

Holomorphic curves in exploded manifolds

Virtual fundamental class

BRETT PARKER

We define Gromov–Witten invariants of exploded manifolds. The technical heart of this paper is a construction of a virtual fundamental class $[\mathcal{K}]$ of any Kuranishi category \mathcal{K} (which is a simplified, more general version of an embedded Kuranishi structure). We also show how to integrate differential forms over $[\mathcal{K}]$ to obtain numerical invariants, and push forward such differential forms over suitable maps. We show that such invariants are independent of any choices, and are compatible with pullbacks, products and tropical completion of Kuranishi categories.

In the case of a compact symplectic manifold, this gives an alternative construction of Gromov–Witten invariants, including gravitational descendants.

53D45

1. Introduction	1877
2. Constructing sections of sheaves on K –categories	1883
3. Kuranishi categories	1890
4. Construction of $[\mathcal{K}]$	1908
5. Representing Gromov–Witten invariants using de Rham cohomology	1914
6. Weak products of Kuranishi categories	1936
7. Tropical completion	1941
8. Construction of Gromov–Witten invariants	1952
References	1958

1 Introduction

We construct Gromov–Witten invariants of exploded manifolds¹ using embedded Kuranishi structures, which we defined and constructed in [26]. As smooth manifolds are a full subcategory of the category of exploded manifolds, this paper gives an alternative construction — and proof of the invariance of — all descendant Gromov–Witten invariants of any compact symplectic manifold.

¹Definitions for exploded manifolds can be found in our [24]. For a short introduction to exploded manifolds, see our [29].

Our construction also provides Gromov–Witten invariants relative normal-crossing divisors. Given a Kähler manifold with normal-crossing divisors, we can explode it, then construct Gromov–Witten invariants of the resulting exploded manifold. There is a similar construction for symplectic manifolds with normal-crossing symplectic divisors; however, we must be more careful here: If we take Definition 2.1 from Tehrani, McLean and Zinger [35], then a simple crossing symplectic divisor is a finite transverse collection of closed, codimension-2, symplectic submanifolds whose intersection is symplectic, with symplectic orientation agreeing with the intersection orientation. After deforming the symplectic structure in such a configuration, it admits a contractible choice of $\bar{\partial}$ -log compatible almost complex structure as in our [25, Section 14], and we can then apply the explosion functor to get an exploded manifold and define relative Gromov–Witten invariants. The isotopy class of the $\bar{\partial}$ -log compatible almost complex structure only depends on the isotopy class of the simple crossing symplectic divisor, so again we may define relative Gromov–Witten invariants in this setting as the Gromov–Witten invariants of the associated exploded manifold.

In some cases, the Gromov–Witten invariants defined in this paper coincide with previously defined Gromov–Witten invariants of symplectic manifolds, defined by Fukaya and Ono [10], McDuff [17], Ruan [33], Liu and Tian [16], Siebert [34] and Li and Tian [15]. I expect that in the algebraic case, the definition of Gromov–Witten invariants given here will agree with the algebraic definition given by Behrend and Fantechi [3], and also log Gromov–Witten invariants defined by Gross and Siebert [11] and Abramovich and Chen [5; 1].

The technical heart of this paper is a construction of a virtual fundamental class associated to an embedded Kuranishi structure. As well as the original work of Fukaya and Ono [10] and Li and Tian [15], there has been much recent work on constructing virtual fundamental cycles from various kinds of Kuranishi structures; see Fukaya, Oh, Ohta and Ono [8], McDuff and Wehrheim [19; 21; 22; 20], Pardon [23], Chen, Li and Wang [4] and Joyce [12; 13]. In the case of smooth manifolds, this paper provides an alternative construction of a virtual fundamental cycle from an embedded Kuranishi structure (constructed in [26]).

An embedded Kuranishi structure consists of a collection of compatible charts (\hat{f}, G, V) where \hat{f} is a family of curves, G is a group of automorphisms of \hat{f} and V is some G -equivariant obstruction bundle over \hat{f} with a natural G -invariant section $\bar{\partial}\hat{f}$. The adjective “embedded” indicates that \hat{f} is a family of curves in some moduli stack $\mathcal{M}_{\bullet}^{\text{st}}$, V is defined on a neighborhood of \hat{f} in $\mathcal{M}_{\bullet}^{\text{st}}$ as a finite-rank subbundle of some natural

obstruction bundle with a natural section $\bar{\partial}$, and \hat{f}/G represents the substack $\bar{\partial}^{-1}V$. We also require compatibility of Kuranishi charts — (\hat{f}, G, V) is compatible with (\hat{f}', G', V') if V is a subbundle of V' (or vice versa) on their common domain of definition. So, we have a kind of transition map between charts

$$\hat{f} \leftarrow \hat{f} \times_{\mathcal{M}_{\bullet}^{\text{st}}} \hat{f}' \rightarrow \hat{f}',$$

where the leftward arrow is a G –equivariant G' –fold cover of an open subfamily of \hat{f} and the rightward arrow is a G' –equivariant, G –fold cover of the subfamily $(\bar{\partial}\hat{f}')^{-1}V \subset \hat{f}'$.

We shall often want to construct a global section of some sheaf over our Kuranishi charts. (For example, we might want to perturb $\bar{\partial}$ to be transverse to 0, or construct a smooth function, or use the Chern–Weil construction to obtain the Chern class of a vectorbundle.) To present a simple and unified construction of such global sections, we introduce the notion of a K –category in Section 2: a K –category is obtained from an embedded Kuranishi structure by discarding all information apart from the charts \hat{f}/G and their transition maps. In Proposition 2.3, we prove that, at the expense of shrinking the charts in a K –category a little, we can construct a global section of any sheaf satisfying three basic axioms, called “patching”, “extension” and “averaging”.

Proposition 2.3 serves well to construct most of our global sections, but there is one important exception: the sheaf of transverse perturbations of the $\bar{\partial}$ equation does not satisfy the averaging axiom. In Section 3.3 we construct a weighted branched cover I of a Kuranishi category as a “sheaf” of finite measure spaces and define a weighted branched section of a sheaf S to be a natural transformation $I \rightarrow S$. We then prove that the corresponding sheaf of weighted branched sections S^I satisfies the patching, extension and averaging axioms if S satisfies the patching and extension axioms. We can then use Proposition 2.3 to construct global sections of S^I .

By the end of Section 4, we construct the virtual fundamental class $[\mathcal{M}_{\bullet}]$ of the moduli stack of holomorphic curves — where \bullet indicates choices such as specifying the genus, number of marked points and homology class of the curves under study. This virtual fundamental class $[\mathcal{M}_{\bullet}]$ is some weighted branched thingy in the moduli stack $\mathcal{M}_{\bullet}^{\text{st}}$, but in Section 5 we show how to integrate differential forms over $[\mathcal{M}_{\bullet}]$, and also how to push forward differential forms along natural evaluation maps to finite-dimensional exploded manifolds or orbifolds. Such differential forms on $\mathcal{M}_{\bullet}^{\text{st}}$ may be obtained by pulling back differential forms from manifolds or orbifolds under natural evaluation maps, or obtained as Chern classes of any naturally defined complex vectorbundle

over $\mathcal{M}_\bullet^{\text{st}}$, so we can define descendant Gromov–Witten invariants using Chern classes of tautological vectorbundles.

To simplify and emphasize the main points of our construction, we introduce the notion of a Kuranishi category, \mathcal{K} , and construct a virtual fundamental class $[\mathcal{K}]$ for any such Kuranishi category. Let us describe our results in terms of $[\mathcal{K}]$.

If θ is a differential form on \mathcal{K} and \mathcal{K} is compact² and oriented, then $\int_{[\mathcal{K}]} \theta$ is defined in Section 5. What type of differential form is θ ? Unlike in the case of smooth manifolds, there are several different types of differential forms that are useful on an exploded manifold \mathbf{B} :

$${}^r_{\text{fg}}\Omega^*(\mathbf{B}) \hookrightarrow {}^r\Omega^*(\mathbf{B}) \hookrightarrow \Omega^*(\mathbf{B}).$$

All three types coincide with smooth differential forms in the case that \mathbf{B} is a smooth manifold. De Rham cohomology defined using $\Omega^*(\mathbf{B})$ is much the same as usual cohomology; see our [31, Definition 1.2 and Corollary 4.2]. Refined cohomology, ${}^rH^*(\mathbf{B})$, defined using refined differential forms in ${}^r\Omega^*(\mathbf{B})$ is usually infinite-dimensional, but admits pushforwards, and acts as expected with fiber-products of exploded manifolds; see [31, Definition 9.1, Theorem 9.2, Lemma 9.3 and Lemma 9.5]. The cohomology ${}^r_{\text{fg}}H(\mathbf{B})$, defined using refined differential forms generated by functions, ${}^r_{\text{fg}}\Omega(\mathbf{B})$, is also compatible with pushforwards and fiber products, but unlike ${}^rH^*(\mathbf{B})$ and $H^*(\mathbf{B})$, is invariant only in families parametrized by connected smooth manifolds, rather than families parametrized by connected exploded manifolds. See Definition 5.4 for ${}^r_{\text{fg}}\Omega^*$. The advantage of differential forms generated by functions is that they are compatible with tropical completion—this is important for defining the contribution of a tropical curve to Gromov–Witten invariants, and for the tropical gluing formula for Gromov–Witten invariants, equation (1) of our [30].

The integral $\int_{\mathcal{K}} \theta$ makes sense for $\theta \in {}^r\Omega^*(\mathcal{K})$, and therefore makes sense for any of the above types of differential forms. If \mathcal{K} is complete³ (Definition 3.5) and $d\theta = 0$, then $\int_{[\mathcal{K}]} \theta$ is independent of all choices in the construction of \mathcal{K} , and depends only on the cohomology class represented by θ . The same holds with the weaker assumption that \mathcal{K} is compact, and the stronger assumption that $\theta \in {}^r_{\text{fg}}\Omega^*\mathcal{K}$.

²A compact Kuranishi category is one in which the subset consisting of holomorphic curves is compact; see Definition 3.5.

³If \mathcal{K} is an embedded Kuranishi structure for the moduli space of holomorphic curve in some complete exploded manifold, it is complete if and only if the corresponding moduli space of curves is compact. See our [28] for cases in which such compactness holds.

Given a complete, relatively oriented map $\pi: \mathcal{K} \rightarrow \mathcal{A}$ to an exploded manifold or orbifold \mathcal{A} , we can integrate forms along the fiber of π to define a map

$$\pi_!: {}^r\Omega^*(\mathcal{K}) \rightarrow {}^r\Omega^*(\mathcal{A}),$$

inducing a map on cohomology independent of all choices involved in the construction of $[\mathcal{K}]$ or $\pi_!$,

$$\pi_!: {}^rH^*(\mathcal{K}) \rightarrow {}^rH^*(\mathcal{A}) \quad \text{and} \quad \pi_!: {}^r_{\text{fg}}H^*(\mathcal{K}) \rightarrow {}^r_{\text{fg}}H^*(\mathcal{A}).$$

As usual, in the case that \mathcal{K} is complete,

$$\int_{[\mathcal{K}]} \pi^* \theta = \int_{\mathcal{A}} \theta \wedge \pi_!(1)$$

for any closed differential form $\theta \in {}^r\Omega^*(\mathcal{A})$.

The following theorem gives that, on the level of cohomology, $\pi_!$ only depends on the cobordism class of \mathcal{K} :

Theorem 1.1 *If \mathcal{K}_0 and \mathcal{K}_1 are cobordant within a stack \mathcal{X} with a map $\pi: \mathcal{X} \rightarrow \mathcal{A}$, then, given any construction of $[\mathcal{K}_j]$, the two composite maps*

$${}^rH^*(\mathcal{X}) \rightarrow {}^rH^*(\mathcal{K}_j) \xrightarrow{\pi_!} {}^rH^*\mathcal{A}$$

are equal, and the same holds for the analogous maps

$${}^r_{\text{fg}}H^*(\mathcal{X}) \rightarrow {}^r_{\text{fg}}H^*(\mathcal{K}_j) \xrightarrow{\pi_!} {}^r_{\text{fg}}H^*\mathcal{A}.$$

We also prove that Gromov–Witten invariants do not change in families of exploded manifolds, because they are compatible with pullbacks. Given a complete submersion $\mathcal{K} \rightarrow \mathcal{Z}$ to an exploded manifold, we can pull back \mathcal{K} over a map $\mathcal{Z}' \rightarrow \mathcal{Z}$ to obtain another Kuranishi category $\mathcal{K}' \rightarrow \mathcal{Z}'$. (For example, \mathcal{K} might come from holomorphic curves in a family of exploded manifolds parametrized by \mathcal{Z} . Then \mathcal{K}' is the Kuranishi category associated to holomorphic curves in the corresponding pulled-back family of exploded manifolds over \mathcal{Z}' .) When the map $\mathcal{K} \rightarrow \mathcal{Z}$ factors through a map $\mathcal{K} \rightarrow \mathcal{A}$, we get the following diagram of maps:

$$\begin{array}{ccc} \mathcal{K}' & \longrightarrow & \mathcal{K} \\ \downarrow \pi' & & \downarrow \pi \\ \mathcal{A} \times_{\mathcal{Z}} \mathcal{Z}' & \xrightarrow{y} & \mathcal{A} \\ \downarrow & & \downarrow \\ \mathcal{Z}' & \longrightarrow & \mathcal{Z} \end{array}$$

Theorem 1.2 *The following diagrams commute:*

$$\begin{array}{ccc}
 {}^rH^*(\mathcal{K}') & \xleftarrow{\quad} & {}^rH^*(\mathcal{K}) \\
 \downarrow \pi'_! & & \downarrow \pi_! \\
 {}^rH^*(A \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^rH^*(A)
 \end{array}
 \qquad
 \begin{array}{ccc}
 {}^r_{\text{fg}}H^*(\mathcal{K}') & \xleftarrow{\quad} & {}^r_{\text{fg}}H^*(\mathcal{K}) \\
 \downarrow \pi'_! & & \downarrow \pi_! \\
 {}^r_{\text{fg}}H^*(A \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^r_{\text{fg}}H^*(A)
 \end{array}$$

To prove gluing theorems, we also need fiber products of Kuranishi categories as well as pullbacks. However, the (fiber) product of Kuranishi categories usually has incompatible charts. In Section 6, this problem is solved by shrinking charts in the product of Kuranishi categories to obtain a “weak” product of Kuranishi categories.

Theorem 1.3 *Suppose that \mathcal{K} is a weak product of some finite collection of complete, oriented Kuranishi categories \mathcal{K}_v with maps $\pi_v: \mathcal{K}_v \rightarrow A_v$, and let $\pi: \mathcal{K} \rightarrow \prod_v A_v$ be the induced map. Then the following diagrams commute:*

$$\begin{array}{ccc}
 {}^rH^*(\mathcal{K}) & \xrightarrow{\pi_!} & {}^rH^*(\prod_v A_v) \\
 \uparrow & \nearrow \Pi_v(\pi_v)_! & \\
 \prod_v {}^rH^*(\mathcal{K}_v) & &
 \end{array}
 \qquad
 \begin{array}{ccc}
 {}^r_{\text{fg}}H^*(\mathcal{K}) & \xrightarrow{\pi_!} & {}^r_{\text{fg}}H^*(\prod_v A_v) \\
 \uparrow & \nearrow \Pi_v(\pi_v)_! & \\
 \prod_v {}^r_{\text{fg}}H^*(\mathcal{K}_v) & &
 \end{array}$$

Theorem 1.3 combines with Theorem 1.2 to show that integration over virtual fundamental classes acts as expected under fiber products. Theorems 1.1, 1.2 and 1.3 follow immediately from Theorems 5.20, 5.22 and 6.2.

Each exploded manifold \mathbf{B} has a tropical part, $\underline{\mathbf{B}}$, which describes the (infinitely) large-scale structure of \mathbf{B} ; the fiber, $\mathbf{B}|_p$, of $\mathbf{B} \rightarrow \underline{\mathbf{B}}$ over any point $p \in \underline{\mathbf{B}}$ is a smooth manifold. The integral of a differential form θ over an exploded manifold \mathbf{B} is defined as $\sum_{p \in \underline{\mathbf{B}}} \int_{\mathbf{B}|_p} \theta$. (Although $\underline{\mathbf{B}}$ is generally uncountable, only finitely many terms in this sum are nonzero when the integral of θ is defined.) Moreover, if $\theta \in {}^r_{\text{fg}}\Omega^*(\mathbf{B})$ is closed, then $\int_{\mathbf{B}|_p} \theta$ only depends on the homology class of θ in ${}^r_{\text{fg}}H^*(\mathbf{B})$. In Section 7 we prove analogous results for integration over $[\mathcal{K}]$.

One issue is that $\mathbf{B}|_p$ is usually not compact or complete, even when \mathbf{B} is complete. We deal with this issue using the tropical completion $\check{\mathbf{B}}|_p$ of \mathbf{B} at p , defined at the start of Section 7. This tropical completion $\check{\mathbf{B}}|_p$ always contains $\mathbf{B}|_p$ as a dense subset, and is complete if \mathbf{B} is compact. We can apply tropical completion to maps, differential forms, and Kuranishi categories. If $\theta \in {}^r_{\text{fg}}\Omega^*(\mathbf{B})$, then $\theta|_p \in {}^r_{\text{fg}}\Omega^*(\check{\mathbf{B}}|_p)$ and $\int_{\mathbf{B}|_p} \theta = \int_{\check{\mathbf{B}}|_p} \theta|_p$.

Lemma 7.10 states that the integral of a closed form $\theta \in {}^r_{\text{fg}}\Omega^*(\mathcal{K})$ breaks up into invariantly defined contributions for each point p in the tropical part of \mathcal{K} , and in particular,

$$\int_{[\mathcal{K}]} \theta = \sum_{p \in \underline{\mathcal{K}}} \int_{[\mathcal{K}]_p} \theta|_p.$$

Similarly, Lemma 7.10 states that, for a complete, relatively oriented map $\pi: \mathcal{K} \rightarrow \mathcal{A}$, cohomology class $\theta \in {}^r_{\text{fg}}H^*\mathcal{K}$ and $p' \in \underline{\mathcal{A}}$,

$$\pi_!(\theta)|_{p'} = \sum_{p \in \pi^{-1}(p')} (\pi|_p)_!(\theta|_p).$$

This paper concludes with Section 8, which summarizes our construction of Gromov–Witten invariants.

Acknowledgements

This research was supported by the Australian Research Council. Special thanks are due to an anonymous referee, who suggested several improvements to both this paper and [32], on which it depends.

2 Constructing sections of sheaves on K –categories

Throughout this paper, we will be using exploded manifolds with the regularity $C^{\infty,1}$, defined in Section 7 of [24]. For all practical purposes, $C^{\infty,1}$ maps are as good as smooth. Let $\mathcal{M}_{\bullet}^{\text{st}}$ indicate some decorated moduli stack of $C^{\infty,1}$ families of not necessarily holomorphic stable curves; see Section 11 of [24] for basic definitions, and Section 2 of [26] for further treatment, including Definition 2.11 of $\mathcal{M}_{\bullet}^{\text{st}}$. By a sheaf (of sets) on $\mathcal{M}_{\bullet}^{\text{st}}$, we mean a contravariant functor S from $\mathcal{M}_{\bullet}^{\text{st}}$ (to the category of sets), so that S is a sheaf when restricted to the category of open subfamilies of any family in $\mathcal{M}_{\bullet}^{\text{st}}$. We shall also be interested in sheaves with more restricted domains.

A stack \mathcal{X} over the category of $C^{\infty,1}$ exploded manifolds is a category \mathcal{X} along with a “nice” functor \mathbf{F} from \mathcal{X} to the category of $C^{\infty,1}$ exploded manifolds; see [7] for an approachable introduction to stacks, and see Section 2 of [26] for a study of the stack of $C^{\infty,1}$ curves. In this paper, “stack” without further qualification will generally mean a stack over the category of $C^{\infty,1}$ exploded manifolds. The “nice” properties of \mathbf{F} are loosely paraphrased as follows: families (parametrized by exploded manifolds)

glue and pull back as we expect bundles (parametrized by exploded manifolds) to glue and pull back. Moreover, morphisms between families also restrict, pull back and glue as expected for fiberwise isomorphisms. We use \hat{f} to indicate an object in \mathcal{X} , and call \hat{f} a family parametrized by $F(\hat{f})$ —this can also be thought of as a map $F(\hat{f}) \rightarrow \mathcal{X}$. The case we are most interested in is when \mathcal{X} is a stack of curves, so \hat{f} can be thought of as a family of curves, and any $f \rightarrow \hat{f}$ with $F(f)$ a point can be thought of as a curve f in \hat{f} . As in Section 2 of [26], by a substack $\mathcal{U} \subset \mathcal{X}$, we mean a full subcategory such that \hat{f} is in \mathcal{U} if and only if every f in \hat{f} is in \mathcal{U} . Then, for any family \hat{f} in \mathcal{X} , there is a corresponding subset $\mathcal{U}(\hat{f}) \subset F(\hat{f})$ such that for any $\hat{g} \rightarrow \hat{f}$, \hat{g} is in \mathcal{U} if and only if $F(\hat{g})$ has image in $\mathcal{U}(\hat{f})$. We call \mathcal{U} an open substack if $\mathcal{U}(\hat{f}) \subset F(\hat{f})$ is open for all \hat{f} in \mathcal{X} . It is proved in Lemma 2.7 of [26] that this topology on the moduli stack of $C^{\infty,1}$ curves matches the topology used in [28] to prove compactness results for the moduli stack of holomorphic curves.

When we take an embedded Kuranishi structure, and throw away all information apart from charts and their embedding into a stack, we obtain a K -category, defined below. See also Remark 7.6 for a definition of a K -category using charts not embedded in a stack.

Definition 2.1 A K -category is a full subcategory \mathcal{K} of a stack \mathcal{K}^{st} and a collection of charts \hat{f}_i/G_i such that the following holds:

- (i) Each \hat{f}_i is a family in \mathcal{K}^{st} , and G_i is a finite group of automorphisms of \hat{f}_i .
- (ii) \hat{f}_i/G_i represents a substack of \mathcal{K}^{st} .
- (iii) Each family \hat{f}_i has some fixed dimension, and whenever $\dim \hat{f}_i \leq \dim \hat{f}_j$,

$$\hat{f}_i \times_{\mathcal{K}^{\text{st}}} \hat{f}_j \rightarrow \hat{f}_i$$

is a G_j -fold cover of an open subfamily of \hat{f}_i , and

$$\hat{f}_i \times_{\mathcal{K}^{\text{st}}} \hat{f}_j \rightarrow \hat{f}_j$$

is a G_i -fold cover of a subfamily of \hat{f}_j —locally defined by the transverse vanishing of some \mathbb{R} -valued $C^{\infty,1}$ functions on $F(\hat{f}_j)$.

- (iv) The families \hat{f}_i cover \mathcal{K}^{st} .
- (v) The set of charts is countable and locally finite—for each i , there are only finitely many j such that $\hat{f}_i \times_{\mathcal{K}^{\text{st}}} \hat{f}_j$ is nonempty.

- (vi) \mathcal{K} is the full subcategory of \mathcal{K}^{st} consisting of families locally isomorphic to some \widehat{f}_i .

Given another K -category $\mathcal{K}^\#$ with charts $\widehat{f}_i^\# / G_i$, use the notation $\mathcal{K} \subset \mathcal{K}^\#$ if

- \mathcal{K}^{st} is a substack of $(\mathcal{K}^\#)^{\text{st}}$,
- \mathcal{K} is a subcategory of $\mathcal{K}^\#$, and
- \widehat{f}_i is a G_i -equivariant open subfamily of $\widehat{f}_i^\#$.

Moreover, say that $\mathcal{K}^\#$ is an extension of \mathcal{K} and use the notation $\mathcal{K} \subseteq_e \mathcal{K}^\#$ if $\mathcal{K} \subset \mathcal{K}^\#$ and the closure of $\widehat{f}_i \subset \widehat{f}_i^\#$ is closed in $(\mathcal{K}^\#)^{\text{st}}$ —equivalently, if \widehat{f}_i' indicates the closure of $\widehat{f}_i \subset \widehat{f}_i^\#$, then, for all j , $\widehat{f}_i' \times_{(\mathcal{K}^\#)^{\text{st}}} \widehat{f}_j^\#$ has closed image in $\widehat{f}_j^\#$.

Say that \mathcal{K} is extendable if there exist extensions $\mathcal{K} \subseteq_e \mathcal{K}' \subseteq_e \mathcal{K}^\#$.⁴

Definition 2.2 A sheaf (of sets) on a K -category, \mathcal{K} , is a contravariant functor, S , from \mathcal{K} (to the category of sets) such that S is a sheaf whenever restricted to the category of open subfamilies of a given family in \mathcal{K} .

A sheaf on an extendable K -category is a sheaf defined on some extension $\mathcal{K}' \subseteq_e \mathcal{K}$.

Extendability is an important property for constructing sections of sheaves on \mathcal{K} . Any extendable \mathcal{K} has a partition of unity, and Proposition 2.3 gives a way of constructing global sections of sheaves on \mathcal{K} . For example, if \mathcal{K} is extendable, each G_i is trivial and each $F(\widehat{f}_i)$ is an n -dimensional manifold, then the \widehat{f}_i provide coordinate charts on an n -dimensional manifold M , and \mathcal{K}^{st} is the moduli stack of maps into M . If we drop the condition that G_i be trivial, then \widehat{f}_i / G_i give charts on an orbifold, M (and depending on your position on orbifolds, \mathcal{K}^{st} is either that orbifold, or is the category of maps into M). If, however, \mathcal{K} is not extendable, M may not be Hausdorff.

Global sections of a sheaf on \mathcal{K} can be constructed using the patching, extension and averaging axioms below. To construct Gromov–Witten invariants, we shall use the sheaf S from Definition 3.12 below. This sheaf obeys the patching and extension axioms, but not the averaging axiom. Accordingly, global sections of this sheaf S may not exist, so we shall use weighted branched sections of S , which we regard as a natural transformation $I \rightarrow S$, where I is a “sheaf”⁵ of finite measure spaces.

The following axioms allow the global construction of sections of a sheaf S over an extendable K -category \mathcal{K} :

⁴We require two extensions so that any extendable K -category \mathcal{K} will have an extendable extension — if $\mathcal{K} \subseteq_e \mathcal{K}' \subseteq_e \mathcal{K}^\#$, then we can obtain an extendable extension of \mathcal{K} by shrinking \mathcal{K}' appropriately.

⁵This “sheaf” requires scare quotes because it has a restricted domain of definition.

- **Patching** Given an open cover $\{O_k\}$ of \hat{f} , and section $\theta_k \in S(O_k)$ for all k , there exists a patched-together section $\theta \in S(\hat{f})$ that agrees with θ_k wherever all θ_k agree.

So, given any morphism $\hat{g} \rightarrow \hat{f}$ in \mathcal{K} and section $\theta' \in S(\hat{g})$, the pullback of θ to $S(\hat{g})$ is θ' if the following condition is satisfied for all k : θ' and θ_k pull back to the same section under the diagram

$$\begin{array}{ccc} \hat{g} \times_{\hat{f}} O_k & \longrightarrow & O_k \\ \downarrow & & \downarrow \\ \hat{g} & \longrightarrow & \hat{f} \end{array}$$

- **Extension** $S(\hat{f})$ is nonempty. Moreover, if $\hat{g} \rightarrow \hat{f}$ is a morphism in \mathcal{K} , then given any $\theta \in S(\hat{g})$ and f in \hat{g} , there exists some $\theta' \in S(\hat{f})$ such that the pullback of θ' to $S(\hat{g})$ agrees with θ in a neighborhood of $f \in \hat{g}$.
- **Averaging** Given any $\theta \in S(\hat{f})$ and a finite G -action on \hat{f} , there is a G -invariant section θ' agreeing with θ wherever θ was already G -invariant—so, given any G -equivariant map $\hat{f}' \rightarrow \hat{f}$ pulling θ back to a G -invariant section, the pullback of θ agrees with the pullback of θ' over $\hat{f}' \rightarrow \hat{f}$.

For example, if $S(\hat{f})$ is the set of $C^{\infty,1}$ real-valued functions on $F(\hat{f})$, then S obeys the above axioms. The extension axiom for S follows from Definition 2.1(iii): given any morphism $\hat{g} \rightarrow \hat{f}$ in \mathcal{K} , the corresponding map $F(\hat{g}) \rightarrow F(\hat{f})$ is always locally an isomorphism onto an exploded submanifold of $F(\hat{f})$. Patching can be achieved using partitions of unity, and averaging achieved by averaging over the action of G .

Proposition 2.3 *Let S be a sheaf on an extendable K -category, $\mathcal{K}_2^\#$, and suppose that S satisfies the patching, extension and averaging axioms. Given the inclusions and extensions of K -categories*

$$\begin{array}{ccc} \mathcal{K}_1 & \xrightarrow{\subseteq_e} & \mathcal{K}_1^\# \\ \downarrow & & \downarrow \\ \mathcal{K}_2 & \xrightarrow{\subseteq_e} & \mathcal{K}_2^\# \end{array}$$

and a global section θ of S defined on $\mathcal{K}_1^\#$, there exists a global section of S defined on \mathcal{K}_2 and agreeing with θ when restricted to \mathcal{K}_1 .

Proof Use notation $\widehat{f}_{i,\mathcal{K}_j}/G_i \subset \widehat{f}_{i,\mathcal{K}_j^\#}/G_i$ for charts on \mathcal{K}_j and $\mathcal{K}_j^\#$, respectively, and index these charts by the natural numbers. We shall construct our section inductively, but at each step we must shrink the domain of definition a little. Accordingly, for each (i, j) where j is a nonnegative integer, choose G_i -invariant open subfamilies $\widehat{f}_{i,j} \subset \widehat{f}_{i,\mathcal{K}_2^\#}$ so that

- (i) if $j \geq i$, then $\widehat{f}_{i,\mathcal{K}_2} \subset \widehat{f}_{i,j}$;
- (ii) if $j < i$, then $\widehat{f}_{i,\mathcal{K}_1} \subset \widehat{f}_{i,j} \subset \widehat{f}_{i,\mathcal{K}_2^\#}$, and $\widehat{f}_{i,j}$ contains every curve in the closure of $\widehat{f}_{i,\mathcal{K}_1}$ within $\mathcal{K}_2^\#$;
- (iii) if $j \neq i$, then $\widehat{f}_{i,j} \subset \widehat{f}_{i,j-1}$, and $\widehat{f}_{i,j-1}$ contains every curve in the closure of $\widehat{f}_{i,j}$ within $\mathcal{K}_2^\#$.

