the homogeneous coordinate ring for its cubic embedding, just as we did for C_{mirror} . Thus we must count triangles according to their areas, not the number of lattice points they contain. By the argument of [42] (see also [6]), we identify the homogeneous coordinate ring of E as that of $\mathcal{T} \times (\Lambda_{\mathbb{C}})$, with the bases of intersection points B_n mapping to standard bases of θ -functions. Thus $E \cong \mathcal{T}(\Lambda_{\mathbb{C}}) = \mathcal{T} \times_{\mathbb{Z}[[q]]} \Lambda_{\mathbb{C}}$ as Weierstrass curves. Since E is smooth, every bounded coherent complex is quasi-isomorphic to a perfect complex, and so $\text{Perf}(E) \simeq D^b \text{Coh}(E)$. Thus we can think of Φ as a functor

$$\Phi: \mathcal{F}(T) \rightarrow \tilde{D}^b \text{Coh}(\mathcal{T}(\Lambda_{\mathbb{C}})).$$

We must now compare the two functors $\mathcal{F}(T, z) \rightarrow \text{Coh}(\mathcal{T}(\Lambda_{\mathbb{C}}))$: first $\Phi \circ e$, and second, ψ followed a base-change functor—call this composite Ψ . Both extend to functors defined on $\text{tw}^\pi \mathcal{F}(T, z)$, and it will suffice to show that these are homotopic. The inclusion $\mathcal{A} \rightarrow \mathcal{F}(T, z)$ induces a quasi-equivalence $\text{tw}^\pi \mathcal{A} \rightarrow \text{tw}^\pi \mathcal{F}(T, z)$, and the definitions of Φ and Ψ both depend on a choice of quasi-inverse $\text{tw}^\pi \mathcal{F}(T, z) \rightarrow \text{tw}^\pi \mathcal{A}$. We take this to be the same in both cases, so that the functors have the same effect on objects—otherwise the assertion of homotopy makes no sense.

Take \mathcal{A} to be the A_∞ -subcategory of $\mathcal{F}(T, z)$ with objects $L_0^\#$ and $L_\infty^\#$, and \mathcal{A}_T the corresponding subcategory of $\mathcal{F}(T)$.

The construction of the functor e depends on choices of functions f_L for each exact Lagrangian L such that $\theta_L = df_L$. Pick such a function f_{L_0} , and choose f_{L_∞} so that the unique intersection point $x \in L_0 \cap L_\infty$ has symplectic action 0. The endomorphism space $CF(L_0^\#, L_0^\#)$ is defined via a Hamiltonian image $\phi_H(L_0)$; endow this with the

function $f_{L_0} \circ \phi_H^{-1}$; similarly for L_∞ . Since e acts as the identity on objects, it defines a quasi-isomorphism $e: \mathcal{A}(\Lambda_{\mathbb{C}}) \rightarrow \mathcal{A}_T$. It induces the canonical isomorphism $H^* \mathcal{A} \rightarrow H^* \mathcal{A}_T$, because the relevant intersection points have action zero.

We have quasi-isomorphisms $e^* \mathcal{A}_T \rightarrow \mathcal{B}_{\mathcal{T}(\Lambda_{\mathbb{C}})}$ and $\mathcal{A} \rightarrow \mathcal{B}_{\mathcal{T}(\Lambda_{\mathbb{C}})}$, both inducing the canonical isomorphisms on cohomology. That is, we have two minimal A_∞ -structures on the graded algebra A , and identifications of both of them with $\mathcal{B}_{\mathcal{T}(\Lambda_{\mathbb{C}})}$. By Theorem C, these two identifications must be induced by an automorphism of $\mathcal{T}(\Lambda_{\mathbb{C}})$ respecting the Weierstrass data (σ, ω) . This can only be the identity map. The result follows. \square