Such families can be constructed as follows. Because $\mathcal{K}_2^\#$ is extendable, there exists a family $\widehat{f}_i^\#$ containing $\widehat{f}_{i,\mathcal{K}_2^\#}$ as an open subfamily, so that the closure of $\widehat{f}_{i,\mathcal{K}_2^\#}$ within $\widehat{f}_i^\#$ contains the closure of $\widehat{f}_{i,\mathcal{K}_2^\#}$ within $\mathcal{K}_2^\#$ (or any of our other K -categories). The same will then hold for the open subfamilies $\widehat{f}_{i,j} \subset \widehat{f}_{i,\mathcal{K}_2^\#}$. Choose $\widehat{f}_{i,0} = \widehat{f}_{i,\mathcal{K}_1^\#}$. Then, choose a continuous function, ρ , on $\widehat{f}_i^\#$ that restricts to be 1 on $\widehat{f}_{i,\mathcal{K}_1} \subset \widehat{f}_{i,0}$, and 0 outside $\widehat{f}_{i,0} \subset \widehat{f}_i^\#$. With this ρ , we then define $\widehat{f}_{i,j} := \rho^{-1}(j/i, 1]$ for $j < i$. Next, define $\widehat{f}_{i,i} := \widehat{f}_{i,\mathcal{K}_2^\#}$, and choose another continuous function λ on $\widehat{f}_i^\#$, equal to 0 on $\widehat{f}_{i,\mathcal{K}_2}$, and 1 outside $\widehat{f}_{i,i} \subset \widehat{f}_i^\#$. Finally, define $\widehat{f}_{i,i+n} := \lambda^{-1}[0, 2^{-n})$ for positive integers n .

Let us construct our section using the following inductive step:

Claim 2.4 Suppose that a section θ_{j-1} of $S(\widehat{f}_{i,j-1})$ has been chosen for all i , and that these sections are compatible — so, given any pair of morphisms $\iota_1: \widehat{f} \rightarrow \widehat{f}_{i,j-1}$ and $\iota_2: \widehat{f} \rightarrow \widehat{f}_{i',j-1}$, we have $\iota_1^* \theta_{j-1} = \iota_2^* \theta_{j-1}$. Suppose moreover that the restriction of θ_{j-1} to $\widehat{f}_{i,\mathcal{K}_1}$ agrees with our original section θ .

Then there exists a compatible choice of sections θ_j of $S(\widehat{f}_{i,j})$ for all i , such that $\theta_j \in S(\widehat{f}_{i,j})$ is the restriction of $\theta_{j-1} \in S(\widehat{f}_{i,j-1})$ for all $i \neq j$, and such that the restrictions of θ_j and θ to $S(\widehat{f}_{i,\mathcal{K}_1})$ are equal.

To prove Claim 2.4, we need only define θ_j on $\widehat{f}_{j,j}$. Let us first construct a local candidate, θ' , for θ_j . For a curve f in $\widehat{f}_{j,j}$, let i be such that $f \in \widehat{f}_{i,j-1}$, and $\widehat{f}_{i,j-1}$ has the maximal dimension⁶ of all such $\widehat{f}_{k,j-1}$ containing f . (If there is no $\widehat{f}_{i,j-1}$ containing f , we are free to choose any section of S on a neighborhood of f that does not intersect any $\widehat{f}_{i,j}$ for $i \neq j$.) By using the extension axiom (or by pulling back θ from $\widehat{f}_{i,j-1}$), there exists a section θ' of S on a neighborhood U_f of f in $\widehat{f}_{j,j}$

⁶Recall, from Definition 2.1(v), that f can only be contained in finitely many \widehat{f}_i .

such that, given any pair of morphisms $\iota_1: \hat{f} \rightarrow U_f$ and $\iota_2: \hat{f} \rightarrow \hat{f}_{i,j-1}$, we have $\iota_1^* \theta' = \iota_2^* \theta_{j-1}$.

Let us check that θ' locally satisfies the required compatibility conditions. Because \hat{f}_i has maximal dimension, we may choose U_f small enough that Definition 2.1(ii) implies that, if \hat{f} is contained in both U_f and $\hat{f}_{i',j}/G_{i'}$ for any other $i' \neq j$, then \hat{f} is contained in $\hat{f}_{i,j-1}/G_i$. It follows that $\theta' \in S(U_f)$ is compatible with $\theta_j = \theta_{j-1} \in S(\hat{f}_{i',j})$ for all $i' \neq j$. We should also ensure that θ' is compatible with θ on $\hat{f}_{j,\mathcal{K}_1}$. Because $\hat{f}_{i,j}$ has maximal dimension, if $\hat{f}_{j,j}$ has larger dimension, f is not in $\hat{f}_{j,\mathcal{K}_1}$. Therefore, we may also choose U_f small enough that if \hat{f} is contained in U_f and $\hat{f}_{j,\mathcal{K}_1}$, then \hat{f} is contained in $\hat{f}_{i,j}/G_i$, ensuring that $\theta' \in S(U_f)$ is also compatible with $\theta \in S(\hat{f}_{j,\mathcal{K}_1})$.

The patching axiom allows us to patch together these sections to a section, $\theta' \in S(\hat{f}_{j,j})$, still compatible with θ_j on $\hat{f}_{i,j}$ for $i \neq j$ and θ on $\hat{f}_{j,\mathcal{K}_1}$. Then the averaging axiom gives a G_j -equivariant section θ_j still compatible with θ_j on $\hat{f}_{i,j}$ and θ on $\hat{f}_{j,\mathcal{K}_1}$. Definition 2.1(ii) ensures that any G_j -equivariant section of $S(\hat{f}_{j,j})$ has the required compatibility property that its pullback does not depend on the choice of morphism.

This completes the proof of Claim 2.4.

Using Claim 2.4 inductively, we can construct θ_i for all i . Any family \hat{f} in \mathcal{K}_2 is everywhere locally isomorphic to some subfamily of $\hat{f}_{i,\mathcal{K}_2} \subset \hat{f}_{i,i}$. We may define $\theta \in S(\hat{f})$ to be the pullback of θ_i . This defines the required global section of S over \mathcal{K}_2 . \square

Remark 2.5 Proposition 2.3 may also be used to construct global sections of sheaves S on \mathcal{K}^{st} because any global section of S on \mathcal{K} automatically pulls back to give a global section of S on \mathcal{K}^{st} .

Lemma 2.6 Consider an extension $\mathcal{K} \subset_e \mathcal{K}^\#$ of extendable K -categories, and a corresponding pair of charts $\hat{f}_i/G_i \subset \hat{f}_i^\#/G_i$. Then any (continuous or $C^{\infty,1}$) function

$$\rho: \hat{f}_i/G_i \rightarrow \mathbb{R}$$

admitting an extension to $\hat{f}_i^\#/G_i$ also extends to all of \mathcal{K} .

Moreover, if $O \subset \mathcal{K}^\#$ is any open subset such that the support of ρ on $\hat{f}_i^\#$ has closure (within $\mathcal{K}^\#$) contained in O , then we may choose the corresponding extension $\rho: \mathcal{K} \rightarrow \mathbb{R}$ also to have support with closure contained in O , and if ρ is nonnegative, we can choose its extension to also be nonnegative.

Proof We shall construct the required function using Proposition 2.3. Define

- $\mathcal{K}_1 \subset \mathcal{K}$ to be the category of families locally isomorphic to an open subset of \widehat{f}_i ,
- $\mathcal{K}_1^\# \subset \mathcal{K}^\#$ to be the category of families locally isomorphic to an open subset of $\widehat{f}_i^\#$,
- $\mathcal{K}_2 = \mathcal{K}$,
- $\mathcal{K}_2^\# = \mathcal{K}^\#$.

Let S be the sheaf with $S(\widehat{f})$ the set of $C^{\infty,1}$ functions on $F(\widehat{f})$ whose support has closure (in $\mathcal{K}^\#$) contained in O . The patching axiom for S is proved using a partition of unity, and the averaging axiom holds, because it is possible to average a $C^{\infty,1}$ function to obtain a G -invariant function. To apply Proposition 2.3, we only need to prove that S satisfies the extension axiom.

The extension axiom follows from Definition 2.1(iii)—every morphism $\widehat{f} \rightarrow \widehat{g}$ is locally an isomorphism onto an exploded submanifold (which must be locally closed because $\mathcal{K}^\#$ is extendable). In particular, this implies that any (continuous or $C^{\infty,1}$) function defined on $F(\widehat{f})$ locally extends to a (continuous or $C^{\infty,1}$) function on $F(\widehat{g})$. If our original function had support with closure contained in O , our local extension may also be chosen to have support with closure contained in O . Therefore, S obeys the extension axiom, and Proposition 2.3 tells us that ρ extends to a global section of S , which is a map

$$\rho: \mathcal{K} \rightarrow \mathbb{R}.$$

The fact that the support of ρ has closure contained in O follows from the analogous fact for each of the individual functions $\widehat{f}_k/G_k \rightarrow \mathbb{R}$, and the local-finiteness condition for $\mathcal{K}^\#$ from Definition 2.1(v).

The proof in the case that ρ is nonnegative is identical, except we use a sheaf with $S(\widehat{f})$ the set of nonnegative $C^{\infty,1}$ functions on $F(\widehat{f})$ whose support has closure contained in O . \square

Remark 2.7 We can use Lemma 2.6 to construct a nonnegative $C^{\infty,1}$ function on \mathcal{K} with zero set any given closed substack C of $\mathcal{K}_e^\# \supset \mathcal{K}$.

In particular, Lemma 2.6 implies that any nonnegative $C^{\infty,1}$ bump function on $\widehat{f}_i^\#/G_i$ whose support has closure contained in the complement of C extends to a nonnegative

$C^{\infty,1}$ function ρ on \mathcal{K} , vanishing on C . As $\widehat{f}_i^\sharp/G_i \setminus C$ is covered by the support of a countable collection of such bump functions, and \mathcal{K} has a countable number of charts, there exists a sequence, ρ_k , of nonnegative functions on \mathcal{K} , vanishing on C and with support covering $\mathcal{K} \setminus C$. Then there exists a sequence, ϵ_k , of positive numbers such that $\sum_k \epsilon_k \rho_k$ converges to a $C^{\infty,1}$ function ρ on \mathcal{K} . (As with smooth functions on smooth manifolds, convergence to a $C^{\infty,1}$ function is equivalent to convergence in a countable sequence of norms, so we can always ensure convergence of a sum by multiplying each term by a suitably small constant. See Definition 7.5 of [24].) Such a ρ is nonnegative and has zero set C , as required.

3 Kuranishi categories

Embedded Kuranishi structures are defined in Section 2.9 of [26]. Each Kuranishi chart $(\mathcal{U}, V, \widehat{f}/G)$ consists of some open substack $\mathcal{U} \subset \mathcal{M}_\bullet^{\text{st}}$, an obstruction bundle V over \mathcal{U} and a family of curves \widehat{f} in \mathcal{U} with automorphism group G . The obstruction bundle, V , is a finite-rank complex vectorbundle over \mathcal{U} , and also a nice⁷ subbundle of the sheaf \mathcal{Y} that is the codomain of the $\bar{\partial}$ equation. Moreover, \widehat{f}/G represents the moduli stack $\bar{\partial}^{-1}V \subset \mathcal{U}$, so $\bar{\partial}\widehat{f}$ defines a G -equivariant section of the restriction, $V(\widehat{f})$, of V to \widehat{f} . Kuranishi charts in an embedded Kuranishi structure have to be compatible in the sense that on $\mathcal{U}_i \cap \mathcal{U}_j$, either V_i is a subbundle of V_j or vice versa. Moreover, our Kuranishi charts have compatible extensions, so we can define an extendable Kuranishi category by taking the charts \widehat{f}_i/G_i .

Definition 3.1 Given an embedded Kuranishi structure $\{(\mathcal{U}_i, V_i, \widehat{f}_i/G_i)\}$ on $\mathcal{M}_\bullet^{\text{st}}$, define its associated Kuranishi category, \mathcal{K} , to be the (full) subcategory of $\mathcal{M}_\bullet^{\text{st}}$ consisting of families locally isomorphic to an open subfamily of some \widehat{f}_i , and define \mathcal{K}^{st} to be the substack of $\mathcal{M}_\bullet^{\text{st}}$ consisting of curves isomorphic to curves in some \widehat{f}_i ; so, a family of curves \widehat{f} in $\mathcal{M}_\bullet^{\text{st}}$ is in \mathcal{K}^{st} if each curve f in \widehat{f} is isomorphic to a curve in some \widehat{f}_i . This category \mathcal{K} comes with the following extra structure:

- the open substacks $\mathcal{U}_i \cap \mathcal{K}^{\text{st}}$, with the vectorbundles V_i ;
- the charts \widehat{f}_i/G_i ;
- the section $\bar{\partial}$ of $V_i(\widehat{f}_i)$ over $F(\widehat{f}_i)$.

In the next definition, we specify the essential properties of this extra structure on \mathcal{K} .

⁷Technically, V satisfies Definitions 2.23, 2.24 and 2.25 of [26].

Definition 3.2 A Kuranishi category is an extendable K -category \mathcal{K} with charts \hat{f}_i/G_i (Definition 2.1) along with:

- open substacks $\mathcal{U}_i \subset \mathcal{K}^{\text{st}}$ containing \hat{f}_i/G_i such that each \mathcal{U}_i only intersects finitely many other \mathcal{U}_j ;
- constant-rank complex vectorbundles V_i on \mathcal{U}_i and, on $\mathcal{U}_i \cap \mathcal{U}_j$, an inclusion of one of V_i or V_j as a subbundle of the other; and
- sections $\bar{\partial}\hat{f}_i: \mathbf{F}(\hat{f}_i) \rightarrow V_i(\hat{f}_i)$

satisfying the following conditions:

- The sections $\bar{\partial}\hat{f}_i$ determine a global section, $\bar{\partial}$, of the sheaf with sections over \hat{f} the sections of a vectorbundle, $V(\hat{f})$, equal to $V_i(\hat{f})$ wherever \hat{f} is locally isomorphic to \hat{f}_i , and with pullbacks induced by the inclusions of vectorbundles above. So, V is a covariant functor, with $V(\hat{f}) \rightarrow \mathbf{F}(\hat{f})$ a complex vectorbundle, and $\bar{\partial}$ is a natural transformation from \mathbf{F} to V , with $\bar{\partial}\hat{f}$ a section of $V(\hat{f})$.
- $\bar{\partial}\hat{f}$ is transverse to $V_j(\hat{f}) \subset V(\hat{f})$ when $V_j(\hat{f})$ is defined and $\dim V_j \leq \dim V(\hat{f})$; moreover, the intersection of $\bar{\partial}\hat{f}$ with V_j is locally isomorphic to \hat{f}_j (and therefore contained in \mathcal{K}).
- On the other hand, if $V(\hat{f}) \subset V_i(\hat{f})$ — so $\hat{f} \in \mathcal{U}_i$ and $\dim V(\hat{f}) \leq \dim V_j(\hat{f})$ — then \hat{f} is in the substack represented by \hat{f}_i/G_i — so, in \mathcal{K} , there is a map of a G_i -fold cover of \hat{f} to \hat{f}_i .

Remark 3.3 A Kuranishi category, \mathcal{K} , contains more information than a choice of good coordinate system from [8]. In particular, the vectorbundles V_i on open subsets \mathcal{U}_i do not appear there, and the transversality condition (ii) is only required to hold at the intersection of $\bar{\partial}$ with 0 — without the extensions of our vectorbundles V_i from our definition, this condition only makes sense at the intersection of $\bar{\partial}$ with 0. We also explicitly require that \mathcal{K} be extendable — I think that this condition may be obtained by shrinking a good coordinate system as in [9], and expect the extra data of Definition 3.2 to be definable from a good coordinate system after shrinking and making choices using Proposition 2.3.

Definition 3.4 For a Kuranishi category, \mathcal{K} , define $\mathcal{K}^{\text{hol}} \subset \mathcal{K}^{\text{st}}$ to be the substack of \mathcal{K}^{st} consisting of all holomorphic curves — those f in \hat{f} such that $\bar{\partial}\hat{f}$ vanishes at f . Use the induced topology from \mathcal{K}^{st} on \mathcal{K}^{hol} , so define an open substack of \mathcal{K}^{hol} to be the intersection of an open substack of \mathcal{K}^{st} with \mathcal{K}^{hol} .

Note that Definition 3.2(ii)–(iii) imply that the intersection of \mathcal{K}^{hol} with \mathcal{U}_i is the quotient of $\{\bar{\partial}\hat{f}_i = 0\}$ by G_i , so each \hat{f}_i/G_i covers an open substack of \mathcal{K}^{hol} .

We need the notion of a Kuranishi category over an exploded manifold or orbifold — by exploded orbifold, we mean a Deligne–Mumford stack over the category of exploded manifolds: a stack \mathcal{Z} with proper diagonal $\mathcal{Z} \rightarrow \mathcal{Z} \times \mathcal{Z}$, and locally represented by A/G , where A is an exploded manifold and G is a finite group acting on A ; see Remark 2.3 of [26].

Definition 3.5 A Kuranishi category over an exploded manifold or orbifold \mathcal{Z} is a Kuranishi category \mathcal{K} along with a submersion $\pi: \mathcal{K} \rightarrow \mathcal{Z}$. Say that \mathcal{K} is proper over \mathcal{Z} if π restricted to \mathcal{K}^{hol} is proper — so, given a map $\mathcal{B} \rightarrow \mathcal{Z}$ from an exploded manifold \mathcal{B} with a compact subset C , the stack $\mathcal{K}^{\text{hol}} \times_{\mathcal{Z}} C$ is compact.

Say that \mathcal{K} is complete over \mathcal{Z} if it is proper over \mathcal{Z} and, for every family \hat{f} in \mathcal{K} , the map of integral affine spaces $\underline{F}(\hat{f}) \rightarrow \underline{\mathcal{Z}}$ is complete — so, the inverse image of any complete subset of $\underline{\mathcal{Z}}$ (with its integral affine connection) is a complete subset of $\underline{F}(\hat{f})$.

Say that \mathcal{K} is compact or complete if it is proper or complete, respectively, over a point.

Note that \mathcal{K} may be complete without \mathcal{K}^{hol} being complete (in the sense of Definition 3.15 of [24]). For example consider \mathcal{K} with a single chart \hat{f} , where $\underline{F}(\hat{f}) = T_{[0,\infty)}^1$, $V(\hat{f}) = \mathbb{C} \times T_{[0,\infty)}^1$ and $\bar{\partial}\hat{f}(\tilde{z}) = ([\tilde{z}], \tilde{z})$ (using notation from Example 3.5 of [24]). Then \mathcal{K}^{hol} is the stack represented by $T_{(0,\infty)}^1$, which is compact but not complete. On the other hand, \mathcal{K} is complete by Definition 3.5 because the tropical part of $T_{[0,\infty)}^1$ is the complete polytope $[0, \infty)$. If we perturb $\bar{\partial}\hat{f}$ to be transverse to the zero section, its intersection with the zero section will be complete, so \mathcal{K}^{hol} wants to be complete.

If \mathcal{K} is the Kuranishi category associated to an embedded Kuranishi structure in some collection of connected components of $\mathcal{M}^{\text{st}}(\hat{\mathcal{B}})$, then Lemma 4.2 and Theorem 6.1 of [28] give conditions under which \mathcal{K}^{hol} is proper over \mathcal{B}_0 . Whenever such a \mathcal{K} is proper over \mathcal{B}_0 , it is also complete over \mathcal{B}_0 , because connected families \hat{f} in \mathcal{K} either consist of curves with domain \mathcal{T} , and are parametrized by (a cover of) the quotient of some open subset of $\hat{\mathcal{B}}$ by some \mathcal{T} -action, or are families of curves in $\hat{\mathcal{B}} \rightarrow \mathcal{B}_0$ with universal tropical structure. Universal tropical structure is defined in Definition 4.1 and Theorem 3.1 of [27]; see also Remark 3.3 of [27]. The \hat{f} in the Kuranishi charts from [26] are constructed to have universal tropical structure. In fact, Theorem 5.3 and

Proposition 5.9 of [26] imply that if \hat{f}/G locally represents a closed substack of the moduli stack of curves (and the curves in \hat{f} don't have domain T), then \hat{f} must have universal tropical structure.

3.1 Orienting Kuranishi charts

We orient our Kuranishi charts using a canonical homotopy of $D\bar{\partial}$ to a complex-linear operator. At holomorphic curves f , $T_f\mathcal{M}_\bullet^{\text{st}}(\mathbf{B})$ has a canonical complex structure, described in Section 2.7 of [26]. In the more general case of curves in a family of targets $\hat{\mathbf{B}} \rightarrow \mathbf{B}_0$, $T_f\mathcal{M}_\bullet^{\text{st}}(\hat{\mathbf{B}})\downarrow_{\mathbf{B}_0}$ has a canonical complex structure.⁸ In this section, we discuss a canonical orientation of our Kuranishi charts relative to \mathbf{B}_0 . This orientation is only defined in a neighborhood of the homomorphic curves, so we construct Gromov–Witten invariants in this neighborhood, where the orientation is defined.

For a Kuranishi chart $(\mathcal{U}, V, \hat{f}/G)$, V is a complex vectorbundle over \mathcal{U} , and at each holomorphic curve f in \hat{f} , a canonical linear homotopy of $D\bar{\partial}$ to a complex operator is transverse to V . As explained in [26], this homotopy gives us a canonical homotopy of $T_f\mathbf{F}(\hat{f})\downarrow_{\mathbf{B}_0}$ to a complex subspace of $T_f\mathcal{M}_\bullet^{\text{st}}(\hat{\mathbf{B}})\downarrow_{\mathbf{B}_0}$. In particular, there is a canonical homotopy class of almost complex structures on $T_f\mathbf{F}(\hat{f})\downarrow_{\mathbf{B}_0}$. Section 8 of [26] constructs a complex structure on $T_f\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$ at all holomorphic curves f , so that the following holds:

- The complex structure extends to the vectorbundle $T\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$ on a neighborhood of the holomorphic curves in $\hat{f}_i^\#$.
- Given an open neighborhood \hat{f} of a holomorphic f in $\hat{f}_i^\#$ and a map $\hat{f} \rightarrow \hat{f}_j^\#$, the short exact sequence

$$0 \rightarrow T_f\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0} = T_f\mathbf{F}(\hat{f})\downarrow_{\mathbf{B}_0} \rightarrow T_f\mathbf{F}(\hat{f}_j)\downarrow_{\mathbf{B}_0} \xrightarrow{D\bar{\partial}} V_j/V_i \rightarrow 0$$

is complex.

- The \mathbb{R} -nil vectors in $T_f\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$ (ie those that act as the zero derivation on \mathbb{R} -valued functions) are given the canonical complex structure.

Such a complex structure on $T_f\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$ was constructed in Proposition 8.10 of [26] by choosing a connection on the inverse image of V_i under the homotopy of $D\bar{\partial}$ to its complex-linear part. In fact, an appropriate sheaf of such choices obeys the patching, extension and averaging axioms.

⁸Given a submersion $\pi: \mathbf{A} \rightarrow \mathbf{B}_0$, we use the notation $T\mathbf{A}\downarrow_{\mathbf{B}_0}$ to indicate the vertical tangent bundle of \mathbf{A} — ie the kernel of $T\pi$; see Section 2.7.2 of [26] for a discussion of $T_f\mathcal{M}_\bullet^{\text{st}}(\hat{\mathbf{B}})\downarrow_{\mathbf{B}_0}$.

Does such a complex structure at holomorphic curves f induce an orientation of $T\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$ elsewhere? To prove it does, we construct a global 2-form, α , so that at holomorphic curves f , α is positive on complex planes within $T_f\mathbf{F}(\hat{f}_i^\#)\downarrow_{\mathbf{B}_0}$.

Definition 3.6 (orienting 2-form) Let the sheaf of orienting 2-forms on \hat{f} be the sheaf of 2-forms, α , on $\mathbf{F}(\hat{f})$ such that α is positive on holomorphic planes within $T_f\mathbf{F}(\hat{f})\downarrow_{\mathbf{B}_0}$ for all holomorphic curves f in \hat{f} .

An orienting 2-form is a global section, α , of the sheaf of orienting 2-forms (over the K -category associated to our embedded Kuranishi structure). Say that α is orienting at a curve f if the following holds:

- Whenever \hat{f}_i contains f , some wedge power of α is a volume form on $T_f\mathbf{F}(\hat{f}_i)\downarrow_{\mathbf{B}_0}$.
- Whenever f is contained in \hat{f}_i and \hat{f}_j , and $V_i \subset V_j$, the short exact sequence

$$0 \rightarrow T_f\mathbf{F}(\hat{f}_i)\downarrow_{\mathbf{B}_0} \rightarrow T_f\mathbf{F}(\hat{f}_j)\downarrow_{\mathbf{B}_0} \xrightarrow{D\bar{\partial}} V_j/V_i \rightarrow 0$$

is oriented when the first two terms are given the orientation from α , and the last term is oriented by its complex structure.

To make such a global choice of α , use Proposition 2.3. Note that 2-forms may be averaged or patched together using a partition of unity, and still satisfy this positivity condition at holomorphic curves. Accordingly, the sheaf of orienting 2-forms satisfies the patching and averaging axioms. Such 2-forms can also be extended to satisfy the positivity condition; therefore, this sheaf also satisfies the extension axiom. By reducing the size of our extensions $\hat{f}_i^\#$ if necessary, Proposition 2.3 constructs a global section α of the above sheaf of orienting 2-forms.

Note that α is always orienting at every holomorphic curve. Because the closure of \hat{f}_i is contained in $\hat{f}_i^\#$, α is orienting on an open neighborhood of each holomorphic curve in \hat{f}_i .

Definition 3.7 An orientation of a Kuranishi category \mathcal{K} is, for all \hat{f} in \mathcal{K} , an orientation of $\mathbf{F}(\hat{f})$ such that for all morphisms $\hat{f} \rightarrow \hat{g}$ and curves $f \in \hat{f}$, the following short exact sequence is oriented:

$$T_f\mathbf{F}(\hat{f}) \rightarrow T_f\mathbf{F}(\hat{g}) \xrightarrow{D\bar{\partial}\hat{g}} V(\hat{g})/V(\hat{f}).$$

Given a Kuranishi category \mathcal{K} over \mathbf{Z} , an orientation of \mathcal{K} relative to \mathbf{Z} is, for all \hat{f} in \mathcal{K} , an orientation of $T\mathbf{F}(\hat{f})\downarrow_{\mathbf{Z}}$ such that the following short exact sequence is oriented:

$$T_f\mathbf{F}(\hat{f})\downarrow_{\mathbf{Z}} \rightarrow T_f\mathbf{F}(\hat{g})\downarrow_{\mathbf{Z}} \xrightarrow{D\bar{\partial}\hat{g}} V(\hat{g})/V(\hat{f}).$$

Remark 3.8 Given an embedded Kuranishi structure on $\mathcal{M}_{\bullet}^{\text{st}}(\hat{\mathbf{B}})$, we may construct an orienting 2-form α , and then restrict the associated Kuranishi category to a neighborhood of \mathcal{K}^{hol} where α is orienting to obtain an extendable Kuranishi category \mathcal{K} that is oriented relative to \mathbf{B}_0 .

3.2 The sheaf S on a Kuranishi category

We shall need the following information to specify the sheaf, S , we use to define Gromov–Witten invariants.

Definition 3.9 (\mathcal{K}_{ϵ} and \mathcal{K}_C) Given a Kuranishi category \mathcal{K} , proper and oriented over \mathbf{Z} , choose continuous functions $\rho_i: \mathcal{K} \rightarrow [-1, 1]$ with the following properties:

- (i) At any holomorphic curve, $\rho_i > \frac{1}{2}$ for some i .
- (ii) For each i , there is some \mathcal{U}_i and associated vectorbundle V_i and chart \hat{f}_i/G_i from Definition 3.2 such that the set where $\rho_i \geq 0$ is in \mathcal{U}_i . (We don't require that the map from the indexing set for ρ_i to the indexing set for charts on \mathcal{K} be injective or surjective.)
- (iii) The subset of $\mathbf{F}(\hat{f}_i)$ where $\rho_i \geq 0$ is compact in the case that \mathbf{Z} is a single point, and more generally, the map from this subset to \mathbf{Z} is proper.
- (iv) Over any compact subset of \mathbf{Z} , there are only finitely many i such that ρ_i is somewhere positive.

Let \mathcal{K}_C be the substack of \mathcal{K}^{st} comprising all curves f in some \hat{f}_i where $\rho_i(f) \geq \frac{1}{2}$ and where $\bar{\partial}f \in V_j$ wherever any $\rho_j > 0$.

For any $0 < \epsilon < \frac{1}{2}$, define $\mathcal{K}_{\epsilon} \subset \mathcal{K}$ to be the (full) subcategory of \mathcal{K} consisting of families \hat{f} such that for some i , \hat{f} is locally isomorphic to an open subfamily of \hat{f}_i where $\rho_i > \epsilon$.

Remark 3.10 Our virtual moduli space shall be contained in \mathcal{K}_C . As well as $\mathcal{K}_C \rightarrow \mathbf{Z}$ being proper, \mathcal{K}_C has the virtue that any family in \mathcal{K}_{ϵ} covers an open substack of \mathcal{K}_C .

We can construct functions ρ_i satisfying Definition 3.9 using Lemma 2.6: Given any holomorphic curve f in \widehat{f}_i , we can choose some G_i -equivariant $\rho_i: F(\widehat{f}_i^\#) \rightarrow [-1, 1]$ so that the closure of the support of $\rho_i + 1$ is in \mathcal{U}_i , and so that $\rho_i(f) > 0$ and the subset of $F(\widehat{f}_i)$ where $\rho_i \geq 0$ is compact. Such a ρ_i can be extended to satisfy condition (ii) by applying Lemma 2.6 to $\rho_i + 1$. Note that condition (iii) only applies to $\{\rho_i \geq 0\}$ on $F(\widehat{f}_i)$, and not $F(\widehat{f}_j)$, so this extension of ρ_i automatically satisfies condition (iii). When Z is compact, a finite collection of such ρ_i will satisfy condition (i), and therefore all conditions of Definition 3.9. When Z is not compact, we can choose an exhaustion of Z by compact subsets Z_k , and construct each ρ_i to be greater than -1 only on some $Z_{k+1} \setminus Z_{k-1}$. Then, because the inverse image of Z_k in \mathcal{K}^{hol} is compact, there is a collection of such ρ_i satisfying conditions (i) and (iv) such that only finitely many ρ_i are greater than -1 over $Z_{k+1} \setminus Z_{k-1}$.

Note that although we can construct ρ_i so that $\{\rho_i \geq 0\} \subset F(\widehat{f}_i)$ is always compact, we only require this subset to be proper over Z so that Definition 3.9 is compatible with pullbacks; see Definition 4.2.

We shall define Gromov–Witten invariants using a sheaf of sections of $V(\widehat{f})$ over \widehat{f} . We need some notion (condition (ii) of Definition 3.12 below) of when these sections are “close enough” to the canonical section defined by $\bar{\partial}\widehat{f}$. This notion is provided by compatibly choosing metrics on these V_i with the property that, where $|\bar{\partial}\widehat{f}| \leq 1$, some ρ_j from Definition 3.9 is greater than $\frac{1}{2}$. Such a choice ensures that sections that are sufficiently close to $\bar{\partial}$ have their zero sets contained where these $\rho_j > \frac{1}{2}$.

Lemma 3.11 *On \mathcal{K}_ϵ consider the sheaf Met , where $\text{Met}(\widehat{f})$ is the set of metrics on the vectorbundle $V(\widehat{f})$ with the property that*

$$(1) \quad \text{on the subset where } |\bar{\partial}\widehat{f}| \leq 1, \text{ some } \rho_j > \frac{1}{2}.$$

Then Met obeys the extension, patching and averaging axioms, so Proposition 2.3 implies that there exists a globally defined metric that is a global section of Met .

Proof As such metrics may be averaged using a partition of unity, Met clearly satisfies the averaging and patching axioms. A section of Met locally exists around any curve $f \in \widehat{f}_i$, because either $\bar{\partial}f > 0$ and we can just choose a metric in which $|\bar{\partial}f| > 1$, or f is holomorphic, and some $\rho_j > \frac{1}{2}$ around f . Now suppose that $V_i \subset V_j$, and a section of Met on \widehat{f}_i has been chosen. Because \widehat{f}_i is locally equal to the transversely cut-out subset of \widehat{f}_j where $\bar{\partial}\widehat{f}_j \in V_i$, a metric on V_i over \widehat{f}_i can be locally extended

to a metric on V_j over \hat{f}_j . Because our condition (1) is always satisfied on an open subset, any such extension will locally meet condition (1), so Met also satisfies the extension axiom. \square

Definition 3.12 (the sheaf S) For a choice of \mathcal{K}_ϵ from Definition 3.9 and metric from Lemma 3.11, define a sheaf S over \mathcal{K}_ϵ as follows: For \hat{f} in \mathcal{K}_ϵ , define $S(\hat{f})$ to be the set of $C^{\infty,1}$ sections ν of $V(\hat{f})$ satisfying the following:

- (i) On an open neighborhood of where $\rho_i \geq 0$ on \hat{f} , $\nu - \bar{\partial}\hat{f}$ is a section of $V_i(\hat{f})$.
- (ii) When using the metric from Lemma 3.11,

$$|\bar{\partial} - \nu| < 1;$$

so, in particular, wherever $\nu = 0$, some $\rho_i > \frac{1}{2}$.

- (iii) ν is transverse to the zero section of $V(\hat{f})$.

To understand the purpose of item (i) above, note that $V_i(\hat{f})$ is defined on an open neighborhood of where $\rho_i \geq 0$, but even here, whenever the dimension of \hat{f} is greater than the dimension of \hat{f}_i , $V_i(\hat{f}) \subsetneq V(\hat{f})$, and $\bar{\partial}\hat{f}$ is transverse to this subbundle.

Given any section ν of $S(\hat{f})$, the intersection of ν with 0 defines a closed, $C^{\infty,1}$ exploded submanifold $\nu^{-1}(0) \subset \mathbf{F}(\hat{f})$. There is a canonical orientation of $\nu^{-1}(0)$ relative to \mathbf{Z} given by the relative orientation of \hat{f} and the complex orientation of $V(\hat{f})$. Intersection with 0 provides a natural transformation from S to a sheaf, E , defined as follows:

Definition 3.13 Define a sheaf of sets, E , on \mathcal{K}_ϵ as follows: Let $E(\hat{f})$ be the set of $C^{\infty,1}$ exploded submanifolds, $X \subset \mathbf{F}(\hat{f})$, satisfying the following conditions:

- (i) $X \subset \mathbf{F}(\hat{f})$ is closed, and locally defined by the transverse vanishing of some collection of $C^{\infty,1}$ functions.
- (ii) X is contained in \mathcal{K}_C (from Definition 3.9).
- (iii) Wherever $\rho_i > 0$, X is contained in $\bar{\partial}\hat{f}^{-1}(V_i(\hat{f}))$.
- (iv) X is oriented relative to \mathbf{Z} .

Pullbacks in E are naturally defined as inverse images: the pullback of X under $\iota: \hat{g} \rightarrow \hat{f}$ is the inverse image, $\mathbf{F}(\iota)^{-1}(X)$, of X under the induced map $\mathbf{F}(\iota): \mathbf{F}(\hat{g}) \rightarrow \mathbf{F}(\hat{f})$.

Is $F(\iota)^{-1}(X)$ really in E ? Remark 3.10 and the second condition above ensure that the image of $F(\iota)$ always intersects X in an open subset. As $F(\iota)$ is locally an isomorphism onto a subset defined by the transverse vanishing of some $C^{\infty,1}$ functions, $F(\iota)^{-1}(X)$ has all the required properties.

We show below that S satisfies the patching and extension axioms, but not the averaging axiom, so global sections of S may not exist. Instead, we will construct a weighted branched section of S whose intersection with 0 is a weighted branched section of E .

Lemma 3.14 *The sheaf S from Definition 3.12 obeys the patching axiom.*

Proof Consider a collection of sections v_k defined on open subsets U_k of \hat{f} . Using a partition of unity h_k subordinate to U_k , these sections v_k may be averaged to produce a section $v := \sum h_k v_k$ of $V(\hat{f})$. Such a section v automatically obeys all the conditions to be a section of $S(\hat{f})$, except that v may not be transverse to 0. Our section v also obeys the condition required by the patching axiom: given any $\hat{g} \rightarrow \hat{f}$, and $v' \in S(\hat{g})$ agreeing with the pullback of all v_k (where defined), the pullback of v is v' . Let X be the closure of the image of all $\hat{g} \rightarrow \hat{f}$ with some $v' \in S(\hat{g})$ satisfying the above condition.

Claim 3.15 *On $X \subset \hat{f}$, v is transverse to the zero section of $V(\hat{f})$.*

Suppose that $v = 0$ at $f \in X \subset \hat{f}$, as otherwise our claim follows trivially. Let V_j have the minimal dimension such that $\rho_j(f) \geq 0$. On a neighborhood of where $\rho_j \geq 0$, $v = \bar{\partial} \bmod V_j$. Therefore, $\bar{\partial}f \in V_j$, and $\bmod V_j$, the derivative of v at f is also equal to the derivative of $\bar{\partial}$ at f . Definition 3.2(ii) states that $\bar{\partial}\hat{f}$ at such an f is transverse to V_j , so we need only check that the derivative of v at f surjects onto V_j .

Suppose that U_k contains f . Then $v = v_k$ at the image of our $\hat{g} \rightarrow \hat{f}$. It follows that $Tv = Tv_k$ when restricted to the closure of the image of $TF(\hat{g}) \rightarrow TF(\hat{f})$. At f , this closure must include $T_f\bar{\partial}^{-1}(V_j(\hat{f}))$, which equals $(T_f v)^{-1}(V_j(f))$ and $(T_f v_k)^{-1}(V_j(\hat{f}))$. As v_k is transverse to 0, $T_f v_k$ restricted to this subspace surjects onto V_j ; therefore, $T_f v$ surjects onto V_j , so v is transverse to 0 at f , as required. This completes the proof of Claim 3.15.

To complete our proof of Lemma 3.14, we need the following:

Claim 3.16 *We can perturb v by a small section v of $V(\hat{f})$ so that $v + v$ is in S and v vanishes on a neighborhood of $X \subset \hat{f}$.*

In particular, the section $v + v$ of S satisfies the conditions required by the patching axiom.

To prove Claim 3.16, consider a curve $f \in \hat{f} \setminus X$. Let $V_{i(f)}$ have the smallest dimension such that $\rho_{i(f)}(f) \geq 0$. Let O_f be an open neighborhood of f such that

- $V_{i(f)}$ is defined on a neighborhood of the closure of O_f , and
- the closure of O_f does not intersect X , and also does not intersect the set where $\rho_j \geq 0$ for any j with $\dim V_j < \dim V_{i(f)}$.

Then $V_{i(f)}$ is defined on O_f , and on O_f we can add a small section of $V_{i(f)}$ to v and still satisfy all conditions of Definition 3.12 apart from possibly (iii). Condition (i) implies that v is equal to $\bar{\partial}$ mod $V_{i(f)}$ on O_f , so condition (ii) of Definition 3.2 implies that v is transverse to $V_{i(f)}$ on O_f , so we can achieve transversality in O_f by adding a small section of $V_{i(f)}$. We can now complete the proof of Claim 3.16 using a standard transversality argument, given below.

Choose an exhaustion of $F(\hat{f})$ by compact subsets, C_k . So, $F(\hat{f}) = \bigcup_k C_k$ where for all k , C_k is compact and contained in the interior of C_{k+1} . Suppose that we have constructed a v_k so that v_k vanishes on a neighborhood of X and $v + v_k$ satisfies conditions (i) and (ii) of Definition 3.12 and is transverse to 0 on a neighborhood of C_k . Let us construct v_{k+1} satisfying these requirements and equal to v_k when restricted to C_{k-1} . Cover C_{k+1} by the open set where $v + v_k$ is transverse to 0 — this open set contains X and C_k — and a finite collection of open subsets O_f satisfying the conditions above, and also contained in $C_{k+2} \setminus C_k$. Then choose a finite collection of sections w_1, \dots, w_N of $V(\hat{f})$, each with support in some O_f and with image in $V_{i(f)}$, so that the map

$$\tilde{v}: \mathbb{R}^N \times F(\hat{f}) \rightarrow V(\hat{f})$$

defined by

$$\tilde{v}(t_1, \dots, t_N, f) := v(f) + v_k(f) + \sum_{i=1}^N t_i w_i(f)$$

is transverse to the zero section of V , at least when restricted to an open neighborhood O of $\mathbb{R}^N \times C_{k+1}$. To achieve such transversality, it suffices that, restricted to each O_f , the sections w_1, \dots, w_N generate $V_{i(f)}$. A finite collection of such sections exists because the closure of O_f is compact and contained in the domain of definition of V_i . Then $\tilde{v}^{-1}(0) \cap O$ is a $C^{\infty,1}$ exploded manifold. The projection $\pi: \tilde{v}^{-1}(0) \cap O \rightarrow \mathbb{R}^N$ factors through a smooth map from a manifold on each stratum, so Sard's theorem

applies and the critical locus of π in \mathbb{R}^N has measure 0. Therefore, there exists a regular point of π , $(t_1, \dots, t_n) \in \mathbb{R}^N$, arbitrarily close to 0. For any such regular point, $v + v_k + \sum_i t_i w_i$ is transverse to 0 on some neighborhood of C_{k+1} .

Because the w_i are compactly supported, condition (ii) of Definition 3.12 is still satisfied by $v + v_k + \sum_i t_i w_i$ so long as (t_1, \dots, t_N) is small enough. Condition (i) holds because each w_i is a section of $V_{i(f)}$ on some O_f . Moreover, $v_{k+1} := v_k + \sum_i t_i w_i$ agrees with v_k on C_{k-1} , and vanishes on a neighborhood of X because each w_i is supported in some O_f .

The section $v = \lim_{k \rightarrow \infty} v_k$ agrees with v_k on C_{k-1} for all k . So, $v + v$ is transverse to 0, and is a section of S . Moreover, v vanishes on a neighborhood of X , so v satisfies the requirements of Claim 3.16. This completes the proof of Claim 3.16 and Lemma 3.14. \square

Lemma 3.17 *The sheaf S from Definition 3.12 obeys the extension axiom.*

Proof As S obeys the patching axiom, we need only verify the local existence of extensions, and the local existence of sections of S .

Sections of S locally exist around any $f \in \hat{f}$. Let V_i have the smallest dimension such that $\rho_i(f) \geq 0$. If $\bar{\partial}f \neq 0$, then in a neighborhood of f , $\bar{\partial}$ is a section of S . If $\bar{\partial}f = 0$, then $T_f \bar{\partial} \hat{f}$ is transverse to V_i , so $\bar{\partial}$ plus some small section of V_i will be transverse to the zero section at f , and also obey all the other conditions of Definition 3.12 on a neighborhood of f .

Now suppose that we have a curve f in \hat{g} and a morphism $\hat{g} \rightarrow \hat{f}$ in \mathcal{K}_ϵ . Again, let V_i have the smallest dimension such that $\rho_i(f) \geq 0$. Given a section $v \in S(\hat{g})$, we must construct an extension of v to v' around f in \hat{f} . As specified by Definitions 2.1 and 3.2, an open neighborhood of f in \hat{g} is isomorphic to an open subset of \hat{f}_j , and an open neighborhood of f in \hat{f} is isomorphic to an open subset of \hat{f}_k , where $V_i \subset V_j \subset V_k$ on a neighborhood of f in \hat{f}_k . Moreover, the morphism $\hat{g} \rightarrow \hat{f}$ is locally an isomorphism onto the transverse intersection of $\bar{\partial} \hat{f}$ with V_j . Note that $\bar{\partial} - v$ is a section of $V_i \subset V_j \subset V_k$. We may therefore extend $\bar{\partial} - v$ to a section $\bar{\partial} - v'$ of V_i on a neighborhood of f within \hat{f} . At f , the resulting section v' is transverse to V_j , and transverse to the zero section of V_j restricted to $(v')^{-1}V_j$. It is therefore transverse to the zero section of $V_k(\hat{f}) = V(\hat{f})$ at f . Therefore, v' is the required local extension of v and S obeys the extension axiom. \square

Remark 3.18 We may impose any extra transversality condition satisfied by a generic section of S , and the patching axiom and extension axiom will still hold, with an identical proof, so long as the required notion of transversality is independent of extensions. For example, given a submersion $\pi: \mathcal{K} \rightarrow X$, we could impose the extra condition that π restricted to the intersection of ν with 0 must be transverse to a given map of another compact exploded manifold to X .

3.3 Weighted branched sections of sheaves on K -categories

The sheaf S from Definition 3.12 obeys the patching and extension axioms; however, it does not obey the averaging axiom, and G -equivariant sections of S may not even exist, so global sections of S can't be constructed. Instead, we construct weighted branched sections of S . Our approach is essentially that originally taken by Fukaya and Ono [10]; however, we use weighted branched sections with a particular branching structure. Our fixed branching structure deals with complications that arise when patching weighted branched sections together. (Even without specifying a particular branching structure, our version of weighted branched sections is subtly different from the definition given by Cieliebak, Mundet i Rivera and Salamon [6] or the intrinsic definition given by McDuff [18].)

In the formalism below, weighted branched sections of a sheaf S over O are labeled by a finite measure space, $(I(O), \mu)$, so each $i \in I(O)$ labels a section, $\nu(i)$, of S with a weight $\mu(i)$. Given a map $O' \rightarrow O$, we want to pull back weighted branched sections to O' , which requires a measure-preserving map $I(O) \rightarrow I(O')$. Sometimes this map will send some i and j to the same point, in which case we require that $\nu(i)$ and $\nu(j)$ coincide on a neighborhood of the image of O' ; our formalism includes an equivalence relation, \equiv , on $I(O)$ to keep track of such requirements.

Definition 3.19 A weighted branched cover of a K -category, \mathcal{K} (Definition 2.1), is a contravariant functor I with the following domain and target:

- The domain of I is a full subcategory $\mathcal{O}_I \subset \mathcal{K}^{\text{st}}$ such that the following holds:
 - All O in \mathcal{O}_I are connected.
 - If O_1 in \mathcal{K}^{st} is connected and there is a morphism $\iota: O_1 \rightarrow O_2$ such that $O_2 \in \mathcal{O}_I$, then $O_1 \in \mathcal{O}_I$.
 - For every family \hat{f} in \mathcal{K} , there is an open cover of \hat{f} contained in \mathcal{O}_I .

- The target of I is the category with objects triples $(I(O), \mu, \equiv)$ where
 - $I(O)$ is a finite set (thought of as a set of sections),
 - μ is a probability measure on $I(O)$ such that every point in $I(O)$ has positive rational measure,
 - \equiv is an equivalence relation on $I(O)$ (thought of as a warning that our sections are related through needing to be glued somewhere),

and morphisms

$$\iota^*: (I(O_2), \mu, \equiv) \rightarrow (I(O_1), \mu, \equiv)$$

consisting of measure-preserving maps $\iota^*: I(O_2) \rightarrow I(O_1)$ such that for any $i, j \in I(O_2)$, $i \equiv j$ if $\iota^*i \equiv \iota^*j$.

We make the additional assumption that, for any curve f and morphism $f \rightarrow O$ in \mathcal{O}_I , there exists a neighborhood $U \subset O$ of f such that $(I(U), \mu, \equiv) \rightarrow (I(f), \mu, \equiv)$ is an isomorphism.

Say that i and j in $I(O)$ are separated at f if their images in $I(f)$ are not equivalent.

Say that I has trivial stabilizers if, whenever $\psi: \hat{f} \rightarrow \hat{f}$ is a nontrivial automorphism in $\mathcal{K} \cap \mathcal{O}_I$ that fixes $[f] \in [F(\hat{f})]$, then for all $i \in I(\hat{f})$, i and ψ^*i are separated at f .

For example, given a finite group G and a G -fold cover of a manifold (in other words, a principal G -bundle), a weighted branched cover I could be defined as follows. Let \mathcal{O}_I be the category of maps from connected manifolds O into our space pulling back our G -fold cover to a trivial cover. Then define $I(O)$ to be the set of sections of this trivial G -fold cover over O with the discrete equivalence relation and the counting measure divided by $|G|$. The curious reader is invited to think about what goes wrong if we try to define $I(O)$ for disconnected O .

We can create another weighted branched cover from the above one by gluing together branches outside some closed set C . Define \mathcal{O}_I as above, but now let $I(O)$ have a single point if O does not intersect C , and let $I(O)$ be the set of sections over O if O intersects C . Give such an $I(O)$ the G -invariant probability measure; use the discrete equivalence relation if O is contained in the interior of C , and use the trivial indiscrete equivalence relation otherwise, so distinct sections of $I(O)$ are separated on the interior of C .

We can build up more complicated examples by taking products of weighted branched covers.

If \hat{f} is in $\mathcal{O}_I \cap \mathcal{K}$, define $S^I(\hat{f})$ to be the set of maps $\nu: I(\hat{f}) \rightarrow S(\hat{f})$ satisfying the following condition: if two sections, i and j , of I are not separated at $f \in \hat{f}$, then $\nu(i)$ and $\nu(j)$ agree on some neighborhood of f . (Equivalently, for any such ν , there exists an open cover $\{U\}$ of \hat{f} such that if $i \equiv j$ in U , then $\nu(i) = \nu(j)$ on U .) Given any morphism $\iota: \hat{g} \rightarrow \hat{f}$ in \mathcal{O}_I , the map $\iota^*: I(\hat{f}) \rightarrow I(\hat{g})$ is surjective, and if $\iota^*i = \iota^*j$, then i and j are not separated anywhere on the image of \hat{g} , so $\nu(i)$ and $\nu(j)$ agree on the image of \hat{g} . It follows that there is a unique pullback map $\iota^*: S^I(\hat{f}) \rightarrow S^I(\hat{g})$ fitting into the commutative diagram

$$\begin{array}{ccc} I(\hat{g}) & \xleftarrow{\iota^*} & I(\hat{f}) \\ \downarrow \iota^* \nu & & \downarrow \nu \\ S(\hat{g}) & \xleftarrow{\iota^*} & S(\hat{f}) \end{array}$$

Extend the definition of S^I to be a sheaf on \mathcal{K} as follows:

Definition 3.20 Given a weighted branched cover I of \mathcal{K} and a sheaf S on \mathcal{K} , define the sheaf S^I as follows:

- For a given family \hat{f} in \mathcal{K} , let $\mathcal{O}_{I,\hat{f}}$ be the category of families in $\mathcal{O}_I \cap \mathcal{K}$ with a given map to \hat{f} .
- Define $I_{\hat{f}}$ as the composition of I with the functor $\mathcal{O}_{I,\hat{f}} \rightarrow \mathcal{O}_I$ that forgets the map to \hat{f} , and define $S_{\hat{f}}$ as the composition of S with the functor $\mathcal{O}_{I,\hat{f}} \rightarrow \mathcal{K}$.
- Define $S^I(\hat{f})$ to be the set of natural transformations⁹ $\nu: I_{\hat{f}} \rightarrow S_{\hat{f}}$ such that for any given O in $\mathcal{O}_{I,\hat{f}}$, and $i, j \in I(O)$ not separated at f , $\nu(i) = \nu(j)$ on a neighborhood of f .
- Given $\iota: \hat{g} \rightarrow \hat{f}$, define $\iota^* \nu$ to be the pullback of ν using the obvious functor $\mathcal{O}_{I,\hat{g}} \rightarrow \mathcal{O}_{I,\hat{f}}$.

If \hat{f} is in $\mathcal{O}_I \cap \mathcal{K}$, $S^I(\hat{f})$ coincides with our easier definition above. This S^I is a sheaf because S is a sheaf. Moreover, we prove below that S^I obeys the patching, extension and averaging axioms if S obeys the patching and extension axioms and I has trivial stabilizers.

⁹For the purposes of saying what a natural transformation is, consider the codomain of $I_{\hat{f}}$ and $S_{\hat{f}}$ to be the category of sets.

Lemma 3.21 *If S obeys the patching axiom, then S^I does too.*

Proof We must prove the patching axiom for S^I . In particular, we must show that given an open cover $\{U_k\}$ of \hat{f} and sections $v_k \in S^I(U_k)$ for all k , there exists a section $v \in S^I(\hat{f})$ satisfying the following property:

Given any morphism $\hat{g} \rightarrow \hat{f}$ in \mathcal{K} and section $v' \in S^I(\hat{g})$ agreeing with the pullback of v_k (where defined) for all k , the pullback of v is v' .

We shall show that the patching axiom for S^I follows from the following special case:

Claim 3.22 *The patching axiom holds in the case when the cover $\{U_k\}$ has two elements, $\{U_1, U_2\}$, and U_2 is in \mathcal{O}_I .*

To prove Claim 3.22, consider a connected component O of $U_1 \cap U_2$. As $O \in \mathcal{O}_I$, the v_i determine maps

$$v_i: I(O) \rightarrow S(O).$$

Below, we shall construct v on O using the patching axiom for S on the individual sections $v_i(k)$, while ensuring that v extends as v_1 on $U_1 \setminus U_2$, and v_2 on $U_2 \setminus U_1$. Repeating the construction for every connected component of $U_1 \cap U_2$, we obtain a section v of S^I on $U_1 \cup U_2$ that agrees with v_i where only v_i is defined. To check the patching axiom for such a section, it suffices to check the patching axiom on O (and the other connected components of $U_i \cap U_j$) individually.

Identify $I(O)$ with the set $\{1, \dots, n\}$. Use the patching axiom for S to patch together $v_1(1)$ and $v_2(1)$ to create a $v(1) \in S(O)$ agreeing with $v_i(1)$ on a neighborhood of the boundary of O within U_i . Such a $v(1)$ obeys the requirements of the patching axiom for S^I on O because it obeys the requirements of the patching axiom for S ; in particular, if $\iota: \hat{g} \rightarrow O$ pulls back v_1 and v_2 to v' (and \hat{g} is connected and hence in \mathcal{O}_I), then the patching axiom for S implies that $\iota^*v_i(1) = v'(\iota^*(1))$.

On some neighborhood $O_{1=2}$ of the set where the sections 1 and 2 of $I(O)$ are not separated, $v_1(1) = v_1(2)$ and $v_2(1) = v_2(2)$. We can cover O by open subsets $O_{2;1}$, $O_{2;2}$ and $O_{1=2}$ so that (see Figure 1)

- the closure of $O_{2;i}$ does not intersect the set where 1 and 2 are not separated,
- the closure of $O_{2;1}$ within U_2 does not intersect the boundary of O within U_2 , and
- the closure of $O_{2;2}$ within U_1 does not intersect the boundary of O within U_1 .




Figure 1

Such a covering exists because the set where the sections 1 and 2 of $I(O)$ are not separated is closed. Now use the patching axiom for S to patch together $v_1(2)|_{O_{2;1}}$, $v_2(2)|_{O_{2;2}}$ and $v(1)|_{O_{1 \equiv 2}}$ to create $v(2)$. Our $v(2)$ obeys the requirements of the patching axiom on O , agrees with $v(1)$ on a neighborhood of where 1 and 2 are not separated, and agrees with $v_i(2)$ on a neighborhood of the boundary of O within U_i .

Inductively continuing this construction gives sections $v(k) \in S(O)$, obeying the requirements of the patching axiom on O , and agreeing with $v_i(j)$ on a neighborhood of the boundary of O within U_i , so that $v(j) = v(k)$ on a neighborhood of the set where the sections j and k of $I(O)$ are not separated. In particular, after constructing $v(j)$ for all $j < k$, we can cover O by open subsets

- $O_{j \equiv k}$ for $j < k$, where $v_1(k) = v_1(j)$ and $v_2(k) = v_2(j)$, covering the set where j and k are not separated;
- $O_{k;1}$ with closure in O not intersecting the set where any $j < k$ is not separated from k , and with closure within U_2 not intersecting the boundary of O within U_2 ; and
- $O_{k;2}$ with closure in O not intersecting the set where any $j < k$ is not separated from k , and with closure within U_1 not intersecting the boundary of O within U_1 .

We can then construct $v(k)$ satisfying the required conditions using the patching axiom with $v(j)|_{O_{j \equiv k}}$ and $v_i(k)|_{O_{k;i}}$. After completing this construction for all $k \in I(O)$,

we have constructed a section $\nu \in S^I(O)$ agreeing with ν_i on a neighborhood of the boundary of O within U_i , and obeying the requirements of the patching axiom on O : for any map $\iota: \hat{g} \rightarrow O$ with connected domain pulling back ν_i to ν' , $\iota^*\nu(k) = \nu'(\iota^*k)$, so $\iota^*\nu = \nu'$. As S^I is a sheaf, there is a section $\nu \in S^I(\hat{f})$ which agrees with our already constructed ν in O , is similarly defined on every other component of $U_1 \cap U_2$ and which agrees with ν_i elsewhere. As the requirements of the patching axiom are local, and such a ν satisfies the requirements of the patching axiom locally, ν satisfies the requirements of the patching axiom. This completes the proof of Claim 3.22.

Now consider an arbitrary open cover of \hat{f} . This arbitrary open cover may be replaced by a countable, locally finite cover $\{U_i\}$ consisting of sets in \mathcal{O}_I contained in one of the open sets of the original cover. Then apply Claim 3.22 inductively to patch together some patched-together section on $\bigcup_{i=1}^n U_i$ and ν_{n+1} on U_{n+1} . This procedure constructs a global section ν of $S^I(\hat{f})$ obeying the requirements of the patching axiom. \square

Lemma 3.23 *If S satisfies the patching and extension axioms then S^I obeys the patching and extension axioms.*

Proof In light of Lemma 3.21, we just need to prove that S^I obeys the extension axiom. As S^I obeys the patching axiom, the extension axiom follows from the local existence of sections of S^I , and the local existence of extensions.

The local existence of sections of S^I is easy: For any O in \mathcal{O}_I , $S(O)$ is nonempty because S obeys the extension axiom. A constant map $I(O) \rightarrow S(O)$ suffices to define a section of $S^I(O)$.

Given a morphism $\iota: \hat{g} \rightarrow \hat{f}$ in \mathcal{K} and $f \in \hat{g}$, can we locally extend $\nu \in S^I(\hat{g})$? This local question can be answered without losing generality by shrinking \hat{g} and \hat{f} until \hat{g} is in \mathcal{O}_I , ν agrees on sections of $I(\hat{g})$ not separated at f , and sections of $I(\hat{f})$ separated at f are separated everywhere. Then, for each section i of $I(\hat{f})$, the extension axiom for S provides an extension $\nu'(i)$ of $\nu(\iota^*i)$. If i and j are not separated at f , we can choose $\nu'(i) = \nu'(j)$, because $\iota^*(i)$ and $\iota^*(j)$ are not separated at f , so $\nu(\iota^*i) = \nu(\iota^*j)$. As all other pairs of sections of $I(\hat{f})$ are separated everywhere, this ν' defines a section in $S^I(\hat{f})$. \square

Lemma 3.24 *If I is a weighted branched cover of \mathcal{K} with trivial stabilizers and S is a sheaf on \mathcal{K} satisfying the patching axiom, then S^I satisfies the patching and averaging axioms.*

Proof Lemma 3.21 implies that S^I satisfies the patching axiom, so we need only verify the averaging axiom. The idea is to take a given section ν and patch together pieces of ν to create a G -equivariant section ν' .

Claim 3.25 Suppose that the following holds:

- $\{U_1, U_2\}$ form a G -invariant open cover of \hat{f} .
- $\nu_i \in S^I(U_i)$.
- ν_1 is G -equivariant.
- Each connected component of U_2 is in \mathcal{O}_I .
- For each connected component O of U_2 , there exists a subset $I_0 \subset I(O)$ such that
 - if $i \in I_0 \subset I(O)$ and $j \equiv i$, then $j \in I_0$;
 - the G -orbit of I_0 contains $I(O)$;
 - if the action of some $g \in G$ sends an element of I_0 to another element of I_0 , then g acts trivially on O .

Then, there exists a G -equivariant section ν' of $S^I(\hat{f})$, agreeing with ν_1 on $U_1 \setminus U_2$ and agreeing with ν_2 where ν_2 is G -equivariant and the same as ν_1 . More precisely, given any G -equivariant map $\iota: \hat{h} \rightarrow \hat{f}$ with image in U_2 , $\iota^*\nu' = \iota^*\nu_2$ so long as $\iota^*\nu_2$ is G -equivariant and equals $\iota^*\nu_1$ on the pullback of U_1 .

To prove Claim 3.25, we shall first define ν' on a connected component O of U_2 . For $j \in I_0$, patch together¹⁰ $\nu_1(j)$ and $\nu_2(j)$ to obtain $\nu'(j) \in S(O)$ agreeing with $\nu_1(j)$ on a neighborhood of the closure of $O \cap U_1$ within U_1 . Similarly to the proof of Claim 3.22, choose these $\nu'(j)$ so that $\nu'(j) = \nu'(k)$ on a neighborhood of the set where j and k are not separated, and so that the requirements of the patching axiom are satisfied by these $\nu'(k)$.

For any morphism g in G and $j \in I_0$, we can then define $\nu'(g^*j)$ to be $g^*\nu'(j)$. Whenever g acts nontrivially on O , g^*j is separated from all $i \in I_0$, and the property of being separated is preserved under the action of G , so there are no further conditions required of our sections $\nu'(g^*j)$, and the resulting ν' defines a G -equivariant section of S^I on the orbit of O . Because ν_1 is equivariant, ν' agrees with ν_1 on an open

¹⁰Here we have abused notation slightly: ν_2 is defined on some open subset of O ; by $\nu_2(j)$ we mean the section which, on connected components, is $\nu_2(j')$, where j' is the restriction of j .

neighborhood of the closure of the G -orbit of $O \cap U_1$ within U_1 . Therefore, ν' may be extended to a G -equivariant section on the union of the G -orbit of O with U_1 by setting $\nu' = \nu_1$ everywhere else. We may similarly extend the definition of ν' to all other components of U_2 .

As our $\nu'(j)$ for $j \in I_0$ obeyed the requirements of the patching axiom, ν' satisfies the required G -equivariant version of the patching axiom. This completes the proof of Claim 3.25.

Construct a G -equivariant ν' inductively using Claim 3.25 as follows. Every curve in \hat{f} has a G -invariant open neighborhood U_f satisfying the requirements of U_2 in Claim 3.25; in particular, U_f can be the G -orbit of a connected open neighborhood O of f such that the following holds:

- O is invariant under the subgroup $H \leq G$ which is the weak stabilizer of f (consisting of elements fixing the image of f in $[F(\hat{f})]$).
- For $g \in G \setminus H$, $g(O) \cap O = \emptyset$.
- O is in \mathcal{O}_I and is small enough that $I(O) \rightarrow I(f)$ is an isomorphism.

For any $i \in I(O)$, i and g^*i are separated for any $g \in H$ acting nontrivially on O (because I has trivial stabilizers). As \equiv is a G -invariant equivalence relation on $I(O)$, there is a subset $I_0 \subset I(O)$ satisfying the requirements of Claim 3.25.