References

- [1] M. Abouzaid, *On the Fukaya categories of higher genus surfaces*, Adv. Math. 217 (2008), no. 3, 1192–1235.
- [2] M. Abouzaid, *A geometric criterion for generating the Fukaya category*, Publ. Math. Inst. Hautes Etudes Sci. No. 112 (2010), 191–240.
- [3] M. Abouzaid, *A topological model for the Fukaya categories of plumbings*, J. Differential Geom. 87 (2011), no. 1, 1–80.
- [4] M. Abouzaid, K. Fukaya, Y.-G. Oh, H. Ohta, K. Ono, in preparation.
- [5] M. Abouzaid, P. Seidel, *An open string version of Viterbo functoriality*, Geom. Topol. 14 (2010), no. 2, 627–718.
- [6] M. Abouzaid, I. Smith, *Homological mirror symmetry for the 4-torus*, Duke Math. J. 152 (2010), no. 3, 373–440.
- [7] M. Aldi, E. Zaslow, *Seidel’s mirror map for abelian varieties*, Adv. Theor. Math. Phys. 10 (2006), no. 4, 591–602.
- [8] M. F. Atiyah, I. G. MacDonald, *Introduction to commutative algebra*, Addison-Wesley, 1969.
- [9] D. Ben-Zvi, J. Francis, D. Nadler, *Integral transforms and Drinfeld centers in derived algebraic geometry*, J. Amer. Math. Soc. 23 (2010), no. 4, 909–966.
- [10] A. Bondal, M. Kapranov, *Representable functors, Serre functors, and reconstructions* (Russian) Izv. Akad. Nauk SSSR Ser. Mat. 53 (1989), no. 6, 1183–1205, 1337; translation in Math. USSR-Izv. 35 (1990), no. 3, 519–541.
- [11] A. Bondal, D. Orlov, *Reconstruction of a variety from the derived category and groups of autoequivalences*, Compositio Math. 125 (2001), no. 3, 327–344.
- [12] A. Căldăraru, *The Mukai pairing, II: the Hochschild-Kostant-Rosenberg isomorphism*, Adv. Math. 194 (2005), no. 1, 34–66.
- [13] B. Conrad, *Grothendieck duality and base change*, Lecture Notes in Mathematics, 1750. Springer-Verlag, Berlin, 2000.

- [14] B. Conrad, *Arithmetic moduli of generalized elliptic curves*, J. Inst. Math. Jussieu 6 (2007), no. 2, 209–278.
- [15] D. Cox, S. Katz, *Mirror symmetry and algebraic geometry*, Mathematical Surveys and Monographs, 68. American Mathematical Society, Providence, RI, 1999.
- [16] P. Deligne, *Courbes elliptiques: formulaire d'après J. Tate*, Modular functions of one variable, IV (Proc. Internat. Summer School, Univ. Antwerp, Antwerp, 1972), pp. 53–73. Lecture Notes in Math., Vol. 476, Springer, Berlin, 1975.
- [17] Deligne, P.; Rapoport, M. *Les schémas de modules de courbes elliptiques*. Modular functions of one variable, II (Proc. Internat. Summer School, Univ. Antwerp, Antwerp, 1972), pp. 143–316. Lecture Notes in Math., Vol. 349, Springer, Berlin, 1973.
- [18] B. Farb, D. Margalit, *A primer on mapping class groups*, Princeton University Press, Princeton, N.J., 2011.
- [19] R. Fiset, *The A-infinity algebra of an Elliptic Curve and the j-invariant*, 2011 preprint, arXiv:1111.6303.
- [20] K. Fukaya, *Cyclic symmetry and adic convergence in Lagrangian Floer theory*, Kyoto J. Math. 50 (2010), no. 3, 521–590.
- [21] S. Ganatra, *Symplectic Cohomology and Duality for the Wrapped Fukaya Category*, MIT Ph.D. thesis, 2012.
- [22] M. Gerstenhaber, *The cohomology structure of an associative ring*, Ann. of Math. (2) 78 (1963), 267–288.
- [23] M. Gross, chapter 8 of *Dirichlet branes and mirror symmetry*, Clay Mathematics Monographs, 4. American Mathematical Society, Providence, RI; Clay Mathematics Institute, Cambridge, MA, 2009.
- [24] R. Hartshorne, *Algebraic geometry*, Graduate Texts in Mathematics, No. 52. Springer-Verlag, New York–Heidelberg, 1977.
- [25] R. Hartshorne, *Deformation theory*, Graduate Texts in Mathematics, 257. Springer, New York, 2010.
- [26] H. Hida, *Geometric modular forms and elliptic curves*, 2ed. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2012.
- [27] P. Horja, *Derived category automorphisms from mirror symmetry*, Duke Math. J. 127 (2005), no. 1, 1–34.
- [28] T. Kadeishvili, *The algebraic structure in the homology of an $A(\infty)$ -algebra* (Russian), Soobshch. Akad. Nauk Gruz. SSR 108 (1982), no. 2, 249–252 (1983).
- [29] N. Katz, B. Mazur, *Arithmetic moduli of elliptic curves*, Annals of Mathematics Studies, 108. Princeton University Press, Princeton, NJ, 1985.
- [30] B. Keller, *Deriving DG categories*, Ann. Sci. Ecole Norm. Sup. (4) 27 (1994), no. 1, 63–102.
- [31] M. Kontsevich, *Homological algebra of mirror symmetry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), 120–139, Birkhäuser, Basel, 1995.