Choose a countable, locally finite open cover $\{O_i\}$ of \hat{f} using sets satisfying the above conditions on U_f . Use Claim 3.25 with $U_1 = \emptyset$, $U_2 = O_1$ and $\nu_2 = \nu$ to construct a G -equivariant section ν' on O_1 . Then inductively apply Claim 3.25 with $U_1 = \bigcup_{i=1}^n O_i$ and ν_1 the already constructed ν' , and $U_2 = O_{n+1}$ and $\nu_2 = \nu$. Claim 3.25 implies that the resulting patched together ν' on $\bigcup_{i=1}^{n+1} O_i$ is G -equivariant and satisfies the condition required by the averaging axiom. As the cover $\{O_i\}$ is locally finite, our sequence of choices ν' converge to a section $\nu' \in S(\hat{f})$ obeying the conditions required by the averaging axiom. \square

4 Construction of $[\mathcal{K}]$

In this section we construct a virtual fundamental class $[\mathcal{K}]$ for a Kuranishi category \mathcal{K} , and show that all choices are cobordant. To state what it means for a choice to be cobordant, we need the notion of a pullback of a Kuranishi category introduced below.

4.1 Pullbacks of Kuranishi categories

Recall that we are considering Kuranishi categories \mathcal{K} over \mathbf{Z} , where \mathbf{Z} is an exploded orbifold — a Deligne–Mumford stack over the category of exploded manifolds. The following is a standard definition of a representable map of stacks:

Definition 4.1 A map

$$\mathbf{Z}' \rightarrow \mathbf{Z}$$

of Deligne–Mumford stacks over the category of exploded manifolds is representable if, given any submersion $\mathcal{A} \rightarrow \mathbf{Z}$ from an exploded manifold \mathcal{A} , the fiber product $\mathcal{A} \times_{\mathbf{Z}} \mathbf{Z}'$ is represented by an exploded manifold.

For more general stacks, say that a submersion $\mathcal{X}' \rightarrow \mathcal{X}$ of stacks is representable if given any map $\mathcal{A} \rightarrow \mathcal{X}$ from an exploded manifold \mathcal{A} , the fiber product $\mathcal{X}' \times_{\mathcal{X}} \mathcal{A}$ is also represented by an exploded manifold.

In the language of orbifolds, a representable map is one that is injective on stabilizers. For example, if \mathbf{Z}' is an exploded manifold, $\mathbf{Z}' \rightarrow \mathbf{Z}$ is always representable. In the definition of pullback of \mathcal{K} over $\mathbf{Z}' \rightarrow \mathbf{Z}$ below, we require that $\mathbf{Z}' \rightarrow \mathbf{Z}$ is representable, so that $\hat{f}_i \times_{\mathbf{Z}} \mathbf{Z}'$ is a family of curves parametrized by an exploded manifold, rather than a more general Deligne–Mumford stack. A notion of pullback for nonrepresentable maps $\mathbf{Z}' \rightarrow \mathbf{Z}$ exists, but there would either be some extra choices, or we would need arbitrary Deligne–Mumford stacks to take the place of \hat{f}/G .

Definition 4.2 Given any Kuranishi category \mathcal{K} over \mathbf{Z} and representable map $\mathbf{Z}' \rightarrow \mathbf{Z}$, the pullback of \mathcal{K} is a Kuranishi category \mathcal{K}' over \mathbf{Z}'

$$\begin{array}{ccc} \mathcal{K}' & \longrightarrow & \mathcal{K} \\ \downarrow & & \downarrow \\ \mathbf{Z}' & \longrightarrow & \mathbf{Z} \end{array}$$

where:

- $(\mathcal{K}')^{\text{st}}$ is defined to be $\mathcal{K}^{\text{st}} \times_{\mathbf{Z}} \mathbf{Z}'$.
- \hat{f}'_i is $\hat{f}_i \times_{\mathbf{Z}} \mathbf{Z}'$ and $G'_i = G_i$. (For the purposes of defining \mathcal{K}' , we discard the indices i for which $\hat{f}_i \times_{\mathbf{Z}} \mathbf{Z}'$ is empty.)
- \mathcal{U}'_i and V'_i are the pullbacks of \mathcal{U}_i and V_i under the map $(\mathcal{K}')^{\text{st}} \rightarrow \mathcal{K}^{\text{st}}$.

- V' is defined as in Definition 3.2(i), so $V'(\hat{f}) = V'_i(\hat{f})$ if \hat{f} is locally isomorphic to \hat{f}'_i . In particular, $V'(\hat{f} \times_{\mathbf{Z}} \mathbf{Z}')$ is the pullback of $V(\hat{f})$.
- $\bar{\partial}'(\hat{f} \times_{\mathbf{Z}} \mathbf{Z}')$ is the pullback of $\bar{\partial}\hat{f}$ under the natural map $\hat{f} \times_{\mathbf{Z}} \mathbf{Z}' \rightarrow \hat{f}$.

Similarly, given a map $\mathcal{K}^{\text{st}} \rightarrow \mathcal{X}$ and a representable submersion $\mathcal{X}' \rightarrow \mathcal{X}$, define the pullback \mathcal{K}' of \mathcal{K} as above with \mathcal{X} taking the place of \mathbf{Z} .

For example, \mathcal{X} might be the moduli stack of $C^{\infty,1}$ curves in \mathbf{B} , and \mathcal{X}' might be the moduli stack of $C^{\infty,1}$ curves in some refinement of \mathbf{B} , or the moduli stack of $C^{\infty,1}$ curves in \mathbf{B} with an extra choice of marked point.

Remark 4.3 Pullback of Kuranishi categories is compatible with many notions:

- Any extension (Definition 2.1) of \mathcal{K} also pulls back to an extension of \mathcal{K}' .
- An orientation of \mathcal{K} over \mathbf{Z} (Definition 3.7) pulls back to an orientation of \mathcal{K}' over \mathbf{Z}' . Similarly, if the fibers of $\mathcal{X}' \rightarrow \mathcal{X}$ are oriented, an orientation of \mathcal{K} pulls back to an orientation of \mathcal{K}' .
- \mathcal{K}' is proper or complete over \mathbf{Z}' (Definition 3.5) if \mathcal{K} is proper or complete over \mathbf{Z} . Similarly, if $\mathcal{X}' \rightarrow \mathcal{X}$ is proper or complete, \mathcal{K}' is proper or complete if \mathcal{K} is.
- All choices involved in defining \mathcal{K}_ϵ (Definition 3.9) pull back to define \mathcal{K}'_ϵ , which happens to coincide with the pullback of \mathcal{K}_ϵ in the above sense. (When we pull back \mathcal{K}_ϵ , we drop the indices j for which the pullback of ρ_j is negative.) Also \mathcal{K}'_C is the pullback of \mathcal{K}_C .
- Any sheaf S defined on \mathcal{K}^{st} pulls back to a sheaf defined on $(\mathcal{K}')^{\text{st}}$, and any global section of such an S over \mathcal{K} pulls back to a global section over \mathcal{K}' .
- Any weighted branched cover I of \mathcal{K}_ϵ (Definition 3.19) pulls back to a weighted branched cover I' of \mathcal{K}'_ϵ .
- For S the sheaf from Definition 3.12, if $\mathbf{Z}' \rightarrow \mathbf{Z}$ is a submersion, any global section of S^I over \mathcal{K}_ϵ pulls back to a global section of $S'^{I'}$ over \mathcal{K}'_ϵ . Similarly, given any representable submersion $\mathcal{X}' \rightarrow \mathcal{X}$, any global section of S^I over \mathcal{K}_ϵ pulls back to a global section of $S'^{I'}$ over \mathcal{K}'_ϵ .

Definition 4.4 Say that two Kuranishi categories \mathcal{K}_0 and \mathcal{K}_1 that are proper (or complete) and oriented over \mathbf{Z} are cobordant over \mathbf{Z} if there exists a Kuranishi category \mathcal{K} proper (or complete) and oriented over $\mathbf{Z} \times \mathbb{R}$ such that \mathcal{K}_i (with its orientation relative to \mathbf{Z}) is the pullback of \mathcal{K} under the inclusion of \mathbf{Z} over $i \in \mathbb{R}$. Call \mathcal{K} a cobordism between \mathcal{K}_i .

If $\mathcal{K}_i^{\text{st}}$ are substacks of a stack \mathcal{X} with a map to \mathbf{Z} , say that \mathcal{K}_i are cobordant within \mathcal{X} if \mathcal{K}^{st} is a substack of $\mathcal{X} \times \mathbb{R}$ and the map $\mathcal{K} \rightarrow \mathbf{Z} \times \mathbb{R}$ is the restriction of the map $\mathcal{X} \times \mathbb{R} \rightarrow \mathbf{Z} \times \mathbb{R}$.

4.2 Construction of $[\mathcal{K}]$

Lemma 4.5 *Suppose that a Kuranishi category \mathcal{K} is proper and oriented over \mathbf{Z} . Then there exists a \mathcal{K}_ϵ and a weighted branched cover I of \mathcal{K}_ϵ with trivial stabilizers, satisfying the conditions of Definitions 3.9 and 3.19.*

Moreover, any two choices of \mathcal{K}_ϵ and I are cobordant in the following sense: Suppose that \mathcal{K}_0 and \mathcal{K}_1 are cobordant (within \mathcal{X}) over \mathbf{Z} . Then, given any two choices of $\mathcal{K}_{i,\epsilon}$ and weighted branched covers I_i , there is some cobordism \mathcal{K} (within \mathcal{X}) between \mathcal{K}_i , along with a construction of \mathcal{K}_ϵ and a weighted branched cover I such that \mathcal{K}_i , $\mathcal{K}_{i,\epsilon}$ and I_i are the pullback of \mathcal{K} , \mathcal{K}_ϵ and I , respectively, under the inclusions of \mathbf{Z} over $i \in \mathbb{R}$.

Proof To construct \mathcal{K}_ϵ , we must choose functions ρ_i as in Definition 3.9. As discussed after Definition 3.9, such functions exist because of Lemma 2.6, the fact that Kuranishi categories are extendable, our assumption that $\mathcal{K}^{\text{hol}} \rightarrow \mathbf{Z}$ is proper, and the fact that each chart \hat{f}_i/G_i covers an open substack of \mathcal{K}^{hol} .

Definition 2.1(ii) states that \hat{f}_i/G_i represents a substack of \mathcal{K}^{st} . Suppose that \hat{f}_i has the largest dimension of any family in \mathcal{K} containing f . Then some G -invariant neighborhood of f in \hat{f}_i covers an open substack U of $\mathcal{K}_\epsilon^{\text{st}} \subset \mathcal{K}^{\text{st}}$. Moreover, because Kuranishi categories are, by definition, extendable, there is a neighborhood U' of f with closure contained in U .

Define a weighted branched cover I_f of \mathcal{K}_ϵ as follows: Define \mathcal{O}_{I_f} to be the full subcategory of $\mathcal{K}_\epsilon^{\text{st}}$ with objects all connected families O such that either

- O does not intersect the closure of U' — in this case $I_f(O)$ is the probability space with a unique element — or
- O intersects the closure of U' , is contained entirely inside U , and the corresponding G -fold cover $O \times_{\hat{f}_i/G_i} \hat{f}_i$ of O defined by the following fiber product diagram is trivial:

$$\begin{array}{ccc} O \times_{\hat{f}_i/G_i} \hat{f}_i & \longrightarrow & \hat{f}_i \\ \downarrow & & \downarrow \\ O & \longrightarrow & \hat{f}_i/G_i \end{array}$$

In this case, define $I_f(O)$ to be the set of sections of the above G_i -fold cover of O with the G_i -invariant probability measure. If O is contained in U' , then use the discrete equivalence relation on $I_f(O)$, otherwise, use the trivial indiscrete equivalence relation on $I_f(O)$.

Given any morphism $\iota: O' \rightarrow O$ in \mathcal{O}_{I_f} , ι^* is either uniquely determined because $I_f(O')$ has a unique point, or ι^* is induced by the natural map between G_i -fold covers. To check that I_f defines a weighted branched cover, we must still check that, given any morphism $g \rightarrow O$ from a curve in \mathcal{O}_{I_f} , there exists a neighborhood O' of the image of g such that the induced map $I_f(O') \rightarrow I_f(g)$ is an isomorphism. If g is in the boundary of U' , then the corresponding map $I_f(O) \rightarrow I_f(g)$ is an isomorphism; if g is not in the boundary of U' , then $I_f(O) \rightarrow I_f(g)$ is an isomorphism so long as O is small enough not to intersect the boundary of U' . Therefore, I_f obeys the final condition we required of weighted branched covers.

Because $\widehat{f_i}/G_i$ represents a substack of \mathcal{K}^{st} , I_f is separating at any curve in U' .

Now, make corresponding choices of I_{f_i} , U_i and U'_i so that $\{U'_i\}$ is a cover of \mathcal{K}_ϵ and so that each U_i intersects U_j for only finitely many j . Then define $I := \prod_i I_{f_i}$ as follows:

- $\mathcal{O}_I := \bigcap_i \mathcal{O}_{I_{f_i}}$. In particular, O is in \mathcal{O}_I if it is connected and is contained in U_i whenever it intersects U'_i . As each curve has a neighborhood that intersects only finitely many U'_i , each family in \mathcal{K}_ϵ still has an open cover contained in \mathcal{O}_I .
- For $O \in \mathcal{O}_I$,

$$I(O) := \prod_i I_{f_i}(O).$$

As O is in only finitely many U_i , I_{f_i} is a probability space with a unique element for all but finitely many i , so $\prod_i I_{f_i}(O)$ is still a finite probability space.

Sections in $I(O)$ are separated if and only if their image is separated in some $I_{f_i}(O)$. Because I_{f_i} has trivial stabilizers on U'_i , I has trivial stabilizers on U'_i for each i , so I has trivial stabilizers on all of \mathcal{K}_ϵ .

It remains to prove that any two choices of $\mathcal{K}_{i,\epsilon}$ and I_i are cobordant. We may reparametrize the original cobordism between \mathcal{K}_i to obtain a Kuranishi category \mathcal{K} over $\mathbf{Z} \times \mathbb{R}$ with the extra property that, for $i \in \{0, 1\}$, there exist neighborhoods N_i of

$i \in \mathbb{R}$ such that the pullback of \mathcal{K} to $\mathbf{Z} \times N_i$ is the pullback of \mathcal{K}_i under the projection $\mathbf{Z} \times N_i \rightarrow \mathbf{Z}$. Using $\mathbf{Z} \times N_i \rightarrow \mathcal{K}_i$, pull back the functions used to define $\mathcal{K}_{i,\epsilon}$. Then, extend these functions by taking the minimum of each function and a smooth function on \mathbb{R} equal to 1 at i and -1 outside of N_i . Following this, use Lemma 2.6 to construct the other functions required to define \mathcal{K}_ϵ , as discussed following Definition 3.9, and choose these other functions to be -1 on the inverse image of $\{0, 1\}$. This construction of \mathcal{K}_ϵ pulls back to the construction of $\mathcal{K}_{i,\epsilon}$, as required.

Now define I . Use the pullbacks of I_i over N_i to define weighted branched covers I'_i as follows: Choose open neighborhoods $N'_i \subset N_i$ of 0 or 1, respectively, so that $\bar{N}'_i \subset N_i$. Let O in \mathcal{K}^{st} be in $\mathcal{O}_{I'_i}$ in the following two cases:

- If O is connected and does not intersect the pullback of \bar{N}'_i , define $I'_i(O)$ to be the probability space with a unique element.
- If O is contained in the inverse image of N_i , and projects to O' within $\mathcal{K}_i^{\text{st}}$ such that $O' \in \mathcal{O}_{I_i}$,
 - define $I'_i(O)$ to be the probability space with a unique element if O does not intersect the pullback of \bar{N}'_i ;
 - define $I'_i(O)$ to be $I_i(O')$ with the same measure, but the trivial indiscrete equivalence relation if O intersects the pullback of the boundary of N'_i ;
 - define $I'_i(O)$ to be $I_i(O')$ with the same measure and equivalence relation if O is contained in the pullback of N'_i .

Now $I'_0 \times I'_1$ is a weighted branched cover of \mathcal{K}_ϵ , and the pullback of $I'_0 \times I'_1$ under the inclusion of \mathbf{Z} over i is I_i .

$I'_0 \times I'_1$ has trivial stabilizers restricted to the pullback of N'_i , but may not have trivial stabilizers elsewhere. We can make our weighted branched cover I have trivial stabilizers by multiplying $I'_0 \times I'_1$ by other weighted branched covers I_f as above. We can achieve this while choosing I_f to be trivial when restricted the inverse image of $i \in \mathbb{R}$, so that I pulls back to give I_i , as required. \square

Remark 4.6 Similarly, there exists a weighted branched cover of any extendable K -category, and as above, any two such weighted branched covers are cobordant.

Definition 4.7 Given a Kuranishi category \mathcal{K} , proper and oriented over \mathbf{Z} , construct the virtual class $[\mathcal{K}]$ as follows:

- (i) Choose \mathcal{K}_ϵ as in Definition 3.9.
- (ii) Choose a separating weighted branched cover I of \mathcal{K}_ϵ as in Lemma 4.5.

- (iii) Consider the sheaf S^I of weighted branched sections of S from Definition 3.12. We have proved that S^I satisfies the patching, extension and averaging axioms. Choose a global section of S^I over \mathcal{K}_ϵ , as allowed by Proposition 2.3. (This may involve increasing ϵ slightly.)
- (iv) Recall that intersection with 0 defines a natural transformation $S \rightarrow E$, where E is the sheaf of oriented subfamilies from Definition 3.13. Intersection with 0 therefore defines a natural transformation $S^I \rightarrow E^I$, so the intersection with 0 of our weighted branched section of S defines a weighted branched section of E . Use the notation $[\mathcal{K}]$ for such a section.

Lemma 4.5 along with Proposition 2.3 and the fact that S^I obeys the patching, extension and averaging axioms imply that any two weighted branched sections of E defined using the above procedure are cobordant.

5 Representing Gromov–Witten invariants using de Rham cohomology

5.1 Differential forms and de Rham cohomologies on stacks and Kuranishi categories

Differential forms on a stack (over the category of smooth manifolds or exploded manifolds) form a sheaf—to each family \hat{f} , we associate the differential forms on $F(\hat{f})$. Any notion for differential forms commuting with pullbacks (such as exterior differentiation, wedge products, sums) also makes sense for such differential forms on stacks. In particular, it is possible to do de Rham cohomology with differential forms on a stack. In the case of a stack (over the category of manifolds) with infinite stabilizers, the resulting cohomology will be smaller than the “correct” cohomology, explained in [2]; however, I do not know how to imitate the constructions of [2] in the infinite-dimensional setting of $\mathcal{M}_\bullet^{\text{st}}$. For exploded manifolds, we use restricted types of differential forms for defining cohomology theories; see [31]. We need several species of differential forms, all of which are just smooth differential forms on smooth manifolds.

Definition 5.1 ($\Omega^*(B)$) Let $\Omega^k(B)$ be the space of $C^{\infty,1}$ differential k -forms θ on an exploded manifold B such that for all integral vectors v , the differential form θ vanishes on v , and for all maps $f: T_{(0,\infty)}^1 \rightarrow B$, the differential form θ vanishes on all vectors in the image of df .

Similarly, a differential form on a stack \mathcal{X} is in $\Omega^*(\mathcal{X})$ if it is in $\Omega^*(F(\hat{f}))$ for all \hat{f} in \mathcal{X} .

Denote by $\Omega_c^k(\mathbf{B}) \subset \Omega^k(\mathbf{B})$ the subspace of forms with complete support.¹¹ Denote the homology of $(\Omega^*(\mathbf{B}), d)$ or $(\Omega^*(\mathcal{X}), d)$, by $H^*(\mathbf{B})$ or $H^*(\mathcal{X})$, respectively.

Remark 5.2 Given a complex vectorbundle W over an extension \mathcal{K}^\sharp of \mathcal{K} , we may represent the Chern classes of W in $H^*(\mathcal{K})$ as closed differential forms. As unitary metrics and connections may be constructed as sections of sheaves satisfying the patching, extension and averaging axioms, they may be constructed on \mathcal{K} using Proposition 2.3, then we may construct the Chern classes of V over \mathcal{K} using the Chern–Weil construction.

We use Ω^* instead of all $C^{\infty,1}$ differential forms in order to use a version of Stokes' theorem, Theorem 3.4 in [31]. We shall also wish to use integration along the fiber. For this we shall need the following more general types of differential forms. (For integration along the fiber, see Theorem 9.2 of [31].)

Definition 5.3 (refined forms) A refined form $\theta \in {}^r\Omega^*(\mathbf{B})$ is a choice $\theta_p \in \bigwedge T_p^*(\mathbf{B})$ for all $p \in \mathbf{B}$ satisfying the following condition: given any point $p \in \mathbf{B}$, there exists an open neighborhood U of p , a complete, surjective, equidimensional submersion

$$r: U' \rightarrow U$$

and a form $\theta' \in \Omega^*(U')$ which is the pullback of θ . In other words, if v is any vector on U' such that $\text{Tr}(v)$ is a vector based at p , then

$$\theta'(v) = \theta_p(\text{Tr}(v)).$$

As refined forms pull back to refined forms, there is an analogous notion of refined forms on any Kuranishi category or stack over the category of $C^{\infty,1}$ exploded manifolds.

A refined form $\theta \in {}^r\Omega^*(\mathbf{B})$ is completely supported if there exists some complete subset V of an exploded manifold \mathbf{C} with a map $\mathbf{C} \rightarrow \mathbf{B}$ such that $\theta_p = 0$ for all p outside the image of V . Use the notation ${}^r\Omega_c^*$ for completely supported refined forms. (There is no analogous notion on $\mathcal{M}_\bullet^{\text{st}}$.)

¹¹A form has complete support if the set where it is nonzero is contained inside a complete subset of \mathbf{B} — in other words, a compact subset with tropical part consisting only of complete polytopes.

Because completely supported forms do not always pull back to be completely supported, there is no analogous notion of forms with complete support on $\mathcal{M}_\bullet^{\text{st}}$. See [2] for a version of compactly supported forms on finite-dimensional differential stacks.

Denote the homology of $({}^r\Omega^*(\mathbf{B}), d)$ by ${}^rH^*(\mathbf{B})$ and $({}^r\Omega_c^*(\mathbf{B}), d)$ by ${}^rH_c^*(\mathbf{B})$.

The Poincaré dual to a map $\mathbf{C} \rightarrow \mathbf{B}$ is correctly viewed as a refined differential form in ${}^rH^*(\mathbf{B})$; see Lemma 9.5 of [31]. As with all types of forms considered in this paper, refined forms admit the usual operations of wedge product, pullbacks, exterior derivatives and contraction with any $C^{\infty,1}$ vectorfield—for a discussion of wedge products, see Section 9 of [31]; contraction with any $C^{\infty,1}$ vectorfield is defined because any equidimensional submersion lifts any such vectorfield uniquely to a $C^{\infty,1}$ vectorfield.

In the coming sections, we will define integration and pushforwards of differential forms using $[\mathcal{K}]$. If \mathcal{K} is complete (Definition 3.5) and contained in the stack \mathcal{X} , then integration over $[\mathcal{K}]$ defines a map

$$H^*(\mathcal{X}) \xrightarrow{\int_{[\mathcal{K}]}} \mathbb{R}$$

or more generally a map

$${}^rH^*(\mathcal{X}) \xrightarrow{\int_{[\mathcal{K}]}} \mathbb{R}.$$

The first map factors through $H^* \rightarrow {}^rH^*$, so the second map contains more information than the first. Given a complex vectorbundle V over \mathcal{X} , we may define more maps ${}^rH^*(\mathcal{X}) \rightarrow \mathbb{R}$ by taking the product with Chern classes of V before integrating.

If \mathcal{K} is complete over \mathbf{Z} and contained in a stack \mathcal{X} , integration along the fiber of the map $[\mathcal{K}] \rightarrow \mathbf{Z}$ defines a map

$${}^rH^*(\mathcal{X}) \rightarrow {}^rH^*(\mathbf{Z}).$$

These maps are compatible with base changes

$$\begin{array}{ccc} \mathcal{X}' & \longrightarrow & \mathcal{X} \\ \downarrow & & \downarrow \\ \mathbf{Z}' & \longrightarrow & \mathbf{Z} \end{array}$$

in the sense that the following diagram commutes:

$$(2) \quad \begin{array}{ccc} {}^rH^*(\mathcal{X}') & \longleftarrow & {}^rH^*(\mathcal{X}) \\ \downarrow & & \downarrow \\ {}^rH^*(\mathbf{Z}') & \longleftarrow & {}^rH^*(\mathbf{Z}) \end{array}$$

In the case that \mathcal{K} is not complete, we need to use the following, more restrictive, types of differential forms to define invariants:

Definition 5.4 (differential forms generated by functions) A differential form is generated by functions if it is locally equal to a form constructed from $C^{\infty,1}$ functions using the operations of exterior differentiation and wedge products. Use the notation ${}^r_{\text{fg}}\Omega^* \subset {}^r\Omega^*$ for the set of refined forms with differential forms generated by functions playing the role of Ω^* in Definition 5.3.

Use ${}^r_{\text{fg}}H^*$ to denote the homology of $({}^r_{\text{fg}}\Omega^*, d)$.

Examples of differential forms generated by functions are the Poincaré dual to a point, the Chern class defined using the Chern–Weil construction, and any smooth differential form on a smooth manifold.

Remark 5.5 Differential forms generated by functions could equivalently be defined as $C^{\infty,1}$ differential forms that vanish on all \mathbb{R} –nil vectors.¹² This follows from the proof of Lemma 4.1 in [31].

In the case that $\mathcal{K} \subset \mathcal{X}$ is compact but not complete, integration over $[\mathcal{K}]$ defines a map

$${}^r_{\text{fg}}H^*(\mathcal{X}) \xrightarrow{\int_{[\mathcal{K}]}} \mathbb{R}.$$

If \mathcal{K} is complete, then the above map factorizes through the map ${}^r_{\text{fg}}H^* \rightarrow {}^rH^*$. In the case that \mathcal{K} is not complete, these invariants only behave well in families parametrized by smooth (not exploded) manifolds. In particular, if $\hat{\mathbf{B}} \rightarrow \mathbf{B}_0$ is a family of targets parametrized by a smooth manifold \mathbf{B}_0 , then Gromov–Witten invariants give a map ${}^r_{\text{fg}}H^*(\mathcal{M}^{\text{st}}_{\bullet}(\hat{\mathbf{B}})) \rightarrow H^*(\mathbf{B}_0)$. These Gromov–Witten invariants are invariant under base changes in the sense that a diagram analogous to diagram (2) commutes.

5.2 Integrating over $[\mathcal{K}]$ for compact \mathcal{K}

Recall that $[\mathcal{K}]$ is a section of E^I , where E is as defined in Definition 3.13 and I is a weighted branched cover of \mathcal{K}_ϵ (Definition 3.19). So $[\mathcal{K}]$ is a natural transformation $I \rightarrow E$. In particular, given any family O in $\mathcal{K}_\epsilon \cap \mathcal{O}_I$, we have a finite probability space $(I(O), \mu)$ and a map from $I(O)$ to the set of complete subfamilies of O contained in \mathcal{K}_C . Each such family $[\mathcal{K}](i)$ has a canonical orientation relative to \mathbf{Z} , so in the case that \mathbf{Z} is a point or oriented, $[\mathcal{K}](i)$ is oriented.

We need two different notions of a partition of unity. By a partition of unity subordinate to a given covering of \mathcal{K}^{st} by open substacks we mean a collection of $C^{\infty,1}$ functions ρ_i

¹²An \mathbb{R} –nil vector is a vector v for which $df(v) = 0$ for all differentiable \mathbb{R} –valued functions.

on \mathcal{K}^{st} with support contained in these open substacks such that $\sum \rho_i = 1$. Remark 2.7 implies that we may construct such partitions of unity as usual (under the assumption that \mathcal{K} is extendable). For integrating differential forms over $[\mathcal{K}]$, we need a different notion of partition of unity on \mathcal{K}_C . When \mathcal{K} is compact in the sense of Definition 3.5, \mathcal{K}_C from Definition 3.9 is compact and has a finite open cover consisting of families in $\mathcal{K}_\epsilon \cap \mathcal{O}_I$. The partition of unity on \mathcal{K}_C defined below has functions living on such an open cover. This definition is similar to the notion of a partition of unity on an étale proper groupoid given in Definition 22 of [2].

Definition 5.6 A partition of unity on \mathcal{K}_C is a family O in \mathcal{K} and a $C^{\infty,1}$ function

$$r: O \rightarrow \mathbb{R}$$

such that the support of r on any connected component O_k of O is compact and such that, for any family F in \mathcal{K}_C , the pullback of r to $O \times_{\mathcal{K}^{\text{st}}} F$ has proper support over F , and has pushforward to F equal to 1:

$$\begin{array}{ccc} O \times_{\mathcal{K}^{\text{st}}} F & \xrightarrow{\pi_2} & F \\ \downarrow \pi_1 & & \downarrow \\ O & \longrightarrow & \mathcal{K}^{\text{st}} \end{array} \quad (\pi_2)_! \pi_1^* r = 1.$$

Say that this partition of unity is compatible with a weighted branched cover I of \mathcal{K}_ϵ if each connected component O_k of O is in $\mathcal{O}_I \cap \mathcal{K}_\epsilon$. If $[\mathcal{K}]$ is a natural transformation $[\mathcal{K}]: I \rightarrow E$, say that a partition of unity is compatible with $[\mathcal{K}]$ if it is compatible with I .

Lemma 5.7 *Given any weighted branched cover I of \mathcal{K}_ϵ , there exists a partition of unity on \mathcal{K}_C compatible with I . Moreover, given any curve $f \in \mathcal{K}_C$, there exists an open neighborhood N of f within \mathcal{K}_C and a partition of unity $r: \coprod_k O_k \rightarrow \mathbb{R}$ satisfying the following:*

- *There is a group G_1 of automorphisms of O_1 such that O_1/G_1 represents a substack of \mathcal{K}^{st} .*
- *On the intersection of O_1 with N , $r = 1/|G_1|$.*
- *On the intersection of O_k with N for $k \neq 1$, $r = 0$.*

Proof Choose a locally finite cover of \mathcal{K}_C by families $O_k \in \mathcal{K}_\epsilon \cap \mathcal{O}_I$ with automorphism groups G_k so that O_k/G_k represents a substack of \mathcal{K}^{st} . We can do this so that O_1/G_1 compactly contains a neighborhood N of the given curve $f \in \mathcal{K}_C$.

Using Lemma 2.6, we may choose $C^{\infty,1}$ functions $r_k: \mathcal{K}_\epsilon \rightarrow [0, 1]$ satisfying the following conditions:

- The support of r_k intersected with \mathcal{K}_C is compactly contained within O_k/G_k .
- The support of r_k intersects the support of r_j for only finitely many j .
- $\sum_k r_k > 0$ on \mathcal{K}_C .
- The support of r_1 contains N , and all other r_j vanish on N .

As \mathcal{K}_C is a closed substack, Remark 2.7 implies that there exists a $C^{\infty,1}$ function $R: \mathcal{K}_\epsilon \rightarrow [0, 1]$ with zero set \mathcal{K}_C .

Then set

$$r = \frac{r_k}{|G_k|(R + \sum_k r_k)} \quad \text{on } O_k.$$

Given a curve f in \mathcal{K}_C , the fiber product of f with O_k consists of $|G_k|$ points if f is also contained in O_k , and is otherwise empty, so

$$\int_{f \times_{\mathcal{K}^{\text{st}}} O_k} r = \frac{r_k(f)}{\sum_k r_k(f)}$$

and

$$\int_{f \times_{\mathcal{K}^{\text{st}}} \mathcal{O}} r = \frac{\sum_k r_k(f)}{\sum_k r_k(f)} = 1,$$

as required. \square

Definition 5.8 Given $\theta \in {}^r\Omega^*(\mathcal{K})$, for \mathcal{K} oriented and compact (Definition 3.5), define the integral of θ over $[\mathcal{K}]$ as follows: Choose a partition of unity $r: \coprod_k O_k \rightarrow \mathbb{R}$ compatible with $[\mathcal{K}]$. Then define

$$\int_{[\mathcal{K}]} \theta := \sum_k \sum_{i \in I(O_k)} \mu(i) \int_{[\mathcal{K}](i)} r\theta,$$

where $[\mathcal{K}](i)$ is the closed, oriented, exploded submanifold of O_k that is the image of i under the map $[\mathcal{K}]: I(O_k) \rightarrow E(O_k)$, and μ is the probability measure on the indexing set $I(O_k)$. Note that $r\theta$ has compact support on $[\mathcal{K}](i)$, so the integral of $r\theta$ is defined as in [31]. Note also that the above sum is finite.

As ${}^r\Omega^*$ contains Ω^* and ${}^r_{\text{fg}}\Omega^*$, the above definition also works for θ in Ω^* or ${}^r_{\text{fg}}\Omega$.

Lemma 5.9 $\int_{[\mathcal{K}]} \theta$ does not depend on the choice of partition of unity.

Proof Let $r: \coprod_k O_k \rightarrow \mathbb{R}$ and $r': O' \rightarrow \mathbb{R}$ be two partitions of unity compatible with $[\mathcal{K}]$. We have

$$\int_{[\mathcal{K}]} \theta := \sum_k \sum_{i \in I(O_k)} \mu(i) \int_{[\mathcal{K}](i)} r \theta$$

and, because r' is a partition of unity,

$$\int_{O' \times_{\mathcal{K}^{\text{st}}} [\mathcal{K}](i)} r' r \theta = \int_{[\mathcal{K}](i)} r \theta.$$

We need to examine $O' \times_{\mathcal{K}^{\text{st}}} [\mathcal{K}](i)$. Let O be a connected component of $O' \times_{\mathcal{K}^{\text{st}}} O_k$, so $O \in \mathcal{O}_I \cap \mathcal{K}_\epsilon$, and comes with a morphism $\iota: O \rightarrow O_k$. For any $j \in \iota^* i$, $[\mathcal{K}](j) = \iota^{-1}([\mathcal{K}](i))$, and therefore $[\mathcal{K}](j)$ is some collection of connected components of $O' \times_{\mathcal{K}^{\text{st}}} [\mathcal{K}](i)$. If $\iota^* i$ always had a unique element j , the union of all such $[\mathcal{K}](j)$ would be $O' \times_{\mathcal{K}^{\text{st}}} [\mathcal{K}](i)$. In the general case, $\iota^* i$ is some finite set with measure $\mu(i)$, so we get

$$\mu(i) \int_{[\mathcal{K}](i)} r \theta = \sum_{j \in \iota^* i} \mu(j) \int_{[\mathcal{K}](j)} r' r \theta,$$

where the sum is over all connected components O of $O' \times_{[\mathcal{K}^{\text{st}}]} O_k$ and j in the inverse image of i within $I(O)$. Taking the sum of this expression over all O_k and $i \in I(O_k)$ gives

$$\int_{[\mathcal{K}]} \theta = \sum_O \sum_{j \in I(O)} \mu(j) \int_{[\mathcal{K}](j)} r' r \theta,$$

where the above sum now is over all connected components O of $O' \times_{\mathcal{K}^{\text{st}}} \coprod_k O_k$. The above expression is symmetric in the two partitions of unity; therefore, the integral is independent of the choice of partition of unity. \square

Lemma 5.10 *If $\theta \in {}^r \Omega^*(\mathcal{K})$ and \mathcal{K} is complete, then*

$$\int_{[\mathcal{K}]} d\theta = 0.$$

Proof The fact that \mathcal{K} is complete implies that the support of each r on each connected component of O is complete. Because exterior differentiation and integration are linear we may use a partition of unity on \mathcal{K}^{st} to reduce to the case that θ has small support. In particular, we can assume that the support of θ is small enough that Lemma 5.7 gives a partition of unity $r: \coprod_k O_k \rightarrow \mathbb{R}$ such that $r\theta$ has support contained in O_1 on

a subset where r is equal to $1/|G|$. Then

$$\int_{[\mathcal{K}]} d\theta = \sum_{i \in I(O_1)} \frac{\mu(i)}{|G|} \int_{[\mathcal{K}](i)} d\theta = 0$$

because $\theta \in {}^r\Omega_c^*([\mathcal{K}](i))$, and such forms satisfy Stokes' theorem; see [31]. \square

Lemma 5.11 *If $\theta \in {}^r_{\text{fg}}\Omega^*(\mathcal{K})$ and \mathcal{K} is compact but not necessarily complete, then*

$$\int_{[\mathcal{K}]} d\theta = 0.$$

Proof As in the proof of Lemma 5.10, this lemma reduces to the case of proving that $\int_{[\mathcal{K}](i)} d\theta = 0$ for any compactly supported $\theta \in {}^r_{\text{fg}}\Omega^*[\mathcal{K}](i)$. By using a partition of unity on $[\mathcal{K}](i)$, we may reduce to the case that θ is supported on an open subset U of $[\mathcal{K}](i)$, and pulls back under a refinement map $U' \rightarrow U$ to a compactly supported form generated by functions on U' . We may use a partition of unity on U' to reduce to the case that θ is compactly supported within a single coordinate chart V on U' . The fact that θ is generated by functions implies that θ is pulled back from a differential form on \mathbb{R}^n under an embedding $[V] \rightarrow \mathbb{R}^n$. Then, the usual Stokes' theorem implies that $\int_V d\theta = 0$. \square

Corollary 5.12 *If \mathcal{K} is complete, then integration over $[\mathcal{K}]$ defines a map*

$${}^rH^*(\mathcal{K}) \rightarrow \mathbb{R}.$$

If \mathcal{K} is compact, integration over $[\mathcal{K}]$ defines a map

$${}^r_{\text{fg}}H^*(\mathcal{K}) \rightarrow \mathbb{R}.$$

If \mathcal{K} is contained in a stack \mathcal{X} , we shall show that the resulting maps ${}^rH^*(\mathcal{X}) \rightarrow \mathbb{R}$ and ${}^r_{\text{fg}}H^*(\mathcal{X}) \rightarrow \mathbb{R}$ only depend on the cobordism class of \mathcal{K} within \mathcal{X} , and in particular are independent of the choices involved in the construction of $[\mathcal{K}]$.

5.3 Pushing forward cohomology classes

To define pushforwards of differential forms along maps $[\mathcal{K}] \rightarrow X$, we need integration along the fiber, constructed in Theorem 9.2 of [31]. Given any oriented submersion of exploded manifolds, $\psi: X \rightarrow Y$, there exists a linear chain map

$$\psi_!: {}^r\Omega_c^*(X) \rightarrow {}^r\Omega_c^*(Y)$$

uniquely determined by the usual property of integration along the fiber, namely,

$$\int_X \psi^* \alpha \wedge \beta = \int_Y \alpha \wedge \psi_! \beta.$$

Above, the notation $*'$ emphasizes that $\psi_!$ does not preserve degree — it shifts it by $\dim Y - \dim X$. Later, in Lemma 7.3, we show that $\psi_!$ sends ${}^r_{\text{fg}}\Omega^*(X)$ into ${}^r_{\text{fg}}\Omega^*(Y)$. We only need that X is oriented relative to Y ; when Y is not oriented, the above expression should either be interpreted locally in Y with a choice of orientation, or α needs to be a form twisted by the orientation line bundle of Y . For integration along the fiber to work, we require our forms to have complete support (or at least complete support relative to the target Y).

Remark 5.13 This pushforward is also defined in the case that X and Y are exploded orbifolds (ie Deligne–Mumford stacks in the category of exploded manifolds). In particular, we can define $\psi_!$ so that given any pullback diagram

$$\begin{array}{ccc} X' & \xrightarrow{\psi'} & Y' \\ \downarrow & & \downarrow \\ X & \xrightarrow{\psi} & Y \end{array}$$

where X' and Y' are exploded manifolds, the following diagram commutes:

$$\begin{array}{ccc} {}^r\Omega^*(X') & \xrightarrow{\psi'_!} & {}^r\Omega^*(Y') \\ \uparrow & & \uparrow \\ {}^r\Omega^*(X) & \xrightarrow{\psi_!} & {}^r\Omega^*(Y) \end{array}$$

If all that can be guaranteed is compact (but not complete) support, integration along the fiber will not give a well-behaved form on the target Y . For example, consider the proper but incomplete map given by the inclusion of $T_{(0,1)}^1$ into T . Integrating the constant function 1 along the fiber of this map gives a discontinuous function on T . To get around this problem, we may use differential forms generated by functions, and restrict to the case that the target of our map is a manifold.

Lemma 5.14 *Given any oriented submersion from an exploded manifold to a smooth manifold,*

$$\psi: X \rightarrow M,$$

and a compactly supported form $\beta \in {}^r_{\text{fg}}\Omega^*X$, integration along the fiber of ψ gives a form $\psi_!\beta \in \Omega_c^*M$ uniquely determined by the property that, for any differential form α on M (twisted by the orientation line bundle if necessary),

$$\int_X \psi^* \alpha \wedge \beta = \int_M \alpha \wedge \psi_! \beta.$$

As usual, integration along the fiber is a chain map, so $d\psi_!\alpha = \psi_!d\alpha$.

Proof Forms on smooth manifolds are uniquely determined by their integrals against all other forms, so if there exists a form $\psi_!\beta$ satisfying the required property, it is unique. As the required property is linear, we can use a partition of unity to reduce to the case that β is compactly contained in a single coordinate chart. As β is a refined form, we may need to refine this coordinate chart and use a further partition of unity to reduce to the case that β is an (unrefined) form compactly supported on some standard coordinate chart isomorphic to $\mathbb{R}^k \times \mathbb{R}^n \times T_P^m$, where the submersion to M is modeled on the projection to \mathbb{R}^k . Integrating along the fiber, and any integral on this coordinate chart, involves a contribution from each 0-dimensional stratum of P , and no other strata contribute.

For each 0-dimensional stratum i of P , let P_i be the union of all strata in P with closure intersecting $i \in P$. We may construct a tropical completion¹³ \check{P}_i of P_i as the union of all rays that start at $i \in P$ and intersect P in more than one point. There is a natural inclusion of $T_{P_i}^m$ into $T_{\check{P}_i}^m$, and any compactly supported form β on $\mathbb{R}^{k+n} \times T_{P_i}^m$ extends uniquely to a completely supported form β_i in $\mathbb{R}^{k+n} \times T_{\check{P}_i}^m$. As the submersion to M is constant on the T_P^m fibers, it also extends uniquely to a submersion $\psi^i: \mathbb{R}^{n+k} \times T_{\check{P}_i}^m \rightarrow M$. As integration along the fiber is defined and satisfies the required properties on completely supported forms, we may define

$$\psi_! \beta := \sum_i \psi^i_! \beta_i.$$

Then

$$\int_M \alpha \wedge \psi_! \beta = \sum_i \int (\psi^i)^* \alpha \wedge \beta_i$$

but the integral of $\psi^* \alpha \wedge \beta$ has contributions from each 0-dimensional stratum i of P , each equal to the integral of $(\psi^i)^* \alpha \wedge \beta_i$, so

$$\int_M \alpha \wedge \psi_! \beta = \int \psi^* \alpha \wedge \beta,$$

¹³ See Section 7 for a more thorough discussion of tropical completion.

as required. The fact that $\psi_!$ commutes with exterior differentiation follows from the analogous fact for $\psi_!^i$. \square

Let $A \rightarrow Z$ be some family of exploded manifolds or orbifolds over Z , oriented relative to Z . For example, in the case of a Kuranishi category describing families of curves in $\hat{B} \rightarrow B_0$, $Z = B_0$, and A might be B_0 , or the n -fold fiber product of \hat{B} over B_0 , or the product of B_0 with a component of (the explosion of) Deligne–Mumford space. In this section, we shall consider pushing forward cohomology classes from ${}^rH^*(\mathcal{K})$ along a map,¹⁴ $\pi: \mathcal{K} \rightarrow A$, compatible with a complete submersion $\mathcal{K} \rightarrow Z$:

$$\begin{array}{ccc} \mathcal{K} & \xrightarrow{\pi} & A \\ & \searrow & \downarrow \\ & & Z \end{array}$$

Definition 5.15 Given a map π compatible with a submersion

$$\begin{array}{ccc} \mathcal{K} & \xrightarrow{\pi} & A \\ & \searrow & \downarrow \\ & & Z \end{array}$$

in the case that \mathcal{K} is complete and oriented over Z , define a pushforward map

$$\pi_!: {}^r\Omega^*(\mathcal{K}) \rightarrow {}^r\Omega^{*'}(A)$$

and in the case where \mathcal{K} is proper over Z but A is a manifold, define

$$\pi_!: {}_{\text{fg}}^r\Omega^*(\mathcal{K}) \rightarrow \Omega^{*'}(A).$$

In either case, $\pi_!$ is defined as follows:

- (i) Choose an oriented vectorbundle $W \rightarrow A$ along with a map

$$x: W \rightarrow A$$

so that x restricted to the zero section is the identity, and satisfies an extra transversality condition described below.

- (ii) Choose a partition of unity $r: \coprod_k O_k \rightarrow \mathbb{R}$ compatible with $[\mathcal{K}]$.

¹⁴By a map from a K -category or Kuranishi category to A , we mean a map $\mathcal{K}^{\text{st}} \rightarrow A$ — this entails a compatible choice of $C^{\infty,1}$ map $F(\hat{f}) \rightarrow A$ for every family \hat{f} in \mathcal{K}^{st} .

- (iii) For any $i \in I(O_k)$, the natural transformation $[\mathcal{K}]: I \rightarrow E$ gives a complete subfamily $[\mathcal{K}](i) \subset O_k$. Use the notation $\pi(i)$ for π restricted to $[\mathcal{K}](i)$. Define a map $\hat{\pi}(i): \pi(i)^*W \rightarrow \mathcal{A}$ by the composition

$$\begin{array}{ccccc} & & \hat{\pi}(i) & & \\ & \searrow & \text{---} & \nearrow & \\ \pi(i)^*W & \longrightarrow & W & \xrightarrow{x} & \mathcal{A} \\ \downarrow & & \downarrow & & \\ [\mathcal{K}](i) & \xrightarrow{\pi(i)} & \mathcal{A} & & \end{array}$$

The extra transversality assumption we require of x is that these maps $\hat{\pi}(i)$ are submersions. This condition is always satisfied in the case that x is a submersion restricted to each fiber.

- (iv) Given any θ in ${}^r\Omega^*(\pi^*W)$ or ${}^r_{\text{fg}}\Omega^*(\pi^*W)$, respectively, with compact support on fibers of $\pi^*W \rightarrow \mathcal{K}$, define

$$\hat{\pi}_!(\theta) := \sum_{O_k} \sum_{i \in I(O_k)} \mu(i) \hat{\pi}(i)_!(r\theta).$$

Note that integration along the fiber of $\hat{\pi}(i)$ requires an orientation relative to \mathcal{A} . Both \mathcal{A} and the family $[\mathcal{K}](i)$ are canonically oriented relative to \mathcal{Z} , and W is an oriented vectorbundle, so $\hat{\pi}(i)$ has a canonical relative orientation.

- (v) Choose a closed form $e \in \Omega^*W$ with fiberwise compact support so that e represents the Thom class of $W \rightarrow \mathcal{A}$.
- (vi) For any θ in ${}^r\Omega^*(\mathcal{K})$ or ${}^r_{\text{fg}}\Omega^*(\mathcal{K})$, respectively, consider θ and e as forms on π^*W , then define

$$\pi_!(\theta) := \hat{\pi}_!(\theta \wedge e).$$

Because the Thom form e is generated by functions, Lemma 7.3 implies that $\pi_!$ sends forms in ${}^r_{\text{fg}}\Omega^*(\mathcal{K})$ to ${}^r_{\text{fg}}\Omega^*(\mathcal{A})$.

Lemma 5.16 *The map $\hat{\pi}_!$ from Definition 5.15 is independent of choice of partition of unity.*

Proof The proof is identical to the proof of Lemma 5.9, except integration over $[\mathcal{K}](i)$ is replaced by pushing forward via $\hat{\pi}(i)_!$. \square

Lemma 5.17 *The map $\hat{\pi}_!$ from Definition 5.15 commutes with exterior differentiation.*

Proof As d and $\hat{\pi}_!$ are linear, we may assume that the support of θ is small enough that Lemma 5.7 gives a compatible partition of unity $r: \coprod_k O_k \rightarrow \mathbb{R}$ such that $r\theta$ has support where $r = 1/|G|$ in O_1 . Then

$$\hat{\pi}_!(d\theta) = \frac{1}{|G|} \sum_{i \in I(O_1)} \mu(i) \hat{\pi}(i)_! d\theta = d \left(\frac{1}{|G|} \sum_{i \in I(O_1)} \mu(i) \hat{\pi}(i)_! \theta \right) = d \hat{\pi}_!(\theta). \quad \square$$

As an immediate corollary, $\pi_!$ induces maps

$$\pi_!: {}^r H^*(\mathcal{K}) \rightarrow {}^r H^*(A) \quad \text{and} \quad \pi_!: {}^r_{\text{fg}} H^*(\mathcal{K}) \rightarrow {}^r_{\text{fg}} H^*(A)$$

in the case that \mathcal{K} is complete over Z , and a map

$$\pi_!: {}^r_{\text{fg}} H^*(\mathcal{K}) \rightarrow H^*(A)$$

in the case that A is a smooth manifold and \mathcal{K} is proper over Z . Of course, if \mathcal{K} is contained in a stack \mathcal{X} , we are usually interested in the precomposition of the above maps with $H^*(\mathcal{X}) \rightarrow H^*(\mathcal{K})$.

Lemma 5.18 *On the level of cohomology, the map $\pi_!$ does not depend on the choice of W , x or e in Definition 5.15.*

Proof For a given W , any two different choices of e differ by $d\alpha$, where α is some fiberwise compactly supported differential form on W . The difference between $\pi_!(\theta)$ defined using these two choices of e is therefore $\hat{\pi}_!(\theta \wedge d\pi^*\alpha)$, which vanishes in cohomology whenever θ is closed because of Lemma 5.17.

Suppose that we have two homotopic choices of x for a given W . We then have a homotopy $\hat{\pi}_t$ between $\hat{\pi}$ defined using each choice of x . If θ is closed, $(\hat{\pi}_t)_!(\theta \wedge \pi^*e)$ then gives a closed form on $A \times [0, 1]$ restricting to $A \times \{0, 1\}$ to give $\pi_!(\theta)$ defined using the two different choices of x . It follows that the cohomology class of $\pi_!(\theta)$ is independent of choice of homotopic x .

Given two choices (W_1, x_1, e_1) and (W_2, x_2, e_2) , consider $W := W_1 \oplus W_2$ with x defined by the projection to W_1 followed by x_1 , and e given by the wedge product of the pullback of e_1 and e_2 to $W_1 \oplus W_2$. We may then factorize $\hat{\pi}(i)$ as

$$\begin{array}{c} \hat{\pi}(i) \\ \curvearrowright \\ \pi(i)^* W_1 \oplus W_2 \xrightarrow{p(i)} \pi(i)^* W_1 \xrightarrow[\hat{\pi}_1(i)]{\quad} A \end{array}$$

so $\hat{\pi}(i)_!$ factorizes as $\hat{\pi}_1(i)_! \circ p(i)_!$. The fact that e_2 is a Thom class implies that

$$p(i)_!(p(i)^*\theta \wedge e_2) = \theta,$$

so if we abuse notation as in Definition 5.15 and regard the form θ on \mathcal{K} as also living on π^*W_1 and $\pi^*W_1 \oplus W_2$, then

$$\hat{\pi}(i)_!(\theta \wedge e_1 \wedge e_2) = \hat{\pi}_1(i)_!(\theta \wedge e_1).$$

It follows that $\pi_!(\theta)$ defined using $\hat{\pi}$ and $\hat{\pi}_1$ is identical. The projection to W_1 followed by x_1 is homotopic to the projection to W_2 followed by x_2 ; therefore, the same argument identifies the cohomology class of $\pi_!(\theta)$ defined using $W_1 \oplus W_2$ with that defined using W_2 . We have now shown that on the level of cohomology, $\pi_!$ does not depend on W , x or e . \square

Lemma 5.19 *If \mathcal{K} is complete, the following equation holds when θ is any closed differential form in ${}^r\Omega^*(A)$:*

$$\int_{[\mathcal{K}]} \pi^*\theta = \int_A \theta \wedge \pi_!(1).$$

In the case that \mathcal{K} is compact and A is a smooth manifold, the above equation holds for θ any closed differential form on A .

Proof Because e represents the Thom class for W ,

$$\int_{[\mathcal{K}](i)} r\pi^*\theta = \int_{\pi(i)^*W} r\pi^*\theta \wedge e,$$

where we have abused notation a little to indicate the lift of $\pi^*\theta$ to π^*W simply as $\pi^*\theta$. This is not to be confused with the pullback of θ to π^*W using $\hat{\pi}$. The definition of integration along the fiber gives that

$$\int_{\pi(i)^*W} \hat{\pi}^*\theta \wedge re = \int_A \theta \wedge \hat{\pi}(i)_!(re).$$

Therefore,

$$\int_{[\mathcal{K}]} \pi^*\theta - \int_A \theta \wedge \pi_!(1) = \sum_{O_k} \sum_{i \in I(O_k)} \mu(i) \int_{\pi(i)^*W} r(\pi^*\theta - \hat{\pi}^*\theta) \wedge e.$$

As the vectorbundle map $W \rightarrow A$ and the map $x: W \rightarrow A$ are homotopic, there is a differential form α in Ω^*W such that $d\alpha$ is the difference between $x^*\theta$ and the lift of θ using the vectorbundle map $W \rightarrow A$.

In particular, on π^*W ,

$$d\pi^*\alpha = \hat{\pi}^*\theta - \pi^*\theta.$$

The same argument as the proof of Lemma 5.10 gives that the integral of $d(\pi^*\alpha \wedge e)$ over the pullback of W to $[\mathcal{K}]$ vanishes, so

$$\sum_{O_k} \sum_{i \in I(O_k)} \mu(i) \int_{\pi(i)^*W} r(\pi^*\theta - \hat{\pi}^*\theta) \wedge \pi^*e = 0.$$

Therefore, we get the required equality,

$$\int_{[\mathcal{K}]} \pi^*\theta = \int_A \theta \wedge \pi_!(1). \quad \square$$

Theorem 5.20 *If \mathcal{K}_0 and \mathcal{K}_1 , complete and oriented over \mathbf{Z} , are cobordant within a stack \mathcal{X} with a map $\pi: \mathcal{X} \rightarrow A$, then, given any construction of $[\mathcal{K}_j]$, the two composite maps*

$${}^rH^*(\mathcal{X}) \xrightarrow{\iota_j^*} {}^rH^*(\mathcal{K}_j) \xrightarrow{\pi_!} {}^rH^*A$$

are equal, and the same holds for the analogous maps

$${}^r_{\text{fg}}H^*(\mathcal{X}) \xrightarrow{\iota_j^*} {}^r_{\text{fg}}H^*(\mathcal{K}_j) \xrightarrow{\pi_!} {}^r_{\text{fg}}H^*A.$$

Moreover given any complex vectorbundle W over \mathcal{X} , and construction of any characteristic class $c(W)$ on \mathcal{K}_j as in Remark 5.2, the maps

$$\theta \mapsto \pi_!(c(W)) \wedge \iota_j^*\theta$$

are equal on the level of cohomology for $j = 0, 1$, in both ${}^rH^$ and ${}^r_{\text{fg}}H^*$.*

If \mathcal{K}_j are only proper over \mathbf{Z} , and A is a manifold, the corresponding maps

$${}^r_{\text{fg}}H^*(\mathcal{X}) \rightarrow H^*(A)$$

are equal.

Proof This theorem follows from the fact that any two choices of $[\mathcal{K}_j]$ and $c(W)$ are cobordant.

More precisely, Lemma 4.5 gives that we can choose a cobordism $\mathcal{K} \subset \mathcal{X} \times \mathbb{R}$ between \mathcal{K}_j , and construct \mathcal{K}_ϵ and I on \mathcal{K} restricting to the given $\mathcal{K}_{j,\epsilon}$ and I_j . By reparametrizing \mathbb{R} if necessary, we may also assume that there are connected open neighborhoods U_0 of 0 and U_1 of 1 such that, on each of these neighborhoods, $\mathcal{K}_{j,\epsilon}$ and I_j are pullbacks of the respective choices under the maps $\mathcal{X} \times U_j \rightarrow \mathcal{X}$.

Proposition 2.3 then implies that by possibly shrinking U_i , and by increasing ϵ to $\epsilon' < \frac{1}{2}$, we may choose a global section ν of S^I over $\mathcal{K}_{\epsilon'}$ so that ν restricted to U_j is the pullback of the corresponding section ν_j under the maps $\mathcal{K}_{\epsilon'}|_{U_j} \rightarrow \mathcal{K}_j$. (Note that increasing $\epsilon' < \frac{1}{2}$ does not affect $[\mathcal{K}]$ at all, because $[\mathcal{K}]$ has image inside $\mathcal{K}_C \subset \mathcal{K}_{\epsilon'}$.) Similarly, as the sheaves of unitary metrics and connections on W satisfy the patching, extension and averaging axioms, we may choose a unitary metric and connection on W over $\mathcal{K}_{\epsilon'}$ that restricted to U_j is the pullback of the corresponding choices used to define $c(W)$. It follows that defining $c(W) \in {}^r_{\text{fg}}\Omega^*\mathcal{K}_{\epsilon'}$ using the Chern–Weil construction on $\mathcal{K}_{\epsilon'}$, we obtain a form that, restricted to U_j , is the pullback of the corresponding forms on $\mathcal{K}_{j,\epsilon'}$.

As allowed by Lemma 5.18, we may assume that the auxiliary choices of (W, x, e) from Definition 5.15 used to define our two choices of $\pi_!$ are the same, and we may also assume that $W = W' \times \mathbb{R}$ and that x factors through $W \times \mathbb{R} \rightarrow W'$ and is a submersion restricted to fibers. If we pull back these choices to define $\pi'_! : {}^r\Omega^*(\mathcal{K}_{\epsilon'}) \rightarrow {}^r\Omega^*(\mathcal{A} \times \mathbb{R})$ as in Definition 5.15, the pulled-back x may not satisfy the transversality from condition (iii) of Definition 5.15 outside of $\mathcal{A} \times U_j$. After possibly shrinking U_j , we can modify this pulled-back x outside of $\mathcal{A} \times U_j$ to satisfy this condition (by changing x in the \mathbb{R} direction so that it becomes a submersion restricted to fibers where necessary).

Restricting attention to $U_j \subset \mathbb{R}$, we may consider $[\mathcal{K}]$ on families of the form $O \times U_j$. Here $I(O \times U_j)$ may be identified with $I(O)$ from the corresponding construction of $[\mathcal{K}_j]$. Moreover, for all $i \in I(O \times U_j) = I(O)$, $[\mathcal{K}](i) = [\mathcal{K}_j](i) \times U_j$.

It follows that given any θ in ${}^r\Omega^*(\mathcal{X})$ or ${}^r_{\text{fg}}\Omega^*(\mathcal{X})$, respectively, and 1-form θ_0 supported inside $U_j \subset \mathbb{R}$,

$$\pi'_!(\theta \wedge \theta_0) = \pi_!(\theta) \wedge \theta_0,$$

where we abuse notation to think of θ also on $\mathcal{X} \times \mathbb{R}$, and θ_0 as also in both ${}^r_{\text{fg}}\Omega^1(\mathcal{X} \times \mathbb{R})$ and ${}^r_{\text{fg}}\Omega^1(\mathcal{A} \times \mathbb{R})$. The above equality holds for the two different choices of $\pi_!$, depending on the choice of U_j . The fact that $\pi'_!$ commutes with d then implies that on the level of cohomology, $\pi_!$ does not depend on these choices. With an even more egregious abuse of notation, it also follows that

$$\pi'_!(\theta \wedge c(W) \wedge \theta_0) = \pi_!(\theta \wedge c(W)) \wedge \theta_0,$$

so on the level of cohomology, the two different maps $\theta \mapsto \pi_!(c(W) \wedge \theta)$ are equal for $j = 0, 1$. \square

Proposition 5.21 Suppose that \mathcal{K} is complete over \mathbf{Z} . Given a representable submersion $\mathbf{Z}' \rightarrow \mathbf{Z}$ (Definition 4.1), let \mathcal{K}' be the pullback of \mathcal{K} . Let π and π' be compatible maps

$$\begin{array}{ccc} \mathcal{K}' & \longrightarrow & \mathcal{K} \\ \downarrow \pi' & & \downarrow \pi \\ \mathbf{A} \times_{\mathbf{Z}} \mathbf{Z}' & \xrightarrow{y} & \mathbf{A} \\ \downarrow & & \downarrow \\ \mathbf{Z}' & \longrightarrow & \mathbf{Z} \end{array}$$

Then the following diagrams commute:

$$\begin{array}{ccc} {}^rH^*(\mathcal{K}') & \longleftarrow & {}^rH^*(\mathcal{K}) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^rH^*(\mathbf{A} \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^rH^*(\mathbf{A}) \end{array} \qquad \begin{array}{ccc} {}^r_{\text{fg}}H^*(\mathcal{K}') & \longleftarrow & {}^r_{\text{fg}}H^*(\mathcal{K}) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^r_{\text{fg}}H^*(\mathbf{A} \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^r_{\text{fg}}H^*(\mathbf{A}) \end{array}$$

Proof Every step of the construction of $[\mathcal{K}]$ pulls back under representable submersions. (In general, sections of S from Definition 3.12 may not pull back to transverse sections, but this is not a problem in the case that $\mathbf{Z}' \rightarrow \mathbf{Z}$ is a submersion.) Therefore, we may take $[\mathcal{K}']$ to be the pullback of $[\mathcal{K}]$.

Explicitly, if I' is the pullback of the weighted branched cover I used to define $[\mathcal{K}]$, for any $O \in \mathcal{O}_I$, there is a map $\iota^*: I(O) \rightarrow I'(O \times_{\mathbf{Z}} \mathbf{Z}')$ (so long as $O \times_{\mathbf{Z}} \mathbf{Z}'$ is connected) and given any $i \in I(O)$,

$$[\mathcal{K}'](\iota^*i) = [\mathcal{K}](i) \times_{\mathbf{Z}} \mathbf{Z}'.$$

When $O \times_{\mathbf{Z}} \mathbf{Z}'$ is not connected, the left-hand side of the above expression must be replaced with $\coprod_i [\mathcal{K}'](\iota^*i)$, where the ι are the maps from each connected component of $O \times_{\mathbf{Z}} \mathbf{Z}'$ to O . For notational convenience, we shall continue with the case that $O \times_{\mathbf{Z}} \mathbf{Z}'$ is connected, and simply use i to indicate ι^*i .