- [32] M. Kontsevich, appendix to C. Frønsdal, *Quantization on curves*, Lett. Math. Phys. 79 (2007), no. 2, 109–129.
- [33] M. Kontsevich, Y. Soibelman, *Notes on A_∞ -algebras, A_∞ -categories and non-commutative geometry*, in *Homological mirror symmetry*, Lecture Notes in Phys., 757, Springer, Berlin, 2009.
- [34] M. Kontsevich, *Symplectic geometry of homological algebra*, 2009 lecture notes, available at www.ihes.fr/~maxim/TEXTS/Symplectic_AT2009.pdf
- [35] Y. Lekili, T. Perutz, *Fukaya categories of the torus and Dehn surgery*, Proc. Natl. Acad. Sci. USA 108 (2011), no. 20, 8106–8113.
- [36] V. Lunts, D. Orlov, *Uniqueness of enhancement for triangulated categories*, J. Amer. Math. Soc. 23 (2010), no. 3, 853–908.
- [37] M. Manetti, *Deformation theory via differential graded Lie algebras*, Algebraic Geometry Seminars, 1998–1999 (Italian) (Pisa), 21–48, Scuola Norm. Sup., Pisa, 1999.
- [38] M. Markl, *Transferring A_∞ (strongly homotopy associative) structures*, Rend. Circ. Mat. Palermo (2) Suppl. No. 79 (2006), 139–151.
- [39] A. Neeman, *The Grothendieck duality theorem via Bousfield’s techniques and Brown representability*, J. Amer. Math. Soc. 9 (1996), 205–236.
- [40] A. Polishchuk, *Massey and Fukaya products on elliptic curves*, Adv. Theor. Math. Phys. 4 (2000), no. 6, 1187–1207.
- [41] A. Polishchuk, *Extensions of homogeneous coordinate rings to A_∞ -algebras*, Homology Homotopy Appl. 5 (2003), no. 1, 407–421.
- [42] A. Polishchuk and E. Zaslow, *Categorical mirror symmetry: the elliptic curve*, Adv. Theor. Math. Phys. 2 (1998), no. 2, 443–470.
- [43] D. Quillen, *On the (co-) homology of commutative rings*, in *Applications of categorical algebra* (Proc. Sympos. Pure Math., Vol. XVII, New York, 1968) pp. 65–87. Amer. Math. Soc., Providence, R.I., 1970.
- [44] A.N. Rudakov, *Exceptional collections, mutations and helices*, in *Helices and vector bundles*, London Math. Soc. Lecture Note Ser., 148, Cambridge Univ. Press, Cambridge, 1990.
- [45] M. Schwarz, *Cohomology operations from S^1 -cobordisms in Floer homology*, Ph.D. thesis, ETH Zürich, 1995.
- [46] P. Seidel, *Fukaya categories and deformations*, Proceedings of the International Congress of Mathematicians, Vol. II (Beijing, 2002), 351–360, Higher Ed. Press, Beijing, 2002.
- [47] P. Seidel, *Homological mirror symmetry for the quartic surface*, 2003 preprint, arXiv:math/0310414v1; revised text arXiv:math/0310414v2 (2011) submitted for publication.
- [48] P. Seidel, *A biased view of symplectic cohomology*, Current developments in mathematics, 2006, 211–253, Int. Press, Somerville, MA, 2008.