The choices of W , x and e from Definition 5.15 are more problematic to pull back. Choose an oriented vectorbundle W over \mathbf{A} and map $x: W \rightarrow \mathbf{A}$ satisfying the conditions of Definition 5.15. Then consider the pullback y^*W of W to $\mathbf{A} \times_{\mathbf{Z}} \mathbf{Z}'$. If $\mathbf{Z}' \rightarrow \mathbf{Z}$ is not proper, then the map x may not correspond compatibly to a map $x': y^*W \rightarrow \mathbf{A} \times_{\mathbf{Z}} \mathbf{Z}'$. On the other hand, the fact that $\mathbf{Z}' \rightarrow \mathbf{Z}$ is a submersion implies

that we may choose $x': y^*W \rightarrow A \times_{\mathbf{Z}} \mathbf{Z}'$ so that, restricted to some neighborhood of the zero section in y^*W , the outer circuit of the following diagram commutes:

$$\begin{array}{ccc} y^*W & \longrightarrow & W \\ x' \left(\downarrow \right. & & \left. \downarrow \right) x \\ A \times_{\mathbf{Z}} \mathbf{Z}' & \xrightarrow{y} & A \end{array}$$

If we also choose x' to be a fiber-preserving, submersive reparametrization of the pullback of x , then x' will still satisfy all conditions of Definition 5.15, including the transversality condition from part (iii). Because the outer circuit of the above diagram commutes on a neighborhood of the zero section, there is a natural map of this neighborhood into a different fiber product, $W_{x \times y} (A \times_{\mathbf{Z}} \mathbf{Z}')$, that uses the map $x: W \rightarrow A$ instead of the vectorbundle projection. Because x and x' coincide with the vectorbundle projections when restricted to the zero section, restricted to a sufficiently small neighborhood N of the zero section, this natural map is an isomorphism onto an open subset of $W_{x \times y} (A \times_{\mathbf{Z}} \mathbf{Z}')$.

Given any compactly contained open subset O of $A \times_{\mathbf{Z}} \mathbf{Z}'$, by modifying x' out of the neighborhood N if necessary, we may assume that there exists a neighborhood U of the zero section in W such that $y^*U \cap (x')^{-1}O \subset N$, so we can consider $U_{x \times y} O$ to be an open subset of $N \subset y^*W$.

Recall from Definition 5.15 that $\pi_!$ is defined using pushforward along the maps $\hat{\pi}(i)$ defined using the diagram

$$\begin{array}{ccccc} & & \hat{\pi}(i) & & \\ & \searrow & & \nearrow & \\ \pi(i)^*W & \longrightarrow & W & \xrightarrow{x} & A \\ \downarrow & & \downarrow & & \\ [\mathcal{K}](i) & \xrightarrow{\pi(i)} & A & & \end{array}$$

Similarly, define $\hat{\pi}'(i)$ using the diagram

$$\begin{array}{ccccc} & & \hat{\pi}'(i) & & \\ & \searrow & & \nearrow & \\ (y \circ \pi'(i))^*W & \longrightarrow & y^*W & \xrightarrow{x'} & A \times_{\mathbf{Z}} \mathbf{Z}' \\ \downarrow & & \downarrow & & \\ [\mathcal{K}'](i) & \xrightarrow{\pi'(i)} & A \times_{\mathbf{Z}} \mathbf{Z}' & & \end{array}$$

Now, consider the diagram

$$\begin{array}{ccc}
 (y \circ \pi'(i))^* W(\hat{y}) & \longrightarrow & \pi(i)^* W \\
 \hat{\pi}'(i) \downarrow & & \downarrow \hat{\pi}(i) \\
 y^* W & \longrightarrow & W \\
 \downarrow x' & & \downarrow x \\
 A \times_{\mathbf{Z}} \mathbf{Z}' & \xrightarrow{y} & A
 \end{array}$$

The top square above is a fiber product diagram, and the bottom square is isomorphic to a fiber product diagram when W is restricted to U , $A \times_{\mathbf{Z}} \mathbf{Z}'$ is restricted to O and $y^* W$ is restricted accordingly. Therefore, the outer circuit of the above diagram is a fiber product diagram when $A \times_{\mathbf{Z}} \mathbf{Z}'$ is restricted to O , $\pi(i)^* W$ is restricted to $\pi(i)^* U$ and $(y \circ \pi'(i))^* W$ is restricted accordingly.

Lemma 9.3 from [31] implies that the following diagram commutes:

$$\begin{array}{ccc}
 {}^r\Omega(\pi(i)^* U \hat{\pi}(i) \times_y O) & \xleftarrow{\hat{y}^*} & {}^r\Omega(\pi(i)^* U) \\
 \downarrow \hat{\pi}'(i)_! & & \downarrow \hat{\pi}(i)_! \\
 {}^r\Omega(O) & \xleftarrow{y^*} & {}^r\Omega(A)
 \end{array}$$

Therefore, the following diagram obeys the restricted commutativity condition described below:

$$\begin{array}{ccc}
 {}^r\Omega((y \circ \pi'(i))^* W) & \xleftarrow{\hat{y}^*} & {}^r\Omega(\pi(i)^* W) \\
 \downarrow \hat{\pi}'(i)_! & & \downarrow \hat{\pi}(i)_! \\
 {}^r\Omega(A \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^r\Omega(A)
 \end{array}$$

If $\theta \in {}^r\Omega(\pi(i)^* W)$ has support compactly contained in $\pi(i)^* U$, then $\hat{\pi}'(i)_! \hat{y}^* \theta$ restricted to O is equal to $y^* \hat{\pi}(i)_! \theta$.

Define $\pi_!$ as in Definition 5.15 using W , x and a Thom form e supported in U . If we define $\pi'_!$ analogously using the pullback of W , x' and the pullback of e , then the following diagram commutes:

$$\begin{array}{ccccc}
 & & {}^r\Omega^*(\mathcal{K}'_\epsilon) & \longleftarrow & {}^r\Omega^*(\mathcal{K}_\epsilon) \\
 & \swarrow \pi'_! & & & \downarrow \pi_! \\
 {}^r\Omega^*(A \times_{\mathbf{Z}} \mathbf{Z}') & \longrightarrow & {}^r\Omega^*(O) & \xleftarrow{y^*} & {}^r\Omega^*(A)
 \end{array}$$

As integration along the fiber and pullbacks preserve the subspace of refined differential forms generated by functions, the same diagram commutes with ${}^r_{\text{fg}}\Omega^*$ replacing ${}^r\Omega^*$.

If we can choose $O \subset A \times_Z Z'$ to be a smooth retraction of $A \times_Z Z'$ so that the restriction maps ${}^rH^*(A \times_Z Z') \rightarrow {}^rH^*(O)$ and ${}^r_{\text{fg}}H^*(A \times_Z Z') \rightarrow {}^r_{\text{fg}}H^*(O)$ are isomorphisms, it follows the required diagrams commute:

$$\begin{array}{ccc} {}^rH^*(\mathcal{K}'_\epsilon) & \longleftarrow & {}^rH^*(\mathcal{K}_\epsilon) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^rH^*(A \times_Z Z') & \xleftarrow{y^*} & {}^rH^*(A) \end{array} \quad \begin{array}{ccc} {}^r_{\text{fg}}H^*(\mathcal{K}'_\epsilon) & \longleftarrow & {}^r_{\text{fg}}H^*(\mathcal{K}_\epsilon) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^r_{\text{fg}}H^*(A \times_Z Z') & \xleftarrow{y^*} & {}^r_{\text{fg}}H^*(A) \end{array}$$

More generally, ${}^rH^*(A \times_A Z')$ is the inverse limit of the system ${}^rH^*(O)$ for all open $O \subset A \times_A Z'$ with compact closure, and the same holds for ${}^r_{\text{fg}}H^*$. (It is obvious that ${}^r\Omega^*(A \times_A Z')$ is the inverse limit of ${}^r\Omega^*(O)$, and the system ${}^r\Omega^*(O)$ satisfies the Mittag-Leffler condition, so taking homology is compatible with inverse limits; the same holds for refined forms generated by functions.) It follows that the above diagrams always commute. \square

Theorem 5.22 Suppose that \mathcal{K} is complete over Z . Let \mathcal{K}' be the pullback of \mathcal{K} over a representable map $Z' \rightarrow Z$ (Definition 4.1). Given compatible maps

$$\begin{array}{ccc} \mathcal{K}' & \longrightarrow & \mathcal{K} \\ \downarrow \pi'_! & & \downarrow \pi \\ A \times_Z Z' & \xrightarrow{y} & A \\ \downarrow & & \downarrow \\ Z' & \longrightarrow & Z \end{array}$$

the following diagrams commute:

$$\begin{array}{ccc} {}^rH^*(\mathcal{K}') & \longleftarrow & {}^rH^*(\mathcal{K}) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^rH^*(A \times_Z Z') & \xleftarrow{y^*} & {}^rH^*(A) \end{array} \quad \begin{array}{ccc} {}^r_{\text{fg}}H^*(\mathcal{K}') & \longleftarrow & {}^r_{\text{fg}}H^*(\mathcal{K}) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^r_{\text{fg}}H^*(A \times_Z Z') & \xleftarrow{y^*} & {}^r_{\text{fg}}H^*(A) \end{array}$$

Proof Extend $Z' \rightarrow Z$ to a map $X \rightarrow Z$ where

- X is a vectorbundle over Z' ,

- restricted to some tubular neighborhood O_1 of the zero section, the map $X \rightarrow Z$ factors through $X \rightarrow Z' \rightarrow Z$,
- restricted to some other open set $O_2 \subset X$, $X \rightarrow Z$ is a submersion with nonempty convex fibers.

Because X is a vectorbundle over Z' , and $Z' \rightarrow Z$ is representable, our extension $X \rightarrow Z$ is also representable. Use \mathcal{K}_X and \mathcal{K}_{O_i} to denote the respective pullbacks of \mathcal{K} . Proposition 5.21 implies that all inner loops in the following diagram commute:

$$(3) \quad \begin{array}{ccccccc} & & \swarrow & & \searrow & & \\ & & \mathcal{K}_{O_1} & \longleftarrow & \mathcal{K}_X & \longrightarrow & \mathcal{K}_{O_2} & \longleftarrow & \mathcal{K} \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ rH^*(O_1 \times_Z A) & \longleftarrow & rH^*(X \times_Z A) & \longrightarrow & rH^*(O_2 \times_Z A) & \longleftarrow & rH^*A \\ & \nwarrow & & \swarrow & & \nwarrow & & \swarrow & \end{array}$$

and the same holds with ${}^r_{fg}H^*$ replacing ${}^rH^*$. Because the restriction maps

$${}^rH^*(X \times_Z A) \rightarrow {}^rH^*(O_i \times_Z A)$$

are isomorphisms, every loop in the above diagram commutes, and the same holds with ${}^r_{fg}H^*$ replacing ${}^rH^*$. Proposition 5.21 also implies the commutativity of the inner square of the diagram

$$(4) \quad \begin{array}{ccc} & \swarrow & \\ {}^rH^*(\mathcal{K}_{O_1}) & \longleftarrow & {}^rH^*(\mathcal{K}') \\ \downarrow & & \downarrow \\ {}^rH^*(O_1 \times_Z A) & \longleftarrow & {}^rH^*(Z' \times_Z A) \\ & \swarrow & \end{array}$$

and the analogous diagram with ${}^r_{fg}H^*$ replacing ${}^rH^*$. The horizontal maps in the above diagram are defined using the projection $O_1 \rightarrow Z'$ and inclusion $Z' \rightarrow O_1$. We shall now show that each pair of horizontal maps in the diagram above are inverse isomorphisms, proving that the above diagram commutes. In each case, the leftward map followed by the rightward is the identity on forms, so it remains to prove that the other composition is the identity on the level of cohomology.

Let $\Psi_t: O_1 \rightarrow O_1$ for $t \in [0, 1]$ be any fiber-preserving smooth homotopy that is the identity at $t = 0$ and the projection onto $Z' \subset O_1$ at $t = 1$. Then, given a form $\theta \in {}^r\Omega^* O_1$, define

$$K\theta = \int_0^1 \Psi_t^* i_{\partial\Psi_t/\partial t} \theta \, dt.$$

Then

$$(dK + Kd)\theta = \Psi_1^* \theta - \Psi_0^* \theta.$$

It follows that if θ is closed, it represents the same cohomology class as $\Psi_1^* \theta$. As the homotopy preserves the fibers of the map to Z , it induces a smooth homotopy from the identity on $O_1 \times_Z A$ to the projection to $Z' \times_Z A$. Making the above argument with this new homotopy gives that the composition of the lower pair of horizontal maps in the above diagram is the identity on the level of cohomology. The same holds with ${}^r_{\text{fg}}H^*$ replacing ${}^rH^*$ because K preserves the subspace of forms generated by functions. It remains to show the same for the upper pair of maps.

Let us consider the homotopy induced by ψ_t on \mathcal{K}_{O_1} , and define an operator analogous to the above K . Start with a family of curves \hat{f} in \mathcal{K}_{O_1} . The map $\mathcal{K}_{O_1}^{\text{st}} \rightarrow (\mathcal{K}')^{\text{st}}$ applied to \hat{f} gives a family of curves \hat{f}_0 with $F(\hat{f}_0) = F(\hat{f})$. We may then pull this family \hat{f}_0 back to get a bigger family of curves $\hat{f}_0 \times_Z O_1$ in $\mathcal{K}_{O_1}^{\text{st}}$ with domain $F(\hat{f}) \times_{Z'} O_1$. Alternatively, using the inclusion $\hat{B}' \subset \hat{O}_1$, we can consider \hat{f}_0 as a family of curves in $\mathcal{K}_{O_1}^{\text{st}}$. As the homotopy $\Psi_t: O_1 \rightarrow O_1$ preserves fibers of the map $O_1 \rightarrow Z'$, it induces a homotopy of $F(\hat{f}_0 \times_{Z'} O_1)$ that is the identity at $t = 0$, and the projection onto $F(\hat{f}_0)$ at $t = 1$. Again call this homotopy Ψ_t , and define $K: {}^r\Omega^* F(\hat{f}_0 \times_{Z'} O_1) \rightarrow {}^r\Omega^{*-1} F(\hat{f}_0 \times_{Z'} O_1)$ as in the above equation. Then, given any $\theta \in {}^rH^*(\mathcal{K}_{\hat{O}_1})$, we can define $K\theta$ on $F(\hat{f})$ by defining $K\theta$ on $F(\hat{f}_0 \times_{Z'} O_1)$ as above, then pulling back this $K\theta$ to $F(\hat{f})$ via the natural inclusion $\hat{f} \rightarrow \hat{f}_0 \times_{Z'} O_1$. Such a $K\theta$ gives a well-defined form in ${}^r\Omega^*(\mathcal{K}_{O_1})$ because given any map of curves $\hat{g} \rightarrow \hat{f}$, the corresponding diagram

$$\begin{array}{ccc} F(\hat{g}_0 \times_{Z'} O_1) & \longrightarrow & F(\hat{f}_0 \times_{Z'} O_1) \\ \downarrow \Psi_t & & \downarrow \Psi_t \\ F(\hat{g}_0 \times_{Z'} O_1) & \longrightarrow & F(\hat{f}_0 \times_{Z'} O_1) \end{array}$$

commutes. Now, $dK\theta$ is the difference between θ and the composition of the topmost pair of maps in diagram (4) applied to θ , so this topmost pair of maps are inverse isomorphisms on the level of cohomology. As all our maps preserve ${}^r_{\text{fg}}\Omega^* \subset {}^r\Omega^*$, the same holds with ${}^r_{\text{fg}}H^*$ replacing ${}^rH^*$.

As the horizontal maps in diagram (4) are inverse isomorphisms, diagram (4) commutes, and combines with the outermost loop of diagram (3) to give our required commutative diagram:

$$\begin{array}{ccc} {}^rH^*(\mathcal{K}') & \xleftarrow{\quad} & {}^rH^*(\mathcal{K}) \\ \downarrow \pi'_! & & \downarrow \pi_! \\ {}^rH^*(A \times_{\mathbf{Z}} \mathbf{Z}') & \xleftarrow{y^*} & {}^rH^*(A) \end{array}$$

The same argument applies to the analogous diagram replacing ${}^rH^*$ with ${}^r_{\text{fg}}H^*$. \square

6 Weak products of Kuranishi categories

In order to discuss gluing theorems for Gromov–Witten invariants, it is useful to have a notion of fiber products of Kuranishi structures or Kuranishi categories. The tropical gluing formula for Gromov–Witten invariants, equation (1) of [30], is proved by expressing a Kuranishi category associated to some tropical curve γ as a (weak) fiber product of Kuranishi categories associated to the vertices, v , of γ . We can decompose such a (weak) fiber product as a pullback, dealt with in Theorem 5.22, and a (weak) product, defined below and dealt with in Theorem 6.2.

Because of our choice that all Kuranishi charts in an embedded Kuranishi structure should be compatible, and the related part (iii) of Definition 2.1, and part (iii) of Definition 3.2, the product of Kuranishi categories is no longer a Kuranishi category, because the obvious candidates for charts on the product will no longer be compatible.¹⁵

Nevertheless, we can construct a weak product of Kuranishi categories by taking the product of charts, and then shrinking these charts appropriately to make them compatible.

Definition 6.1 Say that \mathcal{K} is a weak product of some finite collection \mathcal{K}_v of Kuranishi categories if the following holds:

- (i) $\mathcal{K}^{\text{st}} \subset \prod_v \mathcal{K}_v^{\text{st}}$.
- (ii) $\mathcal{K}^{\text{hol}} = \prod_v \mathcal{K}_v^{\text{hol}}$.
- (iii) For each chart $(\mathcal{U}_i, V_i, \hat{f}_i/G_i)$ on \mathcal{K} , there are corresponding charts

$$(\mathcal{U}_{i_v}, V_{i_v}, \hat{f}_{i_v}/G_{i_v})$$

¹⁵Most other authors work with a version of Kuranishi structures that avoid this problem. Our embedded Kuranishi structure could be regarded as akin to a choice of good coordinate charts in Fukaya, Oh, Ohta and Ono's approach [8].

on \mathcal{K}_v such that

- $\mathcal{U}_i \subset \prod_v \mathcal{U}_{i_v}$,
- $V_i = \bigoplus_v V_{i_v}$ on \mathcal{U}_i ,
- $G_i = \prod_v G_{i_v}$,
- \widehat{f}_i is a G_i -equivariant open subfamily of $\prod_v \widehat{f}_{i_v}$,
- $\bar{\partial} \widehat{f}_i$ is equal to $\prod_v \bar{\partial} \widehat{f}_{i_v}$ restricted to $F(\widehat{f}_i) \subset \prod_v F(\widehat{f}_{i_v})$.

Note that if \mathcal{K}_v are proper or complete and oriented over \mathbf{Z}_v , then \mathcal{K} is proper or complete and oriented over $\prod_v \mathbf{Z}_v$. Given maps $\mathcal{K}_v \rightarrow \mathbf{A}_v$, there is an obvious induced map $\mathcal{K} \rightarrow \prod_v \mathbf{A}_v$.

Theorem 6.2 Suppose that \mathcal{K} is a weak product of some finite collection of complete, oriented Kuranishi categories \mathcal{K}_v with maps $\pi_v: \mathcal{K}_v \rightarrow \mathbf{A}_v$; let $\pi: \mathcal{K} \rightarrow \prod_v \mathbf{A}_v$ be the induced map. Then the following diagrams commute:

$$\begin{array}{ccc} {}^r H^*(\mathcal{K}) & \xrightarrow{\pi!} & {}^r H^*(\prod_v \mathbf{A}_v) \\ \uparrow & \nearrow \Pi_v(\pi_v)! & \\ \prod_v {}^r H^*(\mathcal{K}_v) & & \end{array} \qquad \begin{array}{ccc} {}^r_{\text{fg}} H^*(\mathcal{K}) & \xrightarrow{\pi!} & {}^r_{\text{fg}} H^*(\prod_v \mathbf{A}_v) \\ \uparrow & \nearrow \Pi_v(\pi_v)! & \\ \prod_v {}^r_{\text{fg}} H^*(\mathcal{K}_v) & & \end{array}$$

Proof Using Lemma 2.6, and the fact that $\mathcal{K}^{\text{hol}} = \prod_v \mathcal{K}_v^{\text{hol}}$, we may construct functions $\rho_{i,v}: \mathcal{K}_v \rightarrow \mathbb{R}$ satisfying the requirements of Definition 3.9 so that the corresponding functions $\rho_i := \min_v \rho_{i,v}$ on \mathcal{K} also satisfy the requirements of Definition 3.9, where the chart \widehat{f}_i/G_i associated with ρ_i is locally the product of the charts $\widehat{f}_{i_v}/G_{i_v}$ associated with $\rho_{i,v}$. Indeed, each holomorphic curve $f = \prod_v f_v$ in $\widehat{f}_i \subset \prod_v \widehat{f}_{i_v}$ has some $(G_i$ -invariant) product neighborhood $\prod U_v \subset \widehat{f}_i$; we can choose G_{i_v} -invariant functions $\rho_{f,v}: \prod \widehat{f}_{i_v} \rightarrow [-1, 1]$ with $\rho_{f,v}(f_v) > \frac{1}{2}$ so that $\rho_{f,v} + 1$ is compactly supported within U_v , then use Lemma 2.6 on $\rho_{f,v} + 1$ to extend these functions to functions satisfying Definition 3.9(ii). The corresponding continuous function $\rho_f := \min_v \rho_{f,v}$ is also greater than $\frac{1}{2}$ on f , and satisfies Definition 3.9(ii)–(iii). Our functions ρ_i and $\rho_{i,v}$ will be chosen from such functions ρ_f and $\rho_{f,v}$ for holomorphic f . As we have assumed that \mathcal{K}_v are complete, \mathcal{K}^{hol} is covered by the sets where $\rho_i > 0$ for a finite number of such ρ_i , so we can achieve Definition 3.9(i) and (iv) with this finite collection of functions, ρ_i and $\rho_{i,v}$.

The corresponding $\mathcal{K}_\epsilon^{\text{st}}$ is a substack of $\prod_v \mathcal{K}_{v,\epsilon}^{\text{st}}$, because the subset of \widehat{f}_i where $\rho_i > \epsilon$ is equal to the product of the subsets of \widehat{f}_{i_v} where $\rho_{i,v} > \epsilon$ (but $\prod_v \mathcal{K}_{v,\epsilon}$ contains

more families than \mathcal{K}_ϵ because it also includes the products of the relevant subsets of \widehat{f}_{i_v} , where i may depend on v).

Choose separating weighted branched covers I_v of $\mathcal{K}_{v,\epsilon}$, and let I be the weighted branched cover of \mathcal{K}_ϵ that is the product of the pullback of these I_v via the maps $\mathcal{K}_\epsilon^{\text{st}} \rightarrow \mathcal{K}_{v,\epsilon}^{\text{st}}$. The fact that I_v are separating implies that I is also separating.

We must also choose metrics on V as in Lemma 3.11. We need to do this so that “small enough” perturbations of $\bar{\partial}$ on $\mathcal{K}_{v,\epsilon}$ pull back to “small enough” perturbations of $\bar{\partial}$ on \mathcal{K}_ϵ . The required condition on \mathcal{K}_ϵ is that wherever $|\bar{\partial}\widehat{f}_i| \leq 1$, we have some $\rho_j > \frac{1}{2}$, and the corresponding condition on $\mathcal{K}_{v,\epsilon}$ is that wherever $|\bar{\partial}\widehat{f}_{i_v}| \leq 1$, we have some $\rho_{j_v} > \frac{1}{2}$. If there are n indices v , we need to strengthen these latter conditions so that wherever $|\bar{\partial}\widehat{f}_{i_v}| \leq n$ for all v , there exists some j such that $\rho_{j_v} > \frac{1}{2}$ for all v . Essentially, we need that $\{|\bar{\partial}\widehat{f}_{i_v}| \leq n\}$ is contained in some open neighborhood U_v of $\mathcal{K}_v^{\text{hol}}$ and that $\prod_v U_v$ is covered by the sets where $\rho_{j_v} > \frac{1}{2}$ for all v . This is possible because $\prod_v \mathcal{K}_v^{\text{hol}}$ is compact and the sets where all $\rho_{j_v} > \frac{1}{2}$ form an open cover of $\prod_v \mathcal{K}_v^{\text{hol}} \subset \prod_v \mathcal{K}_v^{\text{st}}$. Once such U_v have been chosen, the proof of Lemma 3.11 applies to show that there exists a global choice of metric on V_v over $\mathcal{K}_{v,\epsilon}$ such that $|\bar{\partial}\widehat{f}_{i_v}| \leq n$ is contained in U_v .

If we use $1/n$ times the product of these metrics on V over \mathcal{K}_ϵ , then this metric satisfies the conditions of Lemma 3.11, and the pullback of any product of sections of V_v that are smaller than 1 is a section of V that is smaller than 1.

We have now made all the choices to define the sheaves S and S_v so that any choice of sections of S_v over \widehat{f}_{i_v} for all v pull back to a section of S over \widehat{f}_i . Recall from Definition 3.12 that a section of S over \widehat{f}_i is a section ν of V_i such that the following holds:

- $\bar{\partial}\widehat{f}_i - \nu$ is contained in $V_j \subset V_i$ on a neighborhood of wherever $\rho_j \geq 0$; a stronger condition is satisfied by sections pulled back from \widehat{f}_{i_v} . These pulled-back sections will be contained in $\bigoplus_v V_{j_v}$ on some neighborhood of where $\rho_{j_v} \geq 0$ for all v . This is stronger because j_v is allowed to depend on v .
- ν is close to $\bar{\partial}\widehat{f}_i$ in the sense that $|\bar{\partial}\widehat{f}_i - \nu| < 1$; we have chosen our metrics so that pullbacks of sections of $\prod_v S_v$ satisfy this condition.
- ν is transverse to 0; pullbacks of sections of $\prod_v S_v$ also clearly satisfy this condition.

Similarly, if a family \widehat{f} in \mathcal{K}_ϵ is a product of families \widehat{f}_v in $\mathcal{K}_{v,\epsilon}$, any choice of section of S_v over \widehat{f}_v for all v pulls back to a section of S over \widehat{f} .

As our weighted branched cover I of \mathcal{K}_ϵ is the pullback of the corresponding weighted branched covers I_v of \mathcal{K}_ϵ , we may also pull back a choice of global section of $S_v^{I_v}$ over $\mathcal{K}_{v,\epsilon}$ for all v to give a global section of S^I over \mathcal{K}_ϵ . In particular, suppose that \hat{f} in \mathcal{K}_ϵ is a product of families \hat{f}_v in $\mathcal{K}_{\epsilon,v} \cap \mathcal{O}_{I_v}$; then \hat{f} is in \mathcal{O}_I , and $I(\hat{f}) = \prod_v I_v(\hat{f}_v)$. On \hat{f}_v a section of $S_v^{I_v}$ is a choice of section $\nu_v(i)$ of $S_v(\hat{f}_v)$ for all $i \in I_v$ such that if i and j are not separated at $f \in \hat{f}_v$, then $\nu_v(i) = \nu_v(j)$ on a neighborhood of f . Taking the product over v of such sections gives sections $\nu(i)$ of $S_v(\hat{f})$ for all i in $I(\hat{f})$, where if i is the product of i_v , then $\nu(i)$ is the product of $\nu_v(i_v)$. Again, $\nu(i)$ will be equal to $\nu(j)$ on a neighborhood of any curve $f \in \hat{f}$ where i is not separated from j , so ν defines a section of $S^I(\hat{f})$.

Any map $\hat{f} \rightarrow \hat{g}$ in \mathcal{K}_ϵ between families of curves that are products of families in $\mathcal{K}_{v,\epsilon}$ corresponds to maps $\hat{f}_v \rightarrow \hat{g}_v$. Global sections of $S_v^{I_v}$ are compatible with pullbacks, so the corresponding sections of $S^I(\hat{f})$ and $S^I(\hat{g})$ are compatible with pullbacks. As any family in \mathcal{K}_ϵ is locally a product of families in $\mathcal{K}_{v,\epsilon}$, it follows that pulling back global sections of $S_v^{I_v}$ gives global sections of S^I .

Intersecting a global section of S^I with the zero section gives the virtual class $[\mathcal{K}]$; $\nu(i)$ intersected with 0 gives $[\mathcal{K}](i)$, an oriented family in $\mathcal{K}_C \subset \mathcal{K}_\epsilon^{\text{st}}$. Unfortunately, the map $\mathcal{K}^{\text{st}} \rightarrow \prod_v \mathcal{K}_v^{\text{st}}$ does not pull back $\prod_v \mathcal{K}_{v,C}$ to be contained in \mathcal{K}_C , but, actually, we have that these families $[\mathcal{K}_v](i)$ are contained in the substack of $\mathcal{K}'_{v,C} \subset \mathcal{K}_{v,C}$ where $\bar{\partial} < 1$, and we have constructed our metrics so that $\prod_v \mathcal{K}'_{v,C}$ pulls back to be contained in \mathcal{K}_C .

Make the auxiliary choices (W, x, e) required by Definition 5.15 for defining $\pi_!$ and $(\pi_v)_!$ so that the diagram

$$(5) \quad \begin{array}{ccc} r\Omega^*(\mathcal{K}) & \xrightarrow{\pi_!} & r\Omega^*(\prod_v \mathcal{A}_v) \\ \uparrow & \nearrow \Pi_v(\pi_v)_! & \\ \prod_v r\Omega^*(\mathcal{K}_v) & & \end{array}$$

commutes. In particular, choose (W_v, x_v, e_v) satisfying the requirements of Definition 5.15, so that $x_v: W_v \rightarrow \prod_v \mathcal{A}_v$ is a submersion restricted to fibers. Then we can take $W = \prod_v W_v$, $x = \prod_v x_v: W \rightarrow \prod_v \mathcal{A}_v$ and $e = \prod_v e_v$. As x is a submersion restricted to fibers, (W, x, e) also satisfies the requirements of Definition 5.15. To avoid needing to think about orientation complications, we may also choose each W_v to have even rank, so that the Thom forms e_v are even-dimensional.

We need to prove that diagram (5) commutes. As the maps in (5) are linear, we may use partitions of unity (on $\mathcal{K}_v^{\text{st}}$) to reduce to the case of differential forms with small support on \mathcal{K}_C and $\mathcal{K}'_{v,C}$. In particular, without losing generality, assume the following:

- $\theta \in {}^r\Omega^*\mathcal{K}$ has support on \mathcal{K}_C compactly contained in \hat{f}/G .
- \hat{f} is in $\mathcal{K}_\epsilon \cap \mathcal{O}_I$ and is a product of families \hat{f}_v in $\mathcal{K}_{v,\epsilon} \cap \mathcal{O}_{I_v}$.
- \hat{f}/G and \hat{f}_v/G_v represent substacks.
- $G = \prod_v G_v$.
- θ is the product of the pullback of $\theta_v \in {}^r\Omega^*(\mathcal{K}_v)$.
- θ_v has support on $\mathcal{K}'_{v,C}$ compactly contained in \hat{f}_v/G_v .
- There are partitions of unity (Definition 5.6) compatible with $[\mathcal{K}]$ and $[\mathcal{K}_v]$ such that one of the connected families \mathcal{O}_1 involved is \hat{f} (resp. \hat{f}_v), and such that the corresponding function r is $1/|G|$ (resp. $1/|G_v|$) when restricted to the support of θ (resp. θ_v) on \hat{f} (resp. \hat{f}_v).

Then $I(\hat{f}) = \prod_v I_v(\hat{f}_v)$, and for $i = \prod_v i_v \in I(\hat{f})$, we have $[\mathcal{K}](i) = \prod_v [\mathcal{K}_v](i_v)$. We have also chosen (W, x) and (W_v, x_v) so that the map $\hat{\pi}(i): \pi(i)^*W \rightarrow \prod_v A_v$ is equal to the product of the maps $\hat{\pi}_v(i_v): \pi_v(i)^*W_v \rightarrow \prod_v A_v$ from Definition 5.15(iii). It follows that $\hat{\pi}(i)_!$ is the product of $\hat{\pi}_v(i_v)_!$. Then, for θ and θ_v satisfying the conditions listed above,

$$\begin{aligned}
 (6) \quad \pi_!(\theta) &= \frac{1}{|G|} \sum_{i \in I(\hat{f})} \mu(i) \hat{\pi}(i)_!(\theta \wedge e) \\
 &= \prod_v \frac{1}{|G_v|} \sum_{i_v \in I_v(\hat{f}_v)} \mu(i_v) \hat{\pi}_v(i_v)_!(\theta_v \wedge e_v) = \prod_v (\pi_v)_!(\theta_v).
 \end{aligned}$$

Note that $\pi_!$ does not depend on the choice of partition of unity, so diagram (5) commutes on forms that are a sum of forms obeying the condition listed above and hence satisfying equation (6). Any choice of form in $\prod_v {}^r\Omega^*(\mathcal{K}_v)$ may be expressed as a sum of forms obeying the conditions listed above and forms that for some v vanish on $\mathcal{K}'_{v,C}$. Both $\pi_!$ and $\prod_v (\pi_v)_!$ vanish on such vanishing forms; therefore, (5) commutes, as required.

On the level of cohomology, $\pi_!$ and $\prod_v (\pi_v)_!$ are independent of any choices; therefore, in general, the following diagrams commute:

$$\begin{array}{ccc}
 {}^rH^*(\mathcal{K}) & \xrightarrow{\pi!} & {}^rH^*(\prod_v A_v) \\
 \uparrow & \nearrow \Pi_v(\pi_v)! & \\
 \prod_v {}^rH^*(\mathcal{K}_v) & &
 \end{array}
 \qquad
 \begin{array}{ccc}
 {}^r_{\text{fg}}H^*(\mathcal{K}) & \xrightarrow{\pi!} & {}^r_{\text{fg}}H^*(\prod_v A_v) \\
 \uparrow & \nearrow \Pi_v(\pi_v)! & \\
 \prod_v {}^r_{\text{fg}}H^*(\mathcal{K}_v) & &
 \end{array}$$

as required. \square

7 Tropical completion

Given a point p in the tropical part \underline{B} of a compact exploded manifold B , the set of points $B|_p \subset B$ with tropical part p is a possibly noncompact manifold. The tropical completion $\check{B}|_p$ of B at p is a canonical way of completing $B|_p$ to a complete exploded manifold. It should be thought of as a minimal way of adding “structure at infinity” to $B|_p$ in a way determined by B . In this section, we show that tropical completion is compatible with our construction of virtual class $[\mathcal{K}]$, as well as the associated integration and pushforward of forms from \mathcal{K} . This is used in the proof of the tropical gluing formula, equation (1) of [30], where the contribution of a tropical curve γ is defined using tropical completion.

Definition 7.1 (tropical completion in a coordinate chart) The tropical completion of a polytope P in \mathbb{R}^m at a point $p \in P$ is a polytope $\check{P}_p \subset \mathbb{R}^m$ which is the union of all rays in \mathbb{R}^m beginning at p and intersecting P in more than one point.

Given an open subset $U \subset \mathbb{R}^n \times T_P^m$ and a point $p \in \underline{U}$, let \bar{U}_p indicate the closure within U of all points with tropical part p . Note that \bar{U}_p is naturally contained in both $\mathbb{R}^n \times T_P^m$ and $\mathbb{R}^n \times T_{\check{P}_p}^m$. Define the tropical completion of the coordinate chart U at $p \in \underline{U}$,

$$U|_p \subset \mathbb{R}^n \times T_{\check{P}_p}^m,$$

to be the smallest open subset of $\mathbb{R}^n \times T_{\check{P}_p}^m$ that contains \bar{U}_p .

Similarly, given a countable collection of points $X \subset \underline{U}$, define

$$U|_X := \coprod_{p \in X} U|_p.$$

Tropical completion in coordinate charts is functorial: given a map $\underline{f}: P \rightarrow Q$ and point $p \in P$, there is a unique map $\underline{f}|_p: \check{P}_p \rightarrow \check{Q}_{f(p)}$ that is equal to \underline{f} on the

restriction to $P \subset \check{P}_p$, and similarly, given a smooth or $C^{\infty,1}$ map of coordinate charts

$$f: U \rightarrow U'$$

and a point p in \underline{U} , there is a unique map

$$f|_p: U|_p \rightarrow U'|_{f(p)}$$

such that $f|_p$ is equal to f when restricted to the inverse image of p in U . Of course, the tropical part of $f|_p$ is equal to the map $\underline{f}|_p$ above. This map $f|_p$ is smooth or $C^{\infty,1}$ if f is, is an isomorphism onto an open subset of $U'|_{f(p)}$ if f is an isomorphism onto an open subset of U' , and is complete if f is proper. There is therefore a functorial construction of the tropical completion $\mathbf{B}|_p$ of any exploded manifold \mathbf{B} at a point $p \in \underline{\mathbf{B}}$ defined by applying tropical completion to coordinate charts.

In particular, let $O \rightarrow \mathbf{B}$ be a cover of \mathbf{B} by a collection of coordinate charts. We can recover \mathbf{B} by gluing together these coordinate charts using the transition maps, encapsulated in $O \times_{\mathbf{B}} O \rightrightarrows O$ — in other words, \mathbf{B} is a pushout as in the following diagram:

$$\begin{array}{ccc} \mathbf{B} & \longleftarrow & O \\ \uparrow & & \uparrow t \\ O & \xleftarrow{s} & O \times_{\mathbf{B}} O \end{array}$$

For such a pushout to exist as an exploded manifold, we require that O and $O \times_{\mathbf{B}} O$ be exploded manifolds, s and t be étale (or local isomorphisms) and that $(s, t): O \times_{\mathbf{B}} O \rightarrow O \times O$ be injective and proper. Such properties are preserved by tropical completion.

Definition 7.2 (tropical completion of an exploded manifold) Given $p \in \underline{\mathbf{B}}$, the tropical completion of \mathbf{B} at p is defined as follows. Choose an open cover of \mathbf{B} by coordinate charts $\pi: O \rightarrow \mathbf{B}$ with transition maps encoded by $s, t: O \times_{\mathbf{B}} O \rightrightarrows O$. Then $\mathbf{B}|_p$ is the pushout created using the tropical completion of s and t at the inverse image of p — so $\mathbf{B}|_p$ has charts and transition maps the tropical completion of those from \mathbf{B} :

$$\begin{array}{ccc} \mathbf{B}|_p & \longleftarrow & O|_{\pi^{-1}p} \\ \uparrow & & \uparrow \\ O|_{\pi^{-1}p} & \xleftarrow{\quad} & (O \times_{\mathbf{B}} O)|_{(\pi \circ s)^{-1}p} \end{array}$$

If $X \subset \mathbf{B}$ is a countable set of points in the tropical part of \mathbf{B} , the tropical completion of \mathbf{B} at X may be defined as above, with X replacing p . The resulting exploded manifold $\mathbf{B}|_X$ is the disjoint union of $\mathbf{B}|_p$ for all $p \in X$.

Clearly, $\mathbf{B}|_p$ is well defined and independent of the particular collection of coordinate charts chosen, because applying tropical completion to extra coordinate charts will just give compatible coordinate charts on $\mathbf{B}|_p$.

If $\mathbf{B}|_p$ indicates the subset of \mathbf{B} with tropical part p , note that $\mathbf{B}|_p$ always contains $\mathbf{B}|_p$ as a dense subset. If \mathbf{B} is basic, then $\mathbf{B}|_p$ also contains a copy of the closure of $\mathbf{B}|_p$ in \mathbf{B} . (In the case that \mathbf{B} is not basic, a single point in the closure of $\mathbf{B}|_p$ may correspond to multiple points in $\mathbf{B}|_p$.)

The construction of tropical completions is functorial in the sense that given a map

$$f: \mathbf{A} \rightarrow \mathbf{B}$$

and a point $p \in \mathbf{A}$, there exists a unique map

$$f|_p: \mathbf{A}|_p \rightarrow \mathbf{B}|_{f(p)}$$

restricting to be f on $\mathbf{A}|_p \subset \mathbf{A}|_p$.

Note that $\mathbf{B}|_p$ is complete if \mathbf{B} is compact, and $f|_p$ is complete if f is proper.

Because \mathbb{R} is also an exploded manifold (with tropical part a single point, and tropical completion at this point still \mathbb{R}), the tropical completion of any \mathbb{R} -valued function on an exploded manifold is still a \mathbb{R} -valued function. Moreover, given any tensor θ on \mathbf{B} , such as an almost complex structure, metric, or differential form, there is a unique tensor $\theta|_p$ on $\mathbf{B}|_p$ restricting to be θ on $\mathbf{B}|_p$. If θ is a refined differential form on \mathbf{B} , the restriction of $\theta|_p$ to $\mathbf{B}|_p$ no longer uniquely specifies $\theta|_p$; in this case $\theta|_p$ is the unique refined differential form on $\mathbf{B}|_p$ such that if θ pulls back to an honest form θ' under $r: U' \rightarrow \mathbf{B}$, then $\theta|_p$ pulls back to $\theta'|_{p'}$ under $r|_{p'}: U'|_{p'} \rightarrow \mathbf{B}|_p$, where $r(p') = p$. Be warned that $\theta|_p$ may not be in $\Omega^*\mathbf{B}|_p$, even if $\theta \in \Omega^*\mathbf{B}$, because $\theta|_p$ may not vanish on the image of all $T^1_{(0,\infty)}$. If, however, $\theta \in {}^r_{\text{fg}}\Omega^*\mathbf{B}$, then $\theta|_p$ is in ${}^r_{\text{fg}}\Omega^*\mathbf{B}|_p$, so differential forms generated by functions are more compatible with tropical completion. In fact, Remark 5.5 implies that $\theta \in {}^r\Omega^*(\mathbf{B})$ is in ${}^r_{\text{fg}}\Omega^*(\mathbf{B})$ if and only if $\theta|_p \in {}^r\Omega^*(\mathbf{B})$ for all $p \in \mathbf{B}$.

Lemma 7.3 *Given any submersion $f: A \rightarrow B$ and form $\theta \in {}^r_{\text{fg}}\Omega_c^*(A)$, the push-forward of θ is in ${}^r_{\text{fg}}\Omega_c^*B$, and for any $p \in B$,*

$$(f_!\theta)|_p = \sum_{p' \in \underline{f}^{-1}p} (f|_{p'})_!\theta|_{p'}.$$

Proof As specified by Theorem 9.2 of [31], $f_!\theta$ is uniquely determined by the property that

$$\int_A f^*\alpha \wedge \theta = \int_B \alpha \wedge f_!\theta$$

for all $\alpha \in {}^r\Omega^*(B)$. For any such α supported only in $B|_p$, the pullback of α is supported only over the inverse image of p , so

$$\int_{B|_p} \alpha|_p \wedge (f_!\theta)|_p = \sum_{p' \in \underline{f}^{-1}p} \int_{A|_{p'}} (f|_{p'})^*\alpha|_p \wedge \theta|_{p'} = \int_{B|_p} \alpha|_p \wedge \sum_{p' \in \underline{f}^{-1}p} (f|_{p'})_!\theta|_{p'}.$$

As any form on the smooth manifold $B|_p$ is determined by its integral against compactly supported forms, and the restriction of $f_!\theta$ to $B|_p$ uniquely determines $(f_!\theta)|_p$, it follows that

$$(f_!\theta)|_p = \sum_{p' \in \underline{f}^{-1}p} (f|_{p'})_!\theta|_{p'},$$

as required. This equation then implies that $(f_!\theta)|_p \in {}^r\Omega^*(B|_p)$ for all $p \in B$, so Remark 5.5 implies that $f_!\theta \in {}^r_{\text{fg}}\Omega^*(B)$. \square

We can define the tropical completion of an orbifold \mathcal{X} similarly, by applying tropical completion to an atlas. In particular, any surjective étale map from a (not necessarily connected) exploded manifold U to \mathcal{X} defines an étale proper¹⁶ groupoid \mathcal{X}' with objects parametrized by U , and morphisms parametrized by $U \times_{\mathcal{X}} U$. We can recover \mathcal{X} from this étale proper groupoid \mathcal{X}' by taking the stack of principal \mathcal{X}' -bundles. Taking tropical completion preserves the structure equations of a groupoid and the property of maps being étale or proper. So, we can apply tropical completion to $U \times_{\mathcal{X}} U \rightrightarrows U$ to obtain a new étale proper groupoid representing the tropical completion of our stack.

¹⁶An étale proper groupoid in the category of exploded manifolds is a groupoid with objects and morphisms parametrized by exploded manifolds \mathcal{X}_0 and \mathcal{X}_1 , all structure maps morphisms in the category of exploded manifolds, source and target maps $s, t: \mathcal{X}_1 \rightarrow \mathcal{X}_0$ étale (ie locally isomorphisms), and the map $(s, t): \mathcal{X}_1 \rightarrow \mathcal{X}_0^2$ proper. The correct generalization of “proper” for exploded manifolds is usually “complete”; however, if s and t are étale, then (s, t) being proper is equivalent to it being complete.

To define tropical completion of an orbifold, we need a notion of the tropical part of an orbifold. The following defines the tropical part of a stack as a set. The tropical part of a stack obviously has a bit more structure, but we shall not need it.

Definition 7.4 As a set, the tropical part $\underline{\mathcal{X}}$ of a stack \mathcal{X} over the category of exploded manifolds is the set of path-connected components of \mathcal{X} . Say that two given points in \mathcal{X} are connected by a path if there is a map of \mathbb{R} into \mathcal{X} that restricts to $\{0, 1\}$ to be (isomorphic to) the two given points.

Define the tropical part $\underline{\mathcal{K}}$ of a K -category \mathcal{K} to be equal to the tropical part of \mathcal{K}^{st} .

The following definition formalizes the idea that tropical completion of an orbifold \mathcal{X} is achieved by applying tropical completion to coordinate charts on \mathcal{X} .

Definition 7.5 (tropical completion of an orbifold) Given an exploded orbifold \mathcal{X} , and a countable subset $A \subset \underline{\mathcal{X}}$, the tropical completion $\mathcal{X}|_A$ of \mathcal{X} at A is defined as follows: Choose an étale surjection $\underline{U} \rightarrow \mathcal{X}$, and abuse notation a little to denote by A the inverse image of A in \underline{U} and within $\underline{U} \times_{\mathcal{X}} \underline{U}$. Apply tropical completion at A to obtain an étale proper groupoid with objects parametrized by $\underline{U}|_A$, morphisms parametrized by $(\underline{U} \times_{\mathcal{X}} \underline{U})|_A$, and groupoid structure maps the tropical completion at A of the structure maps of the original groupoid representing \mathcal{X} . Then $\mathcal{X}|_A$ is the orbifold represented by this groupoid.

Again $\mathcal{X}|_A$ contains $\mathcal{X}|_A$ as a dense suborbifold, and any map $\mathcal{X} \rightarrow \mathcal{Y}$ sending A to A' induces a canonical map $\mathcal{X}|_A \rightarrow \mathcal{Y}|_{A'}$ that restricts to the original map on $\mathcal{X}|_A$. Any (possibly refined) differential form θ on \mathcal{X} defines a differential form $\theta|_A$ on $\mathcal{X}|_A$.

To define the tropical completion of a K -category, we again use tropical completion of charts. Here we run into a problem: \mathcal{K}^{st} is not determined by these charts, just as the smooth structure on the union of an open halfspace with a perpendicular line is not determined by transition data. The problem is that there exist families in \mathcal{K}^{st} such that some points have no neighborhood sent inside a chart. This technical issue is resolved using extensions of our charts, as families in \mathcal{K}^{st} are always locally contained in some extended chart. The following describes the structure we need:

Remark 7.6 The following information is sufficient to define an extendable Kuranishi category:

- a groupoid in the category of $C^{\infty,1}$ exploded manifolds, and with objects parametrized by a countable disjoint union of exploded manifolds $\bigsqcup_i F_i^\sharp$, where each F_i has fixed dimension, and
- open subsets $F_i \subset F'_i \subset F_i^\sharp$,

satisfying the following conditions:

- (i) The morphisms from F_i^\sharp to itself are given by the action of a finite group G_i . So,

$$F_{ii}^\sharp = F_i^\sharp \times_{F_i^\sharp/G_i} F_i^\sharp,$$

where morphisms with source and target F_i^\sharp are parametrized by F_{ii}^\sharp . We require the subsets F_i and F'_i to be G_i -equivariant open subsets of F_i^\sharp .

- (ii) The morphisms with source F_i^\sharp and target F_j^\sharp are parametrized by an exploded manifold F_{ij}^\sharp . For each i , we require that $F_{ij}^\sharp = \emptyset$ for all but finitely many j . The source and target maps,

$$F_i^\sharp \xleftarrow{\phi_{ij}} F_{ij}^\sharp \xrightarrow{\phi_{ji}} F_j^\sharp,$$

must satisfy the following conditions:

- (a) If \bar{F}'_j indicates the closure of F'_j within F_j^\sharp , then $\phi_{ij}(\phi_{ji}^{-1}(\bar{F}'_j)) \subset F_i^\sharp$ is a closed subset of F_i^\sharp . Moreover, F'_i contains the closure of $F_i \subset F_i^\sharp$.
- (b) If $\dim F_i \leq \dim F_j$, then ϕ_{ij} is a G_j -fold cover of an open subset of F_i , and ϕ_{ji} a G_i -fold cover of a exploded submanifold of F_j (locally defined by the transverse vanishing of $C^{\infty,1}$ functions).

Given a K -category \mathcal{K} with extensions $\mathcal{K} \subseteq_e \mathcal{K}' \subseteq_e \mathcal{K}^\sharp$, we may obtain the above data by setting $F_i^\sharp = F(\hat{f}_i^\sharp)$, $F'_i = F(\hat{f}'_i)$, $F_i = F(\hat{f}_i)$ and $F_{ij}^\sharp = F(\hat{f}_i^\sharp \times_{(\mathcal{K}^\sharp)^{\text{st}}} \hat{f}_j^\sharp)$. The groupoid in question is the full subcategory of $(\mathcal{K}^\sharp)^{\text{st}}$ with objects the individual curves in these \hat{f}_i^\sharp .

Given the above data, we define a stack \mathcal{K}^{st} as follows: A family \hat{f} in \mathcal{K}^{st} parametrized by $F(\hat{f})$ is

- (i) a collection of G_i -fold covers $X_i(\hat{f}) \subset X_i^\sharp(\hat{f}) \rightarrow F(\hat{f})$ of subsets of $F(\hat{f})$ such that every point in $F(\hat{f})$ is in the image of some $X_i(\hat{f})$, and in the interior of the image of some $X_i^\sharp(\hat{f})$;¹⁷

¹⁷ $X_i(\hat{f})$ is the fiber product of $F(\hat{f})$ with $F(\hat{f}_i) = F_i$ over \mathcal{K}^{st} .

(ii) G_i -equivariant maps

$$X_i^\#(\hat{f}) \rightarrow \mathbf{F}_i^\#$$

that are $C^{\infty,1}$ on the interior of $X_i^\#(\hat{f})$ and such that the inverse image of $\mathbf{F}_i \subset \mathbf{F}_i^\#$ is $X_i(\hat{f}) \subset X_i^\#(\hat{f})$;

(iii) maps $X_i^\#(\hat{f}) \times_{\mathbf{F}(\hat{f})} X_j^\#(\hat{f}) \rightarrow \mathbf{F}_{ij}^\#$ such that both squares below are fiber-product diagrams (in the category of sets):

$$\begin{array}{ccc} X_i^\#(\hat{f}) & \longrightarrow & \mathbf{F}_i^\# \\ \uparrow & & \uparrow \\ X_i^\#(\hat{f}) \times_{\mathbf{F}(\hat{f})} X_j^\#(\hat{f}) & \longrightarrow & \mathbf{F}_{ij}^\# \\ \downarrow & & \downarrow \\ X_j^\#(\hat{f}) & \longrightarrow & \mathbf{F}_j^\# \end{array}$$

and such that these maps define a map of groupoids (in the category of sets)

$$\begin{array}{ccc} \coprod_{i,j} X_i^\#(\hat{f}) \times_{\mathbf{F}(\hat{f})} X_j^\#(\hat{f}) & \longrightarrow & \coprod_{i,j} \mathbf{F}_{ij}^\# \\ \Downarrow & & \Downarrow \\ \coprod_i X_i^\#(\hat{f}) & \longrightarrow & \coprod_i \mathbf{F}_i^\# \end{array}$$

A morphism $\hat{f} \rightarrow \hat{g}$ in \mathcal{K}^{st} is

- a $C^{\infty,1}$ map $\mathbf{F}(\hat{f}) \rightarrow \mathbf{F}(\hat{g})$, and
- G_i -equivariant maps $X_i^\#(\hat{f}) \rightarrow X_i^\#(\hat{g})$ such that the following diagrams commute:

$$\begin{array}{ccc} X_i^\#(\hat{f}) \rightarrow X_i^\#(\hat{g}) & X_i^\#(\hat{f}) \times_{\mathbf{F}(\hat{f})} X_j^\#(\hat{f}) \rightarrow X_i^\#(\hat{g}) \times_{\mathbf{F}(\hat{f})} X_j^\#(\hat{g}) \\ \searrow \downarrow & \searrow \downarrow \\ & \mathbf{F}_i^\# \qquad \qquad \qquad \mathbf{F}_{ij}^\# \end{array}$$

and such that the following is a fiber-product diagram:

$$\begin{array}{ccc} X_i^\#(\hat{f}) & \longrightarrow & X_i^\#(\hat{g}) \\ \downarrow & & \downarrow \\ \mathbf{F}(\hat{f}) & \longrightarrow & \mathbf{F}(\hat{g}) \end{array}$$

As defined above, \mathcal{K}^{st} is a stack. Moreover, \mathcal{K}^{st} is equivalent to the original \mathcal{K}^{st} in the case that \mathcal{K} was defined as in Definition 2.1. Of course, when \mathcal{K} is a subcategory of a nice moduli stack of curves, it is better to think of \mathcal{K}^{st} as a stack of curves rather than use the above equivalent but ad hoc stack.

To see that \mathcal{K}^{st} is a stack, let us verify Definition 4.3 of [7]. Note that \mathcal{K}^{st} has essentially unique pullbacks, so is a category fibered in groupoids over the category of $C^{\infty,1}$ exploded manifolds. Isomorphisms in \mathcal{K}^{st} form a sheaf because morphisms are defined as maps satisfying local conditions. Each descent datum is effective because a descent datum is sufficient to construct the G_i -fold covers $X_i^{\#}$, and all other information is locally determined, so further choices form a sheaf.

Example 7.7 \mathcal{K}^{st} is an orbifold in the case that each F_i has the same dimension. Then the data above may be thought of as an atlas for \mathcal{K}^{st} , or an étale proper Lie groupoid representing \mathcal{K}^{st} . Condition (b) gives étale, and condition (a) implies properness. Conversely, given any orbifold, we may obtain the above data from a choice of three nested atlases, with the closure of the charts from a given atlas contained in the charts from the next atlas.

Definition 7.8 (tropical completion of a K -category) Given an extendable K -category \mathcal{K} (Definition 2.1) and a countable subset $A \subset \underline{K}$, define the tropical completion $\check{\mathcal{K}}|_A$ of \mathcal{K} at A as follows:

Choose extensions $\mathcal{K}^{\#} \supset_e \mathcal{K}' \supset_e \mathcal{K}$. Define $F_i := F(\hat{f}_i)$, $F'_i := F(\hat{f}'_i)$, $F_i^{\#} := F(\hat{f}_i^{\#})$ and define $F_{ij}^{\#} := F(\hat{f}_i^{\#} \times_{(\mathcal{K}^{\#})^{\text{st}}} \hat{f}_j^{\#})$ to obtain the data from Remark 7.6. Let $\check{\mathcal{K}}|_A$ be the K -category with the data

$$F_i|_A \subset F'_i|_A \subset F_i^{\#}|_A, \quad F_i^{\#}|_A \leftarrow F_{ij}^{\#}|_A \rightarrow F_j^{\#}|_A$$

along with groupoid structure maps corresponding to the tropical completion of the groupoid structures maps from \mathcal{K} . As noted in Remark 7.6, this data is sufficient to define an extendable K -category $\check{\mathcal{K}}|_A$.

Any map $\psi: \mathcal{K} \rightarrow \mathcal{X}$ from \mathcal{K} to an exploded orbifold or manifold \mathcal{X} corresponds to a map $\psi|_A: \check{\mathcal{K}}|_A \rightarrow \mathcal{X}|_{\psi(A)}$. Any differential form θ on \mathcal{K} corresponds to a differential form $\theta|_A$ on $\check{\mathcal{K}}|_A$, any open substack \mathcal{U} of \mathcal{K}^{st} corresponds to an open substack $\mathcal{U}|_A$ of $(\check{\mathcal{K}}|_A)^{\text{st}}$, any vectorbundle V on \mathcal{U} corresponds to a vectorbundle $V|_A$ on $\mathcal{U}|_A$, and any section s of V corresponds to a section $s|_A$ of $V|_A$.

Definition 7.9 (tropical completion of a Kuranishi category) The tropical completion of a Kuranishi category (Definition 3.2) \mathcal{K} at a countable subset $A \subset \underline{\mathcal{K}}$ is the K -category $\mathcal{K}|_A$ along with the vectorbundles $V_i|_A$ on the open substacks $\mathcal{U}_i|_A$ and sections $\bar{\partial}\hat{f}_i|_A: \mathbf{F}(\hat{f}_i|_A) \rightarrow V_i|_A(\hat{f}_i|_A)$.

Note that if $\pi: \mathcal{K} \rightarrow \mathbf{Z}$ is proper (Definition 3.5), then $\pi|_p: \mathcal{K}|_p \rightarrow \mathbf{Z}|_{\pi(p)}$ is complete.

Lemma 7.10 Given any proper, oriented Kuranishi category \mathcal{K} and closed form $\theta \in {}^r_{\text{fg}}\Omega^*\mathcal{K}$,

$$\int_{[\mathcal{K}]} \theta = \sum_{p \in \underline{\mathcal{K}}} \int_{[\mathcal{K}|_p]} \theta|_p.$$

Similarly, if \mathcal{K} is complete and oriented over \mathbf{Z} , and $\pi: \mathcal{K} \rightarrow \mathbf{A}$ is a compatible map, then given any closed $\theta \in {}^r_{\text{fg}}\Omega^*\mathcal{K}$ and $p' \in \underline{\mathbf{A}}$,

$$\pi_!(\theta)|_{p'} = \sum_{p \in \pi^{-1}(p')} (\pi|_p)_!(\theta|_p),$$

where the above equation holds exactly for some choices of construction of $\pi_!$ and $(\pi|_p)_!$, and otherwise holds on the level of cohomology.

Proof The construction of $[\mathcal{K}]$ is compatible with tropical completion. The tropical completions of the functions ρ_i used to define \mathcal{K}_ϵ in Definition 3.9 are appropriate for defining $(\mathcal{K}|_p)_\epsilon$. Tropical completion applied to the metric from Lemma 3.11 also gives an appropriate metric on V over $\mathcal{K}|_p$. With these choices, the sheaf S from Definition 3.12 is compatible with tropical completion — if v is a section of $S(\hat{f})$, then $v|_p$ is a section of $S(\hat{f}|_p)$.

Any weighted branched cover I of \mathcal{K} gives a weighted branched cover $I|_p$ of $\mathcal{K}|_p$ as follows: A connected family O in \mathcal{K}^{st} is in $\mathcal{O}_{I|_p}$ if there exists a family $O' \in \mathcal{K} \cap \mathcal{O}_I$ and a map $O \rightarrow O'|_p$. Note that given any two such families O'_i , a connected component of $O'_1 \times_{\mathcal{K}^{\text{st}}} O'_2$ satisfies the same conditions because tropical completion commutes with fiber products. We can therefore define $I|_p(O)$ as the inverse limit of $I(O')$ over the category of all such maps $O \rightarrow O'|_p$. This $I|_p$ is a functor because any map $O_1 \rightarrow O_2$ induces a corresponding map of inverse limits — simply compose $O_2 \rightarrow O'|_p$ with $O_1 \rightarrow O_2$. Because $I(O')$ only maps to a finite number of other finite probability spaces, inverse limits are easy: there exists some O' with a map $O \rightarrow O'|_p$ such that $I(O') = I|_p(O)$. This $I|_p$ satisfies the requirements of Definition 3.19. Moreover, $I|_p$ is separated if I is, because, given any automorphism ψ of a family O in $\mathcal{K}|_p$,

there exists a family O' in \mathcal{K} with an automorphism ψ' such that O is a connected component of $O'|_p$ and $\psi = \psi'|_p$.

Any global section v of S^I over \mathcal{K}_ϵ corresponds to a global section $v|_p$ of $S^{I|_p}$ over $(\mathcal{K}|_p)_\epsilon$. In particular, given any $O \in \mathcal{O}_{I|_p} \cap (\mathcal{K}|_p)_\epsilon$, there exists a family $O' \in \mathcal{O}_I \cap \mathcal{K}_\epsilon$ with a map $O \rightarrow O'|_p$ such that $I|_p(O) = I(O')$. Then, for any $i \in I|_p(O)$, we have a section $v(i)$ of $S(O')$. The tropical completion of this section is a section $v(i)|_p$ of $S(O'|_p)$, which pulls back under $O \rightarrow O'|_p$ to define $v|_p(i)$ as a section of $S(O)$. This defines the section $v|_p$ of $S^{I|_p}$ over O .

Let us check that $v|_p$ as defined above gives a well-defined global section of $S^{I|_p}$. Given any map $O_2 \rightarrow O$ and map $O_2 \rightarrow O'_2|_p$ with $I|_p(O_2) = I(O'_2)$, there exists some O''_2 with maps such that the following diagram commutes:

$$\begin{array}{ccccc} O_2 & \xrightarrow{\quad} & O & & \\ & \searrow & \downarrow & & \\ O'_2|_p & \xleftarrow{\quad} & O''_2|_p & \xrightarrow{\quad} & O'|_p \end{array}$$

It follows that the section $v|_p$ of $S^{I|_p}(O_2)$ pulled back from O'_2 coincides with the section pulled back from O''_2 and the section pulled back from O' via O . Therefore, $v|_p$ as defined above is a well-defined global section of $S^{I|_p}$.

The resulting $[\mathcal{K}|_p]$ is similarly related to $[\mathcal{K}]$. In particular, for any $O \in \mathcal{O}_{I|_p} \cap (\mathcal{K}|_p)_\epsilon$, we can choose an O' with a map $O \rightarrow O'|_p$ such that $I|_p(O) = I(O')$. Then $[\mathcal{K}|_p](i)$ is the subfamily of O that is the pullback of $[\mathcal{K}](i)|_p \subset O'|_p$ under the map $O \rightarrow O'|_p$.