- [49] P. Seidel, *Homological mirror symmetry for the genus two curve*, J. Algebraic Geom. 20 (2011), no. 4, 727–769.
- [50] P. Seidel, *Fukaya categories and Picard–Lefschetz theory*, Zürich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2008.
- [51] P. Seidel, R. Thomas, *Braid group actions on derived categories of coherent sheaves*, Duke Math. J. 108 (2001), no. 1, 37–108.
- [52] P. Seidel, *Abstract analogues of flux as symplectic invariants*, 2011 preprint, arXiv:1108.0394v1.
- [53] J.-P. Serre, *Local algebra*, Springer Monographs in Mathematics. Springer–Verlag, Berlin, 2000.
- [54] N. Sheridan, *Homological Mirror Symmetry for Calabi–Yau hypersurfaces in projective space*, 2011 preprint, arXiv:1111.0632v2.
- [55] J. Silverman, *The arithmetic of elliptic curves*, 2ed. Graduate Texts in Mathematics, 106. Springer, Dordrecht, 2009.
- [56] R. Swan, *Hochschild cohomology of quasiprojective schemes*, J. Pure Appl. Algebra 110 (1996), no. 1, 57–80.
- [57] J. Tate, *A review of non-Archimedean elliptic functions*, Elliptic curves, modular forms, & Fermat’s last theorem (Hong Kong, 1993), 162–184, Ser. Number Theory, I, Int. Press, Cambridge, MA, 1995.
- [58] R. Thomason, *The classification of triangulated subcategories*, Compositio Math. 105 (1997), no. 1, 1–27.
- [59] R. W. Thomason, T. Trobaugh, *Higher algebraic K-theory of schemes and of derived categories*, in *The Grothendieck Festschrift*, Vol. III, Modern Birkhäuser Classics Vol. 88, 2007, pp 247–435.
- [60] B. Toën, *The homotopy theory of dg-categories and derived Morita theory*, Invent. Math. 167 (2007), no. 3, 615–667.
- [61] C. Viterbo, *Functors and computations in Floer homology with applications. I*, Geom. Funct. Anal. 9 (1999), no. 5, 985–1033.
- [62] C. Weibel, *An introduction to homological algebra*, Cambridge Studies in Advanced Mathematics, 38. Cambridge University Press, Cambridge, 1994.
- [63] C. Weibel, *Algebraic K-theory*, textbook draft, ‘beta’ version dated September 2, 2012, downloaded from <http://www.math.rutgers.edu/~weibel/Kbook.html>.
- [64] A. Yekutieli, *The continuous Hochschild cochain complex of a scheme*, Canad. J. Math. Vol. 54 (2002), no. 6, 1319–1337.
- [65] E. Zaslow, *Seidel’s mirror map for the torus*, Adv. Theor. Math. Phys. 9 (2005), no. 6, 999–1006.

Author affiliations

YL: University of Cambridge

TP: University of Texas at Austin

y1319@cam.ac.uk, perutz@math.utexas.edu