Claim 7.11 *If \mathcal{K} is proper and oriented, and $\theta \in {}^r_{\text{fg}}\Omega^*(\mathcal{K})$ is any (not necessarily closed) form, then*

$$(7) \quad \int_{[\mathcal{K}]} \theta = \sum_{p \in \underline{\mathcal{K}}} \int_{[\mathcal{K}|_p]} \theta|_p$$

when the above method is used to construct the virtual fundamental class $[\mathcal{K}|_p]$ of $\mathcal{K}|_p$ from $[\mathcal{K}]$.

Choose a partition of unity $r: O \rightarrow \mathbb{R}$ compatible with $[\mathcal{K}]$. Then

$$\int_{[\mathcal{K}]} \theta = \sum_{O_k \subset O} \sum_{i \in I(O_k)} \mu(i) \int_{[\mathcal{K}](i)} r\theta,$$

where the sum is over connected components O_k of O . We may apply tropical completion to r and obtain a partition of unity $r|_p: O|_p \rightarrow \mathbb{R}$ compatible with $\mathcal{K}|_p$, where we use p in \underline{O} to mean the inverse image of p under the map $\underline{O} \rightarrow \underline{\mathcal{K}}$.

For any exploded manifold \mathbf{B} , the integral of θ over \mathbf{B} is the sum of the integrals of θ over the manifolds $\mathbf{B}|_p$ (if θ has compact support, only a finite number of these integrals will be nonzero). Furthermore, the integral of θ over $\mathbf{B}|_p$ is equal to the integral of $\theta|_p$ over $\mathbf{B}|_p$. Applying this to our situation, we get

$$\int_{[\mathcal{K}](i)} r\theta = \sum_{p \in \underline{\mathcal{K}}} \int_{[\mathcal{K}](i)|_p} r|_p \theta|_p.$$

The tropical completion of O_k may not be connected. Each connected component O' of $O_k|_p$ is in $\mathcal{O}_{I|_p}$, and there is a corresponding map $\iota^*: I(O_k) \rightarrow I|_p(O')$ such that $[\mathcal{K}|_p](\iota^*i)$ is the corresponding collection of connected components of $[\mathcal{K}](i)|_p$. Therefore,

$$\int_{[\mathcal{K}](i)|_p} r|_p \theta|_p = \sum_{\iota} \int_{[\mathcal{K}|_p](\iota^*i)} r|_p \theta|_p,$$

where the sum is over the different inclusions ι of connected components of $O_k|_p$. Noting that each ι^* is surjective and measure-preserving then gives

$$\begin{aligned} \sum_{O_k \subset O} \sum_{i \in I(O_k)} \mu(i) \int_{[\mathcal{K}](i)|_p} r|_p \theta|_p &= \sum_{O' \subset O|_p} \sum_{i' \in I(O')} \mu(i') \int_{[\mathcal{K}|_p](i')} r|_p \theta|_p \\ &= \int_{[\mathcal{K}|_p]} \theta|_p, \end{aligned}$$

which implies Claim 7.11.

In the case that θ is closed, (7) holds regardless of any choices made in the construction of $[\mathcal{K}]$ and $[\mathcal{K}|_p]$ because these integrals do not depend on such choices. It remains to prove the analogous statement for pushforwards:

Claim 7.12 *If \mathcal{K} is complete and oriented over \mathbf{Z} , and $\pi: \mathcal{K} \rightarrow \mathbf{A}$ is a compatible map, there exists a construction of $\pi_!$ and $(\pi|_p)_!$ such that given any $\theta \in {}^r_{\text{fg}}\Omega^*\mathcal{K}$ and $p' \in \underline{\mathbf{A}}$,*

$$\pi_!(\theta)|_{p'} = \sum_{p \in \pi^{-1}(p')} (\pi|_p)_!(\theta|_p).$$

The proof of Claim 7.12 is very similar to the proof of Claim 7.11 except we must now check that all choices in the construction of $\pi_!$ from Definition 5.15 are compatible

with tropical completion. This is indeed the case: the choices of vectorbundle $W \rightarrow A$, map $x: W \rightarrow A$ and Thom form e on W all may be tropically completed at p' to give $W|_p \rightarrow A|_p$, $x|_p: W|_p \rightarrow A|_p$ and $e|_p$ appropriate for defining $[\pi|_p]_!$ for all $p \in \pi^{-1}p'$.

Definition 5.15 gives that

$$\pi_!(\theta) = \sum_{O_k \subset O} \sum_{i \in I(O_k)} \mu(i) \hat{\pi}(i)_!(r\theta \wedge e),$$

where $\hat{\pi}(i)$ is defined by the composition

$$\begin{array}{ccccc} & & \hat{\pi}(i) & & \\ & \searrow & \text{---} & \nearrow & \\ \pi(i)^*W & \longrightarrow & W & \xrightarrow{x} & A \\ \downarrow & & \downarrow & & \\ [\mathcal{K}](i) & \xrightarrow{\pi(i)} & A & & \end{array}$$

Therefore,

$$\pi_!(\theta)|_{p'} = \sum_{O_k \subset O} \sum_{p \in \pi^{-1}(p')} \sum_{i \in I(O_k)} \mu(i) \hat{\pi}(i)|_p (r|_p \theta|_p \wedge e|_p).$$

As in the proof of Claim 7.11, $O_k|_p$ may have several connected components, $\iota: O' \rightarrow O_k|_p$, and $[\mathcal{K}](i)|_p = \coprod_l [\mathcal{K}|_p](\iota^*i)$, and we similarly have $\hat{\pi}(i)|_p = \coprod_l \hat{\pi}|_p(\iota^*i)$. The same argument as the proof of Claim 7.11 then gives the required result:

$$\pi_!(\theta)|_{p'} = \sum_{p \in \pi^{-1}(p')} (\pi|_p)_!(\theta|_p)$$

In the case that θ is closed, both sides of the above equation are independent of all choices; therefore, the above equation holds regardless of choices. \square

8 Construction of Gromov–Witten invariants

Let us summarize our construction of Gromov–Witten invariants. Start with a complete, basic exploded manifold, \mathbf{B} , with a $\bar{\partial}$ -log compatible almost complex structure, J , tamed by a symplectic form, ω , as described in Definitions 3.1 and 3.5 of [28] and Definitions 3.15 and 4.6 of [24]. Also assume that the tropical part of \mathbf{B} admits a \mathbb{Z} -affine immersion into some $[0, \infty)^N$, so that Lemma 4.2 and Theorem 6.1 of [28] imply that the moduli stack of J -holomorphic curves is compact when restricted to

curves with bounded genus, energy and number of punctures or ends.¹⁸ Some examples of (\mathbf{B}, ω, J) satisfying these conditions include

- any compact symplectic manifold with a tamed almost complex structure J ,
- the explosion¹⁹ of any compact Kähler manifold relative to simple normal crossing divisors, and
- an exploded manifold corresponding to the singular fiber in any simple normal crossing degeneration of a compact Kähler manifold.

Let $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$ be the moduli stack of stable $C^{\infty,1}$ curves in \mathbf{B} with genus g , n labeled ends, and representing a homology class, $\beta: H^*(\mathbf{B}) \rightarrow \mathbb{R}$. As the moduli stack of holomorphic curves within $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$ is compact, Theorem 7.3 of [26] provides a complete embedded Kuranishi structure covering the moduli stack of holomorphic curves within $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$. Next, construct an orienting 2-form α on the associated Kuranishi category, using Definition 3.1, then restrict the associated Kuranishi category to a neighborhood of the holomorphic curves, so that α is orienting. Call this oriented Kuranishi category $\mathcal{K}_{g,n,\beta}$. As $\mathcal{K}_{g,n,\beta}$ is complete, we can construct a virtual fundamental class, $[\mathcal{K}_{g,n,\beta}]$, for $\mathcal{K}_{g,n,\beta}$ using Definition 4.7.

There are several distinct versions of “evaluation maps” from the moduli stack of curves in exploded manifolds. Let

$$\text{ev}: \mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B}) \rightarrow X$$

be any $C^{\infty,1}$ map to an exploded manifold or orbifold X . For example, if \mathbf{B} is a smooth manifold or an exploded manifold with bounded tropical part (such as any exploded manifold in a connected family also containing a smooth manifold), then each end of a curve is sent to a single point in \mathbf{B} , so we could take ev as the usual evaluation map at ends, with $X = \mathbf{B}^n$. If \mathbf{B} has unbounded tropical part, $\underline{\mathbf{B}}$, there

¹⁸An end of a holomorphic curve is a stratum of its domain, C , isomorphic to $T_{(0,\infty)}^1$. The smooth part $[C]$ of C is a nodal curve with a marked point corresponding to each end of C . We can also achieve compactness and construct Gromov–Witten invariants with weaker assumptions when we include more discrete data about the curves at ends. For example, holomorphic curves in T^n have tropical parts that are tropical curves in \mathbb{R}^n . Each end of such a curve corresponds to an end of the corresponding tropical curve, which has some derivative α in \mathbb{Z}^n . To achieve compactness for the moduli stack of holomorphic curves in T^n , we have to keep track of these derivatives α at the ends of curves. See Lemma 4.2 of [28].

¹⁹See Section 5 of [24] for a discussion of the explosion functor, or see Section 4 of [25] for a discussion of the explosion functor as a base change within log geometry. The tropical part of the explosion of a manifold with a normal crossing divisor is the dual intersection complex with a canonical \mathbb{Z} -affine structure. To ensure that this tropical part immerses in some $[0, \infty)^N$, we assume that it is simple, or a union of transversely intersecting, compact, codimension-2, complex submanifolds.

are different possible “evaluation” maps which we could use. In this case, some ends of holomorphic curves have unbounded image in $\underline{\mathbf{B}}$, and evaluation at such punctures lands in a space 2 real dimensions smaller than \mathbf{B} . In such an exploded manifold, evaluation at an end lands in an associated exploded manifold, $\text{End } \mathbf{B}$, constructed in Section 3 of [30], so we can take X as $(\text{End } \mathbf{B})^n$.

We can then define Gromov–Witten invariants as a map

$${}^rH^*(X) \rightarrow \mathbb{R}$$

defined by

$$\theta \mapsto \int_{[\mathcal{K}_{g,n,\beta}]} \text{ev}^* \theta$$

using the integration from Definition 5.8 and the homology from Definition 5.3. We can encode finer information about $[\mathcal{K}_{g,n,\beta}]$ using Definition 5.15 to push forward this virtual fundamental class to define

$$\eta_{g,n,\beta} := \text{ev}_!(1) \in {}^r_{\text{fg}}H^*(X),$$

so

$$\int_{[\mathcal{K}_{g,n,\beta}]} \text{ev}^* \theta = \int_X \theta \wedge \eta_{g,n,\beta}.$$

If $2g + n \geq 3$, there is also a $C^{\infty,1}$ map from $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$ to Deligne–Mumford space, $\overline{\mathcal{M}}_{g,n}$, or its explosion, $\text{Expl } \overline{\mathcal{M}}_{g,n}$ — this stabilization map to $\text{Expl } \overline{\mathcal{M}}_{g,n}$ is constructed in Section 4 of [26], where it is also proved that $\text{Expl } \overline{\mathcal{M}}_{g,n}$ represents the moduli stack of stable exploded curves with genus g and n punctures. So, we could also define Gromov–Witten invariants with the X above being $(\text{End } \mathbf{B})^n \times \text{Expl } \overline{\mathcal{M}}_{g,n}$, or a more sophisticated version of this, $\mathcal{X}_{g,n}(\mathbf{B})$, described in Section 5 of [30]. In the case that \mathbf{B} is a symplectic manifold, Gromov–Witten invariants defined using $X := \mathbf{B}^n \times \overline{\mathcal{M}}_{g,n}$ satisfy Kontsevich and Manin’s axioms for Gromov–Witten invariants from [14]; as explained in [30], the splitting and genus-reduction axioms follow from Theorem 5.2 and Lemma 5.3 of [30].

We can also define descendant Gromov–Witten invariants incorporating Chern classes of tautological line bundles, L_i^* , over $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$. Let us describe these tautological line bundles. Each of our labeled ends corresponds to a $T_{(0,\infty)}^1$ –bundle, L_i , over $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$, with fiber over a curve f the stratum of the domain of f labeled by i ; so over a family, \widehat{f} , $L_i(\widehat{f}) \rightarrow \mathbf{F}(\widehat{f})$ is a union of strata of the domain, $\mathbf{C}(\widehat{f}) \rightarrow \mathbf{F}(\widehat{f})$, of \widehat{f} . The moduli stack of $T_{(0,\infty)}^1$ –bundles is equivalent to the moduli stack of \mathbb{C}^* –bundles or complex line bundles, so there is an associated complex line bundle L_i

on $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$ — the relationship between the fibers of L_i and \mathbf{L}_i is the relationship between \mathbb{C} and $T_{(0,\infty)}^1 \subset T_{[0,\infty)}^1 = \text{Expl}(\mathbb{C}, 0)$. We can also think of the fiber of L_i over f as the tangent space of the smooth part²⁰ of the domain of f at its i^{th} marked point: The smooth part of the domain of \hat{f} , $[\mathbf{C}(\hat{f})]$, is a family of nodal curves over $[\mathbf{F}(\hat{f})]$, and the i^{th} end corresponds to a marked point section $s_i: [\mathbf{F}(\hat{f})] \rightarrow [\mathbf{C}(\hat{f})]$. Our line bundle L_i is the pullback of the vertical tangent bundle of $[\mathbf{C}(\hat{f})]$ under the composition $\mathbf{F}(\hat{f}_i) \rightarrow [\mathbf{F}(\hat{f}_i)] \xrightarrow{s_i} [\mathbf{C}(\hat{f}_i)]$. So, the dual line bundles, L_i^* , are analogous to the usual tautological line bundles over the moduli stack of curves.

With the tautological line bundles L_i^* over $\mathcal{M}_{g,n,\beta}^{\text{st}}(\mathbf{B})$ understood, we can construct their Chern classes using Remark 5.2, and define $\psi_i \in {}^r H^*(\mathcal{K}_{g,n,\beta})$ as the first Chern class of the line bundle L_i^* over \mathcal{K} . We can then define the classes

$$(8) \quad \eta_{g,\{a_i\},\beta} := \text{ev}_! \left(\prod_i \psi_i^{a_i} \right) \subset {}^r H^*(X)$$

for nonnegative integers a_1, \dots, a_n , and define Gromov–Witten invariants as the corresponding maps

$$(9) \quad {}^r H^*(X) \xrightarrow{\theta \mapsto \int_X \eta_{g,\{a_i\},\beta} \wedge \theta} \mathbb{R}.$$

In the case that $\underline{\mathbf{B}}$ is bounded and $X = \mathbf{B}^n$, we can package Gromov–Witten invariants into correlators without losing important information. When $\underline{\mathbf{B}}$ is bounded, ${}^r H^*(\mathbf{B})$ is isomorphic to $H^*(\mathbf{B})$,²¹ so we can use H^* in place of ${}^r H^*$ without losing important information about Gromov–Witten invariants. Moreover, unlike ${}^r H^*$ and ${}^r H^*$, Künneth’s theorem applies to H^* , so $H^*(\mathbf{B}^n) = H^*(\mathbf{B})^{\otimes n}$. Given classes $\theta_1, \dots, \theta_n$ in $H^*(\mathbf{B})$, we can define the numerical Gromov–Witten invariants

$$(10) \quad \begin{aligned} \langle \tau_{a_1}(\theta_1), \dots, \tau_{a_n}(\theta_n) \rangle_{g,n,\beta} &:= \int_{\mathbf{B}^n} \eta_{g,\{a_i\},\beta} \wedge \bigwedge_i \pi_i^* \theta_i \\ &= \int_{[\mathcal{K}_{g,n,\beta}]} \bigwedge_i \psi_i^{a_i} \wedge (\pi_i \circ \text{ev})^* \theta_i, \end{aligned}$$

where $\pi_i: \mathbf{B}^n \rightarrow \mathbf{B}$ is projection onto the i^{th} component. These $\tau_{a_i}(\theta_i)$ could also be thought of as cohomology classes on \mathbf{B} times the stack of complex line bundles, where

²⁰Every exploded manifold \mathbf{B} comes with a natural map $\mathbf{B} \rightarrow [\mathbf{B}]$ to its smooth part; the smooth part of an exploded curve \mathbf{C} is a nodal curve $[\mathbf{C}]$; the smooth part of the explosion of a manifold with normal crossing divisors is the original manifold.

²¹There is no written proof of this fact, but it can be proved by representing each cohomology class in ${}^r H^*(\mathbf{B})$ by a closed form on a refinement of \mathbf{B} ; when $\underline{\mathbf{B}}$ is bounded and admits a \mathbb{Z} -affine immersion into \mathbb{R}^N , each refinement of \mathbf{B} is cobordant to \mathbf{B} , and hence has the same cohomology by [31, Section 11].

our correlator “integrates the pullback” of these classes to $[\mathcal{K}_{g,n,E}]$ using the evaluation map recording the position of the i^{th} puncture and the line bundle L_i^* ; however, our machinery for differential forms on stacks with infinite isotropy groups is inadequate for removing the above scare quotes.

Theorem 5.20 together with [26, Corollary 7.5] imply that $\eta_{g,\{a_i\},\beta}$ does not depend on the choices involved in its construction. Theorem 7.3 of [26] together with Theorem 5.22 provide a kind of invariance in families for the Gromov–Witten invariants (9) and (10).

Let us describe how the correlators (10) are invariant in families. Let $\hat{\mathbf{B}} \rightarrow \mathbf{B}_0$ be a connected family of exploded manifolds containing \mathbf{B} , with a family of $\bar{\partial}$ -log compatible almost complex structures tamed by a family of taming forms, and suppose that there is a \mathbb{Z} -affine map $\hat{\mathbf{B}} \rightarrow [0, \infty)^N$ that is injective when restricted to each stratum of each fiber. (For example, the explosion of a simple normal crossing degeneration satisfies this condition.) In this case, the moduli stack of curves in $\hat{\mathbf{B}}$ with bounded genus, energy and number of punctures is proper over \mathbf{B}_0 .

Before describing Gromov–Witten invariants in our family, let us consider how the cohomology of exploded manifolds varies in families. It is not true that ${}^r_{\text{fg}}H^*$ is invariant in families of exploded manifolds; however, Section 11 of [31] proves that H^* is invariant in connected families of exploded manifolds. More precisely, given any long path²² γ in \mathbf{B}_0 , with inverse image in $\hat{\mathbf{B}}$ joining \mathbf{B} to a fiber, \mathbf{B}' , of $\hat{\mathbf{B}} \rightarrow \mathbf{B}_0$, Definition 11.3 of [31] gives an isomorphism $\Psi_\gamma: H^*(\mathbf{B}) \rightarrow H^*(\mathbf{B}')$. This isomorphism depends on the isotopy class of γ . With Ψ_γ understood, we can write the invariance of our correlators as

$$(11) \quad \langle \tau_{a_1}(\theta_1), \dots, \tau_{a_n}(\theta_n) \rangle_{g,n,\beta} = \langle \tau_{a_1}(\Psi_\gamma(\theta_1)), \dots, \tau_{a_n}(\Psi_\gamma(\theta_n)) \rangle_{g,n,(\Psi_\gamma^{-1})^*\beta}.$$

Let us first prove (11) under the assumption that our long path γ is contained in a small neighborhood of $\mathbf{B} \subset \hat{\mathbf{B}}$, and that there exists a curve f in \mathbf{B} representing the homology class $\beta: H^*(\mathbf{B}) \rightarrow \mathbb{R}$. Proposition 5.9 of [26] implies that f locally extends to a connected family of (not necessarily holomorphic) curves, \hat{f} , with $\mathbf{F}(\hat{f}) \rightarrow \mathbf{B}_0$ a submersion. Then, given any long path γ in the image of $\mathbf{F}(\hat{f})$ joining the image of f with the image of f' , the integral of $f^*\theta$ equals the integral of $(f')^*\Psi_\gamma(\theta)$. So, the homology class, $\beta: H^*(\mathbf{B}) \rightarrow \mathbb{R}$, represented by f locally extends to a map $\hat{\beta}: H^*(\mathbf{B}') \rightarrow \mathbb{R}$ for all fibers \mathbf{B}' in a neighborhood of \mathbf{B} such that for any long path γ in this neighborhood, $\hat{\beta}(\theta) = \hat{\beta}(\Psi_\gamma(\theta))$ (unlike the case of a family of smooth

²²See Definition 11.1 of [31].

manifolds, not all homology classes locally extend in this way). Let $\hat{\mathbf{B}}' \rightarrow \mathbf{B}'_0$ be such a neighborhood of \mathbf{B} where $\hat{\beta}$ exists, and let $\mathcal{M}_{g,n,\hat{\beta}}^{\text{st}}(\hat{\mathbf{B}}')$ be the moduli stack of stable curves with genus g and n punctures representing $\hat{\beta}$.

The moduli stack of holomorphic curves in $\mathcal{M}_{g,n,\hat{\beta}}^{\text{st}}(\hat{\mathbf{B}}')$ is complete over \mathbf{B}'_0 , so using Theorem 7.3 of [26] we can construct an embedded Kuranishi structure with an associated Kuranishi category, $\hat{\mathcal{K}}_{g,n,\hat{\beta}}$, oriented and complete over \mathbf{B}_0 . Then, Definition 4.7 gives us a virtual fundamental class $[\hat{\mathcal{K}}_{g,n,\hat{\beta}}]$.

In this case, evaluation at punctures gives an evaluation map

$$\widehat{\text{ev}}: \mathcal{M}_{g,n,\hat{\beta}}^{\text{st}}(\hat{\mathbf{B}}') \rightarrow \hat{\mathbf{X}},$$

where $\hat{\mathbf{X}}$ is the n^{th} fiber product of $\hat{\mathbf{B}}'$ over \mathbf{B}'_0 . The line bundle L_i still make sense over $\mathcal{M}_{g,n,\hat{\beta}}^{\text{st}}(\hat{\mathbf{B}}')$, so we can define $\psi_i \in {}^r_{\text{fg}}H^*(\hat{\mathcal{K}}_{g,n,\hat{\beta}})$ as the first Chern class of L_i^* using Remark 5.2. Because $\widehat{\text{ev}}: \hat{\mathcal{K}}_{g,n,\hat{\beta}} \rightarrow \hat{\mathbf{X}}$ is proper and relatively oriented, we can define

$$\hat{\eta}_{g,\{a_i\},\hat{\beta}} := \widehat{\text{ev}}! \left(\prod_i \psi_i^{a_i} \right) \in {}^r_{\text{fg}}H^*(\hat{\mathbf{X}})$$

using Definition 5.15. Then Theorem 5.22 implies that $\eta_{g,\{a_i\},\beta}$ is the pullback of $\hat{\eta}_{g,\{a_i\},\hat{\beta}}$ under the corresponding inclusion $\mathbf{B}^n \subset \hat{\mathbf{X}}$, and if γ is a long path in \mathbf{B}'_0 joining \mathbf{B} to \mathbf{B}' , $\eta_{g,\{a_i\},(\Psi_\gamma^{-1})^*\beta}$ is the pullback of $\hat{\eta}_{g,\{a_i\},\hat{\beta}}$ under the inclusion $(\mathbf{B}')^n \subset \hat{\mathbf{X}}$. In this case, (11) holds because $\Psi_\gamma(\theta_i)$ is defined by extending θ_i to a closed form $\hat{\theta}_i$ over the long path γ , and $\Psi_\gamma(\theta_i)$ is the restriction of $\hat{\theta}_i$ to \mathbf{B}' ; so, the integral of $\hat{\eta}_{g,\{a_i\},\hat{\beta}} \wedge \hat{\pi}_i^* \hat{\theta}_i$ over fibers is constant, and in particular, its integral over \mathbf{B}^n equals its integral over $(\mathbf{B}')^n$.

So far, we have proved that (11) holds for long paths in \mathbf{B}'_0 starting at a fiber containing a curve representing β . Now we argue that (11) always holds. As with usual paths, we can reparametrize a long path γ into γ_1 followed by γ_2 such that γ_1 starts at the same point as γ and ends at a chosen point in the middle of γ , and γ_2 starts at this chosen point and ends where γ ends. As usual, whenever γ_1 followed by γ_2 is isotopic to a reparametrization of γ , $\Psi_\gamma = \Psi_{\gamma_2} \circ \Psi_{\gamma_1}$, and reversed paths induce inverse isomorphisms. Let S be the set of points in the domain of γ for which (11) holds — so if γ_1 ends at a point in S , (11) holds for γ_1 . Let γ_1 end at a point in the closure of S , and let \mathbf{B}' be the fiber over this point. If our Gromov–Witten invariant (10) is nonzero, then \mathbf{B}' must contain a holomorphic curve representing $(\Psi_{\gamma_1}^{-1})^*\beta$, because the moduli stack of holomorphic curves with bounded energy, genus, and number of ends is proper over \mathbf{B}_0 . Then, the above argument shows that (11) holds with β

replaced by $(\Psi_{\gamma_1}^{-1})^*\beta$, and γ replaced by any sufficiently small long path starting at B' . As (11) holds for a path if and only if it holds for the reversed path, we can reparametrize γ_1 using two long paths for which (11) holds, so (11) holds for γ_1 . Similarly, (11) holds for any long path that can be reparametrized as γ_1 followed by a sufficiently small long path. It follows that the set S where (11) holds is both open and closed, and therefore includes the entire domain of γ . We have therefore shown that (11) holds so long as our Gromov–Witten invariant is nonzero. If it were zero and (11) failed to hold, then our Gromov–Witten invariant at the other end of γ would be nonzero, and therefore (11) would hold for the reversed long path, and therefore must hold for γ itself.

References

- [1] **D Abramovich, Q Chen**, *Stable logarithmic maps to Deligne–Faltings pairs, II*, Asian J. Math. 18 (2014) 465–488 MR
- [2] **K Behrend**, *Cohomology of stacks*, from “Intersection theory and moduli” (E Arbarello, G Ellingsrud, L Goettsche, editors), ICTP Lect. Notes XIX, Abdus Salam Int. Cent. Theoret. Phys., Trieste (2004) 249–294 MR
- [3] **K Behrend, B Fantechi**, *The intrinsic normal cone*, Invent. Math. 128 (1997) 45–88 MR
- [4] **B Chen, A-M Li, B-L Wang**, *Virtual neighborhood technique for pseudo-holomorphic spheres*, preprint (2013) arXiv
- [5] **Q Chen**, *Stable logarithmic maps to Deligne–Faltings pairs, I*, Ann. of Math. 180 (2014) 455–521 MR
- [6] **K Cieliebak, I Mundet i Riera, D A Salamon**, *Equivariant moduli problems, branched manifolds, and the Euler class*, Topology 42 (2003) 641–700 MR
- [7] **B Fantechi**, *Stacks for everybody*, from “European Congress of Mathematics, I” (C Casacuberta, R M Miró-Roig, J Verdera, S Xambó-Descamps, editors), Progr. Math. 201, Birkhäuser, Basel (2001) 349–359 MR
- [8] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Technical details on Kuranishi structure and virtual fundamental chain*, preprint (2012) arXiv
- [9] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Shrinking good coordinate systems associated to Kuranishi structures*, J. Symplectic Geom. 14 (2016) 1295–1310 MR
- [10] **K Fukaya, K Ono**, *Arnold conjecture and Gromov–Witten invariant*, Topology 38 (1999) 933–1048 MR
- [11] **M Gross, B Siebert**, *Logarithmic Gromov–Witten invariants*, J. Amer. Math. Soc. 26 (2013) 451–510 MR

- [12] **D Joyce**, *D-manifolds and d-orbifolds: a theory of derived differential geometry*, book project (2012) Available at <http://people.maths.ox.ac.uk/~joyce/dmanifolds.html>
- [13] **D Joyce**, *A new definition of Kuranishi space*, preprint (2014) arXiv
- [14] **M Kontsevich, Y Manin**, *Gromov–Witten classes, quantum cohomology, and enumerative geometry*, Comm. Math. Phys. 164 (1994) 525–562 MR
- [15] **J Li, G Tian**, *Virtual moduli cycles and Gromov–Witten invariants of general symplectic manifolds*, from “Topics in symplectic 4-manifolds” (R J Stern, editor), First Int. Press Lect. Ser. I, International, Cambridge, MA (1998) 47–83 MR
- [16] **G Liu, G Tian**, *Floer homology and Arnold conjecture*, J. Differential Geom. 49 (1998) 1–74 MR
- [17] **D McDuff**, *The virtual moduli cycle*, from “Northern California symplectic geometry seminar” (Y Eliashberg, D Fuchs, T Ratiu, A Weinstein, editors), Amer. Math. Soc. Transl. Ser. 2 196, Amer. Math. Soc., Providence, RI (1999) 73–102 MR
- [18] **D McDuff**, *Groupoids, branched manifolds and multisections*, J. Symplectic Geom. 4 (2006) 259–315 MR
- [19] **D McDuff, K Wehrheim**, *Kuranishi atlases with trivial isotropy: the 2013 state of affairs*, preprint (2012) arXiv
- [20] **D McDuff, K Wehrheim**, *Smooth Kuranishi atlases with isotropy*, Geom. Topol. 21 (2017) 2725–2809 MR
- [21] **D McDuff, K Wehrheim**, *The topology of Kuranishi atlases*, Proc. Lond. Math. Soc. 115 (2017) 221–292 MR
- [22] **D McDuff, K Wehrheim**, *The fundamental class of smooth Kuranishi atlases with trivial isotropy*, J. Topol. Anal. 10 (2018) 71–243 MR
- [23] **J Pardon**, *An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves*, Geom. Topol. 20 (2016) 779–1034 MR
- [24] **B Parker**, *Exploded manifolds*, Adv. Math. 229 (2012) 3256–3319 MR
- [25] **B Parker**, *Log geometry and exploded manifolds*, Abh. Math. Semin. Univ. Hambg. 82 (2012) 43–81 MR
- [26] **B Parker**, *Holomorphic curves in exploded manifolds: Kuranishi structure*, preprint (2013) arXiv
- [27] **B Parker**, *Universal tropical structures for curves in exploded manifolds*, preprint (2013) arXiv
- [28] **B Parker**, *Holomorphic curves in exploded manifolds: compactness*, Adv. Math. 283 (2015) 377–457 MR
- [29] **B Parker**, *Notes on exploded manifolds and a tropical gluing formula for Gromov–Witten invariants*, preprint (2016) arXiv

- [30] **B Parker**, *Tropical gluing formulae for Gromov–Witten invariants*, preprint (2017) arXiv
- [31] **B Parker**, *De Rham theory of exploded manifolds*, *Geom. Topol.* 22 (2018) 1–54 MR
- [32] **B Parker**, *Holomorphic curves in exploded manifolds: regularity*, *Geom. Topol.* 23 (2019) 1621–1690
- [33] **Y Ruan**, *Virtual neighborhoods and pseudo-holomorphic curves*, *Turkish J. Math.* 23 (1999) 161–231 MR
- [34] **B Siebert**, *Gromov–Witten invariants of general symplectic manifolds*, preprint (2000) arXiv
- [35] **M F Tehrani, M McLean, A Zinger**, *Normal crossings singularities for symplectic topology*, preprint (2014) arXiv

Monash University
 Melbourne, VIC, Australia
 bretttdparker@gmail.com

Proposed: Richard Thomas
 Seconded: Jim Bryan, Dan Abramovich

Received: 29 June 2017
 Revised: 10 September 2018